Chapter 1

Introduction and Overview

The course has a website at

http://www.theory.caltech.edu/~preskill/ph229

General information can be found there, including a course outline and links to relevant references.

Our topic can be approached from a variety of points of view, but these lectures will adopt the perspective of a theoretical physicist (that is, it's my perspective and I'm a theoretical physicist). Because of the interdisciplinary character of the subject, I realize that the students will have a broad spectrum of backgrounds, and I will try to allow for that in the lectures. Please give me feedback if I am assuming things that you don't know.

1.1 Physics of information

Why is a physicist teaching a course about information? In fact, the *physics* of information and computation has been a recognized discipline for at least several decades. This is natural. Information, after all, is something that is encoded in the state of a physical system; a computation is something that can be carried out on an actual physically realizable device. So the study of information and computation should be linked to the study of the underlying physical processes. Certainly, from an engineering perspective, mastery of principles of physics and materials science is needed to develop state-of-the-art computing hardware. (Carver Mead calls his Caltech research group, dedicated to advancing the art of chip design, the "Physics of Computation" (Physcmp) group).

From a more abstract theoretical perspective, there have been noteworthy milestones in our understanding of how physics constrains our ability to use and manipulate information. For example:

• Landauer's principle. Rolf Landauer pointed out in 1961 that erasure of information is necessarily a *dissipative* process. His insight is that erasure always involves the compression of phase space, and so is irreversible.

For example, I can store one bit of information by placing a single molecule in a box, either on the left side or the right side of a partition that divides the box. Erasure means that we move the molecule to the left side (say) irrespective of whether it started out on the left or right. I can suddenly remove the partition, and then slowly compress the one-molecule "gas" with a piston until the molecule is definitely on the left side. This procedure reduces the entropy of the gas by $\Delta S = k \ln 2$ and there is an associated flow of heat from the box to the environment. If the process is isothermal at temperature T, then work $W = kT \ln 2$ is performed on the box, work that I have to provide. If I am to erase information, someone will have to pay the power bill.

• **Reversible computation**. The logic gates used to perform computation are typically *irreversible*, e.g., the NAND gate

$$(a,b) \to \neg (a \land b) \tag{1.1}$$

has two input bits and one output bit, and we can't recover a unique input from the output bit. According to Landauer's principle, since about one bit is erased by the gate (averaged over its possible inputs), at least work $W = kT \ln 2$ is needed to operate the gate. If we have a finite supply of batteries, there appears to be a theoretical limit to how long a computation we can perform.

But Charles Bennett found in 1973 that any computation can be performed using only reversible steps, and so in principle requires no dissipation and no power expenditure. We can actually construct a reversible version of the NAND gate that preserves all the information about the input: For example, the (Toffoli) gate

$$(a, b, c) \to (a, b, c \oplus a \land b) \tag{1.2}$$

is a reversible 3-bit gate that flips the third bit if the first two both take the value 1 and does nothing otherwise. The third output bit becomes the NAND of a and b if c = 1. We can transform an irreversible computation to a reversible one by replacing the NAND gates by Toffoli gates. This computation could in principle be done with negligible dissipation.

However, in the process we generate a lot of extra junk, and one wonders whether we have only postponed the energy cost; we'll have to pay when we need to erase all the junk. Bennett addressed this issue by pointing out that a reversible computer can run forward to the end of a computation, print out a copy of the answer (a logically reversible operation) and then *reverse* all of its steps to return to its initial configuration. This procedure removes the junk without any energy cost.

In principle, then, we need not pay any power bill to compute. In practice, the (irreversible) computers in use today dissipate orders of magnitude more than $kT \ln 2$ per gate, anyway, so Landauer's limit is not an important engineering consideration. But as computing hardware continues to shrink in size, it may become important to beat Landauer's limit to prevent the components from melting, and then reversible computation may be the only option.

• Maxwell's demon. The insights of Landauer and Bennett led Bennett in 1982 to the reconciliation of Maxwell's demon with the second law of thermodynamics. Maxwell had envisioned a gas in a box, divided by a partition into two parts A and B. The partition contains a shutter operated by the demon. The demon observes the molecules in the box as they approach the shutter, allowing fast ones to pass from A to B, and slow ones from B to A. Hence, A cools and B heats up, with a negligible expenditure of work. Heat flows from a cold place to a hot place at no cost, in apparent violation of the second law.

The resolution is that the demon must collect and store information about the molecules. If the demon has a finite memory capacity, he cannot continue to cool the gas indefinitely; eventually, information must be erased. At that point, we finally pay the power bill for the cooling we achieved. (If the demon does not erase his record, or if we want to do the thermodynamic accounting before the erasure, then we should associate some entropy with the recorded information.)

These insights were largely anticipated by Leo Szilard in 1929; he was truly a pioneer of the physics of information. Szilard, in *his* analysis of the Maxwell demon, invented the concept of a *bit* of information, (the *name* "bit" was introduced later, by Tukey) and associated the entropy $\Delta S = k \ln 2$ with the acquisition of one bit (though Szilard does not seem to have fully grasped Landauer's principle, that it is the *erasure* of the bit that carries an inevitable $\cos t$).

These examples illustrate that work at the interface of physics and information has generated noteworthy results of interest to both physicists and computer scientists.

1.2 Quantum information

The moral we draw is that "information is physical." and it is instructive to consider what physics has to tell us about information. But fundamentally, the universe is quantum mechanical. How does quantum theory shed light on the nature of information?

It must have been clear already in the early days of quantum theory that classical ideas about information would need revision under the new physics. For example, the clicks registered in a detector that monitors a radioactive source are described by a *truly random* Poisson process. In contrast, there is no place for true randomness in deterministic classical dynamics (although of course a complex (chaotic) classical system can exhibit behavior that is in practice indistinguishable from random).

Furthermore, in quantum theory, noncommuting observables cannot simultaneously have precisely defined values (the uncertainty principle), and in fact performing a measurement of one observable A will necessarily influence the outcome of a subsequent measurement of an observable B, if A and Bdo not commute. Hence, the act of acquiring information about a physical system inevitably disturbs the state of the system. There is no counterpart of this limitation in classical physics.

The tradeoff between acquiring information and creating a disturbance is related to quantum randomness. It is because the outcome of a measurement has a random element that we are unable to infer the initial state of the system from the measurement outcome.

That acquiring information causes a disturbance is also connected with another essential distinction between quantum and classical information: quantum information cannot be copied with perfect fidelity (the no-cloning principle annunciated by Wootters and Zurek and by Dieks in 1982). If we *could* make a perfect copy of a quantum state, we could measure an observable of the copy without disturbing the original and we could defeat the principle of disturbance. On the other hand, nothing prevents us from copying classical information perfectly (a welcome feature when you need to back up your hard disk).

These properties of quantum information are important, but the really deep way in which quantum information differs from classical information emerged from the work of John Bell (1964), who showed that the predictions of quantum mechanics cannot be reproduced by any local hidden variable theory. Bell showed that quantum information can be (in fact, typically is) encoded in nonlocal correlations between the different parts of a physical system, correlations with no classical counterpart. We will discuss Bell's theorem in detail later on, and I will also return to it later in this lecture.

The study of quantum information as a coherent discipline began to emerge in the 1980's, and it has blossomed in the 1990's. Many of the central results of classical information theory have quantum analogs that have been discovered and developed recently, and we will discuss some of these developments later in the course, including: compression of quantum information, bounds on classical information encoded in quantum systems, bounds on quantum information sent reliably over a noisy quantum channel.

1.3 Efficient quantum algorithms

Given that quantum information has many unusual properties, it might have been expected that quantum theory would have a profound impact on our understanding of computation. That this is spectacularly true came to many of us as a bolt from the blue unleashed by Peter Shor (an AT&T computer scientist and a former Caltech undergraduate) in April, 1994. Shor demonstrated that, at least in principle, a quantum computer can *factor* a large number efficiently.

Factoring (finding the prime factors of a composite number) is an example of an *intractable* problem with the property:

- The solution can be *easily verified*, once found.
- But the solution is *hard* to find.

That is, if p and q are large prime numbers, the product n = pq can be computed quickly (the number of elementary bit operations required is about $\log_2 p \cdot \log_2 q$). But given n, it is *hard* to find p and q.

The time required to find the factors is strongly believed (though this has never been proved) to be *superpolynomial* in $\log(n)$. That is, as *n* increases, the time needed in the worst case grows faster than any power of $\log(n)$. The best known factoring algorithm (the "number field sieve") requires

time
$$\simeq \exp[c(\ln n)^{1/3}(\ln \ln n)^{2/3}]$$
 (1.3)

where $c = (64/9)^{1/3} \sim 1.9$. The current state of the art is that the 65 digit factors of a 130 digit number can be found in the order of one month by a network of hundreds of work stations. Using this to estimate the prefactor in Eq. 1.3, we can estimate that factoring a 400 digit number would take about 10^{10} years, the age of the universe. So even with vast improvements in technology, factoring a 400 digit number will be out of reach for a while.

The factoring problem is interesting from the perspective of complexity theory, as an example of a problem presumed to be intractable; that is, a problem that can't be solved in a time bounded by a polynomial in the size of the input, in this case $\log n$. But it is also of practical importance, because the difficulty of factoring is the basis of schemes for public key cryptography, such as the widely used RSA scheme.

The exciting new result that Shor found is that a quantum computer can factor in polynomial time, *e.g.*, in time $O[(\ln n)^3]$. So if we had a quantum computer that could factor a 130 digit number in one month (of course we don't, at least not yet!), running Shor's algorithm it could factor that 400 digit number in less than 3 years. The harder the problem, the greater the advantage enjoyed by the quantum computer.

Shor's result spurred my own interest in quantum information (were it not for Shor, I don't suppose I would be teaching this course). It's fascinating to contemplate the implications for complexity theory, for quantum theory, for technology.

1.4 Quantum complexity

Of course, Shor's work had important antecedents. That a quantum system can perform a computation was first explicitly pointed out by Paul Benioff and Richard Feynman (independently) in 1982. In a way, this was a natural issue to wonder about in view of the relentless trend toward miniaturization in microcircuitry. If the trend continues, we will eventually approach the regime where quantum theory is highly relevant to how computing devices function. Perhaps this consideration provided some of the motivation behind Benioff's work. But Feynman's primary motivation was quite different and very interesting. To understand Feynman's viewpoint, we'll need to be more

1.4. QUANTUM COMPLEXITY

explicit about the mathematical description of quantum information and computation.

The indivisible unit of classical information is the bit: an object that can take either one of two values: 0 or 1. The corresponding unit of quantum information is the quantum bit or *qubit*. The qubit is a vector in a two-dimensional complex vector space with inner product; in deference to the classical bit we can call the elements of an orthonormal basis in this space $|0\rangle$ and $|1\rangle$. Then a normalized vector can be represented

$$|\psi\rangle = a|0\rangle + b|1\rangle, \quad |a|^2 + |b|^2 = 1.$$
 (1.4)

where $a, b \in \mathbb{C}$. We can perform a measurement that projects $|\psi\rangle$ onto the basis $|0\rangle, |1\rangle$. The outcome of the measurement is not deterministic — the probability that we obtain the result $|0\rangle$ is $|a|^2$ and the probability that we obtain the result $|1\rangle$ is $|b|^2$.

The quantum state of N qubits can be expressed as a vector in a space of dimension 2^N . We can choose as an orthonormal basis for this space the states in which each qubit has a definite value, either $|0\rangle$ or $|1\rangle$. These can be labeled by binary strings such as

$$|01110010\cdots 1001\rangle$$
 (1.5)

A general normalized vector can be expanded in this basis as

$$\sum_{x=0}^{2^{N}-1} a_x |x\rangle , \qquad (1.6)$$

where we have associated with each string the number that it represents in binary notation, ranging in value from 0 to $2^N - 1$. Here the a_x 's are complex numbers satisfying $\sum_x |a_x|^2 = 1$. If we measure all N qubits by projecting each onto the $\{|0\rangle, |1\rangle\}$ basis, the probability of obtaining the outcome $|x\rangle$ is $|a_x|^2$.

Now, a quantum computation can be described this way. We assemble N qubits, and prepare them in a standard initial state such as $|0\rangle|0\rangle\cdots|0\rangle$, or $|x=0\rangle$. We then apply a unitary transformation U to the N qubits. (The transformation U is constructed as a product of standard quantum gates, unitary transformations that act on just a few qubits at a time). After U is applied, we measure all of the qubits by projecting onto the $\{|0\rangle, |1\rangle\}$ basis. The measurement outcome is the output of the computation. So the final

output is classical information that can be printed out on a piece of paper, and published in Physical Review.

Notice that the algorithm performed by the quantum computer is a *probabilistic* algorithm. That is, we could run exactly the same program twice and obtain different results, because of the randomness of the quantum measurement process. The quantum algorithm actually generates a probability distribution of possible outputs. (In fact, Shor's factoring algorithm is not guaranteed to succeed in finding the prime factors; it just succeeds with a reasonable probability. That's okay, though, because it is easy to verify whether the factors are correct.)

It should be clear from this description that a quantum computer, though it may operate according to different physical principles than a classical computer, cannot do anything that a classical computer can't do. Classical computers can store vectors, rotate vectors, and can model the quantum measurement process by projecting a vector onto mutually orthogonal axes. So a classical computer can surely *simulate* a quantum computer to arbitrarily good accuracy. Our notion of what is *computable* will be the same, whether we use a classical computer or a quantum computer.

But we should also consider how long the simulation will take. Suppose we have a computer that operates on a modest number of qubits, like N = 100. Then to represent the typical quantum state of the computer, we would need to write down $2^N = 2^{100} \sim 10^{30}$ complex numbers! No existing or foreseeable digital computer will be able to do that. And performing a general rotation of a vector in a space of dimension 10^{30} is far beyond the computational capacity of any foreseeable classical computer.

(Of course, N classical bits can take 2^N possible values. But for each one of these, it is very easy to write down a complete description of the configuration — a binary string of length N. Quantum information is very different in that writing down a complete description of just one typical configuration of N qubits is enormously complex.)

So it is true that a classical computer can simulate a quantum computer, but the simulation becomes extremely inefficient as the number of qubits Nincreases. Quantum mechanics is *hard* (computationally) because we must deal with huge matrices – there is too much room in Hilbert space. This observation led Feynman to speculate that a quantum computer would be able to perform certain tasks that are beyond the reach of any conceivable classical computer. (The quantum computer has no trouble simulating *itself*!) Shor's result seems to bolster this view.

1.4. QUANTUM COMPLEXITY

Is this conclusion unavoidable? In the end, our simulation should provide a means of assigning probabilities to all the possible outcomes of the final measurement. It is not really necessary, then, for the classical simulation to track the complete description of the *N*-qubit quantum state. We would settle for a *probabilistic classical algorithm*, in which the outcome is not uniquely determined by the input, but in which various outcomes arise with a probability distribution that coincides with that generated by the quantum computation. We might hope to perform a *local* simulation, in which each qubit has a definite value at each time step, and each quantum gate can act on the qubits in various possible ways, one of which is selected as determined by a (pseudo)-random number generator. This simulation would be much easier than following the evolution of a vector in an exponentially large space.

But the conclusion of John Bell's powerful theorem is *precisely* that this simulation could never work: there is no *local probabilistic algorithm* that can reproduce the conclusions of quantum mechanics. Thus, while there is no known proof, it seems highly likely that simulating a quantum computer is a very hard problem for any classical computer.

To understand better why the mathematical description of quantum information is necessarily so complex, imagine we have a 3N-qubit quantum system $(N \gg 1)$ divided into three subsystems of N qubits each (called subsystems (1),(2), and (3). We randomly choose a quantum state of the 3Nqubits, and then we separate the 3 subsystems, sending (1) to Santa Barbara and (3) to San Diego, while (2) remains in Pasadena. Now we would like to make some measurements to find out as much as we can about the quantum state. To make it easy on ourselves, let's imagine that we have a zillion copies of the state of the system so that we can measure any and all the observables we want.¹ Except for one proviso: we are restricted to carrying out each measurement within one of the subsystems — no collective measurements spanning the boundaries between the subsystems are allowed. Then for a typical state of the 3N-qubit system, our measurements will reveal almost *nothing* about what the state is. Nearly all the information that distinguishes one state from another is in the *nonlocal correlations* between measurement outcomes in subsystem (1) (2), and (3). These are the nonlocal correlations that Bell found to be an essential part of the physical description.

¹We cannot make copies of an unknown quantum state ourselves, but we can ask a friend to prepare many identical copies of the state (he can do it because he knows what the state is), and not tell us what he did.

We'll see that information content can be quantified by entropy (large entropy means little information.) If we choose a state for the 3N qubits randomly, we almost always find that the entropy of each subsystem is very close to

$$S \cong N - 2^{-(N+1)},$$
 (1.7)

a result found by Don Page. Here N is the maximum possible value of the entropy, corresponding to the case in which the subsystem carries no accessible information at all. Thus, for large N we can access only an exponentially small amount of information by looking at each subsystem separately.

That is, the measurements reveal very little information if we don't consider how measurement results obtained in San Diego, Pasadena, and Santa Barbara are correlated with one another — in the language I am using, a measurement of a correlation is considered to be a "collective" measurement (even though it could actually be performed by experimenters who observe the separate parts of the same copy of the state, and then exchange phone calls to compare their results). By measuring the correlations we can learn much more; in principle, we can completely reconstruct the state.

Any satisfactory description of the state of the 3N qubits must characterize these nonlocal correlations, which are exceedingly complex. This is why a classical simulation of a large quantum system requires vast resources. (When such nonlocal correlations exist among the parts of a system, we say that the parts are "entangled," meaning that we can't fully decipher the state of the system by dividing the system up and studying the separate parts.)

1.5 Quantum parallelism

Feynman's idea was put in a more concrete form by David Deutsch in 1985. Deutsch emphasized that a quantum computer can best realize its computational potential by invoking what he called "quantum parallelism." To understand what this means, it is best to consider an example.

Following Deutsch, imagine we have a black box that computes a function that takes a single bit x to a single bit f(x). We don't know what is happening inside the box, but it must be something complicated, because the computation takes 24 hours. There are four possible functions f(x) (because each of f(0) and f(1) can take either one of two possible values) and we'd like to know what the box is computing. It would take 48 hours to find out both f(0) and f(1).

But we don't have that much time; we need the answer in 24 hours, not 48. And it turns out that we would be satisfied to know whether f(x) is constant (f(0) = f(1)) or balanced $(f(0) \neq f(1))$. Even so, it takes 48 hours to get the answer.

Now suppose we have a quantum black box that computes f(x). Of course f(x) might not be invertible, while the action of our quantum computer is unitary and must be invertible, so we'll need a transformation U_f that takes two qubits to two:

$$U_f: |x\rangle|y\rangle \to |x\rangle|y \oplus f(x)\rangle$$
 . (1.8)

(This machine flips the second qubit if f acting on the first qubit is 1, and doesn't do anything if f acting on the first qubit is 0.) We can determine if f(x) is constant or balanced by using the quantum black box twice. But it still takes a day for it to produce one output, so that won't do. Can we get the answer (in 24 hours) by running the quantum black box *just once*. (This is "Deutsch's problem.")

Because the black box is a quantum computer, we can choose the input state to be a *superposition* of $|0\rangle$ and $|1\rangle$. If the second qubit is initially prepared in the state $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, then

$$U_{f}:|x\rangle \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \rightarrow |x\rangle \frac{1}{\sqrt{2}}(|f(x)\rangle - |1 \oplus f(x)\rangle)$$
$$= |x\rangle (-1)^{f(x)} \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle), \qquad (1.9)$$

so we have isolated the function f in an x-dependent phase. Now suppose we prepare the first qubit as $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. Then the black box acts as

$$U_{f}: \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \rightarrow \frac{1}{\sqrt{2}}\left[(-1)^{f(0)}|0\rangle + (-1)^{f(1)}|1\rangle\right]\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) .$$
(1.10)

Finally, we can perform a measurement that projects the first qubit onto the basis

$$|\pm\rangle = \frac{1}{\sqrt{2}} (|0\rangle \pm |1\rangle). \tag{1.11}$$

Evidently, we will always obtain $|+\rangle$ if the function is balanced, and $|-\rangle$ if the function is constant.²

So we have solved Deutsch's problem, and we have found a separation between what a classical computer and a quantum computer can achieve. The classical computer has to run the black box twice to distinguish a balanced function from a constant function, but a quantum computer does the job in one go!

This is possible because the quantum computer is not limited to computing either f(0) or f(1). It can act on a superposition of $|0\rangle$ and $|1\rangle$, and thereby extract "global" information about the function, information that depends on both f(0) and f(1). This is quantum parallelism.

Now suppose we are interested in global properties of a function that acts on N bits, a function with 2^N possible arguments. To compute a complete table of values of f(x), we would have to calculate $f 2^N$ times, completely infeasible for $N \gg 1$ (e.g., 10^{30} times for N = 100). But with a quantum computer that acts according to

$$U_f: |x\rangle|0\rangle \to |x\rangle|f(x)\rangle$$
, (1.12)

we could choose the input register to be in a state

$$\left[\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)\right]^{N} = \frac{1}{2^{N/2}} \sum_{x=0}^{2^{N}-1} |x\rangle , \qquad (1.13)$$

and by computing f(x) only once, we can generate a state

$$\frac{1}{2^{N/2}} \sum_{x=0}^{2^{N-1}} |x\rangle |f(x)\rangle .$$
 (1.14)

Global properties of f are encoded in this state, and we might be able to extract some of those properties if we can only think of an efficient way to do it.

This quantum computation exhibits "massive quantum parallelism;" a simulation of the preparation of this state on a classical computer would

²In our earlier description of a quantum computation, we stated that the final measurement would project each qubit onto the $\{|0\rangle, |1\rangle\}$ basis, but here we are allowing measurement in a different basis. To describe the procedure in the earlier framework, we would apply an appropriate unitary change of basis to each qubit before performing the final measurement.

require us to compute f an unimaginably large number of times (for $N \gg 1$). Yet we have done it with the quantum computer in only one go. It is just this kind of massive parallelism that Shor invokes in his factoring algorithm.

As noted earlier, a characteristic feature of quantum information is that it can be encoded in nonlocal correlations among different parts of a physical system. Indeed, this is the case in Eq. (1.14); the properties of the function fare stored as correlations between the "input register" and "output register" of our quantum computer. This nonlocal information, however, is not so easy to decipher.

If, for example, I were to measure the input register, I would obtain a result $|x_0\rangle$, where x_0 is chosen completely at random from the 2^N possible values. This procedure would prepare a state

$$|x_0\rangle|f(x_0)\rangle.\tag{1.15}$$

We could proceed to measure the output register to find the value of $f(x_0)$. But because Eq. (1.14) has been destroyed by the measurement, the intricate correlations among the registers have been lost, and we get no opportunity to determine $f(y_0)$ for any $y_0 \neq x_0$ by making further measurements. In this case, then, the quantum computation provided no advantage over a classical one.

The lesson of the solution to Deutsch's problem is that we can sometimes be more clever in exploiting the correlations encoded in Eq. (1.14). Much of the art of designing quantum algorithms involves finding ways to make efficient use of the nonlocal correlations.

1.6 A new classification of complexity

The computer on your desktop is not a quantum computer, but still it is a remarkable device: in principle, it is capable of performing any conceivable computation. In practice there are computations that you can't do — you either run out of time or you run out of memory. But if you provide an unlimited amount of memory, and you are willing to wait as long as it takes, then anything that deserves to be called a computation can be done by your little PC. We say, therefore, that it is a "universal computer."

Classical complexity theory is the study of which problems are hard and which ones are easy. Usually, "hard" and "easy" are defined in terms of how much time and/or memory are needed. But how can we make meaningful distinctions between hard and easy without specifying the hardware we will be using? A problem might be hard on the PC, but perhaps I could design a special purpose machine that could solve that problem much faster. Or maybe in the future a much better general purpose computer will be available that solves the problem far more efficiently. Truly meaningful distinctions between hard and easy should be *universal* — they ought not to depend on which machine we are using.

Much of complexity theory focuses on the distinction between "polynomial time" and "exponential time" algorithms. For any algorithm A, which can act on an input of variable length, we may associate a *complexity function* $T_A(N)$, where N is the length of the input in bits. $T_A(N)$ is the longest "time" (that is, number of elementary steps) it takes for the algorithm to run to completion, for any N-bit input. (For example, if A is a factoring algorithm, $T_A(N)$ is the time needed to factor an N-bit number in the worst possible case.) We say that A is polynomial time if

$$T_A(N) \le \operatorname{Poly}(N), \tag{1.16}$$

where Poly (N) denotes a polynomial of N. Hence, polynomial time means that the time needed to solve the problem does not grow faster than a power of the number of input bits.

If the problem is not polynomial time, we say it is exponential time (though this is really a misnomer, because of course that are superpolynomial functions like $N^{\log N}$ that actually increase much more slowly than an exponential). This is a reasonable way to draw the line between easy and hard. But the truly compelling reason to make the distinction this way is that it is machine-independent: it does not matter what computer we are using. The universality of the distinction between polynomial and exponential follows from one of the central results of computer science: one universal (classical) computer can simulate another with at worst "polynomial overhead." This means that if an algorithm runs on your computer in polynomial time, then I can always run it on my computer in polynomial time. If I can't think of a better way to do it, I can always have my computer emulate how yours operates; the cost of running the emulation is only polynomial time. Similarly, your computer can emulate mine, so we will always agree on which algorithms are polynomial time.³

 $^{^{3}}$ To make this statement precise, we need to be a little careful. For example, we should exclude certain kinds of "unreasonable" machines, like a parallel computer with an unlimited number of nodes.

1.7. WHAT ABOUT ERRORS?

Now it is true that information and computation in the physical world are fundamentally quantum mechanical, but this insight, however dear to physicists, would not be of great interest (at least from the viewpoint of complexity theory) were it possible to simulate a quantum computer on a classical computer with polynomial overhead. Quantum algorithms might prove to be of technological interest, but perhaps no more so than future advances in classical algorithms that might speed up the solution of certain problems.

But if, as is indicated (but not proved!) by Shor's algorithm, no polynomialtime simulation of a quantum computer is possible, that changes everything. Thirty years of work on complexity theory will still stand as mathematical truth, as theorems characterizing the capabilities of classical universal computers. But it may fall as physical truth, because a classical Turing machine is not an appropriate model of the computations that can really be performed in the physical world.

If the quantum classification of complexity is indeed different than the classical classification (as is suspected but not proved), then this result will shake the foundations of computer science. In the long term, it may also strongly impact technology. But what is its significance for physics?

I'm not sure. But perhaps it is telling that no conceivable classical computation can accurately predict the behavior of even a modest number of qubits (of order 100). This may suggest that relatively small quantum systems have greater potential than we suspected to surprise, baffle, and delight us.

1.7 What about errors?

As significant as Shor's factoring algorithm may prove to be, there is another recently discovered feature of quantum information that may be just as important: the discovery of quantum error correction. Indeed, were it not for this development, the prospects for quantum computing technology would not seem bright.

As we have noted, the essential property of quantum information that a quantum computer exploits is the existence of nonlocal correlations among the different parts of a physical system. If I look at only part of the system at a time, I can decipher only very little of the information encoded in the system. Unfortunately, these nonlocal correlations are extremely fragile and tend to decay very rapidly in practice. The problem is that our quantum system is inevitably in contact with a much larger system, its environment. It is virtually impossible to perfectly isolate a big quantum system from its environment, even if we make a heroic effort to do so. Interactions between a quantum device and its environment establish nonlocal correlations between the two. Eventually the quantum information that we initially encoded in the device becomes encoded, instead, in correlations between the device and the environment. At that stage, we can no longer access the information by observing only the device. In practice, the information is irrevocably lost. Even if the coupling between device and environment is quite weak, this happens to a macroscopic device remarkably quickly.

Erwin Schrödinger chided the proponents of the mainstream interpretation of quantum mechanics by observing that the theory will allow a quantum state of a cat of the form

$$|\text{cat}\rangle = \frac{1}{\sqrt{2}} (|\text{dead}\rangle + |\text{alive}\rangle).$$
 (1.17)

To Schrödinger, the possibility of such states was a blemish on the theory, because every cat he had seen was either dead or alive, not half dead and half alive.

One of the most important advances in quantum theory over the past 15 years is that we have learned how to answer Schrödinger with growing confidence. The state $|cat\rangle$ is possible in principle, but is rarely seen because it is *extremely* unstable. The cats Schrödinger observed were never well isolated from the environment. If someone were to prepare the state $|cat\rangle$, the quantum information encoded in the superposition of $|dead\rangle$ and $|alive\rangle$ would *immediately* be transferred to correlations between the cat and the environment, and become completely inaccessible. In effect, the environment continually measures the cat, projecting it onto either the state $|alive\rangle$ or $|dead\rangle$. This process is called *decoherence*. We will return to the study of decoherence later in the course.

Now, to perform a complex quantum computation, we need to prepare a delicate superposition of states of a relatively large quantum system (though perhaps not as large as a cat). Unfortunately, this system cannot be perfectly isolated from the environment, so this superposition, like the state $|cat\rangle$, decays very rapidly. The encoded quantum information is quickly lost, and our quantum computer crashes.

To put it another way, contact between the computer and the environment (decoherence) causes *errors* that degrade the quantum information. To operate a quantum computer reliably, we must find some way to prevent or correct these errors.

Actually, decoherence is not our only problem. Even if we could achieve perfect isolation from the environment, we could not expect to operate a quantum computer with perfect accuracy. The quantum gates that the machine executes are unitary transformations that operate on a few qubits at a time, let's say 4×4 unitary matrices acting on two qubits. Of course, these unitary matrices form a continuum. We may have a protocol for applying U_0 to 2 qubits, but our execution of the protocol will not be flawless, so the actual transformation

$$U = U_0 \left(1 + O(\varepsilon) \right) \tag{1.18}$$

will differ from the intended U_0 by some amount of order ε . After about $1/\varepsilon$ gates are applied, these errors will accumulate and induce a serious failure. Classical analog devices suffer from a similar problem, but small errors are much less of a problem for devices that perform discrete logic.

In fact, modern digital circuits are remarkably reliable. They achieve such high accuracy with help from *dissipation*. We can envision a classical gate that acts on a bit, encoded as a ball residing at one of the two minima of a double-lobed potential. The gate may push the ball over the intervening barrier to the other side of the potential. Of course, the gate won't be implemented perfectly; it may push the ball a little too hard. Over time, these imperfections might accumulate, causing an error.

To improve the performance, we *cool* the bit (in effect) after each gate. This is a dissipative process that releases heat to the environment and compresses the phase space of the ball, bringing it close to the local minimum of the potential. So the small errors that we may make wind up heating the environment rather than compromising the performance of the device.

But we can't cool a quantum computer this way. Contact with the environment may enhance the reliability of classical information, but it would destroy encoded quantum information. More generally, accumulation of error will be a problem for classical reversible computation as well. To prevent errors from building up we need to discard the information about the errors, and throwing away information is always a dissipative process.

Still, let's not give up too easily. A sophisticated machinery has been developed to contend with errors in classical information, the theory of error correcting codes. To what extent can we coopt this wisdom to protect quantum information as well?

How does classical error correction work? The simplest example of a classical error-correcting code is a repetition code: we replace the bit we wish to protect by 3 copies of the bit,

$$\begin{array}{rcl} 0 & \rightarrow & (000), \\ 1 & \rightarrow & (111). \end{array} \tag{1.19}$$

Now an error may occur that causes one of the three bits to flip; if it's the first bit, say,

$$(000) \rightarrow (100),$$

 $(111) \rightarrow (011).$ (1.20)

Now in spite of the error, we can still decode the bit correctly, by majority voting.

Of course, if the probability of error in each bit were p, it would be possible for two of the three bits to flip, or even for all three to flip. A double flip can happen in three different ways, so the probability of a double flip is $3p^2(1-p)$, while the probability of a triple flip is p^3 . Altogether, then, the probability that majority voting fails is $3p^2(1-p) + p^3 = 3p^2 - 2p^3$. But for

$$3p^2 - 2p^3 (1.21)$$

the code improves the reliability of the information.

We can improve the reliability further by using a longer code. One such code (though far from the most efficient) is an N-bit repetition code. The probability distribution for the average value of the bit, by the central limit theorem, approaches a Gaussian with width $1/\sqrt{N}$ as $N \to \infty$. If $P = \frac{1}{2} + \varepsilon$ is the probability that each bit has the correct value, then the probability that the majority vote fails (for large N) is

$$P_{error} \sim e^{-N\varepsilon^2},$$
 (1.22)

arising from the tail of the Gaussian. Thus, for any $\varepsilon > 0$, by introducing enough redundancy we can achieve arbitrarily good reliability. Even for $\varepsilon < 0$, we'll be okay if we always assume that majority voting gives the wrong result. Only for $P = \frac{1}{2}$ is the cause lost, for then our block of N bits will be random, and encode no information.

In the 50's, John Von Neumann showed that a classical computer with noisy components can work reliably, by employing sufficient redundancy. He pointed out that, if necessary, we can compute each logic gate many times, and accept the majority result. (Von Neumann was especially interested in how his brain was able to function so well, in spite of the unreliability of neurons. He was pleased to explain why he was so smart.)

But now we want to use error correction to keep a *quantum computer* on track, and we can immediately see that there are difficulties:

1. **Phase errors**. With quantum information, more things can go wrong. In addition to bit-flip errors

there can also be *phase* errors

A phase error is serious, because it makes the state $\frac{1}{\sqrt{2}}[|0\rangle + |1\rangle]$ flip to the orthogonal state $\frac{1}{\sqrt{2}}[|0\rangle - |1\rangle]$. But the classical coding provided no protection against phase errors.

2. Small errors. As already noted, quantum information is continuous. If a qubit is intended to be in the state

$$a|0\rangle + b|1\rangle, \tag{1.25}$$

an error might change a and b by an amount of order ε , and these small errors can accumulate over time. The classical method is designed to correct large (bit flip) errors.

- 3. Measurement causes disturbance. In the majority voting scheme, it seemed that we needed to *measure* the bits in the code to detect and correct the errors. But we can't measure qubits without *disturbing* the quantum information that they encode.
- 4. **No cloning**. With classical coding, we protected information by making extra copies of it. But we know that quantum information cannot be copied with perfect fidelity.

1.8 Quantum error-correcting codes

Despite these obstacles, it turns out that quantum error correction really is possible. The first example of a quantum error-correcting code was constructed about two years ago by (guess who!) Peter Shor. This discovery ushered in a new discipline that has matured remarkably quickly – the theory of quantum error-correcting codes. We will study this theory later in the course.

Probably the best way to understand how quantum error correction works is to examine Shor's original code. It is the most straightforward quantum generalization of the classical 3-bit repetition code.

Let's look at that 3-bit code one more time, but this time mindful of the requirement that, with a quantum code, we will need to be able to correct the errors without measuring any of the encoded information.

Suppose we encode a single qubit with 3 qubits:

$$\begin{aligned} |0\rangle &\to & |\bar{0}\rangle \equiv |000\rangle, \\ |1\rangle &\to & |\bar{1}\rangle \equiv |111\rangle, \end{aligned}$$
 (1.26)

or, in other words, we encode a superposition

$$a|0\rangle + b|1\rangle \rightarrow a|\bar{0}\rangle + b|\bar{1}\rangle = a|000\rangle + b|111\rangle$$
 (1.27)

We would like to be able to correct a bit flip error without destroying this superposition.

Of course, it won't do to measure a single qubit. If I measure the first qubit and get the result $|0\rangle$, then I have prepared the state $|\bar{0}\rangle$ of all three qubits, and we have lost the quantum information encoded in the coefficients a and b.

But there is no need to restrict our attention to single-qubit measurements. I could also perform collective measurements on two-qubits at once, and collective measurements suffice to diagnose a bit-flip error. For a 3-qubit state $|x, y, z\rangle$ I could measure, say, the two-qubit observables $y \oplus z$, or $x \oplus z$ (where \oplus denotes addition modulo 2). For both $|x, y, z\rangle = |000\rangle$ and $|111\rangle$ these would be 0, but if any one bit flips, then at least one of these quantities will be 1. In fact, if there is a single bit flip, the two bits

$$(y \oplus z, x \oplus z), \tag{1.28}$$

just designate in binary notation the position (1, 2 or 3) of the bit that flipped. These two bits constitute a *syndrome* that diagnoses the error that occurred.

For example, if the first bit flips,

$$a|000\rangle + b|111\rangle \rightarrow a|100\rangle + b|011\rangle,$$
 (1.29)

then the measurement of $(y \oplus z, x \oplus z)$ yields the result (0, 1), which instructs us to flip the first bit; this indeed repairs the error.

Of course, instead of a (large) bit flip there could be a small error:

$$\begin{aligned} |000\rangle &\to & |000\rangle + \varepsilon |100\rangle \\ |111\rangle &\to & |111\rangle - \varepsilon |011\rangle. \end{aligned}$$
 (1.30)

But even in this case the above procedure would work fine. In measuring $(y \oplus z, x \oplus z)$, we would project out an eigenstate of this observable. Most of the time (probability $1 - |\varepsilon|^2$) we obtain the result (0, 0) and project the damaged state back to the original state, and so correct the error. Occasionally (probability $|\varepsilon|^2$) we obtain the result (0, 1) and project the state onto Eq. 1.29. But then the syndrome instructs us to flip the first bit, which restores the original state. Similarly, if there is an amplitude of order ε for each of the three qubits to flip, then with a probability of order $|\varepsilon|^2$ the syndrome measurement will project the state to one in which one of the three bits is flipped, and the syndrome will tell us which one.

So we have already overcome 3 of the 4 obstacles cited earlier. We see that it is possible to make a measurement that diagnoses the error without damaging the information (answering (3)), and that a quantum measurement can project a state with a small error to either a state with no error or a state with a large discrete error that we know how to correct (answering (2)). As for (4), the issue didn't come up, because the state $a|\bar{0}\rangle + b|\bar{1}\rangle$ is not obtained by cloning – it is not the same as $(a|0\rangle + b|1\rangle)^3$; that is, it differs from three copies of the unencoded state.

Only one challenge remains: (1) phase errors. Our code does not yet provide any protection against phase errors, for if any one of the three qubits undergoes a phase error then our encoded state $a|\bar{0}\rangle + b|\bar{1}\rangle$ is transformed to $a|\bar{0}\rangle - b|\bar{1}\rangle$, and the encoded quantum information is damaged. In fact, phase errors have become three times more likely than if we hadn't used the code. But with the methods in hand that conquered problems (2)-(4), we can approach problem (1) with new confidence. Having protected against bit-flip errors by encoding bits redundantly, we are led to protect against phase-flip errors by encoding phases redundantly.

Following Shor, we encode a single qubit using nine qubits, according to

$$\begin{aligned} |0\rangle \to |\bar{0}\rangle &\equiv \frac{1}{2^{3/2}} (|000\rangle + |111\rangle) (|000\rangle + |111\rangle) (|000\rangle + |111\rangle) ,\\ |1\rangle \to |\bar{1}\rangle &\equiv \frac{1}{2^{3/2}} (|000\rangle - |111\rangle) (|000\rangle - |111\rangle) (|000\rangle - |111\rangle) .(1.31) \end{aligned}$$

Both $|\bar{0}\rangle$ and $|\bar{1}\rangle$ consist of three clusters of three qubits each, with each cluster prepared in the same quantum state. Each of the clusters has triple bit redundancy, so we can correct a single bit flip in any cluster by the method discussed above.

Now suppose that a phase flip occurs in one of the clusters. The error changes the relative sign of $|000\rangle$ and $|111\rangle$ in that cluster so that

$$|000\rangle + |111\rangle \rightarrow |000\rangle - |111\rangle, |000\rangle - |111\rangle \rightarrow |000\rangle + |111\rangle.$$
 (1.32)

This means that the relative phase of the damaged cluster differs from the phases of the other two clusters. Thus, as in our discussion of bit-flip correction, we can identify the damaged cluster, not by *measuring* the relative phase in each cluster (which would disturb the encoded information) but by *comparing* the phases of pairs of clusters. In this case, we need to measure a six-qubit observable to do the comparison, e.g., the observable that flips qubits 1 through 6. Since flipping twice is the identity, this observable squares to 1, and has eigenvalues ± 1 . A pair of clusters with the same sign is an eigenstate with eigenvalue +1, and a pair of clusters with opposite sign is an eigenstate with eigenvalue -1. By measuring the six-qubit observable for a second pair of clusters, we can determine which cluster has a different sign than the others. Then, we apply a unitary phase transformation to one of the qubits in that cluster to reverse the sign and correct the error.

Now suppose that a unitary error $U = 1 + 0(\varepsilon)$ occurs for each of the 9 qubits. The most general single-qubit unitary transformation (aside from a physically irrelevant overall phase) can be expanded to order ε as

$$U = 1 + i\varepsilon_x \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} + i\varepsilon_y \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} + i\varepsilon_z \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$
 (1.33)

the three terms of order ε in the expansion can be interpreted as a bit flip operator, a phase flip operator, and an operator in which both a bit flip and a phase flip occur. If we prepare an encoded state $a|\bar{0}\rangle + b|\bar{1}\rangle$, allow the unitary errors to occur on each qubit, and then measure the bit-flip and phase-flip syndromes, then most of the time we will project the state back to its original form, but with a probability of order $|\varepsilon|^2$, one qubit will have a large error: a bit flip, a phase flip, or both. From the syndrome, we learn which bit flipped, and which cluster had a phase error, so we can apply the suitable one-qubit unitary operator to fix the error.

Error recovery will fail if, after the syndrome measurement, there are two bit flip errors in each of two clusters (which induces a phase error in the encoded data) or if phase errors occur in two different clusters (which induces a bit-flip error in the encoded data). But the probability of such a double phase error is of order $|\varepsilon|^4$. So for $|\varepsilon|$ small enough, coding improves the reliability of the quantum information.

The code also protects against decoherence. By restoring the quantum state irrespective of the nature of the error, our procedure removes any entanglement between the quantum state and the environment.

Here as always, error correction is a dissipative process, since information about the nature of the errors is flushed out of the quantum system. In this case, that information resides in our recorded measurement results, and heat will be dissipated when that record is erased.

Further developments in quantum error correction will be discussed later in the course, including:

• As with classical coding it turns out that there are "good" quantum codes that allow us to achieve arbitrarily high reliability as long as the error rate per qubit is small enough.

• We've assumed that the error recovery procedure is itself executed flawlessly. But the syndrome measurement was complicated – we needed to measure two-qubit and six-qubit collective observables to diagnose the errors – so we actually might further damage the data when we try to correct it. We'll show, though, that error correction can be carried out so that it still works effectively even if we make occasional errors during the recovery process.

• To operate a quantum computer we'll want not only to store quantum information reliably, but also to process it. We'll show that it is possible to apply quantum gates to encoded information.

Let's summarize the essential ideas that underlie our quantum error correction scheme:

- 1. We *digitized* the errors. Although the errors in the quantum information were small, we performed measurements that projected our state onto either a state with no error, or a state with one of a discrete set of errors that we knew how to convert.
- 2. We measured the errors without measuring the data. Our measurements revealed the nature of the errors without revealing (and hence disturbing) the encoded information.
- 3. The errors are local, and the encoded information is nonlocal. It is important to emphasize the central assumption underlying the construction of the code that errors affecting different qubits are, to a good approximation, *uncorrelated*. We have tacitly assumed that an event that causes errors in two qubits is much less likely than an event causing an error in a single qubit. It is of course a physics question whether this assumption is justified or not we can easily envision processes that will cause errors in two qubits at once. If such correlated errors are common, coding will fail to improve reliability.

The code takes advantage of the presumed local nature of the errors by encoding the information in a nonlocal way - that is the information is stored in *correlations* involving several qubits. There is no way to distinguish $|\bar{0}\rangle$ and $|\bar{1}\rangle$ by measuring a single qubit of the nine. If we measure one qubit we will find $|0\rangle$ with probability $\frac{1}{2}$ and $|1\rangle$ with probability $\frac{1}{2}$ irrespective of the value of the encoded qubit. To access the encoded information we need to measure a 3-qubit observable (the operator that flips all three qubits in a cluster can distinguish $|000\rangle + |111\rangle$ from $|000\rangle - |111\rangle$).

The environment might occasionally kick one of the qubits, in effect "measuring" it. But the encoded information cannot be damaged by disturbing that one qubit, because a single qubit, by itself, actually carries no information at all. Nonlocally encoded information is invulnerable to local influences – this is the central principle on which quantum error-correcting codes are founded.

1.9 Quantum hardware

The theoretical developments concerning quantum complexity and quantum error correction have been accompanied by a burgeoning experimental effort to process coherent quantum information. I'll briefly describe some of this activity here.

To build hardware for a quantum computer, we'll need technology that enables us to manipulate qubits. The hardware will need to meet some stringent specifications:

- 1. **Storage**: We'll need to store qubits for a long time, long enough to complete an interesting computation.
- 2. **Isolation**: The qubits must be well isolated from the environment, to minimize decoherence errors.
- 3. Readout: We'll need to measure the qubits efficiently and reliably.
- 4. Gates: We'll need to manipulate the quantum states of individual qubits, and to induce controlled interactions among qubits, so that we can perform quantum gates.
- 5. **Precision**: The quantum gates should be implemented with high precision if the device is to perform reliably.

1.9.1 Ion Trap

One possible way to achieve these goals was suggested by Ignacio Cirac and Peter Zoller, and has been pursued by Dave Wineland's group at the National Institute for Standards and Technology (NIST), as well as other groups. In this scheme, each qubit is carried by a single ion held in a linear Paul trap. The quantum state of each ion is a linear combination of the ground state $|g\rangle$ (interpreted as $|0\rangle$) and a particular long-lived metastable excited state $|e\rangle$ (interpreted as $|1\rangle$). A coherent linear combination of the two levels,

$$a|g\rangle + be^{i\omega t}|e\rangle, \tag{1.34}$$

can survive for a time comparable to the lifetime of the excited state (though of course the relative phase oscillates as shown because of the energy splitting $\hbar\omega$ between the levels). The ions are so well isolated that spontaneous decay can be the dominant form of decoherence.

It is easy to read out the ions by performing a measurement that projects onto the $\{|g\rangle, |e\rangle\}$ basis. A laser is tuned to a transition from the state $|g\rangle$ to a short-lived excited state $|e'\rangle$. When the laser illuminates the ions, each qubit with the value $|0\rangle$ repeatedly absorbs and reemits the laser light, so that it flows visibly (fluoresces). Qubits with the value $|1\rangle$ remain dark.

Because of their mutual Coulomb repulsion, the ions are sufficiently well separated that they can be individually addressed by pulsed lasers. If a laser is tuned to the frequency ω of the transition and is focused on the *n*th ion, then Rabi oscillations are induced between $|0\rangle$ and $|1\rangle$. By timing the laser pulse properly and choosing the phase of the laser appropriately, we can apply any one-qubit unitary transformation. In particular, acting on $|0\rangle$, the laser pulse can prepare any desired linear combination of $|0\rangle$ and $|1\rangle$.

But the most difficult part of designing and building quantum computing hardware is getting two qubits to interact with one another. In the ion trap, interactions arise because of the Coulomb repulsion between the ions. Because of the mutual Couloumb repulsion, there is a spectrum of coupled normal modes of vibration for the trapped ions. When the ion absorbs or emits a laser photon, the center of mass of the ion recoils. But if the laser is properly tuned, then when a single ion absorbs or emits, a normal mode involving many ions will recoil coherently (the Mössbauer effect).

The vibrational mode of lowest frequency (frequency ν) is the center-ofmass (cm) mode, in which the ions oscillate in lockstep in the harmonic well of the trap. The ions can be laser cooled to a temperature much less than ν , so that each vibrational mode is very likely to occupy its quantum-mechanical ground state. Now imagine that a laser tuned to the frequency $\omega - \nu$ shines on the *n*th ion. For a properly time pulse the state $|e\rangle_n$ will rotate to $|g\rangle_n$, while the *cm* oscillator makes a transition from its ground state $|0\rangle_{cm}$ to its first excited state $|1\rangle_{cm}$ (a *cm* "phonon" is produced). However, the state $|g\rangle_n |0\rangle_{cm}$ is not on resonance for any transition and so is unaffected by the pulse. Thus the laser pulse induces a unitary transformation acting as

$$|g\rangle_n |0\rangle_{cm} \rightarrow |g\rangle_n |0\rangle_{cm}, |e\rangle_n |0\rangle_{cm} \rightarrow -i|g\rangle_n |1\rangle_{cm}.$$

$$(1.35)$$

This operation removes a bit of information that is initially stored in the internal state of the nth ion, and deposits that bit in the *collective* state of motion of *all* the ions.

This means that the state of motion of the *m*th ion $(m \neq n)$ has been influenced by the internal state of the *n*th ion. In this sense, we have succeeded in inducing an interaction between the ions. To complete the quantum gate, we should transfer the quantum information from the *cm* phonon back to

1.9. QUANTUM HARDWARE

the internal state of one of the ions. The procedure should be designed so that the cm mode always returns to its ground state $|0\rangle_{cm}$ at the conclusion of the gate implementation. For example, Cirac and Zoller showed that the quantum XOR (or controlled not) gate

$$|x,y\rangle \to |x,y\oplus x\rangle,$$
 (1.36)

can be implemented in an ion trap with altogether 5 laser pulses. The conditional excitation of a phonon, Eq. (1.35) has been demonstrated experimentally, for a single trapped ion, by the NIST group.

One big drawback of the ion trap computer is that it is an intrinsically slow device. Its speed is ultimately limited by the energy-time uncertainty relation. Since the uncertainty in the energy of the laser photons should be small compared to the characteristic vibrational splitting ν , each laser pulse should last a time long compared to ν^{-1} . In practice, ν is likely to be of order 100 kHz.

1.9.2 Cavity QED

An alternative hardware design (suggested by Pellizzari, Gardiner, Cirac, and Zoller) is being pursued by Jeff Kimble's group here at Caltech. The idea is to trap several neutral atoms inside a small high finesse optical cavity. Quantum information can again be stored in the internal states of the atoms. But here the atoms interact because they all couple to the normal modes of the electromagnetic field in the cavity (instead of the vibrational modes as in the ion trap). Again, by driving transitions with pulsed lasers, we can induce a transition in one atom that is conditioned on the internal state of another atom.

Another possibility is to store a qubit, not in the internal state of an ion, but in the polarization of a photon. Then a trapped atom can be used as the intermediary that causes one photon to interact with another (instead of a photon being used to couple one atom to another). In their "flying qubit" experiment two years ago. The Kimble group demonstrated the operation of a two-photon quantum gate, in which the circular polarization of one photon influences the *phase* of another photon:

$$|L\rangle_{1}|L\rangle_{2} \rightarrow |L\rangle_{1}|L\rangle_{2}$$

$$|L\rangle_{1}|R\rangle_{2} \rightarrow |L\rangle_{1}|R\rangle_{2}$$

$$|R\rangle_{1}|L\rangle_{2} \rightarrow |R\rangle_{1}|L\rangle_{2}$$

$$|R\rangle_{1}|R\rangle_{2} \rightarrow e^{i\Delta}|R\rangle_{1}|R\rangle_{2}$$
(1.37)

where $|L\rangle$, $|R\rangle$ denote photon states with left and right circular polarization. To achieve this interaction, one photon is stored in the cavity, where the $|L\rangle$ polarization does not couple to the atom, but the $|R\rangle$ polarization couples strongly. A second photon transverses the cavity, and for the second photon as well, one polarization interacts with the atom preferentially. The second photon wave pocket acquires a particular phase shift $e^{i\Delta}$ only if both photons have $|R\rangle$ polarization. Because the phase shift is conditioned on the polarization of *both* photons, this is a nontrivial two-qubit quantum gate.

1.9.3 NMR

A third (dark horse) hardware scheme has sprung up in the past year, and has leap frogged over the ion trap and cavity QED to take the current lead in coherent quantum processing. The new scheme uses nuclear magnetic resonance (NMR) technology. Now qubits are carried by certain nuclear spins in a particular molecule. Each spin can either be aligned ($|\uparrow\rangle = |0\rangle$) or antialigned ($|\downarrow\rangle = |1\rangle$) with an applied constant magnetic field. The spins take a long time to relax or decohere, so the qubits can be stored for a reasonable time.

We can also turn on a pulsed rotating magnetic field with frequency ω (where the ω is the energy splitting between the spin-up and spin-down states), and induce Rabi oscillations of the spin. By timing the pulse suitably, we can perform a desired unitary transformation on a single spin (just as in our discussion of the ion trap). All the spins in the molecule are exposed to the rotating magnetic field but only those on resonance respond.

Furthermore, the spins have dipole-dipole interactions, and this coupling can be exploited to perform a gate. The splitting between $|\uparrow\rangle$ and $|\downarrow\rangle$ for one spin actually depends on the state of neighboring spins. So whether a driving pulse is on resonance to tip the spin over is conditioned on the state of another spin.

1.9. QUANTUM HARDWARE

All this has been known to chemists for decades. Yet it was only in the past year that Gershenfeld and Chuang, and independently Cory, Fahmy, and Havel, pointed out that NMR provides a useful implementation of quantum computation. This was not obvious for several reasons. Most importantly, NMR systems are very hot. The typical temperature of the spins (room temperature, say) might be of order a million times larger than the energy splitting between $|0\rangle$ and $|1\rangle$. This means that the quantum state of our computer (the spins in a single molecule) is very noisy - it is subject to strong random thermal fluctuations. This noise will disguise the quantum information. Furthermore, we actually perform our processing not on a single molecule, but on a macroscopic sample containing of order 10^{23} "computers," and the signal we read out of this device is actually averaged over this ensemble. But quantum algorithms are *probabilistic*, because of the randomness of quantum measurement. Hence averaging over the ensemble is not equivalent to running the computation on a single device; averaging may obscure the results.

Gershenfeld and Chuang and Cory, Fahmy, and Havel, explained how to overcome these difficulties. They described how "effective pure states" can be prepared, manipulated, and monitored by performing suitable operations on the thermal ensemble. The idea is to arrange for the fluctuating properties of the molecule to average out when the signal is detected, so that only the underlying coherent properties are measured. They also pointed out that some quantum algorithms (including Shor's factoring algorithm) can be cast in a deterministic form (so that at least a large fraction of the computers give the same answer); then averaging over many computations will not spoil the result.

Quite recently, NMR methods have been used to prepare a maximally entangled state of three qubits, which had never been achieved before.

Clearly, quantum computing hardware is in its infancy. Existing hardware will need to be scaled up by many orders of magnitude (both in the number of stored qubits, and the number of gates that can be applied) before ambitious computations can be attempted. In the case of the NMR method, there is a particularly serious limitation that arises as a matter of principle, because the ratio of the coherent signal to the background declines exponentially with the number of spins per molecule. In practice, it will be very challenging to perform an NMR quantum computation with more than of order 10 qubits.

Probably, if quantum computers are eventually to become practical devices, new ideas about how to construct quantum hardware will be needed.

1.10 Summary

This concludes our introductory overview to quantum computation. We have seen that three converging factors have combined to make this subject exciting.

- 1. Quantum computers can solve hard problems. It seems that a new classification of complexity has been erected, a classification better founded on the fundamental laws of physics than traditional complexity theory. (But it remains to characterize more precisely the class of problems for which quantum computers have a big advantage over classical computers.)
- 2. Quantum errors can be corrected. With suitable coding methods, we can protect a complicated quantum system from the debilitating effects of decoherence. We may never see an actual cat that is half dead and half alive, but perhaps we can prepare and preserve an *encoded cat* that is half dead and half alive.
- 3. Quantum hardware can be constructed. We are privileged to be witnessing the dawn of the age of coherent manipulation of quantum information in the laboratory.

Our aim, in this course, will be to deepen our understanding of points (1), (2), and (3).

Lecture Notes for Ph219/CS219: Quantum Information Chapter 2

John Preskill California Institute of Technology

Updated July 2015

Contents

2	Foundations I: States and Ensembles	3
2.1	Axioms of quantum mechanics	3
2.2	The Qubit	7
	2.2.1 Spin- $\frac{1}{2}$	8
	2.2.2 Photon polarizations	14
2.3	The density operator	16
	2.3.1 The bipartite quantum system	16
	2.3.2 Bloch sphere	21
2.4	Schmidt decomposition	23
	2.4.1 Entanglement	25
2.5	Ambiguity of the ensemble interpretation	26
	2.5.1 Convexity	26
	2.5.2 Ensemble preparation	28
	2.5.3 Faster than light?	30
	2.5.4 Quantum erasure	31
	2.5.5 The HJW theorem	33
2.6	How far apart are two quantum states?	36
	2.6.1 Fidelity and Uhlmann's theorem	36
	2.6.2 Relations among distance measures	38
2.7	Summary	41
2.8	Exercises	43

Foundations I: States and Ensembles

2.1 Axioms of quantum mechanics

In this chapter and the next we develop the theory of *open* quantum systems. We say a system is open if it is imperfectly isolated, and therefore exchanges energy and information with its unobserved environment. The motivation for studying open systems is that all realistic systems are open. Physicists and engineers may try hard to isolate quantum systems, but they never completely succeed.

Though our main interest is in open systems we will begin by recalling the theory of closed quantum systems, which *are* perfectly isolated. To understand the behavior of an open system S, we will regard S combined with its environment E as a closed system (the whole "universe"), then ask how S behaves when we are able to observe S but not E.

Quantum theory is a mathematical model of the physical world. For the case of closed systems we can characterize the model by stating five axioms; these specify how to represent states, observables, measurements, and dynamics, and also how to combine two systems to obtain a composite system.

Axiom 1. States. A state is a complete description of a physical system. In quantum mechanics, a state is a *ray* in a *Hilbert space*.

What is a Hilbert space?

- a) It is a vector space over the complex numbers \mathbb{C} . Vectors will be denoted $|\psi\rangle$ (Dirac's ket notation).
- b) It has an *inner product* $\langle \psi | \varphi \rangle$ that maps an ordered pair of vectors to \mathbb{C} , and that has the properties:
 - i) Positivity: $\langle \psi | \psi \rangle > 0$ for $| \psi \rangle \neq 0$.

ii) Linearity: $\langle \varphi | (a | \psi_1 \rangle + b | \psi_2 \rangle) = a \langle \varphi | \psi_1 \rangle + b \langle \varphi | \psi_2 \rangle.$

iii) Skew symmetry: $\langle \varphi | \psi \rangle = \langle \psi | \varphi \rangle^*$.

(The * denotes complex conjugation.)

c) It is complete in the norm $||\psi|| = \langle \psi |\psi \rangle^{1/2}$.

(Completeness is an important proviso in infinite-dimensional function spaces, since it ensures the convergence of certain eigenfunction expansions. But mostly we will be content to work with finite-dimensional inner-product spaces.)

What is a ray? It is an equivalence class of vectors that differ by multiplication by a nonzero complex scalar. For any nonzero ray, we can by convention choose a representative of the class, denoted $|\psi\rangle$, that has unit norm:

$$\langle \psi | \psi \rangle = 1. \tag{2.1}$$

Thus states correspond to normalized vectors, and the overall phase of the vector has no physical significance: $|\psi\rangle$ and $e^{i\alpha}|\psi\rangle$ describe the same state, where $|e^{i\alpha}| = 1$.

Since every ray corresponds to a possible state, given two states $|\varphi\rangle, |\psi\rangle$, another state can be constructed as the linear superposition of the two, $a|\varphi\rangle + b|\psi\rangle$. The relative phase in this superposition is physically significant; we identify $a|\varphi\rangle + b|\varphi\rangle$ with $e^{i\alpha}(a|\varphi\rangle + b|\psi\rangle)$ but not with $a|\varphi\rangle + e^{i\alpha}b|\psi\rangle$.

We use the notation $\langle \psi |$ (Dirac's bra notation) for a linear function (a *dual vector*) that takes vectors to complex numbers, defined by $|\varphi\rangle \rightarrow \langle \psi |\varphi\rangle$.

Axiom 2. Observables. An observable is a property of a physical system that in principle can be measured. In quantum mechanics, an observable is a *self-adjoint operator*.

An operator is a linear map taking vectors to vectors,

$$\boldsymbol{A}:|\psi\rangle\mapsto\boldsymbol{A}|\psi\rangle,\quad\boldsymbol{A}(a|\psi\rangle+b|\varphi\rangle)=a\boldsymbol{A}|\psi\rangle+b\boldsymbol{A}|\varphi\rangle.$$
(2.2)

(We will often denote operators by boldface letters.) The adjoint A^{\dagger} of the operator A is defined by

$$\langle \varphi | \mathbf{A} \psi \rangle = \langle \mathbf{A}^{\dagger} \varphi | \psi \rangle, \qquad (2.3)$$

for all vectors $|\varphi\rangle, |\psi\rangle$ (where here $\boldsymbol{A}|\psi\rangle$ has been denoted as $|\boldsymbol{A}\psi\rangle$). \boldsymbol{A} is self-adjoint if $\boldsymbol{A} = \boldsymbol{A}^{\dagger}$, or in other words, if $\langle \varphi | \boldsymbol{A} | \psi \rangle = \langle \psi | \boldsymbol{A} | \varphi \rangle^*$ for all vectors $|\varphi\rangle$ and $|\psi\rangle$. If \boldsymbol{A} and \boldsymbol{B} are self adjoint, then so is $\boldsymbol{A} + \boldsymbol{B}$ (because $(\boldsymbol{A} + \boldsymbol{B})^{\dagger} = \boldsymbol{A}^{\dagger} + \boldsymbol{B}^{\dagger}$), but $(\boldsymbol{A}\boldsymbol{B})^{\dagger} = \boldsymbol{B}^{\dagger}\boldsymbol{A}^{\dagger}$, so that $\boldsymbol{A}\boldsymbol{B}$ is self adjoint

only if A and B commute. Note that AB + BA and i(AB - BA) are always self-adjoint if A and B are.

A self-adjoint operator in a Hilbert space \mathcal{H} has a spectral representation – its eigenstates form a complete orthonormal basis in \mathcal{H} . We can express a self-adjoint operator A as

$$\boldsymbol{A} = \sum_{n} a_n \boldsymbol{E}_n. \tag{2.4}$$

Here each a_n is an eigenvalue of A, and E_n is the corresponding orthogonal projection onto the space of eigenvectors with eigenvalue a_n . The E_n 's satisfy

$$\begin{aligned} \boldsymbol{E}_{n} \boldsymbol{E}_{m} &= \delta_{n,m} \boldsymbol{E}_{n}. \\ \boldsymbol{E}_{n}^{\dagger} &= \boldsymbol{E}_{n}. \end{aligned}$$
 (2.5)

The orthogonal projector onto the one-dimensional space spanned by the vector $|\psi\rangle$ may be expressed as $|\psi\rangle\langle\psi|$, where $\langle\psi|$ is the bra that annihilates vectors orthogonal to $|\psi\rangle$. Therefore, an alternative notation for the spectral representation of \boldsymbol{A} is

$$\boldsymbol{A} = \sum_{n} |n\rangle a_n \langle n|, \qquad (2.6)$$

where $\{|n\rangle\}$ is the orthonormal basis of eigenstates of A, with $A|n\rangle = a_n|n\rangle$.

(For unbounded operators in an infinite-dimensional space, the definition of self-adjoint and the statement of the spectral theorem are more subtle, but this need not concern us.)

Axiom 3. Measurement. A measurement is a process in which information about the state of a physical system is acquired by an observer. In quantum mechanics, the measurement of an observable \boldsymbol{A} prepares an eigenstate of \boldsymbol{A} , and the observer learns the value of the corresponding eigenvalue. If the quantum state just prior to the measurement is $|\psi\rangle$, then the outcome a_n is obtained with a priori probability

$$\operatorname{Prob}(a_n) = \|\boldsymbol{E}_n|\psi\rangle\|^2 = \langle\psi|\boldsymbol{E}_n|\psi\rangle ; \qquad (2.7)$$

if the outcome a_n is attained, then the (normalized) quantum state just after the measurement is

$$\frac{\boldsymbol{E}_n|\psi\rangle}{|\boldsymbol{E}_n|\psi\rangle\|}.$$
(2.8)

If the measurement is immediately repeated, then according to this rule the same outcome is obtained again, with probability one. If many identically prepared systems are measured, each described by the state $|\psi\rangle$, then the *expectation value* of the outcomes is

$$\langle a \rangle \equiv \sum_{n} a_{n} \operatorname{Prob}(a_{n}) = \sum_{n} a_{n} \langle \psi | \boldsymbol{E}_{n} | \psi \rangle = \langle \psi | \boldsymbol{A} | \psi \rangle.$$
 (2.9)

Axiom 4. Dynamics. Dynamics describes how a state evolves over time. In quantum mechanics, the time evolution of a closed system is described by a *unitary operator*.

In the Schrödinger picture of dynamics, if the initial state at time t is $|\psi(t)\rangle$, then the final state $|\psi(t')\rangle$ at time t' can be expressed as

$$|\psi(t')\rangle = U(t',t)|\psi(t)\rangle|\psi(t)\rangle, \qquad (2.10)$$

where U(t',t) is the unitary time evolution operator. Infinitesimal time evolution is governed by the *Schrödinger equation*

$$\frac{d}{dt}|\psi(t)\rangle = -i\boldsymbol{H}(t)|\psi(t)\rangle, \qquad (2.11)$$

where H(t) is a self-adjoint operator, called the *Hamiltonian* of the system. (The Hamiltonian has the dimensions of energy; we have chosen units in which Planck's constant $\hbar = h/2\pi = 1$, so that energy has the dimensions of inverse time.) To first order in the infinitesimal quantity dt, the Schrödinger equation can be expressed as

$$|\psi(t+dt)\rangle = (\mathbf{I} - i\mathbf{H}(t)dt)|\psi(t)\rangle.$$
(2.12)

Thus the operator $U(t + dt, t) \equiv I - iH(t)dt$ is unitary; because H is self-adjoint it satisfies $U^{\dagger}U = 1$ to linear order in dt. Since a product of unitary operators is unitary, time evolution governed by the Schrödinger equation over a finite interval is also unitary. In the case where H is time independent we may write $U(t', t) = e^{-i(t'-t)H}$.

Our final axiom relates the description of a composite quantum system AB to the description of its component parts A and B.

Axiom 5. Composite Systems. If the Hilbert space of system A is \mathcal{H}_A and the Hilbert space of system B is \mathcal{H}_B , then the Hilbert space of the composite systems AB is the *tensor product* $\mathcal{H}_A \otimes \mathcal{H}_B$. If system A is prepared in the state $|\psi\rangle_A$ and system B is prepared in the state $|\psi\rangle_A$ and system B is prepared in the state $|\psi\rangle_A$.
What is a tensor product of Hilbert spaces? If $\{|i\rangle_A\}$ denotes an orthonormal basis for \mathcal{H}_A and $\{|\mu\rangle_B\}$ a basis for \mathcal{H}_B , then the states $|i,\mu\rangle_{AB} \equiv |i\rangle_A \otimes |\mu\rangle_B$ are a basis for $\mathcal{H}_A \otimes \mathcal{H}_B$, where the inner product on $\mathcal{H}_A \otimes \mathcal{H}_B$ is defined by

$$_{AB}\langle i,\mu|j,\nu\rangle_{AB} = \delta_{ij}\delta_{\mu\nu}.$$
(2.13)

The tensor product operator $M_A \otimes N_B$ is the operator that applies M_A to system A and N_B to system B. Its action on the orthonormal basis $|i, \mu\rangle_{AB}$ is

$$\boldsymbol{M}_{A} \otimes \boldsymbol{N}_{B}|i,\mu\rangle_{AB} = \boldsymbol{M}_{A}|i\rangle_{A} \otimes \boldsymbol{N}_{B}|\mu\rangle_{B} = \sum_{j,\nu}|j,\nu\rangle_{AB} (M_{A})_{ji} (N_{B})_{\nu\mu}.$$
(2.14)

An operator that acts trivially on system B can be denoted $M_A \otimes I_B$, where I_B is the identity on \mathcal{H}_B , and an operator that acts trivially on system A can be denoted $I_A \otimes N_B$.

These five axioms provide a complete mathematical formulation of quantum mechanics. We immediately notice some curious features. One oddity is that the Schrödinger equation is linear, while we are accustomed to nonlinear dynamical equations in classical physics. This property seems to beg for an explanation. But far more curious is a mysterious dualism; there are two quite distinct ways for a quantum state to change. On the one hand there is unitary evolution, which is deterministic. If we specify the initial state $|\psi(0)\rangle$, the theory predicts the state $|\psi(t)\rangle$ at a later time.

But on the other hand there is measurement, which is probabilistic. The theory does not make definite predictions about the measurement outcomes; it only assigns probabilities to the various alternatives. This is troubling, because it is unclear why the measurement process should be governed by different physical laws than other processes.

The fundamental distinction between evolution and measurement, and in particular the intrinsic randomness of the measurement process, is sometimes called the *measurement problem* of quantum theory. Seeking a more pleasing axiomatic formulation of quantum theory is a worthy task which may eventually succeed. But these five axioms correctly account for all that we currently know about quantum physics, and provide the foundation for all that follows in this book.

2.2 The Qubit

The indivisible unit of classical information is the *bit*, which takes one of the two possible values $\{0, 1\}$. The corresponding unit of quantum information is called the "quantum bit" or *qubit*. It describes a state in the simplest possible quantum system.

The smallest nontrivial Hilbert space is two-dimensional. We may denote an orthonormal basis for a two-dimensional vector space as $\{|0\rangle, |1\rangle\}$. Then the most general normalized state can be expressed as

$$a|0\rangle + b|1\rangle, \tag{2.15}$$

where a, b are complex numbers that satisfy $|a|^2 + |b|^2 = 1$, and the overall phase is physically irrelevant. A *qubit* is a quantum system described by a two-dimensional Hilbert space, whose state can take any value of the form eq.(2.15).

We can perform a measurement that projects the qubit onto the basis $\{|0\rangle, |1\rangle\}$. Then we will obtain the outcome $|0\rangle$ with probability $|a|^2$, and the outcome $|1\rangle$ with probability $|b|^2$. Furthermore, except in the cases a = 0 and b = 0, the measurement irrevocably disturbs the state. If the value of the qubit is initially unknown, then there is no way to determine a and b with that single measurement, or any other conceivable measurement. However, *after* the measurement, the qubit has been prepared in a *known* state – either $|0\rangle$ or $|1\rangle$ – that differs (in general) from its previous state.

In this respect, a qubit differs from a classical bit; we can measure a classical bit without disturbing it, and we can decipher all of the information that it encodes. But suppose we have a classical bit that really does have a definite value (either 0 or 1), but where that value is initially unknown to us. Based on the information available to us we can only say that there is a *probability* p_0 that the bit has the value 0, and a probability p_1 that the bit has the value 1, where $p_0 + p_1 = 1$. When we measure the bit, we acquire additional information; afterwards we know the value with 100% confidence.

An important question is: what is the essential difference between a qubit and a *probabilistic* classical bit? In fact they are *not* the same, for several reasons that we will explore. To summarize the difference in brief: there is only one way to look at a bit, but there is more than one way to look at a qubit.

2.2.1 Spin- $\frac{1}{2}$

First of all, the coefficients a and b in eq.(2.15) encode more than just the probabilities of the outcomes of a measurement in the $\{|0\rangle, |1\rangle\}$ basis. In particular, the *relative phase* of a and b also has physical significance.

The properties of a qubit are easier to grasp if we appeal to a geometrical interpretation of its state. For a physicist, it is natural to interpret eq.(2.15) as the spin state of an object with spin- $\frac{1}{2}$ (like an electron). Then $|0\rangle$ and $|1\rangle$ are the spin up ($|\uparrow\rangle$) and spin down ($|\downarrow\rangle$) states along

a particular axis such as the z-axis. The two real numbers characterizing the qubit (the complex numbers a and b, modulo the normalization and overall phase) describe the *orientation* of the spin in three-dimensional space (the polar angle θ and the azimuthal angle φ).

We will not go deeply here into the theory of symmetry in quantum mechanics, but we will briefly recall some elements of the theory that will prove useful to us. A symmetry is a transformation that acts on a state of a system, yet leaves all observable properties of the system unchanged. In quantum mechanics, observations are measurements of self-adjoint operators. If A is measured in the state $|\psi\rangle$, then the outcome $|a\rangle$ (an eigenvector of A) occurs with probability $|\langle a|\psi\rangle|^2$. A symmetry should leave these probabilities unchanged, when we "rotate" both the system and the apparatus.

A symmetry, then, is a mapping of vectors in Hilbert space

$$|\psi\rangle \mapsto |\psi'\rangle,$$
 (2.16)

that preserves the absolute values of inner products

$$|\langle \varphi | \psi \rangle| = |\langle \varphi' | \psi' \rangle|, \qquad (2.17)$$

for all $|\varphi\rangle$ and $|\psi\rangle$. According to a famous theorem due to Wigner, a mapping with this property can always be chosen (by adopting suitable phase conventions) to be either unitary or antiunitary. The antiunitary alternative, while important for discrete symmetries, can be excluded for continuous symmetries. Then the symmetry acts as

$$|\psi\rangle \mapsto |\psi'\rangle = U|\psi\rangle,$$
 (2.18)

where U is unitary (and in particular, *linear*).

Symmetries form a group: a symmetry transformation can be inverted, and the product of two symmetries is a symmetry. For each symmetry operation R acting on our physical system, there is a corresponding unitary transformation U(R). Multiplication of these unitary operators must respect the group multiplication law of the symmetries – applying $R_1 \circ R_2$ should be equivalent to first applying R_2 and subsequently R_1 . Thus we demand

$$\boldsymbol{U}(R_1)\boldsymbol{U}(R_2) = \text{Phase}(R_1, R_2) \cdot \boldsymbol{U}(R_1 \circ R_2)$$
(2.19)

A phase depending on R_1 and R_2 is permitted in eq.(2.19) because quantum states are *rays*; we need only demand that $U(R_1 \circ R_2)$ act the same way as $U(R_1)U(R_2)$ on rays, not on vectors. We say that U(R) provides a unitary representation, up to a phase, of the symmetry group.

So far, our concept of symmetry has no connection with dynamics. Usually, we demand of a symmetry that it respect the dynamical evolution of the system. This means that it should not matter whether we first transform the system and then evolve it, or first evolve it and then transform it. In other words, the diagram



is commutative, and therefore the time evolution operator e^{itH} commutes with the symmetry transformation U(R):

$$\boldsymbol{U}(R)e^{-it\boldsymbol{H}} = e^{-it\boldsymbol{H}}\boldsymbol{U}(R) ; \qquad (2.20)$$

expanding to linear order in t we obtain

$$\boldsymbol{U}(R)\boldsymbol{H} = \boldsymbol{H}\boldsymbol{U}(R). \tag{2.21}$$

For a continuous symmetry, we can choose R infinitesimally close to the identity, $R = I + \epsilon T$, and then U is close to I:

$$\boldsymbol{U} = \boldsymbol{I} - i\varepsilon \boldsymbol{Q} + O(\varepsilon^2), \qquad (2.22)$$

where Q is an operator determined by T. From the unitarity of U (to order ε) it follows that Q is an observable, $Q = Q^{\dagger}$. Expanding eq.(2.21) to linear order in ε we find

$$[Q, H] = 0;$$
 (2.23)

the observable \boldsymbol{Q} commutes with the Hamiltonian.

Eq.(2.23) is a conservation law. It says, for example, that if we prepare an eigenstate of Q, then time evolution governed by the Schrödinger equation will preserve the eigenstate. Thus we see that symmetries imply conservation laws. Conversely, given a conserved quantity Q satisfying eq.(2.23) we can construct the corresponding symmetry transformations. Finite transformations can be built as a product of many infinitesimal ones:

$$R = (1 + \frac{\theta}{N}T)^N \Rightarrow \boldsymbol{U}(R) = (\boldsymbol{I} + i\frac{\theta}{N}\boldsymbol{Q})^N \to e^{i\theta\boldsymbol{Q}}, \qquad (2.24)$$

taking the limit $N \to \infty$. Once we have decided how infinitesimal symmetry transformations are represented by unitary operators, then it is also

11

determined how finite transformations are represented, for these can be built as a product of infinitesimal transformations. We say that Q is the *generator* of the symmetry.

Let us briefly recall how this general theory applies to spatial rotations and angular momentum. An infinitesimal rotation by $d\theta$ (in the counterclockwise sense) about the axis specified by the unit vector $\hat{n} = (n_1, n_2, n_3)$ can be expressed as

$$R(\hat{n}, d\theta) = I - id\theta \hat{n} \cdot \vec{J}, \qquad (2.25)$$

where (J_1, J_2, J_3) are the components of the angular momentum. A finite rotation is expressed as

$$R(\hat{n},\theta) = \exp(-i\theta\hat{n}\cdot\vec{J}). \tag{2.26}$$

Rotations about distinct axes don't commute. From elementary properties of rotations, we find the commutation relations

$$[J_k, J_\ell] = i\varepsilon_{k\ell m} J_m, \qquad (2.27)$$

where $\varepsilon_{k\ell m}$ is the totally antisymmetric tensor with $\varepsilon_{123} = 1$, and repeated indices are summed. To implement rotations on a quantum system, we find self-adjoint operators J_1, J_2, J_3 in Hilbert space that satisfy these relations.

The "defining" representation of the rotation group is three dimensional, but the simplest nontrivial irreducible representation is two dimensional, given by

$$\boldsymbol{J}_k = \frac{1}{2}\boldsymbol{\sigma}_k, \qquad (2.28)$$

where

$$\boldsymbol{\sigma}_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \boldsymbol{\sigma}_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \boldsymbol{\sigma}_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (2.29)$$

are the Pauli matrices. This is the unique two-dimensional irreducible representation, up to a unitary change of basis. Since the eigenvalues of J_k are $\pm \frac{1}{2}$, we call this the spin- $\frac{1}{2}$ representation. (By identifying J as the angular-momentum, we have implicitly chosen units with $\hbar = 1$.)

The Pauli matrices also have the properties of being mutually anticommuting and squaring to the identity,

$$\boldsymbol{\sigma}_k \boldsymbol{\sigma}_\ell + \boldsymbol{\sigma}_\ell \boldsymbol{\sigma}_k = 2\delta_{k\ell} \boldsymbol{I} ; \qquad (2.30)$$

therefore $(\hat{n} \cdot \vec{\sigma})^2 = n_k n_\ell \sigma_k \sigma_\ell = n_k n_k I = I$ (where repeated indices are summed). By expanding the exponential series, we see that finite rotations are represented as

$$\boldsymbol{U}(\hat{n},\theta) = e^{-i\frac{\theta}{2}\hat{n}\cdot\vec{\boldsymbol{\sigma}}} = \boldsymbol{I}\cos\frac{\theta}{2} - i\hat{n}\cdot\vec{\boldsymbol{\sigma}}\sin\frac{\theta}{2}.$$
 (2.31)

The most general 2×2 unitary matrix with determinant 1 can be expressed in this form. Thus, we are entitled to think of a qubit as a spin- $\frac{1}{2}$ object, and an arbitrary unitary transformation acting on the qubit's state (aside from a possible physically irrelevant rotation of the overall phase) is a *rotation* of the spin.

A peculiar property of the representation $U(\hat{n}, \theta)$ is that it is *double-valued*. In particular a rotation by 2π about any axis is represented non-trivially:

$$\boldsymbol{U}(\hat{n}, \theta = 2\pi) = -\boldsymbol{I}.\tag{2.32}$$

Our representation of the rotation group is really a representation "up to a sign"

$$\boldsymbol{U}(R_1)\boldsymbol{U}(R_2) = \pm \boldsymbol{U}(R_1 \circ R_2). \tag{2.33}$$

But as already noted, this is acceptable, because the group multiplication is respected on *rays*, though not on vectors. These double-valued representations of the rotation group are called *spinor* representations. (The existence of spinors follows from a topological property of the group that it is not simply connected.)

While it is true that a rotation by 2π has no detectable effect on a spin- $\frac{1}{2}$ object, it would be wrong to conclude that the spinor property has no observable consequences. Suppose I have a machine that acts on a pair of spins. If the first spin is up, it does nothing, but if the first spin is down, it rotates the second spin by 2π . Now let the machine act when the first spin is in a *superposition* of up and down. Then

$$\frac{1}{\sqrt{2}} \left(|\uparrow\rangle_1 + |\downarrow\rangle_1 \right) |\uparrow\rangle_2 \mapsto \frac{1}{\sqrt{2}} \left(|\uparrow\rangle_1 - |\downarrow\rangle_1 \right) |\uparrow\rangle_2 . \tag{2.34}$$

While there is no detectable effect on the second spin, the state of the first has flipped to an orthogonal state, which is very much observable.

In a rotated frame of reference, a rotation $R(\hat{n}, \theta)$ becomes a rotation through the same angle but about a rotated axis. It follows that the three components of angular momentum transform under rotations as a vector:

$$\boldsymbol{U}(R)\boldsymbol{J}_{k}\boldsymbol{U}(R)^{\dagger} = R_{k\ell}\boldsymbol{J}_{\ell}.$$
(2.35)

Thus, if a state $|m\rangle$ is an eigenstate of J_3

$$\boldsymbol{J}_3|m\rangle = m|m\rangle,\tag{2.36}$$

then $U(R)|m\rangle$ is an eigenstate of RJ_3 with the same eigenvalue:

$$RJ_{3}(\boldsymbol{U}(R)|m\rangle) = \boldsymbol{U}(R)J_{3}\boldsymbol{U}(R)^{\dagger}\boldsymbol{U}(R)|m\rangle$$
$$= \boldsymbol{U}(R)J_{3}|m\rangle = m\left(\boldsymbol{U}(R)|m\rangle\right). \quad (2.37)$$

2.2 The Qubit 13

Therefore, we can construct eigenstates of angular momentum along the axis $\hat{n} = (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta)$ by applying a counterclockwise rotation through θ , about the axis $\hat{n}' = (-\sin \varphi, \cos \varphi, 0)$, to a J_3 eigenstate. For our spin- $\frac{1}{2}$ representation, this rotation is

$$\exp\left(-i\frac{\theta}{2}\hat{n}'\cdot\vec{\sigma}\right) = \exp\left[\frac{\theta}{2}\begin{pmatrix}0 & -e^{-i\varphi}\\e^{i\varphi} & 0\end{pmatrix}\right] \\ = \begin{pmatrix}\cos\frac{\theta}{2} & -e^{-i\varphi}\sin\frac{\theta}{2}\\e^{i\varphi}\sin\frac{\theta}{2} & \cos\frac{\theta}{2}\end{pmatrix}, \quad (2.38)$$

and applying it to $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, the J_3 eigenstate with eigenvalue 1, we obtain

$$|\psi(\theta,\varphi)\rangle = \begin{pmatrix} e^{-i\varphi/2}\cos\frac{\theta}{2}\\ e^{i\varphi/2}\sin\frac{\theta}{2} \end{pmatrix}, \qquad (2.39)$$

(up to an overall phase). We can check directly that this is an eigenstate of

$$\hat{n} \cdot \vec{\sigma} = \begin{pmatrix} \cos\theta & e^{-i\varphi}\sin\theta\\ e^{i\varphi}\sin\theta & -\cos\theta \end{pmatrix}, \qquad (2.40)$$

with eigenvalue one. We now see that eq.(2.15) with $a = e^{-i\varphi/2} \cos \frac{\theta}{2}$, $b = e^{i\varphi/2} \sin \frac{\theta}{2}$, can be interpreted as a spin pointing in the (θ, φ) direction.

We noted that we cannot determine a and b with a single measurement. Furthermore, even with many identical copies of the state, we cannot completely determine the state by measuring each copy only along the z-axis. This would enable us to estimate |a| and |b|, but we would learn nothing about the relative phase of a and b. Equivalently, we would find the component of the spin along the z-axis

$$\langle \psi(\theta,\varphi)|\sigma_3|\psi(\theta,\varphi)\rangle = \cos^2\frac{\theta}{2} - \sin^2\frac{\theta}{2} = \cos\theta,$$
 (2.41)

but we would not learn about the component in the x-y plane. The problem of determining $|\psi\rangle$ by measuring the spin is equivalent to determining the unit vector \hat{n} by measuring its components along various axes. Altogether, measurements along three different axes are required. *E.g.*, from $\langle \sigma_3 \rangle$ and $\langle \sigma_1 \rangle$ we can determine n_3 and n_1 , but the sign of n_2 remains undetermined. Measuring $\langle \sigma_2 \rangle$ would remove this remaining ambiguity.

If we are permitted to rotate the spin, then only measurements along the z-axis will suffice. That is, measuring a spin along the \hat{n} axis is equivalent to first applying a rotation that rotates the \hat{n} axis to the axis \hat{z} , and then measuring along \hat{z} .

In the special case $\theta = \frac{\pi}{2}$ and $\varphi = 0$ (the \hat{x} -axis) our spin state is

$$|\uparrow_x\rangle = \frac{1}{\sqrt{2}} \left(|\uparrow_z\rangle + |\downarrow_z\rangle\right) \tag{2.42}$$

("spin-up along the x-axis"). The orthogonal state ("spin down along the x-axis") is

$$|\downarrow_x\rangle = \frac{1}{\sqrt{2}} \left(|\uparrow_z\rangle - |\downarrow_z\rangle\right). \tag{2.43}$$

For either of these states, if we measure the spin along the z-axis, we will obtain $|\uparrow_z\rangle$ with probability $\frac{1}{2}$ and $|\downarrow_z\rangle$ with probability $\frac{1}{2}$.

Now consider the combination

$$\frac{1}{\sqrt{2}}\left(|\uparrow_{x}\rangle + |\downarrow_{x}\rangle\right). \tag{2.44}$$

This state has the property that, if we measure the spin along the x-axis, we obtain $|\uparrow_x\rangle$ or $|\downarrow_x\rangle$, each with probability $\frac{1}{2}$. Now we may ask, what if we measure the state in eq.(2.44) along the z-axis?

If these were probabilistic classical bits, the answer would be obvious. The state in eq.(2.44) is in one of two states, and for *each* of the two, the probability is $\frac{1}{2}$ for pointing up or down along the *z*-axis. So of course we should find up with probability $\frac{1}{2}$ when we measure the state $\frac{1}{\sqrt{2}}(|\uparrow_x\rangle + |\downarrow_x\rangle)$ along the *z*-axis.

But not so for qubits! By adding eq.(2.42) and eq.(2.43), we see that the state in eq.(2.44) is really $|\uparrow_z\rangle$ in disguise. When we measure along the z-axis, we always find $|\uparrow_z\rangle$, never $|\downarrow_z\rangle$.

We see that for qubits, as opposed to probabilistic classical bits, probabilities can add in unexpected ways. This is, in its simplest guise, the phenomenon called "quantum interference," an important feature of quantum information.

To summarize the geometrical interpretation of a qubit: we may think of a qubit as a spin- $\frac{1}{2}$ object, and its quantum state is characterized by a unit vector \hat{n} in three dimensions, the spin's direction. A unitary transformation rotates the spin, and a measurement of an observable has two possible outcomes: the spin is either up or down along a specified axis.

It should be emphasized that, while this *formal* equivalence with a spin- $\frac{1}{2}$ object applies to any two-level quantum system, not every two-level system transforms as a spinor under spatial rotations!

2.2.2 Photon polarizations

Another important two-state system is provided by a *photon*, which can have two independent polarizations. These photon polarization states also transform under rotations, but photons differ from our spin- $\frac{1}{2}$ objects in two important ways: (1) Photons are massless. (2) Photons have spin-1 (they are not spinors).

2.2 The Qubit

We will not present here a detailed discussion of the unitary representations of the Poincare group. Suffice it to say that the *spin* of a particle classifies how it transforms under the *little group*, the subgroup of the Lorentz group that preserves the particle's momentum. For a massive particle, we may always boost to the particle's rest frame, and then the little group is the rotation group.

For massless particles, there is no rest frame. The finite-dimensional unitary representations of the little group turn out to be representations of the rotation group in *two* dimensions, the rotations about the axis determined by the momentum. For a photon, this corresponds to a familiar property of classical light — the waves are polarized transverse to the direction of propagation.

Under a rotation about the axis of propagation, the two linear polarization states $(|x\rangle \text{ and } |y\rangle$ for horizontal and vertical polarization) transform as

$$\begin{aligned} |x\rangle &\to \cos \theta |x\rangle + \sin \theta |y\rangle \\ |y\rangle &\to -\sin \theta |x\rangle + \cos \theta |y\rangle. \end{aligned}$$
(2.45)

This two-dimensional representation is actually reducible. The matrix

$$\begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}$$
(2.46)

has the eigenstates

$$|R\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1\\ i \end{pmatrix} \qquad |L\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} i\\ 1 \end{pmatrix},$$
 (2.47)

with eigenvalues $e^{i\theta}$ and $e^{-i\theta}$, the states of right and left circular polarization. That is, these are the eigenstates of the rotation generator

$$\boldsymbol{J} = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} = \boldsymbol{\sigma}_2, \qquad (2.48)$$

with eigenvalues ± 1 . Because the eigenvalues are ± 1 (not $\pm \frac{1}{2}$) we say that the photon has spin-1.

In this context, the quantum interference phenomenon can be described as follows. The polarization states

$$|+\rangle = \frac{1}{\sqrt{2}} (|x\rangle + |y\rangle),$$

$$|-\rangle = \frac{1}{\sqrt{2}} (-|x\rangle + |y\rangle), \qquad (2.49)$$

are mutually orthogonal and can be obtained by rotating the states $|x\rangle$ and $|y\rangle$ by 45°. Suppose that we have a polarization analyzer that allows only one of two orthogonal linear photon polarizations to pass through, absorbing the other. Then an x or y polarized photon has probability $\frac{1}{2}$ of getting through a 45° rotated polarizer, and a 45° polarized photon has probability $\frac{1}{2}$ of getting through an x or y analyzer. But an x photon never passes through a y analyzer.

Suppose that a photon beam is directed at an x analyzer, with a y analyzer placed further downstream. Then about half of the photons will pass through the first analyzer, but every one of these will be stopped by the second analyzer. But now suppose that we place a 45° -rotated analyzer between the x and y analyzers. Then about half of the photons pass through each analyzer, and about one in eight will manage to pass all three without being absorbed. Because of this *interference* effect, there is no consistent interpretation in which each photon carries one classical bit of polarization information. Qubits are different than probabilistic classical bits.

A device can be constructed that rotates the linear polarization of a photon, and so applies the transformation Eq. (2.45) to our qubit; it functions by "turning on" a Hamiltonian for which the circular polarization states $|L\rangle$ and $|R\rangle$ are nondegenerate energy eigenstates. This is not the most general possible unitary transformation. But if we also have a device that alters the relative phase of the two orthogonal linear polarization states

(by turning on a Hamiltonian whose nondegenerate energy eigenstates are the linear polarization states), then the two devices can be employed together to apply an arbitrary 2×2 unitary transformation (of determinant 1) to the photon polarization state.

2.3 The density operator

2.3.1 The bipartite quantum system

Having understood everything about a single qubit, we are ready to address systems with two qubits. Stepping up from one qubit to two is a bigger leap than you might expect. Much that is weird and wonderful about quantum mechanics can be appreciated by considering the properties of the quantum states of two qubits.

The axioms of $\S2.1$ provide a perfectly acceptable general formulation of the quantum theory. Yet under many circumstances, we find that the axioms appear to be violated. The trouble is that our axioms are intended to characterize the quantum behavior of a *closed system* that does not interact with its surroundings. In practice, closed quantum systems do not exist; the observations we make are always limited to a small part of a much larger quantum system.

When we study *open systems*, that is, when we limit our attention to just part of a larger system, then (contrary to the axioms):

- 1. States are *not* rays.
- 2. Measurements are *not* orthogonal projections.
- 3. Evolution is *not* unitary.

To arrive at the laws obeyed by open quantum systems, we must recall our fifth axiom, which relates the description of a composite quantum system to the description of its component parts. As a first step toward understanding the quantum description of an open system, consider a two-qubit world in which we observe only one of the qubits. Qubit A is here in the room with us, and we are free to observe or manipulate it any way we please. But qubit B is locked in a vault where we can't get access to it. The full system AB obeys the axioms of §2.1. But we would like to find a compact way to characterize the observations that can be made on qubit A alone.

We'll use $\{|0\rangle_A, |1\rangle_A\}$ and $\{|0\rangle_B, |1\rangle_B\}$ to denote orthonormal bases for qubits A and B respectively. Consider a quantum state of the two-qubit world of the form

$$|\psi\rangle_{AB} = a|0\rangle_A \otimes |0\rangle_B + b|1\rangle_A \otimes |1\rangle_B. \tag{2.51}$$

In this state, qubits A and B are correlated. Suppose we measure qubit A by projecting onto the $\{|0\rangle_A, |1\rangle_A\}$ basis. Then with probability $|a|^2$ we obtain the result $|0\rangle_A$, and the measurement prepares the state

$$|0\rangle_A \otimes |0\rangle_B ;$$
 (2.52)

with probability $|b|^2$, we obtain the result $|1\rangle_A$ and prepare the state

$$|1\rangle_A \otimes |1\rangle_B. \tag{2.53}$$

In either case, a definite state of qubit *B* is picked out by the measurement. If we subsequently measure qubit *B*, then we are guaranteed (with probability one) to find $|0\rangle_B$ if we had found $|0\rangle_A$, and we are guaranteed to find $|1\rangle_B$ if we had found $|1\rangle_A$. In this sense, the outcomes of the $\{|0\rangle_A, |1\rangle_A\}$ and $\{|0\rangle_B, |1\rangle_B\}$ measurements are perfectly correlated in the state $|\psi\rangle_{AB}$.

But now we would like to consider more general observables acting on qubit A, and we would like to characterize the measurement outcomes for A alone (irrespective of the outcomes of any measurements of the inaccessible qubit B). An observable acting on qubit A only can be expressed as

$$\boldsymbol{M}_A \otimes \boldsymbol{I}_B,$$
 (2.54)

where M_A is a self-adjoint operator acting on A, and I_B is the identity operator acting on B. The expectation value of the observable in the state $|\psi\rangle$ is:

$$\langle \boldsymbol{M}_A \rangle = \langle \psi | \boldsymbol{M}_A \otimes \boldsymbol{I}_B | \psi \rangle$$

= $(a^* \langle 00| + b^* \langle 11|) (\boldsymbol{M}_A \otimes \boldsymbol{I}_B) (a|00\rangle + b|11\rangle)$
= $|a|^2 \langle 0| \boldsymbol{M}_A | 0 \rangle + |b|^2 \langle 1| \mathbf{M}_A | 1 \rangle$ (2.55)

(where we have used the orthogonality of $|0\rangle_B$ and $|1\rangle_B$). This expression can be rewritten in the form

$$\langle \boldsymbol{M}_A \rangle = \operatorname{tr}\left(\boldsymbol{M}_A \boldsymbol{\rho}_A\right), \quad \boldsymbol{\rho}_A = |a|^2 |0\rangle \langle 0| + |b|^2 |1\rangle \langle 1|$$
 (2.56)

and tr(·) denotes the *trace*. The operator ρ_A is called the *density operator* (or *density matrix*) for qubit A. It is self-adjoint, positive (its eigenvalues are nonnegative) and it has unit trace (because $|\psi\rangle$ is a normalized state.)

Because $\langle \boldsymbol{M}_A \rangle$ has the form eq.(2.56) for any observable \boldsymbol{M}_A acting on qubit A, it is consistent to interpret $\boldsymbol{\rho}_A$ as representing an *ensemble* of possible quantum states, each occurring with a specified probability. That is, we would obtain precisely the same result for $\langle \boldsymbol{M}_A \rangle$ if we stipulated that qubit A is in one of two quantum states. With probability $p_0 = |a|^2$ it is in the quantum state $|0\rangle$, and with probability $p_1 = |b|^2$ it is in the state $|1\rangle$. If we are interested in the result of any possible measurement, we can consider \boldsymbol{M}_A to be the projection $\boldsymbol{E}_A(a)$ onto the relevant eigenspace of a particular observable. Then

$$\operatorname{Prob}(a) = p_0 \langle 0 | \boldsymbol{E}_A(a) | 0 \rangle + p_1 \langle 1 | \boldsymbol{E}_A(a) | 1 \rangle, \qquad (2.57)$$

which is the probability of outcome a summed over the ensemble, and weighted by the probability of each state in the ensemble.

We have emphasized previously that there is an essential difference between a coherent superposition of the states $|0\rangle$ and $|1\rangle$, and a probabilistic ensemble, in which $|0\rangle$ and $|1\rangle$ can each occur with specified probabilities. For example, for a spin- $\frac{1}{2}$ object we have seen that if we measure σ_1 in the state $\frac{1}{\sqrt{2}}(|\uparrow_z\rangle + |\downarrow_z\rangle)$, we will obtain the result $|\uparrow_x\rangle$ with probability one. But the ensemble in which $|\uparrow_z\rangle$ and $|\downarrow_z\rangle$ each occur with probability $\frac{1}{2}$ is represented by the density operator

$$\boldsymbol{\rho} = \frac{1}{2} \left(|\uparrow_z\rangle \langle\uparrow_z| + |\downarrow_z\rangle \langle\downarrow_z| \right) = \frac{1}{2} \boldsymbol{I}, \qquad (2.58)$$

and the projection onto $|\uparrow_x\rangle$ then has the expectation value

$$\operatorname{tr}\left(\left|\uparrow_{x}\right\rangle\langle\uparrow_{x}|\boldsymbol{\rho}\right) = \langle\uparrow_{x}|\boldsymbol{\rho}|\uparrow_{x}\rangle = \frac{1}{2}.$$
(2.59)

Similarly, if we measure the spin along any axis labeled by polar angles θ and φ , the probability of obtaining the result "spin up" is

$$\langle |\psi(\theta,\varphi)\rangle \langle \psi(\theta,\varphi)|\rangle = \operatorname{tr} (|\psi(\theta,\varphi)\rangle \langle \psi(\theta,\varphi)|\boldsymbol{\rho})$$

= $\langle \psi(\theta,\varphi)|\frac{1}{2}\boldsymbol{I}|\psi(\theta,\varphi)\rangle = \frac{1}{2}.$ (2.60)

Therefore, if in the two-qubit world an equally weighted coherent superposition of $|00\rangle$ and $|11\rangle$ is prepared, the state of qubit A behaves *incoherently* – along any axis it is an equiprobable mixture of spin up and spin down.

This discussion of the correlated two-qubit state $|\psi\rangle_{AB}$ is easily generalized to an arbitrary state of any bipartite quantum system (a system divided into two parts). The Hilbert space of a bipartite system is $\mathcal{H}_A \otimes \mathcal{H}_B$ where $\mathcal{H}_{A,B}$ are the Hilbert spaces of the two parts. This means that if $\{|i\rangle_A\}$ is an orthonormal basis for \mathcal{H}_A and $\{|\mu\rangle_B\}$ is an orthonormal basis for \mathcal{H}_B , then $\{|i\rangle_A \otimes |\mu\rangle_B\}$ is an orthonormal basis for $\mathcal{H}_A \otimes \mathcal{H}_B$. Thus an arbitrary pure state of $\mathcal{H}_A \otimes \mathcal{H}_B$ can be expanded as

$$|\psi\rangle_{AB} = \sum_{i,\mu} a_{i\mu} |i\rangle_A \otimes |\mu\rangle_B, \qquad (2.61)$$

where $\sum_{i,\mu} |a_{i\mu}|^2 = 1$. The expectation value of an observable $M_A \otimes I_B$ that acts only on subsystem A is

$$\langle \boldsymbol{M}_{A} \rangle = {}_{AB} \langle \psi | \boldsymbol{M}_{A} \otimes \boldsymbol{I}_{B} | \psi \rangle_{AB}$$

$$= \sum_{j,\nu} a_{j\nu}^{*} (_{A} \langle j | \otimes _{B} \langle \nu |) (\boldsymbol{M}_{A} \otimes \boldsymbol{I}_{B}) \sum_{i,\mu} a_{i\mu} (|i\rangle_{A} \otimes |\mu\rangle_{B})$$

$$= \sum_{i,j,\mu} a_{j\mu}^{*} a_{i\mu} \langle j | \boldsymbol{M}_{A} | i \rangle = \operatorname{tr} (\boldsymbol{M}_{A} \boldsymbol{\rho}_{A}),$$

$$(2.62)$$

where

$$\boldsymbol{\rho}_{A} = \operatorname{tr}_{B}\left(|\psi\rangle\langle\psi|\right) \equiv \sum_{i,j,\mu} a_{i\mu} a_{j\mu}^{*} |i\rangle\langle j|$$
(2.63)

is the density operator of subsystem A.

We may say that the density operator ρ_A for subsystem A is obtained by performing a *partial trace* over subsystem B of the density operator (in this case a pure state) for the combined system AB. We may regard a dual vector (or bra) $_B\langle \mu |$ as a linear map that takes vectors in $\mathcal{H}_A \otimes \mathcal{H}_B$ to vectors of \mathcal{H}_A , defined through its action on a basis:

$${}_{B}\langle\mu|i\nu\rangle_{AB} = \delta_{\mu\nu}|i\rangle_{A} ; \qquad (2.64)$$

similarly, the ket $|\mu\rangle_B$ defines a map from the $\mathcal{H}_A \otimes \mathcal{H}_B$ dual basis to the \mathcal{H}_A dual basis, via

$$_{AB}\langle i\nu|\mu\rangle_B = \delta_{\mu\nu} \ _A\langle i|. \tag{2.65}$$

The partial trace operation is a linear map that takes an operator M_{AB} on $\mathcal{H}_A \otimes \mathcal{H}_B$ to an operator on \mathcal{H}_A defined as

$$\operatorname{tr}_{B} \boldsymbol{M}_{AB} = \sum_{\mu} {}_{B} \langle \mu | \boldsymbol{M}_{AB} | \mu \rangle_{B}.$$
(2.66)

We see that the density operator acting on A is the partial trace

$$\boldsymbol{\rho}_A = \operatorname{tr}_B\left(|\psi\rangle\langle\psi|\right). \tag{2.67}$$

From the definition eq.(2.63), we can immediately infer that ρ_A has the following properties:

- 1. ρ_A is self-adjoint: $\rho_A = \rho_A^{\dagger}$.
- 2. ρ_A is positive: For any $|\varphi\rangle$, $\langle \varphi | \rho_A | \varphi \rangle = \sum_{\mu} |\sum_i a_{i\mu} \langle \varphi | i \rangle |^2 \ge 0$.
- 3. tr(ρ_A) = 1: We have tr(ρ_A) = $\sum_{i,\mu} |a_{i\mu}|^2 = 1$, since $|\psi\rangle_{AB}$ is normalized.

It follows that ρ_A can be diagonalized in an orthonormal basis, that the eigenvalues are all real and nonnegative, and that the eigenvalues sum to one.

If we are looking at a subsystem of a larger quantum system, then, even if the state of the larger system is a ray, the state of the subsystem need not be; in general, the state is represented by a density operator. In the case where the state of the subsystem *is* a ray, and we say that the state is *pure*. Otherwise the state is *mixed*. If the state is a pure state $|\psi\rangle_A$, then the density matrix $\rho_A = |\psi\rangle\langle\psi|$ is the *projection* onto the one-dimensional space spanned by $|\psi\rangle_A$. Hence a pure density matrix has the property $\rho^2 = \rho$. A general density matrix, expressed in the basis $\{|a\rangle\}$ in which it is diagonal, has the form

$$\boldsymbol{\rho}_A = \sum_a p_a |a\rangle \langle a|, \qquad (2.68)$$

where $0 < p_a \leq 1$ and $\sum_a p_a = 1$. If the state is not pure, there are two or more terms in this sum, and $\rho^2 \neq \rho$; in fact, tr $\rho^2 = \sum p_a^2 < \sum p_a = 1$.

We say that ρ is an *incoherent* mixture of the states $\{|a\rangle\}$; "incoherent" means that the relative phases of the $|a\rangle$'s are experimentally inaccessible.

Since the expectation value of any observable M acting on the subsystem can be expressed as

$$\langle \boldsymbol{M} \rangle = \operatorname{tr} \, \boldsymbol{M} \boldsymbol{\rho} = \sum_{a} p_a \langle a | \boldsymbol{M} | a \rangle,$$
 (2.69)

we see as before that we may interpret ρ as describing an *ensemble* of pure quantum states, in which the state $|a\rangle$ occurs with probability p_a . We have, therefore, come a long part of the way to understanding how probabilities arise in quantum mechanics when a quantum system A interacts with another system B. A and B become *entangled*, that is, correlated. The entanglement *destroys the coherence* of a superposition of states of A, so that some of the phases in the superposition become inaccessible if we look at A alone. We may describe this situation by saying that the state of system A collapses — it is in one of a set of alternative states, each of which can be assigned a probability.

2.3.2 Bloch sphere

Let's return to the case in which system A is a single qubit, and consider the form of the general density matrix. The most general self-adjoint 2×2 matrix has four real parameters, and can be expanded in the basis $\{I, \sigma_1, \sigma_2, \sigma_3\}$. Since each σ_i is traceless, the coefficient of I in the expansion of a density matrix ρ must be $\frac{1}{2}$ (so that tr(ρ) = 1), and ρ may be expressed as

$$\boldsymbol{\rho}(\vec{P}) = \frac{1}{2} \left(\boldsymbol{I} + \vec{P} \cdot \vec{\sigma} \right) \\
\equiv \frac{1}{2} \left(\boldsymbol{I} + P_1 \boldsymbol{\sigma}_1 + P_2 \boldsymbol{\sigma}_2 + P_3 \boldsymbol{\sigma}_3 \right) \\
= \frac{1}{2} \left(\begin{array}{c} 1 + P_3 & P_1 - iP_2 \\ P_1 + iP_2 & 1 - P_3 \end{array} \right), \quad (2.70)$$

where P_1, P_2, P_3 are real numbers. We can compute $\det \rho = \frac{1}{4} \left(1 - \vec{P}^2 \right)$. Therefore, a necessary condition for ρ to have nonnegative eigenvalues is $\det \rho \geq 0$ or $\vec{P}^2 \leq 1$. This condition is also sufficient; since tr $\rho = 1$, it is not possible for ρ to have two negative eigenvalues. Thus, there is a 1 - 1 correspondence between the possible density matrices of a single qubit and the points on the *unit 3-ball* $0 \leq |\vec{P}| \leq 1$. This ball is usually called the *Bloch sphere* (although it is really a ball, not a sphere).

The boundary $(|\vec{P}| = 1)$ of the ball (which really *is* a sphere) contains the density matrices with vanishing determinant. Since tr $\rho = 1$, these density matrices must have the eigenvalues 0 and 1 — they are onedimensional projectors, and hence pure states. We have already seen that any pure state of a single qubit is of the form $|\psi(\theta,\varphi)\rangle$ and can be envisioned as a spin pointing in the (θ,φ) direction. Indeed using the property

$$\left(\hat{n}\cdot\vec{\boldsymbol{\sigma}}\right)^2 = \boldsymbol{I},\tag{2.71}$$

where \hat{n} is a unit vector, we can easily verify that the pure-state density matrix

$$\boldsymbol{\rho}(\hat{n}) = \frac{1}{2} \left(\boldsymbol{I} + \hat{n} \cdot \boldsymbol{\vec{\sigma}} \right)$$
(2.72)

satisfies the property

$$(\hat{n} \cdot \vec{\sigma}) \,\boldsymbol{\rho}(\hat{n}) = \boldsymbol{\rho}(\hat{n}) \,(\hat{n} \cdot \vec{\sigma}) = \boldsymbol{\rho}(\hat{n}), \qquad (2.73)$$

and, therefore is the projector

$$\boldsymbol{\rho}(\hat{n}) = |\psi(\hat{n})\rangle\langle\psi(\hat{n})| ; \qquad (2.74)$$

that is, \hat{n} is the direction along which the spin is pointing up. Alternatively, from the expression

$$|\psi(\theta,\phi)\rangle = \begin{pmatrix} e^{-i\varphi/2}\cos\left(\theta/2\right) \\ e^{i\varphi/2}\sin\left(\theta/2\right) \end{pmatrix},$$
(2.75)

we may compute directly that

$$\boldsymbol{\rho}(\theta,\phi) = |\psi(\theta,\phi)\rangle\langle\psi(\theta,\phi)| \\
= \begin{pmatrix} \cos^2(\theta/2) & \cos(\theta/2)\sin(\theta/2)e^{-i\varphi} \\ \cos(\theta/2)\sin(\theta/2)e^{i\varphi} & \sin^2(\theta/2) \end{pmatrix} \\
= \frac{1}{2}\boldsymbol{I} + \frac{1}{2}\begin{pmatrix} \cos\theta & \sin\theta e^{-i\varphi} \\ \sin\theta e^{i\varphi} & -\cos\theta \end{pmatrix} = \frac{1}{2}(\boldsymbol{I} + \hat{n} \cdot \vec{\boldsymbol{\sigma}}) \quad (2.76)$$

where $\hat{n} = (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta)$. One nice property of the Bloch parametrization of the pure states is that while $|\psi(\theta, \varphi)\rangle$ has an arbitrary overall phase that has no physical significance, there is no phase ambiguity in the density matrix $\rho(\theta, \varphi) = |\psi(\theta, \varphi)\rangle \langle \psi(\theta, \varphi)|$; all the parameters in ρ have a physical meaning.

From the property

$$\frac{1}{2} \text{tr} \,\boldsymbol{\sigma}_i \boldsymbol{\sigma}_j = \delta_{ij} \tag{2.77}$$

we see that

$$\langle \hat{n} \cdot \vec{\sigma} \rangle_{\vec{P}} = \operatorname{tr} \left(\hat{n} \cdot \vec{\sigma} \rho(\vec{P}) \right) = \hat{n} \cdot \vec{P} .$$
 (2.78)

We say that the vector \vec{P} in Eq. (2.70) parametrizes the *polarization* of the spin. If there are many identically prepared systems at our disposal, we can determine \vec{P} (and hence the complete density matrix $\rho(\vec{P})$) by measuring $\langle \hat{n} \cdot \vec{\sigma} \rangle$ along each of three linearly independent axes.

2.4 Schmidt decomposition

A bipartite pure state can be expressed in a standard form (*the Schmidt decomposition*) that is often very useful.

To arrive at this form, note that an arbitrary vector in $\mathcal{H}_A \otimes \mathcal{H}_B$ can be expanded as

$$|\psi\rangle_{AB} = \sum_{i,\mu} \psi_{i\mu} |i\rangle_A \otimes |\mu\rangle_B \equiv \sum_i |i\rangle_A \otimes |\tilde{i}\rangle_B.$$
(2.79)

Here $\{|i\rangle_A\}$ and $\{|\mu\rangle_B\}$ are orthonormal basis for \mathcal{H}_A and \mathcal{H}_B respectively, but to obtain the second equality in eq.(2.79) we have defined

$$|\tilde{i}\rangle_B \equiv \sum_{\mu} \psi_{i\mu} |\mu\rangle_B.$$
 (2.80)

Note that the $|\tilde{i}\rangle_B$'s need *not* be mutually orthogonal or normalized.

Now let's suppose that the $\{|i\rangle_A\}$ basis is chosen to be the basis in which ρ_A is diagonal,

$$\boldsymbol{\rho}_A = \sum_i p_i |i\rangle \langle i|. \tag{2.81}$$

We can also compute ρ_A by performing a partial trace,

$$\boldsymbol{\rho}_{A} = \operatorname{tr}_{B}(|\psi\rangle\langle\psi|)$$
$$= \operatorname{tr}_{B}(\sum_{i,j}|i\rangle\langle j|\otimes|\tilde{i}\rangle\langle\tilde{j}|) = \sum_{i,j}\langle\tilde{j}|\tilde{i}\rangle(|i\rangle\langle j|) \quad .$$
(2.82)

We obtained the last equality in eq.(2.82) by noting that

$$\operatorname{tr}_{B}\left(|\tilde{i}\rangle\langle\tilde{j}|\right) = \sum_{k} \langle k|\tilde{i}\rangle\langle\tilde{j}|k\rangle$$
$$= \sum_{k} \langle\tilde{j}|k\rangle\langle k|\tilde{i}\rangle = \langle\tilde{j}|\tilde{i}\rangle, \qquad (2.83)$$

where $\{|k\rangle\}$ is a complete orthonormal basis for \mathcal{H}_B . By comparing eq.(2.81) and eq. (2.82), we see that

$${}_{B}\langle \tilde{j}|\tilde{i}\rangle_{B} = p_{i}\delta_{ij}.$$
(2.84)

Hence, it turns out that the $\{|\tilde{i}\rangle_B\}$ are orthogonal after all. We obtain orthonormal vectors by rescaling,

$$|i'\rangle_B = p_i^{-1/2} |\tilde{i}\rangle_B \tag{2.85}$$

(we may assume $p_i \neq 0$, because we will need eq.(2.85) only for *i* appearing in the sum eq.(2.81)), and therefore obtain the expansion

$$|\psi\rangle_{AB} = \sum_{i} \sqrt{p_i} \ |i\rangle_A \otimes |i'\rangle_B, \qquad (2.86)$$

in terms of a *particular* orthonormal basis of \mathcal{H}_A and \mathcal{H}_B .

Eq.(2.86) is the Schmidt decomposition of the bipartite pure state $|\psi\rangle_{AB}$. Any bipartite pure state can be expressed in this form, but the bases used depend on the pure state that is being expanded. In general, we can't simultaneously expand both $|\psi\rangle_{AB}$ and $|\varphi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$ in the form eq.(2.86) using the same orthonormal bases for \mathcal{H}_A and \mathcal{H}_B .

It is instructive to compare the Schmidt decomposition of the bipartite pure state $|\psi\rangle_{AB}$ with its expansion in a generic orthonormal basis

$$|\psi\rangle_{AB} = \sum_{a,\mu} \psi_{a\mu} |a\rangle_A \otimes |\mu\rangle_B.$$
(2.87)

The orthonormal bases $\{|a\rangle_A\}$ and $\{|\mu\rangle_B\}$ are related to the Schmidt bases $\{|i\rangle_A\}$ and $\{|i'\rangle_B\}$ by unitary transformations U_A and U_B , hence

$$|i\rangle_A = \sum_a |a\rangle_A (U_A)_{ai}, \quad |i'\rangle_B = \sum_\mu |\mu\rangle_B (U_B)_{\mu i'}.$$
(2.88)

By equating the expressions for $|\psi\rangle_{AB}$ in eq.(2.86) and eq.(2.87), we find

$$\psi_{a\mu} = \sum_{i} (U_A)_{ai} \ \sqrt{p_i} \ \left(U_B^T\right)_{i\mu}.$$
 (2.89)

We see that by applying unitary transformations on the left and right, any matrix ψ can be transformed to a matrix which is diagonal and nonnegative. (The "diagonal" matrix will be rectangular rather than square if the Hilbert spaces \mathcal{H}_A and \mathcal{H}_B have different dimensions.) Eq.(2.89) is said to be the *singular value decomposition* of ψ , and the weights $\{\sqrt{p_i}\}$ in the Schmidt decomposition are ψ 's singular values.

Using eq.(2.86), we can also evaluate the partial trace over \mathcal{H}_A to obtain

$$\boldsymbol{\rho}_B = \operatorname{tr}_A\left(|\psi\rangle\langle\psi|\right) = \sum_i p_i |i'\rangle\langle i'|.$$
(2.90)

We see that ρ_A and ρ_B have the same nonzero eigenvalues. If \mathcal{H}_A and \mathcal{H}_B do not have the same dimension, the number of zero eigenvalues of ρ_A and ρ_B will differ.

If ρ_A (and hence ρ_B) have no degenerate eigenvalues other than zero, then the Schmidt decomposition of $|\psi\rangle_{AB}$ is essentially uniquely determined by ρ_A and ρ_B . We can diagonalize ρ_A and ρ_B to find the $|i\rangle_A$'s and $|i'\rangle_B$'s, and then we pair up the eigenstates of ρ_A and ρ_B with the same eigenvalue to obtain eq.(2.86). We have chosen the phases of our basis states so that no phases appear in the coefficients in the sum; the only remaining freedom is to redefine $|i\rangle_A$ and $|i'\rangle_B$ by multiplying by opposite phases (which leaves the expression eq.(2.86) unchanged).

But if ρ_A has degenerate nonzero eigenvalues, then we need more information than that provided by ρ_A and ρ_B to determine the Schmidt decomposition; we need to know which $|i'\rangle_B$ gets paired with each $|i\rangle_A$. For example, if both \mathcal{H}_A and \mathcal{H}_B are *d*-dimensional and U_{ij} is any $d \times d$ unitary matrix, then

$$|\psi\rangle_{AB} = \frac{1}{\sqrt{d}} \sum_{i,j=1}^{d} |i\rangle_A U_{ij} \otimes |j'\rangle_B, \qquad (2.91)$$

will yield $\rho_A = \rho_B = \frac{1}{d} I$ when we take partial traces. Furthermore, we are free to apply simultaneous unitary transformations in \mathcal{H}_A and \mathcal{H}_B ; writing

$$|i\rangle_A = \sum_a |a\rangle_A U_{ai}, \quad |i'\rangle_B = \sum_b |b'\rangle_B U_{bi}^*, \tag{2.92}$$

we have

$$|\psi\rangle_{AB} = \frac{1}{\sqrt{d}} \sum_{i} |i\rangle_{A} \otimes |i'\rangle_{B} = \frac{1}{\sqrt{d}} \sum_{i,a,b} |a\rangle_{A} U_{ai} \otimes |b'\rangle_{B} U_{ib}^{\dagger}$$
$$= \frac{1}{\sqrt{d}} \sum_{a} |a\rangle_{A} \otimes |a'\rangle_{B}.$$
(2.93)

This simultaneous rotation preserves the state $|\psi\rangle_{AB}$, illustrating that there is an ambiguity in the basis used when we express $|\psi\rangle_{AB}$ in the Schmidt form.

2.4.1 Entanglement

With any bipartite pure state $|\psi\rangle_{AB}$ we may associate a positive integer, the *Schmidt number*, which is the number of nonzero eigenvalues in ρ_A (or ρ_B) and hence the number of terms in the Schmidt decomposition of $|\psi\rangle_{AB}$. In terms of this quantity, we can define what it means for a bipartite pure state to be *entangled*: $|\psi\rangle_{AB}$ is entangled (or nonseparable) if its Schmidt number is greater than one; otherwise, it is *separable* (or unentangled). Thus, a separable bipartite pure state is a direct product of pure states in \mathcal{H}_A and \mathcal{H}_B ,

$$|\psi\rangle_{AB} = |\varphi\rangle_A \otimes |\chi\rangle_B; \tag{2.94}$$

then the reduced density matrices $\rho_A = |\varphi\rangle\langle\varphi|$ and $\rho_B = |\chi\rangle\langle\chi|$ are pure. Any state that cannot be expressed as such a direct product is entangled; then ρ_A and ρ_B are mixed states.

When $|\psi\rangle_{AB}$ is entangled we say that A and B have quantum correlations. It is not strictly correct to say that subsystems A and B are uncorrelated if $|\psi\rangle_{AB}$ is separable; after all, the two spins in the separable state

$$|\uparrow\rangle_A|\uparrow\rangle_B,\tag{2.95}$$

are surely correlated – they are both pointing in the same direction. But the correlations between A and B in an entangled state have a different character than those in a separable state. One crucial difference is that *entanglement cannot be created locally*. The only way to entangle A and B is for the two subsystems to directly interact with one another.

We can prepare the state eq.(2.95) without allowing spins A and B to ever come into contact with one another. We need only send a (classical!) message to two preparers (Alice and Bob) telling both of them to prepare a spin pointing along the z-axis. But the only way to turn the state eq.(2.95) into an entangled state like

$$\frac{1}{\sqrt{2}} \left(|\uparrow\rangle_A|\uparrow\rangle_B + |\downarrow\rangle_A|\downarrow\rangle_B \right), \tag{2.96}$$

is to apply a *collective* unitary transformation to the state. Local unitary transformations of the form $U_A \otimes U_B$, and local measurements performed by Alice or Bob, *cannot increase the Schmidt number* of the two-qubit state, no matter how much Alice and Bob discuss what they do. To entangle two qubits, we *must* bring them together and allow them to interact.

As we will discuss in Chapter 4, it is also possible to make the distinction between entangled and separable bipartite *mixed* states. We will also discuss various ways in which local operations can modify the form of entanglement, and some ways that entanglement can be put to use.

2.5 Ambiguity of the ensemble interpretation

2.5.1 Convexity

Recall that an operator ρ acting on a Hilbert space \mathcal{H} may be interpreted as a density operator if it has the three properties:

(1) $\boldsymbol{\rho}$ is self-adjoint.

- (2) $\boldsymbol{\rho}$ is nonnegative.
- (3) $\operatorname{tr}(\boldsymbol{\rho}) = 1.$

It follows immediately that, given two density matrices ρ_1 , and ρ_2 , we can always construct another density matrix as a convex linear combination of the two:

$$\boldsymbol{\rho}(\lambda) = \lambda \boldsymbol{\rho}_1 + (1 - \lambda) \boldsymbol{\rho}_2 \tag{2.97}$$

is a density matrix for any real λ satisfying $0 \le \lambda \le 1$. We easily see that $\rho(\lambda)$ satisfies (1) and (3) if ρ_1 and ρ_2 do. To check (2), we evaluate

$$\langle \psi | \boldsymbol{\rho}(\lambda) | \psi \rangle = \lambda \langle \psi | \boldsymbol{\rho}_1 | \psi \rangle + (1 - \lambda) \langle \psi | \boldsymbol{\rho}_2 | \psi \rangle \ge 0; \qquad (2.98)$$

 $\langle \boldsymbol{\rho}(\lambda) \rangle$ is guaranteed to be nonnegative because $\langle \boldsymbol{\rho}_1 \rangle$ and $\langle \boldsymbol{\rho}_2 \rangle$ are. We have, therefore, shown that in a Hilbert space \mathcal{H} of dimension d, the density operators are a *convex subset* of the real vector space of $d \times d$ hermitian operators. (A subset of a vector space is said to be convex if the set contains the straight line segment connecting any two points in the set.)

Most density operators can be expressed as a sum of other density operators in many different ways. But the pure states are special in this regard – it is *not* possible to express a pure state as a convex sum of two other states. Consider a pure state $\rho = |\psi\rangle\langle\psi|$, and let $|\psi_{\perp}\rangle$ denote a vector orthogonal to $|\psi\rangle, \langle\psi_{\perp}|\psi\rangle = 0$. Suppose that ρ can be expanded as in eq.(2.97); then

$$\langle \psi_{\perp} | \boldsymbol{\rho} | \psi_{\perp} \rangle = 0 = \lambda \langle \psi_{\perp} | \boldsymbol{\rho}_{1} | \psi_{\perp} \rangle + (1 - \lambda) \langle \psi_{\perp} | \boldsymbol{\rho}_{2} | \psi_{\perp} \rangle.$$
 (2.99)

Since the right hand side is a sum of two nonnegative terms, and the sum vanishes, both terms must vanish. If λ is not 0 or 1, we conclude that ρ_1 and ρ_2 are orthogonal to $|\psi_{\perp}\rangle$. But since $|\psi_{\perp}\rangle$ can be any vector orthogonal to $|\psi\rangle$, we see that $\rho_1 = \rho_2 = \rho$.

The vectors in a convex set that cannot be expressed as a linear combination of other vectors in the set are called the *extremal points* of the set. We have just shown that the pure states are extremal points of the set of density matrices. Furthermore, *only* the pure states are extremal, because any mixed state can be written $\rho = \sum_i p_i |i\rangle \langle i|$ in the basis in which it is diagonal, and so is a convex sum of pure states.

We have already encountered this structure in our discussion of the special case of the Bloch sphere. We saw that the density operators are a (unit) ball in the three-dimensional set of 2×2 hermitian matrices with unit trace. The ball is convex, and its extremal points are the points on the boundary. Similarly, the $d \times d$ density operators are a convex subset of the (d^2-1) -dimensional set of $d \times d$ hermitian matrices with unit trace, and the extremal points of the set are the pure states.

However, the 2×2 case is atypical in one respect: for d > 2, the points on the boundary of the set of density matrices are not necessarily pure states. The boundary of the set consists of all density matrices with at least one vanishing eigenvalue (since there are nearby matrices with negative eigenvalues). Such a density matrix need not be pure, for d > 2, since the number of nonvanishing eigenvalues can exceed one.

2.5.2 Ensemble preparation

The convexity of the set of density matrices has a simple and enlightening physical interpretation. Suppose that a preparer agrees to prepare one of two possible states; with probability λ , the state ρ_1 is prepared, and with probability $1 - \lambda$, the state ρ_2 is prepared. (A random number generator might be employed to guide this choice.) To evaluate the expectation value of any observable M, we average over *both* the choices of preparation *and* the outcome of the quantum measurement:

$$\langle \boldsymbol{M} \rangle = \lambda \langle \boldsymbol{M} \rangle_1 + (1 - \lambda) \langle \boldsymbol{M} \rangle_2 = \lambda \operatorname{tr}(\boldsymbol{M} \boldsymbol{\rho}_1) + (1 - \lambda) \operatorname{tr}(\boldsymbol{M} \boldsymbol{\rho}_2) = \operatorname{tr}(\boldsymbol{M} \boldsymbol{\rho}(\lambda)).$$
 (2.100)

All expectation values are thus indistinguishable from what we would obtain if the state $\rho(\lambda)$ had been prepared instead. Thus, we have an operational procedure, given methods for preparing the states ρ_1 and ρ_2 , for preparing any convex combination.

Indeed, for any mixed state ρ , there are an infinite variety of ways to express ρ as a convex combination of other states, and hence an infinite variety of procedures we could employ to prepare ρ , all of which have exactly the same consequences for any conceivable observation of the system. But a pure state is different; it can be prepared in only one way. (This is what is "pure" about a pure state.) Every pure state is an eigenstate of some observable, e.g., for the state $\rho = |\psi\rangle\langle\psi|$, measurement of the projection $E = |\psi\rangle\langle\psi|$ is guaranteed to have the outcome 1. (For example, recall that every pure state of a single qubit is "spin-up" along some axis.) Since ρ is the only state for which the outcome of measuring E is 1 with 100% probability, there is no way to reproduce this observable property by choosing one of several possible preparations. Thus, the preparation of a pure state is unambiguous (we can infer a unique preparation if we have many copies of the state to experiment with), but the preparation of a mixed state is always ambiguous.

How ambiguous is it? Since any ρ can be expressed as a sum of pure states, let's confine our attention to the question: in how many ways can a density operator be expressed as a convex sum of pure states? Mathematically, this is the question: in how many ways can ρ be written as a sum of *extremal* states?

As a first example, consider the "maximally mixed" state of a single qubit:

$$\boldsymbol{\rho} = \frac{1}{2}\boldsymbol{I}.\tag{2.101}$$

This can indeed be prepared as an ensemble of pure states in an infinite variety of ways. For example,

$$\boldsymbol{\rho} = \frac{1}{2} |\uparrow_z\rangle \langle\uparrow_z| + \frac{1}{2} |\downarrow_z\rangle \langle\downarrow_z|, \qquad (2.102)$$

so we obtain ρ if we prepare either $|\uparrow_z\rangle$ or $|\downarrow_z\rangle$, each occurring with probability $\frac{1}{2}$. But we also have

$$\boldsymbol{\rho} = \frac{1}{2} |\uparrow_x\rangle \langle\uparrow_x| + \frac{1}{2} |\downarrow_x\rangle \langle\downarrow_x|, \qquad (2.103)$$

so we obtain ρ if we prepare either $|\uparrow_x\rangle$ or $|\downarrow_x\rangle$, each occurring with probability $\frac{1}{2}$. Now the preparation procedures are undeniably *different*. Yet there is no possible way to tell the difference by making observations of the spin.

More generally, the point at the center of the Bloch ball is the sum of any two antipodal points on the sphere – preparing either $|\uparrow_{\hat{n}}\rangle$ or $|\downarrow_{\hat{n}}\rangle$, each occurring with probability $\frac{1}{2}$, will generate $\rho = \frac{1}{2}I$.

Only in the case where ρ has two (or more) degenerate eigenvalues will there be distinct ways of generating ρ from an ensemble of *mutually orthogonal* pure states, but there is no good reason to confine our attention to ensembles of mutually orthogonal pure states. We may consider a point in the interior of the Bloch ball

$$\boldsymbol{\rho}(\vec{P}) = \frac{1}{2} (\boldsymbol{I} + \vec{P} \cdot \vec{\sigma}), \qquad (2.104)$$

with $0 < |\vec{P}| < 1$, and it too can be expressed as

$$\boldsymbol{\rho}(\vec{P}) = \lambda \boldsymbol{\rho}(\hat{n}_1) + (1-\lambda)\boldsymbol{\rho}(\hat{n}_2), \qquad (2.105)$$

if $\vec{P} = \lambda \hat{n}_1 + (1 - \lambda) \hat{n}_2$ (or in other words, if \vec{P} lies somewhere on the line segment connecting the points \hat{n}_1 and \hat{n}_2 on the sphere). Evidently, for any \vec{P} , there is a an expression for $\rho(\vec{P})$ as a convex combination of pure states associated with any chord of the Bloch sphere that passes through the point \vec{P} ; all such chords comprise a two-parameter family.

This highly ambiguous nature of the preparation of a mixed quantum state is one of the characteristic features of quantum information that contrasts sharply with classical probability distributions. Consider, for example, the case of a probability distribution for a single classical bit. The two extremal distributions are those in which either 0 or 1 occurs with 100% probability. Any probability distribution for the bit is a convex sum of these two extremal points. Similarly, if there are d possible states, there are d extremal distributions, and any probability distribution has a *unique* decomposition into extremal ones (the convex set of probability distributions is a *simplex*, the convex hull of its d vertices). If 0 occurs with 21% probability, 1 with 33% probability, and 2 with 46% probability, there is a unique "preparation procedure" that yields this probability distribution.

2.5.3 Faster than light?

Let's now return to our earlier viewpoint — that a mixed state of system A arises because A is *entangled* with system B — to further consider the implications of the ambiguous preparation of mixed states. If qubit A has density matrix

$$\boldsymbol{\rho}_A = \frac{1}{2} |\uparrow_z\rangle \langle\uparrow_z| + \frac{1}{2} |\downarrow_z\rangle \langle\downarrow_z|, \qquad (2.106)$$

this density matrix could arise from an entangled bipartite pure state $|\psi\rangle_{AB}$ with the Schmidt decomposition

$$|\psi\rangle_{AB} = \frac{1}{\sqrt{2}} \left(|\uparrow_z\rangle_A|\uparrow_z\rangle_B + |\downarrow_z\rangle_A|\downarrow_z\rangle_B\right).$$
(2.107)

Therefore, the ensemble interpretation of ρ_A in which either $|\uparrow_z\rangle_A$ or $|\downarrow_z\rangle_A$ is prepared (each with probability $p = \frac{1}{2}$) can be realized by performing a measurement of qubit B. We measure qubit B in the $\{|\uparrow_z\rangle_B, |\downarrow_z\rangle_B\}$ basis; if the result $|\uparrow_z\rangle_B$ is obtained, we have prepared $|\uparrow_z\rangle_A$, and if the result $|\downarrow_z\rangle_B$ is obtained, we have prepared $|\downarrow_z\rangle_A$.

But as we have already noted, in this case, because ρ_A has degenerate eigenvalues, the Schmidt basis is not unique. We can apply simultaneous unitary transformations to qubits A and B (actually, if we apply U to Awe must apply U^* to B as in eq.(2.92)) without modifying the bipartite pure state $|\psi\rangle_{AB}$. Therefore, for any unit 3-vector $\hat{n}, |\psi\rangle_{AB}$ has a Schmidt decomposition of the form

$$|\psi\rangle_{AB} = \frac{1}{\sqrt{2}} \left(|\uparrow_{\hat{n}}\rangle_{A}|\uparrow_{\hat{n}'}\rangle_{B} + |\downarrow_{\hat{n}}\rangle_{A}|\downarrow_{\hat{n}'}\rangle_{B}\right).$$
(2.108)

We see that by measuring qubit B in a suitable basis, we can realize any interpretation of ρ_A as an ensemble of two pure states.

This property suggests a mechanism for faster-than-light communication. Many copies of $|\psi\rangle_{AB}$ are prepared. Alice takes all of the A qubits to the Andromeda galaxy and Bob keeps all of the B qubits on earth. When Bob wants to send a one-bit message to Alice, he chooses to measure either σ_1 or σ_3 for all his spins, thus preparing Alice's spins in either the $\{|\uparrow_z\rangle_A, |\downarrow_z\rangle_A\}$ or $\{|\uparrow_x\rangle_A, |\downarrow_x\rangle_A\}$ ensembles. (*U* is real in this case, so $U = U^*$ and $\hat{n} = \hat{n}'$.) To read the message, Alice immediately measures her spins to see which ensemble has been prepared.

This scheme has a flaw. Though the two preparation methods are surely different, both ensembles are described by precisely the same density matrix ρ_A . Thus, there is no conceivable measurement Alice can make that will distinguish the two ensembles, and no way for Alice to tell what action Bob performed. The "message" is unreadable.

Why, then, do we confidently state that "the two preparation methods are surely different?" To qualm any doubts about that, imagine that Bob either (1) measures all of his spins along the \hat{z} -axis, or (2) measures all of his spins along the \hat{x} -axis, and then calls Alice on the intergalactic telephone. He does not tell Alice whether he did (1) or (2), but he does tell her the results of all his measurements: "the first spin was up, the second was down," etc. Now Alice performs either (1) or (2) on her spins. If both Alice and Bob measured along the same axis, Alice will find that every single one of her measurement outcomes agrees with what Bob found. But if Alice and Bob measured along different (orthogonal) axes, then Alice will find *no correlation* between her results and Bob's. About half of her measurements agree with Bob's and about half disagree. If Bob promises to do either (1) or (2), and assuming no preparation or measurement errors, then Alice will know that Bob's action was different than hers (even though Bob never told her this information) as soon as one of her measurements disagrees with what Bob found. If all their measurements agree, then if many spins are measured, Alice will have very high statistical confidence that she and Bob measured along the same axis. (Even with occasional measurement errors, the statistical test will still be highly reliable if the error rate is low enough.) So Alice does have a way to distinguish Bob's two preparation methods, but in this case there is certainly no faster-than-light communication, because Alice had to receive Bob's phone call before she could perform her test.

2.5.4 Quantum erasure

We had said that the density matrix $\rho_A = \frac{1}{2}I$ describes a spin in an *incoherent* mixture of the pure states $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$. This was to be distinguished from *coherent* superpositions of these states, such as

$$|\uparrow_x,\downarrow_x\rangle = \frac{1}{\sqrt{2}} \left(|\uparrow_z\rangle \pm |\downarrow_z\rangle\right) ; \qquad (2.109)$$

in the case of a coherent superposition, the *relative phase* of the two states has observable consequences (distinguishes $|\uparrow_x\rangle$ from $|\downarrow_x\rangle$). In the case of an incoherent mixture, the relative phase is completely unobservable.

The superposition becomes incoherent if spin A becomes entangled with another spin B, and spin B is inaccessible.

Heuristically, the states $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$ can *interfere* (the relative phase of these states can be observed) only if we have no information about whether the spin state is $|\uparrow_z\rangle_A$ or $|\downarrow_z\rangle_A$. More than that, interference can occur only if there is *in principle no possible way* to find out whether the spin is up or down along the z-axis. Entangling spin Awith spin B destroys interference, (causes spin A to *decohere*) because it is possible in principle for us to determine if spin A is up or down along \hat{z} by performing a suitable measurement of spin B.

But we have now seen that the statement that entanglement causes decoherence requires a qualification. Suppose that Bob measures spin Balong the \hat{x} -axis, obtaining either the result $|\uparrow_x\rangle_B$ or $|\downarrow_x\rangle_B$, and that he sends his measurement result to Alice. Now Alice's spin is a pure state (either $|\uparrow_x\rangle_A$ or $|\downarrow_x\rangle_A$) and in fact a coherent superposition of $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$. We have managed to recover the purity of Alice's spin before the jaws of decoherence could close!

Suppose that Bob allows his spin to pass through a Stern–Gerlach apparatus oriented along the \hat{z} -axis. Well, of course, Alice's spin can't behave like a coherent superposition of $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$; all Bob has to do is look to see which way his spin moved, and he will know whether Alice's spin is up or down along \hat{z} . But suppose that Bob does not look. Instead, he carefully refocuses the two beams without maintaining any record of whether his spin moved up or down, and *then* allows the spin to pass through a second Stern–Gerlach apparatus oriented along the \hat{x} -axis. *This* time he looks, and communicates the result of his σ_1 measurement to Alice. Now the coherence of Alice's spin has been restored!

This situation has been called a quantum eraser. Entangling the two spins creates a "measurement situation" in which the coherence of $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$ is lost because we can find out if spin A is up or down along \hat{z} by observing spin B. But when we measure spin B along \hat{x} , this information is "erased." Whether the result is $|\uparrow_x\rangle_B$ or $|\downarrow_x\rangle_B$ does not tell us anything about whether spin A is up or down along \hat{z} , because Bob has been careful not to retain the "which way" information that he might have acquired by looking at the first Stern–Gerlach apparatus. Therefore, it is possible again for spin A to behave like a coherent superposition of $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$ (and it does, after Alice hears about Bob's result).

We can best understand the quantum eraser from the ensemble viewpoint. Alice has many spins selected from an ensemble described by $\rho_A = \frac{1}{2}I$, and there is no way for her to observe interference between $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$. When Bob makes his measurement along \hat{x} , a particular preparation of the ensemble is realized. However, this has no effect that Alice can perceive – her spin is *still* described by $\rho_A = \frac{1}{2}I$ as before. But, when Alice receives Bob's phone call, she can select a *subensemble* of her spins that are all in the pure state $|\uparrow_x\rangle_A$. The information that Bob sends allows Alice to distill purity from a maximally mixed state.

Another wrinkle on the quantum eraser is sometimes called *delayed* choice. This just means that the situation we have described is really completely symmetric between Alice and Bob, so it can't make any difference who measures first. (Indeed, if Alice's and Bob's measurements are spacelike separated events, there is no invariant meaning to which came first; it depends on the frame of reference of the observer.) Alice could measure all of her spins today (say along \hat{x}) before Bob has made his mind up how he will measure his spins. Next week, Bob can decide to "prepare" Alice's spins in the states $|\uparrow_{\hat{n}}\rangle_A$ and $|\downarrow_{\hat{n}}\rangle_A$ (that is the "delayed choice"). He then tells Alice which were the $|\uparrow_{\hat{n}}\rangle_A$ spins, and she can check her measurement record to verify that

$$\langle \sigma_1 \rangle_{\hat{n}} = \hat{n} \cdot \hat{x} \ . \tag{2.110}$$

The results are the same, irrespective of whether Bob "prepares" the spins before or after Alice measures them.

We have claimed that the density matrix ρ_A provides a complete physical description of the state of subsystem A, because it characterizes all possible measurements that can be performed on A. One might object that the quantum eraser phenomenon demonstrates otherwise. Since the information received from Bob enables Alice to recover a pure state from the mixture, how can we hold that everything Alice can know about A is encoded in ρ_A ?

I prefer to say that quantum erasure illustrates the principle that "information is physical." The state ρ_A of system A is not the same thing as ρ_A accompanied by the information that Alice has received from Bob. This information (which attaches labels to the subensembles) changes the physical description. That is, we should include Alice's "state of knowledge" in our description of her system. An ensemble of spins for which Alice has no information about whether each spin is up or down is a *different* physical state than an ensemble in which Alice knows which spins are up and which are down. This "state of knowledge" need not really be the state of a human mind; any (inanimate) record that labels the subensemble will suffice.

2.5.5 The HJW theorem

So far, we have considered the quantum eraser only in the context of a single qubit, described by an ensemble of equally probable mutually orthogonal states, (*i.e.*, $\rho_A = \frac{1}{2}I$). The discussion can be considerably generalized.

We have already seen that a mixed state of any quantum system can be realized as an ensemble of pure states in an infinite number of different ways. For a density matrix ρ_A , consider one such realization:

$$\boldsymbol{\rho}_A = \sum_i p_i |\varphi_i\rangle \langle \varphi_i|, \quad \sum p_i = 1.$$
 (2.111)

Here the states $\{|\varphi_i\rangle_A\}$ are all normalized vectors, but we do *not* assume that they are mutually orthogonal. Nevertheless, ρ_A can be realized as an ensemble, in which each pure state $|\varphi_i\rangle\langle\varphi_i|$ occurs with probability p_i .

For any such ρ_A , we can construct a "purification" of ρ_A , a bipartite pure state $|\Phi_1\rangle_{AB}$ that yields ρ_A when we perform a partial trace over \mathcal{H}_B . One such purification is of the form

$$|\Phi_1\rangle_{AB} = \sum_i \sqrt{p_i} |\varphi_i\rangle_A \otimes |\alpha_i\rangle_B, \qquad (2.112)$$

where the vectors $|\alpha_i\rangle_B \in \mathcal{H}_B$ are mutually orthogonal and normalized,

$$\langle \alpha_i | \alpha_j \rangle = \delta_{ij}. \tag{2.113}$$

Clearly, then,

$$\operatorname{tr}_B(|\Phi_1\rangle\langle\Phi_1|) = \boldsymbol{\rho}_A. \tag{2.114}$$

Furthermore, we can imagine performing an orthogonal measurement in system B that projects onto the $|\alpha_i\rangle_B$ basis. (The $|\alpha_i\rangle_B$'s might not span \mathcal{H}_B , but in the state $|\Phi_1\rangle_{AB}$, measurement outcomes orthogonal to all the $|\alpha_i\rangle_B$'s never occur.) The outcome $|\alpha_i\rangle_B$ will occur with probability p_i , and will prepare the pure state $|\varphi_i\rangle\langle\varphi_i|$ of system A. Thus, given the purification $|\Phi_1\rangle_{AB}$ of ρ_A , there is a measurement we can perform in system B that realizes the $|\varphi_i\rangle_A$ ensemble interpretation of ρ_A . When the measurement outcome in B is known, we have successfully extracted one of the pure states $|\varphi_i\rangle_A$ from the mixture ρ_A .

What we have just described is a generalization of preparing $|\uparrow_z\rangle_A$ by measuring spin B along \hat{z} (in our discussion of two entangled qubits). But to generalize the notion of a quantum eraser, we wish to see that in the state $|\Phi_1\rangle_{AB}$, we can realize a *different* ensemble interpretation of ρ_A by performing a different measurement of B. So let

$$\boldsymbol{\rho}_A = \sum_{\mu} q_{\mu} |\psi_{\mu}\rangle \langle \psi_{\mu}|, \qquad (2.115)$$

be another realization of the same density matrix ρ_A as an ensemble of pure states. For this ensemble as well, there is a corresponding purification

$$|\Phi_2\rangle_{AB} = \sum_{\mu} \sqrt{q_{\mu}} |\psi_{\mu}\rangle_A \otimes |\beta_{\mu}\rangle_B, \qquad (2.116)$$

where again the $\{|\beta_{\mu}\rangle_{B}$'s $\}$ are orthonormal vectors in \mathcal{H}_{B} . So in the state $|\Phi_{2}\rangle_{AB}$, we can realize the ensemble by performing a measurement in \mathcal{H}_{B} that projects onto the $\{|\beta_{\mu}\rangle_{B}\}$ basis.

Now, how are $|\Phi_1\rangle_{AB}$ and $|\Phi_2\rangle_{AB}$ related? In fact, we can easily show that

$$|\Phi_1\rangle_{AB} = (\boldsymbol{I}_A \otimes \boldsymbol{U}_B) |\Phi_2\rangle_{AB}; \qquad (2.117)$$

the two states differ by a unitary change of basis acting in \mathcal{H}_B alone, or

$$|\Phi_1\rangle_{AB} = \sum_{\mu} \sqrt{q_{\mu}} |\psi_{\mu}\rangle_A \otimes |\gamma_{\mu}\rangle_B, \qquad (2.118)$$

where

$$|\gamma_{\mu}\rangle_{B} = \boldsymbol{U}_{B}|\beta_{\mu}\rangle_{B}, \qquad (2.119)$$

is yet another orthonormal basis for \mathcal{H}_B . We see, then, that there is a *sin-gle* purification $|\Phi_1\rangle_{AB}$ of ρ_A , such that we can realize either the $\{|\varphi_i\rangle_A\}$ ensemble or $\{|\psi_{\mu}\rangle_A\}$ ensemble by choosing to measure the appropriate observable in system B!

Similarly, we may consider many ensembles that all realize ρ_A , where the maximum number of pure states appearing in any of the ensembles is d. Then we may choose a Hilbert space \mathcal{H}_B of dimension d, and a pure state $|\Phi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$, such that any one of the ensembles can be realized by measuring a suitable observable of B. This is the HJWtheorem (for Hughston, Jozsa, and Wootters); it expresses the quantum eraser phenomenon in its most general form.

In fact, the HJW theorem is an easy corollary to the Schmidt decomposition. Both $|\Phi_1\rangle_{AB}$ and $|\Phi_2\rangle_{AB}$ have Schmidt decompositions, and because both yield the same ρ_A when we take the partial trace over B, these decompositions must have the form

$$|\Phi_1\rangle_{AB} = \sum_k \sqrt{\lambda_k} |k\rangle_A \otimes |k_1'\rangle_B,$$

$$|\Phi_2\rangle_{AB} = \sum_k \sqrt{\lambda_k} |k\rangle_A \otimes |k_2'\rangle_B,$$
 (2.120)

where the λ_k 's are the eigenvalues of ρ_A and the $|k\rangle_A$'s are the corresponding eigenvectors. But since $\{|k'_1\rangle_B\}$ and $\{|k'_2\rangle_B\}$ are both orthonormal bases for \mathcal{H}_B , there is a unitary U_B such that

$$|k_1'\rangle_B = \boldsymbol{U}_B |k_2'\rangle_B, \qquad (2.121)$$

from which eq.(2.117) immediately follows.

In the ensemble of pure states described by Eq. (2.111), we would say that the pure states $|\varphi_i\rangle_A$ are mixed *incoherently* — an observer in system A cannot detect the relative phases of these states. Heuristically, the reason that these states cannot interfere is that it is possible in principle to find out which representative of the ensemble is actually realized by performing a measurement in system B, a projection onto the orthonormal basis $\{|\alpha_i\rangle_B\}$. However, by projecting onto the $\{|\gamma_{\mu}\rangle_B\}$ basis instead, and relaying the information about the measurement outcome to system A, we can extract one of the pure states $|\psi_{\mu}\rangle_A$ from the ensemble, even though this state may be a coherent superposition of the $|\varphi_i\rangle_A$'s. In effect, measuring B in the $\{|\gamma_{\mu}\rangle_B\}$ basis "erases" the "which way" information (whether the state of A is $|\varphi_i\rangle_A$ or $|\varphi_j\rangle_A$). In this sense, the HJW theorem characterizes the general quantum eraser. The moral, once again, is that *information is physical* — the information acquired by measuring system B, when relayed to A, changes the physical description of a state of A.

2.6 How far apart are two quantum states?

2.6.1 Fidelity and Uhlmann's theorem

The distinguishability of two pure states $|\psi\rangle$ and $|\varphi\rangle$ is quantified by the deviation from 1 of their *overlap* $|\langle \varphi | \psi \rangle|^2$, also called *fidelity*. For two density operators ρ and σ the fidelity is defined by

$$F(\boldsymbol{\rho}, \boldsymbol{\sigma}) \equiv \left(\operatorname{tr} \sqrt{\boldsymbol{\rho}^{\frac{1}{2}} \boldsymbol{\sigma} \boldsymbol{\rho}^{\frac{1}{2}}} \right)^2 . \qquad (2.122)$$

(Some authors use the name "fidelity" for the square root of this quantity.) The fidelity is nonnegative, vanishes if ρ and σ have support on mutually orthogonal subspaces, and attains its maximum value 1 if and only if the two states are identical. If $\rho = |\psi\rangle\langle\psi|$ is a pure state, then the fidelity is

$$F(\boldsymbol{\rho}, \boldsymbol{\sigma}) = \langle \psi | \boldsymbol{\sigma} | \psi \rangle. \tag{2.123}$$

Suppose we perform an orthogonal measurement on ρ with the two outcomes: "Yes" if the state if $|\psi\rangle$, "No" if the state is orthogonal to $|\psi\rangle$. Then the fidelity is the probability that the outcome is "Yes."

We may also express the fidelity in terms of the L^1 norm,

$$F(\boldsymbol{\rho}, \boldsymbol{\sigma}) = \left\| \boldsymbol{\sigma}^{\frac{1}{2}} \boldsymbol{\rho}^{\frac{1}{2}} \right\|_{1}^{2}, \qquad (2.124)$$

where

$$\|\boldsymbol{A}\|_{1} = \operatorname{tr} \sqrt{\boldsymbol{A}^{\dagger} \boldsymbol{A}}.$$
 (2.125)

The L^1 norm is also sometimes called the *trace norm*. (For Hermitian A, $||A||_1$ is just the sum of the absolute values of its eigenvalues.) The fidelity $F(\rho, \sigma)$ is actually symmetric in its two arguments, although the

symmetry is not manifest in eq. (2.122). To verify the symmetry, note that for any Hermitian A and B, the L^1 norm obeys

$$\|AB\|_1 = \|BA\|_1$$
. (2.126)

This holds because **BAAB** and **ABBA** have the same eigenvalues if $|\psi\rangle$ is an eigenstate of **ABBA** with eigenvalue λ , the **BA** $|\psi\rangle$ is an eigenstate of **BAAB** with eigenvalue λ .

It is useful to know how the fidelity of two density operators is related to the overlap of their *purifications*. We say that $|\Phi\rangle_{AB}$ is a purification of the density operator ρ_A if

$$\boldsymbol{\rho}_A = \operatorname{tr}_B\left(|\Phi\rangle\langle\Phi|\right). \tag{2.127}$$

 \mathbf{If}

$$\boldsymbol{\rho} = \sum_{i} p_{i} |i\rangle \langle i| \tag{2.128}$$

(where $|i\rangle_A$ is an orthonormal basis for system A), then a particular purification of ρ has the form

$$|\Phi_{\rho}\rangle = \sum_{i} \sqrt{p_i} |i\rangle_A \otimes |i\rangle_B \tag{2.129}$$

where $|i\rangle_B$ is an orthonormal basis for system *B*. According to the HJW theorem, a general purification has the form

$$|\Phi_{\rho}(\boldsymbol{V})\rangle = \boldsymbol{I} \otimes \boldsymbol{V} |\Phi_{\rho}\rangle \tag{2.130}$$

where V is unitary, which may also be written

$$|\Phi_{\rho}(\boldsymbol{V})\rangle = \left(\boldsymbol{\rho}^{\frac{1}{2}} \otimes \boldsymbol{V}\right) |\tilde{\Phi}\rangle,$$
 (2.131)

where $|\tilde{\Phi}\rangle$ is the unconventionally normalized maximally entangled state

$$|\tilde{\Phi}\rangle_{AB} = \sum_{i} |i\rangle_A \otimes |i\rangle_B.$$
 (2.132)

If ρ and σ are two density operators on A, the inner product of their purifications on AB can be expressed as

$$\langle \Phi_{\boldsymbol{\sigma}}(\boldsymbol{W}) | \Phi_{\boldsymbol{\rho}}(\boldsymbol{V}) \rangle = \langle \tilde{\Phi} | \boldsymbol{\sigma}^{\frac{1}{2}} \boldsymbol{\rho}^{\frac{1}{2}} \otimes \boldsymbol{W}^{\dagger} \boldsymbol{V} | \tilde{\Phi} \rangle.$$
 (2.133)

Noting that

$$\boldsymbol{U} \otimes \boldsymbol{I} | \tilde{\Phi} \rangle = \sum_{i,j} | j \rangle \otimes | i \rangle U_{ji} = \sum_{i,j} | j \rangle \otimes | i \rangle U_{ij}^T = \boldsymbol{I} \otimes \boldsymbol{U}^T | \tilde{\Phi} \rangle, \quad (2.134)$$

we have

$$\langle \Phi_{\boldsymbol{\sigma}}(\boldsymbol{W}) | \Phi_{\boldsymbol{\rho}}(\boldsymbol{V}) \rangle = \langle \tilde{\Phi} | \boldsymbol{\sigma}^{\frac{1}{2}} \boldsymbol{\rho}^{\frac{1}{2}} \boldsymbol{U} \otimes \boldsymbol{I} | \tilde{\Phi} \rangle = \operatorname{tr} \left(\boldsymbol{\sigma}^{\frac{1}{2}} \boldsymbol{\rho}^{\frac{1}{2}} \boldsymbol{U} \right), \qquad (2.135)$$

where $\boldsymbol{U} = \left(\boldsymbol{W}^{\dagger} \boldsymbol{V} \right)^{T}$.

Now we may use the *polar decomposition*

$$\boldsymbol{A} = \boldsymbol{U}' \sqrt{\boldsymbol{A}^{\dagger} \boldsymbol{A}}, \qquad (2.136)$$

where U' is unitary, to rewrite the inner product as

$$\langle \Phi_{\boldsymbol{\sigma}}(\boldsymbol{W}) | \Phi_{\boldsymbol{\rho}}(\boldsymbol{V}) \rangle = \operatorname{tr}\left(\boldsymbol{U}\boldsymbol{U}'\sqrt{\boldsymbol{\rho}^{\frac{1}{2}}\boldsymbol{\sigma}\boldsymbol{\rho}^{\frac{1}{2}}}\right) = \sum_{a} \lambda_{a} \langle a | \boldsymbol{U}\boldsymbol{U}' | a \rangle, \quad (2.137)$$

where $\{\lambda_a\}$ are the nonnegative eigenvalues of $\sqrt{\rho^{\frac{1}{2}}\sigma\rho^{\frac{1}{2}}}$ and $\{|a\rangle\}$ are the corresponding eigenvectors. It is now evident that the inner product has the largest possible absolute value when we choose $U = U'^{-1}$, and hence we conclude

$$F(\boldsymbol{\rho},\boldsymbol{\sigma}) = \left(\operatorname{tr} \sqrt{\boldsymbol{\rho}^{\frac{1}{2}} \boldsymbol{\sigma} \boldsymbol{\rho}^{\frac{1}{2}}} \right)^2 = \max_{\boldsymbol{V},\boldsymbol{W}} |\langle \Phi_{\boldsymbol{\sigma}}(\boldsymbol{W}) | \Phi_{\boldsymbol{\rho}}(\boldsymbol{V}) \rangle|^2.$$
(2.138)

The fidelity of two density operators is the maximal possible overlap of their purifications, a result called *Uhlmann's theorem*.

One corollary of Uhlmann's theorem is the *monotonicity* of fidelity:

$$F(\boldsymbol{\rho}_{AB}, \boldsymbol{\rho}_{AB}) \le F(\boldsymbol{\rho}_{A}, \boldsymbol{\rho}_{A}), \qquad (2.139)$$

which says that tracing out a subsystem cannot decrease the fidelity of two density operators. Monotonicity means, unsurprisingly, that throwing away a subsystem does not make two quantum states easier to distinguish. It follows from Uhlmann's theorem because any state purifying ρ_{AB} also provides a purification of A; therefore the optimal overlap of the purifications of ρ_{AB} and σ_{AB} is surely achievable by purifications of ρ_A and σ_A .

2.6.2 Relations among distance measures

There are other possible ways besides fidelity for quantifying the difference between quantum states ρ and σ , such as the distance between the states using the L^1 or L^2 norm,

$$\|\boldsymbol{\rho} - \boldsymbol{\sigma}\|_1$$
 or $\|\boldsymbol{\rho} - \boldsymbol{\sigma}\|_2$, (2.140)

where the L^2 norm of an operator is defined by

$$\|\boldsymbol{A}\|_2 = \sqrt{\mathrm{tr}\boldsymbol{A}^{\dagger}\boldsymbol{A}}.$$
 (2.141)

(For Hermitian \mathbf{A} , $\|\mathbf{A}\|_2$ is the square root of the sum of the squares of its eigenvalues.) The L^1 distance is a particularly natural measure of state distinguishability, because (as shown in Exercise 2.5) it can be interpreted as the distance between the corresponding probability distributions achieved by the optimal measurement for distinguishing the states. Although the fidelity, L^1 distance, and L^2 distance are not simply related to one another in general, there are useful inequalities relating these different measures.

If $\{|\lambda_i|, i = 0, 1, 2, \dots d-1\}$ denotes the eigenvalues of $\sqrt{\mathbf{A}^{\dagger} \mathbf{A}}$, then

$$\|\boldsymbol{A}\|_{1} = \sum_{i=0}^{d-1} |\lambda_{i}|; \quad \|\boldsymbol{A}\|_{2} = \sqrt{\sum_{i=0}^{d-1} |\lambda_{i}|^{2}}.$$
 (2.142)

If we regard $\|\mathbf{A}\|_1$ as the inner product of the two vectors (1, 1, 1, ..., 1)and $(|\lambda_0|, |\lambda_1|, ..., |\lambda_{d-1}|)$, then from the Cauchy-Schwarz inequality we find

$$\|\mathbf{A}\|_{1} \le \sqrt{d} \|\mathbf{A}\|_{2}.$$
 (2.143)

Because of the factor of \sqrt{d} on the right hand side, for a high-dimensional system density operators which are close together in the L^2 norm might not be close in the L^1 norm.

There is, however, a dimension-independent inequality relating the L^1 distance between ρ and σ and the L^2 distance between their square roots:

$$\|\sqrt{\boldsymbol{\rho}} - \sqrt{\boldsymbol{\sigma}}\|_2^2 \le \|\boldsymbol{\rho} - \boldsymbol{\sigma}\|_1.$$
(2.144)

To derive this inequality, we first expand the difference of square roots in its basis of eigenvectors,

$$\sqrt{\rho} - \sqrt{\sigma} = \sum_{i} \lambda_{i} |i\rangle \langle i|,$$
 (2.145)

and note that the absolute value of this difference may be written as

$$|\sqrt{\rho} - \sqrt{\sigma}| = \sum_{i} |\lambda_{i}| |i\rangle \langle i| = (\sqrt{\rho} - \sqrt{\sigma}) \mathbf{U} = \mathbf{U} (\sqrt{\rho} - \sqrt{\sigma}), \quad (2.146)$$

where

$$\boldsymbol{U} = \sum_{i} \operatorname{sign}(\lambda_i) |i\rangle \langle i|.$$
(2.147)

Using

$$\boldsymbol{\rho} - \boldsymbol{\sigma} = \frac{1}{2} \left(\sqrt{\boldsymbol{\rho}} - \sqrt{\boldsymbol{\sigma}} \right) \left(\sqrt{\boldsymbol{\rho}} + \sqrt{\boldsymbol{\sigma}} \right) + \frac{1}{2} \left(\sqrt{\boldsymbol{\rho}} + \sqrt{\boldsymbol{\sigma}} \right) \left(\sqrt{\boldsymbol{\rho}} - \sqrt{\boldsymbol{\sigma}} \right)$$
(2.148)

and the cyclicity of the trace, we find

$$\operatorname{tr}\left(\boldsymbol{\rho}-\boldsymbol{\sigma}\right)\boldsymbol{U} = \operatorname{tr}|\sqrt{\boldsymbol{\rho}}-\sqrt{\boldsymbol{\sigma}}|\left(\sqrt{\boldsymbol{\rho}}+\sqrt{\boldsymbol{\sigma}}\right) = \sum_{i}|\lambda_{i}|\langle i|\sqrt{\boldsymbol{\rho}}+\sqrt{\boldsymbol{\sigma}}|i\rangle$$
$$\geq \sum_{i}|\lambda_{i}|\left|\langle i|\sqrt{\boldsymbol{\rho}}-\sqrt{\boldsymbol{\sigma}}|i\rangle\right| = \sum_{i}|\lambda_{i}|^{2} = \|\sqrt{\boldsymbol{\rho}}-\sqrt{\boldsymbol{\sigma}}\|_{2}^{2}.$$
$$(2.149)$$

Finally, using

$$\|\boldsymbol{\rho} - \boldsymbol{\sigma}\|_1 = \operatorname{tr}|\boldsymbol{\rho} - \boldsymbol{\sigma}| \ge \operatorname{tr}(\boldsymbol{\rho} - \boldsymbol{\sigma})\boldsymbol{U}, \qquad (2.150)$$

which is true for any unitary U, we obtain eq.(2.144).

This L^2 distance between square roots can be related to fidelity. First we note that

$$\left\|\sqrt{\rho} - \sqrt{\sigma}\right\|_{2}^{2} = \operatorname{tr}\left(\sqrt{\rho} - \sqrt{\sigma}\right)^{2} = 2 - 2 \operatorname{tr}\left(\sqrt{\rho}\sqrt{\sigma}\right), \qquad (2.151)$$

since tr $\rho = \text{tr } \sigma = 1$. From the polar decomposition $A = U\sqrt{A^{\dagger}A}$ (where U is unitary), we see that tr $\sqrt{A^{\dagger}A} \ge |\text{tr } A|$, and therefore

$$\sqrt{F(\boldsymbol{\rho}, \boldsymbol{\sigma})} = \operatorname{tr}\sqrt{\boldsymbol{\rho}^{\frac{1}{2}} \boldsymbol{\sigma} \boldsymbol{\rho}^{\frac{1}{2}}} \ge \left|\operatorname{tr}\left(\sqrt{\boldsymbol{\rho}} \sqrt{\boldsymbol{\sigma}}\right)\right|$$
 (2.152)

and hence

$$\sqrt{F(\boldsymbol{\rho},\boldsymbol{\sigma})} \ge 1 - \frac{1}{2} \left\| \sqrt{\boldsymbol{\rho}} - \sqrt{\boldsymbol{\sigma}} \right\|_2^2 \ge 1 - \frac{1}{2} \left\| \boldsymbol{\rho} - \boldsymbol{\sigma} \right\|_1.$$
(2.153)

Eq.(2.153) tells us that if ρ and σ are close to one another in the L^1 norm, then their fidelity is close to one.

The L^1 distance also provides an *upper* bound on fidelity:

$$F(\boldsymbol{\rho}, \boldsymbol{\sigma}) \le 1 - \frac{1}{4} \|\boldsymbol{\rho} - \boldsymbol{\sigma}\|_{1}^{2}.$$
(2.154)

To derive eq.(2.154) we first show that it holds with equality for pure states. Any two vectors $|\psi\rangle$ and $|\varphi\rangle$ lie in some two-dimensional subspace, and by choosing a basis and phase conventions appropriately we may write

$$|\psi\rangle = \begin{pmatrix} \cos\theta/2\\ \sin\theta/2 \end{pmatrix}, \quad |\varphi\rangle = \begin{pmatrix} \sin\theta/2\\ \cos\theta/2 \end{pmatrix}, \quad (2.155)$$

40

2.7 Summary

for some angle θ ; then

$$|\psi\rangle\langle\psi| - |\varphi\rangle\langle\varphi| = \begin{pmatrix} \cos\theta & 0\\ 0 & -\cos\theta \end{pmatrix}.$$
 (2.156)

and

$$|\langle \varphi | \psi \rangle|^2 = \sin^2 \theta. \tag{2.157}$$

Therefore,

$$\||\psi\rangle\langle\psi| - |\varphi\rangle\langle\varphi|\|_{1}^{2} = (2\cos\theta)^{2} = 4\left(1 - F(|\psi\rangle\langle\psi|, |\varphi\rangle\langle\varphi|)\right). \quad (2.158)$$

Next, note that L^1 distance, like fidelity, is monotonic:

$$\|\boldsymbol{\rho}_{AB} - \boldsymbol{\sigma}_{AB}\|_{1} \ge \|\boldsymbol{\rho}_{A} - \boldsymbol{\sigma}_{A}\|_{1}.$$
(2.159)

This follows because the L^1 distance is the optimal distance between probability distributions when we measure the two states, and the optimal measurement for distinguishing ρ_A and σ_A is also a possible measurement for ρ_{AB} and σ_{AB} , one that happens to act trivially on B.

Finally, we invoke Uhlmann's theorem. If ρ_{AB} and σ_{AB} are the purifications of ρ_A and σ_A with the largest possible overlap, then

$$F(\boldsymbol{\rho}_{A},\boldsymbol{\sigma}_{A}) = F(\boldsymbol{\rho}_{AB},\boldsymbol{\sigma}_{AB}) = 1 - \frac{1}{4} \|\boldsymbol{\rho}_{AB} - \boldsymbol{\sigma}_{AB}\|_{1}^{2}$$
$$\leq 1 - \frac{1}{4} \|\boldsymbol{\rho}_{A} - \boldsymbol{\sigma}_{A}\|_{1}^{2}, \qquad (2.160)$$

where the first equality uses Uhlmann's theorem, the second uses eq.(2.158), and the final inequality uses monotonicity. By combining eq.(2.160) with eq.(2.153) we have

$$1 - \sqrt{F(\boldsymbol{\rho}, \boldsymbol{\sigma})} \le \frac{1}{2} \|\boldsymbol{\rho} - \boldsymbol{\sigma}\|_1 \le \sqrt{1 - F(\boldsymbol{\rho}, \boldsymbol{\sigma})}; \qquad (2.161)$$

hence ρ and σ are close to one another in the L^1 norm if and only if their fidelity is close to one.

2.7 Summary

Axioms. The arena of quantum mechanics is a Hilbert space \mathcal{H} . The fundamental assumptions are:

- (1) A state is a ray in \mathcal{H} .
- (2) An observable is a self-adjoint operator on \mathcal{H} .
- (3) A measurement is an orthogonal projection.

- (4) *Time evolution* is *unitary*.
- (5) A composite system AB is described by the tensor product $\mathcal{H}_A \otimes \mathcal{H}_B$.

Density operator. But if we confine our attention to only a portion of a larger quantum system, assumptions (1)-(4) need not be satisfied. In particular, a quantum state is described not by a ray, but by a density operator ρ , a nonnegative operator with unit trace. The density operator is *pure* (and the state can be described by a ray) if $\rho^2 = \rho$; otherwise, the state is *mixed*. An observable M has expectation value tr($M\rho$) in this state.

Qubit. A quantum system with a two-dimensional Hilbert space is called a *qubit*. The general density matrix of a qubit is

$$\boldsymbol{\rho}(\vec{P}) = \frac{1}{2} (\boldsymbol{I} + \vec{P} \cdot \vec{\boldsymbol{\sigma}})$$
(2.162)

where \vec{P} is a three-component vector of length $|\vec{P}| \leq 1$. Pure states have $|\vec{P}| = 1$.

Schmidt decomposition. For any pure state $|\psi\rangle_{AB}$ of a bipartite system, there are orthonormal bases $\{|i\rangle_A\}$ for \mathcal{H}_A and $\{|i'\rangle_B\}$ for \mathcal{H}_B such that

$$|\psi\rangle_{AB} = \sum_{i} \sqrt{p_i} \ |i\rangle_A \otimes |i'\rangle_B; \tag{2.163}$$

Eq.(2.163) is called the Schmidt decomposition of $|\psi\rangle_{AB}$. In a bipartite pure state, subsystems A and B separately are described by density operators ρ_A and ρ_B ; it follows from eq.(2.163) that ρ_A and ρ_B have the same nonvanishing eigenvalues (the p_i 's). The number of nonvanishing eigenvalues is called the Schmidt number of $|\psi\rangle_{AB}$. A bipartite pure state is said to be entangled if its Schmidt number is greater than one.

Ensembles. The density operators on a Hilbert space form a convex set, and the pure states are the *extremal points* of the set. A mixed state of a system A can be prepared as an *ensemble* of pure states in many different ways, all of which are experimentally indistinguishable if we observe system A alone. Given any mixed state ρ_A of system A, any preparation of ρ_A as an ensemble of pure states can be realized in principle by performing a measurement in another system B with which A is entangled. In fact given many such preparations of ρ_A , there is a single entangled state of A and B such that any one of these preparations can be realized by measuring a suitable observable in B (the *HJW theorem*). By measuring in system B and reporting the measurement outcome to system A, we can extract from the mixture a pure state chosen from one of the ensembles.
Fidelity. The fidelity $F(\boldsymbol{\rho}, \boldsymbol{\sigma}) = \left\| \boldsymbol{\sigma}^{\frac{1}{2}} \boldsymbol{\rho}^{\frac{1}{2}} \right\|_{1}^{2}$, quantifies the distinguishability of two density operators — it is the maximum overlap achieved by their purifications (*Uhlmann's theorem*). The fidelity $F(\boldsymbol{\rho}, \boldsymbol{\sigma})$ is close to one if and only if the L^{1} distance $\| \boldsymbol{\rho} - \boldsymbol{\sigma} \|_{1}$ is small.

Further important ideas are developed in the Exercises.

2.8 Exercises

2.1 Fidelity of measurement

a) For two states $|\psi_1\rangle$ and $|\psi_2\rangle$ in an N-dimensional Hilbert space, define the relative angle θ between the states by

$$|\langle \psi_2 | \psi_1 \rangle| \equiv \cos \theta , \qquad (2.164)$$

where $0 \le \theta \le \pi/2$. Suppose that the two states are selected *at random*. Find the probability distribution $p(\theta)d\theta$ for the relative angle. **Hint**: We can choose a basis such that

$$\begin{aligned} |\psi_1\rangle &= (1, \vec{0}) \\ |\psi_2\rangle &= (e^{i\varphi}\cos\theta, \psi_2^{\perp}) . \end{aligned}$$
 (2.165)

(2.166)

"Selected at random" means that the probability distribution for the normalized vector $|\psi_2\rangle$ is uniform on the (real) (2N-1)sphere (this is the unique distribution that is invariant under arbitrary unitary transformations). Note that, for fixed θ , $e^{i\varphi}$ parametrizes a circle of radius $\cos \theta$, and $|\psi_2^{\perp}\rangle$ is a vector that lies on a 2N-3 sphere of radius $\sin \theta$.

b) A density operator ρ is said to approximate a pure state $|\psi\rangle$ with fidelity

$$F = \langle \psi | \boldsymbol{\rho} | \psi \rangle . \tag{2.167}$$

Imagine that a state $|\psi_1\rangle$ in an *N*-dimensional Hilbert space is selected at random, and we guess at random that the state is $|\psi_2\rangle$. On the average, what will be the fidelity of our guess?

c) When we measure, we collect information and cause a disturbance – an unknown state is replaced by a different state that is known. Suppose that the state $|\psi\rangle$ is selected at random, and then an orthogonal measurement is performed, projecting onto an orthonormal basis $\{|E_a\rangle\}$. After the measurement, the

state (averaged over all possible outcomes) is described by the density matrix

$$\boldsymbol{\rho} = \sum_{a} \boldsymbol{E}_{a} |\psi\rangle \langle \psi | \boldsymbol{E}_{a} , \qquad (2.168)$$

where $E_a = |E_a\rangle \langle E_a|$; this ρ approximates $|\psi\rangle$ with fidelity

$$F = \sum_{a} \left(\langle \psi | \boldsymbol{E}_{a} | \psi \rangle \right)^{2} . \qquad (2.169)$$

Evaluate F, averaged over the choice of $|\psi\rangle$. **Hint:** Use Bayes's rule and the result from (a) to find the probability distribution for the angle θ between the state $|\psi\rangle$ and the projected state $\mathbf{E}_a |\psi\rangle/||\mathbf{E}_a |\psi\rangle||$. Then evaluate $\langle \cos^2 \theta \rangle$ in this distribution.

Remark: The improvement in F in the answer to (c) compared to the answer to (b) is a crude measure of how much we learned by making the measurement.

2.2 Measurement without disturbance?

Charlie prepares the system A in one of two nonorthogonal states, $|\varphi\rangle_A$ or $|\tilde{\varphi}\rangle_A$, and he challenges Alice to collect some information about which state he prepared without in any way disturbing the state. Alice has an idea about how to meet the challenge.

Alice intends to prepare a second "ancillary" system B in the state $|\beta\rangle_B$, and then apply to the composite system AB a unitary transformation U that acts according to

$$\begin{aligned} \boldsymbol{U} : |\varphi\rangle_A \otimes |\beta\rangle_B &\to |\varphi\rangle_A \otimes |\beta'\rangle_B \\ |\tilde{\varphi}\rangle_A \otimes |\beta\rangle_B &\to |\tilde{\varphi}\rangle_A \otimes |\tilde{\beta}'\rangle_B, \end{aligned}$$
(2.170)

which does indeed leave the state of system A undisturbed. Then she plans to perform a measurement on system B that is designed to distinguish the states $|\beta'\rangle_B$ and $|\tilde{\beta}'\rangle_B$.

- a) What do you think of Alice's idea? **Hint**: What does the unitarity of U tell you about how the states $|\beta'\rangle_B$ and $|\tilde{\beta}'\rangle_B$ are related to one another?
- b) Would you feel differently if the states $|\varphi\rangle_A$ and $|\tilde{\varphi}\rangle_A$ were orthogonal?

2.3 Quantum bit commitment

The Yankees are playing the Dodgers in the World Series. Alice is sure that she knows who will win. Alice doesn't like Bob, so she doesn't want to tell him who the winner will be. But after the Series is over, Alice wants to be able to convince Bob that she knew the outcome all along. What to do?

Bob suggests that Alice write down her choice (0 if the Yankees will win, 1 if the Dodgers will win) on a piece of paper, and lock the paper in a box. She is to give the box to Bob, but she will keep the key for herself. Then, when she is ready to reveal her choice, she will send the key to Bob, allowing him to open the box and read the paper.

Alice rejects this proposal. She doesn't trust Bob, and she knows that he is a notorious safe cracker. Who's to say whether he will be able to open the box and sneak a look, even if he doesn't have the key?

Instead, Alice proposes to certify her honesty in another way, using quantum information. To *commit* to a value $a \in \{0, 1\}$ of her bit, she prepares one of two distinguishable density operators (ρ_0 or ρ_1) of the bipartite system AB, sends system B to Bob, and keeps system A for herself. Later, to *unveil* her bit, Alice sends system A to Bob, and he performs a measurement to determine whether the state of AB is ρ_0 or ρ_1 . This protocol is called *quantum bit commitment*.

We say that the protocol is *binding* if, after commitment, Alice is unable to change the value of her bit. We say that the protocol is *concealing* if, after commitment and before unveiling, Bob is unable to discern the value of the bit. The protocol is *secure* if it is both binding and concealing.

Show that if a quantum bit commitment protocol is concealing, then it is not binding. Thus quantum bit commitment is insecure.

Hint: First argue that without loss of generality, we may assume that the states ρ_0 and ρ_1 are both pure. Then apply the HJW Theorem.

Remark: Note that the conclusion that quantum bit commitment is insecure still applies even if the shared bipartite state (ρ_0 or ρ_1) is prepared during many rounds of quantum communication between Alice and Bob, where in each round one party performs a quantum operation on his/her local system and on a shared message system, and then sends the message system to the other party.

2.4 Completeness of subsystem correlations

Consider a bipartite system AB. Suppose that many copies of the (not necessarily pure) state ρ_{AB} have been prepared. An observer Alice with access only to subsystem A can measure the expectation

value of any observable of the form $M_A \otimes I_B$, while an observer Bob with access only to subsystem B can measure the expectation value of any observable of the form $I_A \otimes N_B$. Neither of these observers, working alone, can expect to gain enough information to determine the joint state ρ_{AB} .

But now suppose that Alice and Bob can communicate, exchanging (classical) information about how their measurement outcomes are *correlated*. Thereby, they can jointly determine the expectation value of any observable of the form $M_A \otimes N_B$ (an observable whose eigenstates are separable direct products states of the form $|\varphi\rangle_A \otimes |\chi\rangle_B$).

The point of this exercise is to show that if Alice and Bob have complete knowledge of the nature of the correlations between subsystems A and B (know the expectation values of any tensor product observable $M_A \otimes N_B$), then in fact they know everything about the bipartite state ρ_{AB} – there will be no surprises when they measure entangled observables, those whose eigenstates are entangled states.

- a) Let $\{M_a, a = 1, 2, ..., d^2\}$ denote a set of d^2 linearly independent self-adjoint operators acting on a Hilbert space \mathcal{H} of dimension d. Show that if ρ is a density operator acting on \mathcal{H} , and tr (ρM_a) is known for each a, then $\langle \varphi | \rho | \varphi \rangle$ is known for any vector $|\varphi\rangle$ in \mathcal{H} .
- b) Show that if $\langle \varphi | \rho | \varphi \rangle$ is known for each vector $| \varphi \rangle$, then ρ is completely known.
- c) Show that if $\{M_a\}$ denotes a basis for self-adjoint operators on \mathcal{H}_A , and $\{N_b\}$ denotes a basis for self-adjoint operators on \mathcal{H}_B , then $\{M_a \otimes N_b\}$ is a basis for the self-adjoint operators on $\mathcal{H}_A \otimes \mathcal{H}_B$.

Remark: It follows from (c) alone that the correlations of the "local" observables determine the expectation values of all the observables. Parts (a) and (b) serve to establish that ρ is completely determined by the expectation values of a complete set of observables.

- d) State and prove the result corresponding to (c) that applies to a *n*-part system with Hilbert space $\mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \otimes \mathcal{H}_n$.
- e) Discuss how the world would be different if quantum states resided in a real Hilbert space rather than a complex Hilbert space. Consider, in particular, whether (c) is true for symmetric operators acting on a real vector space.

2.5 Optimal measurement distinguishing two quantum states

Consider two quantum states described by density operators ρ and σ in an *d*-dimensional Hilbert space, and consider the complete orthogonal measurement $\{E_a, a = 0, 1, 2, \dots d-1\}$, where the E_a 's are one-dimensional projectors satisfying

$$\sum_{a=0}^{d-1} \boldsymbol{E}_a = \boldsymbol{I}.$$
 (2.171)

When the measurement is performed, outcome *a* occurs with probability $p_a = \text{tr } \rho E_a$ if the state is ρ and with probability $q_a = \text{tr } \sigma E_a$ if the state is σ .

The L^1 distance between the two probability distributions is defined as

$$d(p,q) \equiv \|p-q\|_1 \equiv \sum_{a=0}^{d-1} |p_a - q_a| ; \qquad (2.172)$$

this distance is zero if the two distributions are identical, and attains its maximum value two if the two distributions have support on disjoint sets.

a) Show that

$$d(p,q) \leq \sum_{i=0}^{d-1} |\lambda_i| = \|\boldsymbol{\rho} - \boldsymbol{\sigma}\|_1 \equiv d(\boldsymbol{\rho}, \boldsymbol{\sigma}), \qquad (2.173)$$

where the λ_i 's are the eigenvalues of the Hermitian operator $\rho - \sigma$. Hint: Working in the basis in which $\rho - \sigma$ is diagonal, find an expression for $|p_a - q_a|$, and then find an upper bound on $|p_a - q_a|$. Finally, use the completeness property eq.(2.171) to bound d(p,q).

- b) Find a choice for the orthogonal projector $\{E_a\}$ that saturates the upper bound eq.(2.173).
- c) If the states $\rho = |\psi\rangle\langle\psi|$ and $\sigma = |\varphi\rangle\langle\varphi|$ are pure, show that

$$d(p,q) \le 2 \||\psi\rangle - |\varphi\rangle\| \tag{2.174}$$

where $\|\cdot\|$ denotes the Hilbert space norm.

2.6 What probability distributions are consistent with a mixed state?

A density operator ρ , expressed in the orthonormal basis $\{|\alpha_i\rangle\}$ that diagonalizes it, is

$$\boldsymbol{\rho} = \sum_{i} p_{i} |\alpha_{i}\rangle \langle \alpha_{i}|. \qquad (2.175)$$

We would like to realize this density operator as an ensemble of pure states $\{|\varphi_{\mu}\rangle\}$, where $|\varphi_{\mu}\rangle$ is prepared with a specified probability q_{μ} . This preparation is possible if the $|\varphi_{\mu}\rangle$'s can be chosen so that

$$\boldsymbol{\rho} = \sum_{\mu} q_{\mu} |\varphi_{\mu}\rangle \langle \varphi_{\mu}|. \qquad (2.176)$$

We say that a probability vector q (a vector whose components are nonnegative real numbers that sum to 1) is *majorized* by a probability vector p (denoted $q \prec p$), if there exists a *doubly stochastic* matrix D such that

$$q_{\mu} = \sum_{i} D_{\mu i} \ p_{i}. \tag{2.177}$$

A matrix is doubly stochastic if its entries are nonnegative real numbers such that $\sum_{\mu} D_{\mu i} = \sum_{i} D_{\mu i} = 1$. That the columns sum to one assures that D maps probability vectors to probability vectors (*i.e.*, is *stochastic*). That the rows sum to one assures that D maps the uniform distribution to itself. Applied repeatedly, D takes any input distribution closer and closer to the uniform distribution (unless Dis a permutation, with one nonzero entry in each row and column). Thus we can view majorization as a partial order on probability vectors such that $q \prec p$ means that q is more nearly uniform than p(or equally close to uniform, in the case where D is a permutation).

Show that normalized pure states $\{|\varphi_{\mu}\rangle\}$ exist such that eq.(2.176) is satisfied if and only if $q \prec p$, where p is the vector of eigenvalues of ρ .

Hint: Recall that, according to the HJW Theorem, if eq.(2.175) and eq.(2.176) are both satisfied then there is a unitary matrix $V_{\mu i}$ such that

$$\sqrt{q_{\mu}} |\varphi_{\mu}\rangle = \sum_{i} \sqrt{p_{i}} V_{\mu i} |\alpha_{i}\rangle.$$
(2.178)

You may also use (but need not prove) Horn's Lemma: if $q \prec p$, then there exists a unitary (in fact, orthogonal) matrix $U_{\mu i}$ such that q = Dp and $D_{\mu i} = |U_{\mu i}|^2$.

2.7 Alice does Bob a favor

Alice, in Anaheim, and Bob, in Boston, share a bipartite pure state $|\Psi\rangle$, which can be expressed in the Schmidt form

$$|\Psi\rangle = \sum_{i} \sqrt{p_i} |\alpha_i\rangle \otimes |\beta_i\rangle, \qquad (2.179)$$

where $\{|\alpha_i\rangle\}$ is an orthonormal basis for Alice's system A, $\{|\beta_i\rangle\}$ is an orthonormal basis for Bob's system B, and the $\{p_i\}$ are nonnegative real numbers summing to 1. Bob is supposed to perform a complete orthogonal local measurement on B, characterized by the set of projectors $\{\boldsymbol{E}_a^B\}$ — if the measurement outcome is a, then Bob's measurement prepares the state

$$|\Psi\rangle \mapsto |\Psi_a\rangle = \frac{\left(\boldsymbol{I} \otimes \boldsymbol{E}_a^B\right)|\Psi\rangle}{\left\langle\Psi\right| \left(\boldsymbol{I} \otimes \boldsymbol{E}_a^B\right)|\Psi\rangle^{1/2}}.$$
 (2.180)

 $|\Psi_a\rangle$ can also be expressed in the Schmidt form if we choose appropriate orthonormal bases for A and B that depend on the measurement outcome. The new Schmidt basis elements can be written as

$$|\alpha'_{a,i}\rangle = \boldsymbol{U}_a^A |\alpha_i\rangle, \quad |\beta'_{a,i}\rangle = \boldsymbol{U}_a^B |\beta_i\rangle, \quad (2.181)$$

where $\boldsymbol{U}_{a}^{A}, \boldsymbol{U}_{a}^{B}$ are unitary.

Unfortunately, Bob's measurement apparatus is broken, though he still has the ability to perform local unitary transformations on B. Show that Alice can help Bob out by performing a measurement that is "locally equivalent" to Bob's. That is, there are orthogonal projectors $\{\boldsymbol{E}_a^A\}$ and unitary transformations $\boldsymbol{V}_a^A, \boldsymbol{V}_a^B$ such that

$$|\Psi_a\rangle = \boldsymbol{V}_a^A \otimes \boldsymbol{V}_a^B \left(\frac{\left(\boldsymbol{E}_a^A \otimes \boldsymbol{I}\right)|\Psi\rangle}{\left\langle\Psi\right| \left(\boldsymbol{E}_a^A \otimes \boldsymbol{I}\right)|\Psi\rangle^{1/2}}\right)$$
(2.182)

for each a, and furthermore, both Alice's measurement and Bob's yield outcome a with the same probability. This means that instead of Bob doing the measurement, the same effect can be achieved if Alice measures instead, tells Bob the outcome, and both Alice and Bob perform the appropriate unitary transformations. Construct E_a^A (this is most conveniently done by expressing both E_a^A and E_a^B in the Schmidt bases for $|\Psi\rangle$) and express V_a^A and V_a^B in terms of U_a^A and U_a^B .

Remark: This result shows that for any protocol involving local operations and "two-way" classical communication (2-LOCC) that transforms an initial bipartite pure state to a final bipartite pure

state, the same transformation can be achieved by a "one-way" (1-LOCC) protocol in which all classical communication is from Alice to Bob (the Lo-Popescu Theorem). In a two-way LOCC protocol, Alice and Bob take turns manipulating the state for some finite (but arbitrarily large) number of rounds. In each round, one party or the other performs a measurement on her/his local system and broadcasts the outcome to the other party. Either party might use a local "ancilla" system in performing the measurement, but we may include all ancillas used during the protocol in the bipartite pure state $|\Psi\rangle$. Though a party might discard information about the measurement outcome, or fail to broadcast the information to the other party, we are entitled to imagine that the complete information about the outcomes is known to both parties at each step (incomplete information is just equivalent to the special case in which the parties choose not to use all the information). Thus the state is pure after each step.

The solution to the exercise shows that a round of a 2-LOCC protocol in which Bob measures can be simulated by an operation performed by Alice and a local unitary applied by Bob. Thus, we can allow Alice to perform all the measurements herself. When she is through she sends all the outcomes to Bob, and he can apply the necessary product of unitary transformations to complete the protocol.

2.8 The power of noncontextuality

We may regard a quantum state as an assignment of probabilities to projection operators. That is, according to Born's rule, if ρ is a density operator and \boldsymbol{E} is a projector, then $p(\boldsymbol{E}) = \operatorname{tr}(\rho \boldsymbol{E})$ is the probability that the outcome \boldsymbol{E} occurs, if \boldsymbol{E} is one of a complete set of orthogonal projectors associated with a particular quantum measurement. A notable feature of this rule is that the assignment of a probability $p(\boldsymbol{E})$ to \boldsymbol{E} is noncontextual. This means that, while the probability $p(\boldsymbol{E})$ depends on the state ρ , it does not depend on how we choose the rest of the projectors that complete the orthogonal set containing \boldsymbol{E} .

In a hidden variable theory, the probabilistic description of quantum measurement is derived from a more fundamental and complete deterministic description. The outcome of a measurement could be perfectly predicted if the values of the hidden variables were precisely known — then the probability $p(\mathbf{E})$ could take only the values 0 and 1. The standard probabilistic predictions of quantum theory arise when we average over the unknown values of the hidden variables.

ables. The purpose of this exercise is to show that such deterministic assignments conflict with noncontextuality. Thus a hidden variable theory, if it is to agree with the predictions of quantum theory after averaging, must be contextual.

Let $\{I, X, Y, Z\}$ denote the single-qubit observables

.

$$\mathbf{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{X} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$
$$\mathbf{Y} = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \mathbf{Z} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (2.183)$$

and consider the nine two-qubit observables:

The three observables in each row and in each column are mutually commuting, and so can be simultaneously diagonalized. In fact the simultaneous eigenstates of any two operators in a row or column (the third operator is not independent of the other two) are a complete basis for the four-dimensional Hilbert space of the two qubits. Thus we can regard the array eq.(2.184) as a way of presenting six different ways to choose a complete set of one-dimensional projectors for two qubits.

Each of these observables has eigenvalues ± 1 , so that in a deterministic and noncontextual model of measurement (for a fixed value of the hidden variables), each can be assigned a definite value, either +1 or -1.

- a) Any noncontextual deterministic assignment has to be consistent with the multiplicative structure of the observables. For example, the product of the three observables in the top row is the identity $I \otimes I$. Therefore, if we assign a value ± 1 to each operator, the number of -1's assigned to the first row must be even. Compute the product of the three observables in each row and each column to find the corresponding constraints.
- b) Show that there is no way to satisfy all six constraints simultaneously.

Thus a deterministic and noncontextual assignment does not exist.

2.9 Schmidt-decomposable states

We have seen that any vector in a bipartite Hilbert space $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$ can be expressed in the *Schmidt form*: Given the vector $|\psi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$, where \mathcal{H}_A and \mathcal{H}_B are both *d*-dimensional, we can choose orthonormal bases $\{|i\rangle_A\}$ for \mathcal{H}_A and $\{|i\rangle_B\}$ for \mathcal{H}_B so that

$$|\psi\rangle_{AB} = \sum_{i=0}^{d-1} \sqrt{\lambda_i} \; |i\rangle_A \otimes |i\rangle_B, \qquad (2.185)$$

where the λ_i 's are real and nonnegative. (We're not assuming here that the vector has unit norm, so the sum of the λ_i 's is not constrained.) Eq.(2.185) is called the *Schmidt decomposition* of the vector $|\psi\rangle_{AB}$. Of course, the bases in which the vector has the Schmidt form depend on which vector $|\psi\rangle_{AB}$ is being decomposed.

A unitary transformation acting on \mathcal{H}_{AB} is called a *local unitary* if it is a tensor product $U_A \otimes U_B$, where U_A , U_B are unitary transformations acting on \mathcal{H}_A , \mathcal{H}_B respectively. The word "local" is used because if the two parts A and B of the system are widely separated from one another, so that Alice can access only part Aand Bob can access only part B, then Alice and Bob can apply this transformation by each acting locally on her or his part.

a) Now suppose that Alice and Bob choose standard fixed bases $\{|i\rangle_A\}$ and $\{|i\rangle_B\}$ for their respective Hilbert spaces, and are presented with a vector $|\psi_{AB}\rangle$ that is not necessarily in the Schmidt form when expressed in the standard bases. Show that there is a local unitary $U_A \otimes U_B$ that Alice and Bob can apply so that the resulting vector

$$|\psi\rangle'_{AB} = \boldsymbol{U}_A \otimes \boldsymbol{U}_B |\psi\rangle_{AB} \tag{2.186}$$

does have the form eq.(2.185) when expressed in the standard bases.

b) Let's verify that the result of (a) makes sense from the point of view of parameter counting. For a generic vector in the Schmidt form, all λ_i 's are nonvanishing and no two λ_i 's are equal. Consider the orbit that is generated by letting arbitrary local unitaries act on one fixed generic vector in the Schmidt form. What is the dimension of the orbit, that is, how many real parameters are needed to specify one particular vector on the orbit? (**Hint**: To do the counting, consider the local unitaries that differ infinitesimally from the identity $I_A \otimes I_B$. Choose a basis for these, and count the number of independent linear combinations of the basis elements that annihilate the Schmidt-decomposed vector.) Compare the dimension of the orbit to the (real) dimension of \mathcal{H}_{AB} , and check the consistency with the number of free parameters in eq.(2.185).

A vector $|\psi\rangle_{A_1...A_r}$ in a Hilbert space $\mathcal{H}_{A_1} \otimes \cdots \otimes \mathcal{H}_{A_r}$ with r parts is said to be *Schmidt decomposable* if it is possible to choose orthonormal bases for $\mathcal{H}_{A_1}, \ldots \mathcal{H}_{A_r}$ such that vector can be expressed as

$$\psi\rangle_{A_1\dots A_r} = \sum_i \sqrt{\lambda_i} |i\rangle_{A_1} \otimes |i\rangle_{A_2} \otimes \dots \otimes |i\rangle_{A_r}.$$
 (2.187)

Though every vector in a bipartite Hilbert space is Schmidt decomposable, this isn't true for vectors in Hilbert spaces with three or more parts.

- c) Consider a generic Schmidt-decomposable vector in the tripartite Hilbert space of three qubits. Find the dimension of the orbit generated by local unitaries acting on this vector.
- d) By considering the number of free parameters in the Schmidt form eq.(2.187), and the result of (c), find the (real) dimension of the space of Schmidt-decomposable vectors for three qubits. What is the real *codimension* of this space in the three-qubit Hilbert space \mathbb{C}^{8} ?

Lecture Notes for Ph219/CS219: Quantum Information Chapter 3

John Preskill California Institute of Technology

Updated July 2015

Contents

Foundations II: Measurement and Evolution	4
Orthogonal measurement and beyond	4
3.1.1 Orthogonal Measurements	4
3.1.2 Generalized measurements	8
Quantum channels	11
3.2.1 The operator-sum representation	11
3.2.2 Reversibility	13
3.2.3 Quantum channels in the Heisenberg picture	14
3.2.4 Quantum operations	15
3.2.5 Linearity	17
3.2.6 Complete positivity	18
Channel-state duality and the dilation of a channel	19
3.3.1 Channel-state duality	20
3.3.2 Stinespring dilation	22
3.3.3 Axioms revisited	23
Three quantum channels	24
3.4.1 Depolarizing channel	24
3.4.2 Dephasing channel	27
3.4.3 Amplitude-damping channel	31
Master equations for open quantum systems	34
3.5.1 Markovian evolution	34
3.5.2 The Liouvillian	35
3.5.3 Damped harmonic oscillator	38
Non-Markovian noise	41
3.6.1 Gaussian phase noise	41
3.6.2 Spin echo	43
3.6.3 Qubits as noise spectrometers	44
3.6.4 Spin-boson model at nonzero temperature	46
Summary	48
	Foundations II: Measurement and Evolution Orthogonal measurement and beyond 3.1.1 Orthogonal Measurements 3.1.2 Generalized measurements Quantum channels 3.2.1 The operator-sum representation 3.2.2 Reversibility 3.2.3 Quantum channels in the Heisenberg picture 3.2.4 Quantum operations 3.2.5 Linearity 3.2.6 Complete positivity Channel-state duality and the dilation of a channel 3.3.1 Channel-state duality 3.3.2 Stinespring dilation 3.3.3 Axioms revisited Three quantum channels 3.4.1 Depolarizing channel 3.4.2 Dephasing channel 3.4.3 Amplitude-damping channel Master equations for open quantum systems 3.5.1 Markovian evolution 3.5.2 The Liouvillian 3.5.3 Damped harmonic oscillator Non-Markovian noise 3.6.1 Gaussian phase noise 3.6.2 Spin echo 3.6.3 Qubits as noise spectrometers 3.6.4 Spin-boson model at nonzero temperature Summary

3.8 Exercises

49

3

Foundations II: Measurement and Evolution

3.1 Orthogonal measurement and beyond

In Chapter 2 we discussed how to describe the state of an *open* quantum system, one which is part of a larger system. In this Chapter we will extend the theory of open quantum systems further. In particular, we will develop two important concepts: *generalized measurements*, which are performed by making use of an auxiliary system, and *quantum channels*, which describe how open systems evolve.

3.1.1 Orthogonal Measurements

An axiom of quantum theory asserts that a measurement may be described as an orthogonal projection operator. But if a measurement of system S is realized by performing an orthogonal measurement on a larger system that contains S, the resulting operation performed on S alone need not be an orthogonal projection. We would like to find a mathematical description of such "generalized measurements" on system S. But first let's recall how measurement of an arbitrary Hermitian operator can be achieved in principle, following the classic treatment of Von Neumann.

To measure an observable M, we will modify the Hamiltonian of the world by turning on a coupling between that observable and another variable that represents the apparatus. Depending on the context, we will refer to the auxiliary system used in the measurement as the "pointer," the "meter," or the "ancilla." (The word "ancilla" just means something extra which is used to achieve a desired goal.) The coupling establishes a correlation between the eigenstates of the observable and the distinguishable states of the pointer, so that we can prepare an eigenstate of the observable by "observing" the pointer.

This may not seem like a fully satisfying model of measurement because

we have not explained how to measure the pointer. Von Neumann's attitude was that it is possible in principle to correlate the state of a microscopic quantum system with the value of a macroscopic classical variable, and we may take it for granted that we can perceive the value of the classical variable. A quantum measurement, then, is a procedure for amplifying a property of the microscopic world, making it manifest in the macroscopic world.

We may think of the pointer as a particle of mass m that propagates freely apart from its tunable coupling to the quantum system being measured. Since we intend to measure the position of the pointer, it should be prepared initially in a wavepacket state that is narrow in position space — but not too narrow, because a vary narrow wave packet will spread too rapidly. If the initial width of the wave packet is Δx , then the uncertainty in it velocity will be of order $\Delta v = \Delta p/m \sim \hbar/m\Delta x$, so that after a time t, the wavepacket will spread to a width

$$\Delta x(t) \sim \Delta x + \frac{\hbar t}{\Delta x},\tag{3.1}$$

which is minimized for $(\Delta x(t))^2 \sim (\Delta x)^2 \sim \hbar t/m$. Therefore, if the experiment takes a time t, the resolution we can achieve for the final position of the pointer is limited by

$$\Delta x \gtrsim (\Delta x)_{SQL} \sim \sqrt{\frac{\hbar t}{m}},$$
(3.2)

the "standard quantum limit." We will choose our pointer to be sufficiently heavy that this limitation is not serious.

The Hamiltonian describing the coupling of the quantum system to the pointer has the form

$$\boldsymbol{H} = \boldsymbol{H}_0 + \frac{1}{2m} \boldsymbol{P}^2 + \lambda(t) \boldsymbol{M} \otimes \boldsymbol{P}, \qquad (3.3)$$

where $\mathbf{P}^2/2m$ is the Hamiltonian of the free pointer particle (which we will henceforth ignore on the grounds that the pointer is so heavy that spreading of its wavepacket may be neglected), \mathbf{H}_0 is the unperturbed Hamiltonian of the system to be measured, and λ is a coupling constant that we are able to turn on and off as desired. The observable to be measured, \mathbf{M} , is coupled to the momentum \mathbf{P} of the pointer.

If M does not commute with H_0 , then we have to worry about how the observable M evolves during the course of the measurement. To simplify the analysis, let us suppose that either $[M, H_0] = 0$, or else the measurement is carried out quickly enough that the free evolution of the system can be neglected during the measurement procedure. Then the Hamiltonian can be approximated as $\boldsymbol{H} \simeq \lambda(t) \boldsymbol{M} \otimes \boldsymbol{P}$ If the coupling λ switches on suddenly at time zero and switches off suddenly at time T, the resulting time evolution operator is

$$\boldsymbol{U}(T) \simeq \exp\left(-i\lambda T\boldsymbol{M} \otimes \boldsymbol{P}\right). \tag{3.4}$$

Expanding in the basis in which M is diagonal,

$$\boldsymbol{M} = \sum_{a} |a\rangle M_a \langle a|, \qquad (3.5)$$

we express $\boldsymbol{U}(T)$ as

$$\boldsymbol{U}(T) = \sum_{a} |a\rangle \exp\left(-i\lambda t M_{a}\boldsymbol{P}\right) \langle a|.$$
(3.6)

Now we recall that \boldsymbol{P} generates a translation of the *position* of the pointer: $\boldsymbol{P} = -i\frac{d}{dx}$ in the position representation, so that $e^{-ix_0\boldsymbol{P}} = \exp\left(-x_0\frac{d}{dx}\right)$, and by Taylor expanding,

$$e^{-ix_0 P}\psi(x) = \psi(x - x_0);$$
 (3.7)

In other words $e^{-ix_0 \mathbf{P}}$ acting on a wavepacket translates the wavepacket by x_0 . We see that if our quantum system starts in a superposition of \mathbf{M} eigenstates, initially unentangled with the position-space wavepacket $|\psi(x)\rangle$ of the pointer, then after time T the quantum state has evolved to

$$\boldsymbol{U}(T)\left(\sum_{a}\alpha_{a}|a\rangle\otimes|\psi(x)\rangle\right)=\sum_{a}\alpha_{a}|a\rangle\otimes|\psi(x-\lambda TM_{a})\rangle;\qquad(3.8)$$

the position of the pointer has become correlated with the value of the observable M. If the pointer wavepacket is narrow enough for us to resolve all values of the M_a that occur (that is, the width Δx of the packet is small compared to $\lambda T \Delta M_a$, where ΔM_a is the minimal gap between eigenvalues of M), then when we observe the position of the pointer (never mind how!) we will prepare an eigenstate of the observable. With probability $|\alpha_a|^2$, we will detect that the pointer has shifted its position by $\lambda T M_a$, in which case we will have prepared the M eigenstate $|a\rangle$. We conclude that the initial state $|\varphi\rangle$ of the quantum system is projected to $|a\rangle$ with probability $|\langle a|\varphi\rangle|^2$. This is Von Neumann's model of orthogonal measurement.

The classic example is the Stern–Gerlach apparatus. To measure σ_3 for a spin- $\frac{1}{2}$ object, we allow the object to pass through a region of inhomogeneous magnetic field

$$B_3 = \lambda z. \tag{3.9}$$

The magnetic moment of the object is $\mu \vec{\sigma}$, and the coupling induced by the magnetic field is

$$\boldsymbol{H} = -\lambda \mu \boldsymbol{z} \boldsymbol{\sigma}_3. \tag{3.10}$$

In this case σ_3 is the observable to be measured, coupled to the position z rather than the momentum of the pointer; thus, because z generates a translation of P_z , the coupling imparts an *impulse* to the pointer which is correlated with its spin. We can perceive whether the object is pushed up or down, and so project out the spin state $|\uparrow_z\rangle$ or $|\downarrow_z\rangle$. By rotating the magnet, we could measure the observable $\hat{n} \cdot \vec{\sigma}$ instead.

Thinking more abstractly, suppose that $\{E_a, a = 0, 1, 2, ..., N-1\}$ is a complete set of orthogonal projectors satisfying

$$\boldsymbol{E}_{a}\boldsymbol{E}_{b} = \delta_{ab}\boldsymbol{E}_{a}, \quad \boldsymbol{E}_{a} = \boldsymbol{E}_{a}^{\dagger}, \quad \sum_{a} \boldsymbol{E}_{a} = \boldsymbol{I}.$$
 (3.11)

To perform an orthogonal measurement with these outcomes, we introduce an N-dimensional pointer system with fiducial orthonormal basis states $\{|a\rangle, a = 0, 1, 2, ..., N-1\}$, and, by coupling the system to the pointer, perform the unitary transformation

$$\boldsymbol{U} = \sum_{a,b} \boldsymbol{E}_a \otimes |b+a\rangle \langle b|. \tag{3.12}$$

Thus the pointer advances by an amount a if the state of the system is within the support of the projector E_a . (The addition in $|b + a\rangle$ is understood to be modulo N; we may envision the pointer as a circular dial with N uniformly spaced tick marks.) The unitarity of U is easy to verify:

$$\boldsymbol{U}\boldsymbol{U}^{\dagger} = \left(\sum_{a,b} \boldsymbol{E}_{a} \otimes |b+a\rangle\langle b|\right) \left(\sum_{c,d} \boldsymbol{E}_{c} \otimes |d\rangle\langle d+c|\right)$$
$$= \sum_{a,b,c,d} \delta_{ac}\boldsymbol{E}_{a} \otimes \delta_{bd}|b+a\rangle\langle d+c|$$
$$= \sum_{a} \boldsymbol{E}_{a} \otimes \sum_{b} |b+a\rangle\langle b+a| = \boldsymbol{I} \otimes \boldsymbol{I}.$$
(3.13)

This unitary transformation acts on an initial product state of system and pointer according to

$$\boldsymbol{U}: |\Psi\rangle = |\psi\rangle \otimes |0\rangle \mapsto |\Psi'\rangle = \sum_{a} \boldsymbol{E}_{a} |\psi\rangle \otimes |a\rangle; \qquad (3.14)$$

if the pointer is then measured in the fiducial basis, the measurement postulate implies that the outcome a occurs with probability

$$\operatorname{Prob}(a) = \langle \Psi' | \left(\boldsymbol{I} \otimes |a\rangle \langle a| \right) | \Psi' \rangle = \langle \psi | \boldsymbol{E}_a | \psi \rangle, \qquad (3.15)$$

and that when this outcome occurs the normalized post-measurement state is

$$\frac{E_a|\psi\rangle}{\|E_a|\psi\rangle\|}.\tag{3.16}$$

If the measurement is performed and its outcome is not known, the initial pure state of the system becomes a mixture of these post-measurement states:

$$\sum_{a} \operatorname{Prob}(a) \frac{\boldsymbol{E}_{a} |\psi\rangle \langle \psi | \boldsymbol{E}_{a}}{\langle \psi | \boldsymbol{E}_{a} |\psi\rangle} = \sum_{a} \boldsymbol{E}_{a} |\psi\rangle \langle \psi | \boldsymbol{E}_{a}.$$
(3.17)

In fact, the system is described by this density operator once it becomes entangled with the pointer, whether we bother to observe the pointer or not. If the initial state of the system before the measurement is a mixed state with density matrix ρ , then by expressing ρ as an ensemble of pure states we conclude that the measurement modifies the state according to

$$\boldsymbol{\rho} \mapsto \sum_{a} \boldsymbol{E}_{a} \boldsymbol{\rho} \boldsymbol{E}_{a}.$$
 (3.18)

We see that if, by coupling the system to our pointer, we can execute suitable unitary transformations correlating the system and the pointer, and if we can observe the pointer in its fiducial basis, then we are empowered to perform any conceivable orthogonal measurement on the system.

3.1.2 Generalized measurements

In this discussion of orthogonal measurement, the fiducial basis of the pointer had two different roles — we assumed that the fiducial pointer states become correlated with the system projectors $\{E_a\}$, and also that the measurement of the pointer projects onto the fiducial basis. In principle we could separate these two roles. Perhaps the unitary transformation applied to system and pointer picks out a different preferred basis than the basis in which the pointer is easily measured. Or perhaps the pointer which becomes entangled with the system is itself microscopic, and we may entangle it with a second macroscopic pointer in order to measure the microscopic pointer in whatever basis we prefer.

Suppose, to be concrete, that the system A is a single qubit, and so is the pointer B. They interact, resulting in the unitary map

$$\boldsymbol{U}: (\alpha|0\rangle + \beta|1\rangle)_A \otimes |0\rangle_B \mapsto \alpha|0\rangle_A \otimes |0\rangle_B + \beta|1\rangle_A \otimes |1\rangle_B.$$
(3.19)

Measuring the pointer by projecting onto the basis $\{|0\rangle, |1\rangle\}$ would induce an orthogonal measurement of the system, also in the $\{|0\rangle, |1\rangle\}$ basis. But suppose that we measure the pointer in a different basis instead, such as $\{|\pm\rangle = \frac{1}{\sqrt{2}} (|0\rangle \pm |1\rangle)\}$. Then the measurement postulate dictates that the two outcomes + and - occur equiprobably, and that the corresponding post-measurement states of the system are

$$\alpha|0\rangle \pm \beta|1\rangle. \tag{3.20}$$

In contrast to an orthogonal measurement of the system, these two postmeasurement states are not orthogonal, unless $|\alpha| = |\beta|$. Furthermore, also in contrast to an orthogonal measurement, if two such measurements are performed in rapid succession, the outcomes need not be the same. We use the term *generalized measurement* to mean a measurement, like this one, which is not necessarily an orthogonal projection acting on the system.

It is convenient to describe this measurement procedure by expanding the entangled state of system and pointer in the basis in which the pointer is measured; hence we rewrite eq.(3.21) as

$$\boldsymbol{U}:|\psi\rangle_A\otimes|0\rangle_B\mapsto\boldsymbol{M}_+|\psi\rangle_A\otimes|+\rangle_B+\boldsymbol{M}_-|\psi\rangle_A\otimes|-\rangle_B,\qquad(3.21)$$

where

$$\boldsymbol{M}_{+} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = \frac{1}{\sqrt{2}} \boldsymbol{I}, \quad \boldsymbol{M}_{-} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = \frac{1}{\sqrt{2}} \boldsymbol{\sigma}_{3}. \quad (3.22)$$

Evidently, by measuring B in the basis $\{|\pm\rangle\}$, we prepare A in one of the states $M_{\pm}|\psi\rangle$, up to a normalization factor.

Now let's generalize this idea to the case where the pointer system B is N-dimensional, and the measurement of the pointer projects onto an orthonormal basis $\{|a\rangle, a = 0, 1, 2, ..., N-1\}$. Again we'll assume that the system A and pointer B are initially in a product state, then an entangling unitary transformation U correlates the system with the pointer. By expanding the action of U in the basis for B we obtain

$$\boldsymbol{U}:|\psi\rangle_A\otimes|0\rangle_B\mapsto\sum_a\boldsymbol{M}_a|\psi\rangle_A\otimes|a\rangle_B.$$
(3.23)

Since U is unitary, it preserves the norm of any input, which means that

$$1 = \left\|\sum_{a} \boldsymbol{M}_{a} |\psi\rangle \otimes |a\rangle\right\|^{2} = \sum_{a,b} \langle \psi | \boldsymbol{M}_{a}^{\dagger} \boldsymbol{M}_{b} |\psi\rangle \langle a |b\rangle = \sum_{a} \langle \psi | \boldsymbol{M}_{a}^{\dagger} \boldsymbol{M}_{a} |\psi\rangle$$
(3.24)

for any $|\psi\rangle$; hence

$$\sum_{a} \boldsymbol{M}_{a}^{\dagger} \boldsymbol{M}_{a} = \boldsymbol{I}.$$
(3.25)

The complete orthogonal measurement projecting onto the pointer basis $\{|a\rangle_B\}$ is equivalent to the incomplete orthogonal measurement on AB with projectors $\{I \otimes |a\rangle\langle a|\}$; the measurement postulate asserts that outcome *a* occurs with probability

$$\operatorname{Prob}(a) = \|\boldsymbol{M}_a|\psi\rangle\|^2, \qquad (3.26)$$

and that if outcome a occurs the post-measurement state of the system is

$$\frac{\boldsymbol{M}_a|\psi\rangle}{\|\boldsymbol{M}_a|\psi\rangle\|}.$$
(3.27)

The completeness relation $\sum_{a} M_{a}^{\dagger} M_{a} = I$ ensures that the probabilities sum to one, but the possible post-measurement states need not be mutually orthogonal, nor are the measurements necessarily repeatable. If we perform the measurement twice in succession and obtain outcome *a* the first time, the conditional probability of obtaining outcome *b* in the second measurement is

$$\operatorname{Prob}(b|a) = \frac{\|\boldsymbol{M}_{b}\boldsymbol{M}_{a}|\psi\rangle\|^{2}}{\|\boldsymbol{M}_{a}|\psi\rangle\|^{2}}.$$
(3.28)

The two measurements agree if $\operatorname{Prob}(b|a) = \delta_{ba}$, which is satisfied for arbitrary initial states of the system only if $M_b M_a = \delta_{ba} M_a$ up to a phase factor, *i.e.* in the case where the measurement is projective.

We see that if the initial state of the system is the density operator ρ (realized as an ensemble of pure states), there is an operator $E_a = M_a^{\dagger} M_a$ associated with each possible measurement outcome a, such that the probability of outcome a is

$$\operatorname{Prob}(a) = \operatorname{tr}(\boldsymbol{\rho}\boldsymbol{E}_a). \tag{3.29}$$

The measurement operators $\{E_a\}$ form a complete set of Hermitian nonnegative operators; that is, they satisfy the properties:

- 1. Hermiticity. $E_a = E_a^{\dagger}$.
- 2. Positivity. $\langle \psi | \mathbf{E}_a | \psi \rangle \geq 0$ for any vector $| \psi \rangle$; we abbreviate this property by simply writing $\mathbf{E}_a \geq 0$.
- 3. Completeness. $\sum_{a} E_{a} = I$.

Such a partition of unity by nonnegative operators is called a *positive* operator-valued measure, or POVM. (The word term measure is a bit heavy-handed in this finite-dimensional context; it becomes more apt when the index a can be continuously varying.)

We have seen how a POVM can arise when an orthogonal measurement is performed on a meter after the meter interacts with the system. In fact any POVM can be realized this way. We need only observe that a nonnegative Hermitian operator E_a has a nonnegative square root $\sqrt{E_a}$; more generally, the operator

$$\boldsymbol{M}_a = \boldsymbol{U}_a \sqrt{\boldsymbol{E}_a} \tag{3.30}$$

obeys $M_a^{\dagger}M_a = E_a$ where U_a is an arbitrary unitary operator eq.(3.30) is called the *polar decomposition* of the operator M_a . Plugging into eq.(3.23) yields the unitary interaction which realizes the POVM $\{E_a\}$. In this formulation, the post-measurement state corresponding to outcome a,

$$\boldsymbol{U}_{a}\left(\frac{\sqrt{\boldsymbol{E}_{a}}|\psi\rangle}{\|\sqrt{\boldsymbol{E}_{a}}|\psi\rangle\|}\right),\tag{3.31}$$

is arbitrary, since we are free to choose the unitary U_a however we please for each possible outcome. The POVM attributes a probability to each measurement outcome, but provides no guidance regarding the state after the measurement. Indeed, after the measurement we have the freedom to discard the state and replace it by whatever freshly prepared state we desire.

3.2 Quantum channels

3.2.1 The operator-sum representation

We now proceed to the next step in our program of understanding the behavior of one part of a bipartite quantum system. We have seen that a pure state of the bipartite system AB may behave like a mixed state when we observe subsystem A alone, and that an orthogonal measurement of the bipartite system can realize a (nonorthogonal) POVM on A alone. Next we ask, if a state of the bipartite system undergoes unitary evolution, how do we describe the evolution of A alone?

In effect, we have already answered this question in our discussion of generalized measurements. If system A starts out in a pure state $|\psi\rangle$ (unentangled with B), and then interacts with B, the joint state of AB has the form eq.(3.23); the resulting density operator for A is found by tracing out B. Equivalently, we may imagine measuring system B in the basis $\{|a\rangle\}$, but failing to record the measurement outcome, so we are forced to average over all the possible post-measurement states, weighted by their probabilities. The result is that the initial density operator $\rho = |\psi\rangle\langle\psi|$ is subjected to a linear map \mathcal{E} , which acts as

$$\mathcal{E}(\boldsymbol{\rho}) = \sum_{a} \boldsymbol{M}_{a} \boldsymbol{\rho} \boldsymbol{M}_{a}^{\dagger}, \qquad (3.32)$$

where the operators $\{M_a\}$ obey the completeness relation eq.(3.25). Being linear, \mathcal{E} acts as in eq.(3.32) not just on pure states, but on any density operator.

A linear map of the form eq.(3.32), where the $\{M_a\}$ obey eq.(3.25), is called a quantum channel. The word "channel" is drawn from communication theory — we are to imagine a sender who transmits the state ρ though a communication link to another party who receives the modified state $\mathcal{E}(\rho)$. Sometimes the word superoperator is used as a synonym for quantum channel, where "super" conveys that the map takes operators to operators rather than vectors to vectors. Yet another name for the same object is trace-preserving completely positive map, or TPCP map for short. The justification for this name will emerge shortly. Eq.(3.32) is said to be an operator-sum representation of the quantum channel, and the operators $\{M_a\}$ are called the Kraus operators or operation elements of the channel.

A quantum channel maps density operators to density operators; that is, has the following easily verified properties:

- 1. Linearity. $\mathcal{E}(\alpha \rho_1 + \beta \rho_2) = \alpha \mathcal{E}(\rho_1) + \beta \mathcal{E}(\rho_2).$
- 2. Preserves Hermiticity. $\boldsymbol{\rho} = \boldsymbol{\rho}^{\dagger}$ implies $\mathcal{E}(\boldsymbol{\rho}) = \mathcal{E}(\boldsymbol{\rho})^{\dagger}$.
- 3. Preserves positivity. $\rho \geq 0$ implies $\mathcal{E}(\rho) \geq 0$.
- 4. Preserves trace. $\operatorname{tr} (\mathcal{E}(\boldsymbol{\rho})) = \operatorname{tr} (\boldsymbol{\rho}).$

These properties partially explain the locution "trace-preserving completely positive map," except that we are still missing the reason for the modifier "completely." That's coming soon.

We've seen how a quantum channel acting on system A arises from a unitary transformation acting on A and B followed by a partial trace on B. As in our discussion of generalized measurements, we can also run this argument backwards to see that any quantum channel may be realized this way. Given a quantum channel \mathcal{E} acting on A with Kraus operators $\{M_a\}$, we may introduce the auxiliary system B with Hilbert space dimension matching the number of Kraus operators. A unitary transformation may then be constructed whose action on $|\psi\rangle_A \otimes |0\rangle_B$ is as in eq.(3.23), from which the quantum channel \mathcal{E} is obtained by tracing out B.

The operator-sum representation of a given quantum channel \mathcal{E} is not unique, because we can perform the partial trace on B in any basis we please. When expressed in terms of rotated basis states $\{|\mu\rangle\}$ such that

$$|a\rangle = \sum_{\mu} |\mu\rangle V_{\mu a} \tag{3.33}$$

for unitary V, the joint state of AB becomes

$$\sum_{a,\mu} \boldsymbol{M}_a |\psi\rangle_A \otimes |\mu\rangle_B V_{\mu a} = \sum_{\mu} \boldsymbol{N}_{\mu} |\psi\rangle_A \otimes |\mu\rangle_B$$
(3.34)

where the new Kraus operators are

$$\boldsymbol{N}_{\mu} = \sum_{a} V_{\mu a} \boldsymbol{M}_{a}. \tag{3.35}$$

13

We will see soon that any two operator-sum representations of the same quantum channel are always related by such a unitary change of basis for the Kraus operators.

Quantum channels are important because they provide us with a formalism for discussing *decoherence*, the evolution of pure states into mixed states. Unitary evolution of ρ_A is the special case in which there is only one term in the operator sum. If there are two or more terms, then there are pure initial states of A which become entangled with B under evolution governed by the joint unitary transformation U_{AB} , and therefore the state of A becomes mixed when we trace out B.

Two channels \mathcal{E}_1 and \mathcal{E}_2 can be composed to obtain another channel $\mathcal{E}_2 \circ \mathcal{E}_1$; if \mathcal{E}_1 describes evolution from yesterday to today, and \mathcal{E}_2 describes evolution from today to tomorrow, then $\mathcal{E}_2 \circ \mathcal{E}_1$ describes the evolution from yesterday to tomorrow. Specifically, if \mathcal{E}_1 has an operator-sum representation with N Kraus operators $\{M_a\}$, and \mathcal{E}_2 has an operator-sum representation with M Kraus operators $\{N_\mu\}$, then $\mathcal{E}_2 \circ \mathcal{E}_1$ has an operator-sum representation with NM Kraus operators $\{N_\mu\}$, then $\mathcal{E}_2 \circ \mathcal{E}_1$ has an operator-sum representation with NM Kraus operators $\{N_\mu\}$, then $\mathcal{E}_2 \circ \mathcal{E}_1$ has an operator-sum representation with NM Kraus operators $\{N_\mu\}$. Because we can compose them in this way, we say that quantum channels form a dynamical semigroup.

3.2.2 Reversibility

A unitary transformation U has a unitary inverse U^{\dagger} . Thus if today's quantum state was obtained by applying U to yesterday's state, we can in principle recover yesterday's state by applying U^{\dagger} to today's state. Unitary time evolution is reversible.

Is the same true for general quantum channels? If channel \mathcal{E}_1 with Kraus operators $\{M_a\}$ is inverted by channel \mathcal{E}_2 with Kraus operators $\{N_\mu\}$, then for any pure state $|\psi\rangle$ we have

$$\mathcal{E}_2 \circ \mathcal{E}_1(|\psi\rangle\langle\psi|) = \sum_{\mu,a} \mathbf{N}_{\mu} \mathbf{M}_a |\psi\rangle\langle\psi| \mathbf{M}_a^{\dagger} \mathbf{N}_{\mu}^{\dagger} = |\psi\rangle\langle\psi|.$$
(3.36)

Since the left-hand side is a sum of positive terms, eq.(3.36) can hold only if each of these terms is proportional to $|\psi\rangle\langle\psi|$, hence

$$\boldsymbol{N}_{\mu}\boldsymbol{M}_{a} = \lambda_{\mu a}\boldsymbol{I} \tag{3.37}$$

for each μ and a. Using the completeness relation, we find

$$\boldsymbol{M}_{b}^{\dagger}\boldsymbol{M}_{a} = \boldsymbol{M}_{b}^{\dagger}\left(\sum_{\mu}\boldsymbol{N}_{\mu}^{\dagger}\boldsymbol{N}_{\mu}\right)\boldsymbol{M}_{a} = \sum_{\mu}\lambda_{\mu b}^{*}\lambda_{\mu a}\boldsymbol{I} \equiv \beta_{ba}\boldsymbol{I}.$$
 (3.38)

where each β_{aa} is real and positive unless $M_a = 0$. The polar decomposition of M_a yields

$$\boldsymbol{M}_{a} = \boldsymbol{U}_{a} \sqrt{\boldsymbol{M}_{a}^{\dagger} \boldsymbol{M}_{a}} = \sqrt{\beta_{aa}} \boldsymbol{U}_{a}$$
(3.39)

for some unitary U_a , and it then follows that

$$\boldsymbol{M}_{b}^{\dagger}\boldsymbol{M}_{a} = \sqrt{\beta_{aa}\beta_{bb}} \; \boldsymbol{U}_{b}^{\dagger}\boldsymbol{U}_{a} = \beta_{ba}\boldsymbol{I}, \qquad (3.40)$$

and hence

$$\boldsymbol{U}_a = \frac{\beta_{ba}}{\sqrt{\beta_{aa}\beta_{bb}}} \boldsymbol{U}_b \tag{3.41}$$

for each a and b. We conclude that each Kraus operator M_a is proportional to a single unitary matrix, and hence that \mathcal{E}_1 is a unitary map. A quantum channel can be inverted by another quantum channel only if it is unitary.

We have found that decoherence is irreversible. Once system A becomes entangled with system B, we can't undo the damage to A if we don't have access to B. Decoherence causes quantum information to leak to a system's environment, and because we cannot control the environment this information cannot be recovered.

3.2.3 Quantum channels in the Heisenberg picture

We have described quantum channels using the *Schrödinger picture* in which the quantum state evolves with time. Sometimes it is convenient to use the *Heisenberg picture*, in which the state is stationary and the operators evolve instead.

When time evolution is unitary, in the Schrödinger picture the state vector at time t is obtained from the state vector at time 0 by

$$|\psi(t)\rangle = \boldsymbol{U}(t)|\psi(0)\rangle \tag{3.42}$$

where $\boldsymbol{U}(t)$ is unitary, and correspondingly a density operator evolves according to

$$\boldsymbol{\rho}(t) = \boldsymbol{U}(t)\boldsymbol{\rho}(0)\boldsymbol{U}(t)^{\dagger}. \tag{3.43}$$

In the Heisenberg picture the density operator ρ is fixed, and an operator A evolves according to

$$\boldsymbol{A}(t) = \boldsymbol{U}(t)^{\dagger} \boldsymbol{A}(0) \boldsymbol{U}(t). \tag{3.44}$$

15

This evolution law is chosen so that the two pictures agree on the expectation values of observables at any time:

$$\langle \boldsymbol{A} \rangle_{t,\text{Schr}} = \text{tr} \left(\boldsymbol{A}(0)\boldsymbol{\rho}(t) \right) = \text{tr} \left(\boldsymbol{A}(t)\boldsymbol{\rho}(0) \right) = \langle \boldsymbol{A} \rangle_{t,\text{Heis}},$$
 (3.45)

where we have used the cyclic property of the trace.

Likewise, if the ${\mathcal E}$ is a quantum channel which acts on density operators according to

$$\boldsymbol{\rho}' = \mathcal{E}(\boldsymbol{\rho}) = \sum_{a} \boldsymbol{M}_{a} \boldsymbol{\rho} \boldsymbol{M}_{a}^{\dagger}, \qquad (3.46)$$

we may use an alternative description in which the state is fixed, but operators evolve as

$$\mathbf{A}' = \mathcal{E}^*(\mathbf{A}) = \sum_a \mathbf{M}_a^{\dagger} \mathbf{A} \mathbf{M}_a, \qquad (3.47)$$

so that

$$\operatorname{tr}\left(\boldsymbol{A}\ \mathcal{E}(\boldsymbol{\rho})\right) = \operatorname{tr}\left(\mathcal{E}^{*}(\boldsymbol{A})\boldsymbol{\rho}\right). \tag{3.48}$$

We say that \mathcal{E}^* is the *dual* or *adjoint* of \mathcal{E} .

Note that the dual of a channel need not be a channel, that is, might not be trace preserving. Instead, the completeness property of the Kraus operators $\{M_a\}$ implies that

$$\mathcal{E}^*(\mathbf{I}) = \mathbf{I} \tag{3.49}$$

if \mathcal{E} is a channel. We say that a map is *unital* if it preserves the identity operator, and conclude that the dual of a channel is a unital map.

Not all quantum channels are unital, but some are. If the Kraus operators of \mathcal{E} satisfy

$$\sum_{a} \boldsymbol{M}_{a}^{\dagger} \boldsymbol{M}_{a} = \boldsymbol{I} = \sum_{a} \boldsymbol{M}_{a} \boldsymbol{M}_{a}^{\dagger}, \qquad (3.50)$$

then \mathcal{E} is unital and its dual \mathcal{E}^* is also a unital channel. A unital quantum channel maps a maximally mixed density operator to itself; it is the quantum version of a doubly stochastic classical map, which maps probability distributions to probability distributions and preserves the uniform distribution.

3.2.4 Quantum operations

Generalized measurements and quantum channels are actually special cases of a more general notion called a *quantum operation*. As already noted, a generalized measurement can be realized by entangling a system with a meter and performing orthogonal measurement on the meter, while a quantum channel arises if we measure the meter but completely forget the measurement outcome. In a quantum operation, we imagine measuring the meter, then retaining some of the information about the outcome while discarding the rest.

We may consider a generalized measurement described by Kraus operators $\{M_{a\mu}\}$ which carry two labels, a and μ . These obey the usual completeness relation

$$\sum_{a,\mu} \boldsymbol{M}_{a\mu}^{\dagger} \boldsymbol{M}_{a\mu} = \boldsymbol{I}.$$
(3.51)

Suppose that, after a measurement that projects onto a definite value for both a and μ , we remember a but forget μ . Then, if the quantum state is ρ before the measurement, the post-measurement state (up to a normalization factor) is

$$\mathcal{E}_{a}(\boldsymbol{\rho}) \equiv \sum_{\mu} \boldsymbol{M}_{a\mu} \boldsymbol{\rho} \boldsymbol{M}_{a\mu}^{\dagger}, \qquad (3.52)$$

where the outcome a occurs with probability

$$\operatorname{Prob}(a) = \operatorname{tr} \mathcal{E}_a(\boldsymbol{\rho}). \tag{3.53}$$

Eq.(3.52) looks like the operator-sum representation for a quantum channel, except that now instead of the completeness relation the Kraus operators obey an inequality constraint

$$\sum_{\mu} \boldsymbol{M}_{a\mu}^{\dagger} \boldsymbol{M}_{a\mu} \leq \boldsymbol{I}.$$
(3.54)

(We write $A \leq I$ as a shorthand for the statement that I - A is a nonnegative operation; that is, the eigenvalues of the Hermitian operator Aare no larger than 1.) Our earlier notion of a generalized measurement is recovered when μ takes just one value (all information about the outcome is retained), and the operation becomes a channel when a takes just one value (all information about the outcome is discarded).

The state needs to be renormalized to restore the unit trace condition; therefore under an operation the state really evolves nonlinearly according to

$$\boldsymbol{\rho} \mapsto \frac{\mathcal{E}_a(\boldsymbol{\rho})}{\operatorname{tr} \mathcal{E}_a(\boldsymbol{\rho})}.$$
(3.55)

It is often convenient, though, to regard the operation as a linear map that takes ρ to a subnormalized state. For example, we may want to consider a sequence of *n* consecutive measurements with outcomes $\{a_1, a_2, \ldots, a_{n-1}, a_n\}$, where the *i*th measurement transforms the state according to the operation \mathcal{E}_{a_i} . Rather than renormalizing the state after each measurement, we can wait until after the final measurement in the sequence before renormalizing. The final state can then be written

$$\boldsymbol{\rho} \mapsto \frac{\mathcal{E}_{a_n} \circ \mathcal{E}_{a_{n-1}} \circ \cdots \circ \mathcal{E}_{a_2} \circ \mathcal{E}_{a_1}(\boldsymbol{\rho})}{\operatorname{tr} \mathcal{E}_{a_n} \circ \mathcal{E}_{a_{n-1}} \circ \cdots \circ \mathcal{E}_{a_2} \circ \mathcal{E}_{a_1}(\boldsymbol{\rho})}$$
(3.56)

where the normalizing factor in the denominator is just the probability of the observed sequence of measurement outcomes.

3.2.5 Linearity

A quantum channel specifies how an initial density operator evolves to a final density operator. Why on general grounds should we expect evolution of a quantum state to be described by a linear map? One possible answer is that nonlinear evolution would be incompatible with interpreting the density operator as an ensemble of possible states.

Suppose that \mathcal{E} maps an initial state at time t = 0 to a final state at time t = T, and suppose that at time t = 0 the initial state ρ_i is prepared with probability p_i . Then the time-evolved state at t = T will be $\mathcal{E}(\rho_i)$ with probability p_i .

On the other hand we argued in Chapter 2 that an ensemble in which σ_i is prepared with probability q_i can be described by the convex combination of density operators

$$\boldsymbol{\sigma} = \sum_{i} q_i \boldsymbol{\sigma}_i. \tag{3.57}$$

Therefore the initial state is described by $\sum_i p_i \rho_i$, which evolves to

$$\boldsymbol{\rho}' = \mathcal{E}\left(\sum_{i} p_i \boldsymbol{\rho}_i\right). \tag{3.58}$$

But we can also apply eq.(3.57) to the ensemble of final states, concluding that the final state may alternatively be described by

$$\boldsymbol{\rho}' = \sum_{i} p_i \mathcal{E}(\boldsymbol{\rho}_i). \tag{3.59}$$

Equating the two expressions for ρ' we find that \mathcal{E} must act linearly, at least on convex combinations of states.

Similar reasoning applies to quantum operations, if we regard the normalization of an operation \mathcal{E}_a as indicating the probability of the corresponding measurement outcome. Suppose again that the initial state ρ_i is prepared with *a priori* probability p_i and subsequently measured. If the state is ρ_i then measurement outcome *a* occurs with conditional probability p(a|i), and the post-measurement state is $\mathcal{E}_a(\rho_i)/p(a|i)$; hence the state ensemble after the measurement is described by the density operator

$$\boldsymbol{\rho}' = \sum_{i} p(i|a) \frac{\mathcal{E}_a(\boldsymbol{\rho}_i)}{p(a|i)},\tag{3.60}$$

where p(i|a) is the *a posteriori* probability that state ρ_i was prepared, taking into account the information gained by doing the measurement. On the other hand, applying the operation \mathcal{E}_a to the convex combination of the initial states $\{\rho_i\}$ yields

$$\boldsymbol{\rho}' = \frac{\mathcal{E}_a\left(\sum_i p_i \boldsymbol{\rho}_i\right)}{p_a}.$$
(3.61)

Invoking Bayes' rule

$$p_i p(a|i) = p_a p(i|a) \tag{3.62}$$

we see that the operation \mathcal{E}_a is required to be a linear map:

$$\mathcal{E}_a\left(\sum_i p_i \boldsymbol{\rho}_i\right) = \sum_i p_i \mathcal{E}_a(\boldsymbol{\rho}_i). \tag{3.63}$$

3.2.6 Complete positivity

A quantum channel is a linear map taking density operators to density operators. In particular, if its input is a nonnegative operator than so is its output. We therefore say that a channel is a *positive map*.

But a channel has a stronger property than mere positivity; it is *completely positive*. This means that the channel remains positive even when we consider it to be acting on just part of a larger system.

If a channel \mathcal{E} maps linear operators on Hilbert space \mathcal{H}_A to linear operators on Hilbert space $\mathcal{H}_{A'}$, we will usually express this more economically by saying \mathcal{E} maps A to A'. We may extend the input Hilbert space to $\mathcal{H}_A \otimes \mathcal{H}_B$, and consider the extended channel $\mathcal{E} \otimes I$ mapping ABto A'B. We say that \mathcal{E} is completely positive if any such extension of \mathcal{E} is positive.

Clearly, quantum channels are completely positive, because if \mathcal{E} has an operator-sum representation with Kraus operators $\{M_a\}$, then $\mathcal{E} \otimes I$ has an operator-sum representation with Kraus operators $\{M_a \otimes I\}$. Likewise, quantum operations, though not necessarily trace preserving, are also completely positive.

It is perfectly reasonable to demand that a channel be completely positive if it is to describe the time evolution of a quantum system — even though the channel acts on just part of the world, it should map an initial state of the whole world to a final state of the whole world. It is therefore important to note that not all positive maps are completely positive; complete positivity is a stronger condition.

A simple example is the transpose map T, mapping the d-dimensional system A to itself. In a particular basis $\{|i\rangle\}$, T acts as

$$T:|i\rangle\langle j|\mapsto |j\rangle\langle i| \tag{3.64}$$

and hence

$$T: \boldsymbol{\rho} \mapsto \boldsymbol{\rho}^T. \tag{3.65}$$

The map T is evidently positive because

$$\langle \psi | \boldsymbol{\rho}^T | \psi \rangle = \sum_{i,j} \psi_j^* \left(\boldsymbol{\rho}^T \right)_{ji} \psi_i = \sum_{i,j} \psi_i \left(\boldsymbol{\rho} \right)_{ij} \psi_j^* = \langle \psi^* | \boldsymbol{\rho} | \psi^* \rangle \qquad (3.66)$$

for any vector $|\psi\rangle$; therefore ρ^T is nonnegative if ρ is.

But T is not completely positive. Consider the (unconventionally normalized) maximally entangled state on AB, where B is also d-dimensional:

$$|\tilde{\Phi}\rangle_{AB} = \sum_{i=0}^{d-1} |i\rangle_A \otimes |i\rangle_B \equiv \sum_i |i,i\rangle.$$
(3.67)

The extension of T acts on this state as

$$T \otimes I : |\tilde{\Phi}\rangle\langle\tilde{\Phi}| = \sum_{i,j} |i\rangle\langle j| \otimes |i\rangle\langle j| \mapsto \sum_{i,j} |j\rangle\langle i| \otimes |i\rangle\langle j| \equiv \sum_{i,j} |j,i\rangle\langle i,j|;$$
(3.68)

that is, it maps $|\tilde{\Phi}\rangle\langle\tilde{\Phi}|$ to the SWAP operator which interchanges the systems A and B:

$$SWAP : |\psi\rangle_A \otimes |\varphi\rangle_B = \sum_{i,j} \psi_i \varphi_j |i,j\rangle \Longrightarrow \sum_{i,j} \varphi_j \psi_i |j,i\rangle = |\varphi\rangle_A \otimes |\psi\rangle_B$$

$$(3.69)$$

Since the square of SWAP is the identity, its eigenvalues are ± 1 . States which are symmetric under interchange of A and B have eigenvalue 1, while antisymmetric states have eigenvalue -1. Thus SWAP has negative eigenvalues, which means that $T \otimes I$ is not positive and therefore T is not completely positive.

3.3 Channel-state duality and the dilation of a channel

We have now seen that a quantum channel acting on A, which arises from a unitary map on an extension of A, is a completely positive linear map of density operators to density operators. We have also argued that linearity and complete positivity are properties that should hold for any reasonable evolution law on quantum states. It is therefore satisfying to note that any trace-preserving completely positive linear map is a quantum channel — it has an operator sum representation and a unitary realization. When considering the (in general nonunitary) evolution of A, we are entitled to imagine that A is part of an extended system which evolves unitarily.

3.3.1 Channel-state duality

To prove this statement we will use a trick which is interesting in its own right and also has other applications. For the purpose of studying the properties of a map \mathcal{E} taking A to A', we introduce an auxiliary system R with the same dimension as A, which we call the *reference* system. If \mathcal{E} is completely positive, it maps a maximally entangled state on RA to a nonnegative operator on RA'. Conversely, we may associate with any nonnegative operator on RA' a corresponding CP map taking A to A'. This correspondence between maps and states, called the *Choi-Jamiolkowski isomorphism* or *channel-state duality*, is a very useful tool.

To see how it works, consider how $I \otimes \mathcal{E}$ acts on the maximally entangled state

$$|\tilde{\Phi}\rangle_{RA} = \sum_{i=0}^{d-1} |i\rangle_R \otimes |i\rangle_A.$$
(3.70)

where A and R both have dimension d. This vector has norm \sqrt{d} instead of norm 1; we choose this unconventional normalization, highlighted by the tilde over Φ , to avoid annoying factors of d in the formulas that follow. If $\mathcal{E}_{A\to A'}$ is completely positive, then $I \otimes \mathcal{E}$ maps $|\tilde{\Phi}\rangle\langle \tilde{\Phi}|$ (up to normalization) to a density operator on RA', which like any density operator can be realized by an ensemble of pure states; hence

$$(I \otimes \mathcal{E}) \left(\left(|\tilde{\Phi}\rangle \langle \tilde{\Phi} | \right)_{RA} \right) = \sum_{a} \left(|\tilde{\Psi}_{a}\rangle \langle \tilde{\Psi}_{a} | \right)_{RA'}.$$
(3.71)

Here the normalization of $|\tilde{\Psi}_a\rangle$ may depend on a; in order to make the equation look less cluttered, we've absorbed the probability of each pure state occurring in the ensemble into that state's normalization.

Now we notice that

$$|\varphi\rangle_A = \sum_i \varphi_i |i\rangle_A = \sum_i \varphi_i \left({}_R \langle i | \tilde{\Phi} \rangle_{RA} \right) = {}_R \langle \varphi^* | \tilde{\Phi} \rangle_{RA}; \qquad (3.72)$$

using the linearity of \mathcal{E} , eq.(3.71) then implies

$$\mathcal{E}\left(\left(|\varphi\rangle\langle\varphi|\right)_{A}\right) = \sum_{a} \left(\langle\varphi^{*}|\tilde{\Psi}_{a}\rangle\langle\tilde{\Psi}_{a}|\varphi^{*}\rangle\right)_{A'}.$$
(3.73)

(This scheme for extracting the action on $|\varphi\rangle_A$ using the dual vector $_R\langle\varphi^*|$ is called the *relative-state method.*) Given a vector $|\tilde{\Phi}\rangle_{RA'}$, where R is d dimensional, we may define an operator M_a mapping \mathcal{H}_A to $\mathcal{H}_{A'}$ (where A is d dimensional) by

$$\boldsymbol{M}_{a}|\varphi\rangle_{A} = {}_{R}\langle\varphi^{*}|\tilde{\Psi}_{a}\rangle_{RA'}; \qquad (3.74)$$

it is easy to check that M_a is linear. Thus eq.(3.73) provides an operatorsum representation of \mathcal{E} acting on the pure state $(|\varphi\rangle\langle\varphi|)_A$ (and hence by linearity acting on any density operator):

$$\mathcal{E}(\boldsymbol{\rho}) = \sum_{a} \boldsymbol{M}_{a} \boldsymbol{\rho} \boldsymbol{M}_{a}^{\dagger}. \tag{3.75}$$

We have now established the desired isomorphism between states and CP maps: Eq.(3.71) tells us how to obtain a state on RA' from the channel $\mathcal{E}_{A\to A'}$, and eq.(3.71) tells us how to recover the CP map from the state. Furthermore, the $\{M_a\}$ must obey the completeness relation $\sum_a M_a^{\dagger} M_a = I$ if \mathcal{E} is trace-preserving.

Put succinctly, the argument went as follows. Because $\mathcal{E}_{A\to A'}$ is completely positive, $I \otimes \mathcal{E}$ takes a maximally entangled state on RA to a density operator on RA', up to normalization. This density operator can be expressed as an ensemble of pure states, and each of these pure states is associated with a Kraus operator in the operator-sum representation of \mathcal{E} .

From this viewpoint, we see that the freedom to choose the Kraus operators representing a channel in many different ways is really the same thing as the freedom to choose the ensemble of pure states representing a density operator in many different ways. According to the HJW theorem, two different ensemble realizations of the same density operator,

$$(I \otimes \mathcal{E}) \left(\left(|\tilde{\Phi}\rangle \langle \tilde{\Phi}| \right)_{RA} \right) = \sum_{a} \left(|\tilde{\Psi}_{a}\rangle \langle \tilde{\Psi}_{a}| \right)_{RA'} = \sum_{\mu} \left(|\tilde{\gamma}_{\mu}\rangle \langle \tilde{\gamma}_{\mu}| \right)_{RA'}, \quad (3.76)$$

are related by a unitary change of basis,

$$|\tilde{\gamma}_{\mu}\rangle = \sum_{a} V_{\mu a} |\tilde{\Psi}_{a}\rangle. \tag{3.77}$$

Correspondingly, two different sets of Kraus operators $\{M_a\}$ and $\{N_\mu\}$ representing the same channel are related by

$$\boldsymbol{N}_{\mu} = \sum_{a} V_{\mu a} \boldsymbol{M}_{a} \tag{3.78}$$

where $V_{\mu a}$ is a unitary matrix.

21

Channel-state duality also makes clear how many Kraus operators are needed to describe a channel. A channel \mathcal{E} mapping A to A', where A has dimension d and A' has dimension d', is equivalent to a density operator on RA', where R has dimension d, and the minimal number of Kraus operators needed to represent the channel is the same as the minimal number of pure states needed in an ensemble representation of the density operator. This is the density operator's rank (number of nonzero eigenvalues), which is no larger than dd'. Of course, there may be operator-sum representations of \mathcal{E} which use many more than this minimal number of Kraus operators, just as an ensemble representation of a density operator might use many more than the minimal number of pure states.

The number of free parameters needed to specify a channel mapping A to A' is the number $(dd')^2$ needed to specify a density operator on RA', except that there are d^2 constraints because the map is trace preserving for each of d^2 linearly independent inputs. Therefore the number of real free parameters is

$$d^2 \left(d'^2 - 1 \right). \tag{3.79}$$

This is 12 parameters for a general channel taking qubits to qubits. In contrast, a unitary map on qubits has only 3 parameters, aside from the physically irrelevant overall phase.

3.3.2 Stinespring dilation

Once we have an operator-sum representation of the channel $\mathcal{E}_{A\to A'}$, it is easy to see how \mathcal{E} can be realized by a unitary map acting on an extended system. We introduce an extra system E, the channel's *environment*, which has dimension equal to the number of Kraus operators and orthonormal basis $\{|a\rangle\}$. Then we define an inner-product preserving map (an *isometry*) which takes A to A'E according to

$$\boldsymbol{U}_{A\to A'E}: |\psi\rangle \mapsto \sum_{a} \boldsymbol{M}_{a} |\psi\rangle \otimes |a\rangle.$$
(3.80)

The completeness relation satisfied by the $\{M_a\}$ implies $U^{\dagger}U = I_A$. Though U may not actually be unitary, it might as well be, because we can easily extend an isometry to a unitary transformation by expanding the input Hilbert space. This isometry, which yields $\mathcal{E}_{A \to A'}$ when we trace out the environment, is called the *Stinespring dilation* of the channel

Another way to think about the construction of the Stinespring dilation is that we have used E to construct a *purification* of the density operator arising from channel-state duality:

$$|\bar{\Phi}\rangle_{RA'E} = \sum_{a} |\tilde{\Phi}_{a}\rangle_{RA'} \otimes |a\rangle_{E}.$$
(3.81)

Apart from a normalization factor of \sqrt{d} , this is the pure state of RA'E that results when the dilation acts on the maximally entangled state $|\tilde{\Phi}\rangle_{RA}$; we may recover the dilation from $|\bar{\Phi}\rangle$ using

$$\boldsymbol{U}_{A\to A'E}|\psi\rangle_A = {}_{R}\langle\psi^*|\bar{\Phi}\rangle_{RA'E}.$$
(3.82)

This succinct way to characterize a channel using a pure state is sometimes quite convenient, and we'll make heavy use of it when studying quantum channels in Chapter 10.

3.3.3 Axioms revisited

In Chapter 2 we stated the axioms of quantum mechanics in a form appropriate for closed systems. With the theory of open systems now in hand, we can give an alternative formulation with revised notions of how states, measurements, and evolution are described.

- **States**. A state is a *density operator*, a nonnegative Hermitian operator in Hilbert space with unit trace.
- **Measurement**. A measurement is a *positive operator-valued measure* (*POVM*), a partition of unity by nonnegative operators. When the measurement $\{E_a\}$ is performed on the state ρ , the outcome *a* occurs with probability tr $(E_a\rho)$.
- **Dynamics**. Time evolution is described by a *trace-preserving completely* positive map (TPCP map).

One could regard either the open-system or closed-system version as the fundamental formulation of the theory; it's really a matter of taste. We have already seen how the open-system axioms are obtained starting from the closed-system axioms. Alternatively, starting with the opensystem axioms, pure states arise as the extremal points in the space of density operators, or from the observation that every density operator has a purification in an extended system. Similarly, orthogonal measurements and unitary evolution arise naturally because every POVM can be realized by an orthogonal measurement in an extended system, and every tracepreserving completely positive map has an isometric Stinespring dilation. The notion that an open system may always be regarded as part of a larger closed system is fondly known as the *church of the larger Hilbert space*.

23

3.4 Three quantum channels

The best way to familiarize ourselves with the concept of a quantum channel is to study a few examples. We will now consider three examples (all interesting and useful) of channels acting on a single qubit. If we wish we may imagine that the channel \mathcal{E} describes the fate of quantum information that is transmitted with some loss of fidelity from a sender to a receiver. Or, if we prefer, we may imagine that the transmission is in time rather than space; that is, \mathcal{E} describes the time evolution of a quantum system that interacts with its environment.

3.4.1 Depolarizing channel

The *depolarizing channel* is a model of a decohering qubit that has particularly nice symmetry properties. We can describe it by saying that, with probability 1 - p the qubit remains intact, while with probability pan "error" occurs. The error can be of any one of three types, where each type of error is equally likely. If $\{|0\rangle, |1\rangle\}$ is an orthonormal basis for the qubit, the three types of errors can be characterized as:

1. Bit flip error:
$$|0\rangle \mapsto |1\rangle \\ |1\rangle \mapsto |0\rangle$$
 or $|\psi\rangle \mapsto \sigma_1 |\psi\rangle, \sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$,

2. Phase flip error:
$$|0\rangle \mapsto |0\rangle \\ |1\rangle \mapsto -|1\rangle$$
 or $|\psi\rangle \mapsto \sigma_3 |\psi\rangle, \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$,

3. Both:
$$|0\rangle \mapsto +i|1\rangle \\ |1\rangle \mapsto -i|0\rangle$$
 or $|\psi\rangle \mapsto \sigma_2 |\psi\rangle, \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$.

If an error occurs, then $|\psi\rangle$ evolves to an ensemble of the three states $\sigma_1 |\psi\rangle, \sigma_2 |\psi\rangle, \sigma_3 |\psi\rangle$, all occurring with equal likelihood.

Unitary representation. The depolarizing channel mapping qubit A to A can be realized by an isometry mapping A to AE, where E is a fourdimensional environment, acting as

$$U_{A\to AE} : |\psi\rangle_A \mapsto \sqrt{1-p} \ |\psi\rangle_A \otimes |0\rangle_E + \sqrt{\frac{p}{3}} \left(\sigma_1 |\psi\rangle_A \otimes |1\rangle_E + \sigma_2 |\psi\rangle_A \otimes |2\rangle_E + \sigma_3 |\psi\rangle_A \otimes |3\rangle_E \right).$$
(3.83)

The environment evolves to one of four mutually orthogonal states that "keep a record" of what transpired; if we could only measure the environment in the basis $\{|a\rangle_E, a = 0, 1, 2, 3\}$, we would know what kind of error had occurred (and we would be able to intervene and reverse the error).

Operator-sum representation. To obtain an operator-sum representation of the channel, we evaluate the partial trace over the environment in the $\{|a\rangle_E\}$ basis. Then

$$\boldsymbol{M}_{a} = {}_{E} \langle a | \boldsymbol{U}, \qquad (3.84)$$

so that

$$M_0 = \sqrt{1-p} \ I, \ M_1 = \sqrt{\frac{p}{3}} \ \sigma_1, \ M_2 = \sqrt{\frac{p}{3}} \ \sigma_2, \ M_3 = \sqrt{\frac{p}{3}} \ \sigma_3.$$
 (3.85)

Using $\sigma_i^2 = I$, we can readily check the normalization condition

$$\sum_{a} \boldsymbol{M}_{a}^{\dagger} \boldsymbol{M}_{a} = \left((1-p) + 3\frac{p}{3} \right) \boldsymbol{I} = \boldsymbol{I}.$$
(3.86)

A general initial density matrix ρ of the qubit evolves as

$$\boldsymbol{\rho} \mapsto \boldsymbol{\rho}' = (1-p)\boldsymbol{\rho} + \frac{p}{3} \left(\boldsymbol{\sigma}_1 \boldsymbol{\rho} \boldsymbol{\sigma}_1 + \boldsymbol{\sigma}_2 \boldsymbol{\rho} \boldsymbol{\sigma}_2 + \boldsymbol{\sigma}_3 \boldsymbol{\rho} \boldsymbol{\sigma}_3 \right).$$
(3.87)

where we are summing over the four (in principle distinguishable) possible final states of the environment.

Relative-state representation. We can also characterize the channel by introducing a reference qubit R and describing how a maximally-entangled state of the two qubits RA evolves, when the channel acts only on A. There are four mutually orthogonal maximally entangled states, which may be denoted

$$\begin{aligned} |\phi^{+}\rangle &= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), \\ |\phi^{-}\rangle &= \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle), \\ |\psi^{+}\rangle &= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle), \\ |\psi^{-}\rangle &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle). \end{aligned}$$
(3.88)

If the initial state is $|\phi^+\rangle_{RA}$, then when the depolarizing channel acts on qubit A, the entangled state evolves as

$$|\phi^{+}\rangle\langle\phi^{+}|\mapsto(1-p)|\phi^{+}\rangle\langle\phi^{+}| + \frac{p}{3}\left(|\psi^{+}\rangle\langle\psi^{+}| + |\psi^{-}\rangle\langle\psi^{-}| + |\phi^{-}\rangle\langle\phi^{-}|\right).$$
(3.89)
The "worst possible" quantum channel has p = 3/4, for in that case the initial entangled state evolves as

$$|\phi^{+}\rangle\langle\phi^{+}|\mapsto\frac{1}{4}\left(|\phi^{+}\rangle\langle\phi^{+}|+|\phi^{-}\rangle\langle\phi^{-}|+|\psi^{+}\rangle\langle\psi^{+}|+|\psi^{-}\rangle\langle\psi^{-}|\right) = \frac{1}{4}\boldsymbol{I};$$
(3.90)

it becomes the maximally mixed density matrix on RA. By the relativestate method, then, we see that a pure state $|\psi\rangle$ of qubit A evolves as

$$(|\psi\rangle\langle\psi|)_A \mapsto {}_R\langle\psi^*|2\left(\frac{1}{4}\boldsymbol{I}_{RA}\right)|\psi^*\rangle_R = \frac{1}{2}\boldsymbol{I}_A, \qquad (3.91)$$

where the factor of two has been inserted because here we have used the standard normalization of the entangled states, instead of the unconventional normalization used in our earlier discussion of the relative-state method. We see that, for p = 3/4, the qubit is mapped to the maximally mixed density operator on A, irrespective of the value of the initial state $|\psi\rangle_A$. It is as though the channel threw away the initial quantum state, and replaced it by completely random junk.

An alternative way to express the evolution of the maximally entangled state is

$$|\phi^+\rangle\langle\phi^+|\mapsto \left(1-\frac{4}{3}p\right)|\phi^+\rangle\langle\phi^+|+\frac{4}{3}p\left(\frac{1}{4}\boldsymbol{I}_{RA}\right).$$
(3.92)

Thus instead of saying that an error occurs with probability p, with errors of three types all equally likely, we could instead say that an error occurs with probability 4/3p, where the error completely "randomizes" the state (at least we can say that for $p \leq 3/4$). The existence of two natural ways to define an "error probability" for this channel can sometimes cause confusion.

One useful measure of how well the channel preserves the original quantum information is called the "entanglement fidelity" F_e . It quantifies how "close" the final density matrix is to the original maximally entangled state $|\phi^+\rangle$ after the action of $I \otimes \mathcal{E}$:

$$F_e = \langle \phi^+ | \boldsymbol{\rho}' | \phi^+ \rangle. \tag{3.93}$$

For the depolarizing channel, we have $F_e = 1 - p$, and we can interpret F_e as the probability that no error occurred.

Bloch-sphere representation. It is also instructive to see how the depolarizing channel acts on the Bloch sphere. An arbitrary density matrix for a single qubit can be written as

$$\boldsymbol{\rho}(\vec{P}) = \frac{1}{2} \left(\boldsymbol{I} + \vec{P} \cdot \vec{\boldsymbol{\sigma}} \right), \qquad (3.94)$$

where \vec{P} is the "spin polarization" of the qubit. The depolarizing channel maps this state to

$$\boldsymbol{\rho}' = \left(1 - \frac{4}{3}p\right)\boldsymbol{\rho} + \frac{4}{3}p\boldsymbol{I} = \boldsymbol{\rho}(\vec{P}') \tag{3.95}$$

where

$$\vec{P}' = \left(1 - \frac{4}{3}p\right)\vec{P} \tag{3.96}$$

Hence the Bloch sphere contracts uniformly under the action of the channel (for $p \leq 3/4$); the spin polarization shrinks by the factor $1 - \frac{4}{3}p$ (which is why we call it the depolarizing channel).

Reversibility? Why do we say that the channel is not invertible? Evidently we can reverse a uniform contraction of the sphere with a uniform inflation. But the trouble is that the inflation of the Bloch sphere is not a channel, because it is not positive. Inflation will take some values of \vec{P} with $|\vec{P}| \leq 1$ to values with $|\vec{P}| > 1$, and so will take a density operator to an operator with a negative eigenvalue. Decoherence can shrink the ball, but no physical process can blow it up again! A channel running backwards in time is *not* a channel.

3.4.2 Dephasing channel

Our next example is the *dephasing channel*, also called the *phase-damping channel*. This case is particularly instructive, because it provides a revealing caricature of decoherence in realistic physical situations, with all inessential mathematical details stripped away.

Unitary representation. An isometric representation of the channel is

$$|0\rangle_{A} \mapsto \sqrt{1-p} |0\rangle_{A} \otimes |0\rangle_{E} + \sqrt{p} |0\rangle_{A} \otimes |1\rangle_{E}, |1\rangle_{A} \mapsto \sqrt{1-p} |1\rangle_{A} \otimes |0\rangle_{E} + \sqrt{p} |1\rangle_{A} \otimes |2\rangle_{E}.$$
(3.97)

In this case, unlike the depolarizing channel, qubit A does not make any transitions in the $\{|0\rangle, |1\rangle\}$ basis. Instead, the environment "scatters" off of the qubit occasionally (with probability p), being kicked into the state $|1\rangle_E$ if A is in the state $|0\rangle_A$ and into the state $|2\rangle_E$ if A is in the state $|1\rangle_A$. Furthermore, also unlike the depolarizing channel, the channel picks out a preferred basis for qubit A; the basis $\{|0\rangle, |1\rangle\}$ is the only basis in which bit flips never occur.

Kraus operators. Evaluating the partial trace over E in the $\{|0\rangle_E, |1\rangle_E, |2\rangle_E\}$ basis, we obtain the Kraus operators

$$\boldsymbol{M}_0 = \sqrt{1-p} \boldsymbol{I}, \quad \boldsymbol{M}_1 = \sqrt{p} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad \boldsymbol{M}_2 = \sqrt{p} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}.$$
 (3.98)

It is easy to check that $M_0^2 + M_1^2 + M_2^2 = I$. In this case, three Kraus operators are not really needed; a representation with two Kraus operators is possible. Expressing

$$\boldsymbol{M}_{1} = \frac{\sqrt{p}}{2} \left(\boldsymbol{I} + \boldsymbol{\sigma}_{3} \right), \quad \boldsymbol{M}_{2} = \frac{\sqrt{p}}{2} \left(\boldsymbol{I} - \boldsymbol{\sigma}_{3} \right), \quad (3.99)$$

we find

$$\mathcal{E}(\boldsymbol{\rho}) = \sum_{a} \boldsymbol{M}_{a} \boldsymbol{\rho} \boldsymbol{M}_{a} = \left(1 - \frac{1}{2}p\right) \boldsymbol{\rho} + \frac{1}{2}p \ \boldsymbol{\sigma}_{3} \boldsymbol{\rho} \boldsymbol{\sigma}_{3}, \qquad (3.100)$$

so an alternative description of the channel is that σ_3 is applied with probability p/2 and nothing happens with probability (1 - p/2). An initial density matrix ρ evolves to

$$\mathcal{E}\left(\begin{array}{cc}\rho_{00} & \rho_{01}\\\rho_{10} & \rho_{11}\end{array}\right) = \left(\begin{array}{cc}\rho_{00} & (1-p)\rho_{01}\\(1-p)\rho_{10} & \rho_{11}\end{array}\right);$$
(3.101)

the on-diagonal terms in ρ remain invariant while the off-diagonal terms decay.

Continuous dephasing. We may also consider dephasing that occurs continuously in time. Suppose that the probability of a scattering event per unit time is Γ , so that $p = \Gamma \Delta t \ll 1$ when a brief time interval Δt elapses. The evolution over a time $t = n\Delta t$ is governed by \mathcal{E}^n (\mathcal{E} repeated *n* times in succession), so that the off-diagonal terms in the density operator become suppressed by

$$(1-p)^n = (1 - \Gamma t/n)^n \to e^{-\Gamma t},$$
 (3.102)

taking the limit $n \to 0$ with t fixed. Thus, if we prepare an initial pure state $\alpha |0\rangle + \beta |1\rangle$, then after a time $t \gg \Gamma^{-1}$, the density operator evolves as

$$\begin{pmatrix} |\alpha|^2 & \alpha\beta^*\\ \alpha^*\beta & |\beta|^2 \end{pmatrix} \mapsto \begin{pmatrix} |\alpha|^2 & 0\\ 0 & |\beta|^2 \end{pmatrix};$$
(3.103)

The state decoheres, in the preferred basis $\{|0\rangle, |1\rangle\}$.

Bloch-sphere representation. We can compute how the polarization of the density operator evolves using the representation of the channel eq.(3.100), finding

$$\boldsymbol{\rho}(\vec{P}) = \frac{1}{2} \left(\boldsymbol{I} + \vec{P} \cdot \vec{\boldsymbol{\sigma}} \right) \mapsto \boldsymbol{\rho}(\vec{P}') \tag{3.104}$$

where

$$P'_{1,2} = (1-p)P_{1,2}, \quad P'_3 = P_3;$$
 (3.105)

the Bloch ball shrinks to a prolate spheroid aligned with the z axis. Under continuous dephasing, the ball deflates in the x-y plane, degenerating to the z axis in the limit of large Γt .

You might wonder whether there is a quantum channel which causes just one component of the polarization to decay, mapping the Bloch ball to an oblate spheroid which touches the Bloch sphere along its equator. In fact no such map can be completely positive (the *no-pancake theorem*).

Interpretation. We might interpret the phase-damping channel as describing a heavy "classical" particle (e.g., an interstellar dust grain) interacting with a background gas of light particles (e.g., the 3K microwave photons). We can imagine that the dust is initially prepared in a superposition of position eigenstates $|\psi\rangle = \frac{1}{\sqrt{2}}(|x\rangle + |-x\rangle)$ (or more realistically a superposition of position-space wavepackets with little overlap). We might be able to monitor the behavior of the dust particle, but it is hopeless to keep track of the quantum state of all the photons that scatter from the particle; for our purposes, the quantum state of the particle is described by the density matrix ρ obtained by tracing over the photon degrees of freedom.

Our analysis of the phase damping channel indicates that if photons are scattered by the dust particle at a rate Γ , then the off-diagonal terms in ρ decay like $\exp(-\Gamma t)$, and so become completely negligible for $t \gg \Gamma^{-1}$. At that point, the coherence of the superposition of position eigenstates is completely lost – there is no chance that we can recombine the wavepackets and induce them to interfere. If we attempt to do a double-slit interference experiment with dust grains, we will not see any interference pattern if it takes a time $t \gg \Gamma^{-1}$ for the grain to travel from the source to the screen.

The dust grain is heavy. Because of its large inertia, its state of motion is little affected by the scattered photons. Thus, there are two disparate time scales relevant to its dynamics. On the one hand, there is a damping time scale, the time for a significant amount of the particle's momentum to be transfered to the photons, which is a long time for such a heavy particle. On the other hand, there is the decoherence time scale. In this model, the time scale for decoherence is of order Γ , the time for a *single* photon to be scattered by the dust grain, which is far shorter than the damping time scale. For a macroscopic object, decoherence is *fast*.

As we have already noted, the phase-damping channel picks out a preferred basis for decoherence, which in our "interpretation" we have assumed to be the position-eigenstate basis. Physically, decoherence prefers the spatially localized states of the dust grain because the *interactions* of photons and grains are localized in space. Grains in distinguishable positions tend to scatter the photons of the environment into mutually orthogonal states.

Even if the separation between the "grains" were so small that it could not be resolved very well by the scattered photons, the decoherence process would still work in a similar way. Perhaps photons that scatter off grains at positions x and -x are not mutually orthogonal, but instead have an overlap

$$\langle \gamma + | \gamma - \rangle = 1 - \varepsilon, \quad \varepsilon \ll 1.$$
 (3.106)

The phase-damping channel would still describe this situation, but with p replaced by $p\varepsilon$ (if p is still the probability of a scattering event). Thus, the decoherence rate would become $\Gamma_{dec} = \varepsilon \Gamma_{scat}$, where Γ_{scat} is the scattering rate.

The intuition we distill from this simple model applies to a wide variety of physical situations. A coherent superposition of macroscopically distinguishable states of a "heavy" object decoheres very rapidly compared to its damping rate. The spatial locality of the interactions of the system with its environment gives rise to a preferred "local" basis for decoherence. The same principle applies to Schrödinger's unfortunate cat, delicately prepared in a coherent superposition of its dead state and its alive state, two states that are easily distinguished by spatially localized probes. The cat quickly interacts with its environment, which is "scattered" into one of two mutually orthogonal states perfectly correlated with the cat's state in the $\{|\text{dead}\rangle, |\text{alive}\rangle\}$ basis, thus transforming the cat into an incoherent mixture of those two basis states.

Visibility. On the other hand, for microscopic systems the time scale for decoherence need not be short compared to dynamical time scales. Consider for example a single atom, initially prepared in a uniform superposition of its ground state $|0\rangle$ and an excited state $|1\rangle$ with energy $\hbar\omega$ above the ground state energy. Neglecting decoherence, after time t the atom's state will be

$$|\psi(t)\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle + e^{-i\omega t}|1\rangle\right). \tag{3.107}$$

If dephasing occurs in the $\{|0\rangle, |1\rangle\}$ basis with rate Γ , the off-diagonal terms in the density operator decay, yielding the density operator

$$\boldsymbol{\rho}(t) = \frac{1}{2} \begin{pmatrix} 1 & e^{i\omega t} e^{-\Gamma t} \\ e^{-i\omega t} e^{-\Gamma t} & 1 \end{pmatrix}.$$
(3.108)

If after time t we measure the atom in the basis

$$|\pm\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle \pm |1\rangle\right), \qquad (3.109)$$

the probability of the + outcome is

$$\operatorname{Prob}(+,t) = \langle +|\boldsymbol{\rho}(t)|+\rangle = \frac{1}{2} \left(1 + e^{-\Gamma t} \cos \omega t\right).$$
(3.110)

In principle this time dependence of the probability can be measured by varying the time t between the preparation and measurement, and by repeating the experiment many times for each t to estimate the probability with high statistical confidence. The decoherence rate Γ can be determined experimentally by fitting the declining *visibility* of the coherent oscillations of Prob(+, t) to a decaying exponential function of t.

3.4.3 Amplitude-damping channel

The *amplitude-damping channel* is a schematic model of the decay of an excited state of a (two-level) atom due to spontaneous emission of a photon. By detecting the emitted photon ("observing the environment") we can perform a POVM that gives us information about the initial preparation of the atom.

Unitary representation. We denote the atomic ground state by $|0\rangle_A$ and the excited state of interest by $|1\rangle_A$. The "environment" is the electromagnetic field, assumed initially to be in its vacuum state $|0\rangle_E$. After we wait a while, there is a probability p that the excited state has decayed to the ground state and a photon has been emitted, so that the environment has made a transition from the state $|0\rangle_E$ ("no photon") to the state $|1\rangle_E$ ("one photon"). This evolution is described by a unitary transformation that acts on atom and environment according to

$$|0\rangle_A \otimes |0\rangle_E \mapsto |0\rangle_A \otimes |0\rangle_E |1\rangle_A \otimes |0\rangle_E \mapsto \sqrt{1-p} |1\rangle_A \otimes |0\rangle_E + \sqrt{p} |0\rangle_A \otimes |1\rangle_E.$$
(3.111)

(Of course, if the atom starts out in its ground state, and the environment in its vacuum state, then no transition occurs.) Kraus operators. By evaluating the partial trace over the environment in the basis $\{|0\rangle_E, |1\rangle_E\}$, we find the Kraus operators

$$\boldsymbol{M}_{0} = \begin{pmatrix} 1 & 0\\ 0 & \sqrt{1-p} \end{pmatrix}, \quad \boldsymbol{M}_{1} = \begin{pmatrix} 0 & \sqrt{p}\\ 0 & 0 \end{pmatrix}, \quad (3.112)$$

and we can check that

$$\boldsymbol{M}_{0}^{\dagger}\boldsymbol{M}_{0} + \boldsymbol{M}_{1}^{\dagger}\boldsymbol{M}_{1} = \begin{pmatrix} 1 & 0 \\ 0 & 1-p \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & p \end{pmatrix} = \boldsymbol{I}.$$
 (3.113)

The operator M_1 induces a "quantum jump," the decay from $|1\rangle_A$ to $|0\rangle_A$, and M_0 describes how the state changes if no jump occurs. The density matrix evolves as

$$\boldsymbol{\rho} \mapsto \mathcal{E}(\boldsymbol{\rho}) = \boldsymbol{M}_{0} \boldsymbol{\rho} \boldsymbol{M}_{0}^{\dagger} + \boldsymbol{M}_{1} \boldsymbol{\rho} \boldsymbol{M}_{1}^{\dagger} \\
= \begin{pmatrix} \rho_{00} & \sqrt{1-p} \rho_{01} \\ \sqrt{1-p} \rho_{10} & (1-p) \rho_{11} \end{pmatrix} + \begin{pmatrix} p\rho_{11} & 0 \\ 0 & 0 \end{pmatrix} \\
= \begin{pmatrix} \rho_{00} + p\rho_{11} & \sqrt{1-p} \rho_{01} \\ \sqrt{1-p} \rho_{10} & (1-p) \rho_{11} \end{pmatrix}.$$
(3.114)

Time dependence. If Γ is the spontaneous decay rate per unit time, then the decay occurs with probability $p = \Gamma \Delta t \ll 1$ in a small time interval Δt . We find the density operator after time $t = n\Delta t$ by applying the channel *n* times in succession. The ρ_{11} matrix element then decays as

$$\rho_{11} \mapsto (1-p)^n \rho_{11} = (1-\Gamma t/n)^n \to e^{-\Gamma t},$$
(3.115)

the expected exponential decay law, while the off-diagonal entries decay by the factor $(1-p)^{n/2} = e^{-\Gamma t/2}$; hence we find

$$\boldsymbol{\rho}(t) = \begin{pmatrix} \rho_{00} + (1 - e^{-\Gamma t}) \rho_{11} & e^{-\Gamma t/2} \rho_{01} \\ e^{-\Gamma t/2} \rho_{10} & e^{-\Gamma t} \rho_{11} \end{pmatrix}$$
(3.116)

It is customary to use " T_1 " to denote the exponential decay time for the excited population, and to use " T_2 " to denote the exponential decay time for the off-diagonal terms in the density operator. In some systems where dephasing is very rapid T_2 is much shorter than T_1 , but we see that for the amplitude-damping channel these two times are related and comparable:

$$T_2 = 2\Gamma^{-1} = 2T_1. \tag{3.117}$$

By the time that $t \gg T_1$, the atom is in its ground state with high probability $(\rho_{00}(t) \approx 1)$.

Watching the environment. So far we have described the evolution of the qubit under the assumption that the state of the environment is not observed. But now suppose we surround the atom with photon detectors, so we know whether a photon has been emitted or not. Rather than a channel, then, we consider a POVM performed on the atom.

Returning to the joint unitary dynamics of system and environment, we see that a coherent superposition of the atomic ground and excited states evolves as

$$(\alpha|0\rangle_A + \beta|1\rangle_A) \otimes |0\rangle_E \mapsto (\alpha|0\rangle_A + \beta\sqrt{1-p} |1\rangle_A) \otimes |0\rangle_E + \sqrt{p} |0\rangle_A \otimes |1\rangle_E; \quad (3.118)$$

To describe the system evolving continuously in time, we may consider applying this unitary map $n \gg 1$ times in succession, but where photons emitted at different times are perfectly distinguishable and hence orthogonal. The resulting POVM has n+1 Kraus operators, associated with the vacuum state of the environment and n different possible single photon states:

$$\boldsymbol{M}_{0} = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{(1-p)^{n}} \end{pmatrix}, \quad \boldsymbol{M}_{k} = \begin{pmatrix} 0 & \sqrt{(1-p)^{k-1}p} \\ 0 & 0 \end{pmatrix}, \quad (3.119)$$

for k = 1, 2, ..., n. Taking the continuous-time limit we find that if no spontaneous decay occurs for time t, the corresponding Kraus operator is

$$\boldsymbol{M}_0 = \begin{pmatrix} 1 & 0\\ 0 & e^{-\Gamma t/2} \end{pmatrix}. \tag{3.120}$$

If we detect a photon (and so project out a single-photon state of the environment), then we have prepared the state $|0\rangle_A$ of the atom. Not only that, we have prepared a state in which we know with certainty that the initial atomic state was the excited state $|1\rangle_A$; if the atom had started out in the ground state than it could not have decayed and no photon could have been detected.

On the other hand, if we detect no photon, and our photon detector has perfect efficiency, then we have projected out the vacuum state of the environment, and so have prepared the atomic state

$$\boldsymbol{M}_{0}\left(\alpha|0\rangle + \beta|1\rangle\right) = \alpha|0\rangle + e^{-\Gamma t/2}\beta|1\rangle, \qquad (3.121)$$

- - -

up to a normalization factor. As time goes by, the *a posteriori* quantum state has larger and larger overlap with the ground state, because if it had started out in the excited state it should have decayed by now. In the limit $t \to \infty$ our POVM becomes an orthogonal measurement: either a photon is detected, in which case the initial state of the atom must have

been $|1\rangle$, or no photon is detected, in which case the initial state must have been $|0\rangle$. It's odd but true: we can project out the state $|0\rangle$ of the atom by *not* detecting anything.

3.5 Master equations for open quantum systems

3.5.1 Markovian evolution

Quantum channels provide a general description of the evolution of density operators, including the evolution of pure states to mixed states (decoherence). In the same sense, unitary transformations provide a general description of coherent quantum evolution. But in the case of coherent evolution, we often find it very convenient to characterize the dynamics of a quantum system with a *Hamiltonian*, which describes the evolution over an infinitesimal time interval. The dynamics is then encoded in a differential equation, the *Schrödinger equation*, and we may calculate the evolution over a finite time interval by integrating the equation, that is, by piecing together the evolution over many infinitesimal intervals. Likewise, it is often possible to describe the (not necessarily coherent) evolution of a density operator, at least to a good approximation, by a differential equation which is called the *master equation*.

It is not obvious that there should be a differential equation that describes the decoherence of an open system. Such a description is possible only if the evolution of the quantum system is *Markovian*, that is, *local* in time. For the evolution of the density operator $\rho(t)$ to be governed by a (first-order) differential equation in t, $\rho(t + dt)$ must be completely determined by $\rho(t)$.

In the case of an open system A, we are to imagine that its evolution is actually unitary on the extended system AE, where E is the environment. But though the evolution of AE may be governed by a Schrödinger equation, that's not enough to ensure that the time evolution is Markovian for A by itself. The trouble is that information can flow from A to E and then return at a later time. In that case the density operator $\rho_A(t + dt)$ is not fully determined by $\rho_A(t)$; we need to know ρ_A at earlier times as well.

This quandary arises because information flow is a two-way street. An open system (whether classical or quantum) is *dissipative* because information and energy can flow from the system to the environment. But that means that information can also flow back from environment to system, resulting in non-Markovian *fluctuations* of the system. This inescapable connection underlies the fluctuation-dissipation theorem, a widely applicable tool of statistical physics.

For any open system these fluctuations are inevitable, and an exact

Markovian description of quantum dynamics is impossible. Nevertheless, a Markovian description can be a very good approximation if there is a clean separation between the typical correlation time of the fluctuations and the time scale of the evolution that we want to follow. Crudely speaking, we may denote by $(\Delta t)_{env}$ the time it takes for the environment to "forget" information it acquired from the system — after time $(\Delta t)_{env}$ we can regard that information as lost forever, and neglect the possibility that the information may return to influence the subsequent evolution of the system.

To describe the evolution we "coarse-grain" in time, perceiving the dynamics through a filter that screens out the high frequency components of the motion with $\omega \gg (\Delta t_{\text{coarse}})^{-1}$. An approximately Markovian description should be possible for $(\Delta t)_{\text{env}} \ll (\Delta t)_{\text{coarse}}$; we may neglect the memory of the reservoir if we are unable to resolve its effects. This Markovian approximation is *useful* if the time scale of the dynamics that we want to observe is long compared to $(\Delta t)_{\text{coarse}}$, for example if the *damping* time scale $(\Delta t)_{\text{damp}}$ satisfies

$$(\Delta t)_{\text{damp}} \gg (\Delta t)_{\text{coarse}} \gg (\Delta t)_{\text{env}}.$$
 (3.122)

This is a good approximation in some physical settings, like an atom interacting with the radiation field, but more dubious in other cases, like an electron spin interacting with nuclear spins in a semiconductor.

We could attempt to derive the master equation starting with the Schrödinger equation for AE, treating the coupling between A and E in time-dependent perturbation theory, and carefully introducing a frequency cutoff, but we won't do that here. Instead let's take it for granted that the dynamics is Markovian, and use the theory of quantum channels to infer the form of the master equation.

3.5.2 The Liouvillian

For a closed quantum system, time evolution is governed by a self-adjoint Hamiltonian H according to

$$|\psi(t+dt)\rangle = (\boldsymbol{I} - idt\boldsymbol{H}) |\psi(t)\rangle, \qquad (3.123)$$

and correspondingly the density operator evolves as

$$\boldsymbol{\rho}(t+dt) = \boldsymbol{\rho}(t) - idt[\boldsymbol{H}, \boldsymbol{\rho}(t)]. \tag{3.124}$$

In the case of an open quantum system, Markovian evolution for the infinitesimal time interval dt may be expressed as

$$\boldsymbol{\rho}(t+dt) = \mathcal{E}_{dt}(\boldsymbol{\rho}(t)), \qquad (3.125)$$

where \mathcal{E}_{dt} is a quantum channel. By adopting this Markovian form, we take the view that, after each infinitesimal time increment in the joint evolution of the system and its environment, the state of the environment is discarded and replaced by a fresh state of the environment unentangled with the system. We already made this assumption implicitly when discussing continuous-time dephasing and spontaneous decay in §3.4.

Expanding \mathcal{E}_{dt} to linear order,

$$\mathcal{E}_{dt} = \mathbf{I} + dt\mathcal{L} \tag{3.126}$$

we find

$$\dot{\boldsymbol{\rho}} = \mathcal{L}(\boldsymbol{\rho}), \tag{3.127}$$

where the linear map \mathcal{L} generating time evolution is called the *Liouvillian* or *Lindbladian*. This evolution equation has the formal solution

$$\boldsymbol{\rho}(t) = \lim_{n \to \infty} \left(1 + \frac{\mathcal{L}t}{n} \right)^n \left(\boldsymbol{\rho}(0) \right) = e^{\mathcal{L}t}(\boldsymbol{\rho}(0)) \tag{3.128}$$

if \mathcal{L} is time independent.

The channel has an operator-sum representation

$$\boldsymbol{\rho}(t+dt) = \mathcal{E}_{dt}(\boldsymbol{\rho}(t)) = \sum_{a} \boldsymbol{M}_{a} \boldsymbol{\rho}(t) \boldsymbol{M}_{a}^{\dagger} = \boldsymbol{\rho}(t) + O(dt), \qquad (3.129)$$

where, if we retain only terms up to linear order in dt, we may assume without loss of generality that $\mathbf{M}_0 = \mathbf{I} + O(dt)$, and that \mathbf{M}_a is of order \sqrt{dt} for a > 0. Each of the Kraus operators $\mathbf{M}_{1,2,\dots}$ describes a possible "quantum jump" that the system might undergo, which occurs during time interval dt with probability O(dt), and \mathbf{M}_0 describes how the system evolves when no jump occurs. We may write

$$M_0 = \mathbf{I} + (-i\mathbf{H} + \mathbf{K})dt,$$

$$M_a = \sqrt{dt} \mathbf{L}_a, \quad a = 1, 2, 3,$$
(3.130)

where H and K are both hermitian and L_a, H , and K are all zeroth order in dt. In fact, we can determine K by invoking the Kraus-operator completeness relation; keeping terms up to linear order in O(dt), we find

$$\mathbf{I} = \sum_{\mu} \mathbf{M}_{a}^{\dagger} \mathbf{M}_{a} = \mathbf{I} + dt \left(2\mathbf{K} + \sum_{\mu > 0} \mathbf{L}_{a}^{\dagger} \mathbf{L}_{a} \right) + \cdots, \qquad (3.131)$$

or

$$\boldsymbol{K} = -\frac{1}{2} \sum_{a>0} \boldsymbol{L}_a^{\dagger} \boldsymbol{L}_a. \tag{3.132}$$

Substituting into eq. (3.129), we obtain the Lindblad master equation:

$$\dot{\boldsymbol{\rho}} = \mathcal{L}(\boldsymbol{\rho}) = -i[\boldsymbol{H}, \boldsymbol{\rho}] + \sum_{a>0} \left(\boldsymbol{L}_a \boldsymbol{\rho} \boldsymbol{L}_a^{\dagger} - \frac{1}{2} \boldsymbol{L}_a^{\dagger} \boldsymbol{L}_a \boldsymbol{\rho} - \frac{1}{2} \boldsymbol{\rho} \boldsymbol{L}_a^{\dagger} \boldsymbol{L}_a \right). \quad (3.133)$$

This is the general Markovian evolution law for quantum states, assuming time evolution is a trace-preserving completely positive linear map. The first term in $\mathcal{L}(\boldsymbol{\rho})$ is the familiar Hermitian Schrödinger term generating unitary evolution. The other terms describe the possible transitions that the system may undergo due to interactions with the environment. The operators \boldsymbol{L}_a are called *Lindblad operators* or *quantum jump operators*. Each $\boldsymbol{L}_a \boldsymbol{\rho} \boldsymbol{L}_a^{\dagger}$ term induces one of the possible quantum jumps, while the terms $-1/2\boldsymbol{L}_a^{\dagger}\boldsymbol{L}_a\boldsymbol{\rho} - 1/2\boldsymbol{\rho} \boldsymbol{L}_a^{\dagger}\boldsymbol{L}_a$ are needed to normalize properly the case in which no jumps occur.

As for any nonunitary quantum channel, we have the freedom to redefine the Kraus operators in the operator-sum representation of \mathcal{E}_{dt} , replacing $\{M_a\}$ by operators $\{N_\mu\}$ which differ by a unitary change of basis. In particular, invoking this freedom for the jump operators (while leaving M_0 untouched), we may replace $\{L_a\}$ by $\{L'_\mu\}$ where

$$\boldsymbol{L}'_{\mu} = \sum_{a} V_{\mu a} \boldsymbol{L}_{a} \tag{3.134}$$

and $V_{\mu a}$ is a unitary matrix. We say that these two ways of choosing the jump operators are two different *unravelings* of the same Markovian dynamics.

The master equation describes what happens when the system interacts with an unobserved environment, but we may also consider what happens if the environment is continuously monitored. In that case each quantum jump is detected; we update the quantum state of the system whenever a jump occurs, and an initial pure state remains pure at all later times. Specifically, a jump of type a occurs during the interval (t, t + dt) with probability

$$\operatorname{Prob}(a) = dt \langle \psi(t) | \boldsymbol{L}_{a}^{\dagger} \boldsymbol{L}_{a} | \psi(t) \rangle, \qquad (3.135)$$

and when a type-a jump is detected the updated state is

$$|\psi(t+dt)\rangle = \frac{\boldsymbol{L}_a|\psi(t)\rangle}{\|\boldsymbol{L}_a|\psi(t)\rangle\|},\tag{3.136}$$

while when no jump occurs the state evolves as

$$|\psi(t+dt)\rangle = \frac{\boldsymbol{M}_0|\psi(t)\rangle}{\|\boldsymbol{M}_0|\psi(t)\rangle\|}.$$
(3.137)

This stochastic Schrödinger evolution can be numerically simulated; each simulated quantum trajectory is different, but averaging over a sample of many such trajectories reproduces the evolution of the density operator as described by the master equation. Simulating the stochastic Schrödinger equation may have advantages over simulating the master equation, since it is less costly to follow the evolution of a *d*-dimensional state vector than a $d \times d$ density matrix.

3.5.3 Damped harmonic oscillator

As an example to illustrate the master equation, consider the case of a harmonic oscillator coupled to the electromagnetic field via

$$\boldsymbol{H}' = \sum_{k} g_{k} (\boldsymbol{a} \boldsymbol{b}_{k}^{\dagger} + \boldsymbol{a}^{\dagger} \boldsymbol{b}_{k}), \qquad (3.138)$$

where \boldsymbol{a} is the annihilation operator of the oscillator, $\boldsymbol{b}_k^{\dagger}$ creates a photon in mode k, and g_k is a coupling constant. Let's also suppose that the environment is at zero temperature; then the excitation level of the oscillator can cascade down by successive emission of photons, but no absorption of photons will occur. If each photon, once emitted, never interacts again with the oscillator, the evolution is Markovian, and there is only one Lindblad jump operator:

$$\boldsymbol{L} = \sqrt{\Gamma} \boldsymbol{a}. \tag{3.139}$$

Here Γ is the rate for the oscillator to decay from the first excited (n = 1)state to the ground (n = 0) state, which can be computed as $\Gamma = \sum_i \Gamma_i$, where Γ_i is the rate for emission into mode *i*. The rate for the decay from level *n* to n-1 is $n\Gamma$. (The *n*th level of excitation of the oscillator may be interpreted as a state of *n* noninteracting particles; the rate is $n\Gamma$ because any one of the *n* particles can decay.)

The master equation in the Lindblad form becomes

$$\dot{\boldsymbol{\rho}} = -i[\boldsymbol{H}_0, \boldsymbol{\rho}] + \Gamma(\boldsymbol{a}\boldsymbol{\rho}\boldsymbol{a}^{\dagger} - \frac{1}{2}\boldsymbol{a}^{\dagger}\boldsymbol{a}\boldsymbol{\rho} - \frac{1}{2}\boldsymbol{\rho}\boldsymbol{a}^{\dagger}\boldsymbol{a}) \qquad (3.140)$$

where $H_0 = \omega a^{\dagger} a$ is the Hamiltonian of the oscillator. The jump term describes the *damping* of the oscillator due to photon emission. To study the effect of the jumps, it is convenient to adopt the *interaction picture*; we define interaction picture operators ρ_I and a_I by

$$\boldsymbol{\rho}(t) = e^{-i\boldsymbol{H}_0 t} \boldsymbol{\rho}_I(t) e^{i\boldsymbol{H}_0 t},$$

$$\boldsymbol{a} = e^{-i\boldsymbol{H}_0 t} \boldsymbol{a}_I(t) e^{i\boldsymbol{H}_0 t},$$
(3.141)

so that

$$\dot{\boldsymbol{\rho}}_{I} = \Gamma(\boldsymbol{a}_{I}\boldsymbol{\rho}_{I}\boldsymbol{a}_{I}^{\dagger} - \frac{1}{2}\boldsymbol{a}_{I}^{\dagger}\boldsymbol{a}_{I}\boldsymbol{\rho}_{I} - \frac{1}{2}\boldsymbol{\rho}_{I}\boldsymbol{a}_{I}^{\dagger}\boldsymbol{a}_{I})$$
$$= \Gamma(\boldsymbol{a}\boldsymbol{\rho}_{I}\boldsymbol{a}^{\dagger} - \frac{1}{2}\boldsymbol{a}^{\dagger}\boldsymbol{a}\boldsymbol{\rho}_{I} - \frac{1}{2}\boldsymbol{\rho}_{I}\boldsymbol{a}^{\dagger}\boldsymbol{a}) \qquad (3.142)$$

where we use $a_I(t) = ae^{-i\omega t}$ to replace a_I by a in the second line. By observing the evolution in a "rotating frame," we have frozen the unperturbed motion of the oscillator, isolating the effect of the damping.

The variable $\tilde{a}(t) = e^{-iH_0t}ae^{+iH_0t} = e^{i\omega t}a$ remains constant in the absence of damping. Including damping, its expectation value

$$\langle \tilde{\boldsymbol{a}}(t) \rangle = \operatorname{tr}\left(\tilde{\boldsymbol{a}}(t)\boldsymbol{\rho}(t)\right) = \operatorname{tr}\left(\boldsymbol{a}\boldsymbol{\rho}_{I}(t)\right)$$
 (3.143)

evolves according to

$$\frac{d}{dt}\langle \tilde{\boldsymbol{a}}(t) \rangle = \operatorname{tr}\left(\boldsymbol{a}\dot{\boldsymbol{\rho}}_{I}\right) \tag{3.144}$$

and from eq. (3.142) we have

$$\frac{d}{dt}\langle \tilde{\boldsymbol{a}}(t)\rangle = \operatorname{tr}\left(\boldsymbol{a}\dot{\boldsymbol{\rho}}_{I}\right) = \Gamma \operatorname{tr}\left(\boldsymbol{a}^{2}\boldsymbol{\rho}_{I}\boldsymbol{a}^{\dagger} - \frac{1}{2}\boldsymbol{a}\boldsymbol{a}^{\dagger}\boldsymbol{a}\boldsymbol{\rho}_{I} - \frac{1}{2}\boldsymbol{a}\boldsymbol{\rho}_{I}\boldsymbol{a}^{\dagger}\boldsymbol{a}\right)$$
$$= \Gamma \operatorname{tr}\left(\frac{1}{2}[\boldsymbol{a}^{\dagger},\boldsymbol{a}]\boldsymbol{a}\boldsymbol{\rho}_{I}\right) = -\frac{\Gamma}{2}\operatorname{tr}(\boldsymbol{a}\boldsymbol{\rho}_{I}) = -\frac{\Gamma}{2}\langle \tilde{\boldsymbol{a}}(t)\rangle. \quad (3.145)$$

Integrating this equation, we obtain

$$\langle \tilde{\boldsymbol{a}}(t) \rangle = e^{-\Gamma t/2} \langle \tilde{\boldsymbol{a}}(0) \rangle.$$
 (3.146)

Similarly, the occupation number of the oscillator $n\equiv a^{\dagger}a=\tilde{a}^{\dagger}\tilde{a}$ decays according to

$$\frac{d}{dt} \langle \boldsymbol{n} \rangle = \frac{d}{dt} \langle \tilde{\boldsymbol{a}}^{\dagger} \tilde{\boldsymbol{a}} \rangle = \operatorname{tr}(\boldsymbol{a}^{\dagger} \boldsymbol{a} \dot{\boldsymbol{\rho}}_{I})$$

$$= \Gamma \operatorname{tr} \left(\boldsymbol{a}^{\dagger} \boldsymbol{a} \boldsymbol{a} \boldsymbol{\rho}_{I} \boldsymbol{a}^{\dagger} - \frac{1}{2} \boldsymbol{a}^{\dagger} \boldsymbol{a} \boldsymbol{a}^{\dagger} \boldsymbol{a} \boldsymbol{\rho}_{I} - \frac{1}{2} \boldsymbol{a}^{\dagger} \boldsymbol{a} \boldsymbol{\rho}_{I} \boldsymbol{a}^{\dagger} \boldsymbol{a} \right)$$

$$= \Gamma \operatorname{tr} \left(\boldsymbol{a}^{\dagger} [\boldsymbol{a}^{\dagger}, \boldsymbol{a}] \boldsymbol{a} \boldsymbol{\rho}_{I} \right) = -\Gamma \operatorname{tr} \left(\boldsymbol{a}^{\dagger} \boldsymbol{a} \boldsymbol{\rho}_{I} \right) = -\Gamma \langle \boldsymbol{n} \rangle, \qquad (3.147)$$

which integrates to

$$\langle \boldsymbol{n}(t) \rangle = e^{-\Gamma t} \langle \boldsymbol{n}(0) \rangle.$$
 (3.148)

Thus Γ is indeed the damping rate of the oscillator. (If we interpret the *n*th excitation state of the oscillator as a state of *n* noninteracting particles, each with a decay probability Γ per unit time, then eq. (3.148) is just the exponential law satisfied by the population of decaying particles.)

3 Foundations II: Measurement and Evolution

More interesting is what the master equation tells us about decoherence. In our amplitude damping model, it is the annihilation operator \boldsymbol{a} and its adjoint that appear in the coupling \boldsymbol{H}' of oscillator to environment, so we can anticipate that the oscillator's state will decohere in the basis of \boldsymbol{a} eigenstates. The *coherent state*

$$|\alpha\rangle = e^{-|\alpha|^2/2} e^{\alpha a^{\dagger}} |0\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle, \qquad (3.149)$$

is the normalized eigenstate of \boldsymbol{a} with complex eigenvalue α . The operator \boldsymbol{a} is not Hermitian, and two coherent states with distinct eigenvalues α and β are not orthogonal; rather

$$\begin{aligned} |\langle \alpha | \beta \rangle|^2 &= e^{-|\alpha|^2} e^{-|\beta|^2} e^{2Re(\alpha^*\beta)} \\ &= \exp(-|\alpha - \beta|^2), \end{aligned}$$
(3.150)

so the overlap is very small when $|\alpha - \beta|$ is large.

The solution to the master equation eq.(3.142) is worked out in Exercise 3.9, where we find that an initial coherent state remains coherent, but with a decaying amplitude; after time t the state $|\alpha\rangle$ evolves as

$$|\alpha\rangle \mapsto |\alpha e^{-\Gamma t}\rangle \tag{3.151}$$

(in the rotating frame). We may also consider what happens when the initial state is a superposition of coherent states (a "cat state")

$$|\psi\rangle = N_{\alpha,\beta}(|\alpha\rangle + |\beta\rangle), \qquad (3.152)$$

(where $N_{\alpha,\beta}$ is a normalization factor), or

$$\boldsymbol{\rho} = N_{\alpha,\beta}^2 \left(|\alpha\rangle\langle\alpha| + |\alpha\rangle\langle\beta| + |\beta\rangle\langle\alpha| + |\beta\rangle\langle\beta| \right).$$
(3.153)

The off-diagonal terms in this density operator evolve as

$$|\alpha\rangle\langle\beta|\mapsto e^{i\phi(\alpha,\beta)}e^{-\Gamma t|\alpha-\beta|^2/2}|\alpha e^{-\Gamma t/2}\rangle\langle\beta e^{-\Gamma t/2}|,\qquad(3.154)$$

where $e^{i\phi(\alpha,\beta)}$ is a phase factor. Thus the off-diagonal terms decay exponentially with time, at a rate

$$\Gamma_{\text{decohere}} = \frac{1}{2} \Gamma |\alpha - \beta|^2 \qquad (3.155)$$

proportional to the distance squared $|\alpha - \beta|^2$, and hence much larger than the damping rate for $|\alpha - \beta|^2 \gg 1$. This behavior is easy to interpret. The expectation value of the occupation number \boldsymbol{n} in a coherent state is $\langle \alpha | \boldsymbol{a}^{\dagger} \boldsymbol{a} | \alpha \rangle = |\alpha|^2$. Therefore, if α, β have comparable moduli but significantly different phases (as for a superposition of minimum uncertainty wave packets centered at x and -x), the decoherence rate is of the order of the rate for emission of a *single* photon. This rate is very large compared to the rate for a significant fraction of the oscillator energy to be dissipated.

We can also consider an oscillator coupled to an environment with a nonzero temperature. Again, the decoherence rate is roughly the rate for a single photon to be emitted or absorbed, but the rate may be much faster than at zero temperature. Because the photon modes with frequency comparable to the oscillator frequency ω have a thermal occupation number

$$n_{\gamma} \approx \frac{T}{\hbar\omega},$$
 (3.156)

(for $T \gg \hbar \omega$), the interaction rate is further enhanced by the factor n_{γ} . We have then

$$\frac{\Gamma_{\rm dec}}{\Gamma_{\rm damp}} \sim n_{\rm osc} n_{\gamma} \sim \frac{E}{\hbar\omega} \frac{T}{\hbar\omega} \sim \frac{m\omega^2 x^2}{\hbar\omega} \frac{T}{\hbar\omega} \sim x^2 \frac{mT}{\hbar^2} \sim \frac{x^2}{\lambda_T^2}, \qquad (3.157)$$

where x is the amplitude of oscillation and λ_T is the thermal de Broglie wavelength of the oscillating object. For macroscopic objects, decoherence is really *fast*.

3.6 Non-Markovian noise

3.6.1 Gaussian phase noise

The master equation describes the evolution of a quantum system subject to Markovian noise, but in some experimental systems the Markovian approximation is not very accurate. In this section we will discuss some of the features of decoherence for a system subjected to non-Markovian noise.

As a simple example, consider a single qubit with energy eigenstates $|0\rangle$ and $|1\rangle$, where the energy splitting between the two states fluctuates. For example, the qubit could be a spin- $\frac{1}{2}$ particle in a magnetic field pointing along the z-axis, where the magnetic field is not perfectly controlled in the laboratory. The Hamiltonian for this system is

$$\boldsymbol{H} = -\frac{1}{2}\omega_{01}\boldsymbol{\sigma}_3 - \frac{1}{2}f(t)\boldsymbol{\sigma}_3 , \qquad (3.158)$$

where f(t) is the fluctuating component of the magnetic field. This is a model of *classical noise*, arising not because the system interacts with an unobserved environment, but rather because a term in the system's Hamiltonian fluctuates.

The function f is treated stochastically; that is, we consider an ensemble of possible functions $\{f\}$, each with an assigned probability weight $\mu(f)$. We imagine that the actual f(t) in each run of the experiment is selected by sampling from this distribution, and predict the observed behavior of the system by averaging over the distribution $\mu(f)$. The model is particularly simple because the unperturbed Hamiltonian $H_0 = \frac{1}{2}\omega_{01}\sigma_3$ commutes with the noise term $H_f = \frac{1}{2}f(t)\sigma_3$, and in fact we can transform H_0 away by going to the interaction picture.

The fluctuations induce dephasing of the qubit in the energy eigenstate basis. To analyze the dephasing, we will make a further simplifying assumption, that the noise is *Gaussian*. Whether the noise is classical of quantum, this Gaussian approximation often applies in laboratory situations where the system is weakly coupled to many different fluctuating variables in the environment. We denote averaging over the distribution $\mu(f)$ by [·], and assume the distribution to be stationary with mean zero; that is [f(t)] = 0, and [f(t)f(t')] = K(t - t') is a function only of the difference t - t', which is called the *covariance* of the distribution. The Gaussian distribution can be characterized by its generating functional Z[J], which can be expressed in terms of the covariance as

$$Z[J] \equiv \left[e^{\int dt J(t)f(t)}\right]_f = \exp\left(\frac{1}{2}\int dt dt' J(t)K(t-t')J(t')\right). \quad (3.159)$$

An initial density operator $\rho(0)$ evolves in time T to

$$\boldsymbol{\rho}(T) = \left[\exp\left(i\int_{o}^{T}\frac{1}{2}f(t)\boldsymbol{\sigma}_{3}\right)\boldsymbol{\rho}(0)\exp\left(-i\int_{o}^{T}\frac{1}{2}f(t)\boldsymbol{\sigma}_{3}\right) \right]. \quad (3.160)$$

The energy eigenstates $|0\rangle\langle 0|$ or $|1\rangle\langle 1|$ are not affected, but using eq.(3.159) we see that the coefficients of the off-diagonal entries $|0\rangle\langle 1|$ and $|1\rangle\langle 0|$ decay by the factor

$$\exp\left(-\frac{1}{2}\int_0^T dt \int_0^T dt' K(t-t')\right)$$
$$= \exp\left(-\frac{1}{2}\int_0^T dt \int_0^T dt' \int_{-\infty}^\infty \frac{d\omega}{2\pi} e^{-i\omega(t-t')} \tilde{K}(\omega)\right); \qquad (3.161)$$

here we have introduced the Fourier transform $\tilde{K}(\omega)$ of the covariance K(t), which is said to be the spectral density or power spectrum of the

noise. Doing the t and t' integrals we obtain

$$\exp\left(-\frac{1}{2}\int_{-\infty}^{\infty}\frac{d\omega}{2\pi}\tilde{K}(\omega)W_{T}(\omega)\right)$$
(3.162)

where $W_T(\omega)$ is the smooth window function

$$W_T(\omega) = \left| \int_0^T dt \ e^{-i\omega t} \right|^2 = \frac{4}{\omega^2} \sin^2(\omega T/2), \qquad (3.163)$$

which has most of its support on the interval $[0, 2\pi/T]$.

Assuming that $\tilde{K}(\omega = 0)$ is finite, we expect that for T sufficiently large, $\tilde{K}(\omega)$ can be regarded as approximately constant in the region where $W_T(\omega)$ is supported. Using $\int_{-\infty}^{\infty} dx \frac{\sin^2 x}{x^2} = \pi$, we then obtain the decay factor $e^{-\Gamma_2 T}$, where the dephasing rate Γ_2 is

$$\Gamma_2 = \tilde{K}(\omega = 0). \tag{3.164}$$

(Here we've assumed that $\tilde{K}(\omega)$ is continuous at $\omega = 0$ — otherwise we should average its limiting values as ω approaches zero from positive and negative values.) If the spectral density is flat ("white noise"), this formula for Γ_2 applies at any time T, but in general, the time scale for which dephasing can be described by a rate Γ_2 depends on the shape of the noise's spectral density. In effect, an experimentalist who measures the dephasing time $T_2 = \Gamma_2^{-1}$ of a qubit is probing the noise power at low frequency.

Crudely speaking, we expect $\tilde{K}(\omega)$ to be roughly constant in the interval $[0, \omega_c]$, where $\omega_c = 2\pi/\tau_c$, and τ_c is a characteristic "autocorrelation" or "memory" time of the noise. That is, τ_c is chosen so that the correlation function K(t - t') is small for $|t - t'| \gg \tau_c$. Thus we see that in order to speak of a "dephasing rate" Γ_2 (and a corresponding dephasing time $T_2 = \Gamma_2^{-1}$) we must consider evolution that has been "coarse-grained" in time. For the purpose of describing evolution over a time period $T \gg \tau_c$, the non-Markovian noise model can be replaced by a corresponding effective Markovian model in which the memory of the fluctuations can be neglected, as in our analysis of dephasing in §3.4.2. But for $T \ll \tau_c$ such a description is not applicable.

3.6.2 Spin echo

Strategies for mitigating the damaging effects of the noise become possible when the noise autocorrelation time τ_c is long compared to the time scale over which the experimentalist can manipulate the system. For example, when observing the dephasing of a spin evolving for time T, we may apply a fast pulse that flips the spin about the x-axis at time T/2. Then the effects of low-frequency phase noise during the second half of the evolution will tend to compensate for the effects of the phase noise during the first half. This trick is called the *spin echo* phenomenon.

If we use this trick, the damping factor applied to $|0\rangle\langle 1|$ is again given by

$$\exp\left(-\frac{1}{2}\int_{-\infty}^{\infty}\frac{dw}{2\pi}\tilde{K}(\omega)W_{T}(\omega)\right)$$
(3.165)

but with a modified window function

$$W_T(\omega) = \left| \int_0^T dt J(t) e^{i\omega t} \right|^2, \qquad (3.166)$$

where J(t) is a modulating function that expresses the effect of the spin echo pulse sequence. For example, if we flip the spin at time T/2, then J(t) is +1 in the interval [0, T/2] and -1 in the interval [T/2, T]; therefore

$$W_T(w) = \frac{1}{\omega^2} \left| 1 - 2e^{i\omega T/2} + e^{i\omega T} \right|^2$$

= $\frac{1}{\omega^2} \left| \frac{1 - e^{i\omega T/2}}{1 + e^{i\omega T/2}} \left(1 - e^{i\omega T} \right) \right|^2$
= $\tan^2(\omega T/4) \cdot \frac{4}{\omega^2} \sin^2(\omega t/2).$ (3.167)

In effect, the spin echo modifies $\tilde{K}(\omega)$ by the multiplicative factor $\tan^2(\omega T/4)$, which suppresses the low frequency noise.

The suppression can be improved further by using more pulses. In practice, pulses have bounded strength and nonzero duration, which places limitations on the effectiveness of this strategy.

3.6.3 Qubits as noise spectrometers

Now let's consider a different model of classical noise, in which the fluctuating term does not commute with the unperturbed Hamiltonian:

$$H = -\frac{1}{2}\omega_{01}\boldsymbol{\sigma}_3 + f(t)\boldsymbol{\sigma}_1.$$
(3.168)

In this model the fluctuating field can induce transitions among the energy eigenstates, at a rate that can be computed using lowest-order interactionpicture perturbation theory if the noise is weak. The probability that a qubit prepared in the state $|1\rangle$ at time 0 is observed in the state $|0\rangle$ at time T, averaged over the fluctuating classical field, is

$$\operatorname{Prob}(1 \to 0) = \left[\left| -i \int_{0}^{T} dt \ f(t) e^{-i\omega_{01}t} \langle 0 | \boldsymbol{\sigma}_{1} | 1 \rangle \right|^{2} \right]$$
$$= \int_{0}^{T} dt \int_{0}^{T} dt' \ e^{-i\omega_{01}(t-t')} \left[f(t) f(t') \right]$$
$$= \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \tilde{K}(\omega) W_{T}(\omega - \omega_{01}) . \qquad (3.169)$$

This expression is similar to the formula eq.(3.162) for the off-diagonal term in the density operator obtained in the dephasing model, except that now the center of the window function has been shifted to the frequency ω_{01} of the transition.

As before, if we consider the observation time T to be large compared to the autocorrelation time τ_c of the noise, then the support of the window function is narrow, and $\tilde{K}(\omega)$ is approximately constant in the window. Thus, after a suitable coarse-graining of the time evolution, we may identify a rate for the decay of the qubit

$$\Gamma_{\downarrow} = \tilde{K}(\omega = \omega_{01}). \tag{3.170}$$

Similarly, for the transition from ground state to excited state, we find

$$\Gamma_{\uparrow} = \tilde{K}(\omega = -\omega_{01}). \tag{3.171}$$

Thus negative frequency noise transfers energy from the noise reservoir to the qubit, exciting the qubit, while positive frequency noise transfers energy from qubit to the noise reservoir, returning the excited qubit to the ground state. (Dephasing of a qubit, in contrast, involves a negligible exchange of energy and therefore is controlled by low frequency noise.) We conclude that an experimentalist capable of varying the energy splitting ω_{01} and measuring the qubit's transition rate can determine how the noise power depends on the frequency.

For the case we have considered in which the noise source is classical, f(t) and f(t') are real commuting variables; therefore K(t) is an even function of t and correspondingly $\tilde{K}(\omega)$ is an even function of ω . Classical noise is spectrally symmetric, and the rates for excitation and decay are equal.

On the other hand, noise driven by a quantized thermal "bath" can be spectrally asymmetric. When the qubit comes to thermal equilibrium with the bath, up and down transitions occur at equal rates. If p_0 denotes the probability that the qubit is in the ground state $|0\rangle$ and p_1 denotes the probability that the qubit is in the excited state $|1\rangle$, then in equilibrium

$$p_0\Gamma_{\uparrow} = p_1\Gamma_{\downarrow} \Rightarrow \frac{K(-\omega_{01})}{\tilde{K}(\omega_{01})} = \frac{p_1}{p_0} = e^{-\beta\omega_{01}}; \qquad (3.172)$$

the ratio of noise strengths at positive and negative frequencies is given (for a thermal bath) by a Boltzmann factor; this property of the noise is called the Kubo-Martin-Schwinger (KMS) condition. The noise becomes classical in the high-temperature limit $\beta \omega_{01} \ll 1$, and is in the deeply quantum regime for $\beta \omega_{01} \gg 1$.

3.6.4 Spin-boson model at nonzero temperature

To turn our model of classical dephasing noise into a quantum model, we replace the stochastic classical field f(t) by an operator acting on a quantized bath. The noise will still be Gaussian if the bath is a system of harmonic oscillators, uncoupled to one another and each coupled linearly to the dephasing qubit. The Hamiltonian for the system A and bath B is

$$\boldsymbol{H}_{A} + \boldsymbol{H}_{B} + \boldsymbol{H}_{AB} = -\frac{1}{2}\omega_{01}\boldsymbol{\sigma}_{3} + \sum_{k}\omega_{k}\boldsymbol{a}_{k}^{\dagger}\boldsymbol{a}_{k} - \frac{1}{2}\boldsymbol{\sigma}_{3}\left(\sum_{k}g_{k}\boldsymbol{a}_{k} + g_{k}^{*}\boldsymbol{a}_{k}^{\dagger}\right),$$
(3.173)

which is called the *spin-boson model*, as it describes a single $spin-\frac{1}{2}$ particle coupled to many bosonic variables. This is a model of dephasing because the coupling of the spin to the bath is diagonal in the spin's energy eigenstate basis. (Otherwise the physics of the model would be harder to analyze.) Despite its simplicity, the spin-boson model provides a reasonably realistic description of dephasing for a qubit weakly coupled to many degrees of freedom in the environment.

If there are many oscillators, the sum over k can be approximated by a frequency integral:

$$\sum_{k} |g_k|^2 \approx \int_0^\infty d\omega J(\omega), \qquad (3.174)$$

where $J(\omega)$ is said to be the spectral function of the oscillator bath. Let's assume that the bath is in thermal equilibrium at temperature β^{-1} . In principle, the coupling to the system could tweak the equilibrium distribution of the bath, but we assume that this effect is negligible, because the bath is much bigger than the system. The fluctuations of the bath are Gaussian, and the average over the ensemble of classical functions in our previous analysis can be replaced by the thermal expectation value:

$$[\boldsymbol{f}(t)\boldsymbol{f}(0)] \mapsto \langle \boldsymbol{f}(t)\boldsymbol{f}(0) \rangle_{\beta} \equiv \operatorname{tr}\left(e^{-\beta \boldsymbol{H}_{B}}\boldsymbol{f}(t)\boldsymbol{f}(0)\right), \qquad (3.175)$$

where now f(t) denotes the operator

$$\boldsymbol{f}(t) = e^{it\boldsymbol{H}_B}\boldsymbol{f}(0)e^{-it\boldsymbol{H}_B} = \sum_k \left(g_k\boldsymbol{a}_k e^{-i\omega_k t} + g_k^*\boldsymbol{a}_k^{\dagger} e^{i\omega_k t}\right). \quad (3.176)$$

We see that

$$K_{\beta}(t) \equiv \langle \boldsymbol{f}(t)\boldsymbol{f}(0)\rangle_{\beta} = \sum_{k} |g_{k}|^{2} \langle e^{-i\omega_{k}t}\boldsymbol{a}_{k}\boldsymbol{a}_{k}^{\dagger} + e^{i\omega_{k}t}\boldsymbol{a}_{k}^{\dagger}\boldsymbol{a}_{k}\rangle_{\beta}.$$
 (3.177)

From the Planck distribution, we find

$$\langle \boldsymbol{a}_{k}^{\dagger}\boldsymbol{a}_{k}\rangle_{\beta} = \frac{1}{e^{\beta\omega} - 1} = \frac{1}{2}\coth(\beta\omega_{k}/2) - \frac{1}{2},$$
$$\langle \boldsymbol{a}_{k}\boldsymbol{a}_{k}^{\dagger}\rangle_{\beta} = \langle \boldsymbol{a}_{k}^{\dagger}\boldsymbol{a}_{k} + 1\rangle_{\beta} = \frac{1}{2}\coth(\beta\omega_{k}/2) + \frac{1}{2}, \qquad (3.178)$$

and by Fourier transforming we obtain the spectral density of the noise

$$\tilde{K}_{\beta}(\omega) \equiv \int_{-\infty}^{\infty} dt \ e^{i\omega t} K_{\beta}(t)$$

$$= \sum_{k} |g_{k}|^{2} \left(2\pi \delta(\omega - \omega_{k}) \langle \boldsymbol{a}_{k} \boldsymbol{a}_{k}^{\dagger} \rangle_{\beta} + 2\pi \delta(\omega + \omega_{k}) \langle \boldsymbol{a}_{k}^{\dagger} \boldsymbol{a}_{k} \rangle_{\beta} \right),$$
(3.179)

which may be written as

$$K_{\beta}(\omega) = \pi J(\omega) \left(\coth(\beta \omega/2) + 1 \right), \quad \omega > 0,$$

$$\tilde{K}_{\beta}(\omega) = \pi J(\omega) \left(\coth(\beta \omega/2) - 1 \right), \quad \omega < 0.$$
(3.180)

Thus, as we anticipated, noise power spectrum exhibits the spectral asymmetry required by the KMS condition — the spectral density $\tilde{K}_{\beta}(-\omega)$ of the noise at negative frequency is supressed relative to the spectral density $\tilde{K}_{\beta}(\omega)$ at positive frequency by the Boltzmann factor $e^{-\beta\omega}$.

Since the window function $W_T(\omega)$ is an even function of ω , only the even part of $\tilde{K}_{\beta}(\omega)$ contributes to the attenuation of $|0\rangle\langle 1|$; the attenuation factor

$$\exp\left(-\frac{1}{2}\int_{-\infty}^{\infty}\frac{d\omega}{2\pi}\tilde{K}_{\beta}(\omega)W_{T}(\omega)\right),\qquad(3.181)$$

therefore becomes

$$\exp\left(-\int_0^\infty d\omega J(\omega)\frac{2\sin^2(\omega T/2)}{\omega^2}\coth(\beta\omega/2)\right).$$
 (3.182)

A dephasing *rate* can be identified if the spectral function $J(\omega)$ behaves suitably at low frequency; the attenuation factor is $e^{-\Gamma_2 T}$ in the limit $T \to \infty$ where

$$\Gamma_2 = \lim_{\omega \to 0} \tilde{K}_{\beta}(\omega) = \lim_{\omega \to 0} 2\pi J(\omega) / (\beta \omega), \qquad (3.183)$$

assuming that this limit exists. The noise is said to be *Ohmic* if $J(\omega) \approx A\omega$ is linear in ω at low frequency, and in that case the dephasing rate becomes $\Gamma_2 = 2\pi A\beta^{-1}$ in the limit of long time T.

3.7 Summary

POVM. If we restrict our attention to a subspace of a larger Hilbert space, then an orthogonal (Von Neumann) measurement performed on the larger space cannot in general be described as an orthogonal measurement on the subspace. Rather, it is a *generalized measurement* or POVM — the outcome *a* occurs with a probability

$$\operatorname{Prob}(a) = \operatorname{tr}\left(\boldsymbol{E}_{a}\boldsymbol{\rho}\right) , \qquad (3.184)$$

where ρ is the density matrix of the subsystem, each E_a is a positive hermitian operator, and the E_a 's satisfy

$$\sum_{a} \boldsymbol{E}_{a} = \boldsymbol{I} \ . \tag{3.185}$$

A POVM in \mathcal{H}_A can be realized as a unitary transformation on the tensor product $\mathcal{H}_A \otimes \mathcal{H}_B$, followed by an orthogonal measurement in \mathcal{H}_B .

Quantum channel. Unitary evolution on $\mathcal{H}_A \otimes \mathcal{H}_B$ will not in general appear to be unitary if we restrict our attention to \mathcal{H}_A alone. Rather, evolution in \mathcal{H}_A will be described by a *quantum channel*, (which can be inverted by another channel only if unitary). A general channel \mathcal{E} has an operator-sum representation:

$$\mathcal{E}: \boldsymbol{\rho} \to \mathcal{E}(\boldsymbol{\rho}) = \sum_{a} \boldsymbol{M}_{a} \boldsymbol{\rho} \boldsymbol{M}_{a}^{\dagger} , \qquad (3.186)$$

where

$$\sum_{a} \boldsymbol{M}_{a}^{\dagger} \boldsymbol{M}_{a} = \boldsymbol{I}.$$
 (3.187)

In fact, any reasonable (linear, trace preserving, and completely positive) mapping of density operators to density operators has such an operatorsum representation.

Decoherence. Decoherence — the decay of quantum information due to the interaction of a system with its environment — can be described by a quantum channel. If the environment frequently "scatters" off the system, and the state of the environment is not monitored, then offdiagonal terms in the density operator of the system decay rapidly in a preferred basis (typically a spatially localized basis selected by the nature of the coupling of the system to the environment). The time scale for decoherence is set by the scattering rate, which may be much larger than the damping rate for the system.

Master Equation. When the relevant dynamical time scale of an open quantum system is long compared to the time for the environment to "forget" quantum information, the evolution of the system is effectively local in time (the Markovian approximation). Much as general unitary evolution is generated by a Hamiltonian, a general Markovian superoperator is generated by a *Liouvillian* \mathcal{L} as described by the *master equation*:

$$\dot{\boldsymbol{\rho}} \equiv \mathcal{L}(\boldsymbol{\rho}) = -i[\boldsymbol{H}, \boldsymbol{\rho}] + \sum_{a} \left(\boldsymbol{L}_{a} \boldsymbol{\rho} \boldsymbol{L}_{a}^{\dagger} - \frac{1}{2} \boldsymbol{L}_{a}^{\dagger} \boldsymbol{L}_{a} \boldsymbol{\rho} - \frac{1}{2} \boldsymbol{\rho} \boldsymbol{L}_{a}^{\dagger} \boldsymbol{L}_{a} \right). \quad (3.188)$$

Here each Lindblad operator (or quantum jump operator) \mathcal{L}_a describes a "quantum jump" that could in principle be detected if we monitored the environment faithfully. By solving the master equation, we can compute the decoherence rate of an open system.

Non-Markovian noise. Non-Markovian noise can be characterized by its power spectrum, and the effects of the noise on dephasing over a long time period are determined by the behavior of the power spectrum at low frequency. Quantum noise in thermal equilibrium at temperature β^{-1} has a spectral asymmetry — the noise at negative frequency $(-\omega)$ is suppressed compared to the noise at positive frequency ω by a Boltzmann factor $e^{-\beta\omega}$ (the *KMS condition*).

Further important ideas are developed in the Exercises.

3.8 Exercises

3.1 Which state did Alice make?

Consider a game in which Alice prepares one of two possible states: either ρ_1 with a priori probability p_1 , or ρ_2 with a priori probability $p_2 = 1 - p_1$. Bob is to perform a measurement and on the basis of the outcome to guess which state Alice prepared. If Bob's guess is right, he wins; if he guesses wrong, Alice wins.

In this exercise you will find Bob's best strategy, and determine his optimal probability of error.

Let's suppose (for now) that Bob performs a POVM with two possible outcomes, corresponding to the two nonnegative Hermitian operators E_1 and $E_2 = I - E_1$. If Bob's outcome is E_1 , he guesses that Alice's state was ρ_1 , and if it is E_2 , he guesses ρ_2 . Then the probability that Bob guesses wrong is

$$p_{\text{error}} = p_1 \operatorname{tr} (\rho_1 E_2) + p_2 \operatorname{tr} (\rho_2 E_1) .$$
 (3.189)

a) Show that

$$p_{\text{error}} = p_1 + \sum_i \lambda_i \langle i | \boldsymbol{E}_1 | i \rangle , \qquad (3.190)$$

where $\{|i\rangle\}$ denotes the orthonormal basis of eigenstates of the Hermitian operator $p_2 \rho_2 - p_1 \rho_1$, and the λ_i 's are the corresponding eigenvalues.

b) Bob's best strategy is to perform the two-outcome POVM that minimizes this error probability. Find the nonnegative operator E_1 that minimizes p_{error} , and show that error probability when Bob performs this optimal two-outcome POVM is

$$(p_{\text{error}})_{\text{optimal}} = p_1 + \sum_{\text{neg}} \lambda_i$$
 (3.191)

where \sum_{neg} denotes the sum over all of the *negative* eigenvalues of $p_2 \rho_2 - p_1 \rho_1$.

c) It is convenient to express this optimal error probability in terms of the L^1 norm of the operator $p_2 \rho_2 - p_1 \rho_1$,

$$||p_2 \rho_2 - p_1 \rho_1||_1 = \text{tr} |p_2 \rho_2 - p_1 \rho_1| = \sum_{\text{pos}} \lambda_i - \sum_{\text{neg}} \lambda_i$$
, (3.192)

the difference between the sum of positive eigenvalues and the sum of negative eigenvalues. Use the property tr $(p_2 \rho_2 - p_1 \rho_1) = p_2 - p_1$ to show that

$$(p_{\text{error}})_{\text{optimal}} = \frac{1}{2} - \frac{1}{2} ||p_2 \rho_2 - p_1 \rho_1||_1 .$$
 (3.193)

Check whether the answer makes sense in the case where $\rho_1 = \rho_2$ and in the case where ρ_1 and ρ_2 have support on orthogonal subspaces.

d) Now suppose that Alice decides at random (with $p_1 = p_2 = 1/2$) to prepare one of two pure states $|\psi_1\rangle, |\psi_2\rangle$ of a single qubit, with

$$|\langle \psi_1 | \psi_2 \rangle| = \sin(2\alpha) , \quad 0 \le \alpha \le \pi/4 . \tag{3.194}$$

With a suitable choice of basis, the two states can be expressed as

$$|\psi_1\rangle = \begin{pmatrix} \cos \alpha \\ \sin \alpha \end{pmatrix}$$
, $|\psi_2\rangle = \begin{pmatrix} \sin \alpha \\ \cos \alpha \end{pmatrix}$. (3.195)

Find Bob's optimal two-outcome measurement, and compute the optimal error probability.

e) Bob wonders whether he can find a better strategy if his POVM $\{E_i\}$ has more than two possible outcomes. Let p(a|i) denote the probability that state *a* was prepared, given that the measurement outcome was *i*; it can be computed using the relations

$$p_i p(1|i) = p_1 p(i|1) = p_1 \text{ tr } \boldsymbol{\rho}_1 \boldsymbol{E}_i ,$$

$$p_i p(2|i) = p_2 p(i|2) = p_2 \text{ tr } \boldsymbol{\rho}_2 \boldsymbol{E}_i ; \qquad (3.196)$$

here p(i|a) denotes the probability that Bob finds measurement outcome *i* if Alice prepared the state ρ_a , and p_i denotes the probability that Bob finds measurement outcome *i*, averaged over Alice's choice of state. For each outcome *i*, Bob will make his decision according to which of the two quantities

$$p(1|i)$$
, $p(2|i)$ (3.197)

is the larger; the probability that he makes a mistake is the smaller of these two quantities. This probability of error, given that Bob obtains outcome i, can be written as

$$p_{\text{error}}(i) = \min\left(p(1|i), p(2|i)\right) = \frac{1}{2} - \frac{1}{2}\left|p(2|i) - p(1|i)\right| .$$
(3.198)

Show that the probability of error, averaged over the measurement outcomes, is

$$p_{\text{error}} = \sum_{i} p_{i} \ p_{\text{error}}(i) = \frac{1}{2} - \frac{1}{2} \sum_{i} |\operatorname{tr} \left(p_{2} \boldsymbol{\rho}_{2} - p_{1} \boldsymbol{\rho}_{1} \right) \boldsymbol{E}_{i}| \quad .$$
(3.199)

f) By expanding in terms of the basis of eigenstates of $p_2 \rho_2 - p_1 \rho_1$, show that

$$p_{\text{error}} \ge \frac{1}{2} - \frac{1}{2} ||p_2 \rho_2 - p_1 \rho_1||_1 .$$
 (3.200)

(**Hint**: Use the completeness property $\sum_i E_i = I$.) Since we have already shown that this bound can be saturated with a two-outcome POVM, the POVM with many outcomes is no better.

3.2 Eavesdropping and disturbance

Alice wants to send a message to Bob. Alice is equipped to prepare either one of the two states $|u\rangle$ or $|v\rangle$. These two states, in a suitable basis, can be expressed as

$$|u\rangle = \begin{pmatrix} \cos \alpha \\ \sin \alpha \end{pmatrix}$$
, $|v\rangle = \begin{pmatrix} \sin \alpha \\ \cos \alpha \end{pmatrix}$, (3.201)

where $0 < \alpha < \pi/4$. Suppose that Alice decides at random to send either $|u\rangle$ or $|v\rangle$ to Bob, and Bob is to make a measurement to determine what she sent. Since the two states are not orthogonal, Bob cannot distinguish the states perfectly.

a) Bob realizes that he can't expect to be able to identify Alice's qubit every time, so he settles for a procedure that is successful only some of the time. He performs a POVM with three possible outcomes: $\neg u$, $\neg v$, or DON'T KNOW. If he obtains the result $\neg u$, he is certain that $|v\rangle$ was sent, and if he obtains $\neg v$, he is certain that $|u\rangle$ was sent. If the result is DON'T KNOW, then his measurement is inconclusive. This POVM is defined by the operators

$$\boldsymbol{E}_{\neg u} = A(\boldsymbol{I} - |u\rangle\langle u|) , \quad \boldsymbol{E}_{\neg v} = A(\boldsymbol{I} - |v\rangle\langle v|) ,$$
$$\boldsymbol{E}_{\mathrm{DK}} = (1 - 2A)\boldsymbol{I} + A(|u\rangle\langle u| + |v\rangle\langle v|) , \quad (3.202)$$

where A is a positive real number. How should Bob choose A to minimize the probability of the outcome DK, and what is this minimal DK probability (assuming that Alice chooses from $\{|u\rangle, |v\rangle\}$ equiprobably)? **Hint:** If A is too large, $E_{\rm DK}$ will have negative eigenvalues, and Eq.(3.202) will not be a POVM.

b) Eve also wants to know what Alice is sending to Bob. Hoping that Alice and Bob won't notice, she intercepts each qubit that Alice sends, by performing an orthogonal measurement that projects onto the basis $\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \}$. If she obtains the outcome $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, she sends the state $|u\rangle$ on to Bob, and if she obtains the outcome $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$, she sends $|v\rangle$ on to Bob. Therefore each time Bob's POVM has a conclusive outcome, Eve knows with certainty what that outcome is. But Eve's tampering causes detectable errors; sometimes Bob obtains a "conclusive" outcome that actually differs from what Alice sent. What is the probability of such an error, when Bob's outcome is conclusive?

3.3 Minimal disturbance

Consider a game in which Alice decides at random (equiprobably) whether to prepare one of two possible pure states of a single qubit, either

$$|\psi\rangle = \begin{pmatrix} \cos \alpha \\ \sin \alpha \end{pmatrix}$$
, or $|\tilde{\psi}\rangle = \begin{pmatrix} \sin \alpha \\ \cos \alpha \end{pmatrix}$, (3.203)

and sends the state to Bob. By performing an orthogonal measurement in the basis $\{|0\rangle, |1\rangle\}$, Bob can identify the state with minimal

error probability

$$(p_{\text{error}})_{\text{optimal}} = \sin^2 \alpha = \frac{1}{2}(1 - \sin \theta) , \qquad (3.204)$$

where we have defined θ by

$$\langle \psi | \tilde{\psi} \rangle \equiv \cos \theta = \sin(2\alpha) .$$
 (3.205)

But now let's suppose that Eve wants to *eavesdrop* on the state as it travels from Alice to Bob. Like Bob, she wishes to extract optimal information that distinguishes $|\psi\rangle$ from $|\tilde{\psi}\rangle$, and she also wants to minimize the disturbance introduced by her eavesdropping, so that Alice and Bob are not likely to notice that anything is amiss.

Eve realizes that the optimal POVM can be achieved by measurement operators

$$\boldsymbol{M}_0 = |\phi_0\rangle\langle 0| , \qquad \boldsymbol{M}_1 = |\phi_1\rangle\langle 1| , \qquad (3.206)$$

where the vectors $|\phi_0\rangle$, and $|\phi_1\rangle$ are arbitrary. If Eve performs this measurement, then Bob receives the state

$$\boldsymbol{\rho}' = \cos^2 \alpha |\phi_0\rangle \langle \phi_0| + \sin^2 \alpha |\phi_1\rangle \langle \phi_1| , \qquad (3.207)$$

if Alice sent $|\psi\rangle$, and the state

$$\tilde{\rho}' = \sin^2 \alpha |\phi_0\rangle \langle \phi_0| + \cos^2 \alpha |\phi_1\rangle \langle \phi_1| , \qquad (3.208)$$

if Alice sent $|\tilde{\psi}\rangle$.

Eve wants the average fidelity of the state received by Bob to be as large as possible. The quantity that she wants to minimize, which we will call the "disturbance" D, measures how close this average fidelity is to one:

$$D = 1 - \frac{1}{2}(F + \tilde{F}) , \qquad (3.209)$$

where

$$F = \langle \psi | \boldsymbol{\rho}' | \psi \rangle , \qquad \tilde{F} = \langle \tilde{\psi} | \tilde{\boldsymbol{\rho}}' | \tilde{\psi} \rangle . \qquad (3.210)$$

The purpose of this exercise is to examine how effectively Eve can reduce the disturbance by choosing her measurement operators properly.

a) Show that $F + \tilde{F}$ can be expressed as

$$F + \tilde{F} = \langle \phi_0 | \boldsymbol{A} | \phi_0 \rangle + \langle \phi_1 | \boldsymbol{B} | \phi_1 \rangle , \qquad (3.211)$$

where

$$A = \begin{pmatrix} 1 - 2\cos^2 \alpha \sin^2 \alpha & \cos \alpha \sin \alpha \\ \cos \alpha \sin \alpha & 2\cos^2 \alpha \sin^2 \alpha \end{pmatrix},$$

$$B = \begin{pmatrix} 2\cos^2 \alpha \sin^2 \alpha & \cos \alpha \sin \alpha \\ \cos \alpha \sin \alpha & 1 - 2\cos^2 \alpha \sin^2 \alpha \end{pmatrix}. (3.212)$$

b) Show that if $|\phi_0\rangle$ and $|\phi_1\rangle$ are chosen optimally, the minimal disturbance that can be attained is

$$D_{\min}(\cos^2 \theta) = \frac{1}{2}(1 - \sqrt{1 - \cos^2 \theta + \cos^4 \theta}) . \qquad (3.213)$$

[**Hint**: We can choose $|\phi_0\rangle$ and $|\phi_1\rangle$ to maximize the two terms in eq. (3.211) independently. The maximal value is the maximal eigenvalue of \boldsymbol{A} , which since the eigenvalues sum to 1, can be expressed as $\lambda_{\max} = \frac{1}{2} \left(1 + \sqrt{1 - 4 \cdot \det \boldsymbol{A}} \right)$.] Of course, Eve could reduce the disturbance further were she willing to settle for a less than optimal probability of guessing Alice's state correctly.

c) Sketch a plot of the function $D_{\min}(\cos^2 \theta)$. Interpret its value for $\cos \theta = 1$ and $\cos \theta = 0$. For what value of θ is D_{\min} largest? Find D_{\min} and $(p_{\text{error}})_{\text{optimal}}$ for this value of θ .

3.4 The price of quantum state encryption

Alice and Bob are working on a top secret project. I can't tell you exactly what the project is, but I will reveal that Alice and Bob are connected by a perfect quantum channel, and that Alice uses the channel to send quantum states to Bob. Alice and Bob are worried that an eavesdropper (Eve) might intercept some of Alice's transmissions. By measuring the intercepted quantum state, Eve could learn something about what Alice is sending, and perhaps make an inference about the nature of the project.

To protect against eavesdropping, Alice and Bob decide to *encrypt* the quantum states that Alice sends. They share a *secret key*, a string of random bits about which the eavesdropper knows nothing. By consuming 2n bits of secret key, Alice can encrypt, and Bob can decrypt, an arbitrary *n*-qubit state ρ . For every possible state ρ , the encrypted state looks exactly the same to Eve, so she cannot find out anything about ρ .

Here is how the encryption procedure works: We may express the 2n bit string x as $x = x_0 x_1 x_2 \cdots x_{n-1}$, where $x_i \in \{0, 1, 2, 3\}$, and denote a tensor product of n Pauli operators as

$$\boldsymbol{\sigma}(x) = \boldsymbol{\sigma}_{x_0} \otimes \boldsymbol{\sigma}_{x_2} \otimes \cdots \otimes \boldsymbol{\sigma}_{x_{n-1}}$$
(3.214)

(where $\sigma_0 = I$). Note that $\sigma(x)^2 = I^{\otimes n}$, the identity operator acting on *n* qubits. To encrypt, Alice consults her random string to determine *x* (which is chosen uniformly at random), and applies $\sigma(x)$ to the state, obtaining $\sigma(x)\rho\sigma(x)$. To decrypt, Bob, consults the same string and applies $\sigma(x)$ to recover ρ .

a) Since Eve does not know the secret key, to her the encrypted state is indistinguishable from

$$\mathcal{E}(\boldsymbol{\rho}) = \frac{1}{2^{2n}} \sum_{x} \boldsymbol{\sigma}(x) \boldsymbol{\rho} \boldsymbol{\sigma}(x) . \qquad (3.215)$$

Show that, for any *n*-qubit state ρ

$$\mathcal{E}(\boldsymbol{\rho}) = \frac{1}{2^n} \boldsymbol{I}^{\otimes n} . \tag{3.216}$$

Since $\mathcal{E}(\boldsymbol{\rho})$ is independent of $\boldsymbol{\rho}$, no information about $\boldsymbol{\rho}$ is accessible to Eve.

b) Alice wonders if it is possible to encrypt the state using a shorter key. Alice and Bob could use their shared randomness to sample an arbitrary probability distribution. That is, they could agree on a set of N unitary matrices $\{\boldsymbol{U}_a, a = 1, 2, 3, \ldots, N\}$, and Alice could encrypt by applying \boldsymbol{U}_a with probability p_a . Then Bob could decrypt by applying \boldsymbol{U}_a^{-1} . To Eve, the encrypted state would then appear to be

$$\mathcal{E}'(\boldsymbol{\rho}) = \sum_{a} p_a \boldsymbol{U}_a \boldsymbol{\rho} \boldsymbol{U}_a^{-1} . \qquad (3.217)$$

Show that, if $\mathcal{E}'(\boldsymbol{\rho}) = \boldsymbol{I}^{\otimes n}$, then $p_a \leq 2^{-2n}$ for each a.

Hint: Note that \mathcal{E} has an operator sum representation with Kraus operators $\{\sigma(x)/2^n\}$ and that \mathcal{E}' has an operator sum representation with Kraus operators $\{\sqrt{p_a} \ U_a\}$. Furthermore $\mathcal{E} = \mathcal{E}'$. Therefore, there exists an $M \times M$ unitary matrix V_{ax} (where $M = \max(N, 2^{2n})$) such that $\sqrt{p_a}U_a = \sum_x V_{ax}\sigma(x)/2^n$. Now express $p_a \operatorname{tr}(U_a U_a^{\dagger})$ in terms of V.

Remark: The result shows that encryption requires $N \ge 2^{2n}$, and that at least 2n bits of key are required to specify U_a . Thus the encryption scheme in which $\sigma(x)$ is applied is the most efficient possible scheme. (For encryption to be effective, it is enough for $\mathcal{E}(\rho)$ to be independent of ρ ; it is not necessary that $\mathcal{E}(\rho) = I^{\otimes n}/2^n$. But the same result applies under the weaker assumption that $\mathcal{E}(\rho)$ is independent of ρ .)

3.5 Unital maps and majorization

Recall that the action of a trace-preserving completely positive (TPCP) map \mathcal{E} can be expressed as

$$\mathcal{E}(\boldsymbol{\rho}) = \sum_{a} \boldsymbol{M}_{a} \boldsymbol{\rho} \boldsymbol{M}_{a}^{\dagger} , \qquad (3.218)$$

where

$$\sum_{a} \boldsymbol{M}_{a}^{\dagger} \boldsymbol{M}_{a} = \boldsymbol{I} . \qquad (3.219)$$

A TPCP map is said to be *unital* if $\mathcal{E}(I) = I$, or equivalently if

$$\sum_{a} \boldsymbol{M}_{a} \boldsymbol{M}_{a}^{\dagger} = \boldsymbol{I} . \qquad (3.220)$$

If A is a nonnegative Hermitian operator with unit trace (tr A = 1), let $\lambda(A)$ denote the vector of eigenvalues of A, which can be regarded as a probability vector. If A and B are nonnegative Hermitian operators with unit trace, we say that $A \prec B$ ("A is majorized by B") if $\lambda(A) \prec \lambda(B)$. (Recall that for two probability vectors pand q, we say that $p \prec q$ if there is a doubly stochastic matrix Dsuch that p = Dq.)

Show that if ρ is a density operator and \mathcal{E} is a unital map, then

$$\mathcal{E}(\boldsymbol{\rho}) \prec \boldsymbol{\rho} \ . \tag{3.221}$$

Hint: Express $\rho = U\Delta U^{\dagger}$ where Δ is diagonal and U is unitary, and express $\rho' \equiv \mathcal{E}(\rho) = V\Delta' V^{\dagger}$, where Δ' is diagonal and V is unitary. Then try to show that the diagonal entries of Δ' can be expressed as a doubly stochastic matrix acting on the diagonal entries of Δ .

Remark: A unital map is the natural quantum generalization of a doubly stochastic map (a doubly stochastic map can be regarded as the special case of a unital map that preserves the basis in which ρ is diagonal). The result of the exercise shows that a unital map takes an input density operator to an output density operator that is no less random than the input.

3.6 What transformations are possible for bipartite pure states?

Alice and Bob share a bipartite pure state $|\Psi\rangle$. Using a 2-LOCC protocol, they wish to transform it to another bipartite pure state $|\Phi\rangle$. Furthermore, the protocol must be *deterministic* — the state $|\Phi\rangle$ is obtained with probability one irrespective of the outcomes of the measurements that Alice and Bob perform.

Suppose that these initial and final states have Schmidt decompositions

$$|\Psi\rangle = \sum_{i} \sqrt{(p_{\Psi})_{i}} |\alpha_{i}\rangle \otimes |\beta_{i}\rangle , \quad |\Phi\rangle = \sum_{i} \sqrt{(p_{\Phi})_{i}} |\alpha_{i}'\rangle \otimes |\beta_{i}'\rangle .$$
(3.222)

Show that if the deterministic transformation $|\Psi\rangle \mapsto |\Phi\rangle$ is possible, then $p_{\Psi} \prec p_{\Phi}$.

Hints: Using the Lo-Popescu Theorem from Exercise 2.9, we can reduce the 2-LOCC to an equivalent 1-LOCC. That is, if the deterministic transformation is possible, then there is a generalized measurement that can be applied by Alice, and an operation depending on Alice's measurement outcome that can be applied by Bob, such that for each possible measurement outcome Alice's measurement followed by Bob's operation maps $|\Psi\rangle$ to $|\Phi\rangle$. Recall that a generalized measurement is defined by a set of operators $\{M_a\}$ such that $\sum_a M_a^{\dagger} M_a = I$, and that the action of the measurement on a pure state $|\psi\rangle$ if outcome *a* occurs is

$$|\psi\rangle \mapsto \frac{M_a |\psi\rangle}{\sqrt{\langle \psi | M_a^{\dagger} M_a |\psi\rangle}}$$
 (3.223)

Think about how the 1-LOCC protocol transforms Alice's density operator. You might want to use the *polar decomposition*: a matrix \boldsymbol{A} can be expressed as $\sqrt{\boldsymbol{A}\boldsymbol{A}^{\dagger}} \boldsymbol{U}$, where \boldsymbol{U} is unitary.

Remark: The converse is also true. Thus majorization provides the necessary and sufficient condition for the deterministic transformation of one bipartite pure state to another (*Nielsen's Theorem*). In this respect, majorization defines a partial order on bipartite pure states such that we may say that $|\Psi\rangle$ is no less entangled than $|\Phi\rangle$ if $p_{\Psi} \prec p_{\Phi}$.

3.7 Fidelity and overlap

The *overlap* of two probability distributions $\{p_i\}$ and $\{\tilde{p}_i\}$ is defined as

$$\operatorname{Overlap}(\{p_i\}, \{\tilde{p}_i\}) \equiv \sum_i \sqrt{p_i \cdot \tilde{p}_i} . \qquad (3.224)$$

Suppose that we try to distinguish the two states ρ and $\tilde{\rho}$ by performing the POVM $\{E_i\}$. Then the two corresponding probability distributions have the overlap

$$\operatorname{Overlap}(\boldsymbol{\rho}, \tilde{\boldsymbol{\rho}}; \{\boldsymbol{E}_i\}) \equiv \sum_i \sqrt{\operatorname{tr} \boldsymbol{\rho} \boldsymbol{E}_i} \cdot \sqrt{\operatorname{tr} \tilde{\boldsymbol{\rho}} \boldsymbol{E}_i} . \qquad (3.225)$$

It turns out that the minimal overlap that can be achieved by any POVM is related to the fidelity $F(\boldsymbol{\rho}, \tilde{\boldsymbol{\rho}}) = \left\| \tilde{\boldsymbol{\rho}}^{\frac{1}{2}} \boldsymbol{\rho}^{\frac{1}{2}} \right\|_{1}^{2}$:

$$\min_{\{\boldsymbol{E}_i\}} \left[\text{Overlap}(\boldsymbol{\rho}, \tilde{\boldsymbol{\rho}}; \{\boldsymbol{E}_i\}) \right] = \sqrt{F(\boldsymbol{\rho}, \tilde{\boldsymbol{\rho}})} . \tag{3.226}$$

In this exercise, you will show that the square root of the fidelity is a lower bound on the overlap, but not that the bound can be saturated.

b) The space of linear operators acting on a Hilbert space is itself a Hilbert space, where the inner product (A, B) of two operators A and B is (A, B) = (A, B) = (A, B)

$$(\boldsymbol{A}, \boldsymbol{B}) \equiv \operatorname{tr}\left(\boldsymbol{A}^{\dagger}\boldsymbol{B}\right) .$$
 (3.227)

For this inner product, the Schwarz inequality becomes

$$|\operatorname{tr} \boldsymbol{A}^{\dagger} \boldsymbol{B}| \leq \left(\operatorname{tr} \boldsymbol{A}^{\dagger} \boldsymbol{A}\right)^{1/2} \left(\operatorname{tr} \boldsymbol{B}^{\dagger} \boldsymbol{B}\right)^{1/2} , \qquad (3.228)$$

Choosing $\mathbf{A} = \boldsymbol{\rho}^{\frac{1}{2}} \mathbf{E}_{i}^{\frac{1}{2}}$ and $\mathbf{B} = \mathbf{U} \tilde{\boldsymbol{\rho}}^{\frac{1}{2}} \mathbf{E}_{i}^{\frac{1}{2}}$ (for an arbitrary unitary \mathbf{U}), use this form of the Schwarz inequality to show that

 $\operatorname{Overlap}(\boldsymbol{\rho}, \tilde{\boldsymbol{\rho}}; \{\boldsymbol{E}_i\}) \geq |\operatorname{tr} \boldsymbol{\rho}^{\frac{1}{2}} \boldsymbol{U} \tilde{\boldsymbol{\rho}}^{\frac{1}{2}}| . \tag{3.229}$

c) Now use the polar decomposition

$$\boldsymbol{A} = \boldsymbol{V} \sqrt{\boldsymbol{A}^{\dagger} \boldsymbol{A}} \tag{3.230}$$

(where V is unitary) to write

$$\tilde{\boldsymbol{\rho}}^{\frac{1}{2}} \boldsymbol{\rho}^{\frac{1}{2}} = \boldsymbol{V} \sqrt{\boldsymbol{\rho}^{\frac{1}{2}} \tilde{\boldsymbol{\rho}} \boldsymbol{\rho}^{\frac{1}{2}}} , \qquad (3.231)$$

and by choosing the unitary U in eq. (3.229) to be $U = V^{-1}$, show that

$$\operatorname{Overlap}(\boldsymbol{\rho}, \tilde{\boldsymbol{\rho}}; \{\boldsymbol{E}_i\}) \ge \sqrt{F(\boldsymbol{\rho}, \tilde{\boldsymbol{\rho}})} . \tag{3.232}$$

d) We can obtain an explicit formula for the fidelity in the case of two states of a single qubit. Using the Bloch parametrization

$$\boldsymbol{\rho}(\vec{P}) = \frac{1}{2} \left(I + \vec{\boldsymbol{\sigma}} \cdot \vec{P} \right) , \qquad (3.233)$$

show that the fidelity of two single-qubit states with polarization vectors \vec{P} and \vec{Q} is

$$F(\vec{P}, \vec{Q}) = \frac{1}{2} \left(1 + \vec{P} \cdot \vec{Q} + \sqrt{(1 - \vec{P}^2)(1 - \vec{Q}^2)} \right) . \quad (3.234)$$

Hint: First note that the eigenvalues of a 2×2 matrix can be expressed in terms of the trace and determinant of the matrix. Then evaluate the determinant and trace of $\left(\rho^{\frac{1}{2}}\tilde{\rho}\rho^{\frac{1}{2}}\right)$, and calculate the fidelity using the corresponding expression for the eigenvalues.

3.8 Semicausal and semilocal maps in the Heisenberg picture

In the Schrödinger picture, a completely positive (CP) map \mathcal{E} leaves observables fixed and takes an input density operator to an output density operator, $\mathcal{E} : \boldsymbol{\rho}_{in} \mapsto \boldsymbol{\rho}_{out} = \mathcal{E}(\boldsymbol{\rho}_{in})$. In the Heisenberg picture, the dual map \mathcal{E}^* leaves density operators fixed and takes an input observable to an output observable, $\mathcal{E}^* : \boldsymbol{a}_{in} \mapsto \boldsymbol{a}_{out} = \mathcal{E}^*(\boldsymbol{a}_{in})$. If \mathcal{E} has the operator sum representation $\mathcal{E}(\boldsymbol{\rho}) = \sum_{\mu} \boldsymbol{M}_a \boldsymbol{\rho} \boldsymbol{M}_a^{\dagger}$, then its dual has operator sum representation

$$\mathcal{E}^*(\boldsymbol{a}) = \sum_a \boldsymbol{M}_a^{\dagger} \boldsymbol{a} \boldsymbol{M}_a \; . \tag{3.235}$$

a) If \mathcal{E} is a TPCP map, show that its dual \mathcal{E}^* can be represented as

$$\mathcal{E}^*(a) = {}_C \langle 0 | \boldsymbol{U}_{AC}^{\dagger} \left(\boldsymbol{a}_A \otimes \boldsymbol{I}_C \right) \boldsymbol{U}_{AC} | 0 \rangle_C , \qquad (3.236)$$

where U_{AC} is a unitary transformation on AC, a_A is an observable on A, I_C is the identity on C, and $|0\rangle_C$ is a fixed pure state in \mathcal{H}_C . (You may use the corresponding property of the TPCP map \mathcal{E} .)

b) Consider a CP map \mathcal{E} acting on a bipartite quantum system AB. We way that \mathcal{E} is *semicausal* if the map does not convey any information from B to A. That is, suppose that Alice and Bob share an initial state ρ_{AB} . Then if Bob performs an operation on B before the map \mathcal{E} acts, and Alice makes a measurement on A after the map \mathcal{E} acts, Alice's measurement collects no information about the operation that Bob performed. Show that if \mathcal{E} is semicausal, then there is an operation $\tilde{\mathcal{E}}$ on A such that

$$\mathcal{E}^*(\boldsymbol{a}_A \otimes \boldsymbol{I}_B) = \tilde{\mathcal{E}}^*(\boldsymbol{a}_A) \otimes \boldsymbol{I}_B . \qquad (3.237)$$

c) We say that \mathcal{E} is *semilocal* if it can be performed by means of local operations and one-way quantum communication from A to B. That is, there is a message system C that can be passed from Alice to Bob. We may assume that the initial state of ABC is a product $\rho_{AB} \otimes \rho_C$ — the state of the message is uncorrelated with the joint state held by Alice and Bob. To apply \mathcal{E} to ρ_{AB} , Alice applies an operation to AC, and sends

C to Bob. Then Bob applies an operation to BC, and discards C. Show that if \mathcal{E} is semilocal, then there are CP maps \mathcal{G}_{AC} from A to AC and \mathcal{F}_{BC} from BC to B such that

$$\mathcal{E}^* = (\mathcal{G}^*_{AC} \otimes I_B) \circ (I_A \otimes \mathcal{F}^*_{BC}) ; \qquad (3.238)$$

here \circ denotes composition of maps, with the map on the right acting first.

d) Using the Heisenberg-picture characterizations of semicausal and semilocal maps found in (b) and (c), show that a semilocal map is semicausal, and express $\tilde{\mathcal{E}}$ in terms of \mathcal{F} and \mathcal{G} .

Remark. The result (d) is intuitively obvious — communication from Alice to Bob cannot convey a signal from Bob to Alice. What is less obvious is that the converse is also true: every semicausal map is semilocal.

3.9 Damped harmonic oscillator at zero temperature

Let's suppose the oscillations of a quantum harmonic oscillator with circular frequency ω are damped because the oscillator can emit photons with energy $\hbar\omega$. When a photon is emitted, the oscillator makes a transition from the energy eigenstate with energy $E_n = n\hbar\omega$ to the energy eigenstate with energy $E_{n-1} = (n-1)\hbar\omega$, and the photon carries away the lost energy. The probability that a photon is emitted in an infinitesimal time interval dt is Γdt ; we say that Γ is the emission rate. Therefore, the coupled evolution of the oscillator and the electromagnetic field for time interval dt can be described as:

$$\begin{aligned} |\Psi(0)\rangle &= |\psi\rangle \otimes |0\rangle \mapsto \\ |\Psi(dt)\rangle &= \sqrt{\Gamma dt} \ \boldsymbol{a}|\psi\rangle \otimes |1\rangle + \left(\boldsymbol{I} - \frac{1}{2}\Gamma dt \ \boldsymbol{a}^{\dagger}\boldsymbol{a}\right)|\psi\rangle \otimes |0\rangle. \end{aligned}$$

$$(3.239)$$

Here $|\psi\rangle$ is the initial normalized state vector of the oscillator and $\{|0\rangle, |1\rangle\}$ are orthonormal states of the electromagnetic field; $|0\rangle$ denotes the state in which no photon has been emitted and $|1\rangle$ denotes the state containing one photon. The operator \boldsymbol{a} reduces the excitation level of the oscillator by one unit, and the $\boldsymbol{a}^{\dagger}\boldsymbol{a}$ factor in the second term is needed to ensure that the evolution is unitary.

a) Check unitarity by verifying that $\langle \Psi(dt) | \Psi(dt) \rangle = 1$, to linear order in the small quantity dt.

3.8 Exercises

Because the states $\{|0\rangle, |1\rangle\}$ of the electromagnetic field are orthogonal, the quantum state of the oscillator may decohere. Summing over these basis states, we see that the initial pure state $|\psi\rangle\langle\psi|$ of the oscillator evolves in time dt as

$$\begin{split} |\psi\rangle\langle\psi| &\mapsto \langle 0|\Psi(dt)\rangle\langle\Psi(dt)|0\rangle + \langle 1|\Psi(dt)\rangle\langle\Psi(dt)|1\rangle \\ &= \Gamma dt \ \boldsymbol{a}|\psi\rangle\langle\psi|\boldsymbol{a}^{\dagger} + \left(\boldsymbol{I} - \frac{1}{2}\Gamma dt \ \boldsymbol{a}^{\dagger}\boldsymbol{a}\right)|\psi\rangle\langle\psi|\left(\boldsymbol{I} - \frac{1}{2}\Gamma dt \ \boldsymbol{a}^{\dagger}\boldsymbol{a}\right); \end{split}$$

More generally, the initial (not necessarily pure) density operator ρ of the oscillator evolves as

$$\boldsymbol{\rho} \mapsto \Gamma dt \ \boldsymbol{a} \boldsymbol{\rho} \boldsymbol{a}^{\dagger} + \left(\boldsymbol{I} - \frac{1}{2} \Gamma dt \ \boldsymbol{a}^{\dagger} \boldsymbol{a} \right) \boldsymbol{\rho} \left(\boldsymbol{I} - \frac{1}{2} \Gamma dt \ \boldsymbol{a}^{\dagger} \boldsymbol{a} \right). \tag{3.240}$$

Now suppose that the initial state of the oscillator is a coherent state

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle, \qquad (3.241)$$

where α is a complex number. For this problem, we will ignore the usual dynamics of the oscillator that causes α to rotate uniformly in time: $\alpha \mapsto \alpha e^{-i\omega t}$; equivalently, we will assume that the dynamics is described in a "rotating frame" such that the rotation of α is transformed away. We will only be interested in how the states of the oscillator are affected by the damping described by eq.(3.240).

b) Show that, to linear order in dt,

$$\left(\boldsymbol{I} - \frac{1}{2}\Gamma dt \; \boldsymbol{a}^{\dagger}\boldsymbol{a}\right) \approx e^{-\Gamma dt|\alpha|^{2}/2} |\alpha \; e^{-\Gamma dt/2}\rangle.$$
 (3.242)

Note that there are two things to check in eq.(3.242): that the value of α decays with time, and that the normalization of the state decays with time.

c) Verify that, also to linear order in dt,

$$\Gamma dt \ \boldsymbol{a} |\alpha\rangle \langle \alpha | \boldsymbol{a}^{\dagger} \approx \Gamma dt |\alpha|^2 \ |\alpha \ e^{-\Gamma dt/2} \rangle \langle \alpha \ e^{-\Gamma dt/2} |, \qquad (3.243)$$

and thus show that, to linear order in dt, $|\alpha\rangle\langle\alpha|$ evolves as

$$|\alpha\rangle\langle\alpha|\mapsto |\alpha \ e^{-\Gamma dt/2}\rangle\langle\alpha \ e^{-\Gamma dt/2}|. \tag{3.244}$$
By considering many consecutive small time increments, argue that, in a finite time t, the initial coherent state evolves as

$$|\alpha\rangle \mapsto |\alpha \ e^{-\Gamma t/2}\rangle.$$
 (3.245)

Thus, the state remains a (pure) coherent state at all times, with the value of α decaying exponentially with time. Since the energy stored in the oscillator is proportional to α^2 , which decays like $e^{-\Gamma t}$, we say that Γ is the *damping rate* of the oscillator.

Now consider what happens if the initial state of the oscillator is a superposition of two coherent states:

$$|\psi\rangle = N_{\alpha,\beta} \left(|\alpha\rangle + |\beta\rangle\right). \tag{3.246}$$

Here $N_{\alpha,\beta}$ is a real nonnegative normalization constant (note that, though the states $|\alpha\rangle$ and $\beta\rangle$ are both normalized, they are not orthogonal).

d) Evaluate $\langle \beta | \alpha \rangle$, and determine $N_{\alpha,\beta}$.

For example we might choose $\alpha = \xi_0/\sqrt{2}$ and $\beta = -\xi_0/\sqrt{2}$, so that the two superposed coherent states are minimum uncertainty wavepackets (with width $\Delta \xi = 1/\sqrt{2}$) centered at dimensionless positions $\pm \xi_0$. If $|\alpha - \beta| \gg 1$, then the two wavepackets are well separated compared to their width, and we might say that oscillator state $|\psi\rangle$ is "in two places at once." How quickly will such a superposition of two separated wavepackets decohere?

The initial density operator of the oscillator is

$$\boldsymbol{\rho} = N_{\alpha,\beta}^2 \big(|\alpha\rangle \langle \alpha| + |\alpha\rangle \langle \beta| + |\beta\rangle \langle \alpha| + |\beta\rangle \langle \beta| \big). \tag{3.247}$$

We already know from part (c) how the "diagonal" terms $|\alpha\rangle\langle\alpha|$ and $|\beta\rangle\langle\beta|$ evolve, but what about the "off-diagonal" terms $|\alpha\rangle\langle\beta|$ and $|\beta\rangle\langle\alpha|$?

e) Using arguments similar to those used in parts (b) and (c), show that in time t, the operator $|\alpha\rangle\langle\beta|$ evolves as

$$|\alpha\rangle\langle\beta|\mapsto e^{i\phi(\alpha,\beta)}e^{-\Gamma t|\alpha-\beta|^2/2}|\alpha e^{-\Gamma t/2}\rangle\langle\beta e^{-\Gamma t/2}|,\qquad(3.248)$$

and find the phase factor $e^{i\phi(\alpha,\beta)}$. Thus the off-diagonal terms decay exponentially with time, at a rate

$$\Gamma_{\text{decohere}} = \frac{1}{2} \Gamma |\alpha - \beta|^2 \qquad (3.249)$$

proportional to the distance squared $|\alpha - \beta|^2$.

f) Consider an oscillator with mass m = 1 g, circular frequency $\omega = 1 s^{-1}$ and (very good) quality factor $Q \equiv \omega/\Gamma = 10^9$. Thus the damping time is very long: over 30 years. A superposition of minimum uncertainty wavepackets is prepared, centered at positions $x = \pm 1$ cm. Estimate the decoherence rate. (Wow! For macroscopic objects, decoherence is really fast!)

3.10 One-qubit decoherence

The matrices $I, \sigma_1, \sigma_2, \sigma_3$, where $\sigma_{1,2,3}$ are the Pauli matrices and bfI is the identity matrix, are a basis for the four-dimensional space of 2×2 matrices. Let us denote I as σ_0 .

a) Let \mathcal{E} be a quantum operation (a completely positive map) acting on the density operator ρ of a single qubit. Show that we may express $\mathcal{E}(\rho)$ as

$$\mathcal{E}(\boldsymbol{\rho}) = \sum_{\mu,\nu=0}^{3} \mathcal{E}_{\mu\nu} \,\boldsymbol{\sigma}_{\mu} \boldsymbol{\rho} \boldsymbol{\sigma}_{\nu} \,, \qquad (3.250)$$

where the $\mathcal{E}_{\mu\nu}$'s are complex numbers satisfying $\mathcal{E}_{\mu\nu} = \mathcal{E}^*_{\nu\mu}$. **Hint**: The operation \mathcal{E} has an operator-sum representation with operation elements $\{M_a\}$. Each M_a can be expanded in the basis $\{\sigma_{\mu}, \mu = 0, 1, 2, 3\}$.

- b) Find four independent conditions that must be satisfied by the $\mathcal{E}_{\mu\nu}$'s in order that the operation \mathcal{E} be trace-preserving (a *channel*).
- c) A Hermitian 2×2 operator can be expressed as

$$\boldsymbol{\rho}(P) = \frac{1}{2} \sum_{\mu=0}^{3} P_{\mu} \boldsymbol{\sigma}_{\mu} , \qquad (3.251)$$

where P_0, P_1, P_2, P_3 are real numbers. Show that a linear map that takes Hermitian operators to Hermitian operators acts as

$$\mathcal{E}(\boldsymbol{\rho}(P)) = \boldsymbol{\rho}(P') , \qquad (3.252)$$

where P' = MP and M is a real matrix. What is the (real) dimension of the space of such linear maps?

d) Suppose that tr $\rho = 1$ and that \mathcal{E} is trace preserving, so that $P_0 = P'_0 = 1$. Show that

$$\vec{P}' = M\vec{P} + \vec{v}$$
, (3.253)

where \vec{P} and $\vec{P'}$ are real three-component polarization vectors, M is a real matrix, and \vec{v} is a real three-component vector. What is the (real) dimension of the space of such trace-preserving maps?

- e) Express \vec{v} in terms of the \mathcal{E}_{0k} 's. Hint: Use the result of (b).
- f) On a Hilbert space of dimension d, the space of linear maps from Hermitian operators to Hermitian operators has real dimension d^4 . What is the dimension of the space of trace-preserving maps? **Hint**: Count the number of independent conditions that must be imposed to ensure that the map is trace preserving.

3.11 Orthogonal or not?

Consider a generalized measurement (POVM) on an *d*-dimensional Hilbert space. There are *d* possible outcomes for the measurement corresponding to the *d* nonnegative operators E_a , $a = 0, 1, 2, \ldots, d-1$, where $\sum_{a=0}^{d-1} E_a = I$. Suppose that each E_a is one-dimensional (has one nonzero eigenvalue). Is this POVM necessarily an orthogonal measurement? Explain your answer.

3.12 Heterodyne measurement of an oscillator

The coherent states $\{|\alpha\rangle, \alpha \in \mathbf{C}\}$ are an overcomplete basis for a one-dimensional harmonic oscillator, satisfying

$$\langle \beta | \alpha \rangle = \exp\left(-\frac{1}{2}|\beta|^2 + \beta^* \alpha - \frac{1}{2}|\alpha|^2\right) \tag{3.254}$$

a) Show that

$$\int d^2 \alpha \, \boldsymbol{E}_{\alpha} = \boldsymbol{I} \, , \qquad (3.255)$$

where

$$\boldsymbol{E}_{\alpha} = \frac{1}{\pi} |\alpha\rangle \langle \alpha | . \qquad (3.256)$$

Hint: Evaluate matrix elements of both sides of the equation between coherent states.

b) Since the E_{α} 's provide a partition of unity, they define a POVM (an "ideal heterodyne measurement" of the oscillator). Suppose that a coherent state $|\beta\rangle$ is prepared, and that an ideal heterodyne measurement is performed, so that the coherent state $|\alpha\rangle$ is obtained with probability distribution $P(\alpha) d^2\alpha =$ $\langle\beta|E_{\alpha}|\beta\rangle d^2\alpha$. With what fidelity does the measurement outcome $|\alpha\rangle$ approximate the initial coherent state $|\beta\rangle$, averaged over the possible outcomes?

3.13 Master equation for the depolarizing channel

a) Consider a depolarizing qubit that is subjected to "Pauli errors" at a rate $\tilde{\Gamma}$, where σ_1 , σ_2 , and σ_3 errors are all equally likely. The depolarization can be described by a master equation with Lindblad operators $\sqrt{\tilde{\Gamma}/3} \sigma_1$, $\sqrt{\tilde{\Gamma}/3} \sigma_2$, and $\sqrt{\tilde{\Gamma}/3} \sigma_3$. Show that this master equation has the form

$$\dot{\boldsymbol{\rho}} = -i[\boldsymbol{H}, \boldsymbol{\rho}] - \Gamma\left(\boldsymbol{\rho} - \frac{1}{2}\boldsymbol{I}\right) . \qquad (3.257)$$

How is Γ related to $\tilde{\Gamma}$?

b) Up to an irrelevant term proportional to the identity, the most general 2×2 Hermitian matrix is

$$\boldsymbol{H} = \frac{\omega}{2} \boldsymbol{n} \cdot \boldsymbol{\sigma} , \qquad (3.258)$$

where \boldsymbol{n} is a unit vector. Use this form of \boldsymbol{H} and the Bloch parametrization

$$\boldsymbol{\rho} = \frac{1}{2} (\boldsymbol{I} + \vec{P} \cdot \vec{\boldsymbol{\sigma}}) , \qquad (3.259)$$

to show that the master equation eq. (3.257) can be rewritten as

$$\vec{P} = \omega(\boldsymbol{n} \times \vec{P}) - \Gamma \vec{P}$$
 . (3.260)

Thus the polarization precesses uniformly with circular frequency ω about the *n*-axis as it contracts with lifetime Γ^{-1} .

c) Alice and Bob play a game in which Alice decides to "turn on" one of the two Hamiltonians

$$\boldsymbol{H} = \frac{\omega}{2} \sigma_3 , \qquad \boldsymbol{H}' = 0 , \qquad (3.261)$$

and Bob is to guess which Hamiltonian Alice chose. Bob has a supply of qubits, and he can observe whether the qubits "precess" in order to distinguish \boldsymbol{H} from \boldsymbol{H}' . However, his qubits are also subject to depolarization at the rate Γ as in eq. (3.257). Suppose that Bob prepares his qubits at time 0 with polarization $\vec{P}_0 = (1,0,0)$; after time t elapses, (1) find the polarization $\vec{P}(t)$ if the Hamiltonian is \boldsymbol{H} and (2) find the polarization $\vec{P}'(t)$ if the Hamiltonian is \boldsymbol{H}' .

d) What is Bob's optimal measurement for distinguishing the polarizations $\vec{P}(t)$ and $\vec{P}'(t)$ (assuming that Alice is as likely to choose \boldsymbol{H} as \boldsymbol{H}' ? What is his optimal probability of error $(p_{\rm e})_{\rm opt}(t)$?

3 Foundations II: Measurement and Evolution

e) The probability of error is smallest if Bob waits for a time t_{best} before measuring. Find t_{best} as a function of Γ and ω . Does your answer make sense in the limits $\Gamma \gg \omega$ and $\Gamma \ll \omega$?

Lecture Notes for Ph219/CS219: Quantum Information and Computation Chapter 4

John Preskill California Institute of Technology

November 2, 2001

Contents

4	Quantum Entanglement	4
4.1	Nonseparability of EPR pairs	4
	4.1.1 Hidden quantum information	4
	4.1.2 Einstein locality and hidden variables	8
4.2	The Bell inequality	10
	4.2.1 Three quantum coins	10
	4.2.2 Quantum entanglement vs. Einstein locality	13
4.3	More Bell inequalities	17
	4.3.1 CHSH inequality	17
	4.3.2 Maximal violation	18
	4.3.3 Quantum strategies outperform classical strategies	20
	4.3.4 All entangled pure states violate Bell inequalities	22
	4.3.5 Photons	24
	4.3.6 Experiments and loopholes	26
4.4	Using entanglement	27
	4.4.1 Dense coding	28
	4.4.2 Quantum teleportation	30
	4.4.3 Quantum teleportation and maximal entanglement	32
	4.4.4 Quantum software	35
4.5	Quantum cryptography	36
	4.5.1 EPR quantum key distribution	36
	4.5.2 No cloning	39
4.6	Mixed-state entanglement	41
	4.6.1 Positive-partial-transpose criterion for separability	43
4.7	Nonlocality without entanglement	45
4.8	Multipartite entanglement	48
	4.8.1 Three quantum boxes	49
	4.8.2 Cat states	55
	4.8.3 Entanglement-enhanced communication	57

	Contents	3
4.9	Manipulating entanglement	59
4.10	Summary	59
4.11	Bibliographical notes	59
4.12	Exercises	59

4 Quantum Entanglement

4.1 Nonseparability of EPR pairs

4.1.1 Hidden quantum information

The deep ways that quantum information differs from classical information involve the properties, implications, and uses of *quantum entanglement*. Recall from §2.4.1 that a bipartite pure state is *entangled* if its Schmidt number is greater than one. Entangled states are interesting because they exhibit correlations that have no classical analog. We will study these correlations in this chapter.

Recall, for example, the maximally entangled state of two qubits (or $EPR \ pair$) defined in §3.4.1:

$$|\phi^+\rangle_{AB} = \frac{1}{\sqrt{2}} (|00\rangle_{AB} + |11\rangle_{AB}) .$$
 (4.1)

"Maximally entangled" means that when we trace over qubit B to find the density operator ρ_A of qubit A, we obtain a multiple of the identity operator

$$\boldsymbol{\rho}_A = \operatorname{tr}_B(|\phi^+\rangle\langle\phi^+|) = \frac{1}{2}\boldsymbol{I}_A , \qquad (4.2)$$

(and similarly $\rho_B = \frac{1}{2} I_B$). This means that if we measure spin A along any axis, the result is completely random — we find spin up with probability 1/2 and spin down with probability 1/2. Therefore, if we perform any local measurement of A or B, we acquire no information about the preparation of the state, instead we merely generate a random bit. This situation contrasts sharply with case of a single qubit in a pure state; there we can store a bit by preparing, say, either $|\uparrow_{\hat{n}}\rangle$ or $|\downarrow_{\hat{n}}\rangle$, and we can recover that bit reliably by measuring along the \hat{n} -axis. With two qubits, we ought to be able to store two bits, but in the state $|\phi^+\rangle_{AB}$ this information is *hidden*; at least, we can't acquire it by measuring A or B.

In fact, $|\phi^+\rangle$ is one member of a basis of four mutually orthogonal states for the two qubits, all of which are maximally entangled — the basis

$$\begin{aligned} |\phi^{\pm}\rangle &= \frac{1}{\sqrt{2}} (|00\rangle \pm |11\rangle) ,\\ |\psi^{\pm}\rangle &= \frac{1}{\sqrt{2}} (|01\rangle \pm |10\rangle) , \end{aligned}$$
(4.3)

introduced in §3.4.1. Imagine that Alice and Bob play a game with Charlie. Charlie prepares one of these four states, thus encoding two bits in the state of the two-qubit system. One bit is the *parity* bit $(|\phi\rangle \text{ or } |\psi\rangle)$: are the two spins aligned or antialigned? The other is the *phase* bit (+ or -): what superposition was chosen of the two states of like parity. Then Charlie sends qubit A to Alice and qubit B to Bob. To win the game, Alice (or Bob) has to identify which of the four states Charlie prepared.

Of course, if Alice and Bob bring their qubits together, they can identify the state by performing an orthogonal measurement that projects onto the $\{|\phi^+\rangle, |\phi^-\rangle, |\psi^+\rangle, |\psi^-\rangle\}$ basis. But suppose that Alice and Bob are in different cities, and that they are unable to communicate at all. Acting locally, neither Alice nor Bob can collect any information about the identity of the state.

What they can do locally is manipulate this information. Alice may apply σ_3 to qubit A, flipping the relative phase of $|0\rangle_A$ and $|1\rangle_A$. This action flips the phase bit stored in the entangled state:

$$\begin{aligned} |\phi^+\rangle \leftrightarrow |\phi^-\rangle , \\ |\psi^+\rangle \leftrightarrow |\psi^-\rangle . \end{aligned}$$
 (4.4)

On the other hand, she can apply σ_1 , which flips her spin $(|0\rangle_A \leftrightarrow |1\rangle_A)$, and also flips the parity bit of the entangled state:

$$\begin{aligned} |\phi^+\rangle \leftrightarrow |\psi^+\rangle , \\ |\phi^-\rangle \leftrightarrow -|\psi^-\rangle . \end{aligned}$$

$$\tag{4.5}$$

Bob can manipulate the entangled state similarly. In fact, as we discussed in §2.4, either Alice or Bob can perform a local unitary transformation that changes one maximally entangled state to any other maximally entangled state.^{*} What their local unitary transformations *cannot* do is alter

^{*} But of course, this does not suffice to perform an arbitrary unitary transformation on the four-dimensional space $\mathcal{H}_A \otimes \mathcal{H}_B$, which contains states that are not maximally entangled. The maximally entangled states are *not* a subspace — a superposition of maximally entangled states typically is *not* maximally entangled.

 $\rho_A = \rho_B = \frac{1}{2}I$ — the information they are manipulating is information that neither one can read.

But now suppose that Alice and Bob are able to exchange (classical) messages about their measurement outcomes; together, then, they can learn about how their measurements are correlated. The entangled basis states are conveniently characterized as the simultaneous eigenstates of two commuting observables:

$$\boldsymbol{\sigma}_{1}^{(A)} \otimes \boldsymbol{\sigma}_{1}^{(B)} ,$$

$$\boldsymbol{\sigma}_{3}^{(A)} \otimes \boldsymbol{\sigma}_{3}^{(B)} ;$$
 (4.6)

the eigenvalue of $\sigma_3^{(A)} \otimes \sigma_3^{(B)}$ is the parity bit, and the eigenvalue of $\sigma_1^{(A)} \otimes \sigma_1^{(B)}$ is the phase bit. Since these operators commute, they can in principle be measured simultaneously. But they cannot be measured simultaneously if Alice and Bob perform localized measurements. Alice and Bob could both choose to measure their spins along the z-axis, preparing a simultaneous eigenstate of $\sigma_3^{(A)}$ and $\sigma_3^{(B)}$. Since $\sigma_3^{(A)}$ and $\sigma_3^{(B)}$ both commute with the parity operator $\sigma_3^{(A)} \otimes \sigma_3^{(B)}$, their orthogonal measurements do not disturb the parity bit, and they can combine their results to infer the parity bit. But $\sigma_3^{(A)}$ and $\sigma_3^{(B)}$ do not commute with phase operator $\sigma_1^{(A)} \otimes \sigma_1^{(A)}$, so their measurement disturbs the phase bit. On the other hand, they could both choose to measure their spins along the *x*-axis; then they would learn the phase bit at the cost of disturbing the parity bit. But they can't have it both ways. To have hope of acquiring the parity bit without disturbing the phase bit, they would need to learn about the product $\sigma_3^{(A)} \otimes \sigma_3^{(B)}$ without finding out anything about $\sigma_3^{(A)}$ and $\sigma_3^{(B)}$ separately. That cannot be done locally.

Now let us bring Alice and Bob together, so that they can operate on their qubits jointly. How might they acquire both the parity bit and the phase bit of their pair? By applying an appropriate unitary transformation, they can rotate the entangled basis $\{|\phi^{\pm}\rangle, |\psi^{\pm}\rangle\}$ to the unentangled basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$. Then they can measure qubits A and B separately to acquire the bits they seek. How is this transformation constructed?

This is a good time to introduce notation that will be used heavily in later chapters, the quantum circuit notation. Qubits are denoted by horizontal lines, and the single-qubit unitary transformation U is denoted:



A particular single-qubit unitary we will find useful is the Hadamard transform

$$\boldsymbol{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1\\ 1 & -1 \end{pmatrix} = \frac{1}{\sqrt{2}} (\boldsymbol{\sigma}_1 + \boldsymbol{\sigma}_3) , \qquad (4.7)$$

which has the properties

$$\boldsymbol{H}^2 = \boldsymbol{I} , \qquad (4.8)$$

and

$$H\sigma_1 H = \sigma_3 ,$$

$$H\sigma_3 H = \sigma_1 .$$
(4.9)

(We can envision H (up to an overall phase) as a $\theta = \pi$ rotation about the axis $\hat{n} = \frac{1}{\sqrt{2}}(\hat{n}_1 + \hat{n}_3)$ that rotates \hat{x} to \hat{z} and vice-versa; we have

$$\boldsymbol{U}(\hat{n},\theta) = \boldsymbol{I}\cos\frac{\theta}{2} + i\hat{n}\cdot\vec{\boldsymbol{\sigma}}\sin\frac{\theta}{2} = i\frac{1}{\sqrt{2}}(\boldsymbol{\sigma}_1 + \boldsymbol{\sigma}_3) = i\boldsymbol{H} .)$$
(4.10)

Also useful is the two-qubit transformation known as the reversible XOR or controlled-NOT transformation; it acts as

$$\mathbf{CNOT}: |a, b\rangle \to |a, a \oplus b\rangle , \qquad (4.11)$$

on the basis states a, b = 0, 1, where $a \oplus b$ denotes addition modulo 2. The **CNOT** is denoted:



Thus this transformation flips the second bit if the first is 1, and acts trivially if the first bit is 0; it has the property

$$(\mathbf{CNOT})^2 = \mathbf{I} \otimes \mathbf{I} \ . \tag{4.12}$$

We call a the control (or source) bit of the **CNOT**, and b the target bit.

By composing these "primitive" transformations, or quantum *gates*, we can build other unitary transformations. For example, the "circuit"



4 Quantum Entanglement

(to be read from left to right) represents the product of H applied to the first qubit followed by **CNOT** with the first bit as the source and the second bit as the target. It is straightforward to see that this circuit transforms the standard basis to the entangled basis,

$$\begin{split} |00\rangle &\to \frac{1}{\sqrt{2}} (|0\rangle + |1\rangle) |0\rangle \to |\phi^{+}\rangle, \\ |01\rangle &\to \frac{1}{\sqrt{2}} (|0\rangle + |1\rangle) |1\rangle \to |\psi^{+}\rangle, \\ |10\rangle &\to \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle) |0\rangle \to |\phi^{-}\rangle, \\ |11\rangle &\to \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle) |1\rangle \to |\psi^{-}\rangle, \end{split}$$
(4.13)

so that the first bit becomes the phase bit in the entangled basis, and the second bit becomes the parity bit.

Similarly, we can invert the transformation by running the circuit backwards (since both **CNOT** and H square to the identity); if we apply the inverted circuit to an entangled state, and then measure both bits, we can learn the value of both the phase bit and the parity bit.

Of course, H acts on only one of the qubits; the "nonlocal" part of our circuit is the controlled-NOT gate — this is the operation that establishes or removes entanglement. If we could only perform an "interstellar **CNOT**," we would be able to create entanglement among distantly separated pairs, or extract the information encoded in entanglement. But we can't. To do its job, the **CNOT** gate must act on its target without revealing the value of its source. Local operations and classical communication will not suffice.

4.1.2 Einstein locality and hidden variables

Einstein was disturbed by quantum entanglement. Eventually, he along with Podolsky and Rosen (EPR) sharpened their discomfort into what they regarded as a paradox. As later reinterpreted by Bohm, the situation they described is really the same as that discussed in §2.5.3. Given a maximally entangled state of two qubits shared by Alice and Bob, Alice can choose one of several possible measurements to perform on her spin that will realize different possible ensemble interpretations of Bob's density matrix; for example, she can prepare either σ_1 or σ_3 eigenstates.

We have seen that Alice and Bob are unable to exploit this phenomenon for faster-than-light communication. Einstein knew this but he was still dissatisfied. He felt that in order to be considered a *complete* description of physical reality a theory should meet a stronger criterion, that might be called *Einstein locality* (also sometimes known as *local realism*): Suppose that A and B are spacelike separated systems. Then in a *complete* description of physical reality an action performed on system A must not modify the description of system B.

But if A and B are entangled, a measurement of A is performed and a *particular* outcome is known to have been obtained, then the density matrix of B does change. Therefore, by Einstein's criterion, the description of a quantum system by a wavefunction or density operator cannot be considered complete.

Einstein seemed to envision a more complete description that would remove the indeterminacy of quantum mechanics. A class of theories with this feature are called *local hidden-variable theories*. In a hidden-variable theory, measurement is actually fundamentally deterministic, but appears to be probabilistic because some degrees of freedom are not precisely known. For example, perhaps when a spin is prepared in what quantum theory would describe as the pure state $|\uparrow_{\hat{z}}\rangle$, there is actually a deeper theory in which the state prepared is parametrized as (\hat{z}, λ) where λ ($0 \leq \lambda \leq 1$) is the hidden variable. Suppose that with present-day experimental technique, we have no control over λ , so when we prepare the spin state, λ might take any value — the probability distribution governing its value is uniform on the unit interval.

Now suppose that when we measure the spin along an axis \hat{n} rotated by θ from the \hat{z} axis, the outcome will be

$$|\uparrow_{\hat{n}}\rangle$$
, for $0 \le \lambda \le \cos^2 \frac{\theta}{2}$,
 $|\downarrow_{\hat{n}}\rangle$, for $\cos^2 \frac{\theta}{2} < \lambda \le 1$. (4.14)

If we know λ , the outcome is deterministic, but if λ is completely unknown, then the probability distribution governing the measurement will agree with the predictions of quantum theory. In a hidden-variable theory, the randomness of the measurement outcome is not intrinsic; rather, it results from ignorance — our description of the system is not the most complete possible description.

Now, what about entangled states? When we say that a hidden-variable theory is *local*, we mean that it satisfies the Einstein locality constraint. A measurement of A does not modify the values of the variables that govern the measurements of B. Rather, when Alice measures her half of an entangled state that she shares with Bob, she gains information about the values of the hidden variables, sharpening her ability to predict what Bob will find when he measures the other half. This seems to be what Einstein had in mind when he envisioned a more complete description.

4.2 The Bell inequality

4.2.1 Three quantum coins

Is a local hidden-variable theory merely a reformulation of quantum mechanics, or is it a testable hypothesis? John Bell's fruitful idea was to test Einstein locality by considering the quantitative properties of the correlations between measurement outcomes obtained by two parties, Alice and Bob, who share an entangled state. Let us consider an example of the sort of correlations that Alice and Bob would like to explain.

The system that Alice and Bob are studying might be described this way: Alice, in Pasadena, has in her possession three coins laid out on a table, labeled 1, 2, 3. Each coin has either its heads (H) or tails (T) side facing up, but it is hidden under an opaque cover, so that Alice is not able to tell whether it is an H or a T. Alice can uncover any one of the three coins, and so learn its value (H or T). However, as soon as that one coin is uncovered, the other two covered coins instantly disappear in a puff of smoke, and Alice never gets an opportunity to uncover the other coins. She has many copies of the three-coin set, and eventually she learns that, no matter which coin she exposes, she is just as likely to find an H as a T. Bob, in Chicago, has a similar set of coins, also labeled 1, 2, 3. He too finds that each one of his coins, when revealed, is as likely to be an H as a T.

In fact, Alice and Bob have many identical copies of their shared set of coins, so they conduct an extensive series of experiments to investigate how their coin sets are correlated with one another. They quickly make a remarkable discovery: Whenever Alice and Bob uncover coins with the same label (whether 1, 2, or 3), they always find coins with the same value — either both are H or both are T. They conduct a million trials, just to be sure, and it works every single time! Their coin sets are perfectly correlated.

Alice and Bob suspect that they have discovered something important, and they frequently talk on the phone to brainstorm about the implications of their results. One day, Alice is in an especially reflective mood:

- Alice: You know, Bob, sometimes it's hard for me to decide which of the three coins to uncover. I know that if I uncover coin 1, say, then coins 2 and 3 will disappear, and I'll never have a chance to find out the values of those coins. Once, just once, I'd love to be able to uncover two of the three coins, and find out whether each is an H or a T. But I've tried and it just isn't possible there's no way to look at one coin and prevent the other from going poof!
- **Bob**: [Long pause] Hey ... wait a minute Alice, I've got an idea ... Look, I think there is a way for you to find the value of two of your

coins, after all! Let's say you would like to uncover coin 1 and coin 2. Well, I'll uncover my coin 2 here in Chicago, and I'll call you to tell you what I found, let's say its an H. We know, then, that you are certain to find an H if you uncover your coin 2 also. There's absolutely no doubt about that, because we've checked it a million times. Right?

Alice: Right ...

- **Bob**: But now there's no reason for you to uncover your coin 2; you know what you'll find anyway. You can uncover coin 1 instead. And then you'll know the value of both coins.
- Alice: Hmmm... yeah, maybe. But we won't be sure, will we? I mean, yes, it always worked when we uncovered the same coin before, but this time you uncovered your coin 2, and your coins 1 and 3 disappeared, and I uncovered my coin 1, and my coins 2 and 3 disappeared. There's no way we'll ever be able to check anymore what would have happened if we had both uncovered coin 2.
- **Bob**: We don't have to check that anymore, Alice; we've already checked it a million times. Look, your coins are in Pasadena and mine are in Chicago. Clearly, there's just no way that my decision to uncover my coin 2 can have any *influence* on what you'll find when you uncover your coin 2. That's not what's happening. It's just that when I uncover my coin 2 we're collecting the information we need to predict with certainty what will happen when you uncover your coin 2. Since we're already certain about it, why bother to do it!
- Alice: Okay, Bob, I see what you mean. Why don't we do an experiment to see what really happens when you and I uncover different coins?
- **Bob**: I don't know, Alice. We're not likely to get any funding to do such a dopey experiment. I mean, does anybody really care what happens when I uncover coin 2 and you uncover coin 1?
- Alice: I'm not sure, Bob. But I've heard about a theorist named Bell. They say that he has some interesting ideas about the coins. He might have a theory that makes a prediction about what we'll find. Maybe we should talk to him.
- **Bob**: Good idea! And it doesn't really matter whether his theory makes any sense or not. We can still propose an experiment to test his prediction, and they'll probably fund us.

So Alice and Bob travel to CERN to have a chat with Bell. They tell Bell about the experiment they propose to do. Bell listens closely, but for a long time he remains silent, with a faraway look in his eyes. Alice and Bob are not bothered by his silence, as they rarely understand anything that theorists say anyway. But finally Bell speaks.

- **Bell**: I think I have an idea When Bob uncovers his coin in Chicago, that can't exert any *influence* on Alice's coin in Pasadena. Instead, what Bob finds out by uncovering his coin reveals some *information* about what will happen when Alice uncovers her coin.
- **Bob**: Well, that's what I've been saying ...
- **Bell**: Right. Sounds reasonable. So let's assume that Bob is right about that. Now Bob can uncover any one of his coins, and know for sure what Alice will find when she uncovers the corresponding coin. He isn't *disturbing* her coin in any way; he's just finding out about it. We're forced to conclude that there must be some *hidden variables* that specify the condition of Alice's coins. And if those variables are completely known, then the value of each of Alice's coins can be unambiguously predicted.
- Bob: [Impatient with all this abstract stuff] Yeah, but so what?
- **Bell**: When your correlated coin sets are prepared, the values of the hidden variables are not completely specified; that's why any one coin is as likely to be an H as a T. But there must be some probability distribution P(x, y, z) (with $x, y, z \in \{H, T\}$) that characterizes the preparation and governs Alice's three coins. These probabilities must be nonnegative, and they sum to one:

$$\sum_{x,y,z \in \{H,T\}} P(x,y,z) = 1 .$$
(4.15)

Alice can't uncover all three of her coins, so she can't measure P(x, y, z) directly. But with Bob's help, she can in effect uncover any two coins of her choice. Let's denote with $P_{\text{same}}(i, j)$, the probability that coins *i* and *j* (i, j = 1, 2, 3) have the same value, either both *H* or both *T*. Then we see that

$$P_{\text{same}}(1,2) = P(HHH) + P(HHT) + P(TTH) + P(TTT) ,$$

$$P_{\text{same}}(2,3) = P(HHH) + P(THH) + P(HTT) + P(TTT) ,$$

$$P_{\text{same}}(1,3) = P(HHH) + P(HTH) + P(THT) + P(TTT) ,$$

(4.16)

and it immediately follows from eq. (4.15) that

$$P_{\text{same}}(1,2) + P_{\text{same}}(2,3) + P_{\text{same}}(1,3)$$

= 1 + 2 P(HHH) + 2 P(TTT) \ge 1. (4.17)

13

So that's my prediction: P_{same} should obey the inequality

$$P_{\text{same}}(1,2) + P_{\text{same}}(2,3) + P_{\text{same}}(1,3) \ge 1$$
. (4.18)

You can test it my doing your experiment that "uncovers" two coins at a time.

- **Bob**: Well, I guess the math looks right. But I don't really get it. Why does it work?
- Alice: I think I see Bell is saying that if there are three coins on a table, and each one is either an H or a T, then at least two of the three have to be the *same*, either both H or both T. Isn't that it, Bell?

Bell stares at Alice, a surprised look on his face. His eyes glaze, and for a long time he is speechless. Finally, he speaks:

Bell: Yes

So Alice and Bob are amazed and delighted to find that Bell is that rarest of beasts — a theorist who makes sense. With Bell's help, their proposal is approved and they do the experiment, only to obtain a shocking result. After many careful trials, they conclude, to very good statistical accuracy that

$$P_{\text{same}}(1,2) \simeq P_{\text{same}}(2,3) \simeq P_{\text{same}}(1,3) \simeq \frac{1}{4}$$
, (4.19)

and hence

$$P_{\text{same}}(1,2) + P_{\text{same}}(2,3) + P_{\text{same}}(1,3) \simeq 3 \cdot \frac{1}{4} = \frac{3}{4} < 1$$
. (4.20)

The correlations found by Alice and Bob flagrantly violate Bell's inequality!

Alice and Bob are good experimenters, but dare not publish so disturbing a result unless they can find a plausible theoretical interpretation. Finally, they become so desperate that they visit the library to see if quantum mechanics can offer any solace ...

4.2.2 Quantum entanglement vs. Einstein locality

What Alice and Bob read about is quantum entanglement. Eventually, they learn that their magical coins are governed by a maximally entangled state of two qubits. What Alice and Bob really share are many copies of the state $|\psi^{-}\rangle$. When Alice uncovers a coin, she is measuring her qubit

14 4 Quantum Entanglement

along one of three possible axes, no two of which are orthogonal. Since the measurements don't commute, Alice can uncover only one of her three coins. Similarly, when Bob uncovers his coin, he measures his member of the entangled pair along any one of three axes, so he too is limited to uncovering just one of his three coins. But since Alice's measurements commute with Bob's, they can uncover one coin each, and study how Alice's coins are correlated with Bob's coins.

To help Alice and Bob interpret their experiment, let's see what quantum mechanics predicts about these correlations. The state $|\psi^{-}\rangle$ has the convenient property that it remains invariant if Alice and Bob each apply the same unitary transformation,

$$\boldsymbol{U} \otimes \boldsymbol{U} |\psi\rangle = |\psi\rangle$$
 . (4.21)

For infinitesimal unitaries, this becomes the property

$$\left(\vec{\boldsymbol{\sigma}}^{(A)} + \vec{\boldsymbol{\sigma}}^{(B)}\right) |\psi^{-}\rangle = 0 \tag{4.22}$$

(the state has vanishing total angular momentum, as you can easily check by an explicit computation). Now consider the expectation value

$$\langle \psi^{-} | \left(\vec{\sigma}^{(A)} \cdot \hat{a} \right) \left(\vec{\sigma}^{(B)} \cdot \hat{b} \right) | \psi^{-} \rangle , \qquad (4.23)$$

where \hat{a} and \hat{b} are unit 3-vectors. Acting on $|\psi^{-}\rangle$, we can replace $\vec{\sigma}^{(B)}$ by $-\vec{\sigma}^{(A)}$; therefore, the expectation value can be expressed as a property of Alice's system, which has density operator $\rho_{A} = \frac{1}{2}I$:

$$- \langle \psi^{-} | \left(\vec{\boldsymbol{\sigma}}^{(A)} \cdot \hat{a} \right) \left(\vec{\boldsymbol{\sigma}}^{(A)} \cdot \hat{b} \right) | \psi^{-} \rangle$$

= $-a_{i}b_{j} \operatorname{tr} \left(\boldsymbol{\rho}_{A} \boldsymbol{\sigma}_{i}^{(A)} \boldsymbol{\sigma}_{j}^{(A)} \right) = -a_{i}b_{j}\delta_{ij} = -\hat{a} \cdot \hat{b} = -\cos\theta ,$ (4.24)

where θ is the angle between the axes \hat{a} and \hat{b} . Thus we find that the measurement outcomes are always perfectly anticorrelated when we measure both spins along the same axis \hat{a} , and we have also obtained a more general result that applies when the two axes are different.

The projection operator onto the spin up (spin down) states along \hat{n} is

 $\boldsymbol{E}(\hat{n},\pm) = \frac{1}{2}(\boldsymbol{I} \pm \hat{n} \cdot \boldsymbol{\vec{\sigma}});$ we therefore obtain

$$P(++) = \langle \psi^{-} | \mathbf{E}^{(A)}(\hat{a}, +) \mathbf{E}^{(B)}(\hat{b}, +) | \psi^{-} \rangle = \frac{1}{4} (1 - \cos \theta) ,$$

$$P(--) = \langle \psi^{-} | \mathbf{E}^{(A)}(\hat{a}, -) \mathbf{E}^{(B)}(\hat{b}, -) | \psi^{-} \rangle = \frac{1}{4} (1 - \cos \theta) ,$$

$$P(+-) = \langle \psi^{-} | \mathbf{E}^{(A)}(\hat{a}, +) \mathbf{E}^{(B)}(\hat{b}, -) | \psi^{-} \rangle = \frac{1}{4} (1 + \cos \theta) ,$$

$$P(-+) = \langle \psi^{-} | \mathbf{E}^{(A)}(\hat{a}, -) \mathbf{E}^{(B)}(\hat{b}, +) | \psi^{-} \rangle = \frac{1}{4} (1 + \cos \theta) ;$$

$$(4.25)$$

here P(++) is the probability that Alice and Bob both obtain the spinup outcome when Alice measures along \hat{a} and Bob measures along \hat{b} , etc. The probability that their outcomes are the same is

$$P_{\text{same}} = P(++) + P(--) = \frac{1}{2}(1 - \cos\theta) , \qquad (4.26)$$

and the probability that their outcomes are opposite is

$$P_{\text{opposite}} = P(+-) + P(-+) = \frac{1}{2}(1 + \cos\theta) . \qquad (4.27)$$

Now suppose that Alice measures her spin along one of the three symmetrically distributed axes in the x - z plane,

$$\hat{a}_{1} = (0, 0, 1) ,$$

$$\hat{a}_{2} = \left(\frac{\sqrt{3}}{2}, 0, -\frac{1}{2}\right) ,$$

$$\hat{a}_{3} = \left(-\frac{\sqrt{3}}{2}, 0, -\frac{1}{2}\right) ,$$
(4.28)

so that

$$\hat{a}_1 \cdot \hat{a}_2 = \hat{a}_2 \cdot \hat{a}_3 = \hat{a}_3 \cdot \hat{a}_1 = -\frac{1}{2}$$
 (4.29)

And suppose that Bob measures along one of three axes that are diametrically opposed to Alice's:

$$\hat{b}_1 = -\hat{a}_1 , \quad \hat{b}_2 = -\hat{a}_2 , \quad \hat{b}_3 = -\hat{a}_3 .$$
 (4.30)

When Alice and Bob choose opposite axes, then $\theta = 180^{\circ}$ and $P_{\text{same}} = 1$. But otherwise $\theta = \pm 60^{\circ}$ so that $\cos \theta = 1/2$ and $P_{\text{same}} = 1/4$. This is just the behavior that Alice and Bob found in their experiment, in violation of Bell's prediction.

4 Quantum Entanglement

Bell's logic seemed compelling but something went wrong, so we are forced to reconsider his tacit assumptions. First, Bell assumed that there is a joint probability distribution that governs the possible outcomes of all measurements that Alice and Bob might perform. This is the hiddenvariable hypothesis. He imagines that if the values of the hidden variables are exactly known, then the outcome of any measurement can be predicted with certainty — measurement outcomes are described probabilistically because the values of the hidden variables are drawn from an ensemble of possible values. Second, Bell assumed that Bob's decision about what to measure in Chicago has no effect on the hidden variables that govern Alice's measurement in Pasadena. This is the assumption that the hidden variables are local. If we accept these two assumptions, there is no escaping Bell's conclusion. We have found that the correlations predicted by quantum theory are incompatible with theses assumptions.

What are the implications? Perhaps the moral of the story is that it can be dangerous to reason about what might have happened, but didn't actually happen — what are sometimes called *counterfactuals*. Of course, we do this all the time in our everyday lives, and we usually get away with it; reasoning about counterfactuals seems to be acceptable in the classical world, but sometimes it gets us into trouble in the quantum world. We claimed that Alice knew what would happen when she measured along \hat{a}_1 , because Bob measured along $-\hat{a}_1$, and every time we have ever checked, their measurement outcomes are always perfectly correlated. But Alice did not measure along \hat{a}_1 ; she measured along \hat{a}_2 instead. We got into trouble by trying to assign probabilities to the outcomes of measurements along \hat{a}_1, \hat{a}_2 , and \hat{a}_3 , even though Alice can perform just one of those measurements. In quantum theory, assuming that there is a probability distribution that governs the outcomes of all three measurements that Alice might have made, even though she was able to carry out only one of these measurements, leads to mathematical inconsistencies, so we had better not do it. We have affirmed Bohr's principle of *complementary* – we are forbidden to consider simultaneously the possible outcomes of two mutually exclusive experiments.

One who rejects the complementarity principle may prefer to say that violations of the Bell inequalities (confirmed experimentally) have exposed an essential nonlocality built into the quantum description of Nature. *If* we do insist that it is legitimate to talk about outcomes of mutually exclusive experiments *then* we are forced to conclude that Bob's choice of measurement actually exerted a subtle *influence* on the outcome of Alice's measurement. Thus advocates of this viewpoint speak of "quantum nonlocality."

By ruling out local hidden variables, Bell demolished Einstein's dream that the indeterminacy of quantum theory could be eradicated by adopting a more complete, yet still local, description of Nature. If we accept locality as an inviolable principle, then we are forced to accept randomness as an unavoidable and intrinsic feature of quantum measurement, rather than a consequence of incomplete knowledge.

To some, the peculiar correlations unmasked by Bell's inequality call out for a deeper explanation than quantum mechanics seems to provide. They see the EPR phenomenon as a harbinger of new physics awaiting discovery. But they may be wrong. We have been waiting over 65 years since EPR, and so far no new physics.

The human mind seems to be poorly equipped to grasp the correlations exhibited by entangled quantum states, and so we speak of the weirdness of quantum theory. But whatever your attitude, experiment forces you to accept the existence of the weird correlations among the measurement outcomes. There is no big mystery about how the correlations were established — we saw that it was necessary for Alice and Bob to get together at some point to create entanglement among their qubits. The novelty is that, even when A and B are distantly separated, we cannot accurately regard A and B as two separate qubits, and use classical information to characterize how they are correlated. They are more than just correlated, they are a single *inseparable* entity. They are *entangled*.

4.3 More Bell inequalities

4.3.1 CHSH inequality

Experimental tests of Einstein locality typically are based on another form of the Bell inequality, which applies to a situation in which Alice can measure either one of two observables a and a', while Bob can measure either b or b'. Suppose that the observables a, a', b, b' take values in $\{\pm 1\}$, and are functions of hidden random variables.

If $a, a' = \pm 1$, it follows that either a+a' = 0, in which case $a-a' = \pm 2$, or else a - a' = 0, in which case $a + a' = \pm 2$; therefore

$$C \equiv (a + a')b + (a - a')b' = \pm 2$$
. (4.31)

(Here is where the local hidden-variable assumption sneaks in — we have imagined that values in $\{\pm 1\}$ can be assigned simultaneously to all four observables, even though it is impossible to measure both of a and a', or both of b and b'.) Evidently

$$|\langle \boldsymbol{C} \rangle| \le \langle |\boldsymbol{C}| \rangle = 2, \tag{4.32}$$

so that

$$|\langle \boldsymbol{a}\boldsymbol{b}\rangle + \langle \boldsymbol{a}'\boldsymbol{b}\rangle + \langle \boldsymbol{a}\boldsymbol{b}'\rangle - \langle \boldsymbol{a}'\boldsymbol{b}'\rangle| \le 2.$$
(4.33)

This result is called the *CHSH inequality* (for Clauser-Horne-Shimony-Holt). It holds for any random variables a, a', b, b' taking values in ± 1 that are governed by a joint probability distribution.

To see that quantum mechanics violates the CHSH inequality, let a, a' denote the Hermitian operators

$$\boldsymbol{a} = \vec{\boldsymbol{\sigma}}^{(A)} \cdot \hat{a} , \quad \boldsymbol{a}' = \vec{\boldsymbol{\sigma}}^{(A)} \cdot \hat{a}' , \qquad (4.34)$$

acting on a qubit in Alice's possession, where \hat{a}, \hat{a}' are unit 3-vectors. Similarly, let $\boldsymbol{b}, \boldsymbol{b}'$ denote

$$\boldsymbol{b} = \vec{\boldsymbol{\sigma}}^{(B)} \cdot \hat{\boldsymbol{b}} , \quad \boldsymbol{b}' = \vec{\boldsymbol{\sigma}}^{(B)} \cdot \hat{\boldsymbol{b}}' , \qquad (4.35)$$

acting on Bob's qubit. Each observable has eigenvalues ± 1 so that an outcome of a measurement of the observable takes values in ± 1 .

Recall that if Alice and Bob share the maximally-entangled state $|\psi^{-}\rangle$, then

$$\langle \psi^{-} | \left(\vec{\boldsymbol{\sigma}}^{(A)} \cdot \hat{a} \right) \left(\vec{\boldsymbol{\sigma}}^{(B)} \cdot \hat{b} \right) | \psi^{-} \rangle = -\hat{a} \cdot \hat{b} = -\cos\theta , \qquad (4.36)$$

where θ is the angle between \hat{a} and \hat{b} . Consider the case where $\hat{a}', \hat{b}, \hat{a}, \hat{b}'$ are coplanar and separated by successive 45° angles. so that the quantummechanical predictions are

$$\langle \boldsymbol{a}\boldsymbol{b}\rangle = \langle \boldsymbol{a}'\boldsymbol{b}\rangle = \langle \boldsymbol{a}\boldsymbol{b}'\rangle = -\cos\frac{\pi}{4} = -\frac{1}{\sqrt{2}},$$
$$\langle \boldsymbol{a}'\boldsymbol{b}'\rangle = -\cos\frac{3\pi}{4} = \frac{1}{\sqrt{2}}.$$
(4.37)

The CHSH inequality then becomes

$$4 \cdot \frac{1}{\sqrt{2}} = 2\sqrt{2} \le 2 , \qquad (4.38)$$

which is clearly violated by the quantum-mechanical prediction.

4.3.2 Maximal violation

In fact the case just considered provides the largest possible quantummechanical violation of the CHSH inequality, as we can see by the following argument. Suppose that a, a', b, b' are Hermitian operators with eigenvalues ± 1 , so that

$$a^2 = a'^2 = b^2 = b'^2 = I$$
, (4.39)

and suppose that "Alice's observables" $\boldsymbol{a}, \boldsymbol{a}'$ commute with "Bob's observables" $\boldsymbol{b}, \boldsymbol{b}'$:

$$0 = [a, b] = [a, b'] = [a', b] = [a', b'].$$
(4.40)

Defining

$$C = ab + a'b + ab' - a'b'$$
, (4.41)

we evaluate

$$C^{2} = \begin{array}{cccc} I & +aa' & +bb' & -aa'bb' \\ +a'a & +I & +a'abb' & -bb' \\ +b'b & +aa'b'b & +I & -aa' \\ -a'ab'b & -b'b & -a'a & +I \end{array},$$
(4.42)

using eq. (4.39). All the quadratic terms cancel pairwise, so that we are left with

$$C^{2} = 4I - aa'bb' + a'abb' + aa'b'b - a'ab'b = 4I - [a, a'][b, b'] .$$
(4.43)

Now recall that the $\sup\ norm\parallel M\parallel_{\sup}$ of a bounded operator M is defined by

$$\|\boldsymbol{M}\|_{\sup} = \sup_{|\psi\rangle} \left(\frac{\|\boldsymbol{M}|\psi\rangle\|}{\||\psi\rangle\|}\right) ; \qquad (4.44)$$

that is, the sup norm of M is the maximum eigenvalue of $\sqrt{M^{\dagger}M}$. It is easy to verify that the sup norm has the properties

$$\| \boldsymbol{M} \boldsymbol{N} \|_{\sup} \leq \| \boldsymbol{M} \|_{\sup} \cdot \| \boldsymbol{N} \|_{\sup} ,$$

$$\| \boldsymbol{M} + \boldsymbol{N} \|_{\sup} \leq \| \boldsymbol{M} \|_{\sup} + \| \boldsymbol{N} \|_{\sup} .$$
(4.45)

A Hermitian operator with eigenvalues ± 1 has unit sup norm, so that

$$\|\boldsymbol{C}^2\|_{\sup} \leq 4 + 4 \|\boldsymbol{a}\|_{\sup} \cdot \|\boldsymbol{a}'\|_{\sup} \cdot \|\boldsymbol{b}\|_{\sup} \cdot \|\boldsymbol{b}'\|_{\sup} = 8.$$
(4.46)

Because C is Hermitian,

$$\| \boldsymbol{C}^2 \|_{\sup} = \| \boldsymbol{C} \|_{\sup}^2 , \qquad (4.47)$$

and therefore

$$\| \boldsymbol{C} \|_{\sup} \leq 2\sqrt{2} , \qquad (4.48)$$

which is known as Cirel'son's inequality.

The CHSH inequality is the statement $|\langle \boldsymbol{C} \rangle| \leq 2$. Quantum mechanically, the absolute value of the expectation value of the Hermitian operator \boldsymbol{C} can be no larger than its largest eigenvalue,

$$|\langle \boldsymbol{C} \rangle| \le \| \boldsymbol{C} \|_{\sup} \le 2\sqrt{2} . \tag{4.49}$$

We saw that this upper bound is saturated in the case where a', b, a, b' are separated by successive 45° angles. Thus the violation of the CHSH inequality that we found is the largest violation allowed by quantum theory.

4.3.3 Quantum strategies outperform classical strategies

The CHSH inequality is a limitation on the strength of the correlations between the two parts of a bipartite classical system, and the Cirel'son inequality is a limitation on the strength of the correlations between the two parts of a bipartite quantum system. We can deepen our appreciation of how quantum correlations differ from classical correlations by considering a game for which quantum strategies outperform classical strategies.

Alice and Bob are playing a game with Charlie. Charlie prepares two bits $x, y \in \{0, 1\}$; then he sends x to Alice and y to Bob. After receiving the input bit x, Alice is to produce an output bit $a \in \{0, 1\}$, and after receiving y, Bob is to produce output bit $b \in \{0, 1\}$. But Alice and Bob are not permitted to communicate, so that Alice does not know y and Bob does not know x.

Alice and Bob win the game if their output bits are related to the input bits according to

$$a \oplus b = x \wedge y , \qquad (4.50)$$

where \oplus denotes the sum modulo 2 (the XOR gate) and \wedge denotes the product (the AND gate). Can Alice and Bob find a strategy that enables them to win the game every time, no matter how Charlie chooses the input bits?

No, it is easy to see that there is no such strategy. Let a_0, a_1 denote the value of Alice's output if her input is x = 0, 1 and let b_0, b_1 denote Bob's output if his input is y = 0, 1. For Alice and Bob to win for all possible inputs, their output bits must satisfy

$$a_0 \oplus b_0 = 0 ,$$

$$a_0 \oplus b_1 = 0 ,$$

$$a_1 \oplus b_0 = 0 ,$$

$$a_1 \oplus b_1 = 1 .$$
(4.51)

But this is impossible, since by summing the four equations we obtain 0=1.

Suppose that Charlie generates the input bits at random. Then there is a very simple strategy that enables Alice and Bob to win the game three times our of four: they always choose the output a = b = 0 so that they lose only if the input is x = y = 1. The CHSH inequality can be regarded as the statement that, if Alice and Bob share no quantum entanglement, then there is no better strategy.

To connect this statement with our previous formulation of the CHSH inequality, define random variables taking values ± 1 as

Then the CHSH inequality says that for any joint probability distribution governing $a, a', b, b' \in \{\pm 1\}$, the expectation values satisfy

$$\langle \boldsymbol{a}\boldsymbol{b}\rangle + \langle \boldsymbol{a}\boldsymbol{b}'\rangle + \langle \boldsymbol{a}'\boldsymbol{b}\rangle - \langle \boldsymbol{a}'\boldsymbol{b}'\rangle \leq 2$$
. (4.53)

Furthermore, if we denote by p_{xy} the probability that eq. (4.51) is satisfied when the input bits are (x, y), then

for example $\langle ab \rangle = p_{00} - (1 - p_{00}) = 2p_{00} - 1$, because the value of ab is +1 when Alice and Bob win and -1 when they lose. The CHSH inequality eq. (4.53) becomes

$$2(p_{00} + p_{01} + p_{10} + p_{11}) - 4 \le 2, \qquad (4.55)$$

or

$$\langle p \rangle \equiv \frac{1}{4} \left(p_{00} + p_{01} + p_{10} + p_{11} \right) \le \frac{3}{4} , \qquad (4.56)$$

where $\langle p \rangle$ denotes the probability of winning averaged over a uniform ensemble for the input bits. Thus, if the input bits are random, Alice and Bob cannot attain a probability of winning higher than 3/4.

It is worthwhile to consider how the assumption that Alice and Bob take actions governed by "*local* hidden variables" limits their success in playing the game. Although Alice and Bob do not share any quantum entanglement, they are permitted to share a table of random numbers that

they may consult to produce their output bits. Thus we may imagine that hidden variables drawn from an ensemble of possible values guide Alice and Bob to make correlated decisions. These correlations are limited by locality — Alice does not know Bob's input and Bob does not know Alice's. In fact, we have learned that for playing this game their shared randomness is of no value — their best strategy does not use the shared randomness at all.

But if Alice and Bob share quantum entanglement, they can devise a better strategy. Based on the value of her input bit, Alice decides to measure one of two Hermitian observables with eigenvalues ± 1 : a if x = 0and a' is x = 1. Similarly, Bob measures b if y = 0 and b' if y = 1. Then the quantum-mechanical expectation values of these observables satisfy the Cirel'son inequality

$$\langle \boldsymbol{a}\boldsymbol{b}\rangle + \langle \boldsymbol{a}\boldsymbol{b}'\rangle + \langle \boldsymbol{a}'\boldsymbol{b}\rangle - \langle \boldsymbol{a}'\boldsymbol{b}'\rangle \le 2\sqrt{2} , \qquad (4.57)$$

and the probability that Alice and Bob win the game is constrained by

$$2(p_{00} + p_{01} + p_{10} + p_{11}) - 4 \le 2\sqrt{2} , \qquad (4.58)$$

or

$$\langle p \rangle \equiv \frac{1}{4} \left(p_{00} + p_{01} + p_{10} + p_{11} \right) \le \frac{1}{2} + \frac{1}{2\sqrt{2}} \approx .853 .$$
 (4.59)

Furthermore, we have seen that this inequality can be saturated if Alice and Bob share a maximally entangled state of two qubits, and the observables a, a', b, b' are chosen appropriately.

Thus we have found that Alice and Bob can play the game more successfully with quantum entanglement than without it. At least for this purpose, shared quantum entanglement is a more powerful resource than shared classical randomness. But even the power brought by entanglement has limits, limits embodied by the Cirel'son inequality.

4.3.4 All entangled pure states violate Bell inequalities

Separable states do not violate Bell inequalities. For example, in the case of a separable *pure* state, if a is an observable acting on Alice's qubit, and b is an observable acting on Bob's, then

$$\langle \boldsymbol{a}\boldsymbol{b}\rangle = \langle \boldsymbol{a}\rangle\langle \boldsymbol{b}\rangle.$$
 (4.60)

No Bell-inequality violation can occur, because we have already seen that a (local) hidden-variable theory *does* exist that correctly reproduces the predictions of quantum theory for a pure state of a single qubit. A general separable state is just a probabilistic mixture of separable pure states, so that the correlations between the subsystems are entirely classical, and the Bell inequalities apply.

On the other hand, we have seen that a maximally entangled state such as $|\psi^-\rangle$ is Bell-inequality violating. But what about pure states that are only partially entangled such as

$$|\phi\rangle = \alpha|00\rangle + \beta|11\rangle ? \tag{4.61}$$

Any pure state of two qubits can be expressed this way in the Schmidt basis; with suitable phase conventions, α and β are real and nonnegative.

Suppose that Alice and Bob both measure along an axis in the x-z plane, so that their observables are

$$\boldsymbol{a} = \boldsymbol{\sigma}_{3}^{(A)} \cos \theta_{A} + \boldsymbol{\sigma}_{1}^{(A)} \sin \theta_{A} ,$$

$$\boldsymbol{b} = \boldsymbol{\sigma}_{3}^{(B)} \cos \theta_{B} + \boldsymbol{\sigma}_{1}^{(B)} \sin \theta_{B} .$$
(4.62)

The state $|\phi\rangle$ has the properties

$$\langle \phi | \boldsymbol{\sigma}_3 \otimes \boldsymbol{\sigma}_3 | \phi \rangle = 1 , \quad \langle \phi | \boldsymbol{\sigma}_1 \otimes \boldsymbol{\sigma}_1 | \phi \rangle = 2\alpha\beta , \langle \phi | \boldsymbol{\sigma}_3 \otimes \boldsymbol{\sigma}_1 | \phi \rangle = 0 = \langle \phi | \boldsymbol{\sigma}_1 \otimes \boldsymbol{\sigma}_3 | \phi \rangle ,$$
 (4.63)

so that the quantum-mechanical expectation value of ab is

$$\langle \boldsymbol{a}\boldsymbol{b}\rangle = \langle \phi | \boldsymbol{a}\boldsymbol{b} | \phi \rangle = \cos\theta_A \cos\theta_B + 2\alpha\beta\sin\theta_A \sin\theta_B$$
(4.64)

(and we recover $\cos(\theta_A - \theta_B)$ in the maximally entangled case $\alpha = \beta = 1/\sqrt{2}$). Now let us consider, for simplicity, the (nonoptimal!) special case

$$\theta_A = 0, \quad \theta'_A = \frac{\pi}{2}, \quad \theta'_B = -\theta_B,$$
(4.65)

so that the quantum predictions are:

$$\langle \boldsymbol{a}\boldsymbol{b} \rangle = \cos\theta_B = \langle \boldsymbol{a}\boldsymbol{b}' \rangle , \langle \boldsymbol{a}'\boldsymbol{b} \rangle = 2\alpha\beta\sin\theta_B = -\langle \boldsymbol{a}'\boldsymbol{b}' \rangle .$$
 (4.66)

Plugging into the CHSH inequality, we obtain

$$|\cos\theta_B - 2\alpha\beta\sin\theta_B| \le 1 , \qquad (4.67)$$

and we easily see that violations occur for θ_B close to 0 or π . Expanding to linear order in θ_B , the left-hand side is

$$\simeq 1 - 2\alpha\beta\theta_B , \qquad (4.68)$$

which surely exceeds 1 for $\alpha\beta > 0$ and θ_B negative and small.

We have shown that *any* entangled pure state of two qubits violates some Bell inequality. It is not hard to generalize the argument to an arbitrary bipartite pure state. For bipartite pure states, then, "entangled" is equivalent to "Bell-inequality violating." For bipartite mixed states, however, we will see later that the situation is more subtle.

4.3.5 Photons

Experiments that test the Bell inequality usually are done with entangled photons, not with spin- $\frac{1}{2}$ objects. What are the quantum-mechanical predictions for photons?

Recall from §2.2.2 that for a photon traveling in the \hat{z} direction, we use the notation $|x\rangle$, $|y\rangle$ for the states that are linearly polarized along the xand y axes respectively. In terms of these basis states, the states that are linearly polarized along "horizontal" and "vertical" axes that are rotated by angle θ relative to the x and y axes can be expressed as

$$|H(\theta)\rangle = \begin{pmatrix} \cos\theta\\\sin\theta \end{pmatrix}$$
, $|V(\theta)\rangle = \begin{pmatrix} -\sin\theta\\\cos\theta \end{pmatrix}$. (4.69)

We can construct a 2×2 matrix whose eigenstates are $|H(\theta)\rangle$ and $|V(\theta)\rangle$, with respective eigenvalues ± 1 ; it is

$$\boldsymbol{\tau}(\theta) \equiv |H(\theta)\rangle \langle H(\theta)| - |V(\theta)\rangle \langle V(\theta)| = \begin{pmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{pmatrix} .$$
(4.70)

The generator of rotations about the \hat{z} axis is $J = \sigma_2$, and the eigenstates of J with eigenvalues ± 1 are the circularly polarized states

$$|+\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1\\i \end{pmatrix}$$
, $|-\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} i\\1 \end{pmatrix}$. (4.71)

Suppose that an excited atom emits two photons that come out back to back, with vanishing angular momentum and even parity. The two-photon states

$$\begin{aligned} |+\rangle_A|-\rangle_B \\ |-\rangle_A|+\rangle_B \end{aligned} (4.72)$$

are invariant under rotations about \hat{z} . The photons have opposite values of J_z , but the same *helicity* (angular-momentum along the axis of propagation), since they are propagating in opposite directions. Under a

reflection in the y-z plane, the polarization states are modified according to

$$|x\rangle \to -|x\rangle , \quad |y\rangle \to |y\rangle , \qquad (4.73)$$

or

$$+\rangle \rightarrow +i|-\rangle , \quad |-\rangle \rightarrow -i|+\rangle ;$$
 (4.74)

therefore, the parity eigenstates are *entangled* states

$$\frac{1}{\sqrt{2}}(|+\rangle_A|-\rangle_B \pm |-\rangle_A|+\rangle_B) . \tag{4.75}$$

The state with $J_z = 0$ and even parity, then, expressed in terms of the linear polarization states, is

$$-\frac{i}{\sqrt{2}}(|+-\rangle_{AB}+|-+\rangle_{AB})$$
$$=\frac{1}{\sqrt{2}}(|xx\rangle_{AB}+|yy\rangle_{AB}) \equiv |\phi^{+}\rangle_{AB} . \qquad (4.76)$$

Because of invariance under rotations about \hat{z} , the state has this form irrespective of how we orient the x and y axes.

Alice or Bob can use a polarization analyzer to project the polarization state of a photon onto the basis $\{|H(\theta)\rangle, |V(\theta)\rangle\}$, and hence measure $\tau(\theta)$. For two photons in the state $|\phi^+\rangle$, if Alice orients her polarizer with angle θ_A and Bob with angle θ_B , then the correlations of their measurement outcomes are encoded in the expectation value

$$\langle \phi^+ | \boldsymbol{\tau}^{(A)}(\theta_A) \boldsymbol{\tau}^{(B)}(\theta_B) | \phi^+ \rangle.$$
 (4.77)

Using rotational invariance:

$$= \langle \phi^{+} | \boldsymbol{\tau}^{(A)}(0) \boldsymbol{\tau}^{(B)}(\theta_{B} - \theta_{A}) | \phi^{+} \rangle$$

$$= \frac{1}{2} \langle x | \boldsymbol{\tau}^{(B)}(\theta_{B} - \theta_{A}) | x \rangle - \frac{1}{2} \langle y | \boldsymbol{\tau}^{(B)}(\theta_{B} - \theta_{A}) | y \rangle$$

$$= \cos 2(\theta_{B} - \theta_{A}) . \qquad (4.78)$$

Recall that for the measurement of qubits on the Bloch sphere, we found the similar expression $\cos \theta$, where θ is the angle between Alice's polarization axis and Bob's. Here we have $\cos 2\theta$ instead, because photons have spin-1 rather than spin- $\frac{1}{2}$.

If Alice measures one of the two observables $\boldsymbol{a} = \boldsymbol{\tau}^{(A)}(\theta_A)$ or $\boldsymbol{a}' = \boldsymbol{\tau}^{(A)}(\theta'_A)$ and Bob measures either $\boldsymbol{b} = \boldsymbol{\tau}^{(B)}(\theta_B)$ or $\boldsymbol{b}' = \boldsymbol{\tau}^{(B)}(\theta_B)$, then under the local hidden-variable assumption the CHSH inequality applies.

If we plug in the quantum predictions for the expectation values, we obtain

$$\left|\cos 2(\theta_B - \theta_A) + \cos 2(\theta_B - \theta'_A) + \cos 2(\theta'_B - \theta_A) - \cos 2(\theta'_B - \theta'_A)\right| \le 2.$$
(4.79)

The maximal violation that saturates Cirel'son's inequality — left-hand side equal to $2\sqrt{2}$ — occurs when θ'_A , θ_B , θ_A and θ'_B are separated by successive $22\frac{1}{2}^{\circ}$ angles, so that

$$\frac{1}{\sqrt{2}} = \cos 2(\theta_B - \theta_A) = \cos 2(\theta_B - \theta'_A)$$
$$= -\cos 2(\theta'_B - \theta_A) = -\cos 2(\theta'_B - \theta'_A) . \quad (4.80)$$

4.3.6 Experiments and loopholes

Locality loophole. Experiments with entangled pairs of photons have tested the CHSH inequality in the form eq. (4.79). The experiments confirm the quantum predictions, and demonstrate convincingly that the CHSH inequality is violated. These experiments, then, seem to show that Nature cannot be accurately described by a local hidden-variable theory.

Or do they? A skeptic might raise objections. For example, in the derivation of the CHSH inequality, we assumed that after Alice decides to measure either a or a', no information about Alice's decision reaches Bob's detector before Bob measures (and likewise, we assumed that if Bob measures first, no information about Bob's decision reaches Alice before she measures). Otherwise, the marginal probability distribution for Bob's outcomes could be updated after Alice's measurement and before Bob's, so that the CHSH inequality need not apply. The assumption that no such update can occur is justified by relativistic causality if Alice's decision and measurement are events spacelike separated from Bob's decision and measurement. The skeptic would insist that the experiment fulfill this condition, which is called the *locality loophole*.

In 1982, Aspect and collaborators conducted an experiment that addressed the locality loophole. Two entangled photons were produced in the decay of an excited calcium atom, and each photon was directed by a switch to one of two polarization analyzers, chosen pseudo-randomly. The photons were detected about 12m apart, corresponding to a light travel time of about 40 ns. This time was considerably longer than either the cycle time of the switch, or the difference in the times of arrival of the two photons. Therefore the "decision" about which observable to measure was made after the photons were already in flight, and the events that selected the axes for the measurement of photons A and B were spacelike separated. The results were consistent with the quantum predictions, and violated the CHSH inequality by five standard deviations. Since Aspect, many other experiments have confirmed this finding, including ones in which detectors A and B are kilometers apart.

Detection loophole. Another objection that the skeptic might raise is called the *detection loophole*. In experiments with photons, the detection efficiency is low. Most entangled photon pairs do not result in detections at both A and B. Among the things that can go wrong: a photon might be absorbed before reaching the detector, a photon might miss the detector, or a photon might arrive in the detector but fail to trigger it. Data is accepted by the experiment only when two photons are detected in co-incidence, and in testing the CHSH inequality, we must assume that the data collected is a fair sample of all the entangled pairs.

But, what if the local hidden variables govern not just *what* polarization state is detected, but also *whether* the detector fires at all? Then the data collected might be a biased sample, and the CHSH inequality need not apply.

In Exercise 4.??, we will show that the detection loophole can be closed if the photons are detected with an efficiency above 82.84%. Current experiments with photons don't approach the necessary efficiency. Experiments that use ion traps have tested the CHSH inequality with detection efficiency close to 100%, but these experiments do not address the locality loophole. No experiment that simultaneously closes the locality and detection loopholes has yet been done.

Free-will loophole. Suppose that an experiment is done in which the photon detection efficiency is perfect, and in which Alice and Bob appear to make spacelike-separated decisions. A skeptic might still resist the conclusion that local hidden-variable theories are ruled out, by invoking the *free-will loophole*. Conceivably, the decisions that Alice and Bob make about what to measure are themselves governed by the local hidden variables. Then their decisions might be correlated with the values of the hidden variables that determine the measurement outcomes, so that they are unable to obtain a fair sample of the distribution of the hidden variables, and the CHSH inequality might be violated.

All of us have to decide for ourselves how seriously to take this objection.

4.4 Using entanglement

After Bell's work, quantum entanglement became a subject of intensive study, among those interested in the foundations of quantum theory. Gradually, a new viewpoint evolved: entanglement is not just a unique tool for exposing the weirdness of quantum mechanics, but also a potentially valuable *resource*. By exploiting entangled quantum states, we can perform tasks that are otherwise difficult or impossible.

4.4.1 Dense coding

Our first example is an application of entanglement to communication. Alice wants to send messages to Bob. She might send classical bits (like dots and dashes in Morse code), but let's suppose that Alice and Bob are linked by a *quantum* channel. For example, Alice can prepare qubits (like photons) in any polarization state she pleases, and send them to Bob, who measures the polarization along the axis of his choice. Is there any advantage to sending qubits instead of classical bits?

In principle, if their quantum channel has perfect fidelity, and Alice and Bob perform the preparation and measurement with perfect efficiency, then they are no *worse* off using qubits instead of classical bits. Alice can prepare, say, either $|\uparrow_z\rangle$ or $|\downarrow_z\rangle$, and Bob can measure along \hat{z} to infer the choice she made. This way, Alice can send one classical bit with each qubit. But in fact, that is the best she can do. Sending one qubit at a time, no matter how she prepares it and no matter how Bob measures it, no more than one classical bit can be carried by each qubit (even if the qubits are entangled with one another). This statement, a special case of the Holevo bound on the classical information capacity of a quantum channel, will be derived in Chapter 5.

But now, let's change the rules a bit — let's suppose that Alice and Bob share an entangled pair of qubits in the state $|\phi^+\rangle_{AB}$. The pair was prepared last year; one qubit was shipped to Alice and the other to Bob, in the hope that the shared entanglement would come in handy someday. Now, use of the quantum channel is very expensive, so Alice can afford to send only one qubit to Bob. Yet it is of the utmost importance for Alice to send Bob *two* classical bits of information.

Fortunately, Alice remembers about the entangled state $|\phi^+\rangle_{AB}$ that she shares with Bob, and she carries out a protocol that she and Bob had arranged for just such an emergency. On her member of the entangled pair, she can perform one of four possible unitary transformations:

1) I (she does nothing),

- 2) σ_1 (180° rotation about \hat{x} -axis),
- 3) σ_2 (180° rotation about \hat{y} -axis),
- 4) σ_3 (180° rotation about \hat{z} -axis).

As we have seen, by doing so, she transforms $|\phi^+\rangle_{AB}$ to one of 4 mutually orthogonal states:

- 1) $|\phi^+\rangle_{AB}$,
- **2)** $|\psi^+\rangle_{AB}$,
- **3)** $|\psi^-\rangle_{AB}$ (up to a phase),
- **4)** $|\phi^{-}\rangle_{AB}$.

Now, she sends her qubit to Bob, who receives it and then performs an orthogonal collective measurement on the pair that projects onto the maximally entangled basis. The measurement outcome unambiguously distinguishes the four possible actions that Alice could have performed. Therefore the single qubit sent from Alice to Bob has successfully carried 2 bits of classical information! Hence this procedure is called "dense coding."

A nice feature of this protocol is that, if the message is highly confidential, Alice need not worry that an eavesdropper will intercept the transmitted qubit and decipher her message. The transmitted qubit has density matrix $\rho_A = \frac{1}{2} I_A$, and so carries no information at all. All the information is in the correlations between qubits A and B, and this information is inaccessible unless the adversary is able to obtain both members of the entangled pair. (Of course, the adversary *can* "jam" the channel, preventing the information from reaching Bob.)

From one point of view, Alice and Bob really *did* need to use the channel twice to exchange two bits of information. For example, we can imagine that Alice prepared the state $|\phi^+\rangle$ herself. Last year, she sent half of the state to Bob, and now she sends him the other half. So in effect, Alice has sent two qubits to Bob in one of four mutually orthogonal states, to convey two classical bits of information as the Holevo bound allows.

Still, dense coding is rather weird, for several reasons. First, Alice sent the first qubit to Bob long before she knew what her message was going to be. Second, each qubit by itself carries no information at all; all the information is encoded in the correlations between the qubits. Third, it would work just as well for Bob to prepare the entangled pair and send half to Alice; then two classical bits are transmitted from Alice to Bob by sending a single qubit from Bob to Alice and back again.

Anyway, when an emergency arose and two bits had to be sent immediately while only one use of the channel was possible, Alice and Bob could exploit the pre-existing entanglement to communicate more efficiently. They used entanglement as a resource.

4 Quantum Entanglement

4.4.2 Quantum teleportation

In dense coding, quantum information could be exploited to enhance the transmission of classical information. Specifically, if Alice and Bob share entanglement, then sending one qubit is sufficient to convey two classical bits. Now one wonders about the converse. If Alice and Bob share entanglement, can sending two classical bits suffice to convey a qubit?

Imagine that Charlie has prepared for Alice a qubit in the state $|\psi\rangle$, but Alice doesn't know anything about what state Charlie prepared. Bob needs this qubit desperately, and Alice wants to help him. But that darn quantum channel is down again! Alice can send only *classical* information to Bob.

She could try measuring $\vec{\sigma} \cdot \hat{n}$, projecting her qubit to either $|\uparrow_{\hat{n}}\rangle$ or $|\downarrow_{\hat{n}}\rangle$. She could send the one-bit measurement outcome to Bob who could then proceed to prepare the state that Alice found. But you showed in Exercise ?? that Bob's state $|\varphi\rangle$ will not be a perfect copy of Alice's; on the average it will match Alice's qubit with fidelity

$$F = |\langle \varphi | \psi \rangle|^2 = \frac{2}{3}, \qquad (4.81)$$

This fidelity is better than could have been achieved if Bob had merely chosen a state at random $(F = \frac{1}{2})$, but it is not nearly as good as the fidelity that Bob requires. Furthermore, as we will see in Chapter 5, there is no protocol in which Alice measures the qubit and sends classical information to Bob that achieves a fidelity better than 2/3.

Fortunately, Alice and Bob recall that they share the maximally entangled state $|\phi^+\rangle_{AB}$, which they prepared last year. Why not use the entanglement as a *resource*? If they are willing to consume the shared entanglement and communicate classically, can Alice send her qubit to Bob with fidelity better than 2/3?

In fact they can achieve fidelity F = 1, by carrying out the following protocol: Alice unites the unknown qubit $|\psi\rangle_C$ she wants to send to Bob with her half of the $|\phi^+\rangle_{AB}$ pair that she shares with Bob. She measures the two commuting observables

$$\boldsymbol{\sigma}_1^{(C)} \otimes \boldsymbol{\sigma}_1^{(A)} , \quad \boldsymbol{\sigma}_3^{(C)} \otimes \boldsymbol{\sigma}_3^{(A)} , \qquad (4.82)$$

thus performing *Bell measurement* — a projection of the two qubits onto one of the four maximally entangled states $|\phi^{\pm}\rangle_{CA}$, $|\psi^{\pm}\rangle_{CA}$. She sends her measurement outcome (two bits of classical information) to Bob over the classical channel. Upon receiving this information, Bob performs one of four operations on his qubit

Alice measures
$$|\phi^+\rangle_{CA} \rightarrow$$
 Bob applies $I^{(B)}$,
Alice measures $|\psi^+\rangle_{CA} \rightarrow$ Bob applies $\sigma_1^{(B)}$,
Alice measures $|\psi^-\rangle_{CA} \rightarrow$ Bob applies $\sigma_2^{(B)}$,
Alice measures $|\phi^-\rangle_{CA} \rightarrow$ Bob applies $\sigma_3^{(B)}$.
(4.83)

This action transforms Bob's qubit (his member of the entangled pair that he initially shared with Alice) into a perfect copy of $|\psi\rangle_C$. This magic trick is called *quantum teleportation*.

How does it work? We merely note that for $|\psi\rangle=a|0\rangle+b|1\rangle,$ we may write

$$\begin{split} |\psi\rangle_{C}|\phi^{+}\rangle_{AB} &= (a|0\rangle_{C} + b|1\rangle_{C})\frac{1}{\sqrt{2}}(|00\rangle_{AB} + |11\rangle_{AB}) \\ &= \frac{1}{\sqrt{2}}(a|000\rangle_{CAB} + a|011\rangle_{CAB} + b|100\rangle_{CAB} + b|111\rangle_{CAB}) \\ &= \frac{1}{2}a(|\phi^{+}\rangle_{CA} + |\phi^{-}\rangle_{CA})|0\rangle_{B} + \frac{1}{2}a(|\psi^{+}\rangle_{CA} + |\psi^{-}\rangle_{CA})|1\rangle_{B} \\ &+ \frac{1}{2}b(|\psi^{+}\rangle_{CA} - |\psi^{-}\rangle_{CA})|0\rangle_{B} + \frac{1}{2}b(|\phi^{+}\rangle_{CA} - |\phi^{-}\rangle_{CA})|1\rangle_{B} \\ &= \frac{1}{2}|\phi^{+}\rangle_{CA}(a|0\rangle_{B} + b|1\rangle_{B}) \\ &+ \frac{1}{2}|\psi^{+}\rangle_{CA}(a|1\rangle_{B} + b|0\rangle_{B}) \\ &+ \frac{1}{2}|\psi^{-}\rangle_{CA}(a|1\rangle_{B} - b|0\rangle_{B}) \\ &+ \frac{1}{2}|\phi^{-}\rangle_{CA}(a|0\rangle_{B} - b|1\rangle_{B}) \\ &= \frac{1}{2}|\phi^{+}\rangle_{CA}|\psi\rangle_{B} + \frac{1}{2}|\psi^{+}\rangle_{CA}\sigma_{1}|\psi\rangle_{B} \\ &+ \frac{1}{2}|\psi^{-}\rangle_{CA}(-i\sigma_{2})|\psi\rangle_{B} + \frac{1}{2}|\phi^{-}\rangle_{CA}\sigma_{3}|\psi\rangle_{B}. \end{split}$$
(4.84)

Thus we see that when Alice performs the Bell measurement on qubits C and A, all four outcomes are equally likely. Once Bob learns Alice's measurement outcome, he possesses the pure state $\sigma |\psi\rangle$, where σ is a known Pauli operator, one of $\{I, \sigma_1, \sigma_2, \sigma_3\}$. The action prescribed in eq. (4.83) restores Bob's qubit to the initial state $|\psi\rangle$.

Quantum teleportation is a curious procedure. Initially, Bob's qubit is completely uncorrelated with the unknown qubit $|\psi\rangle_C$, but Alice's Bell
measurement establishes a correlation between A and C. The measurement outcome is in fact completely random, so Alice (and Bob) actually acquire no information at all about $|\psi\rangle$ by making this measurement. And that is a good thing, as we know that if they were to collect any information about the state they would unavoidably disturb the state.

How then does the quantum state manage to travel from Alice to Bob? It is a bit puzzling. On the one hand, we can hardly say that the two classical bits that were transmitted carried this information — the bits were random. So we are tempted to say that the shared entangled pair made the teleportation possible. But remember that the entangled pair was actually prepared last year, long before Alice ever dreamed that she would be sending the qubit to Bob ...

We should also note that the teleportation procedure is fully consistent with the no-cloning principle. True, a copy of the state $|\psi\rangle_B$ appeared in Bob's hands. But the original $|\psi\rangle_C$ had to be destroyed by Alice's measurement before the copy could be created.

Our findings about dense coding and quantum teleportation can be summarized as statements about how one type of resource can simulate another. Let us introduce the terminology *ebit* for an entangled pair of qubits shared by two parties (*e* for *entangled*), and *cbit* for a classical bit (*c* for *classical*). We teleport one qubit from Alice to Bob by consuming one ebit and sending two cbits, and we send two cbits from Alice and Bob via dense coding by consuming one ebit and transporting one qubit. Thus we may say

$$\begin{array}{rcl} 1 \mbox{ ebit } + \ 2 \mbox{ cbits } & \rightarrow & 1 \mbox{ qubit }, \\ 1 \mbox{ ebit } + \ 1 \mbox{ qubit } & \rightarrow & 2 \mbox{ cbits }, \end{array} \tag{4.85}$$

meaning that the resources on the left suffice to simulate the resources on the right. Entanglement is essential in these protocols. Without ebits, a qubit is worth only one cbit, and without ebits, a "teleported" qubit has fidelity $F \leq 2/3$.

4.4.3 Quantum teleportation and maximal entanglement

The teleportation concept has an air of mystery. One would like to understand more deeply why it works. A helpful clue is that to teleport with fidelity F = 1 the entangled state consumed in the protocol must be *maximally* entangled. And the crucial feature of bipartite maximally entangled states is that *either Alice or Bob* can transform one maximally entangled state to another by applying a local unitary transformation.

To see more clearly how quantum teleportation works, consider teleporting an N-dimensional system using an $N \times N$ maximally entangled state of the form

$$|\Phi\rangle = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle \otimes |i\rangle . \qquad (4.86)$$

A useful property of this state is

$$C_{A}\langle \Phi | \Phi \rangle_{AB} = \frac{1}{N} \sum_{i,j} (C \langle i | \otimes_{A} \langle i |) (|j\rangle_{A} \otimes |j\rangle_{B})$$
$$= \frac{1}{N} \sum_{i} |i\rangle_{B} C \langle i | \equiv \frac{1}{N} (\mathbf{T})_{BC}$$
(4.87)

Here we have defined the *transfer operator* $(T)_{BC}$ which has the property

$$\boldsymbol{T}_{BC}|\varphi\rangle_{C} = \boldsymbol{T}_{BC}\left(\sum_{i}a_{i}|i\rangle_{C}\right) = \sum_{i}a_{i}|i\rangle_{B} = |\varphi\rangle_{B} ; \qquad (4.88)$$

it maps a state in C to the identical state in B. This property has no invariant meaning independent of the choice of basis in B and C; rather T_{BC} just describes an arbitrary way to relate the orthonormal bases of the two systems. Of course, Alice and Bob would need to align their bases in some way to verify that teleportation has really succeeded.

Now recall that any other $N \times N$ maximally entangled state has a Schmidt decomposition of the form

$$\frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i'\rangle \otimes |i\rangle , \qquad (4.89)$$

and so can be expressed as

$$|\Phi(\boldsymbol{U})\rangle \equiv \boldsymbol{U} \otimes \boldsymbol{I} |\Phi\rangle , \qquad (4.90)$$

where

$$\boldsymbol{U}|i\rangle = |i'\rangle = \sum_{j} |j\rangle U_{ji}$$
 (4.91)

Writing

$$|\Phi(\boldsymbol{U})\rangle_{AB} = \frac{1}{\sqrt{N}} \sum_{i,j} |j\rangle_A \otimes |i\rangle_B \ U_{ji} \ , \tag{4.92}$$

we can easily verify that

$${}_{CA}\langle \Phi(\boldsymbol{U})|\Phi(\boldsymbol{V}^{T})\rangle_{AB} = \frac{1}{N} \left(\boldsymbol{V}\boldsymbol{U}^{-1}\right)_{B} \boldsymbol{T}_{BC} , \qquad (4.93)$$

where V^T denotes the transpose of V in the standard basis $(V_{ij}^T = V_{ji})$; in particular, then, the transfer operator can be expressed as

$$\frac{1}{N}\boldsymbol{T}_{BC} = {}_{CA}\langle \Phi(\boldsymbol{U}) | \Phi((\boldsymbol{U}^T)) \rangle_{AB} , \qquad (4.94)$$

for any unitary U.

Now suppose that Alice and Bob share $|\Phi\rangle_{AB}$, and that Charlie has prepared the state $|\psi\rangle_C$ and has deposited it in Alice's laboratory. Alice performs a measurement that projects CA onto a maximally entangled basis, obtaining the outcome $|\Phi(U_a)\rangle_{CA}$ for some unitary U_a . Then we know from eq. (4.94) that *if* Alice and Bob had shared the state $|\Phi((U_a^T)\rangle_{AB}$ instead of $|\Phi\rangle_{AB}$, then Alice's measurement would have prepared in Bob's lab a perfect replica of the state $|\psi\rangle$. Unfortunately, they did not have the foresight to share the right state to begin with. But it's not too late! Bob realizes that

$$|\Phi(\boldsymbol{U}_{a}^{T})\rangle = \boldsymbol{I}_{A} \otimes (\boldsymbol{U}_{a})_{B} |\Phi\rangle_{AB} , \qquad (4.95)$$

and of course $(\boldsymbol{U}_a)_B$ commutes with Alice's measurement. Hence, when Bob hears from Alice that her measurement outcome was $|\Phi((\boldsymbol{U}_a^T)\rangle_{AB})$, he applies $(\boldsymbol{U}_a)_B$ to his half of the state he had shared with Alice. Then the protocol is equivalent to one in which they had shared the right maximally entangled state to begin with, and Bob's state has been transformed into $|\psi\rangle_B!$

This approach to teleportation has some conceptual advantages. For one, we can easily see that Alice is not required to perform an orthogonal measurement. To achieve teleportation with fidelity F = 1 it suffices that she perform a POVM with operation elements M_a , where each M_a has the property

$$M_a^{\dagger} M_a \propto |\Phi(U_a)\rangle \langle \Phi(U_a)|$$
 (4.96)

for some unitary U_a . Also, we can easily see how the teleportation protocol should be modified if the initial maximally entangled state that Alice and Bob share is not $|\Phi\rangle_{AB}$ but rather

$$|\Phi(\mathbf{V}^T)\rangle_{AB} = \mathbf{I}_A \otimes \mathbf{V}_B |\Phi\rangle_{AB} .$$
 (4.97)

If Alice's measurement outcome is $|\Phi(U_a)\rangle_{CA}$, then eq. (4.93) tells us that the state Bob receives is

$$\boldsymbol{V}\boldsymbol{U}_{a}^{-1}|\psi\rangle_{B}. \qquad (4.98)$$

To recover $|\psi\rangle_B$, Bob must apply $U_a V^{-1}$.

The operator ordering in eq. (4.98) may seem counterintuitive at first it seems as though Alice's measurement (U_a) precedes the preparation of the shared entangled state (V). But this "time reversal" has a straightforward interpretation. If Alice's measurement outcome is $|\Phi(U_a)\rangle_{CA}$, then Bob would have received a perfect copy of $|\psi\rangle$ if the initial entangled state had been $I_A \otimes (U_a)_B |\Phi\rangle_{AB}$. To simulate the situation in which the entangled state had been chosen properly from the start, Bob first applies V^{-1} to undo the "twist" in $|\Phi(V^T)\rangle_{AB}$, recovering $|\Phi\rangle_{AB}$, and then applies U_a to transform the entangled state to the desired one.

There is a more fanciful interpretation of eq. (4.98) which, while not necessary, is nonetheless irresistable. We might "explain" how quantum information is transferred from Alice and Bob by following the world line of a qubit traveling in spacetime. The qubit moves forward in time from Charlie's preparation to Alice's measurement, then backward in time from the measurement to the initial preparation of the entangled pair, and finally forward in time again from the preparation of the pair to Bob's laboratory. Since this world line visits Alice's measurement before arriving at the preparation of the entanglement, U_a^{-1} acts "first" and V acts "later on."

4.4.4 Quantum software

Teleportation has some interesting applications. For example, imagine that Alice and Bob wish to apply the "quantum gate" V to the unknown state $|\psi\rangle_C$. But applying V requires sophisitcated hardware that they can't afford.

A more economical alternative is to purchase *quantum software* from a vendor. The software is a bipartite state that the vendor certifies to be

$$|\Phi(\boldsymbol{V}^T)\rangle_{AB} = \boldsymbol{I}_A \otimes \boldsymbol{V}_B |\phi\rangle_{AB} . \tag{4.99}$$

Alice's hardware is powerful enough for her to perform a measurement that projects onto the basis { $|\Phi(\boldsymbol{U}_a)\rangle_{CA}$ }; once the outcome *a* is known, the state $\boldsymbol{V}\boldsymbol{U}_a^{-1}|\psi\rangle_B$ has been prepared. Bob can then complete the execution of \boldsymbol{V} to $|\psi\rangle$ by applying $\boldsymbol{V}\boldsymbol{U}_a\boldsymbol{V}^{-1}$

This procedure may seem silly — why assume that Bob is able to apply VU_aV^{-1} but unable to apply V? In fact it is not so silly, and has important applications to fault-tolerant quantum computation that we will explore further in Chapter 8. In some cases, executing VU_aV^{-1} really is a lot easier than applying V. Furthermore, Alice and Bob might be able to prepare the quantum software themselves, instead of buying it, even though they can't apply V reliably. This is possible because it is easier to verify that a *known* quantum state has been properly prepared than to

verify that a known unitary transformation has been successfully applied to an unknown state. If the hardware that applies V cannot be trusted, then we prefer to use it to prepare software offline, and then subject the software to quality assurance, rather than risk causing irrevocable damage to our unknown state through a faulty execution of V.

Each application of V consumes one copy of the quantum software. Thus, this protocol for executing V with the help of quantum software uses entanglement as a resource.

4.5 Quantum cryptography

4.5.1 EPR quantum key distribution

Everyone has secrets, including Alice and Bob. Alice needs to send a highly private message to Bob, but Alice and Bob have a very nosy friend, Eve, who they know will try to listen in. Can they communicate with assurance that Eve is unable to eavesdrop?

Obviously, they should use some kind of code. Trouble is, aside from being very nosy, Eve is also very smart. Alice and Bob are not confident that they are clever enough to devise a code that Eve cannot break.

Except there is one coding scheme that is surely unbreakable. If Alice and Bob share a *private key*, a string of random bits known only to them, then Alice can convert her message to ASCII (a string of bits no longer than the key) *add* each bit of her message (module 2) to the corresponding bit of the key, and send the result to Bob. Receiving this string, Bob can add the key to it to extract Alice's message.

This scheme is secure because even if Eve should intercept the transmission, she will not learn anything because the transmitted string itself carries no information — the message is encoded in a correlation between the transmitted string and the key (which Eve doesn't know).

There is still a problem, though, because Alice and Bob need to establish a shared random key, and they must ensure that Eve can't know the key. They could meet to exchange the key, but that might be impractical. They could entrust a third party to transport the key, but what if the intermediary is secretly in cahoots with Eve? They could use "public key" distribution protocols, but the security of such protocols is founded on assumptions about the computational resources available to a potential adversary. Indeed, we will see in Chapter 6 that public key protocols are vulnerable to attack by an eavesdropper who is equipped with a quantum computer.

Can Alice and Bob exploit *quantum* information (and specifically entanglement) to solve the key exchange problem? They can! *Quantum key distribution* protocols can be devised that are invulnerable to any attack allowed by the laws of physics.

Let's suppose that Alice and Bob share a supply of entangled pairs, each prepared in the state $|\phi^+\rangle$. To establish a shared private key, they may carry out this protocol:

For each qubit in her/his possession, Alice and Bob decide to measure either σ_1 or σ_3 . The decision is pseudo-random, each choice occuring with probability 1/2. Then, after the measurements are performed, both Alice and Bob publicly announce what observables they measured, but do not reveal the outcomes they obtained. For those cases (about half) in which they measured their qubits along different axes, their results are discarded (as Alice and Bob obtained uncorrelated outcomes). For those cases in which they measured along the same axis, their results, though random, are *perfectly correlated*. Hence, they have established a shared random key.

But, is this protocol really invulnerable to a sneaky attack by Eve? In particular, Eve might have clandestinely tampered with the pairs at some time in the past. Then the pairs that Alice and Bob possess might be (unbeknownst to Alice and Bob) not perfect $|\phi^+\rangle$'s, but rather pairs that are entangled with qubits in Eve's possession. Eve can then wait until Alice and Bob make their public announcements, and proceed to measure her qubits in a manner designed to acquire maximal information about the results that Alice and Bob obtained. Alice and Bob must protect themselves against this type of attack.

If Eve has indeed tampered with Alice's and Bob's pairs, then the most general possible state for an AB pair and a set of E qubits has the form

$$\begin{aligned} |\Upsilon\rangle_{ABE} &= |00\rangle_{AB} |e_{00}\rangle_E + |01\rangle_{AB} |e_{01}\rangle_E \\ &+ |10\rangle_{AB} |e_{10}\rangle_E + |11\rangle_{AB} |e_{11}\rangle_E , \end{aligned}$$
(4.100)

where Eve's states $|e_{ij}\rangle_E$ are neither normalized nor mutually orthogonal. Now recall that the defining property or $|\phi^+\rangle$ is that it is an eigenstate with eigenvalue +1 of both $\sigma_1^{(A)}\sigma_1^{(B)}$ and $\sigma_3^{(A)}\sigma_3^{(B)}$. Suppose that A and B are able to verify that the pairs in their possession have this property. To satisfy $\sigma_3^{(A)}\sigma_3^{(B)} = 1$, we must have

$$|\Upsilon\rangle_{AB} = |00\rangle_{AB}|e_{00}\rangle_E + |11\rangle_{AB}|e_{11}\rangle_E , \qquad (4.101)$$

and to also satisfy $\boldsymbol{\sigma}_1^{(A)} \boldsymbol{\sigma}_1^{(B)} = 1$, we must have

$$|\Upsilon\rangle_{ABE} = \frac{1}{\sqrt{2}} (|00\rangle_{AB} + |11\rangle_{AB})|e\rangle_E = |\phi^+\rangle_{AB}|e\rangle_E .$$
(4.102)

We see that it is possible for the AB pairs to be eigenstates of $\sigma_1^{(A)}\sigma_1^{(B)}$ and $\sigma_3^{(A)}\sigma_3^{(B)}$ only if they are completely unentangled with Eve's qubits. Therefore, Eve will not be able to learn anything about Alice's and Bob's measurement results by measuring her qubits. The random key is secure.

To verify the properties $\sigma_1^{(A)}\sigma_1^{(B)} = 1 = \sigma_3^{(A)}\sigma_3^{(B)}$, Alice and Bob can sacrifice a portion of their shared key, and publicly compare their measurement outcomes. They should find that their results are indeed perfectly correlated. If so they will have high statistical confidence that Eve is unable to intercept the key. If not, they have detected Eve's nefarious activity. They may then discard the key, and make a fresh attempt to establish a secure key.

As I have just presented it, the quantum key distribution protocol seems to require entangled pairs shared by Alice and Bob, but this is not really so. We might imagine that Alice prepares the $|\phi^+\rangle$ pairs herself, and then measures one qubit in each pair before sending the other to Bob. This is completely equivalent to a scheme in which Alice prepares one of the four states

$$|\uparrow_{z}\rangle, |\downarrow_{z}\rangle, |\uparrow_{x}\rangle, |\downarrow_{x}\rangle, \qquad (4.103)$$

(chosen at random, each occuring with probability 1/4) and sends the qubit to Bob. Bob's measurement and the verification are then carried out as before. This scheme (known as the BB84 quantum key distribution protocol) is just as secure as the entanglement-based scheme.[†]

Another intriguing variation is called the "time-reversed EPR" scheme. Here both Alice and Bob prepare one of the four states in eq. (4.103), and they both send their qubits to Charlie. Then Charlie performs a Bell measurement on the pair — that is, he measures $\sigma_1^{(A)} \sigma_1^{(B)}$ and $\sigma_3^{(A)} \sigma_3^{(B)}$, orthogonally projecting out one of $|\phi^{\pm}\rangle|\psi^{\pm}\rangle$, and he publicly announces the result. Since all four of these states are simultaneous eigenstates of $\sigma_1^{(A)} \sigma_1^{(B)}$ and $\sigma_3^{(A)} \sigma_3^{(B)}$, when Alice and Bob both prepared their spins along the same axis (as they do about half the time) they share a single bit.[‡] Of course, Charlie could be allied with Eve, but Alice and Bob can verify that Charlie and Eve have acquired no information as before, by comparing a portion of their key. This scheme has the advantage that Charlie could operate a central switching station by storing qubits received from many parties, and then perform his Bell measurement when two of the parties request a secure communication link. (Here we assume that Charlie has a stable quantum memory in which qubits can be stored

[†] Except that in the EPR scheme, Alice and Bob can wait until just before they need to talk to generate the key, thus reducing the risk that Eve might at some point burglarize Alice's safe to learn what states Alice prepared (and so infer the key).

[‡] Until Charlie makes his measurement, the states prepared by Bob and Alice are totally uncorrelated. A definite correlation (or anti-correlation) is established after Charlie performs his measurement.

accurately for as long as necessary.) A secure key can be established even if the quantum communication line is down temporarily, as long as both parties had the foresight to send their qubits to Charlie on an earlier occasion (when the quantum channel was open).

So far, we have made the unrealistic assumption that the quantum communication channel is perfect, but of course in the real world errors will occur. Therefore even if Eve has been up to no mischief, Alice and Bob will sometimes find that their verification test will fail. But how are they to distinguish errors due to imperfections of the channel from errors that occur because Eve has been eavesdropping?

To address this problem, Alice and Bob can enhance their protocol in two ways. First they implement (classical) error correction to reduce the effective error rate. For example, to establish each bit of their shared key they could actually exchange a block of three random bits. If the three bits are not all the same, Alice can inform Bob which of the three is different than the other two; Bob can flip that bit in his block, and *then* use majority voting to determine a bit value for the block. This way, Alice and Bob share the same key bit even if an error occured for one bit in the block of three.

However, error correction alone does not suffice to ensure that Eve has acquired negligible information about the key — error correction must be supplemented by (classical) privacy amplification. For example, after performing error correction so that they are confident that they share the same key, Alice and Bob might extract a bit of "superkey" as the *parity* of n key bits. To know *anything* about the parity of n bits, Eve would need to know *something* about each of the bits. Therefore, the parity bit is considerably more secure, on the average, than each of the individual key bits.

If the error rate of the channel is low enough, one can show that quantum key distribution, supplemented by error correction and privacy amplification, is invulnerable to any attack that Eve might muster (in the sense that the information acquired by Eve can be guaranteed to be arbitrarily small). We will return to the problem of proving the security of quantum key distribution in Chapter 7.

4.5.2 No cloning

The security of quantum key distribution is based on an essential difference between quantum information and classical information. It is not possible to acquire information that *distinguishes* between nonorthogonal quantum states without *disturbing* the states.

For example, in the BB84 protocol, Alice sends to Bob any one of the four states $|\uparrow_z\rangle|\downarrow_z\rangle|\uparrow_x\rangle|\downarrow_x\rangle$, and Alice and Bob are able to verify that

none of their states are perturbed by Eve's attempt at eavesdropping. Suppose, more generally, that $|\varphi\rangle$ and $|\psi\rangle$ are two nonorthogonal states in $\mathcal{H}(\langle \psi | \varphi \rangle \neq 0)$ and that a unitary transformation U is applied to $\mathcal{H} \otimes \mathcal{H}_E$ (where \mathcal{H}_E is a Hilbert space accessible to Eve) that leaves both $|\psi\rangle$ and $|\varphi\rangle$ undisturbed. Then

$$U: \quad |\psi\rangle \otimes |0\rangle_E \to |\psi\rangle \otimes |e\rangle_E , |\varphi\rangle \otimes |0\rangle_E \to |\varphi\rangle \otimes |f\rangle_E , \qquad (4.104)$$

and unitarity implies that

$$\langle \psi | \phi \rangle = (_E \langle 0 | \otimes \langle \psi |) (| \varphi \rangle \otimes | 0 \rangle_E) = (_E \langle e | \otimes \langle \psi |) (| \varphi \rangle \otimes | f \rangle_E) = \langle \psi | \varphi \rangle \langle e | f \rangle .$$
 (4.105)

Hence, for $\langle \psi | \varphi \rangle \neq 0$, we have $\langle e | f \rangle = 1$, and therefore since the states are normalized, $|e\rangle = |f\rangle$. This means that no measurement in \mathcal{H}_E can reveal any information that distinguishes $|\psi\rangle$ from $|\varphi\rangle$. In the BB84 case this argument shows that, if Eve does not disturb the states sent by Alice, then the state in \mathcal{H}_E is the same irrespective of which of the four states $|\uparrow_z\rangle, |\downarrow_z\rangle, |\uparrow_x\rangle, |\downarrow_x\rangle$ is sent by Alice, and therefore Eve learns nothing about the key shared by Alice and Bob. On the other hand, if Alice is sending to Bob one of the two orthogonal states $|\uparrow_z\rangle$ or $|\downarrow_z\rangle$, there is nothing to prevent Eve from acquiring a copy of the information (as with classical bits).

We have noted earlier that if we have many identical copies of a qubit, then it is possible to measure the mean value of noncommuting observables like σ_1, σ_2 , and σ_3 to completely determine the density matrix of the qubit. Inherent in the conclusion that nonorthogonal state cannot be distinguished without disturbing them, then, is the implicit provision that it is not possible to make a perfect copy of a qubit. (If we could, we would make as many copies as we need to find $\langle \sigma_1 \rangle, \langle \sigma_2 \rangle$, and $\langle \sigma_3 \rangle$ to any specified accuracy.) Let's now make this point explicit: there is no such thing as a perfect quantum Xerox machine.

Orthogonal quantum states (like classical information) *can* be reliably copied. For example, the unitary transformation that acts as

$$U: \quad |0\rangle_A |0\rangle_E \to |0\rangle_A |0\rangle_E , |1\rangle_A |0\rangle_E \to |1\rangle_A |1\rangle_E , \qquad (4.106)$$

copies the first qubit onto the second if the first qubit is in one of the states $|0\rangle_A$ or $|1\rangle_A$. But if instead the first qubit is in the state $|\psi\rangle = a|0\rangle_A + b|1\rangle_A$, then

$$U: \quad (a|0\rangle_A + b|1\rangle_A)|0\rangle_E \rightarrow a|0\rangle_A|0\rangle_E + b|1\rangle_A|1\rangle_E .$$
(4.107)

This is *not* the state $|\psi\rangle \otimes |\psi\rangle$ (a tensor product of the original and the copy); rather it is something very different – an entangled state of the two qubits.

To consider the most general possible quantum Xerox machine, we allow the full Hilbert space to be larger than the tensor product of the space of the original and the space of the copy. Then the most general "copying" unitary transformation acts as

$$U: \quad |\psi\rangle_A|0\rangle_E|0\rangle_F \to |\psi\rangle_A|\psi\rangle_E|e\rangle_F \quad |\varphi\rangle_A|0\rangle_E|0\rangle_F \to |\varphi\rangle_A|\varphi\rangle_E|f\rangle_F .$$
(4.108)

Unitarity then implies that

$$\langle \psi | \varphi \rangle = \langle \psi | \varphi \rangle \langle \psi | \varphi \rangle \langle e | f \rangle ; \qquad (4.109)$$

therefore, if $\langle \psi | \varphi \rangle \neq 0$, then

$$1 = \langle \psi | \varphi \rangle \langle e | f \rangle. \tag{4.110}$$

Since the states are normalized, we conclude that

$$|\langle \psi | \varphi \rangle| = 1, \tag{4.111}$$

so that $|\psi\rangle$ and $|\varphi\rangle$ actually represent the same ray. No unitary machine can make a copy of both $|\varphi\rangle$ and $|\psi\rangle$ if $|\varphi\rangle$ and $|\psi\rangle$ are *distinct*, *nonorthogonal* states. This result is called the no-cloning theorem.

4.6 Mixed-state entanglement

The crucial property of quantum entanglement is that it cannot be created *locally*. Up to now in this chapter we have limited our attention to the properties of entangled *pure* states, but it is important to recognize that mixed states can be entangled, too.

Recall that a bipartite pure state $|\Psi\rangle_{AB}$ is *separable* if and only if it is a product state $|\Psi\rangle_{AB} = |\alpha\rangle_A \otimes |\beta\rangle_B$. We say that a bipartite mixed state ρ_{AB} is separable if and only if it can be realized as an ensemble of separable pure states,

$$\boldsymbol{\rho}_{AB} = \sum_{i} p_i \left(|\alpha_i\rangle \langle \alpha_i | \right)_A \otimes \left(|\beta_i\rangle \langle \beta_i | \right)_B \quad , \tag{4.112}$$

where the p_i 's are positive and sum to one. Alternatively, we may say that ρ_{AB} is separable if and only if it can be expressed as

$$\boldsymbol{\rho}_{AB} = \sum_{i,j} p_{ij} \boldsymbol{\rho}_{A,i} \otimes \boldsymbol{\rho}_{B,j} , \qquad (4.113)$$

4 Quantum Entanglement

where each $\rho_{A,i}$ and $\rho_{B,j}$ is a density operator, and the p_{ij} 's are positive and sum to one. Thus if a state is separable, the correlations between the state of part A and the state of part B are entirely classical, and embodied by the joint probability distribution p_{ij} . The two criteria eq. (4.112) and eq. (4.113) are equivalent because $\rho_{A,i}$ and $\rho_{B,j}$ can be realized as an ensemble of pure states.

Of course, it may be possible to realize a separable mixed state as an ensemble of entangled pure states as well. A simple example is that the random state $\rho = \frac{1}{4} \mathbf{I} \otimes \mathbf{I}$ of two qubits can be expressed as either

$$\boldsymbol{\rho} = \frac{1}{4} \left(|00\rangle \langle 00| + |01\rangle \langle 10| + |10\rangle \langle 01| + |11\rangle \langle 11| \right)$$
(4.114)

(an ensemble of product states) or

$$\boldsymbol{\rho} = \frac{1}{4} \left(|\phi^+\rangle \langle \phi^+| + |\phi^-\rangle \langle \phi^-| + |\psi^+\rangle \langle \psi^+| + |\psi^-\rangle \langle \psi^-| \right)$$
(4.115)

(an ensemble of maximally entangled states). The state is separable if and only if there is *some* way to represent is as an ensemble of product states. As for a pure state, if a mixed state is not separable, we say that it is *inseparable* or *entangled*.

Consider two distantly separated parties Alice and Bob who carry out a protocol involving local operations and classical communication. That is, Alice is permitted to perform quantum operations on her system A, Bob is permitted to perform quantum operations on his system B, and Alice and Bob are permitted to exchange classical bits as many times as they want. But no exchange of qubits is permitted. Then if Alice and Bob share a separable state to start with, their state will still be separable at the end of the protocol. The reason is that neither a local operation nor exchange of a classical bit can increase the Schmidt number of a bipartite pure state from the value 1 to a value greater than 1. Of course, Alice and Bob might have a mixed state, but in each step of the protocol an ensemble of product states is transformed to another ensemble of product states. Alice and Bob cannot create entanglement locally if they have none to begin with. In discussions of entanglement, the concept of a protocol that uses only Local Operations and Classical Communication is so prevalent that we will find it convenient to use the abbreviation *LOCC*.

On the other hand, with LOCC, Alice and Bob can prepare any separable state. To prepare ρ_{AB} in eq. (4.112), Alice generates random numbers to sample the probability distribution $\{p_i\}$; if outcome *i* is found, she informs Bob, and Alice prepares the $|\alpha_i\rangle_A$ while Bob prepares $|\beta_i\rangle_B$.

4.6 Mixed-state entanglement

4.6.1 Positive-partial-transpose criterion for separability

Now, consider a bipartite density operator ρ_{AB} on $\mathcal{H}_A \otimes \mathcal{H}_B$, presented as (say) a matrix in some basis. We would like to know whether ρ_{AB} is separable. How do we decide? It is not obvious how to devise an efficient algorithm that will definitively answer whether ρ_{AB} can be realized as an ensemble of product states. However, it is useful to note that there are necessary conditions for separability that are easy to check.

Recall that relative to a specified orthonormal basis $\{|i\rangle\}$ for a Hilbert space \mathcal{H} , a transpose operation T can be defined — the transpose acts on a basis for the linear operators according to

$$T:|i\rangle\langle j| \to (|i\rangle\langle j|)^T = |j\rangle\langle i|; \qquad (4.116)$$

its action on a matrix M_{ij} expressed in this basis is

$$\left(M^T\right)_{ij} = M_{ji} \ . \tag{4.117}$$

Evidently transposition preserves the trace of the matrix M. If M is Hermitian, then its transpose is its complex conjugate, which has the same (real) eigenvalues. Therefore, the transpose of a density operator is another density operator with the same eigenvalues — the transpose is a trace-preserving positive map.

But we saw in §3.?? that the transpose, although positive, is not completely positive; that is, the *partial transpose* $I \otimes T$ can map a bipartite positive operator to an operator that is not positive. For example, the maximally entangled state

$$|\Phi\rangle = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle_A \otimes |i\rangle_B \tag{4.118}$$

has density operator

$$\boldsymbol{\rho} = \frac{1}{N} \sum_{i,j} |ii\rangle \langle jj| . \qquad (4.119)$$

Its *partial transpose* is

$$(I \otimes T)(\boldsymbol{\rho}) = \frac{1}{N} \sum_{i,j} |ij\rangle\langle ji| = \frac{1}{N} (\text{SWAP}) \quad ; \quad (4.120)$$

the SWAP operator has eigenstates with eigenvalue +1 (symmetric states) and eigenstates with eigenvalue -1 (antisymmetric states) — hence it is not positive. We will use the notation

$$\boldsymbol{\rho}^{PT} = (I \otimes T)(\boldsymbol{\rho}) \tag{4.121}$$

4 Quantum Entanglement

for the partial transpose of the bipartite density operator ρ .

While the partial transpose is not a positive map in general, it *is* positive acting on separable states. The partial transpose of ρ_{AB} in eq. (4.113) is

$$\boldsymbol{\rho}_{AB}^{PT} = \sum_{i,j} p_{ij} \boldsymbol{\rho}_{A,i} \otimes \boldsymbol{\rho}_{B,j}^T ; \qquad (4.122)$$

since $\boldsymbol{\rho}_{B,j}^{T}$ is a density operator, so is $\boldsymbol{\rho}_{AB}^{PT}$. Thus we arrive at a useful necessary condition for separability.

Positive partial-transpose criterion for separability: If ρ_{AB} is separable, then ρ_{AB}^{PT} is nonnegative.

We will say that a bipartite density operator is PPT (for "positive partial transpose) if its partial transpose is nonnegative.

Thus, if we are presented with a density operator ρ_{AB} , we may compute the eigenvalues of ρ_{AB}^{PT} ; if negative eigenvalues are found, then ρ_{AB} is known to be inseparable. But because the PPT condition is necessary but not sufficient for separability, if ρ_{AB}^{PT} is found to be nonnegative, then whether ρ_{AB} is separable remains unsettled. The PPT criterion is sometimes called the *Peres-Horodecki* criterion for separability.

Let's apply the PPT criterion to a two-qubit state of the form

$$\boldsymbol{\rho}(\lambda) = \lambda |\phi^+\rangle \langle \phi^+| + \frac{1}{4}(1-\lambda)\boldsymbol{I} . \qquad (4.123)$$

This state may also be expressed as

$$\rho(F) = F |\phi^+\rangle \langle \phi^+| + \frac{1}{3}(1-F) \left(|\phi^-\rangle \langle \phi^-| + |\psi^+\rangle \langle \psi^+| + |\psi^-\rangle \langle \psi^-| \right), (4.124)$$

where $(1-F) = \frac{3}{4}(1-\lambda)$, and as we saw in §3.??, it results from subjecting half of the state $|\phi^+\rangle$ to the depolarizing channel with error probability p = 1 - F. This state is sometimes called a *Werner state* with fidelity *F*. Now

$$\left(|\phi^+\rangle\langle\phi^+|\right)^{PT} = \frac{1}{2}(\text{SWAP}) = \frac{1}{2}\boldsymbol{I} - |\psi^-\rangle\langle\psi^-| , \qquad (4.125)$$

where the second equality follows from the property that $|\phi^{\pm}\rangle, |\psi^{+}\rangle$ (which are symmetric) are eigenvalues of SWAP with eigenvalue 1, and $|\psi^{-}\rangle$ (which is antisymmetric) is an eigenstate of SWAP with eigenvalue -1. Since also $I^{PT} = I$, we see that the partial transpose of a Werner state is

$$\boldsymbol{\rho}(\lambda)^{PT} = \lambda \left(\frac{1}{2}\boldsymbol{I} - |\psi^{-}\rangle\langle\psi^{-}|\right) + \frac{1}{4}(1-\lambda)\boldsymbol{I} \\ = \frac{1}{4}(1+\lambda)\boldsymbol{I} - \lambda|\psi^{-}\rangle\langle\psi^{-}| . \qquad (4.126)$$

This operator has a negative eigenvalue if $\lambda > 1/3$, and we conclude that the Werner state is inseparable for $\lambda > 1/3$. Therefore, if half of the maximally entangled state $|\phi^+\rangle$ is subjected to the depolarizing channel with error probability p < 1/2, the resulting state remains entangled.

Although we won't prove it here, it turns out that for the case of twoqubit states, the PPT criterion is both necessary and sufficient for separability. Thus the Werner state with $\lambda < 1/3$ (or F < 1/2) is separable.

While we found that a bipartite pure state is entangled if and only if it violates some Bell inequality, this equivalence does not hold for mixed states. You will show in Exercise 4.?? that for a Werner state with $\lambda = 1/2$ (or any smaller value of λ) there is a local hidden-variable theory that fully accounts for the correlations between measurements of Alice's qubit and Bob's. Thus, Werner states with $1/3 < \lambda < 1/2$ are inseparable states that violate no Bell inequality.

Oddly, though a Werner state with $1/3 < \lambda < 1/2$ is not Bell-inequality violating, it is nonetheless a shared resource more powerful than classical randomness. You will also show in Exercise 4.?? that by consuming a Werner state Alice and Bob can teleport a qubit in an unknown state with fidelity

$$F_{\text{teleport}} = \frac{1}{2}(1+\lambda) \ . \tag{4.127}$$

This fidelity exceeds the maximal fidelity $F_{\text{teleport}} = 2/3$ that can be achieved without any shared entanglement, for any $\lambda > 1/3$ — that is, for any inseparable Werner state, whether Bell-inequality violating or not. Even if well described by local hidden variables, an entangled mixed state can be useful.

It seems rather strange that shared entangled states described by local hidden-variable theory should be a more powerful resource than classical shared randomness. Further observations to be discussed in §5.?? will deepen our grasp of the situation. There we will find that if Alice and Bob share many copies of the Werner state $\rho(\lambda)$ with $1/3 < \lambda < 1/2$, then while local hidden variables provide an adequate description of the correlations if Alice and Bob are restricted to measuring the pairs one at a time, violations of Bell inequalities still arise if they are permitted to perform more general kinds of measurements. These observations illustrate that mixed-state entanglement is a surprisingly subtle and elusive concept.

4.7 Nonlocality without entanglement

Quantum entanglement typifies the principle that there are bipartite quantum operations that cannot be implemented using only local operations and classical communication (LOCC). For example, if Alice and Bob share no prior entanglement, they cannot perform Bell measurement or prepare the entangled state $|\phi^+\rangle_{AB}$ unless they get together. Now we will encounter an interesting surprise: some things that Alice and Bob are unable to do with LOCC do *not* involve quantum entanglement, at least not directly.

Consider a game played by Alice, Bob, and Charlie. Charlie prepares a state $|\psi_i\rangle_{AB}$ selected from an ensemble of mutually orthogonal bipartite states, and distributes $|\psi_i\rangle_{AB}$ to Alice and Bob. To win the game, Alice and Bob must identify the state that Charlie prepared. Of course, if Alice and Bob were permitted to unite, they could perform an orthogonal measurement that would identify the state with certainty, and they would be able to win every time. But the rules of the game require Alice and Bob to stay separated, and they are forbidden to exchange quantum information — only LOCC is allowed. Thus, if Charlie's ensemble includes entangled states, Alice and Bob won't be able to win in general.

To make things easier for Alice and Bob, let's impose a new rule: Charlie is required to prepare a product state

$$|\psi\rangle_{AB} = |\alpha_i\rangle_A \otimes |\beta_i\rangle_B . \tag{4.128}$$

Now, since Alice has a pure state, and so does Bob, we might expect them to be able to devise a winning strategy. But on further reflection, this is not so obvious. Though the states $\{|\psi_i\rangle_{AB}\}$ in Charlie's ensemble are mutually orthogonal, the states $\{|\alpha_i\rangle_A\}$ that Alice could receive need not be mutually orthogonal, and the same is true of the states $\{|\beta_i\rangle_B\}$ that Bob could receive.

Indeed, even under the new rules, there is no winning strategy for Alice and Bob in general. Though Charlie sends a pure state to Alice and a pure state to Bob, there is no way for Alice and Bob, using LOCC, to fully decipher the message that Charlie has sent to them. This phenomenon is called *nonlocality without entanglement*.

The best way to understand nonlocality without entanglement is to consider an example. Suppose that Alice and Bob share a pair of *qutrits* (3-level quantum systems), and denote the three elements of an orthonormal basis for the qutrit by $\{|0\rangle, |1\rangle, |2\rangle\}$. In a streamlined notation, Charlie's ensemble of nine mutually orthogonal states is

$$\begin{aligned} |\psi\rangle_{1,2} &= |0, 0 \pm 1\rangle , \\ |\psi\rangle_{3,4} &= |0 \pm 1, 2\rangle , \\ |\psi\rangle_{5,6} &= |2, 1 \pm 2\rangle , \\ |\psi\rangle_{7,8} &= |1 \pm 2, 0\rangle , \\ |\psi\rangle_{9} &= |1, 1\rangle . \end{aligned}$$
(4.129)

(Here, $|0, 0 \pm 1\rangle$ denotes $|0\rangle_A \otimes \frac{1}{\sqrt{2}}(|0\rangle_B \pm |1\rangle_B)$, etc.) For ease of visualization, it is very convenient to represent this basis pictorially, as a tiling of a square by rectangles:



In the picture, the mutual orthogonality of the elements of Charlie's basis is reflected in the property that the rectangles are nonoverlapping.

When Charlie prepares one of these 9 states and distributes it, Alice receives one of the states

$$|0\rangle, |1\rangle, |2\rangle, |0\pm1\rangle, |1\pm2\rangle, \qquad (4.130)$$

and similarly for Bob. These states are *not* mutually orthogonal, and so cannot be perfectly distinguished by the recipient.

For example, Alice might perform an incomplete orthogonal measurement that distinguishes the state $|2\rangle$ from its orthogonal complement. Pictorially, this measurement "cuts" the square into two nonoverlapping parts. If Charlie prepared one of $|\psi_{5,6}\rangle$, $|\psi_{7,8}\rangle$, then Alice's outcome could be $|2\rangle\langle 2|$; in that case the state prepared by her measurement can be represented as:



After learning Alice's measurement outcome, Bob can perform an orthogonal measurement that projects on the basis

$$\{|0\rangle, |1+2\rangle, |1-2\rangle\}$$
. (4.131)

If his outcome is $|1 + 2\rangle$ or $|1 - 2\rangle$, then Alice and Bob have successfully identified Charlie's state as $|\psi_5\rangle$ or $|\psi_6\rangle$. But if Bob's outcome is $|0\rangle$, then Alice and Bob remain uncertain whether Charlie prepared $|\psi_7\rangle$ of $|\psi_8\rangle$. On the other hand, if Charlie prepared one of $|\psi_{1,2}\rangle$, $|\psi_{3,4}\rangle$, $|\psi_{7,8}\rangle$, $|\psi_9\rangle$, then Alice's outcome could be $|0\rangle\langle 0| + |1\rangle\langle 1|$; in that case the state prepared by her measurement can be represented as:



Once again, Alice and Bob have lost any hope of distinguishing $|\psi_7\rangle$ from $|\psi_8\rangle$, but in a few more rounds of LOCC, they can successfully identify any of the other five states. Bob projects onto $|2\rangle$ or its complement; if he finds $|2\rangle\langle 2|$, then Alice projects onto $|0 \pm 1\rangle$ to complete the protocol. If Bob's outcome is $|0\rangle\langle 0| + |1\rangle\langle 1|$, then Alice projects onto $\{|0\rangle, |1\rangle\}$; finally Bob measures in either the $|0\pm 1\rangle$ basis (if Alice found $|0\rangle$) or the $\{|0\rangle, |1\rangle\}$ basis (if Alice found $|1\rangle$).

By choosing one of nine mutually orthogonal product states, Charlie has sent two trits of classical information to Alice and Bob. But their LOCC protocol, which fails to distinguish $|\psi_7\rangle$ from $|\psi_8\rangle$, has not been able to recover all of the information in Charlie's message. Of course, this is just one possible protocol, but one can prove (we won't here) that no LOCC protocol can extract two trits of classical information. The trouble is that with LOCC, Alice and Bob cannot fully "dissect" the square into nonoverlapping rectangles. This is nonlocality without entanglement.

4.8 Multipartite entanglement

Up until now, we have mostly limited our attention to quantum states shared by two parties. We will conclude this chapter with some observations about the properties of entanglement shared by three or more parties: *multipartite entanglement*.

Consider for example the case of a pure state $|\psi\rangle_{A_1,A_2,\ldots,A_n}$ shared by n parties A_1, A_2, \ldots, A_n , and suppose that there is no way to divide the parties into two smaller camps, where each camp shares a pure state.

Thus the state is entangled, and furthermore, it can't be expressed as a product of states each involving fewer than n parties. Hence we might say that the state exhibits n-party entanglement. If the parties start out with an n-fold product state $|\psi_1\rangle_{A_1} \otimes |\psi_2\rangle_{A_2} \otimes \cdots |\psi_n\rangle_{A_n}$, then there is no way for them to assemble the state $|\psi\rangle_{A_1,A_2,\ldots,A_n}$ using LOCC alone — quantum communication is required. Indeed, no matter how we divide the n parties into two subsystems A and B, quantum communication between A and B is needed.

What if we disallow quantum communication, but we do equip the parties with *pairwise* entanglement that has been established in advance? Then for the purpose of constructing the state $|\psi\rangle_{A_1,A_2,...A_n}$, it clearly sufficies for the first party A_1 to share bipartite entanglement with each of the other n-1 parties. Party A_1 can build the state $|\psi\rangle_{A_1,A_2,...A_n}$ in her own laboratory, and then teleport the corresponding share of the state to each of the n-1 other parties. In this sense, then, bipartite entanglement and LOCC is as powerful a resource as multiparty entanglement.

Nonetheless, multipartite entangled states exhibit some qualitatively new phenomena that we don't encounter in the study of bipartite entanglement, such as nonprobabilistic tests of Einstein locality, and entanglementenhanced multiparty communication.

4.8.1 Three quantum boxes

In the wake of the wildly successful experiment with the three coins on the table, Alice and Bob are now world famous. They are both tenured professors, Alice at Caltech, and Bob at Chicago. They are much too important to spend much time in the lab, but they have many graduate students and remain scientifically active.

Their best student, Charlie, who did all the hard work on the coin experiment, has graduated and is now an assistant professor at Princeton. Alice and Bob would like to nurture Charlie's career, and help him earn tenure. One day, Alice and Bob are chatting on the phone:

- Alice: You know, Bob, we really ought to help Charlie. Can you think of a neat experiment that the three of us can do together?
- **Bob**: Well, I dunno, Alice. There are a lot of experiments I'd like to do with our entangled pairs of qubits. But in each experiment, there's one qubit for me and one for you. It looks like Charlie's the odd man out.
- Alice: [Long pause] Bob Have you ever thought of doing an experiment with three qubits?

4 Quantum Entanglement

Bob's jaw drops and his pulse races. In a sudden epiphany, his whole future career seems mapped out before him. Truth be told, Bob was beginning to wonder if pairs of qubits were getting to be old hat. Now he knows that for the next five years, he will devote himself slavishly to performing the definitive three-qubit experiment. By that time, he, Alice, and Charlie will have trained another brilliant student, and will be ready for a crack at four qubits. Then another student, and another qubit. And so on to retirement.

Here is the sort of three-qubit experiment that Alice and Bob decide to try: Alice instructs her technician in her lab at Caltech to prepare carefully a state of three quantum boxes. (But Alice doesn't know exactly how the technician does it.) She keeps one box for herself, and she ships the other two by quantum express, one to Bob and one to Charlie. Each box has a ball inside that can be either black or white, but the box is sealed tight shut. The only way to find out what is inside is to open the box, but there are two different ways to open it — the box has two doors, clearly marked X and Y. When either door opens, a ball pops out whose color can be observed. It isn't possible to open both doors at once.

Alice, Bob, and Charlie decide to study how the boxes are correlated. They conduct many carefully controlled trials. Each time, one of the three, chosen randomly, opens door X, while the other two open door Y. Lucky as ever, Alice, Bob, and Charlie make an astonishing discovery. They find that every single time they open the boxes this way, the number of black balls they find is *always* odd.

That is, Alice, Bob and Charlie find that when they open door X on one box and door Y on the other two, the colors of the balls in the boxes are guaranteed to be one of

$$0_A 0_B 1_C , \quad 0_A 1_B 0_C , \quad 1_A 0_B 0_C , \quad 1_A 1_B 1_C ,$$

$$(4.132)$$

(0 for white, 1 for black); They never see any of

$$1_A 1_B 0_C$$
, $1_A 0_B 1_C$, $0_A 1_B 1_C$, $0_A 0_B 0_C$.
(4.133)

It makes no difference which of the three boxes is opened through door X.

After a while, Alice, Bob, and Charlie catch on that after opening two of the boxes, they can always predict what will happen before they open the third box. If the first two balls are the same color, the last ball is sure to be black, and if the first two are different colors, the last ball is sure to be white. They've tried it a zillion times, and it always works! Even after all the acclaim showered upon the three-coin experiment, Alice, Bob, and Charlie have never quite shaken their attachment to Einstein locality. One day they are having a three-way conference call:

- Alice: You know, guys, sometimes I just can't decide whether to open door X or door Y of my box. I know I have to choose carefully ... If I open door X, that's sure to disturb the box; so I'll never know what would have happened if I had opened door Y instead. And if I open door Y, I'll never know what I would have found if I had opened door X. It's frustrating!
- **Bob**: Alice, you're so wrong! Our experiment shows that you can have it both ways. Don't you see? Let's say that you want to know what will happen when you open door X. Then just ask Charlie and me to open door Y of our boxes and to tell you what we find. You'll know absolutely for sure, without a doubt, what's going to happen when you open door X. We've tested that over and over again, and it always works. So why bother to open door X? You can go ahead and open door Y instead, and see what you find. That way, you really do know the result of opening both doors!
- **Charlie**: But how can you be sure? If Alice opens door Y, she passes up the opportunity to open door X. She can't really ever have it both ways. After she opens door Y, we can never check whether opening door X would have given the result we expected.
- **Bob**: Oh come on, how can it be otherwise? Look, you don't really believe that what you do to your box in Princeton and I do to mine in Chicago can exert any *influence* on what Alice finds when she opens her box in Pasadena, do you? When we open our boxes, we can't be changing anything in Alice's box; we're just finding the information we need to predict with certainty what Alice is going to find.
- **Charlie**: Well, maybe we should do some more experiments to find out if you're right about that.

Indeed, the discovery of the three-box correlation has made Alice and Bob even more famous than before, but Charlie hasn't gotten the credit he deserves — he still doesn't have tenure. No wonder he wants to do more experiments! He continues:

Charlie: Here's something we can try. In all the experiments we've done up to now, we have always opened door Y on two boxes and door X on the other box. Maybe we should try something different.

Like, maybe we should see what happens if we open the same door on all three boxes. We could try opening three X doors.

- **Bob**: Oh, come on! I'm tired of three boxes. We already know all about three boxes. It's time to move on, and I think Diane is ready to help out. Let's do four boxes!
- Alice: No, I think Charlie's right. We can't really say that we know everything there is to know about three boxes until we've experimented with other ways of opening the doors.
- **Bob**: Forget it. They'll never fund us! After we've put all that effort into opening two Y's and an X, now we're going to say we want to open three X's? They'll say we've done whiffnium and now we're proposing whaffnium ... We'll sound ridiculous!
- Alice: Bob has a point. I think that the only way we can get funding to do this experiment is if we can make a prediction about what will happen. Then we can say that we're doing the experiment to test the prediction. Now, I've heard about some theorists named Greenberger, Horne, Zeilinger, and Mermin (GHZM). They've been thinking a lot about our three-box experiments; maybe they'll be able to suggest something.
- **Bob**: Well, these boxes are my life, and they're just a bunch of theorists. I doubt that they'll have anything interesting or useful to say. But I suppose it doesn't really matter whether their theory makes any sense ... If we can test it, then even I will accept that we have a reason for doing another three-box experiment.

And so it happens that Alice, Bob, and Charlie make the pilgrimage to see GHZM. And despite Bob's deep skepticism, GHZM make a very interesting suggestion indeed:

GHZM: Bob says that opening a box in Princeton and a box in Chicago can't possibly have any influence on what will happen when Alice opens a box in Pasadena. Well, let's suppose that he's right. Now you guys are going to do an experiment in which you all open your X doors. No one can say what's going to happen, but we can reason this way: Let's just assume that if you had opened three Y doors, you would have found three white balls. Then we can use Bob's argument to see that if you open three X doors instead, you will have to find three black balls. It goes like this: if Alice opens X, Bob opens Y, and Charlie opens Y, then you know for certain that the number of black balls has to be odd. So, if we know that Bob and Charlie both would find white when they open door Y, then Alice has to find black when she opens door X. Similarly, if Alice and Charlie both would find white when they open Y, then Bob has to find black when he opens X, and if Alice and Bob both would find white when they open Y, then Charlie must find black when he opens X. So we see that[§]

$$Y_A Y_B Y_C = 000 \longrightarrow X_A X_B X_C = 111 . \tag{4.134}$$

Don't you agree?

- **Bob**: Well, maybe that's logical enough, but what good is it? We don't know what we're going to find inside a box until we open it. You've assumed that we know $Y_A Y_B Y_C = 000$, but we never know that ahead of time.
- **GHZM**: Sure, but wait. Yes, you're right that we can't know ahead of time what we would find if we opened door Y on each box. But there are only eight possibilities for three boxes, and we can easily list them all. And for each of those eight possibilities for $Y_A Y_B Y_C$ we can use the same reasoning as before to infer the value of $X_A X_B X_C$. We obtain a table, like this:

$$Y_A Y_B Y_C = 000 \longrightarrow X_A X_B X_C = 111$$

$$Y_A Y_B Y_C = 001 \longrightarrow X_A X_B X_C = 001$$

$$Y_A Y_B Y_C = 010 \longrightarrow X_A X_B X_C = 010$$

$$Y_A Y_B Y_C = 100 \longrightarrow X_A X_B X_C = 100$$

$$Y_A Y_B Y_C = 011 \longrightarrow X_A X_B X_C = 100$$

$$Y_A Y_B Y_C = 101 \longrightarrow X_A X_B X_C = 010$$

$$Y_A Y_B Y_C = 110 \longrightarrow X_A X_B X_C = 001$$

$$Y_A Y_B Y_C = 111 \longrightarrow X_A X_B X_C = 111$$

$$(4.135)$$

Bob: Okay, but so what?

GHZM: There's something interesting about the table, Bob! Look at the values for $X_A X_B X_C \ldots$ Every single entry has an *odd* number of 1's. That's our prediction: when you all open door X on your boxes, you'll always find an odd number of black balls! Could be one, or could be three, but always *odd*.

Naturally, Alice, Bob, and Charlie are delighted by this insight from GHZM. They proceed to propose the experiment, which is approved and

 $[\]S$ Here 0 stands for white and 1 stands for black; Y_A is what Alice finds when she opens door Y on her box, and so on.

generously funded. Finally the long awaited day arrives when they are to carry out the experiment for the first time. And when Alice, Bob, and Charlie each open door X on their boxes, can you guess what they find? Three white balls. Whaaaa??!!

Suspecting an error, Alice and Bob and Charlie repeat the experiment, very carefully, over and over and over again. And in every trial, every single time, they find an even number of black balls when they open door X on all three boxes. Sometimes none, sometimes two, but never one and never three. What they find, every single time, is just the opposite of what GHZM had predicted would follow from the principle of Einstein locality!

Desperation once again drives Alice, Bob, and Charlie into the library, seeking enlightenment. After some study of a quantum mechanics textbook, and a thorough interrogation of Alice's lab technician, they realize that their three boxes had been prepared in a GHZM quantum state

$$|\psi\rangle_{ABC} = \frac{1}{\sqrt{2}} \left(|000\rangle_{ABC} + |111\rangle_{ABC}\right) , \qquad (4.136)$$

a simultaneous eigenstate with eigenvalue one of the three observables

$$Z_A \otimes Z_B \otimes I_C$$
, $I_A \otimes Z_B \otimes Z_C$, $X_A \otimes X_B \otimes X_C$.
(4.137)

And since ZX = iY, they realize that this state has the properties

$$Y_A \otimes Y_B \otimes X_C = -1$$

$$X_A \otimes Y_B \otimes Y_C = -1$$

$$Y_A \otimes X_B \otimes Y_C = -1$$

$$X_A \otimes X_B \otimes X_C = 1.$$
(4.138)

In opening the box through door X or door Y, Alice, Bob, and Charlie are measuring the observable X or Y, where the outcome 1 signifies a white ball, and the outcome -1 a black ball. Thus if the three qubit state eq. (4.136) is prepared, eq. (4.138) says that an odd number of black balls will be found if door Y is opened on two boxes and door X on the third, while an even number of black balls will be found if door X is opened on all three boxes. This behavior, unambiguously predicted by quantum mechanics, is just what had seemed so baffling to Alice, Bob, and Charlie, and to their fellow die-hard advocates of Einstein locality.

After much further study of the quantum mechanics textbook, Alice, Bob, and Charlie gradually come to recognize the flaw in their reasoning. They learn of Bohr's principle of complementarity, of the irreconcilable incompatibility of noncommuting observables. And they recognized that to arrive at their prediction, they had *postulated* an outcome for the measurement of YYY, and then proceeded to infer the consequences for a measurement of XXX. By failing to heed the insistent admonitions of Niels Bohr, they had fallen prey to the most pernicious of fallacies.

As they had hoped, the experiment of the three boxes brings even further acclaim to Alice and Bob, and tenure to Charlie. Of course, the three-coin experiment had already convincingly struck down Einstein locality; even so, the three-box experiment had a different character. In the coin experiment, Alice and Bob could uncover any two of the three coins, finding any one of four possible configurations: HH, HT, TH, TT. Only by carrying out many trials could they amass a convincing statistical case for the violation of the Bell inequality. In contrast, in the three-box experiment, Alice, Bob, and Charlie had found a result inconsistent with Einstein locality in every single trial in which they opened door X on all three boxes!

4.8.2 Cat states

The GHZM state studied by Alice, Bob, and Charlie is a natural threequbit generalization of the maximally entangled Bell pair. A Bell pair can be characterized as the simultaneous eigenstate of the two commuting operators ZZ, whose eigenvalue is the "parity bit" of the pair, and XX, whose eigenvalue is the phase bit. (Here we use a compressed notation in which the tensor product symbol \otimes is suppressed — *e.g.*, XX denotes the operator that simultaneously applies X to both Alice's qubit and Bob's.) The GHZM state is the simultaneous eigenstate of ZZI, IZZ, and XXX.

An n-qubit generalization of the GHZM state can be defined, which is the simultaneous eigenstate of the n commuting operators

$$ZZIII \dots I,$$

$$IZZII \dots I,$$

$$IIZZI \dots I,$$

$$III \dots IZZ,$$

$$XX \dots XX.$$

$$(4.139)$$

Each such simultaneous eigenstate has the form

$$\frac{1}{\sqrt{2}} \left(|x\rangle \pm |\neg x\rangle \right) \,, \tag{4.140}$$

where $\neg x$ denotes the complement of the binary string x. Since for large n this state is a coherent superposition of two "macroscopically distinguish-

able" states, it is called an *n*-qubit *cat state*, in homage to Schrödinger's cat. The *n*-qubit cat state has n - 1 parity bits, and just one phase bit.

Some noteworthy properties of cat states are:

- Each qubit is maximally entangled with the rest. That is, if we trace over the other n-1 qubits, the qubit's density operator is $\rho = \frac{1}{2}I$. For this reason, it is sometimes said that a cat state is a maximally entangled state of n qubits.
- But this is a rather misleading locution. Because its parity and phase bits are treated quite asymmetrically, the cat is not so profoundly entangled as some other multiqubit states that we will encounter in Chapter 7. For example, for the cat state with x = 000...0, if we trace over n - 2 qubits, the density operator of the remaining two is

$$\boldsymbol{\rho}_{2-\text{qubit}} = \frac{1}{2} \Big(|00\rangle \langle 00| + |11\rangle \langle 11| \Big) , \qquad (4.141)$$

which has rank two rather than four. Correspondingly, we can acquire a bit of information about a cat state (one of its parity bits) by observing only two of the qubits in the state. Other multiqubit states, which might be regarded as more highly entangled than cat states, have the property that the density operator of two (or more) qubits is proportional to the identity, if we trace over the rest.

- Suppose that Charlie prepares one of the 2ⁿ possible cat states and distributes it to n parties. Using LOCC, the parties can determine all n 1 parity bits of the state each pary measures Z and all broadcast their results. But by measuring Z they destroy the phase bit. Alternatively, they can all measure X to determine the phase bit, but at the cost of destroying all the parity bits.
- Each party, by applying one of $\{I, X, Y, Z\}$ can transform a given cat state to any one of four other cat states; that is, the party can modify the phase bit and one of the n-1 parity bits. All nparties, working together, can transform one cat state to any one of the 2^n mutually orthogonal cat states; for example, one party can manipulate the phase bit while each of the the other n-1 parties controls a parity bit.
- If the parties unite, the phase bit and all parity bits can be simultaneously measured.

If the parties start out with a product state, the three-qubit cat state (for example) can be prepared by executing the quantum circuit:



For the *n*-party case, a similar circuit with n - 1 CNOT gates does the job. Thus, to prepare the state, it suffices for the first party to visit each of the other n - 1 parties. By running the circuit in reverse, a cat state can be transformed to a product state that can be measured locally.

4.8.3 Entanglement-enhanced communication

An intriguing property of the n-qubit cat state is that its phase bit can be manipulated by each one of the n parties that share the state. One wonders how this shared resource might be exploited.

We will describe a setting in which possession of a cat state reduces the amount of communication that is required to accomplish a distributed information processing task. Suppose that each one of n parties labeled by index $i = 1, 2, 3, \ldots, n$ resides on a separate planet, and that party i possesses some data (a string of bits x_i) known only to that party. The goal of the parties is to compute a function f (with a one-bit output) that depends on all the data:

$$f(x_1, x_2, x_3, \dots, x_n) \in \{0, 1\} . \tag{4.142}$$

In this universe, computation is cheap, and communication is expensive. Each party has unlimited computational power at her disposal, but since no party knows the full input of the function f, no one can compute f unless the parties communicate. For this purpose, they are equipped with a broadcast channel — if any party speaks, all the others can hear her. However, use of the broadcast channel is very expensive, so that the parties wish to compute f while making minimal use of the channel.

With this motivation, we define the classical communication complexity CCC[f] of the function f:

CCC[f] = the minimum of bits that must be broadcast (in the worst case) for all the parties to know the value of $f(x_1, x_2, x_3, \ldots, x_n)$.

Here "in the worst case" means that we maximize the number of bits of communication required over all possible values for the input strings $x_1, x_2, x_3, \ldots, x_n$.

4 Quantum Entanglement

We are interested in whether using quantum information can reduced the amount of communication required to compute a function. Hence we contrast the function's classical communication complexity with its quantum communication complexity. There are actually several different natural ways to generalize a classical communication setting to a quantum setting. In one, to which we return in Chapter 6, the parties are allowed to exchange qubits instead of classical bits. Here, we consider a scenario in which all communication is via the classical broadcast channel, but the parties are allowed to share entangled states that have been prepared in advance, and to manipulate their shared entanglement locally. Thus we define the quantum communication complexity QCC[f] as

QCC[f] = the minimum of bits that must be broadcast (in the worst case) for all the parties to know the value of $f(x_1, x_2, x_3, \ldots, x_n)$, where the parties are permitted to share prior quantum entanglement.

One way to argue that multipartite entanglement can be a useful resource is to establish that there are functions f such that

$$QCC[f] < CCC[f] . (4.143)$$

Here is an example of such a function: Each party holds an m-bit string, and they are to compute

$$\sum_{i=1}^{n} x_i \pmod{2^m} . \tag{4.144}$$

Except that they have been promised that the answer is either 0 or 2^{m-1} ; therefore, their function has just a one-bit output.

First consider what strategy the parties should play if they share no entanglement. Suppose that parties 2 through n broadcast their data, and that the first party computes f and broadcasts the result. But note that it is not necessary for the parties to broadcast all of their bits, since some of the bits cannot affect the answer. Indeed, the k least significant bits are irrelevant as long as

$$(n-1)\left(2^k-1\right) < 2^{m-1} , \qquad (4.145)$$

which is satisfied provided that

$$(n-1)2^k \le 2^{m-1} . (4.146)$$

It suffices then, for parties 2 through n to broadcast their m - k most significant bits, where

$$m-k \ge \log_2(n-1)+1$$
; (4.147)

including one more bit for the first party to broadcast the answer, we conclude that

$$CCC[f] \le (n-1) \Big(\log_2(n-1) + 1 \Big) + 1$$
. (4.148)

In fact, this protocol is close to optimal — it can be proved that

$$CCC[f] > n \log_2 n - n$$
 . (4.149)

But the amount of communication required can be reduced if the parties share an *n*-qubit cat state, for they can imprint the answer on their shared phase bit! Each applies to her qubit the transformation

$$\begin{array}{l} |0\rangle \rightarrow |0\rangle , \\ |1\rangle \rightarrow e^{2\pi i (x_i/2^m)} |1\rangle . \end{array}$$

$$(4.150)$$

Thus the cat state

$$|\text{cat}\rangle = \frac{1}{\sqrt{2}} \Big(|000\dots0\rangle + |111\dots1\rangle \Big) \tag{4.151}$$

is transformed to

$$|\operatorname{cat}'\rangle = \frac{1}{\sqrt{2}} \Big(|000\dots0\rangle + \eta|111\dots1\rangle\Big) , \qquad (4.152)$$

where the phase η is

$$\eta = \exp\left(2\pi i \left(\sum_{i=1}^{n} x_i\right)/2^m\right) = (-1)^{f(x_1, x_2, \dots, x_n)} .$$
(4.153)

Thus the fu

4.9 Manipulating entanglement

4.10 Summary

Summary 1. Summary 2. Summary 3.

4.11 Bibliographical notes

4.12 Exercises

4.1 Hardy's theorem

Bob (in Boston) and Claire (in Chicago) share many identically prepared copies of the two-qubit state

$$|\psi\rangle = \sqrt{(1-2x)} |00\rangle + \sqrt{x} |01\rangle + \sqrt{x} |10\rangle , \qquad (4.154)$$

where x is a real number between 0 and 1/2. They conduct many trials in which each measures his/her qubit in the basis $\{|0\rangle, |1\rangle\}$, and they learn that if Bob's outcome is 1 then Claire's is always 0, and if Claire's outcome is 1 then Bob's is always 0.

Bob and Claire conduct further experiments in which Bob measures in the basis $\{|0\rangle, |1\rangle\}$ and Claire measures in the orthonormal basis $\{|\varphi\rangle, |\varphi^{\perp}\rangle\}$. They discover that if Bob's outcome is 0, then Claire's outcome is always φ and never φ^{\perp} . Similarly, if Claire measures in the basis $\{|0\rangle, |1\rangle\}$ and Bob measures in the basis $\{|\varphi\rangle, |\varphi^{\perp}\rangle\}$, then if Claire's outcome is 0, Bob's outcome is always φ and never φ^{\perp} .

a) Express the basis $\{|\varphi\rangle, |\varphi^{\perp}\rangle\}$ in terms of the basis $\{|0\rangle, |1\rangle\}$.

Bob and Claire now wonder what will happen if they both measure in the basis $\{|\varphi\rangle, |\varphi^{\perp}\rangle\}$. Their friend Albert, a firm believer in local realism, predicts that it is impossible for both to obtain the outcome φ^{\perp} (a prediction known as *Hardy's theorem*). Albert argues as follows:

When both Bob and Claire measure in the basis $\{|\varphi\rangle, |\varphi^{\perp}\rangle\}$, it is reasonable to consider what might have happened if one or the other had measured in the basis $\{|0\rangle, |1\rangle\}$ instead.

So suppose that Bob and Claire both measure in the basis $\{|\varphi\rangle, |\varphi^{\perp}\rangle\}$, and that they both obtain the outcome φ^{\perp} . Now if Bob had measured in the basis $\{|0\rangle, |1\rangle\}$ instead, we can be certain that his outcome would have been 1, since experiment has shown that if Bob had obtained 0 then Claire could not have obtained φ^{\perp} . Similarly, if Claire had measured in the basis $\{|0\rangle, |1\rangle\}$, then she certainly would have obtained the outcome 1. We conclude that if Bob and Claire both measured in the basis $\{|0\rangle, |1\rangle\}$, both would have obtained the outcome 1. But this is a contradiction, for experiment has shown that it is not possible for both Bob and Claire to obtain the outcome 1 if they both measure in the basis $\{|0\rangle, |1\rangle\}$.

We are therefore forced to conclude that if Bob and Claire both measure in the basis $\{|\varphi\rangle, |\varphi^{\perp}\rangle\}$, it is impossible for both to obtain the outcome φ^{\perp} . Though impressed by Albert's reasoning, Bob and Claire decide to investigate what predictions can be inferred from quantum mechanics.

- b) If Bob and Claire both measure in the basis $\{|\varphi\rangle, |\varphi^{\perp}\rangle\}$, what is the quantum-mechanical prediction for the probability P(x)that both obtain the outcome φ^{\perp} ?
- c) Find the "maximal violation" of Hardy's theorem: show that the maximal value of P(x) is $P[(3-\sqrt{5})/2] = (5\sqrt{5}-11)/2 \approx .0902$.
- d) Bob and Claire conduct an experiment that confirms the prediction of quantum mechanics. What was wrong with Albert's reasoning?

4.2 Closing the detection loophole

Recall that the CHSH inequality

$$|\langle ab\rangle + \langle a'b\rangle + \langle ab'\rangle - \langle a'b'\rangle| \le 2 \tag{4.155}$$

holds if the random variables a, b, a'b' take values ± 1 and are governed by a joint probability distribution. The maximal violation of this inequality by the quantum-mechanical predictions occurs when the left-hand-side is $2\sqrt{2}$, which is achieved if Alice and Bob share the maximally entangled state $|\phi^+\rangle$, a, a' are measurements of Alice's qubit along axes \hat{x} and \hat{z} , and b, b' are measurements of Bob's qubit along axes $(\hat{x} + \hat{z})/\sqrt{2}$ and $(\hat{x} - \hat{z})/\sqrt{2}$.

Alice and Bob have done a beautiful experiment measuring the polarizations of entangled photon pairs, and have confirmed the CHSH inequality violation predicted by quantum mechanics. Albert is skeptical. He points out that the detectors used by Alice and Bob in their experiment are not very efficient. Usually, when Alice detects a photon, Bob does not, and when Bob detects a photon, Alice does not. Therefore, they discard the data for most of the photon pairs, and retain the results only in the case when two photons are detected in coincidence. In their analysis of the data, Alice and Bob assume that their results are based on a fair sample of the probability distribution governing the measured variables. But Albert argues that their conclusions could be evaded if *whether* a photon is detected is correlated with the *outcome* of the polarization measurement.

Alice and Bob wonder how much they will need to improve their detector efficiency to do an experiment that will impress Albert.

Alice can choose to orient her detector along any axis, and if she aligns the detector with the axis a, then ideally the detector will

4 Quantum Entanglement

click when her qubit's spin is pointing up along a, but because of detector inefficiencies it sometimes fails to click even though the qubit points up. For pair number i, let $x_i \in \{0, 1\}$ be a variable indicating whether Alice's detector would click when aligned with a — if there would be a click then $x_i = 1$, and if there would be no click then $x_i = 0$. Since the detector is imperfect, x_i may be 0 even though the qubit points up along a. Similarly, $x'_i \in \{0, 1\}$ indicates whether Alice's detector would click if aligned with $a', y_i \in \{0, 1\}$ indicates whether Bob's detector would click if aligned with b and $y'_i \in \{0, 1\}$ indicates whether Bob's detector would click if aligned with b and $y'_i \in \{0, 1\}$ indicates whether Bob's detector would click if aligned with b and $y'_i \in \{0, 1\}$ indicates whether Bob's detector would click if aligned with b and $y'_i \in \{0, 1\}$ indicates whether Bob's detector would click if aligned with b and $y'_i \in \{0, 1\}$ indicates whether Bob's detector would click if aligned with b and $y'_i \in \{0, 1\}$ indicates whether Bob's detector would click if aligned with b and $y'_i \in \{0, 1\}$ indicates whether Bob's detector would click if aligned with b.

Alice and Bob are free to decide how to align their detectors in each measurement; therefore they can fairly sample the values of x, x', y, y' and infer from their measurements the values of

$$P_{++}(ab) = N^{-1} \sum_{i=1}^{N} x_i y_i ,$$

$$P_{++}(a'b) = N^{-1} \sum_{i=1}^{N} x'_i y_i ,$$

$$P_{++}(ab') = N^{-1} \sum_{i=1}^{N} x_i y'_i ,$$

$$P_{++}(a'b') = N^{-1} \sum_{i=1}^{N} x'_i y'_i ,$$
(4.156)

where N is the total number of pairs tested. Here $e.g P_{++}(ab)$ is the probability that both detectors click when Alice and Bob orient their detectors along a and b (including the effects of detector inefficiency).

a) If $x, x', y, y' \in \{0, 1\}$, show that

$$xy + xy' + x'y - x'y' \le x + y . \tag{4.157}$$

b) Show that

$$P_{++}(ab) + P_{++}(a'b) + P_{++}(ab') - P_{++}(a'b') \le P_{+}(a) + P_{+}(b) ;$$
(4.158)

here $P_{+.}(a)$ denotes the probability that Alice's detector clicks if oriented along a, and $P_{+}(b)$ denotes the probability that Bob's detector clicks if oriented along b. 4.12 Exercises

c) Now compare with the predictions of quantum mechanics, where Alice's detector has efficiency η_A and Bob's detector has efficiency η_B . This means that Alice's detector clicks with probability $P = \eta_A P_{\text{perf}}$, where P_{perf} is the probability of a click if her detector were perfect, and similarly for Bob. Choosing the a, a', b, b' that maximally violate the CHSH inequality, show that the quantum-mechanical predictions violate eq. (4.158) only if

$$\frac{\eta_A \eta_B}{\eta_A + \eta_B} > \frac{1}{1 + \sqrt{2}}$$
 (4.159)

Thus, if $\eta_A = \eta_B$, Alice and Bob require detectors with efficiency above 82.84% to overcome Albert's objection.

4.3 Teleportation with continuous variables

One complete orthonormal basis for the Hilbert space of two particles on the real line is the (separable) position eigenstate basis $\{|q_1\rangle \otimes |q_2\rangle\}$. Another is the entangled basis $\{|Q, P\rangle\}$, where

$$|Q,P\rangle = \frac{1}{\sqrt{2\pi}} \int dq \ e^{iPq} |q\rangle \otimes |q+Q\rangle \ ; \tag{4.160}$$

these are the simultaneous eigenstates of the relative position operator $Q \equiv q_2 - q_1$ and the total momentum operator $P \equiv p_1 + p_2$.

a) Verify that

$$\langle Q', P' | Q, P \rangle = \delta(Q' - Q)\delta(P' - P) .$$
(4.161)

b) Since the states $\{|Q, P\rangle\}$ are a basis, we can expand a position eigenstate as

$$|q_1\rangle \otimes |q_2\rangle = \int dQdP \; |Q,P\rangle \langle Q,P| \left(|q_1\rangle \otimes |q_2\rangle\right) \;. \tag{4.162}$$

Evaluate the coefficients $\langle Q, P | (|q_1\rangle \otimes |q_2\rangle)$.

c) Alice and Bob have prepared the entangled state $|Q, P\rangle_{AB}$ of two particles A and B; Alice has kept particle A and Bob has particle B. Now Alice has received an unknown singleparticle wavepacket $|\psi\rangle_C = \int dq \ |q\rangle_C \ C\langle q|\psi\rangle_C$ that she intends to teleport to Bob. Design a protocol that they can execute to achieve the teleportation. What should Alice measure? What information should she send to Bob? What should Bob do when he receives this information, so that particle B will be prepared in the state $|\psi\rangle_B$?

4.4 Teleportation with mixed states.

An operational way to define entanglement is that an entangled state can be used to teleport an unknown quantum state with better fidelity than could be achieved with local operations and classical communication only. In this exercise, you will show that there are mixed states that are entangled in this sense, yet do not violate any Bell inequality. Hence, for mixed states (in contrast to pure states) "entangled" and "Bell-inequality-violating" are not equivalent.

Consider a "noisy" entangled pair with density matrix.

$$\boldsymbol{\rho}(\lambda) = (1-\lambda)|\psi^{-}\rangle\langle\psi^{-}| + \lambda \frac{1}{4}\mathbf{1}.$$
(4.163)

- a) Find the fidelity F that can be attained if the state $\rho(\lambda)$ is used to teleport a qubit from Alice to Bob. [Hint: Recall that you showed in an earlier exercise that a "random guess" has fidelity $F = \frac{1}{2}$.]
- b) For what values of λ is the fidelity found in (a) better than what can be achieved if Alice measures her qubit and sends a classical message to Bob? [Hint: Earlier, you showed that F = 2/3 can be achieved if Alice measures her qubit. In fact this is the best possible F attainable with classical communication.]
- c) Compute

$$\operatorname{Prob}(\uparrow_{\hat{n}}\uparrow_{\hat{m}}) \equiv \operatorname{tr}\left(\boldsymbol{E}_{A}(\hat{n})\boldsymbol{E}_{B}(\hat{m})\boldsymbol{\rho}(\lambda)\right),$$
(4.164)

where $\boldsymbol{E}_A(\hat{n})$ is the projection of Alice's qubit onto $|\uparrow_{\hat{n}}\rangle$ and $\boldsymbol{E}_B(\hat{m})$ is the projection of Bob's qubit onto $|\uparrow_{\hat{m}}\rangle$.

d) Consider the case $\lambda = 1/2$. Show that in this case the state $\rho(\lambda)$ violates no Bell inequalities. Hint: It suffices to construct a local hidden-variable model that correctly reproduces the spin correlations found in (c), for $\lambda = 1/2$. Suppose that the hidden variable $\hat{\alpha}$ is uniformly distributed on the unit sphere, and that there are functions f_A and f_B such that

$$\operatorname{Prob}_{A}(\uparrow_{\hat{n}}) = f_{A}(\hat{\alpha} \cdot \hat{n}), \quad \operatorname{Prob}_{B}(\uparrow_{\hat{m}}) = f_{B}(\hat{\alpha} \cdot \hat{m}).$$
(4.165)

The problem is to find f_A and f_B (where $0 \le f_{A,B} \le 1$) with the properties

$$\int_{\hat{\alpha}} f_A(\hat{\alpha} \cdot \hat{n}) = 1/2, \quad \int_{\hat{\alpha}} f_B(\hat{\alpha} \cdot \hat{m}) = 1/2,$$
$$\int_{\hat{\alpha}} f_A(\hat{\alpha} \cdot \hat{n}) f_B(\hat{\alpha} \cdot \hat{m}) = \operatorname{Prob}(\uparrow_{\hat{n}} \uparrow_{\hat{m}}). \quad (4.166)$$

4.5 Quantum key distribution

Alice and Bob want to execute a quantum key distribution protocol. Alice is equipped to prepare either one of the two states $|u\rangle$ or $|v\rangle$. These two states, in a suitable basis, can be expressed as

$$|u\rangle = \begin{pmatrix} \cos \alpha \\ \sin \alpha \end{pmatrix}$$
, $|v\rangle = \begin{pmatrix} \sin \alpha \\ \cos \alpha \end{pmatrix}$, (4.167)

where $0 < \alpha < \pi/4$. Alice decides at random to send either $|u\rangle$ or $|v\rangle$ to Bob, and Bob is to make a measurement to determine what she sent. Since the two states are not orthogonal, Bob cannot distinguish the states perfectly.

a) Bob realizes that he can't expect to be able to identify Alice's qubit every time, so he settles for a procedure that is successful only some of the time. He performs a POVM with three possible outcomes: $\neg u$, $\neg v$, or DON'T KNOW. If he obtains the result $\neg u$, he is certain that $|v\rangle$ was sent, and if he obtains $\neg v$, he is certain that $|u\rangle$ was sent. If the result is DON'T KNOW, then his measurement is inconclusive. This POVM is defined by the operators

$$\boldsymbol{f}_{\neg u} = A(\boldsymbol{I} - |u\rangle\langle u|) , \quad \boldsymbol{f}_{\neg v} = A(\boldsymbol{I} - |v\rangle\langle v|) ,$$
$$\boldsymbol{f}_{\mathrm{DK}} = (1 - 2A)\boldsymbol{I} + A(|u\rangle\langle u| + |v\rangle\langle v|) , \quad (4.168)$$

where A is a positive real number. How should Bob choose A to minimize the probability of the outcome DK, and what is this minimal DK probability (assuming that Alice chooses from $\{|u\rangle, |v\rangle\}$ equiprobably)? [Hint: If A is too large, $f_{\rm DK}$ will have negative eigenvalues, and eq.(4.168) will not be a POVM.]

- b) Design a quantum key distribution protocol using Alice's source and Bob's POVM.
- c) Of course, Eve also wants to know what Alice is sending to Bob. Hoping that Alice and Bob won't notice, she intercepts each qubit that Alice sends, by performing an orthogonal measurement that projects onto the basis $\left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$. If she obtains the outcome $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, she sends the state $|u\rangle$ on to Bob, and if she obtains the outcome $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$, she sends $|v\rangle$ on to Bob. Therefore each time Bob's POVM has a conclusive outcome, Eve knows with certainty what that outcome is. But Eve's tampering causes detectable errors; sometimes Bob obtains a "conclusive"

outcome that actually differs from what Alice sent. What is the probability of such an error?

4.6 Minimal disturbance

In Exercise 2.1, we studied a game in which Alice decides at random (equiprobably) whether to prepare one of two possible pure states of a single qubit, either

$$|\psi\rangle = \begin{pmatrix} \cos \alpha \\ \sin \alpha \end{pmatrix}$$
, or $|\tilde{\psi}\rangle = \begin{pmatrix} \sin \alpha \\ \cos \alpha \end{pmatrix}$, (4.169)

and sends the state to Bob. By performing an orthogonal measurement in the basis $\{|0\rangle, |1\rangle\}$, Bob can identify the state with minimal error probability

$$(p_{\text{error}})_{\text{optimal}} = \sin^2 \alpha = \frac{1}{2}(1 - \sin \theta) , \qquad (4.170)$$

where we have defined θ by

$$\langle \psi | \psi \rangle \equiv \cos \theta = \sin(2\alpha) .$$
 (4.171)

But now let's suppose that Eve wants to *eavesdrop* on the state as it travels from Alice to Bob. Like Bob, she wishes to extract optimal information that distinguishes $|\psi\rangle$ from $|\tilde{\psi}\rangle$, and she also wants to minimize the disturbance introduced by her eavesdropping, so that Alice and Bob are not likely to notice that anything is amiss.

Eve realizes that the optimal POVM can be achieved by measurement operators

$$\boldsymbol{M}_0 = |\phi_0\rangle\langle 0|$$
, $\boldsymbol{M}_1 = |\phi_1\rangle\langle 1|$, (4.172)

where the vectors $|\phi_0\rangle$, and $|\phi_1\rangle$ are arbitrary. If Eve performs this measurement, then Bob receives the state

$$\rho' = \cos^2 \alpha |\phi_0\rangle \langle \phi_0| + \sin^2 \alpha |\phi_1\rangle \langle \phi_1| , \qquad (4.173)$$

if Alice sent $|\psi\rangle$, and the state

$$\tilde{\rho}' = \sin^2 \alpha |\phi_0\rangle \langle \phi_0| + \cos^2 \alpha |\phi_1\rangle \langle \phi_1| , \qquad (4.174)$$

if Alice sent $|\tilde{\psi}\rangle$.

Eve wants the average fidelity of the state received by Bob to be as large as possible. The quantity that she wants to minimize, which we will call the "disturbance" D, measures how close this average fidelity is to one:

$$D = 1 - \frac{1}{2}(F + \tilde{F}) , \qquad (4.175)$$

where

$$F = \langle \psi | \rho' | \psi \rangle$$
, $\tilde{F} = \langle \tilde{\psi} | \tilde{\rho}' | \tilde{\psi} \rangle$. (4.176)

The purpose of this exercise is to examine how effectively Eve can reduce the disturbance by choosing her measurement operators properly.

a) Show that $F + \tilde{F}$ can be expressed as

$$F + \tilde{F} = \langle \phi_0 | A | \phi_0 \rangle + \langle \phi_1 | B | \phi_1 \rangle , \qquad (4.177)$$

where

$$A = \begin{pmatrix} 1 - 2\cos^2 \alpha \sin^2 \alpha & \cos \alpha \sin \alpha \\ \cos \alpha \sin \alpha & 2\cos^2 \alpha \sin^2 \alpha \end{pmatrix},$$
$$B = \begin{pmatrix} 2\cos^2 \alpha \sin^2 \alpha & \cos \alpha \sin \alpha \\ \cos \alpha \sin \alpha & 1 - 2\cos^2 \alpha \sin^2 \alpha \end{pmatrix}. (4.178)$$

b) Show that if $|\phi_0\rangle$ and $|\phi_1\rangle$ are chosen optimally, the minimal disturbance that can be attained is

$$D_{\min}(\cos^2 \theta) = \frac{1}{2}(1 - \sqrt{1 - \cos^2 \theta + \cos^4 \theta}) .$$
(4.179)

[Hint: We can choose $|\phi_0\rangle$ and $|\phi_1\rangle$ to maximize the two terms in eq. (4.177) independently. The maximal value is the maximal eigenvalue of A, which since the eigenvalues sum to 1, can be expressed as $\lambda_{\max} = \frac{1}{2} \left(1 + \sqrt{1 - 4 \cdot \det A} \right)$.] Of course, Eve could reduce the disturbance further were she willing to settle for a less than optimal probability of guessing Alice's state correctly.

c) Sketch a plot of the function $D_{\min}(\cos^2 \theta)$. Interpret its value for $\cos \theta = 1$ and $\cos \theta = 0$. For what value of θ is D_{\min} largest? Find D_{\min} and $(p_{\text{error}})_{\text{optimal}}$ for this value of θ .

4.7 Approximate cloning

The no-cloning theorem shows that we can't build a unitary machine that will make a perfect copy of an unknown quantum state. But
suppose we are willing to settle for an *imperfect* copy — what fidelity might we achieve?

Consider a machine that acts on three qubit states according to

$$|000\rangle_{ABC} \rightarrow \sqrt{\frac{2}{3}}|00\rangle_{AB}|0\rangle_{C} + \sqrt{\frac{1}{3}}|\psi^{+}\rangle_{AB}|1\rangle_{C}$$
$$|100\rangle_{ABC} \rightarrow \sqrt{\frac{2}{3}}|11\rangle_{AB}|1\rangle_{C} + \sqrt{\frac{1}{3}}|\psi^{+}\rangle_{AB}|0\rangle_{C} . \quad (4.180)$$

a) Is such a device physically realizable, in principle?

If the machine operates on the initial state $|\psi\rangle_A|00\rangle_{BC}$, it produces an pure entangled state $|\Psi\rangle_{ABC}$ of the three qubits. But if we observe qubit A alone, its final state is the density operator $\rho'_A = \operatorname{tr}_{BC}(|\Psi\rangle_{ABC} |\Psi\rangle_{ABC} \langle\Psi|)$. Similarly, the qubit B, observed in isolation, has the final state ρ'_B . It is easy to see that $\rho'_A = \rho'_B$ —these are the identical, but imperfect, copies of the input pure state $|\psi\rangle_A$.

- b) The mapping from the initial state $|\psi\rangle_A {}_A \langle \psi|$ to the final state ρ'_A of qubit A defines a superoperator \$. Find an operator-sum representation of \$.
- c) For $|\psi\rangle_A = a|0\rangle_A + b|1\rangle_A$, find ρ'_A , and compute its fidelity $F \equiv {}_A \langle \psi | \rho'_A | \psi \rangle_A$.

4.8 We're so sorry, Uncle Albert

Consider the n-qubit "cat" state

$$|\psi\rangle_n = \sqrt{\frac{1}{2}} \left(|000...0\rangle + |111...1\rangle\right) .$$
 (4.181)

This state can be characterized as the simultaneous eigenstate (with eigenvalue 1) of the n operators

$$\sigma_{3} \otimes \sigma_{3} \otimes I \otimes I \otimes \cdots \otimes I \otimes I \otimes I \otimes I$$

$$I \otimes \sigma_{3} \otimes \sigma_{3} \otimes I \otimes \cdots \otimes I \otimes I \otimes I$$

$$\cdots$$

$$I \otimes I \otimes I \otimes I \otimes \cdots \otimes I \otimes \sigma_{3} \otimes \sigma_{3}$$

$$\sigma_{1} \otimes \sigma_{1} \otimes \sigma_{1} \otimes \cdots \otimes \sigma_{1} \otimes \sigma_{1} \otimes \sigma_{1}$$

$$(4.182)$$

a) Show that $|\psi\rangle_n$ is an eigenstate of the operator

$$(\boldsymbol{\sigma}_1 + i\boldsymbol{\sigma}_2)^{\otimes n} + (\boldsymbol{\sigma}_1 - i\boldsymbol{\sigma}_2)^{\otimes n}$$
, (4.183)

and compute its eigenvalue.

- b) If we believe in local hidden variables, then we believe that, for each of the *n* qubits, both σ_1 and σ_2 have definite values once the hidden variables are specified. If so, then what can we say about the *modulus* of $(\sigma_1 + i\sigma_2)^{\otimes n}$ or $(\sigma_1 - i\sigma_2)^{\otimes n}$, assuming definite values for the hidden variables?
- c) From (b), derive an upper bound on

$$\frac{1}{2} \left| (\boldsymbol{\sigma}_1 + i\boldsymbol{\sigma}_2)^{\otimes n} + (\boldsymbol{\sigma}_1 - i\boldsymbol{\sigma}_2)^{\otimes n} \right|$$
(4.184)

that follows from the local hidden-variable hypothesis.

d) Compare with (a). What would Einstein say?

4.9 Entanglement manipulation

- a) Twenty-five players on the New York Yankees, and twenty-five players on the San Diego Padres, want to share a 50-qubit cat state. The Yankees prepare a 26-qubit cat state, and give one of the qubits to Alice; so do the Padres. Now Alice is to sew the states together and prepare the 50-qubit state. What should she do? [Hint: Think about stabilizers.]
- b) After joining the Yankees, Alice assumed custody of one of the qubits in their 25-qubit cat state. But today she has been traded! Alice is ordered to pull her qubit out of the cat state, leaving an undamaged 24-qubit cat state for the other players. What should she do? [Hint: Think about stabilizers.]

4.10 Peres-Horodecki criterion in d dimensions

Recall that a *Werner state* of a pair of qubits can be expressed as

$$\boldsymbol{\rho}(\lambda) = \lambda |\phi^+\rangle \langle \phi^+| + \frac{1}{4}(1-\lambda)\boldsymbol{I} , \qquad (4.185)$$

and that the *partial transpose* ρ_{AB}^{PT} of a bipartite density operator ρ_{AB} is defined as

$$\boldsymbol{\rho}_{AB}^{PT} \equiv (I_A \otimes T_B)(\boldsymbol{\rho}_{AB}) \tag{4.186}$$

where T is the transpose operation that acts in the computational basis $\{|i\rangle\}$ as

$$T\left(|i\rangle\langle j|\right) = |j\rangle\langle i| . \tag{4.187}$$

We saw in class that the partial transpose of the Werner state $\rho(\lambda)$ is negative for $\lambda > 1/3$; therefore, by the *Peres-Horodecki criterion*, the Werner state is inseparable for $\lambda > 1/3$.

a) One natural way to generalize the Werner state to a pair of *d*dimensional systems is to consider

$$\boldsymbol{\rho}_{\Phi}(\lambda) = \lambda |\Phi\rangle \langle \Phi| + \frac{1}{d^2} (1 - \lambda) \boldsymbol{I} , \qquad (4.188)$$

where $|\Phi\rangle$ is the maximally entangled state

$$|\Phi\rangle = \frac{1}{\sqrt{d}} \sum_{i=1}^{d} |i\rangle \otimes |i\rangle . \qquad (4.189)$$

Show that

$$(|\Phi\rangle\langle\Phi|)^{PT} = \frac{1}{d} \left(\boldsymbol{I} - 2\boldsymbol{E}_{\text{antisym}} \right) , \qquad (4.190)$$

where $\boldsymbol{E}_{\text{antisym}}$ is the projector onto the space that is antisymmetric under interchange of the two systems A and B.

- b) For what values of λ does the state $\rho_{\Phi}(\lambda)$ have a negative partial transpose?
- c) If the Werner state for two qubits is chosen to be

$$\boldsymbol{\rho}(\lambda) = \lambda |\psi^{-}\rangle \langle \psi^{-}| + \frac{1}{4}(1-\lambda)\boldsymbol{I} , \qquad (4.191)$$

then another natural way to generalize the Werner state to a pair of d-dimensional systems is to consider

$$\boldsymbol{\rho}_{\text{anti}}(\lambda) = \lambda \left(\frac{1}{\frac{1}{2}d(d-1)}\right) \boldsymbol{E}_{\text{antisym}} + \frac{1}{d^2}(1-\lambda)\boldsymbol{I} \quad (4.192)$$

For what values of λ does $\pmb{\rho}_{\rm anti}(\lambda)$ have a negative partial transpose?

Lecture Notes for Ph219/CS219: Quantum Information Chapter 5

John Preskill California Institute of Technology

Updated July 2015

Contents

5	Classical and quantum circuits	3
5.1	Classical Circuits	3
	5.1.1 Universal gates	3
	5.1.2 Most functions require large circuits	5
	5.1.3 Circuit complexity	6
	5.1.4 Randomized computation	12
5.2	Reversible computation	13
	5.2.1 Landauer's principle	13
	5.2.2 Reversible gates	14
	5.2.3 Saving space: the pebble game	20
5.3	Quantum Circuits	22
	5.3.1 Accuracy	26
	5.3.2 BQP \subseteq PSPACE	29
	5.3.3 Most unitary transformations require large quantum	ı
	circuits	31
5.4	Universal quantum gates	33
	5.4.1 Notions of universality	33
	5.4.2 Two-qubit gates are exactly universal	36
	5.4.3 Finite universal gate sets	39
	5.4.4 The Solovay-Kitaev approximation	42
5.5	Summary	45
5.6	Exercises	46

 $\mathbf{5}$

Classical and quantum circuits

5.1 Classical Circuits

The concept of a quantum computer was introduced in Chapter 1. Here we will specify our model of quantum computation more precisely, and we will point out some basic properties of the model; later we will investigate the power of the model. But before we explain what a quantum computer does, we should say what a classical computer does.

5.1.1 Universal gates

A (deterministic) classical computer evaluates a function: given n-bits of input it produces m-bits of output that are uniquely determined by the input; that is, it finds the value of the function

$$f: \{0,1\}^n \to \{0,1\}^m,\tag{5.1}$$

for a particular specified n-bit argument x. A function with an m-bit output is equivalent to m functions, each with a one-bit output, so we may just as well say that the basic task performed by a computer is the evaluation of

$$f: \{0,1\}^n \to \{0,1\}.$$
 (5.2)

A function talking an *n*-bit input to a one-bit output is called a Boolean function. We may think of such a function as a binary string of length 2^n , where each bit of the string is the output f(x) for one of the 2^n possible values of the input x. Evidently, there are 2^{2^n} such strings; that's a lot of functions! Already for n = 5 there are $2^{32} \approx 4.3 \times 10^9$ Boolean functions — you've encountered only a tiny fraction of these in your lifetime.

It is sometimes useful to regard a Boolean function as a subset Σ of the *n*-bit strings containing those values of the input x such that f(x) = 1; we say these strings are "accepted" by f. The complementary set $\overline{\Sigma}$ contains values of x such that f(x) = 0, which we say are "rejected" by f.

The evaluation of a Boolean function f can be reduced to a sequence of simple logical operations. To see how, denote the *n*-bit strings accepted by f as $\Sigma = \{x^{(1)}, x^{(2)}, x^{(3)}, \ldots\}$ and note that for each $x^{(a)}$ we can define a function $f^{(a)} : \{0, 1\}^n \to \{0, 1\}$ such that

$$f^{(a)}(x) = \begin{cases} 1 & x = x^{(a)} \\ 0 & \text{otherwise} \end{cases}$$
(5.3)

Then f can be expressed as

$$f(x) = f^{(1)}(x) \lor f^{(2)}(x) \lor f^{(3)}(x) \lor \dots,$$
(5.4)

the logical OR (\lor) of all the $f^{(a)}$'s. In binary arithmetic the \lor operation of two bits may be represented

$$x \lor y = x + y - x \cdot y; \tag{5.5}$$

it has the value 0 if x and y are both zero, and the value 1 otherwise.

Now consider the evaluation of $f^{(a)}$. We express the *n*-bit string x as

$$x = x_{n-1}x_{n-2}\dots x_2x_1x_0. (5.6)$$

In the case where $x^{(a)} = 11 \dots 111$, we may write

$$f^{(a)}(x) = x_{n-1} \wedge x_{n-2} \wedge \ldots \wedge x_2 \wedge x_1 \wedge x_0; \tag{5.7}$$

it is the logical AND (^) of all n bits. In binary arithmetic, the AND is the product

$$x \wedge y = x \cdot y. \tag{5.8}$$

For any other $x^{(a)}$, $f^{(a)}$ is again obtained as the AND of n bits, but where the NOT (\neg) operation is first applied to each x_i such that $x_i^{(a)} = 0$; for example

$$f^{(a)}(x) = \dots (\neg x_3) \land x_2 \land x_1 \land (\neg x_0)$$
(5.9)

if

$$x^{(a)} = \dots 0110. \tag{5.10}$$

The NOT operation is represented in binary arithmetic as

$$\neg x = 1 - x. \tag{5.11}$$

We have now constructed the function f(x) from three elementary logical connectives: NOT, AND, OR. The expression we obtained is called the "disjunctive normal form" (DNF) of f(x). We have also implicitly used another operation INPUT (x_i) , which inputs the *i*th bit of x. These considerations motivate the circuit model of computation. A computer has a few basic components that can perform elementary operations on bits or pairs of bits, such as NOT, AND, OR. It can also input a variable bit or prepare a constant bit. A computation is a finite sequence of such operations, a *circuit*, applied to a specified string of input bits. Each operation is called a *gate*. The result of the computation is the final value of all remaining bits, after all the elementary operations have been executed. For a Boolean function (with a one-bit output), if there are multiple bits still remaining at the end of the computation, one is designated as the output bit. A circuit can be regarded as a directed acyclic graph, where each vertex in the graph is a gate, and the directed edges indicate the flow of bits through the circuit, with the direction specifying the order in which gates are applied. By acyclic we mean that no directed closed loops are permitted.

We say that the gate set {NOT, AND, OR, INPUT} is "universal," meaning that any function can be evaluated by building a circuit from these components. It is a remarkable fact about the world that an arbitrary computation can be performed using such simple building blocks.

5.1.2 Most functions require large circuits

Our DNF construction shows that any Boolean function with an *n*-bit input can be evaluated using no more than 2^n OR gates, $n2^n$ AND gates, $n2^n$ NOT gates, and $n2^n$ INPUT gates, a total of $(3n + 1)2^n$ gates. Of course, some functions can be computed using much smaller circuits, but for *most* Boolean functions the smallest circuit that evaluates the function really does have an exponentially large (in *n*) number of gates. The point is that if the circuit size (*i.e.*, number of gates) is subexponential in *n*, then there are many, many more functions than circuits.

How many circuits are there with G gates acting on an n-bit input? Consider the gate set from which we constructed the DNF, where we will also allow inputting of a constant bit (either 0 or 1) in case we want to use some scratch space when we compute. Then there are n + 5 different gates: NOT, AND, OR, INPUT(0), INPUT(1), and INPUT(x_i) for i = $0, 1, 2, \ldots n - 1$. Each two-qubit gate acts on a pair of bits which are outputs from preceding gates; this pair can be chosen in fewer than G^2 ways. Therefore the total number of size-G circuits can be bounded as

$$N_{\text{circuit}}(G) \le \left((n+5)G^2 \right)^G.$$
(5.12)

If $G = c \frac{2^n}{2n}$, where c is a constant independent of n, then

$$\log_2 N_{\text{circuit}}(G) \le G \left(\log_2(n+5) + 2 \log_2 G \right) = c 2^n \left(1 + \frac{1}{2n} \log_2 \left(\frac{c^2(n+5)}{4n^2} \right) \right) \le c 2^n,$$
(5.13)

where the second inequality holds for n sufficiently large. Comparing with the number of Boolean functions, $N_{\text{function}}(n) = 2^{2^n}$, we find

$$\log_2\left(\frac{N_{\text{circuit}}(G)}{N_{\text{function}}(n)}\right) \le (c-1)2^n \tag{5.14}$$

for n sufficiently large. Therefore, for any c < 1, the number of circuits is smaller than the number of functions by a huge factor. We did this analysis for one particular universal gate set, but the counting would not have been substantially different if we had used a different gate set instead.

We conclude that for any positive ε , then, most Boolean functions require circuits with at least $(1 - \varepsilon)\frac{2^n}{2n}$ gates. The circuit size is so large because most functions have no structure that can be exploited to construct a more compact circuit. We can't do much better than consulting a "look-up table" that stores a list of all accepted strings, which is essentially what the DNF does.

5.1.3 Circuit complexity

So far, we have only considered a computation that acts on an input with a fixed (n-bit) size, but we may also consider *families* of circuits that act on inputs of variable size. Circuit families provide a useful scheme for analyzing and classifying the *complexity* of computations, a scheme that will have a natural generalization when we turn to quantum computation.

Boolean functions arise naturally in the study of complexity. A Boolean function f may be said to encode a solution to a "decision problem" — the function examines the input and issues a YES or NO answer. Often, what might not be stated colloquially as a question with a YES/NO answer can be "repackaged" as a decision problem. For example, the function that defines the FACTORING problem is:

$$f(x,y) = \begin{cases} 1 & \text{if integer } x \text{ has a divisor } z \text{ such that } 1 < z < y, \\ 0 & \text{otherwise;} \end{cases}$$
(5.15)

knowing f(x, y) for all y < x is equivalent to knowing the *least* nontrivial factor of x (if there is one).

Another example of a decision problem is the HAMILTONIAN PATH problem: let the input be an ℓ -vertex graph, represented by an $\ell \times \ell$ adjacency matrix (a 1 in the ij entry means there is an edge linking vertices i and j); the function is

$$f(x) = \begin{cases} 1 & \text{if graph } x \text{ has a Hamiltonian path,} \\ 0 & \text{otherwise.} \end{cases}$$
(5.16)

A path on the graph is Hamiltonian if it visits each vertex exactly once.

For the FACTORING problem the size of the input is the number of bits needed to specify x and y, while for the HAMILTONIAN PATH problem the size of the input is the number of bits needed to specify the graph. Thus each problem really defines a family of Boolean functions with variable input size. We denote such a function family as

$$f: \{0,1\}^* \to \{0,1\},\tag{5.17}$$

where the * indicates that the input size is variable. When x is an nbit string, by writing f(x) we mean the Boolean function in the family which acts on an n-bit input is evaluated for input x. The set L of strings accepted by a function family

$$L = \{x \in \{0, 1\}^* : f(x) = 1\}$$
(5.18)

is called a *language*.

We can quantify the hardness of a problem by stating how the computational resources needed to solve the problem scale with the input size *n*. In the circuit model of computation, it is natural to use the circuit size (number of gates) to characterize the required resources. Alternatively, we might be interested in how much *time* it takes to do the computation if many gates can be executed in parallel; the *depth* of a circuit is the number of time steps required, assuming that gates acting on distinct bits can operate simultaneously (that is, the depth is the maximum length of a directed path from the input to the output of the circuit). The *width* of a circuit, the maximum number of gates (including identity gates acting on "resting" bits) that act in any one time step, quantifies the storage space used to execute the computation.

We would like to divide the decision problems into two classes: easy and hard. But where should we draw the line? For this purpose, we consider decision problems with variable input size, where the number of bits of input is n, and examine how the size of the circuit that solves the problem scales with n.

If we use the scaling behavior of a circuit family to characterize the difficulty of a problem, there is a subtlety. It would be cheating to hide the difficulty of the problem in the *design* of the circuit. Therefore, we should

restrict attention to circuit families that have acceptable "uniformity" properties — it must be "easy" to build the circuit with n + 1 bits of input once we have constructed the circuit with an *n*-bit input.

Associated with a family of functions $\{f_n\}$ (where f_n has *n*-bit input) are circuits $\{C_n\}$ that compute the functions. We say that a circuit family $\{C_n\}$ is "polynomial size" if the size $|C_n|$ of C_n grows with *n* no faster than a power of *n*,

size
$$(C_n) \le \operatorname{poly}(n),$$
 (5.19)

where poly denotes a polynomial. Then we define:

$P = \{ \text{decision problems solved by polynomial-size} \\ \text{uniform circuit families} \}$

(P for "polynomial time"). Decision problems in P are "easy." The rest are "hard." Notice that C_n computes $f_n(x)$ for every possible *n*-bit input, and therefore, if a decision problem is in P we can find the answer even for the "worst-case" input using a circuit of size no greater than poly(*n*). As already noted, we implicitly assume that the circuit family is "uniform" so that the *design* of the circuit can itself be solved by a polynomialtime algorithm. Under this assumption, solvability in polynomial time by a circuit family is equivalent to solvability in polynomial time by a universal Turing machine.

Of course, to determine the size of a circuit that computes f_n , we must know what the elementary components of the circuit are. Fortunately, though, whether a problem lies in P does not depend on what gate set we choose, as long as the gates are universal, the gate set is finite, and each gate acts on a constant number of bits. One universal gate set can simulate another efficiently.

The way of distinguishing easy and hard problems may seem rather arbitrary. If $|C_n| \sim n^{1000}$ we might consider the problem to be intractable in practice, even though the scaling is "polynomial," and if $|C_n| \sim n^{\log \log \log n}$ we might consider the problem to be easy in practice, even though the scaling is "superpolynomial." Furthermore, even if $|C_n|$ scales like a modest power of n, the constants in the polynomial could be very large. Such pathological cases seem to be uncommon, however. Usually polynomial scaling is a reliable indicator that solving the problem is feasible.

Of particular interest are decision problems that can be answered by exhibiting an example that is easy to verify. For example, given x and y < x, it is hard (in the worst case) to determine if x has a factor less than y. But if someone kindly provides a z < y that divides x, it is easy for us to check that z is indeed a factor of x. Similarly, it is hard to determine if a graph has a Hamiltonian path, but if someone kindly provides a path, it is easy to verify that the path really is Hamiltonian. This concept that a problem may be hard to solve, but that a solution can be easily verified once found, can be formalized. The complexity class of decision problems for which the answer can be checked efficiently, called NP, is defined as

Definition. NP. A language L is in NP iff there is a polynomial-size verifier V(x, y) such that

If
$$x \in L$$
, then there exists y such that $V(x, y) = 1$ (completeness),
If $x \notin L$, then, for all y, $V(x, y) = 0$ (soundness).

The verifier V is the uniform circuit family that checks the answer. Completeness means that for each input in the language (for which the answer is YES), there is an appropriate "witness" such that the verifier accepts the input if that witness is provided. Soundness means that for each input not in the language (for which the answer is NO) the verifier rejects the input no matter what witness is provided. It is implicit that the witness is of polynomial length, |y| = poly(|x|); since the verifier has a polynomial number of gates, including input gates, it cannot make use of more than a polynomial number of bits of the witness. NP stands for "nondeterministic polynomial time;" this name is used for historical reasons, but it is a bit confusing since the verifier is actually a deterministic circuit (evaluates a function).

If is obvious that $P \subseteq NP$; if the problem is in P then the polynomialtime verifier can decide whether to accept x on its own, without any help from the witness. But some problems in NP seem to be hard, and are believed not to be in P. Much of complexity theory is built on a fundamental conjecture:

$$Conjecture : P \neq NP.$$
(5.20)

Proving or refuting this conjecture is the most important open problem in computer science, and one of the most important problems in mathematics.

Why should we believe $P \neq NP$? If P = NP that would mean we could easily find the solution to any problem whose solution is easy to check. In effect, then, we could automate creativity; in particular, computers would be able to discover all the mathematical theorems which have short proofs. The conjecture $P \neq NP$ asserts that our machines will never achieve such awesome power — that the mere existence of a succinct proof of a statement does not ensure that we can find the proof by any systematic procedure in any reasonable amount of time.

An important example of a problem in NP is CIRCUIT-SAT. In this case the input is a Boolean circuit C, and the problem is to determine

whether any input x is accepted by C. The function to be evaluated is

$$f(C) = \begin{cases} 1 & \text{if there exists } x \text{ with } C(x) = 1, \\ 0 & \text{otherwise.} \end{cases}$$
(5.21)

This problem is in NP because if the circuit C has polynomial size, then if we are provided with an input x accepted by C it is easy to check that C(x) = 1.

The problem CIRCUIT-SAT is particularly interesting because it has a remarkable property — if we have a machine that solves CIRCUIT-SAT we can use it to solve any other problem in NP. We say that every problem in NP is (efficiently) *reducible* to CIRCUIT-SAT. More generally, we say that problem B reduces to problem A if a machine that solves A can be used to solve B as well.

That is, if A and B are Boolean function families, then "B reduces to A" means there is a function family R, computed by poly-size circuits, such that B(x) = A(R(x)). Thus B accepts x iff A accepts R(x). In particular, then, if we have a poly-size circuit family that solves A, we can hook A up with R to obtain a poly-size circuit family that solves B.

A problem B in NP reduces to CIRCUIT-SAT because problem B has a poly-size verifier V(x, y), such that B accepts x iff there exists some witness y such that V accepts (x, y). For each fixed x, asking whether such a witness y exists is an instance of CIRCUIT-SAT. So a poly-size circuit family that solves CIRCUIT-SAT can also be used to solve problem B.

We say a problem A in NP is NP-complete if every problem in NP is reducible to A. Hence, CIRCUIT-SAT is NP-complete. The NP-complete problems are the "hardest" problems in NP, in the sense that if we can solve any NP-complete problem then we can solve every NP problem. Furthermore, to show that problem A is NP-complete, it is enough to show that B reduces to A where B is NP-complete. If C is any problem in NP and B is NP-complete then there is a poly-size reduction R such that C(x) = B(R(x)), and if B is reducible to A then there is another poly-size reduction R' such that B(y) = A(R'(y)). Hence C(x) = A(R'(R(x))), and since the composition $R' \circ R$ of two poly-size reductions is also poly-size, we see that an arbitrary problem C in NP reduces to A, and therefore Ais NP-complete. NP-completeness is a useful concept because hundreds of "natural" computational problems turn out to be NP-complete. For example, one can exhibit a polynomial reduction of CIRCUIT-SAT to HAMILTONIAN PATH, and it follows that HAMILTONIAN PATH is also NP-complete.

Another noteworthy complexity class is called co-NP. While NP problems can be decided by exhibiting an *example* if the answer is YES, co-NP problems can be answered by exhibiting a *counter-example* if the answer is NO. More formally:

Definition. co-NP. A language L is in co-NP iff there is a polynomialsize verifier $\overline{V}(x, y)$ such that

If
$$x \notin L$$
, then there exists y such that $\bar{V}(x, y) = 1$,
If $x \in L$, then, for all y, $\bar{V}(x, y) = 0$.

For NP the witness y testifies that x is in the language while for co-NP the witness testifies that x is *not* in the language. Thus if language L is in NP, then its complement \overline{L} is in co-NP and vice-versa. We see that whether we consider a problem to be in NP or in co-NP depends on how we choose to frame the question — while "Is there a Hamiltonian path?" is in NP, the complementary question "Is there no Hamiltonian path?" is in co-NP.

Though the distinction between NP and co-NP may seem arbitrary, it is nevertheless interesting to ask whether a problem is in *both* NP and co-NP. If so, then we can easily verify the answer (once a suitable witness is in hand) regardless of whether the answer is YES or NO. It is believed (though not proved) that NP \neq co-NP. For example, we can show that a graph has a Hamiltonian path by exhibiting an example, but we don't know how to show that it has *no* Hamiltonian path that way!

If we assume that $P \neq NP$, it is known that there exist problems in NP of intermediate difficulty (the class NPI), which are neither in P nor NP-complete. Furthermore, assuming that that NP \neq co-NP, it is known that no co-NP problems are NP-complete. Therefore, problems in the intersection of NP and co-NP, if not in P, are good candidates for inclusion in NPI.

In fact, a problem in NP \cap co-NP believed not to be in P is the FAC-TORING problem. As already noted, FACTORING is in NP because, if we are offered a factor of x, we can easily check its validity. But it is also in co-NP, because it is known that if we are given a prime number we can efficiently verify its primality. Thus, if someone tells us the prime factors of x, we can efficiently check that the prime factorization is right, and can *exclude* that any integer less than y is a divisor of x. Therefore, it seems likely that FACTORING is in NPI.

We are led to a crude (conjectured) picture of the structure of NP \cup co-NP. NP and co-NP do not coincide, but they have a nontrivial intersection. P lies in NP \cap co-NP but the intersection also contains problems not in P (like FACTORING). No NP-complete or co-NP-complete problems lie in NP \cap co-NP.

5.1.4 Randomized computation

It is sometimes useful to consider *probabilistic* circuits that have access to a random number generator. For example, a gate in a probabilistic circuit might act in either one of two ways, and flip a fair coin to decide which action to execute. Such a circuit, for a single fixed input, can sample many possible computational paths. An algorithm performed by a probabilistic circuit is said to be "randomized."

If we run a randomized computation many times on the same input, we won't get the same answer every time; rather there is a probability distribution of outputs. But the computation is useful if the probability of getting the right answer is high enough. For a decision problem, we would like a randomized computation to accept an input x which is in the language L with probability at least $\frac{1}{2} + \delta$, and to accept an input x which is not in the language with probability no greater than $\frac{1}{2} - \delta$, where $\delta > 0$ is a constant independent of the input size. In that case we can *amplify* the probability of success by performing the computation many times and taking a majority vote on the outcomes. For $x \in L$, if we run the computation N times, the probability of rejecting in more than half the runs is no more than $e^{-2N\delta^2}$ (the *Chernoff bound*). Likewise, for $x \notin L$, the probability of accepting in the majority of N runs is no more than $e^{-2N\delta^2}$.

Why? There are all together 2^N possible sequences of outcomes in the N trials, and the probability of any particular sequence with N_W wrong answers is

$$\left(\frac{1}{2} - \delta\right)^{N_W} \left(\frac{1}{2} + \delta\right)^{N - N_W}.$$
(5.22)

The majority is wrong only if $N_W \ge N/2$, so the probability of any sequence with an incorrect majority is no larger than

$$\left(\frac{1}{2} - \delta\right)^{N/2} \left(\frac{1}{2} + \delta\right)^{N/2} = \frac{1}{2^N} \left(1 - 4\delta^2\right)^{N/2}.$$
 (5.23)

Using $1 - x \le e^{-x}$ and multiplying by the total number of sequences 2^N , we obtain the Chernoff bound:

Prob(wrong majority)
$$\leq (1 - 4\delta^2)^{N/2} \leq e^{-2N\delta^2}$$
. (5.24)

If we are willing to accept a probability of error no larger than ε , then, it suffices to run the computation a number of times

$$N \ge \frac{1}{2\delta^2} \ln\left(1/\varepsilon\right). \tag{5.25}$$

Because we can make the error probability very small by repeating a randomized computation a modest number of times, the value of the constant δ does not really matter for the purpose of classifying complexity, as long as it is positive and independent of the input size. The standard convention is to specify $\delta = 1/6$, so that $x \in L$ is accepted with probability at least 2/3 and $x \notin L$ is accepted with probability no more than 1/3. This criterion defines the class BPP ("bounded-error probabilistic polynomial time") containing decision problems solved by polynomial-size randomized uniform circuit families.

It is clear that BPP contains P, since a deterministic computation is a special case of a randomized computation, in which we never consult the source of randomness. It is widely believed that BPP=P, that randomness does not enhance our computational power, but this has not been proven. It is not even known whether BPP is contained in NP.

We may define a randomized class analogous to NP, called MA ("Merlin-Arthur"), containing languages that can be checked when a randomized verifier is provided with a suitable witness:

Definition. MA. A language L is in MA iff there is a polynomial-size randomized verifier V(x, y) such that

If $x \in L$, then there exists y such that $\operatorname{Prob}(V(x, y) = 1) \ge 2/3$, If $x \notin L$, then, for all y, $\operatorname{Prob}(V(x, y) = 1) \le 1/3$.

The colorful name evokes a scenario in which the all-powerful Merlin uses his magical powers to conjure the witness, allowing the mortal Arthur, limited to polynomial time computation, to check the answer. Obviously BPP is contained in MA, but we expect BPP \neq MA just as we expect P \neq NP.

5.2 Reversible computation

In devising a model of a quantum computer, we will generalize the circuit model of classical computation. But our quantum logic gates will be unitary transformations, and hence will be invertible, while classical logic gates like the AND gate are not invertible. Before we discuss quantum circuits, it is useful to consider some features of *reversible* classical computation.

5.2.1 Landauer's principle

Aside from providing a bridge to quantum computation, classical reversible computing is interesting in its own right, because of *Landauer's principle*. Landauer observed that erasure of information is necessarily a *dissipative* process. His insight is that erasure always involves the compression of phase space, and so is thermodynamically, as well as logically, irreversible.

For example, I can store one bit of information by placing a single molecule in a box, either on the left side or the right side of a partition that divides the box. Erasure means that we move the molecule to the right side (say) irrespective of whether it started out on the left or right. I can suddenly remove the partition, and then slowly compress the one-molecule "gas" with a piston until the molecule is definitely on the right side. This procedure changes the entropy of the gas by $\Delta S = -k \ln 2$ (where k is Boltzmann's constant) and there is an associated flow of heat from the box to its environment. If the process is quasi-static and isothermal at temperature T, then work $W = -kT\Delta S = kT \ln 2$ is performed on the box, work that I have to provide. If I erase information, someone has to pay the power bill.

Landauer also observed that, because irreversible logic elements erase information, they too are necessarily dissipative, and therefore require an unavoidable expenditure of energy. For example, an AND gate maps two input bits to one output bit, with 00, 01, and 10 all mapped to 0, while 11 is mapped to one. If the input is destroyed and we can read only the output, then if the output is 0 we don't know for sure what the input was — there are three possibilities. If the input bits are chosen uniformly at random, than on average the AND gate destroys $\frac{3}{4} \log_2 3 \approx 1.189$ bits of information. Indeed, if the input bits are uniformly random any 2-to-1 gate must "erase" at least one bit on average. According to Landauer's principle, then, we need to do an amount of work at least $W = kT \ln 2$ to operate a 2-to-1 logic gate at temperature T.

But if a computer operates reversibly, then in principle there need be no dissipation and no power requirement. We can compute for free! At present this idea is not of great practical importance, because the power consumed in today's integrated circuits exceeds kT per logic gate by at least three orders of magnitude. As the switches on chips continue to get smaller, though, reversible computing might eventually be invoked to reduce dissipation in classical computing hardware.

5.2.2 Reversible gates

A reversible computer evaluates an invertible function taking n bits to n bits

$$f: \{0,1\}^n \to \{0,1\}^n$$
 (5.26)

An invertible function has a unique input for each output, and we can run the computation backwards to recover the input from the output. We may regard an invertible function as a permutation of the 2^n strings of n bits — there are $(2^n)!$ such functions. If we did not insist on invertibility, there would be $(2^{2^n})^n = (2^n)^{2^n}$ functions taking n bits to n bits (the number of ways to choose n Boolean functions); using the Stirling approximation, $(2^n)! \approx (2^n/e)^{2^n}$, we see that the fraction of all functions which are invertible is quite small, about e^{-2^n} .

Any irreversible computation can be "packaged" as an evaluation of an invertible function. For example, for any $f : \{0,1\}^n \to \{0,1\}$, we can construct $\tilde{f} : \{0,1\}^{n+1} \to \{0,1\}^{n+1}$ such that

$$f(x,y) = (x, y \oplus f(x)).$$
 (5.27)

Here y is a bit and \oplus denotes the XOR gate (addition mod 2) — the *n*-bit input x is preserved and the last bit flips iff f(x) = 1. Applying \tilde{f} a second time undoes this bit flip; hence \tilde{f} is invertible, and equal to its own inverse. If we set y = 0 initially and apply \tilde{f} , we can read out the value of f(x) in the last output bit.

Just as for Boolean functions, we can ask whether a complicated reversible computation can be executed by a circuit built from simple components — are there universal reversible gates? It is easy to see that one-bit and two-bit reversible gates do not suffice; we will need three-bit gates for universal reversible computation.

Of the four 1-bit \rightarrow 1-bit gates, two are reversible; the trivial gate and the NOT gate. Of the $(2^4)^2 = 256$ possible 2-bit \rightarrow 2-bit gates, 4! = 24 are reversible. One of special interest is the controlled-NOT (CNOT) or reversible XOR gate that we already encountered in Chapter 4:

$$XOR: (x, y) \mapsto (x, x \oplus y), \tag{5.28}$$

$$\begin{array}{ccc} x & & & \\ y & & & \\ \end{array} \begin{array}{c} & & \\ & & \\ \end{array} \begin{array}{c} x \\ & & \\ \end{array} \begin{array}{c} x \\ & \\ \end{array} \begin{array}{c} y \\ \\ \end{array} \begin{array}{c} x \\ \\ \end{array} \begin{array}{c} y \\ \end{array} \begin{array}{c} y \\ \end{array} \begin{array}{c} x \\ \\ \end{array} \begin{array}{c} y \\ \end{array} \end{array}$$

This gate flips the second bit if the first is 1, and does nothing if the first bit is 0 (hence the name controlled-NOT). Its square is trivial; hence it inverts itself. Anticipating the notation that will be convenient for our discussion of quantum gates, we will sometimes use $\Lambda(\mathbf{X})$ to denote the CNOT gate. More generally, by $\Lambda(\mathbf{G})$ we mean a gate that applies the operation \mathbf{G} to a "target" conditioned on the value of a "control bit;" \mathbf{G} is applied if the control bit is 1 and the identity is applied if the control bit is 0. In the case of the CNOT gate, \mathbf{G} is the Pauli operator \mathbf{X} , a bit flip. The CNOT gate performs a NOT on the second bit if the first bit x is set to 1, and it performs the copy operation if y is initially set to zero:

$$CNOT: (x,0) \mapsto (x,x). \tag{5.29}$$

With the circuit



constructed from three XOR's, we can swap two bits:

$$(x,y) \to (x,x \oplus y) \to (y,x \oplus y) \to (y,x).$$
(5.30)

With these swaps we can shuffle bits around in a circuit, bringing them together if we want to act on them with a "local gate" at a fixed location.

To see that the one-bit and two-bit gates are nonuniversal, we observe that all these gates are *linear*. Each reversible two-bit gate has an action of the form

$$\begin{pmatrix} x \\ y \end{pmatrix} \mapsto \begin{pmatrix} x' \\ y' \end{pmatrix} = M \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} a \\ b \end{pmatrix};$$
(5.31)

the pair of bits $\binom{a}{b}$ can take any one of four possible values, and the matrix M is one of the six invertible matrices with binary entries

$$M = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}.$$
(5.32)

(All addition is performed modulo 2.) Combining the six choices for M with the four possible constants, we obtain 24 distinct gates, exhausting all the reversible $2 \rightarrow 2$ gates.

Since the linear transformations are closed under composition, any circuit composed from reversible $2 \rightarrow 2$ (and $1 \rightarrow 1$) gates will compute a linear function

$$x \mapsto Mx + a. \tag{5.33}$$

But for $n \geq 3$, there are invertible functions on *n*-bits that are nonlinear. An important example is the 3-bit *Toffoli gate* (or controlled-controlled-NOT) $\Lambda^2(\mathbf{X})$

$$\Lambda^2(\boldsymbol{X}): (x, y, z) \to (x, y, z \oplus xy); \tag{5.34}$$



it flips the third bit if the first two are 1 and does nothing otherwise, thus invoking the (nonlinear) multiplication of the two bits x and y. The $\Lambda^2(\cdot)$ notation indicates that the operation acting on the target bit is triggered only if both control bits are set to 1. Like the CNOT gate $\Lambda(\mathbf{X})$, $\Lambda^2(\mathbf{X})$ is its own inverse.

Unlike the reversible 2-bit gates, the Toffoli gate serves as a universal gate for Boolean logic, if we can provide constant input bits and ignore output bits. If we fix x = y = 1, then the Toffoli gate performs NOT acting on the third bit, and if z is set to zero initially, then the Toffoli gate outputs $z = x \wedge y$ in the third bit. Since NOT/AND/OR are a universal gate set, and we can construct OR from NOT and AND $(x \lor y = \neg(\neg x \land \neg y))$, this is already enough to establish that the Toffoli gate is universal. Note also that if we fix x = 1 the Toffoli gate functions like a CNOT gate acting on y and z; we can use it to copy.

The Toffoli gate $\Lambda^2(\mathbf{X})$ is also universal in the sense that we can build a circuit to compute any reversible function using Toffoli gates alone (if we can fix input bits and ignore output bits). It will be instructive to show this directly, without relying on our earlier argument that NOT/AND/OR is universal for Boolean functions. Specifically, we can show the following: From the NOT gate and the Toffoli gate $\Lambda^2(\mathbf{X})$, we can construct any invertible function on n bits, provided we have one extra bit of scratchpad space available.

The first step is to show that from the three-bit Toffoli-gate $\Lambda^2(\mathbf{X})$ we can construct an *n*-bit Toffoli gate $\Lambda^{n-1}(\mathbf{X})$ that acts as

$$(x_1, x_2, \dots, x_{n-1}, y) \to (x_1, x_2, \dots, x_{n-1}, y \oplus x_1 x_2 \dots x_{n-1})$$
 (5.35)

using one extra bit of scratch space. For example, we construct $\Lambda^3(\mathbf{X})$ from $\Lambda^2(\mathbf{X})$'s with the circuit



The purpose of the last $\Lambda^3(\mathbf{X})$ gate is to reset the scratch bit back to its original value 0. Actually, with one more gate we can obtain an implementation of $\Lambda^3(\mathbf{X})$ that works irrespective of the initial value of the scratch bit:



We can see that the scratch bit really is necessary, because $\Lambda^3(\mathbf{X})$ is an odd permutation (in fact a transposition) of the 4-bit strings it transposes 1111 and 1110. But $\Lambda^2(\mathbf{X})$ acting on any three of the four bits is an even permutation; *e.g.*, acting on the last three bits it transposes both 0111 with 0110 and 1111 with 1110. Since a product of even permutations is also even, we cannot obtain $\Lambda^3(\mathbf{X})$ as a product of $\Lambda^2(\mathbf{X})$'s that act only on the four bits.

This construction of $\Lambda^3(\mathbf{X})$ from four $\Lambda^2(\mathbf{X})$'s generalizes immediately to the construction of $\Lambda^{n-1}(\mathbf{X})$ from two $\Lambda^{n-2}(\mathbf{X})$'s and two $\Lambda^2(\mathbf{X})$'s (just expand x_1 to several control bits in the above diagram). Iterating the construction, we obtain $\Lambda^{n-1}(\mathbf{X})$ from a circuit with $2^{n-2} + 2^{n-3} - 2$ $\Lambda^2(\mathbf{X})$'s. Furthermore, just one bit of scratch suffices. (With more scratch space, we can build $\Lambda^{n-1}(\mathbf{X})$ from $\Lambda^2(\mathbf{X})$'s much more efficiently — see Exercise 5.1.)

The next step is to note that, by conjugating $\Lambda^{n-1}(\mathbf{X})$ with NOT gates, we can in effect modify the value of the control string that "triggers" the gate. For example, the circuit



flips the value of y if $x_1x_2x_3 = 010$, and it acts trivially otherwise. Thus this circuit transposes the two strings 0100 and 0101. In like fashion, with $\Lambda^{n-1}(\mathbf{X})$ and NOT gates, we can devise a circuit that transposes any two *n*-bit strings that differ in only one bit. (The location of the bit where they differ is chosen to be the *target* of the $\Lambda^{n-1}(\mathbf{X})$ gate.)

But in fact a transposition that exchanges any two *n*-bit strings can be expressed as a product of transpositions that interchange strings that differ in only one bit. If a_0 and a_s are two strings that are Hamming distance *s* apart (differ in *s* places), then there is a sequence of strings

$$a_0, a_1, a_2, a_3, \dots, a_s,$$
 (5.36)

such that each string in the sequence is Hamming distance one from its neighbors. Therefore, each of the transpositions

$$(a_0a_1), (a_1a_2), (a_2a_3), \dots, (a_{s-1}a_s),$$
 (5.37)

can be implemented as a $\Lambda^{n-1}(X)$ gate conjugated by NOT gates. By composing transpositions we find

$$(a_0a_s) = (a_{s-1}a_s)(a_{s-2}a_{s-1})\dots(a_2a_3)(a_1a_2)(a_0a_1)(a_1a_2)(a_2a_3)\dots(a_{s-2}a_{s-1})(a_{s-1}a_s);$$
(5.38)

we can construct the Hamming-distance-s transposition from 2s - 1 Hamming-distance-one transpositions. It follows that we can construct (a_0a_s) from $\Lambda^{n-1}(\mathbf{X})$'s and NOT gates.

Finally, since every permutation is a product of transpositions, we have shown that every invertible function on *n*-bits (every permutation of the *n*-bit strings) is a product of $\Lambda^{n-1}(\mathbf{X})$'s and NOT's, using just one bit of scratch space.

Of course, a NOT can be performed with a $\Lambda^2(\mathbf{X})$ gate if we fix two input bits at 1. Thus the Toffoli gate $\Lambda^2(\mathbf{X})$ is universal for reversible computation, if we can fix input bits and discard output bits.

5 Classical and quantum circuits

5.2.3 Saving space: the pebble game

We have seen that with Toffoli and NOT gates we can compute any invertible function using very little scratch space, and also that by fixing constant input bits and ignoring output bits, we can simulate any (irreversible) Boolean circuit using only reversible Toffoli gates. In the latter case, though, we generate two bits of junk every time we simulate an AND gate. Our memory gradually fills with junk, until we reach the stage where we cannot continue with the computation without erasing some bits to clear some space. At that stage, we will finally have to pay the power bill for the computing we have performed, just as Landauer had warned.

Fortunately, there is a general procedure for simulating an irreversible circuit using reversible gates, in which we can erase the junk without using any power. We accumulate and save all the junk output bits as the simulation proceeds, and when we reach the end of the computation we make a copy of the output. The COPY operation, which is logically reversible, can be done with a CNOT or Toffoi gate. Then we run the full computation in reverse, executing the circuit in the opposite order and replacing each gate by its inverse. This procedure cleans up all the junk bits, and returns all registers to their original settings, without any irreversible erasure steps. Yet the result of the computation has been retained, because we copied it before reversing the circuit.

Because we need to run the computation both forward and backward, the reversible simulation uses roughly twice as many gates as the irreversible circuit it simulates. Far worse than that, this simulation method requires a substantial amount of memory, since we need to be able to store about as many bits as the number of gates in the circuit before we finally start to clear the memory by reversing the computation.

It is possible, though, at the cost of modestly increasing the simulation time, to substantially reduce the space required. The trick is to clear space during the course of the simulation by running a part of the computation backward. The resulting tradeoff between time and space is worth discussing, as it illustrates both the value of "uncomputing" and the concept of a recursive simulation.

We imagine dividing the computation into steps of roughly equal size. When we run step k forward, the first thing we do is make a copy of the output from the previous step, then we execute the gates of step k, retaining all the junk accumulated by those gates. We cannot run step k forward unless we have previously completed step k - 1. Furthermore, we will not be able to run step k backward if we have already run step k - 1backward. The trouble is that we will not be able to reverse the COPY step at the very beginning of step k unless we have retained the output from step k - 1.

To save space in our simulation we want to minimize at all times the number of steps that have already been computed but have not yet been uncomputed. The challenge we face can be likened to a game — the *reversible pebble game*. The steps to be executed form a one-dimension directed graph with sites labeled $1, 2, 3, \ldots, T$. Execution of step k is modeled by placing a pebble on the kth site of the graph, and executing step k in reverse is modeled as removal of a pebble from site k. At the beginning of the game, no sites are covered by pebbles, and in each turn we add or remove a pebble. But we cannot place a pebble at site k (except for k = 1) unless site k - 1 is already covered by a pebble, and we cannot remove a pebble from site k (except for k = 1) unless site k - 1 is covered. The object is to cover site T (complete the computation) without using more pebbles than necessary (generating a minimal amount of garbage).

We can construct a recursive procedure that enables us to reach site $t = 2^n$ using n + 1 pebbles and leaving only one pebble in play. Let $F_1(k)$ denote placing a pebble at site k, and $F_1(k)^{-1}$ denote removing a pebble from site k. Then

$$F_2(1,2) = F_1(1)F_1(2)F_1(1)^{-1}, (5.39)$$

leaves a pebble at site k = 2, using a maximum of two pebbles at intermediate stages. Similarly

$$F_3(1,4) = F_2(1,2)F_2(3,4)F_2(1,2)^{-1},$$
(5.40)

reaches site k = 4 using three pebbles, and

$$F_4(1,8) = F_3(1,4)F_3(5,8)F_3(1,4)^{-1},$$
(5.41)

reaches k = 8 using four pebbles. Proceeding this way we construct $F_n(1, 2^n)$ which uses a maximum of n + 1 pebbles and leaves a single pebble in play.

Interpreted as a routine for simulating $T_{irr} = 2^n$ steps of an irreversible computation, this strategy for playing the pebble game represents a reversible simulation requiring space S_{rev} scaling like

$$S_{\rm rev} \approx S_{\rm step} \log_2 \left(T_{\rm irr} / T_{\rm step} \right),$$
 (5.42)

where T_{step} is the number of gates is a single step, and S_{step} is the amount of memory used in a single step. How long does the simulation take? At each level of the recursive procedure described above, two steps forward are replaced by two steps forward and one step back. Therefore, an irreversible computation with $T_{\text{irr}}/T_{\text{step}} = 2^n$ steps is simulated in $T_{\text{rev}}/T_{\text{step}} = 3^n$ steps, or

$$T_{\rm rev} = T_{\rm step} \left(T_{\rm irr} / T_{\rm step} \right)^{\log 3 / \log 2} = T_{\rm step} (T_{\rm irr} / T_{\rm step})^{1.58}, \tag{5.43}$$

a modest power law slowdown.

We can improve this slowdown to

$$T_{\rm rev} \sim (T_{\rm irr})^{1+\varepsilon},$$
 (5.44)

for any $\varepsilon > 0$. Instead of replacing two steps forward with two forward and one back, we replace ℓ forward with ℓ forward and $\ell - 1$ back. A recursive procedure with n levels reaches site ℓ^n using a maximum of $n(\ell - 1) + 1$ pebbles. Now we have $T_{irr} \propto \ell^n$ and $T_{rev} \propto (2\ell - 1)^n$, so that

$$T_{\rm rev} = T_{\rm step} (T_{\rm irr}/T_{\rm step})^{\log(2\ell-1)/\log\ell};$$
(5.45)

the power characterizing the slowdown is

$$\frac{\log(2\ell-1)}{\log\ell} = \frac{\log 2\ell + \log\left(1 - \frac{1}{2\ell}\right)}{\log\ell} \simeq 1 + \frac{\log 2}{\log\ell} \equiv 1 + \varepsilon, \quad (5.46)$$

and the space requirement scales as

$$S_{\rm rev}/S_{\rm step} \approx \ell n \approx 2^{1/\varepsilon} \log_{\ell} \left(T_{\rm irr}/T_{\rm step} \right) \approx \varepsilon 2^{1/\varepsilon} \log_2 \left(T_{\rm irr}/T_{\rm step} \right),$$
 (5.47)

where $1/\varepsilon = \log_2 \ell$. The required space still scales as $S_{\text{rev}} \sim \log T_{\text{irr}}$, yet the slowdown is no worse than $T_{\text{rev}} \sim (T_{\text{irr}})^{1+\varepsilon}$. By using more than the minimal number of pebbles, we can reach the last step faster.

You might have worried that, because reversible computation is "harder" than irreversible computation, the classification of complexity depends on whether we compute reversibly or irreversibly. But don't worry — we've now seen that a reversible computer can simulate an irreversible computer pretty easily.

5.3 Quantum Circuits

Now we are ready to formulate a mathematical model of a quantum computer. We will generalize the circuit model of classical computation to the quantum circuit model of quantum computation.

A classical computer processes bits. It is equipped with a finite set of gates that can be applied to sets of bits. A quantum computer processes qubits. We will assume that it too is equipped with a discrete set of fundamental components, called *quantum gates*. Each quantum gate is a unitary transformation that acts on a fixed number of qubits. In a quantum computation, a finite number n of qubits are initially set to the value $|00...0\rangle$. A circuit is executed that is constructed from a finite number of quantum gates acting on these qubits. Finally, an orthogonal measurement of all the qubits (or a subset of the qubits) is performed,

projecting each measured qubit onto the basis $\{|0\rangle, |1\rangle\}$. The outcome of this measurement is the result of the computation.

Several features of this model invite comment:

(1) Preferred decomposition into subsystems. It is implicit but important that the Hilbert space of the device has a preferred decomposition into a tensor product of low-dimensional subsystems, in this case the qubits. Of course, we could have considered a tensor product of, say, qutrits instead. But anyway we assume there is a natural decomposition into subsystems that is respected by the quantum gates — the gates act on only a few subsystems at a time. Mathematically, this feature of the gates is crucial for establishing a clearly defined notion of quantum complexity. Physically, the fundamental reason for a natural decomposition into subsystems is *locality*; feasible quantum gates must act in a bounded spatial region, so the computer decomposes into subsystems that interact only with their neighbors.

(2) Finite instruction set. Since unitary transformations form a continuum, it may seem unnecessarily restrictive to postulate that the machine can execute only those quantum gates chosen from a discrete set. We nevertheless accept such a restriction, because we do not want to invent a new physical implementation each time we are faced with a new computation to perform. (When we develop the theory of fault-tolerant quantum computing we will see that only a discrete set of quantum gates can be well protected from error, and we'll be glad that we assumed a finite gate set in our formulation of the quantum circuit model.)

(3) Unitary gates and orthogonal measurements. We might have allowed our quantum gates to be trace-preserving completely positive maps, and our final measurement to be a POVM. But since we can easily simulate a TPCP map by performing a unitary transformation on an extended system, or a POVM by performing an orthogonal measurement on an extended system, the model as formulated is of sufficient generality.

(4) Simple preparations. Choosing the initial state of the n input qubits to be $|00...0\rangle$ is merely a convention. We might want the input to be some nontrivial classical bit string instead, and in that case we would just include NOT gates in the first computational step of the circuit to flip some of the input bits from 0 to 1. What is important, though, is that the initial state is easy to prepare. If we allowed the input state to be a complicated entangled state of the n qubits, then we might be hiding the difficulty of executing the quantum algorithm in the difficulty of preparing the input state. We start with a product state instead, regarding it as uncontroversial that preparation of a product state is easy.

(5) Simple measurements. We might allow the final measurement to be a collective measurement, or a projection onto a different basis. But any such measurement can be implemented by performing a suitable unitary transformation followed by a projection onto the standard basis $\{|0\rangle, |1\rangle\}^n$. Complicated collective measurements can be transformed into measurements in the standard basis only with some difficulty, and it is appropriate to take into account this difficulty when characterizing the complexity of an algorithm.

(6) Measurements delayed until the end. We might have allowed measurements at intermediate stages of the computation, with the subsequent choice of quantum gates conditioned on the outcome of those measurements. But in fact the same result can always be achieved by a quantum circuit with all measurements postponed until the end. (While we can postpone the measurements in principle, it might be very useful in practice to perform measurements at intermediate stages of a quantum algorithm.)

A quantum gate, being a unitary transformation, is reversible. In fact, a classical reversible computer is a special case of a quantum computer. A classical reversible gate

$$x \to y = f(x), \tag{5.48}$$

implementing a permutation of k-bit strings, can be regarded as a unitary transformation U acting on k qubits, which maps the "computational basis" of product states

$$\{|x_i\rangle, i = 0, 1, \dots 2^k - 1\}$$
(5.49)

to another basis of product states $\{|y_i\rangle\}$ according to

$$U|x_i\rangle = |y_i\rangle. \tag{5.50}$$

Since U maps one orthonormal basis to another it is manifestly unitary. A quantum computation constructed from such reversible classical gates takes $|0...0\rangle$ to one of the computational basis states, so that the outcome of the final measurement in the $\{|0\rangle, |1\rangle\}$ basis is deterministic.

There are four main issues concerning our model that we would like to address in this Chapter. The first issue is *universality*. The most general unitary transformation that can be performed on n qubits is an element of $U(2^n)$. Our model would seem incomplete if there were transformations in $U(2^n)$ that we were unable to reach. In fact, we will see that there are many ways to choose a discrete set of *universal quantum gates*. Using a universal gate set we can construct circuits that compute a unitary transformation coming as close as we please to any element in $U(2^n)$.

Thanks to universality, there is also a machine independent notion of *quantum complexity*. We may define a new complexity class BQP ("bounded-error quantum polynomial time") — the class of languages that can be decided with high probability by polynomial-size uniform quantum circuit families. Since one universal quantum computer can simulate another efficiently, the class does not depend on the details of our hardware (on the universal gate set that we have chosen).

Notice that a quantum computer can easily simulate a probabilistic classical computer: it can prepare $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and then project to $\{|0\rangle, |1\rangle\}$, generating a random bit. Therefore BQP certainly contains the class BPP. But as we discussed in Chapter 1, it seems quite reasonable to expect that BQP is actually larger than BPP, because a probabilistic classical computer cannot easily simulate a quantum computer. The fundamental difficulty is that the Hilbert space of n qubits is huge, of dimension 2^n , and hence the mathematical description of a typical vector in the space is exceedingly complex.

Our second issue is to better characterize the resources needed to simulate a quantum computer on a classical computer. We will see that, despite the vastness of Hilbert space, a classical computer can simulate an *n*-qubit quantum computer even if limited to an amount of memory space that is polynomial in *n*. This means the BQP is contained in the complexity class PSPACE, the decision problems that can be solved with polynomial space, but may require exponential time. We also know that NP is contained in PSPACE, because we can determine whether a verifier V(x, y) accepts the input *x* for any witness *y* by running the verifier for all possible witnesses. Though there are an exponential number of candidate witnesses to interrogate, each one can be checked in polynomial time and space.

The third important issue we should address is *accuracy*. The class BQP is defined formally under the idealized assumption that quantum gates can be executed with perfect precision. Clearly, it is crucial to relax this assumption in any realistic implementation of quantum computation. A polynomial size quantum circuit family that solves a hard problem would not be of much interest if the quantum gates in the circuit were required to have exponential accuracy. In fact, we will show that this is not the case. An idealized *T*-gate quantum circuit can be simulated with acceptable accuracy by noisy gates, provided that the error probability per gate scales like 1/T.

The fourth important issue is *coverage*. We saw that polynomial-size classical circuits can reach only a tiny fraction of all Boolean functions, because there are many more functions than circuits. A similar issue arises for unitary transformations — the unitary group acting on n qubits is vast, and there are not nearly enough polynomial-size quantum circuits to explore it thoroughly. Most quantum states of n qubits can never

be realized in Nature, because they cannot be prepared using reasonable resources.

Despite this limited reach of polynomial-size quantum circuits, quantum computers nevertheless pose a serious challenge to the strong Church– Turing thesis, which contends that any physically reasonable model of computation can be simulated by probabilistic *classical* circuits with at worst a polynomial slowdown. We have good reason to believe that classical computers are unable in general to simulate quantum computers efficiently, in complexity theoretic terms that

$$BPP \neq BQP, \tag{5.51}$$

yet this remains an unproven conjecture. Proving BPP \neq BQP is a great challenge, and no proof should be expected soon. Indeed, a corollary would be

$$BPP \neq PSPACE, \tag{5.52}$$

which would settle a long-standing and pivotal open question in classical complexity theory.

It might be less unrealistic to hope for a proof that BPP \neq BQP follows from another standard conjecture of complexity theory such as P \neq NP, though no such proof has been found so far. The most persuasive evidence we have suggesting that BPP \neq BQP is that there are some problems which *seem* to be hard for classical circuits yet can be solved efficiently by quantum circuits.

It seems likely, then, that the classification of complexity will be different depending on whether we use a classical computer or a quantum computer to solve problems. If such a separation really holds, it is the quantum classification that should be regarded as the more fundamental, for it is better founded on the physical laws that govern the universe.

5.3.1 Accuracy

Let's discuss the issue of accuracy. We imagine that we wish to implement a computation in which the quantum gates U_1, U_2, \ldots, U_T are applied sequentially to the initial state $|\varphi_0\rangle$. The state prepared by our ideal quantum circuit is

$$|\varphi_T\rangle = \boldsymbol{U}_T \boldsymbol{U}_{T-1} \dots \boldsymbol{U}_2 \boldsymbol{U}_1 |\varphi_0\rangle. \tag{5.53}$$

But in fact our gates do not have perfect accuracy. When we attempt to apply the unitary transformation U_t , we instead apply some "nearby" unitary transformation \tilde{U}_t . If we wish to include environmental decoherence in our model of how the actual unitary deviates from the ideal one, we may regard \tilde{U}_t as a transformation acting jointly on the system and environment, where the ideal unitary is a product $U_t \otimes V_t$, with U_t acting on the computer and V_t acting on the environment.

The errors cause the actual state of the computer to wander away from the ideal state. How far does it wander? After one step, the ideal state would be

$$|\varphi_1\rangle = \boldsymbol{U}_1 |\varphi_0\rangle. \tag{5.54}$$

But if the actual transformation $\tilde{m{U}}_1$ where applied instead the state would be

$$\hat{\boldsymbol{U}}_1|\varphi_0\rangle = |\varphi_1\rangle + |E_1\rangle, \qquad (5.55)$$

where

$$|E_1\rangle = (\tilde{\boldsymbol{U}}_1 - \boldsymbol{U}_1)|\varphi_0\rangle \tag{5.56}$$

is an unnormalized vector. (We could also suppose that the initial state deviates from $|\varphi_0\rangle$, which would contribute an additional error to the computation that does not depend on the size of the circuit. We'll ignore that error because we are trying to understand how the error scales with the circuit size.)

Now, if $\tilde{\boldsymbol{U}}_t$ denotes the actual gate applied at step t, $|\tilde{\varphi}_t\rangle$ denotes the actual state after t steps, and $|\varphi_t\rangle$ denotes the ideal state, then we may write

$$\begin{split} |\tilde{\varphi}_t\rangle &= \tilde{\boldsymbol{U}}_t |\tilde{\varphi}_{t-1}\rangle = \boldsymbol{U}_t |\varphi_{t-1}\rangle + \left(\tilde{\boldsymbol{U}}_t - \boldsymbol{U}_t\right) |\varphi_{t-1}\rangle + \tilde{\boldsymbol{U}}_t \left(|\tilde{\varphi}_{t-1}\rangle - |\varphi_{t-1}\rangle\right) \\ &= |\varphi_t\rangle + |E_t\rangle + \tilde{\boldsymbol{U}}_t \left(|\tilde{\varphi}_{t-1}\rangle - |\varphi_{t-1}\rangle\right), \end{split}$$
(5.57)

where $|E_t\rangle = \left(\tilde{\boldsymbol{U}}_t - \boldsymbol{U}_t\right)|\varphi_{t-1}\rangle$. Hence,

$$\begin{aligned} |\tilde{\varphi}_2\rangle &= \tilde{U}_2 |\tilde{\varphi}_1\rangle = |\varphi_2\rangle + |E_2\rangle + \tilde{U}_2 |E_1\rangle, \\ |\tilde{\varphi}_3\rangle &= \tilde{U}_3 |\tilde{\varphi}_2\rangle = |\varphi_3\rangle + |E_3\rangle + \tilde{U}_3 |E_2\rangle + \tilde{U}_3 \tilde{U}_2 |E_1\rangle, \end{aligned}$$
(5.58)

and so forth, and after T steps we obtain

$$|\tilde{\varphi}_T\rangle = |\varphi_T\rangle + |E_T\rangle + \tilde{\boldsymbol{U}}_T |E_{T-1}\rangle + \tilde{\boldsymbol{U}}_T \tilde{\boldsymbol{U}}_{T-1} |E_{T-2}\rangle + \ldots + \tilde{\boldsymbol{U}}_T \tilde{\boldsymbol{U}}_{T-1} \ldots \tilde{\boldsymbol{U}}_2 |E_1\rangle.$$
 (5.59)

Thus we have expressed the difference between $|\tilde{\varphi}_T\rangle$ and $|\varphi_T\rangle$ as a sum of T remainder terms. The worst case yielding the largest deviation of $|\tilde{\varphi}_T\rangle$ from $|\varphi_T\rangle$ occurs if all remainder terms line up in the same direction, so that the errors interfere constructively. Therefore, we conclude that

$$\| \left| \tilde{\varphi}_{T} \right\rangle - \left| \varphi_{T} \right\rangle \| \leq \| \left| E_{T} \right\rangle \| + \| \left| E_{T-1} \right\rangle \|$$

$$+ \ldots + \| \left| E_{2} \right\rangle \| + \| \left| E_{1} \right\rangle \|,$$
 (5.60)

where we have used the property $|| U|E_t \rangle || = || |E_t \rangle ||$ for any unitary U.

Let $\| \mathbf{A} \|_{\sup}$ denote the sup norm of the operator \mathbf{A} — that is, the largest eigenvalue of $\sqrt{\mathbf{A}^{\dagger}\mathbf{A}}$. We then have

$$|| |E_t\rangle ||=|| \left(\tilde{\boldsymbol{U}}_t - \boldsymbol{U}_t\right) |\varphi_{t-1}\rangle || \leq || \tilde{\boldsymbol{U}}_t - \boldsymbol{U}_t ||_{\sup}$$
(5.61)

(since $|\varphi_{t-1}\rangle$ is normalized). Now suppose that, for each value of t, the error in our quantum gate is bounded by

$$\| \tilde{\boldsymbol{U}}_t - \boldsymbol{U}_t \|_{\sup} \leq \varepsilon; \qquad (5.62)$$

then after T quantum gates are applied, we have

$$\| \left| \tilde{\varphi}_T \right\rangle - \left| \varphi_T \right\rangle \| \leq T\varepsilon; \tag{5.63}$$

in this sense, the accumulated error in the state grows linearly with the length of the computation.

The distance bounded in eq.(5.62) can equivalently be expressed as $\| \mathbf{W}_t - \mathbf{I} \|_{\sup}$, where $\mathbf{W}_t = \tilde{\mathbf{U}}_t \mathbf{U}_t^{\dagger}$. Since \mathbf{W}_t is unitary, each of its eigenvalues is a phase $e^{i\theta}$, and the corresponding eigenvalue of $\mathbf{W}_t - \mathbf{I}$ has modulus

$$|e^{i\theta} - 1| = (2 - 2\cos\theta)^{1/2}, \tag{5.64}$$

so that eq.(5.62) is the requirement that each eigenvalue satisfies

$$\cos\theta > 1 - \varepsilon^2/2, \tag{5.65}$$

(or $|\theta| \lesssim \varepsilon$, for ε small). The origin of eq.(5.63) is clear. In each time step, $|\tilde{\varphi}\rangle$ rotates relative to $|\varphi\rangle$ by (at worst) an angle of order ε , and the distance between the vectors increases by at most of order ε .

How much accuracy is good enough? In the final step of our computation, we perform an orthogonal measurement, and the probability of outcome a, in the ideal case, is

$$p(a) = |\langle a | \varphi_T \rangle|^2. \tag{5.66}$$

Because of the errors, the actual probability is

$$\tilde{p}(a) = |\langle a | \tilde{\varphi}_T \rangle|^2.$$
(5.67)

It is shown in Exercise 2.5 that the L^1 distance between the ideal and actual probability distributions satisfies

$$\frac{1}{2}\|\tilde{p}-p\|_1 = \frac{1}{2}\sum_a |\tilde{p}(a)-p(a)| \leq \||\tilde{\varphi}_T\rangle - |\varphi_T\rangle\| \leq T\varepsilon.$$
(5.68)

Therefore, if we keep $T\varepsilon$ fixed (and small) as T gets large, the error in the probability distribution also remains fixed (and small).

If we use a quantum computer to solve a decision problem, we want the actual quantum circuit to get the right answer with success probability $\frac{1}{2} + \tilde{\delta}$, where $\tilde{\delta}$ is a positive constant. If the ideal quantum circuit contains T gates and has success probability $\frac{1}{2} + \delta$, where $\delta > 0$, eq.(5.68) shows that $\tilde{\delta}$ is also positive provided $\varepsilon < \delta/T$. We should be able to solve hard problems using quantum computers as long as we can improve the accuracy of the gates linearly with the circuit size. This is still a demanding requirement, since performing very accurate quantum gates is a daunting challenge for the hardware builder. Fortunately, we will be able to show, using the theory of quantum fault tolerance, that *physical* gates with constant accuracy (independent of T) suffice to achieve *logical* gates acting on encoded quantum states with accuracy improving like 1/T, as is required for truly scalable quantum computing.

5.3.2 $BQP \subseteq PSPACE$

A randomized classical computer can simulate any quantum circuit if we grant the classical computer enough time and storage space. But how much memory does the classical computer require? Naively, since the simulation of an *n*-qubit circuit involves manipulating matrices of size 2^n , it may seem that an amount of memory space exponential in *n* is needed. But we will now show that the classical simulation of a quantum computer can be done to acceptable accuracy (albeit very slowly!) in polynomial space. This means that the quantum complexity class BQP is contained in the class PSPACE of problems that can be solved with polynomial space on a classical computer.

The object of the randomized classical simulation is to sample from a probability distribution that closely approximates the distribution of measurement outcomes for the specified quantum circuit. We will actually exhibit a classical simulation that performs a potentially harder task estimating the probability p(a) for each possible outcome a of the final measurement, which can be expressed as

$$p(a) = |\langle a|\boldsymbol{U}|0\rangle|^2, \qquad (5.69)$$

where

$$\boldsymbol{U} = \boldsymbol{U}_T \boldsymbol{U}_{T-1} \dots \boldsymbol{U}_2 \boldsymbol{U}_1, \tag{5.70}$$

is a product of T quantum gates. Each U_t , acting on the n qubits, can be represented by a $2^n \times 2^n$ unitary matrix, characterized by the complex matrix elements

$$\langle y|\boldsymbol{U}_t|x\rangle,$$
 (5.71)

where $x, y \in \{0, 1, ..., 2^n - 1\}$. Writing out the matrix multiplication explicitly, we have

$$\langle a|\boldsymbol{U}|0\rangle = \sum_{\{x_t\}} \langle a|\boldsymbol{U}_T|x_{T-1}\rangle \langle x_{T-1}|\boldsymbol{U}_{T-1}|x_{T-2}\rangle \dots$$
$$\dots \langle x_2|\boldsymbol{U}_2|x_1\rangle \langle x_1|\boldsymbol{U}_1|0\rangle.$$
(5.72)

Eq.(5.72) is a sort of "path integral" representation of the quantum computation – the probability amplitude for the final outcome a is expressed as a coherent sum of amplitudes for each of a vast number $(2^{n(T-1)})$ of possible computational paths that begin at 0 and terminate at a after Tsteps.

Our classical simulator is to add up the $2^{n(T-1)}$ complex numbers in eq.(5.72) to compute $\langle a|\mathbf{U}|0\rangle$. The first problem we face is that finite size classical circuits do integer arithmetic, while the matrix elements $\langle y|\mathbf{U}_t|x\rangle$ need not be rational numbers. The classical simulator must therefore settle for an approximate calculation to reasonable accuracy. Each term in the sum is a product of T complex factors, and there are $2^{n(T-1)}$ terms in the sum. The accumulated errors are sure to be small if we express the matrix elements to m bits of accuracy, with m large compared to $nT \log T$. Therefore, we can replace each complex matrix element by pairs of signed integers — the binary expansions, each m bits long, of the real and imaginary parts of the matrix element.

Our simulator will need to compute each term in the sum eq.(5.72) and accumulate a total of all the terms. But each addition requires only a modest amount of scratch space, and furthermore, since only the accumulated subtotal need be stored for the next addition, not much space is needed to sum all the terms, even though there are exponentially many.

So it only remains to consider the evaluation of a typical term in the sum, a product of T matrix elements. We will require a classical circuit that evaluates

$$\langle y|\boldsymbol{U}_t|x\rangle;$$
 (5.73)

this circuit receives the 2n-bit input (x, y), and outputs the 2m-bit value of the (complex) matrix element. Given a circuit that performs this function, it will be easy to build a circuit that multiplies the complex numbers together without using much space.

This task would be difficult if U_t were an arbitrary $2^n \times 2^n$ unitary transformation. But now we may appeal to the properties we have demanded of our quantum gate set — the gates from a discrete set, and each gate acts on a bounded number of qubits. Because there are a fixed finite number of gates, there are only a fixed number of gate subroutines that our simulator needs to be able to call. And because the gate acts on only a few qubits, nearly all of its matrix elements vanish (when n is large), and the value $\langle y | \boldsymbol{U}_t | x \rangle$ can be determined (to the required accuracy) by a simple circuit requiring little memory.

For example, in the case of a single-qubit gate acting on the first qubit, we have

$$\langle y_{n-1} \dots y_1 y_0 | \boldsymbol{U}_t | x_{n-1} \dots x_1 x_0 \rangle = 0 \text{ if } y_{n-1} \dots y_1 \neq x_{n-1} \dots x_1.$$
 (5.74)

A simple circuit can compare x_1 with y_1, x_2 with y_2 , *etc.*, and output zero if the equality is not satisfied. In the event of equality, the circuit outputs one of the four complex numbers

$$\langle y_0 | \boldsymbol{U}_t | x_0 \rangle,$$
 (5.75)

to m bits of precision. A simple classical circuit can encode the 8m bits of this 2×2 complex-valued matrix. Similarly, a simple circuit, requiring only space polynomial in m, can evaluate the matrix elements of any gate of fixed size.

We see, then, that a classical computer with memory space scaling like $nT \log T$ suffices to simulate a quantum circuit with T gates acting on n qubits. If we wished to consider quantum circuits with superpolynomial size T, we would need a lot of memory, but for a quantum circuit families with size poly(n), a polynomial amount of space is enough. We have shown that BQP \subseteq PSPACE.

But it is also evident that the simulation we have described requires exponential time, because we need to evaluate the sum of $2^{n(T-1)}$ complex numbers (where each term in the sum is a product of T complex numbers). Though most of these terms vanish, there are still an exponentially large number of nonvanishing terms to sum.

5.3.3 Most unitary transformations require large quantum circuits

We saw that any Boolean function can be computed by an exponentialsize classical circuit, and also that exponential-size circuits are needed to compute most functions. What are the corresponding statements about unitary transformations and quantum circuits? We will postpone for now consideration of how large a quantum circuit *suffices* to reach any unitary transformation, focusing instead on showing that exponential-size quantum circuits are *required* to reach most unitaries.

The question about quantum circuits is different than the corresponding question about classical circuits because there is a finite set of Boolean functions acting on n input bits, and a continuum of unitary transformations acting on n qubits. Since the quantum circuits are countable (if the quantum computer's gate set is finite), and the unitary transformations are not, we can't reach arbitrary unitaries with finite-size circuits. We'll be satisfied to accurately *approximate* an arbitrary unitary.

As noted in our discussion of quantum circuit accuracy, to ensure that we have a good approximation in the L^1 norm to the probability distribution for any measurement performed after applying a unitary transformation, it suffices for the actual unitary \tilde{U} to be close to the ideal unitary U in the sup norm. Therefore we will say that \tilde{U} is δ -close to Uif $\|\tilde{U} - U\|_{\sup} \leq \delta$. How large should the circuit size T be if we want to approximate any n-qubit unitary to accuracy δ ?

If we imagine drawing a ball of radius δ (in the sup norm) centered at each unitary achieved by some circuit with T gates, we want the balls to cover the unitary group U(N), where $N = 2^n$. The number N_{balls} of balls needed satisfies

$$N_{\text{balls}} \ge \frac{\text{Vol}(U(N))}{\text{Vol}(\delta-\text{ball})},$$
(5.76)

where $\operatorname{Vol}(U(N))$ means the total volume of the unitary group and $\operatorname{Vol}(\delta-\operatorname{ball})$ means the volume of a single ball with radius δ . The geometry of U(N) is actually curved, but we may safely disregard that subtlety — all we need to know is that U(N) contains a ball centered at the identity element with a small but constant radius C (independent of N). Ignoring the curvature, because U(N) has real dimension N^2 , the volume of this ball (a lower bound on the volume of U(N)) is $\Omega_{N^2}C^{N^2}$, where Ω_{N^2} denotes the volume of a unit ball in flat space; likewise, the volume of a δ -ball is $\Omega_{N^2}\delta^{N^2}$. We conclude that

$$N_{\text{balls}} \ge \left(\frac{C}{\delta}\right)^{N^2}.$$
 (5.77)

On the other hand, if our universal set contains a constant number of quantum gates (independent of n), and each gate acts on no more than k qubits, where k is a constant, then the number of ways to choose the quantum gate at step t of a circuit is no more than constant $\times \binom{n}{k} = \text{poly}(n)$. Therefore the number N_T of quantum circuits with T gates acting on n qubits is

$$N_T \le (\operatorname{poly}(n))^T \,. \tag{5.78}$$

We conclude that if we want to reach every element of U(N) to accuracy δ with circuits of size T, hence $N_T \geq N_{\text{balls}}$, we require

$$T \ge 2^{2n} \frac{\log(C/\delta)}{\log(\operatorname{poly}(n))};\tag{5.79}$$

the circuit size must be exponential. With polynomial-size quantum circuits, we can achieve a good approximation to unitaries that occupy only an exponentially small fraction of the volume of $U(2^n)$!

Reaching any desired quantum state by applying a suitable quantum circuit to a fixed initial (e.g., product) state is easier than reaching any desired unitary, but still hard, because the volume of the 2^n -dimensional *n*-qubit Hilbert space is exponential in *n*. Hence, circuits with size exponential in *n* are required. Future quantum engineers will know the joy of exploring Hilbert space, but no matter how powerful their technology, most quantum states will remain far beyond their grasp. It's humbling.

5.4 Universal quantum gates

We must address one more fundamental question about quantum computation; how do we construct an adequate set of quantum gates? In other words, what constitutes a universal quantum computer?

We will find a pleasing answer. Any generic two-qubit gate suffices for universal quantum computation. That is, for all but a set of measure zero of 4×4 unitary matrices, if we can apply that matrix to any pair of qubits, then we can construct a circuit acting on n qubits which computes a transformation coming as close as we please to any element of $U(2^n)$.

Mathematically, this is not a particularly deep result, but physically it is significant. It means that, in the quantum world, as long as we can devise a generic interaction between any two qubits, and we can implement that interaction accurately, we can build up any quantum computation, no matter how complex. Nontrivial computation is ubiquitous in quantum theory.

Aside from this general result, it is also of some interest to exhibit particular universal gate sets that might be particularly easy to implement physically. We will discuss a few examples.

5.4.1 Notions of universality

In our standard circuit model of quantum computation, we imagine that our circuit has a finite set of "hard-wired" quantum gates

$$\mathcal{G} = \{ \boldsymbol{U}_1, \boldsymbol{U}_2, \dots, \boldsymbol{U}_m \}, \tag{5.80}$$

where U_j acts on k_j qubits, and $k_j \leq k$ (a constant) for each j. Normally we also assume that the gate U_j can be applied to any k_j of the nqubits in the computer. Actually, placing some kind of geometric locality constraints on the gates would not drastically change our analysis of complexity, as long as we can construct (a good approximation to a) a **SWAP** gate (which swaps the positions of two neighboring qubits) using our gate set. If we want to perform U_j on k_j qubits that are widely separated, we may first perform a series of **SWAP** gates to bring the qubits together,
then perform the gate, and finally perform **SWAP** gates to return the qubits to their original positions.

When we say the gate set \mathcal{G} is *universal* we mean that the unitary transformations that can be constructed as quantum circuits using this gate set are *dense* in the unitary group $U(2^n)$, up to an overall phase. That is for any $V \in U(2^n)$ and any $\delta > 0$, there is a unitary \tilde{V} achieved by a finite circuit such that

$$\|V - e^{i\phi}V\|_{\sup} \le \delta \tag{5.81}$$

for some phase $e^{i\phi}$. (It is natural to use the sup norm to define the deviation of the circuit from the target unitary, but we would reach similar conclusions using any reasonable topology on $U(2^n)$.) Sometimes it is useful to relax this definition of universality; for example we might settle for *encoded universality*, meaning that the circuits are dense not in $U(2^n)$ but rather some subgroup U(N), where N is exponential (or at least superpolynomial) in n.

There are several variations on the notion of universality that are noteworthy, because they illuminate the general theory or are useful in applications.

(1) Exact universality. If we are willing to allow uncountable gate sets, then we can assert that for certain gate sets we can construct a circuit that achieves an arbitrary unitary transformation exactly. We will see that two-qubit gates are exactly universal — any element of $U(2^n)$ can be constructed as a finite circuit of two qubit gates. Another example is that the two-qubit CNOT gate, combined with arbitrary single-qubit gates, is exactly universal (Exercise 5.2).

In fact the CNOT gate is not special in this respect. Any "entangling" two-qubit gate, when combined with arbitrary single-qubit gates, is universal (Exercise 5.6). We say a two-qubit gate is entangling if it maps some product state to a state which is not a product state.

An example of a two-gate which is not entangling is a "local gate" a product unitary $V = A \otimes B$; another example is the **SWAP** gate, or any gate "locally equivalent" to **SWAP**, *i.e.*, of the form

$$\boldsymbol{V} = (\boldsymbol{A} \otimes \boldsymbol{B}) \left(\mathbf{SWAP} \right) \left(\boldsymbol{C} \otimes \boldsymbol{D} \right).$$
(5.82)

In fact these are the only non-entangling two-qubit gates. Every two-qubit unitary which is *not* local or locally equivalent to \mathbf{SWAP} is entangling, and hence universal when combined with arbitrary single-qubit gates.

(2) Generic universality. Gates acting on two or more qubits which are not local are typically universal. For example, almost any two-qubit gate is universal, if the gate can be applied to any pair of the n qubits. By "almost any" we mean except for a set of measure zero in U(4).

(3) Particular finite universal gate sets. It is shown in the Exercises 5.3-5.5 that each one of the following gate sets is universal:

$$\mathcal{G} = \{ \boldsymbol{H}, \Lambda(\boldsymbol{S}) \}, \quad \{ \boldsymbol{H}, \boldsymbol{T}, \Lambda(\boldsymbol{X}) \}, \quad \{ \boldsymbol{H}, \boldsymbol{S}, \Lambda^2(\boldsymbol{X}) \}, \quad (5.83)$$

where H, S, T are the single-qubit gates

$$\boldsymbol{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1\\ 1 & -1 \end{pmatrix}, \ \boldsymbol{S} = \begin{pmatrix} e^{-i\pi/4} & 0\\ 0 & e^{i\pi/4} \end{pmatrix}, \ \boldsymbol{T} = \begin{pmatrix} e^{-i\pi/8} & 0\\ 0 & e^{i\pi/8} \end{pmatrix}.$$
(5.84)

In Bloch sphere language, the "Hadamard gate" $\boldsymbol{H} = \frac{1}{\sqrt{2}} (\boldsymbol{X} + \boldsymbol{Z})$ is a rotation by π about the axis $\hat{x} + \hat{z}$, $\boldsymbol{S} = \exp\left(-i\frac{\pi}{4}\boldsymbol{Z}\right)$ is a rotation by $\pi/2$ about the \hat{z} axis, and $\boldsymbol{T} = \exp\left(-i\frac{\pi}{8}\boldsymbol{Z}\right)$ is a rotation by $\pi/4$ about the \hat{z} axis. By $\Lambda(\boldsymbol{S})$ we mean the two-qubit in which \boldsymbol{S} is applied to the target qubit iff the control qubit is $|1\rangle$. More generally, we use the notation $\Lambda(\boldsymbol{U})$, where \boldsymbol{U} is a single-qubit gate, to denote the two-qubit gate

$$\Lambda(\boldsymbol{U}) = |0\rangle\langle 0| \otimes \boldsymbol{I} + |1\rangle\langle 1| \otimes \boldsymbol{U};$$
(5.85)

likewise we use $\Lambda^2(U)$ to denote the three-qubit gate

$$\Lambda^{2}(\boldsymbol{U}) = (\boldsymbol{I} - |11\rangle\langle 11|) \otimes \boldsymbol{I} + |11\rangle\langle 11| \otimes \boldsymbol{U}, \qquad (5.86)$$

etc.

That particular finite gates sets are universal is especially important in the theory of quantum fault tolerance, in which highly accurate logical gates acting on encoded quantum states are constructed from noisy physical gates. As we'll discuss in Chapter 8, only a discrete set of logical gates can be well protected against noise, where the set depends on how the quantum information is encoded. The goal of fault-tolerant gate constructions is to achieve a universal set of such protected gates.

(4) Efficient circuits of universal gates. The above results concern only the "reachability" of arbitrary *n*-qubit unitaries; they say nothing about the circuit size needed for a good approximation. Yet the circuit size is highly relevant if we want to approximate one universal gate set by using another one, or if we want to approximate the steps in an ideal quantum algorithm to acceptable accuracy.

We already know that circuits with size exponential in n are needed to approximate arbitrary n-qubit unitaries using a finite gate set. However, we will see that, for any fixed k, a k-qubit unitary can be approximated to accuracy ε using a circuit whose size scales with the error like polylog $(1/\varepsilon)$. This result, the *Solovay-Kitaev theorem*, holds for any universal gate set which is "closed under inverse" — that is, such that the inverse of each gate in the set can be constructed exactly using a finite circuit. The Solovay-Kitaev theorem (which we prove in §5.4.4) tells us that one universal gate set can accurately approximate another one at a modest cost; therefore a characterization of the complexity of a computation based on quantum circuit size is not very sensitive to how the universal gate set is chosen. For example, suppose I build a unitary transformation U using T gates chosen from gate set \mathcal{G}_1 , and I want to approximate U to constant accuracy ε using gates chosen from gate set \mathcal{G}_2 . It will suffice to approximate each gate from \mathcal{G}_1 to accuracy ε/T , which can be achieved using a circuit of polylog (T/ε) gates from \mathcal{G}_2 . Therefore U can be approximated with all together O(T polylog(T)) \mathcal{G}_2 gates.

Another consequence of the Solovay-Kitaev theorem concerns our conclusion that polynomial-size circuits can reach (to constant accuracy) only a tiny fraction of $U(2^n)$. How is the conclusion modified if we build circuits using arbitrary k-qubit unitaries (where k is constant) rather than gates chosen from a finite gate set? Because approximating the k-qubit unitaries using the finite gate set inflates the circuit size by only a polylog(T) factor, if we can achieve an accuracy- δ approximation using a circuit of size T built from arbitrary k-qubit unitaries, then we can also achieve an accuracy-(2δ) approximation using a circuit of size T polylog(T/δ) built from a finite gate set. Thus the criterion eq.(5.79) for reaching all unitaries to accuracy δ using circuits of size T constructed from the finite gate set is replaced by

$$T \operatorname{polylog}(T/\delta) \ge 2^{2n} \frac{\log(C/2\delta)}{\log n}.$$
 (5.87)

if we use circuits constructed from arbitrary k-qubit gates. The required circuit size is smaller than exponential by only a poly(n) factor. The group $U(2^n)$ is unimaginably vast not because we are limited to a discrete set of gates, but rather because we are unable to manipulate more than a constant number of qubits at a time.

5.4.2 Two-qubit gates are exactly universal

We will show in two steps that an arbitrary element of $U(2^n)$ can be achieved by a finite circuit of two-qubit gates. First we will show how to express an element of U(N) as a product of "2 × 2" unitaries; then we will show how to obtain any 2 × 2 unitary from a circuit of two-qubit unitaries.

What is a 2×2 unitary? Fix a standard orthonormal basis $\{|0\rangle, |1\rangle, |2\rangle, \dots, |N-1\rangle\}$ for an N-dimensional space. We say a unitary transformation U is 2×2 if it acts nontrivially only in the two-dimensional subspace spanned by two basis elements $|i\rangle$ and $|j\rangle$; that is, U decomposes

as a direct sum

$$U = U^{(2)} \oplus I^{(N-2)},$$
 (5.88)

where $U^{(2)}$ is a 2 × 2 unitary matrix acting on the span of $|i\rangle$ and $|j\rangle$, and $I^{(N-2)}$ is the identity matrix acting on the complementary (N-2)dimensional subspace.

We should be careful not to confuse a 2×2 unitary with a two-qubit unitary acting on the *n*-qubit space of dimension $N = 2^n$. A two-qubit unitary U decomposes as a tensor *product*

$$U = U^{(4)} \otimes I^{(2^{n-2})},$$
 (5.89)

where $U^{(4)}$ is a 4×4 unitary matrix acting on a pair of qubits, and $I^{(2^{n-2})}$ is the identity matrix acting on the remaining n-2 qubits. We can regard the two-qubit unitary as a direct sum of 2^{n-2} 4×4 blocks, with each block labeled by a basis state of the (n-2)-qubit Hilbert space, and $U^{(4)}$ acting on each block.

Let's see how to express $U \in U(N)$ as a product of 2×2 unitaries. Consider the action of U on the basis state $|0\rangle$:

$$\boldsymbol{U}|0\rangle = \sum_{i=0}^{N-1} a_i |i\rangle.$$
(5.90)

We can see that $U|0\rangle$ can be written as $W_0|0\rangle$, where W_0 is a product of $(N-1) \ 2 \times 2$ unitaries which act as follows:

$$|0\rangle \mapsto a_0|0\rangle + b_0|1\rangle,$$

$$b_0|1\rangle \mapsto a_1|1\rangle + b_1|2\rangle,$$

$$b_1|2\rangle \mapsto a_2|2\rangle + b_2|3\rangle,$$

$$\cdots$$

$$b_{N-2}|N-2\rangle \mapsto a_{N-2}|N-2\rangle + a_{N-1}|N-1\rangle.$$

(5.91)

Next define $U_1 = W_0^{-1}U$, and note that $U_1|0\rangle = |0\rangle$, so U_1 acts nontrivially only in the (N-1)-dimensional span of $\{|1\rangle, |2\rangle, \dots, |N-1\rangle\}$. By the same construction as above, we construct W_1 as a product of (N-2) 2 × 2 unitaries such that $W_1|0\rangle = |0\rangle$ and $W_1|1\rangle = U_1|1\rangle$, then define $U_2 = W_1^{-1}U_1$ such that U_2 preserves both $|0\rangle$ and $|1\rangle$. Proceeding in this way we construct W_2, W_3, \dots, W_{N-2} such that

$$W_{N-2}^{-1}W_{N-3}^{-1}\dots W_{1}^{-1}W_{0}^{-1}U = I;$$
 (5.92)

that is, we may express \boldsymbol{U} as

$$\boldsymbol{U} = \boldsymbol{W}_0 \boldsymbol{W}_1 \dots \boldsymbol{W}_{N-3} \boldsymbol{W}_{N-2} , \qquad (5.93)$$

a product of $(N-1) + (N-2) + \dots + 2 + 1 = \frac{1}{2}N(N-1) \ 2 \times 2$ unitaries.

Now it remains to show that we can construct any 2×2 unitary as a circuit of two-qubit unitaries. It will be helpful to notice that the threequbit gate $\Lambda^2(\boldsymbol{U}^2)$ can be constructed as a circuit of $\Lambda(\boldsymbol{U})$, $\Lambda(\boldsymbol{U}^{\dagger})$, and $\Lambda(\boldsymbol{X})$ gates. Using the notation



for the $\Lambda(U)$ gate, the circuit



does the job. We can check that the power of \boldsymbol{U} applied to the third qubit is

$$y - (x \oplus y) + x = y - (x + y - 2xy) + x = 2xy.$$
 (5.94)

That is, U^2 is applied if x = y = 1, and the identity is applied otherwise; hence this circuit achieves the $\Lambda^2(U^2)$ gate. Since every unitary V has a square root U such that $V = U^2$, the construction shows that, using two-qubit gates, we can achieve $\Lambda^2(V)$ for any single-qubit V.

Generalizing this construction, we can find a circuit that constructs $\Lambda^m(U^2)$ using $\Lambda^{m-1}(U)$, $\Lambda^{m-1}(X)$, $\Lambda(U)$, and $\Lambda(U^{\dagger})$ gates. If we replace the $\Lambda(X)$ gates in the previous circuit by $\Lambda^{m-1}(X)$ gates, and replace the last $\Lambda(U)$ gate by $\Lambda^{n-1}(U)$, then, if we denote the *m* control bits by $x_1, x_2, x_3, \ldots x_m$, the power of U applied to the last qubit is

$$x_{m} + x_{1}x_{2}x_{3}\dots x_{m-1} - (x_{m} \oplus x_{1}x_{2}x_{3}\dots x_{m-1})$$

= $x_{m} + x_{1}x_{2}x_{3}\dots x_{m-1}$
- $(x_{m} + x_{1}x_{2}x_{3}\dots x_{m-1} - 2x_{1}x_{2}x_{3}\dots x_{m-1}x_{m})$
= $2x_{1}x_{2}x_{3}\dots x_{m-1}x_{m},$ (5.95)

where we have used the identity $x \oplus y = x + y - 2xy$. Now U^2 is applied if $x_1 = x_2 = \cdots = x_m = 1$ and the identity is applied otherwise; this circuit achieves the $\Lambda^m(U^2)$ gate.

Using the construction recursively, we see that with two-qubit gates we can construct $\Lambda^2(\mathbf{V})$ for any \mathbf{V} , then with these gates and two-qubit gates we can construct $\Lambda^3(\mathbf{V})$ for any \mathbf{V} , which allows us to construct $\Lambda^4(\mathbf{V})$ for any \mathbf{V} and so on. We have shown, therefore, how to construct the *n*-qubit gate $\Lambda^{n-1}(\mathbf{V})$ for any \mathbf{V} using a circuit of two-qubit gates.

To complete the argument showing that any element of $U(2^n)$ is a product of two-qubit unitaries, it will suffice to show that arbitrary 2×2 unitaries can be constructed from $\Lambda^{n-1}(\mathbf{V})$ and two-qubit gates. Note that $\Lambda^{n-1}(\mathbf{V})$ is, in fact, a 2×2 unitary — it applies \mathbf{V} in the twodimensional space spanned by the two computational basis states

$$\{|111...110\rangle, |111...111\rangle\}.$$
 (5.96)

If we wish to apply V in the space spanned by computational states $\{|x\rangle, |y\rangle\}$ instead, we can use a permutation Σ of the computational basis states with the action

$$\boldsymbol{\Sigma} : |x\rangle \mapsto |111\dots110\rangle , |y\rangle \mapsto |111\dots111\rangle ,$$
 (5.97)

constructing

$$\boldsymbol{\Sigma}^{-1} \circ \boldsymbol{\Lambda}^{n-1}(\boldsymbol{V}) \circ \boldsymbol{\Sigma} . \tag{5.98}$$

This is to be read from right to left, with Σ acting first and Σ^{-1} acting last. But we have already seen in §5.2.2 how to construct an arbitrary permutation of computational basis states using $\Lambda^{n-1}(\mathbf{X})$ gates and (singlequbit) NOT gates, and we now know how to construct $\Lambda^{n-1}(\mathbf{X})$ (a special case of $\Lambda^{n-1}(\mathbf{U})$) from two-qubit gates. Therefore, using two-qubit gates, we have constructed the general 2 × 2 unitary (in the computational basis) as in eq.(5.98). That completes the proof that any element of $U(2^n)$ can be achieved by a circuit of two-qubit gates. In fact we have proven a somewhat stronger result: that the two-qubit gates { $\Lambda(\mathbf{U})$ }, where \mathbf{U} is an arbitrary single-qubit gate, constitute an exactly universal gate set.

5.4.3 Finite universal gate sets

Denseness on the circle. A finite gate set is universal if circuits constructed using that gate set are dense in $U(2^n)$. As a first simple example of denseness, consider the group U(1) of rotations of the circle, *e.g.* the rotations of the Bloch sphere about the \hat{z} axis:

$$\left\{ \boldsymbol{U}(\theta) = \exp\left(i\frac{\theta}{2}\boldsymbol{\sigma}_3\right), \quad \theta \in [0, 4\pi) \right\} .$$
 (5.99)

We claim that the positive integer powers of $U(4\pi\alpha)$ are dense in U(1) if $\alpha \in [0, 1)$ is irrational. Equivalently, the points

$$\{n\alpha \pmod{1}, n = 1, 2, 3, \dots, \}$$
 (5.100)

are dense in the unit interval.

To see why, first note that the points $\{n\alpha \pmod{1}\}\$ are all distinct, since $n\alpha = m\alpha + k$ for integers k and $n \neq m$ would imply that α is a rational number $\alpha = k/(n-m)$. Now consider open intervals of width ε centered on each of the N points $\{n\alpha \pmod{1}, n = 1, 2, 3, \ldots, N\}$. For $N\varepsilon > 1$, at least two of these intervals must intersect — if all intervals were disjoint then their total length $N\varepsilon$ would exceed the length of the interval. Hence there exist distinct positive integers n and m less than $1/\varepsilon$ such that $|n - m|\alpha \pmod{1} < \varepsilon$; in other words, the positive integer $r = |n - m| < 1/\varepsilon$ satisfies $r\alpha \pmod{1} < \varepsilon$. Now the positive integer multiples of $r\alpha \pmod{1}$ are equally spaced points on the unit interval separated by less than ε . Therefore, for sufficiently large M, the intervals of width ε centered on the points $\{kr\alpha \pmod{1}, k = 1, 2, 3, \ldots, M\}$ fill the unit interval. Since ε can be any positive real number, we conclude that the points $\{n\alpha \pmod{1}, n = 1, 2, 3, \ldots\}$ are dense in the interval.

Powers of a generic gate. Generalizing this argument, consider the positive integer powers of a generic element of U(N). In a suitable basis, $U \in U(N)$ is diagonal, with eigenvalues

$$\{e^{i\theta_1/2}, e^{i\theta_2/2}, \dots, e^{i\theta_N/2}\}.$$
 (5.101)

Since rational numbers are countable and real numbers are not, for a generic U (that is, for all elements of U(N) except for a set of measure zero) each θ_i/π and θ_i/θ_j is an irrational number. For each positive integer k, the eigenvalues $\{e^{-ik\theta_i/2}, i = 1, 2, ..., N\}$ of U^k define a point on the N-dimensional torus (the product of N circles), and as k ranges over all positive integers, these points densely fill the whole N-torus. We conclude that for any generic U, the elements $\{U^k, k = 1, 2, ..., \}$ are dense in the group $U(1)^N$, *i.e.*, come as close as we please to every unitary matrix which is diagonal in the same basis as U.

Note that this argument does not provide any upper bound on how large k must be for U^k to be ε -close to any specified element of $U(1)^N$. In fact, the required value of k could be extremely large if, for some m and i, $|m\theta_i \pmod{4\pi}| \ll \varepsilon$. It might be hard (that is, require many gates) to approximate a specified unitary transformation with circuits of commuting quantum gates, because the unitary achieved by the circuit only depends on how many times each gate is applied, not on the order in which the gates are applied. It is much easier (requires fewer gates) to achieve a good approximation using circuits of noncommuting gates. If the gates are noncommuting, then the order in which the gates are applied matters, and many more unitaries can be reached by circuits of specified size than if the gates are noncommuting.

Reaching the full Lie algebra. Suppose we can construct the two gates $U = \exp(iA), V = \exp(iB) \in U(N)$, where A and B are $N \times N$ Hermitian matrices. If these are generic gates, positive powers of U come as close as we please to $e^{i\alpha A}$ for any real α and positive powers of V come as close as we please to $e^{i\beta B}$ for any real β . That is enough to ensure that there is a finite circuit constructed from U and V gates that comes as close as we please to e^{iC} , where C is any Hermitian element of the Lie algebra generated by A and B.

We say that a unitary transformation U is *reachable* if for any $\varepsilon > 0$ there is a finite circuit achieving \tilde{U} which is ε -close to U in the sup norm. Noting that

$$\lim_{n \to \infty} (e^{i\alpha \mathbf{A}/n} e^{i\beta \mathbf{B}/n})^n = \lim_{n \to \infty} \left(1 + \frac{i}{n} (\alpha \mathbf{A} + \beta \mathbf{B}) + O\left(\frac{1}{n^2}\right) \right)^n$$
$$= e^{i(\alpha \mathbf{A} + \beta \mathbf{B})}, \tag{5.102}$$

we see that any $e^{i(\alpha A + \beta B)}$ is reachable if each $e^{i\alpha A/n}$ and $e^{i\beta B/n}$ is reachable. Furthermore, because

$$\lim_{n \to \infty} \left(e^{i\mathbf{A}/\sqrt{n}} e^{i\mathbf{B}/\sqrt{n}} e^{-i\mathbf{A}/\sqrt{n}} e^{-i\mathbf{B}/\sqrt{n}} \right)^n$$
$$= \lim_{n \to \infty} \left(1 - \frac{1}{n} \left(\mathbf{A}\mathbf{B} - \mathbf{B}\mathbf{A} \right) + O\left(\frac{1}{n^{3/2}}\right) \right)^n = e^{-[\mathbf{A}, \mathbf{B}]}, \qquad (5.103)$$

we see that $e^{-[\mathbf{A}, \mathbf{B}]}$ is also reachable.

For example, positive integer powers of a generic element of SU(2) allow us to reach a U(1) subgroup; if we orient our axes on the Bloch sphere appropriately, this is the subgroup generated by the Pauli operator Z. Positive integer powers of a second generic element of SU(2) allow us to reach a different U(1) subgroup, generated by $\tilde{X} = X + \gamma Z$ (for some real γ) with an appropriate choice of axes. Because [Z, X] = iY, the elements $\{Z, \tilde{X}, -i[Z, \tilde{X}]\}$ span the three-dimensional SU(2) Lie algebra. It follows that circuits built from any two generic elements of SU(2) suffice to reach any element of SU(2).

This observation applies to higher-dimensional Lie algebras as well. For example, the SU(4) Lie algebra is 15 dimensional. It contains various lower-dimensional subalgebras, such as the Lie algebras of the SU(4) subgroups $U(1)^3$, $SU(2) \times SU(2) \times U(1)$, $SU(3) \times U(1)$, etc. But two generic elements of the SU(4) Lie algebra already suffice to general the full Lie algebra. The generated algebra closes on one of the lower-dimensional subalgebras only if some nested commutators vanish "by accident," a criterion satisfied by only a set of measure zero among all pairs of SU(4)generators. Actually, as we have already noted, a generic element of SU(4) allows us to reach the torus $U(1)^3$, and no nontrivial subgroup of SU(4) contains two generic $U(1)^3$ subgroups. Therefore, circuits built from two generic two-qubit gates suffice to reach any two-qubit gate (up to an overall phase). And since we can reach any element of $U(2^n)$ with two-qubit gates, a pair of generic two-qubit gates provides a universal gate set, assuming we can apply the gates to any pair of qubits.

But in fact just one generic two-qubit gate is already enough, if we are free to choose not just the pair of qubits on which the gate acts but also the ordering of the qubits. That is, a generic two-qubit gate does not commute with the operator **SWAP** which interchanges the two qubits. If U is a generic two-qubit gate, then

$$\boldsymbol{V} = \mathbf{SWAP} \circ \boldsymbol{U} \circ \mathbf{SWAP} \tag{5.104}$$

(the same gate applied to the same two qubits, but in the opposite order) is another two-qubit gate not commuting with U. Positive powers of U reach one $U(1)^3$ subgroup of SU(4) while positive powers of V reach a different $U(1)^3$, so that circuits built from U and V reach all of SU(4).

Even nongeneric universal gates, in particular gates whose eigenvalues are all rational multiples of π , can suffice for universality. One example discussed in the homework is the gate set {CNOT, H, T}, where H rotates the Bloch sphere by the angle π about the axis $\frac{1}{\sqrt{2}}(\hat{x} + \hat{x})$, and Trotates the Bloch sphere by the angle $\pi/4$ about the \hat{z} axis. If we replaced T with the $\pi/2$ rotation T^2 , then the gate set would not be universal; in that case the only achievable single-qubit rotations would be those in a finite subgroup of SU(2), the symmetry group of the cube. But SU(2)has few such finite nonabelian subgroups (the only finite nonabelian subgroups of the rotation group SO(3) are the symmetry groups of regular polygons and of regular three-dimensional polyhedra, the platonic solids). If the gate set reaches beyond these finite subgroups it will reach either a U(1) subgroup of SU(2) or all of SU(2).

5.4.4 The Solovay-Kitaev approximation

Up until now our discussion of universal gates has focused on reachability and has ignored complexity. But when we have a finite universal gate set, we want to know not only *whether* we can approximate a desired unitary transformation to accuracy ε , but also *how hard* it is to achieve that approximation. How large a circuit suffices? The question really has two parts. (1) Given a unitary transformation U, how large a quantum circuit is needed to construct \tilde{U} such that $\|\tilde{U} - e^{i\phi}U\|_{\sup} \leq \varepsilon$? (2) How large a *classical* circuit is needed to find the quantum circuit that achieves \tilde{U} ? We will see that, for any universal set of gates (closed under inverse) used to approximate elements of a unitary group of constant dimension, the answer to both questions is $\operatorname{polylog}(1/\varepsilon)$. We care about the answer to the second question because it would not be very useful to know that U can be well approximated by small quantum circuits if these circuits are very hard to find.

We will prove this result by devising a recursive algorithm which achieves successively better and better approximations. We say that a finite repertoire of unitary transformations \mathcal{R} is an " ε -net" in U(N) if every element of U(N) is no more than distance ε away (in the sup norm) from some element of \mathcal{R} , and we say that \mathcal{R} is "closed under inverse" if the inverse of every element of \mathcal{R} is also in \mathcal{R} . The key step of the recursive algorithm is to show that if \mathcal{R} is an ε -net, closed under inverse, then we can construct a new repertoire \mathcal{R}' , also closed under inverse, with the following properties: (1) each element of \mathcal{R}' is achieved by a circuit of at most 5 gates from \mathcal{R} . (2) \mathcal{R}' is an ε -net, where

$$\varepsilon' = C\varepsilon^{3/2} , \qquad (5.105)$$

and C is a constant.

Before explaining how this step works, let's see why it ensures that we can approximate any unitary using a quantum circuit with size polylogarithmic in the accuracy. Suppose to start with that we have found an ε_0 -net \mathcal{R}_0 , closed under inverse, where each element of \mathcal{R}_0 can be achieved by a circuit with no more than L_0 gates chosen from our universal gate set. If $\varepsilon_0 < 1/C^2$, then we can invoke the recursive step to find an ε_1 -net \mathcal{R}_1 , where $\varepsilon_1 < \varepsilon_0$, and each element of \mathcal{R}_1 can be achieved by a circuit of $L_1 = 5L_0$ gates. By repeating this step k times, we can make the error ε_k much smaller than the level-0 error ε_0 . Iterating the relation

$$C^2 \varepsilon_k = \left(C^2 \varepsilon_{k-1} \right)^{3/2} \tag{5.106}$$

k times we obtain

$$C^2 \varepsilon_k = \left(C^2 \varepsilon_0 \right)^{(3/2)^k} , \qquad (5.107)$$

and by taking logs of both sides we find

$$\left(\frac{3}{2}\right)^k = \frac{\log\left(1/C^2\varepsilon_k\right)}{\log\left(1/C^2\varepsilon_0\right)} \ . \tag{5.108}$$

After k recursive steps the circuit size for each unitary in the ε_k -net \mathcal{R}_k

is no larger than L_k where

$$L_k/L_0 = 5^k = \left(\left(\frac{3}{2}\right)^k\right)^{\log 5/\log(3/2)} = \left(\frac{\log(1/C^2\varepsilon_k)}{\log(1/C^2\varepsilon_0)}\right)^{\log 5/\log(3/2)}.$$
(5.109)

Thus the circuit size scales with the accuracy ε_k as $\left[\log(1/\varepsilon_k)\right]^{3.97}$.

Now let's see how the ε' -net \mathcal{R}' is constructed from the ε -net \mathcal{R} . For any $U \in SU(N)$ there is an element $\tilde{U} \in \mathcal{R}$ such that $||U - \tilde{U}||_{\sup} \leq \varepsilon$, or equivalently $||U\tilde{U}^{-1} - I||_{\sup} \leq \varepsilon$. Now we will find W, constructed as a circuit of 4 elements of \mathcal{R} , such that $||U\tilde{U}^{-1} - W|_{\sup} \leq \varepsilon'$, or equivalently $||U - W\tilde{U}|_{\sup} \leq \varepsilon'$. Thus U is approximated to accuracy ε' by $W\tilde{U}$, which is achieved by a circuit of 5 elements of \mathcal{R} .

We may write $U\tilde{U}^{-1} = e^{iA}$, where $A = O(\varepsilon)$. (By $A = O(\varepsilon)$ we mean $||A||_{\sup} = O(\varepsilon)$, *i.e.*, $||A||_{\sup}$ is bounded above by a constant times ε for ε sufficiently small.) It is possible to find Hermitian B, C, both $O(\varepsilon^{1/2})$, such that [B, C] = -iA. Furthermore, because \mathcal{R} is an ε -net, there is an element $e^{i\tilde{B}}$ of \mathcal{R} which is ε -close to e^{iB} , and an element $e^{i\tilde{C}}$ of \mathcal{R} which is ε -close to e^{iC} . It follows that $B - \tilde{B} = O(\varepsilon)$ and $C - \tilde{C} = O(\varepsilon)$.

Now we consider the circuit

$$\boldsymbol{W} = e^{i\boldsymbol{\tilde{B}}} e^{i\boldsymbol{\tilde{C}}} e^{-i\boldsymbol{\tilde{B}}} e^{-i\boldsymbol{\tilde{C}}} = \boldsymbol{I} - [\boldsymbol{\tilde{B}}, \boldsymbol{\tilde{C}}] + O(\varepsilon^{3/2}); \qquad (5.110)$$

the remainder term is cubic order in \tilde{B} and \tilde{C} , hence $O(\varepsilon^{3/2})$. First note that the inverse of this circuit, $e^{i\tilde{C}}e^{i\tilde{B}}e^{-i\tilde{C}}e^{-i\tilde{B}}$, can also be constructed as a size-4 circuit of gates from \mathcal{R} . Furthermore,

$$\boldsymbol{W} = I - [\boldsymbol{B} + O(\varepsilon), \boldsymbol{C} + O(\varepsilon)] + O(\varepsilon^{3/2}) = I + i\boldsymbol{A} + O(\varepsilon^{3/2})$$
$$= e^{i\boldsymbol{A}} + O(\varepsilon^{3/2}); \qquad (5.111)$$

thus W, a circuit of 4 gates from \mathcal{R} , approximates $U\tilde{U}^{-1}$ to $O(\varepsilon^{3/2})$ accuracy, as we wanted to show.

Finally, let's consider the classical computational cost of finding the quantum circuit which approximates a unitary transformation. The classical algorithm receives a unitary transformation U as input, and produces as output a quantum circuit evaluating \tilde{U} , which approximates U to accuracy ε . To improve the accuracy of the approximation to ε' , we need to call the accuracy- ε algorithm three times, to find circuits evaluating \tilde{U} , $e^{i\tilde{B}}$, and $e^{i\tilde{C}}$. Therefore, if the classical cost of the accuracy- ε algorithm is t, the classical cost of the improved accuracy- ε' algorithm is t' = 3t + constant, where the additive constant is needed to cover the cost

of tasks that do not scale with ε , such as finding the matrices **B** and **C** satisfying $[\mathbf{B}, \mathbf{C}] = -i\mathbf{A}$. After k iterations, the classical cost scales like

$$O(3^k) = O\left(\left[\log(1/\varepsilon_k) \right]^{\log 3/\log(3/2)} \right) = O\left(\left[\log(1/\varepsilon_k) \right]^{2.71} \right), \quad (5.112)$$

polylogarithmic in the accuracy achieved by the level-k version of the algorithm.

What we have accomplished is a bit surprising. By composing unitary transformations with $O(\varepsilon)$ errors we have obtained unitary transformations with smaller $O(\varepsilon^{3/2})$ errors. How could we achieve such sharp results with such blunt tools? The secret is that we have constructed our circuit so that the $O(\varepsilon)$ errors cancel, leaving only the higher-order errors. This would not have worked if \mathcal{R} had not been closed under inverse. If instead of the inverses of $e^{i\tilde{B}}$ and $e^{i\tilde{C}}$ we had been forced to use $O(\varepsilon)$ approximations to these inverses, the cancellations would not have occurred, and our quest for an improved approximation would have failed. But if our universal gate set allows us to construct the exact inverse of each element of the gate set, then we can use the Solovay-Kitaev approach to recursively improve the approximation.

This scheme works for any universal gate set that is closed under inverse. For particular gate sets improved approximations are possible. For example, the gate set $\{H, T\}$ can be used to approximate an arbitrary single-qubit unitary to accuracy ε using $O(\log(1/\varepsilon))$ gates, a substantial improvement over $O([\log(1/\varepsilon)]^{3.97})$ established by the general argument, and the circuits achieving this improved overhead cost can be efficiently constructed.

5.5 Summary

Classical circuits. The complexity of a problem can be characterized by the size of a uniform family of logic circuits that solve the problem: The problem is hard if the size of the circuit is a superpolynomial function of the size of the input, and easy otherwise. One classical universal computer can simulate another efficiently, so the classification of complexity is machine independent. The 3-bit Toffoli gate is universal for classical reversible computation. A reversible computer can simulate an irreversible computer without a significant slowdown and without unreasonable memory resources.

Quantum Circuits. Although there is no proof, it seems likely that polynomial-size quantum circuits cannot be simulated in general by polynomial-size randomized classical circuits (BQP \neq BPP); however, polynomial space is sufficient (BQP \subseteq PSPACE). A noisy quantum circuit can simulate an ideal quantum circuit of size T to acceptable accuracy if each quantum gate has an accuracy of order 1/T. Any *n*-qubit unitary transformation can be constructed from two-qubit gates. A generic two-qubit quantum gate, if it can act on any two qubits in a device, is sufficient for universal quantum computation. One universal quantum computer can simulate another to accuracy ε with a polylog $(1/\varepsilon)$ overheard cost; therefore the complexity class BQP is machine independent.

Do the Exercises to learn more about universal sets of quantum gates.

5.6 Exercises

5.1 Linear simulation of Toffoli gate.

In §5.2.2 we constructed the *n*-bit Toffoli gate $\Lambda^{n-1}(\mathbf{X})$ from 3-bit Toffoli gates ($\Lambda^2(\mathbf{X})$'s). The circuit required only one bit of scratch space, but the number of gates was exponential in *n*. With more scratch, we can substantially reduce the number of gates.

- a) Find a circuit family with $2n 5 \Lambda^2(\mathbf{X})$'s that evaluates $\Lambda^{n-1}(\mathbf{X})$. (Here n-3 scratch bits are used, which are set to 0 at the beginning of the computation and return to the value 0 at the end.)
- b) Find a circuit family with $4n 12 \Lambda^2(\mathbf{X})$'s that evaluates $\Lambda^{n-1}(\mathbf{X})$, which works irrespective of the initial values of the scratch bits. (Again the n-3 scratch bits return to their initial values, but they don't need to be set to zero at the beginning.)

5.2 An exactly universal quantum gate set.

The purpose of this exercise is to complete the demonstration that the controlled-NOT gate $\Lambda(\mathbf{X})$ and arbitrary single-qubit gates constitute an exactly universal set.

a) If U is any unitary 2×2 matrix with determinant one, find unitary A, B, and C such that

$$ABC = I \tag{5.113}$$

$$AXBXC = U. (5.114)$$

Hint: From the Euler angle construction, we know that

$$\boldsymbol{U} = \boldsymbol{R}_{z}(\psi)\boldsymbol{R}_{y}(\theta)\boldsymbol{R}_{z}(\phi), \qquad (5.115)$$

where, e.g., $\mathbf{R}_z(\phi)$ denotes a rotation about the z-axis by the angle ϕ . We also know that, e.g.,

$$\boldsymbol{X}\boldsymbol{R}_{z}(\phi)\boldsymbol{X} = \boldsymbol{R}_{z}(-\phi). \tag{5.116}$$

- b) Consider a two-qubit controlled phase gate which applies $U = e^{i\alpha}\mathbf{1}$ to the second qubit if the first qubit has value $|1\rangle$, and acts trivially otherwise. Show that it is actually a one-qubit gate.
- c) Draw a circuit using $\Lambda(\mathbf{X})$ gates and single-qubit gates that implements $\Lambda(\mathbf{U})$, where \mathbf{U} is an arbitrary 2×2 unitary transformation.

Since the argument in §5.4.2 shows that the gate set $\{\Lambda(U)\}$ is exactly universal, we have shown that $\Lambda(X)$ together with singlequbit gates are an exactly universal set.

5.3 Universal quantum gates I

In this exercise and the two that follow, we will establish that several simple sets of gates are universal for quantum computation.

The Hadamard transformation H is the single-qubit gate that acts in the standard basis $\{|0\rangle, |1\rangle\}$ as

$$\boldsymbol{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1\\ 1 & -1 \end{pmatrix} ; \qquad (5.117)$$

in quantum circuit notation, we denote the Hadamard gate as

$$-H$$

The single-qubit *phase gate* \boldsymbol{S} acts in the standard basis as

$$\boldsymbol{S} = \begin{pmatrix} 1 & 0\\ 0 & i \end{pmatrix} , \qquad (5.118)$$

and is denoted



A two-qubit controlled phase gate $\Lambda(\mathbf{S})$ acts in the standard basis $\{|00\rangle, 01\rangle, |10\rangle, |11\rangle\}$ as the diagonal 4×4 matrix

$$\Lambda(\mathbf{S}) = \operatorname{diag}(1, 1, 1, i) \tag{5.119}$$

and can be denoted



Despite this misleading notation, the gate $\Lambda(S)$ actually acts symmetrically on the two qubits:



We will see that the two gates \boldsymbol{H} and $\Lambda(\boldsymbol{S})$ comprise a *universal gate* set – any unitary transformation can be approximated to arbitrary accuracy by a quantum circuit built out of these gates.

a) Consider the two-qubit unitary transformations U_1 and U_2 defined by quantum circuits



and



Let $|ab\rangle$ denote the element of the standard basis where a labels the upper qubit in the circuit diagram and b labels the lower qubit. Write out U_1 and U_2 as 4×4 matrices in the standard basis. Show that U_1 and U_2 both act trivially on the states

$$|00\rangle, \quad \frac{1}{\sqrt{3}} (|01\rangle + |10\rangle + |11\rangle) \quad . \tag{5.120}$$

b) Thus U_1 and U_2 act nontrivially only in the two-dimensional space spanned by

$$\left\{\frac{1}{\sqrt{2}}\left(|01\rangle - |10\rangle\right), \frac{1}{\sqrt{6}}\left(|01\rangle + |10\rangle - 2|11\rangle\right)\right\} .$$
 (5.121)

Show that, expressed in this basis, they are

$$\boldsymbol{U}_1 = \frac{1}{4} \begin{pmatrix} 3+i & \sqrt{3}(-1+i) \\ \sqrt{3}(-1+i) & 1+3i \end{pmatrix} , \qquad (5.122)$$

and

$$\boldsymbol{U}_2 = \frac{1}{4} \begin{pmatrix} 3+i & \sqrt{3}(1-i) \\ \sqrt{3}(1-i) & 1+3i \end{pmatrix} .$$
 (5.123)

c) Now express the action of U_1 and U_2 on this two-dimensional subspace in the form

$$\boldsymbol{U}_1 = \sqrt{i} \left(\frac{1}{\sqrt{2}} - i \frac{1}{\sqrt{2}} \hat{\boldsymbol{n}}_1 \cdot \boldsymbol{\vec{\sigma}} \right) , \qquad (5.124)$$

and

$$\boldsymbol{U}_2 = \sqrt{i} \left(\frac{1}{\sqrt{2}} - i \frac{1}{\sqrt{2}} \hat{\boldsymbol{n}}_2 \cdot \boldsymbol{\vec{\sigma}} \right) . \tag{5.125}$$

What are the unit vectors \hat{n}_1 and \hat{n}_2 ?

d) Consider the transformation $U_2^{-1}U_1$ (Note that U_2^{-1} can also be constructed from the gates \boldsymbol{H} and $\Lambda(\boldsymbol{S})$.) Show that it performs a rotation with half-angle $\theta/2$ in the two-dimensional space spanned by the basis eq.(5.121), where $\cos(\theta/2) = 1/4$.

5.4 Universal quantum gates II

We have seen in Exercise 5.3 how to compose the quantum gates \boldsymbol{H} and $\Lambda(\boldsymbol{S})$ to perform, in a two-dimensional subspace of the fourdimensional Hilbert space of two qubits, a rotation with $\cos(\theta/2) = 1/4$. In this exercise, we will show that the angle θ is not a rational multiple of π . Equivalently, we will show that

$$e^{i\theta/2} \equiv \cos(\theta/2) + i\sin(\theta/2) = \frac{1}{4} \left(1 + i\sqrt{15}\right)$$
 (5.126)

is not a root of unity: there is no finite integer power n such that $(e^{i\theta/2})^n = 1$.

Recall that a *polynomial of degree* n is an expression

$$P(x) = \sum_{k=0}^{n} a_k x^k$$
 (5.127)

with $a_n \neq 0$. We say that the polynomial is *rational* if all of the a_k 's are rational numbers, and that it is *monic* if $a_n = 1$. A polynomial is *integral* if all of the a_k 's are integers, and an integral polynomial is *primitive* if the greatest common divisor of $\{a_0, a_1, \ldots, a_n\}$ is 1.

a) Show that the monic rational polynomial of minimal degree that has $e^{i\theta/2}$ as a root is

$$P(x) = x^2 - \frac{1}{2}x + 1 . (5.128)$$

The property that $e^{i\theta/2}$ is not a root of unity follows from the result (a) and the

Theorem If a is a root of unity, and P(x) is a monic rational polynomial of minimal degree with P(a) = 0, then P(x) is integral. Since the minimal monic rational polynomial with root $e^{i\theta/2}$ is not integral, we conclude that $e^{i\theta/2}$ is not a root of unity. In the rest of this exercise, we will prove the theorem.

b) By "long division" we can prove that if A(x) and B(x) are rational polynomials, then there exist rational polynomials Q(x)and R(x) such that

$$A(x) = B(x)Q(x) + R(x) , \qquad (5.129)$$

where the "remainder" R(x) has degree less than the degree of B(x). Suppose that $a^n = 1$, and that P(x) is a rational polynomial of minimal degree such that P(a) = 0. Show that there is a rational polynomial Q(x) such that

$$x^{n} - 1 = P(x)Q(x) . (5.130)$$

- c) Show that if A(x) and B(x) are both primitive integral polynomials, then so is their product C(x) = A(x)B(x). Hint: If $C(x) = \sum_k c_k x^k$ is not primitive, then there is a prime number p that divides all of the c_k 's. Write $A(x) = \sum_l a_l x^l$, and $B(x) = \sum_m b_m x^m$, let a_r denote the coefficient of lowest order in A(x) that is not divisible by p (which must exist if A(x) is primitive), and let b_s denote the coefficient of lowest order in B(x) that is not divisible by p. Express the product $a_r b_s$ in terms of c_{r+s} and the other a_l 's and b_m 's, and reach a contradiction.
- d) Suppose that a monic integral polynomial P(x) can be factored into a product of two monic rational polynomials, P(x) = A(x)B(x). Show that A(x) and B(x) are integral. **Hint:** First note that we may write $A(x) = (1/r) \cdot \tilde{A}(x)$, and $B(x) = (1/s) \cdot \tilde{B}(x)$, where r, s are positive integers, and $\tilde{A}(x)$ and $\tilde{B}(x)$ are primitive integral; then use (c) to show that r = s = 1.
- e) Combining (b) and (d), prove the theorem.

What have we shown? Since $U_2^{-1}U_1$ is a rotation by an irrational multiple of π , the powers of $U_2^{-1}U_1$ are dense in a U(1) subgroup. Similar reasoning shows that $U_1U_2^{-1}$ is a rotation by the same angle about a different axis, and therefore its powers are dense in another U(1) subgroup. Products of elements of these two noncommuting

U(1) subgroups are dense in the SU(2) subgroup that contains both U_1 and U_2 .

Furthermore, products of

$$\Lambda(\boldsymbol{S})\boldsymbol{U}_{2}^{-1}\boldsymbol{U}_{1}\Lambda(\boldsymbol{S})^{-1} \quad \text{and} \quad \Lambda(\boldsymbol{S})\boldsymbol{U}_{1}\boldsymbol{U}_{2}^{-1}\Lambda(\boldsymbol{S})^{-1}$$
(5.131)

are dense in another SU(2), acting on the span of

$$\left\{\frac{1}{\sqrt{2}}(|01\rangle - |10\rangle), \frac{1}{\sqrt{6}}(|01\rangle + |10\rangle - 2i|11\rangle)\right\} .$$
 (5.132)

Together, these two SU(2) subgroups close on the SU(3) subgroup that acts on the three-dimensional space orthogonal to $|00\rangle$. Conjugating this SU(3) by $\boldsymbol{H} \otimes \boldsymbol{H}$ we obtain another SU(3) acting on the three-dimensional space orthogonal to $|+,+\rangle$, where $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. The only subgroup of SU(4) that contains both of these SU(3) subgroups is SU(4) itself.

Therefore, the circuits constructed from the gate set $\{H, \Lambda(S)\}$ are dense in SU(4) — we can approximate any two-qubit gate to arbitrary accuracy, which we know suffices for universal quantum computation. Whew!

5.5 Universal quantum gates III

We have shown in Exercises 5.3 and 5.4 that the gate set $\{\boldsymbol{H}, \Lambda(\boldsymbol{S})\}$ is universal. Thus any gate set from which both \boldsymbol{H} and $\Lambda(\boldsymbol{S})$ can be constructed is also universal. In particular, we can see that $\{\boldsymbol{H}, \boldsymbol{S}, \Lambda^2(\boldsymbol{X})\}$ and $\{\boldsymbol{H}, \boldsymbol{T}, \Lambda(\boldsymbol{X})\}$ are universal gates sets, where $\boldsymbol{T} = \exp\left(-i\frac{\pi}{8}\boldsymbol{Z}\right)$.

a) It is sometimes convenient to characterize a quantum gate by specifying the action of the gate when it conjugates a Pauli operator. Show that H and S have the properties

$$HXH = Z$$
, $HYH = -Y$, $HZH = X$, (5.133)

and

$$SXS^{-1} = Y$$
, $SYS^{-1} = -X$, $SZS^{-1} = Z$. (5.134)

b) Note that, since $S^{-1} = S^3$, the gate $K = HS^{-1}HSH$ can be constructed using H and S. Show that

$$KXK = Y$$
, $KYK = X$, $KZK = -Z$. (5.135)

- c) Construct circuits for $\Lambda^2(\mathbf{Y})$ and $\Lambda^2(\mathbf{Z})$ using the gate set $\{\mathbf{H}, \mathbf{S}, \Lambda^2(\mathbf{X})\}$. Then complete the proof of universality for this gate set by constructing $\Lambda(\mathbf{S}) \otimes \mathbf{I}$ using $\Lambda^2(\mathbf{X}), \Lambda^2(\mathbf{Y})$, and $\Lambda^2(\mathbf{Z})$.
- d) Show that $\{H, T, \Lambda(X)\}$ is a universal gate set by constructing a circuit for $\Lambda(S)$ from $\Lambda(X)$ and T. Hint: Observe that $T^2 = e^{-i\pi/4}S$, then use the construction suggested in Exercise 5.2, noting that $T^{-1}T^{-1}T^2 = I$ and $T^{-1}XT^{-1}XT^2 = T^2$.

The Toffoli gate $\Lambda^2(\mathbf{X})$ is universal for reversible classical computation. What must be added to realize the full power of quantum computing? We have just seen that the single-qubit gates \mathbf{H} and \mathbf{S} , together with the Toffoli gate, are adequate for reaching any unitary transformation. But in fact, just \mathbf{H} and $\Lambda^2(\mathbf{X})$ suffice to efficiently simulate any quantum computation.

Of course, since H and $\Lambda^2(X)$ are both real orthogonal matrices, a circuit composed from these gates is necessarily real — there are complex *n*-qubit unitaries that cannot be constructed with these tools. But a 2^n -dimensional complex vector space is isomorphic to a 2^{n+1} -dimensional real vector space. A complex vector can be encoded by a real vector according to

$$|\psi\rangle = \sum_{x} \psi_{x} |x\rangle \mapsto |\tilde{\psi}\rangle = \sum_{x} (\operatorname{Re} \psi_{x}) |x,0\rangle + (\operatorname{Im} \psi_{x}) |x,1\rangle , \quad (5.136)$$

and the action of the unitary transformation U can be represented by a real orthogonal matrix U_R defined as

$$\begin{aligned} \boldsymbol{U}_R: \quad & |x,0\rangle \mapsto (\operatorname{Re}\,\boldsymbol{U})|x\rangle \otimes |0\rangle + (\operatorname{Im}\,\boldsymbol{U})|x\rangle \otimes |1\rangle , \\ & |x,1\rangle \mapsto -(\operatorname{Im}\,\boldsymbol{U})|x\rangle \otimes |0\rangle + (\operatorname{Re}\,\boldsymbol{U})|x\rangle \otimes |1\rangle \, (5.137) \end{aligned}$$

To show that the gate set $\{\boldsymbol{H}, \Lambda^2(\boldsymbol{X})\}$ is "universal," it suffices to demonstrate that the real encoding $\Lambda(\boldsymbol{S})_R$ of $\Lambda(\boldsymbol{S})$ can be constructed from $\Lambda^2(\boldsymbol{X})$ and \boldsymbol{H} .

- d) Verify that $\Lambda(\mathbf{S})_R = \Lambda^2(\mathbf{X}\mathbf{Z}).$
- e) Use $\Lambda^2(\mathbf{X})$ and \mathbf{H} to construct a circuit for $\Lambda^2(\mathbf{X}\mathbf{Z})$.

Thus, the classical Toffoli gate does not need much help to unleash the power of quantum computing. In fact, *any* nonclassical singlequbit gate (one that does not preserve the computational basis), combined with the Toffoli gate, is sufficient.

5.6 Universality from any entangling two-qubit gate

We say that a two-qubit unitary quantum gate is *local* if it is a tensor product of single-qubit gates, and that the two-qubit gates U and V are *locally equivalent* if one can be transformed to the other by local gates:

$$\boldsymbol{V} = (\boldsymbol{A} \otimes \boldsymbol{B}) \boldsymbol{U} (\boldsymbol{C} \otimes \boldsymbol{D}) . \tag{5.138}$$

It turns out (you are not asked to prove this) that every two-qubit gate is locally equivalent to a gate of the form:

$$\boldsymbol{V}(\theta_x, \theta_y, \theta_z) = \exp\left[i\left(\theta_x \boldsymbol{X} \otimes \boldsymbol{X} + \theta_y \boldsymbol{Y} \otimes \boldsymbol{Y} + \theta_z \boldsymbol{Z} \otimes \boldsymbol{Z}\right)\right],$$
(5.139)

where

$$-\pi/4 < \theta_x \le \theta_y \le \theta_z \le \pi/4 . \tag{5.140}$$

a) Show that $V(\pi/4, \pi/4, \pi/4)$ is (up to an overall phase) the **SWAP** operation that interchanges the two qubits:

SWAP
$$(|\psi\rangle \otimes |\phi\rangle) = |\phi\rangle \otimes |\psi\rangle$$
. (5.141)

b) Show that $V(0, 0, \pi/4)$ is locally equivalent to the CNOT gate $\Lambda(\mathbf{X})$.

As shown in Exercise 5.2, the CNOT gate $\Lambda(\mathbf{X})$ together with arbitrary single-qubit gates form an exactly universal gate set. But in fact there is nothing special about the CNOT gate in this regard. Any two-qubit gate \mathbf{U} , when combined with arbitrary single-qubit gates, suffices for universality unless \mathbf{U} is either local or locally equivalent to **SWAP**.

To demonstrate that U is universal when assisted by local gates it suffices to construct $\Lambda(\mathbf{X})$ using a circuit containing only local gates and U gates.

Lemma If U is locally equivalent to $V(\theta_x, \theta_y, \theta_z)$, then $\Lambda(X)$ can be constructed from a circuit using local gates and U gates except in two cases: (1) $\theta_x = \theta_y = \theta_z = 0$ (U is local), (2) $\theta_x = \theta_y = \theta_z = \pi/4$ (U is locally equivalent to **SWAP**)..

You will prove the Lemma in the rest of this exercise.

c) Show that:

$$(\mathbf{I} \otimes \mathbf{X}) \mathbf{V}(\theta_x, \theta_y, \theta_z) (\mathbf{I} \otimes \mathbf{X}) \mathbf{V}(\theta_x, \theta_y, \theta_z) = \mathbf{V}(2\theta_x, 0, 0) , (\mathbf{I} \otimes \mathbf{Y}) \mathbf{V}(\theta_x, \theta_y, \theta_z) (\mathbf{I} \otimes \mathbf{Y}) \mathbf{V}(\theta_x, \theta_y, \theta_z) = \mathbf{V}(0, 2\theta_y, 0) , (\mathbf{I} \otimes \mathbf{Z}) \mathbf{V}(\theta_x, \theta_y, \theta_z) (\mathbf{I} \otimes \mathbf{Z}) \mathbf{V}(\theta_x, \theta_y, \theta_z) = \mathbf{V}(0, 0, 2\theta_z) .$$

$$(5.142)$$

- d) Show that $V(0,0,\theta)$ is locally equivalent to the controlled rotation $\Lambda[\mathbf{R}(\hat{n},4\theta)]$, where $\mathbf{R}(\hat{n},4\theta) = \exp[-2i\theta(\hat{n}\cdot\vec{\sigma})]$, for an arbitrary axis of rotation \hat{n} . (Here $\vec{\sigma} = (\mathbf{X}, \mathbf{Y}, \mathbf{Z})$.)
- e) Now use the results of (c) and (d) to prove the Lemma.

Chapter 6

Quantum Computation

6.1 Classical Circuits

The concept of a quantum computer was introduced in Chapter 1. Here we will specify our model of quantum computation more precisely, and we will point out some basic properties of the model. But before we explain what a quantum computer does, perhaps we should say what a classical computer does.

6.1.1 Universal gates

A classical (deterministic) computer evaluates a function: given n-bits of input it produces m-bits of output that are uniquely determined by the input; that is, it finds the value of

$$f: \{0,1\}^n \to \{0,1\}^m, \tag{6.1}$$

for a particular specified n-bit argument. A function with an m-bit value is equivalent to m functions, each with a one-bit value, so we may just as well say that the basic task performed by a computer is the evaluation of

$$f: \{0,1\}^n \to \{0,1\}. \tag{6.2}$$

We can easily count the number of such functions. There are 2^n possible inputs, and for each input there are two possible outputs. So there are altogether 2^{2^n} functions taking n bits to one bit.

The evaluation of any such function can be reduced to a sequence of elementary logical operations. Let us divide the possible values of the input

$$x = x_1 x_2 x_3 \dots x_n, \tag{6.3}$$

into one set of values for which f(x) = 1, and a complementary set for which f(x) = 0. For each $x^{(a)}$ such that $f(x^{(a)}) = 1$, consider the function $f^{(a)}$ such that

$$f^{(a)}(x) = \begin{cases} 1 & x = x^{(a)} \\ 0 & \text{otherwise} \end{cases}$$
(6.4)

Then

$$f(x) = f^{(1)}(x) \vee f^{(2)}(x) \vee f^{(3)}(x) \vee \dots$$
(6.5)

f is the logical OR (\lor) of all the $f^{(a)}$'s. In binary arithmetic the \lor operation of two bits may be represented

$$x \lor y = x + y - x \cdot y; \tag{6.6}$$

it has the value 0 if x and y are both zero, and the value 1 otherwise.

Now consider the evaluation of $f^{(a)}$. In the case where $x^{(a)} = 111...1$, we may write

$$f^{(a)}(x) = x_1 \wedge x_2 \wedge x_3 \dots \wedge x_n; \tag{6.7}$$

it is the logical AND (\wedge) of all *n* bits. In binary arithmetic, the AND is the product

$$x \wedge y = x \cdot y. \tag{6.8}$$

For any other $x^{(a)}$, $f^{(a)}$ is again obtained as the AND of *n* bits, but where the NOT (\neg) operation is first applied to each x_i such that $x_i^{(a)} = 0$; for example

$$f^{(a)}(x) = (\neg x_1) \land x_2 \land x_3 \land (\neg x_4) \land \dots$$
(6.9)

if

$$x^{(a)} = 0110\dots$$
 (6.10)

The NOT operation is represented in binary arithmetic as

$$\neg x = 1 - x. \tag{6.11}$$

We have now constructed the function f(x) from three elementary logical connectives: NOT, AND, OR. The expression we obtained is called the "disjunctive normal form" of f(x). We have also implicitly used another operation, COPY, that takes one bit to two bits:

$$COPY: x \to xx. \tag{6.12}$$

We need the COPY operation because each $f^{(a)}$ in the disjunctive normal form expansion of f requires its own copy of x to act on.

In fact, we can pare our set of elementary logical connectives to a smaller set. Let us define a NAND ("NOT AND") operation by

$$x \uparrow y = \neg (x \land y) = (\neg x) \lor (\neg y). \tag{6.13}$$

In binary arithmetic, the NAND operation is

$$x \uparrow y = 1 - xy. \tag{6.14}$$

If we can COPY, we can use NAND to perform NOT:

$$x \uparrow x = 1 - x^2 = 1 - x = \neg x. \tag{6.15}$$

(Alternatively, if we can prepare the constant y = 1, then $x \uparrow 1 = 1 - x = \neg x$.) Also,

$$(x \uparrow y) \uparrow (x \uparrow y) = \neg (x \uparrow y) = 1 - (1 - xy) = xy = x \land y,$$
(6.16)

and

$$(x \uparrow x) \uparrow (y \uparrow y) = (\neg x) \uparrow (\neg y) = 1 - (1 - x)(1 - y)$$
$$= x + y - xy = x \lor y.$$
(6.17)

So if we can COPY, NAND performs AND and OR as well. We conclude that the single logical connective NAND, together with COPY, suffices to evaluate any function f. (You can check that an alternative possible choice of the universal connective is NOR:

$$x \downarrow y = \neg (x \lor y) = (\neg x) \land (\neg y).) \tag{6.18}$$

If we are able to prepare a constant bit (x = 0 or x = 1), we can reduce the number of elementary operations from two to one. The NAND/NOT gate

$$(x, y) \to (1 - x, 1 - xy),$$
 (6.19)

computes NAND (if we ignore the first output bit) and performs copy (if we set the second input bit to y = 1, and we subsequently apply NOT to both output bits). We say, therefore, that NAND/NOT is a universal gate. If we have a supply of constant bits, and we can apply the NAND/NOT gates to any chosen pair of input bits, then we can perform a sequence of NAND/NOT gates to evaluate any function $f : \{0,1\}^n \to \{0,1\}$ for any value of the input $x = x_1 x_2 \dots x_n$.

These considerations motivate the circuit model of computation. A computer has a few basic components that can perform elementary operations on bits or pairs of bits, such as COPY, NOT, AND, OR. It can also prepare a constant bit or input a variable bit. A computation is a finite sequence of such operations, a *circuit*, applied to a specified string of input bits.¹ The result of the computation is the final value of all remaining bits, after all the elementary operations have been executed.

It is a fundamental result in the theory of computation that just a few elementary gates suffice to evaluate any function of a finite input. This result means that with very simple hardware components, we can build up arbitrarily complex computations.

So far, we have only considered a computation that acts on a particular fixed input, but we may also consider *families* of circuits that act on inputs of variable size. Circuit families provide a useful scheme for analyzing and classifying the *complexity* of computations, a scheme that will have a natural generalization when we turn to quantum computation.

6.1.2 Circuit complexity

In the study of complexity, we will often be interested in functions with a one-bit output

$$f: \{0, 1\}^n \to \{0, 1\}.$$
 (6.20)

¹The circuit is required to be *acyclic*, meaning that no *directed* closed loops are permitted.

Such a function f may be said to encode a solution to a "decision problem" — the function examines the input and issues a YES or NO answer. Often, a question that would not be stated colloquially as a question with a YES/NO answer can be "repackaged" as a decision problem. For example, the function that defines the FACTORING problem is:

$$f(x,y) = \begin{cases} 1 & \text{if integer } x \text{ has a divisor less than } y, \\ 0 & \text{otherwise;} \end{cases}$$
(6.21)

knowing f(x, y) for all y < x is equivalent to knowing the *least* nontrivial factor of y. Another important example of a decision problem is the HAMIL-TONIAN path problem: let the input be an ℓ -vertex graph, represented by an $\ell \times \ell$ adjacency matrix (a 1 in the ij entry means there is an edge linking vertices i and j); the function is

$$f(x) = \begin{cases} 1 & \text{if graph } x \text{ has a Hamiltonian path,} \\ 0 & \text{otherwise.} \end{cases}$$
(6.22)

(A path is Hamiltonian if it visits each vertex exactly once.)

We wish to gauge how hard a problem is by quantifying the resources needed to solve the problem. For a decision problem, a reasonable measure of hardness is the *size* of the smallest circuit that computes the corresponding function $f : \{0,1\}^n \to \{0,1\}$. By size we mean the number of elementary gates or components that we must wire together to evaluate f. We may also be interested in how much *time* it takes to do the computation if many gates are permitted to execute in parallel. The *depth* of a circuit is the number of time steps required, assuming that gates acting on distinct bits can operate simultaneously (that is, the depth is the maximum length of a directed path from the input to the output of the circuit). The *width* of a circuit is the maximum number of gates that act in any one time step.

We would like to divide the decision problems into two classes: easy and hard. But where should we draw the line? For this purpose, we consider infinite families of decision problems with variable input size; that is, where the number of bits of input can be any integer n. Then we can examine how the size of the circuit that solves the problem scales with n.

If we use the scaling behavior of a circuit family to characterize the difficulty of a problem, there is a subtlety. It would be cheating to hide the difficulty of the problem in the *design* of the circuit. Therefore, we should restrict attention to circuit families that have acceptable "uniformity" properties — it must be "easy" to build the circuit with n + 1 bits of input once we have constructed the circuit with an *n*-bit input.

Associated with a family of functions $\{f_n\}$ (where f_n has *n*-bit input) are circuits $\{C_n\}$ that compute the functions. We say that a circuit family $\{C_n\}$ is "polynomial size" if the size of C_n grows with *n* no faster than a power of n,

size
$$(C_n) \le \text{poly } (n),$$
 (6.23)

where poly denotes a polynomial. Then we define:

 $P = \{ \text{decision problem solved by polynomial-size circuit families} \}$

(*P* for "polynomial time"). Decision problems in *P* are "easy." The rest are "hard." Notice that C_n computes $f_n(x)$ for every possible *n*-bit input, and therefore, if a decision problem is in *P* we can find the answer even for the "worst-case" input using a circuit of size no greater than poly(*n*). (As noted above, we implicitly assume that the circuit family is "uniform" so that the *design* of the circuit can itself be solved by a polynomial-time algorithm. Under this assumption, solvability in polynomial time by a circuit family is equivalent to solvability in polynomial time by a universal Turing machine.)

Of course, to determine the size of a circuit that computes f_n , we must know what the elementary components of the circuit are. Fortunately, though, whether a problem lies in P does not depend on what gate set we choose, as long as the gates are universal, the gate set is finite, and each gate acts on a set of bits of bounded size. One universal gate set can *simulate* another.

The vast majority of function families $f : \{0,1\}^n \to \{0,1\}$ are not in P. For most functions, the output is essentially random, and there is no better way to "compute" f(x) than to consult a look-up table of its values. Since there are 2^n *n*-bit inputs, the look-up table has *exponential* size, and a circuit that encodes the table must also have exponential size. The problems in P belong to a very special class — they have enough structure so that the function f can be computed efficiently.

Of particular interest are decision problems that can be answered by exhibiting an example that is easy to verify. For example, given x and y < x, it is hard (in the worst case) to determine if x has a factor less than y. But if someone kindly provides a z < y that divides x, it is easy for us to check that z is indeed a factor of x. Similarly, it is hard to determine if a graph has a Hamiltonian path, but if someone kindly provides a path, it is easy to verify that the path really is Hamiltonian.

This concept that a problem may be hard to solve, but that a solution can be easily verified once found, can be formalized by the notion of a "nondeterministic" circuit. A nondeterministic circuit $\tilde{C}_{n,m}(x^{(n)}, y^{(m)})$ associated with the circuit $C_n(x^{(n)})$ has the property:

$$C_n(x^{(n)}) = 1$$
 iff $\tilde{C}_{n,m}(x^{(n)}, y^{(m)}) = 1$ for some $y^{(m)}$. (6.24)

(where $x^{(n)}$ is *n* bits and $y^{(m)}$ is *m* bits.) Thus for a particular $x^{(n)}$ we can use $\tilde{C}_{n,m}$ to verify that $C_n(x^{(n)} = 1, if$ we are fortunate enough to have the right $y^{(m)}$ in hand. We define:

NP: {decision problems that admit a polynomial-size *nondeterministic* circuit family}

(*NP* for "nondeterministic polynomial time"). If a problem is in *NP*, there is no guarantee that the problem is easy, only that a solution is easy to check once we have the right information. Evidently $P \subseteq NP$. Like *P*, the *NP* problems are a small subclass of all decision problems.

Much of complexity theory is built on a fundamental conjecture:

$$Conjecture: P \neq NP; \tag{6.25}$$

there exist hard decision problems whose solutions are easily verified. Unfortunately, this important conjecture still awaits proof. But after 30 years of trying to show otherwise, most complexity experts are firmly confident of its validity.

An important example of a problem in NP is CIRCUIT-SAT. In this case the input is a circuit C with n gates, m input bits, and one output bit. The problem is to find if there is any m-bit input for which the output is 1. The function to be evaluated is

$$f(C) = \begin{cases} 1 & \text{if there exists } x^{(m)} \text{ with } C(x^{(m)}) = 1, \\ 0 & \text{otherwise.} \end{cases}$$
(6.26)

This problem is in NP because, given a circuit, it is easy to simulate the circuit and evaluate its output for any particular input.

I'm going to state some important results in complexity theory that will be relevant for us. There won't be time for proofs. You can find out more by consulting one of the many textbooks on the subject; one good one is *Computers and Intractability: A Guide to the Theory of NP-Completeness*, by M. R. Garey and D. S. Johnson.

Many of the insights engendered by complexity theory flow from Cook's Theorem (1971). The theorem states that *every* problem in NP is *polynomially reducible* to CIRCUIT-SAT. This means that for any PROBLEM $\in NP$, there is a polynomial-size circuit family that maps an "instance" $x^{(n)}$ of PROBLEM to an "instance" $y^{(m)}$ of CIRCUIT-SAT; that is

CIRCUIT – SAT
$$(y^{(m)}) = 1$$
 iff PROBLEM $(x^{(n)}) = 1.$
(6.27)

It follows that if we had a magical device that could efficiently solve CIRCUIT-SAT (a CIRCUIT-SAT "oracle"), we could couple that device with the polynomial reduction to efficiently solve PROBLEM. Cook's theorem tells us that if it turns out that CIRCUIT-SAT $\in P$, then P = NP.

A problem that, like CIRCUIT-SAT, has the property that every problem in NP is polynomially reducible to it, is called NP-complete (NPC). Since Cook, many other examples have been found. To show that a PROB-LEM $\in NP$ is NP-complete, it suffices to find a polynomial reduction to PROBLEM of another problem that is already known to be NP-complete. For example, one can exhibit a polynomial reduction of CIRCUIT-SAT to HAMILTONIAN. It follows from Cook's theorem that HAMILTONIAN is also NP-complete.

If we assume that $P \neq NP$, it follows that there exist problems in NP of intermediate difficulty (the class NPI). These are neither P nor NPC.

Another important complexity class is called co-NP. Heuristically, NP decision problems are ones we can answer by exhibiting an *example* if the answer is YES, while co-NP problems can be answered with a *counter-example* if the answer is NO. More formally:

$$\{C\} \in NP : C(x) = 1 \text{ iff } C(x, y) = 1 \text{ for some } y$$
 (6.28)

$$\{C\} \in co - NP : C(x) = 1 \text{ iff } C(x, y) = 1 \text{ for all } y.$$
(6.29)

Clearly, there is a symmetry relating the classes NP and co-NP — whether we consider a problem to be in NP or co-NP depends on how we choose to frame the question. ("Is there a Hamiltonian circuit?" is in NP. "Is there no Hamiltonian circuit?" is in co-NP). But the interesting question is: is a problem in *both* NP and co-NP? If so, then we can easily verify the answer (once a suitable example is in hand) regardless of whether the answer is YES or NO. It is believed (though not proved) that $NP \neq co-NP$. (For example, we can show that a graph has a Hamiltonian path by exhibiting an example, but we don't know how to show that it has *no* Hamiltonian path that way!) Assuming that $NP \neq co-NP$, there is a theorem that says that no co-NP problems are contained in NPC. Therefore, problems in the intersection of NP and co-NP, if not in P, are good candidates for inclusion in NPI.

In fact, a problem in $NP \cap co-NP$ that is believed not in P is the FACTORING problem. As already noted, FACTORING is in NP because, if we are offered a factor of x, we can easily check its validity. But it is also in co-NP, because it is known that if we are given a prime number then (at least in principle), we can efficiently verify its primality. Thus, if someone tells us the prime factors of x, we can efficiently check that the prime factorization is right, and can *exclude* that any integer less than y is a divisor of x. Therefore, it seems likely that FACTORING is in NPI.

We are led to a crude (conjectured) picture of the structure of $NP \cup co-NP$. NP and co-NP do not coincide, but they have a nontrivial intersection. P lies in $NP \cap co-NP$ (because P = co-P), but the intersection also contains problems not in P (like FACTORING). Neither NPC nor co-NPC intersects with $NP \cap co-NP$.

There is much more to say about complexity theory, but we will be content to mention one more element that relates to the discussion of quantum complexity. It is sometimes useful to consider *probabilistic* circuits that have access to a random number generator. For example, a gate in a probabilistic circuit might act in either one of two ways, and flip a fair coin to decide which action to execute. Such a circuit, for a single fixed input, can sample many possible computational paths. An algorithm performed by a probabilistic circuit is said to be "randomized."

If we attack a decision problem using a probabilistic computer, we attain a probability distribution of outputs. Thus, we won't necessarily always get the right answer. But if the probability of getting the right answer is larger than $\frac{1}{2} + \delta$ for every possible input ($\delta > 0$), then the machine is useful. In fact, we can run the computation many times and use majority voting to achieve an error probability less than ε . Furthermore, the number of times we need to repeat the computation is only polylogarithmic in ε^{-1} .

If a problem admits a probabilistic circuit family of polynomial size that always gives the right answer with probability larger than $\frac{1}{2} + \delta$ (for any input, and for fixed $\delta > 0$), we say the problem is in the class *BPP* ("bounded-error

probabilistic polynomial time"). It is evident that

$$P \subseteq BPP, \tag{6.30}$$

but the relation of NP to BPP is not known. In particular, it has not been proved that BPP is contained in NP.

6.1.3 Reversible computation

In devising a model of a quantum computer, we will generalize the circuit model of classical computation. But our quantum logic gates will be unitary transformations, and hence will be invertible, while classical logic gates like the NAND gate are not invertible. Before we discuss quantum circuits, it is useful to consider some features of reversible classical computation.

Aside from the connection with quantum computation, another incentive for studying reversible classical computation arose in Chapter 1. As Landauer observed, because irreversible logic elements erase information, they are necessarily dissipative, and therefore, require an irreducible expenditure of power. But if a computer operates reversibly, then in principle there need be no dissipation and no power requirement. We can compute for free!

A reversible computer evaluates an invertible function taking n bits to n bits

$$f: \{0,1\}^n \to \{0,1\}^n, \tag{6.31}$$

the function must be invertible so that there is a unique input for each output; then we are able in principle to run the computation backwards and recover the input from the output. Since it is a 1-1 function, we can regard it as a permutation of the 2^n strings of n bits — there are $(2^n)!$ such functions.

Of course, any irreversible computation can be "packaged" as an evaluation of an invertible function. For example, for any $f : \{0,1\}^n \to \{0,1\}^m$, we can construct $\tilde{f} : \{0,1\}^{n+m} \to \{0,1\}^{n+m}$ such that

$$\tilde{f}(x;0^{(m)}) = (x;f(x)),$$
(6.32)

(where $0^{(m)}$ denotes *m*-bits initially set to zero). Since \tilde{f} takes each $(x; 0^{(m)})$ to a distinct output, it can be extended to an invertible function of n + m bits. So for any f taking n bits to m, there is an invertible \tilde{f} taking n + m to n + m, which evaluates f(x) acting on $(x, 0^{(m)})$

Now, how do we build up a complicated reversible computation from elementary components — that is, what constitutes a universal gate set? We will see that one-bit and two-bit reversible gates do not suffice; we will need three-bit gates for universal reversible computation.

Of the four 1-bit \rightarrow 1-bit gates, two are reversible; the trivial gate and the NOT gate. Of the $(2^4)^2 = 256$ possible 2-bit \rightarrow 2-bit gates, 4! = 24 are reversible. One of special interest is the controlled-NOT or reversible XOR gate that we already encountered in Chapter 4:

XOR:
$$(x, y) \mapsto (x, x \oplus y),$$
 (6.33)

$$\begin{array}{c} x & & \\ y & & \\ & & \\ \end{array} \begin{array}{c} x \\ x \oplus y \end{array}$$

This gate flips the second bit if the first is 1, and does nothing if the first bit is 0 (hence the name controlled-NOT). Its square is trivial, that is, it inverts itself. Of course, this gate performs a NOT on the second bit if the first bit is set to 1, and it performs the copy operation if y is initially set to zero:

$$XOR: (x,0) \mapsto (x,x). \tag{6.34}$$

With the circuit



constructed from three X0R's, we can swap two bits:

$$(x,y) \to (x,x \oplus y) \to (y,x \oplus y) \to (y,x). \tag{6.35}$$

With these swaps we can shuffle bits around in a circuit, bringing them together if we want to act on them with a particular component in a fixed location.

To see that the one-bit and two-bit gates are nonuniversal, we observe that all these gates are *linear*. Each reversible two-bit gate has an action of the form

$$\begin{pmatrix} x \\ y \end{pmatrix} \to \begin{pmatrix} x' \\ y' \end{pmatrix} = \mathcal{M} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} a \\ b \end{pmatrix}, \tag{6.36}$$

where the constant $\begin{pmatrix} a \\ b \end{pmatrix}$ takes one of four possible values, and the matrix \mathcal{M} is one of the six invertible matrices

$$\mathcal{M} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}.$$
(6.37)

(All addition is performed modulo 2.) Combining the six choices for \mathcal{M} with the four possible constants, we obtain 24 distinct gates, which exhausts all the reversible $2 \rightarrow 2$ gates.

Since the linear transformations are closed under composition, any circuit composed from reversible $2 \rightarrow 2$ (and $1 \rightarrow 1$) gates will compute a linear function

$$x \to \mathcal{M}x + a.$$
 (6.38)

But for $n \geq 3$, there are invertible functions on *n*-bits that are nonlinear. An important example is the 3-bit *Toffoli gate* (or controlled-NOT) $\theta^{(3)}$

$$\theta^{(3)}: (x, y, z) \to (x, y, z \oplus xy);$$

$$x \xrightarrow{\qquad } x \\ y \xrightarrow{\qquad } y \\ z \xrightarrow{\qquad } z \oplus xy$$

$$(6.39)$$

it flips the third bit if the first two are 1 and does nothing otherwise. Like the XOR gate, it is its own inverse.

Unlike the reversible 2-bit gates, the Toffoli gate serves as a universal gate for Boolean logic, if we can provide fixed input bits and ignore output bits. If z is initially 1, then $x \uparrow y = 1 - xy$ appears in the third output — we can perform NAND. If we fix x = 1, the Toffoli gate functions like an XOR gate, and we can use it to copy.

The Toffoli gate $\theta^{(3)}$ is universal in the sense that we can build a circuit to compute any reversible function using Toffoli gates alone (if we can fix input bits and ignore output bits). It will be instructive to show this directly, without relying on our earlier argument that NAND/NOT is universal for Boolean functions. In fact, we can show the following: From the NOT gate and the Toffoli gate $\theta^{(3)}$, we can construct any invertible function on n bits, provided we have one extra bit of scratchpad space available.

The first step is to show that from the three-bit Toffoli-gate $\theta^{(3)}$ we can construct an *n*-bit Toffoli gate $\theta^{(n)}$ that acts as

$$(x_1, x_2, \dots, x_{n-1}, y) \to (x_1, x_2, \dots, x_{n-1}y \oplus x_1x_2 \dots x_{n-1}).$$

(6.40)

The construction requires one extra bit of scratch space. For example, we construct $\theta^{(4)}$ from $\theta^{(3)}$'s with the circuit



The purpose of the last $\theta^{(3)}$ gate is to reset the scratch bit back to its original value zero. Actually, with one more gate we can obtain an implementation of $\theta^{(4)}$ that works irrespective of the initial value of the scratch bit:



Again, we can eliminate the last gate if we don't mind flipping the value of the scratch bit.

We can see that the scratch bit really is necessary, because $\theta^{(4)}$ is an odd permutation (in fact a transposition) of the 24 4-bit strings — it transposes 1111 and 1110. But $\theta^{(3)}$ acting on any three of the four bits is an even permutation; *e.g.*, acting on the last three bits it transposes 0111 with 0110, and 1111 with 1110. Since a product of even permutations is also even, we cannot obtain $\theta^{(4)}$ as a product of $\theta^{(3)}$'s that act on four bits only.

The construction of $\theta^{(4)}$ from four $\theta^{(3)}$'s generalizes immediately to the construction of $\theta^{(n)}$ from two $\theta^{(n-1)}$'s and two $\theta^{(3)}$'s (just expand x_1 to several control bits in the above diagram). Iterating the construction, we obtain $\theta^{(n)}$ from a circuit with $2^{n-2} + 2^{n-3} - 2 \theta^{(3)}$'s. Furthermore, just one bit of scratch space is sufficient.²) (When we need to construct $\theta^{(k)}$, any available extra bit will do, since the circuit returns the scratch bit to its original value. The next step is to note that, by conjugating $\theta^{(n)}$ with NOT gates, we can in effect modify the value of the control string that "triggers" the gate. For example, the circuit



flips the value of y if $x_1x_2x_3 = 010$, and it acts trivially otherwise. Thus this circuit transposes the two strings 0100 and 0101. In like fashion, with $\theta^{(n)}$ and NOT gates, we can devise a circuit that transposes any two *n*-bit strings that differ in only one bit. (The location of the bit where they differ is chosen to be the *target* of the $\theta^{(n)}$ gate.)

But in fact a transposition that exchanges any two *n*-bit strings can be expressed as a product of transpositions that interchange strings that differ in only one bit. If a_0 and a_s are two strings that are Hamming distance s apart (differ in s places), then there is a chain

$$a_0, a_1, a_2, a_3, \dots, a_s,$$
 (6.41)

such that each string in the chain is Hamming distance one from its neighbors. Therefore, each of the transpositions

$$(a_0a_1), (a_1a_2), (a_2a_3), \dots (a_{s-1}a_s),$$
 (6.42)

²With more scratch space, we can build $\theta^{(n)}$ from $\theta^{(3)}$'s much more efficiently — see the exercises.

can be implemented as a $\theta^{(n)}$ gate conjugated by NOT gates. By composing transpositions we find

$$(a_0a_s) = (a_{s-1}a_s)(a_{s-2}a_{s-1})\dots(a_2a_3)(a_1a_2)(a_0a_1)(a_1a_2)(a_2a_3) \dots(a_{s-2}a_{s-1})(a_{s-1}a_s);$$
(6.43)

we can construct the Hamming-distance-s transposition from 2s-1 Hammingdistance-one transpositions. It follows that we can construct (a_0a_s) from $\theta^{(n)}$'s and NOT gates.

Finally, since every permutation is a product of transpositions, we have shown that every invertible function on *n*-bits (every permutation on *n*-bit strings) is a product of $\theta^{(3)}$'s and NOT's, using just one bit of scratch space.

Of course, a NOT can be performed with a $\theta^{(3)}$ gate if we fix two input bits at 1. Thus the Toffoli gate $\theta^{(3)}$ is universal for reversible computation, if we can fix input bits and discard output bits.

6.1.4 Billiard ball computer

Two-bit gates suffice for universal irreversible computation, but three-bit gates are needed for universal reversible computation. One is tempted to remark that "three-body interactions" are needed, so that building reversible hardware is more challenging than building irreversible hardware. However, this statement may be somewhat misleading.

Fredkin described how to devise a universal reversible computer in which the fundamental interaction is an elastic collision between two billiard balls. Balls of radius $\frac{1}{\sqrt{2}}$ move on a square lattice with unit lattice spacing. At each integer valued time, the center of each ball lies at a lattice site; the presence or absence of a ball at a particular site (at integer time) encodes a bit of information. In each unit of time, each ball moves unit distance along one of the lattice directions. Occasionally, at integer-valued times, 90° elastic collisions occur between two balls that occupy sites that are distance $\sqrt{2}$ apart (joined by a lattice diagonal).

The device is programmed by nailing down balls at certain sites, so that those balls act as perfect reflectors. The program is executed by fixing initial positions and directions for the moving balls, and evolving the system according to Newtonian mechanics for a finite time. We read the output by observing the final positions of all the moving balls. The collisions are nondissipative, so that we can run the computation backward by reversing the velocities of all the balls.
To show that this machine is a universal reversible computer, we must explain how to operate a universal gate. It is convenient to consider the three-bit *Fredkin gate*

$$(x, y, z) \to (x, xz + \bar{x}y, xy + \bar{x}z), \tag{6.44}$$

which swaps y and z if x = 0 (we have introduced the notation $\bar{x} = \neg x$). You can check that the Fredkin gate can simulate a NAND/NOT gate if we fix inputs and ignore outputs.

We can build the Fredkin gate from a more primitive object, the *switch* gate. A switch gate taking two bits to three acts as

$$(x,y) \to (x,xy,\bar{x}y). \tag{6.45}$$
$$\begin{array}{c} x\\ y \end{array} = \underbrace{S} = \begin{array}{c} x\\ xy\\ \bar{x}y \end{array}$$

The gate is "reversible" in that we can run it backwards acting on a constrained 3-bit input taking one of the four values

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} (6.46)$$

Furthermore, the switch gate is itself universal; fixing inputs and ignoring outputs, it can do NOT (y = 1, third output) AND (second output), and COPY (y = 1, first and second output). It is not surprising, then, that we can compose switch gates to construct a universal reversible $3 \rightarrow 3$ gate. Indeed, the circuit

builds the Fredkin gate from four switch gates (two running forward and two running backward). Time delays needed to maintain synchronization are not explicitly shown.

In the billiard ball computer, the switch gate is constructed with two reflectors, such that (in the case x = y = 1) two moving balls collide twice. The trajectories of the balls in this case are:

A ball labeled x emerges from the gate along the same trajectory (and at the same time) regardless of whether the other ball is present. But for x = 1, the position of the other ball (if present) is shifted down compared to its final position for x = 0 — this is a switch gate. Since we can perform a switch gate, we can construct a Fredkin gate, and implement universal reversible logic with a billiard ball computer.

An evident weakness of the billiard-ball scheme is that initial errors in the positions and velocities of the ball will accumulate rapidly, and the computer will eventually fail. As we noted in Chapter 1 (and Landauer has insistently pointed out) a similar problem will afflict any proposed scheme for dissipationless computation. To control errors we must be able to compress the phase space of the device, which will necessarily be a dissipative process.

6.1.5 Saving space

But even aside from the issue of error control there is another key question about reversible computation. How do we manage the scratchpad space needed to compute reversibly?

In our discussion of the universality of the Toffoli gate, we saw that in principle we can do any reversible computation with very little scratch space. But in practice it may be impossibly difficult to figure out how to do a particular computation with minimal space, and in any case economizing on space may be costly in terms of the run time.

There is a general strategy for simulating an irreversible computation on a reversible computer. Each irreversible NAND or COPY gate can be simulated by a Toffoli gate by fixing inputs and ignoring outputs. We accumulate and save all "garbage" output bits that are needed to reverse the steps of the computation. The computation proceeds to completion, and then a copy of the output is generated. (This COPY operation is logically reversible.) Then the computation runs in reverse, cleaning up all garbage bits, and returning all registers to their original configurations. With this procedure the reversible circuit runs only about twice as long as the irreversible circuit that it simulates, and all garbage generated in the simulation is disposed of without any dissipation and hence no power requirement.

This procedure works, but demands a huge amount of scratch space – the space needed scales linearly with the length T of the irreversible computation being simulated. In fact, it is possible to use space far more efficiently (with only a minor slowdown), so that the space required scales like log T instead

of T. (That is, there is a general-purpose scheme that requires space $\propto \log T$; of course, we might do even better in the simulation of a particular computation.)

To use space more effectively, we will divide the computation into smaller steps of roughly equal size, and we will run these steps backward when possible during the course of the computation. However, just as we are unable to perform step k of the computation unless step k - 1 has already been completed, we are unable to run step k in reverse if step k - 1 has previously been executed in reverse.³ The amount of space we require (to store our garbage) will scale like the maximum value of the number of forward steps minus the number of backward steps that have been executed.

The challenge we face can be likened to a game — the reversible pebble game.⁴ The steps to be executed form a one-dimension directed graph with sites labeled 1, 2, 3...T. Execution of step k is modeled by placing a pebble on the kth site of the graph, and executing step k in reverse is modeled as removal of a pebble from site k. At the beginning of the game, no sites are covered by pebbles, and in each turn we add or remove a pebble. But we cannot place a pebble at site k (except for k = 1) unless site k - 1 is already covered by a pebble, and we cannot remove a pebble from site k (except for k = 1) unless site k (except for k

In fact, with n pebbles we can reach site $T = 2^n - 1$, but we can go no further.

We can construct a recursive procedure that enables us to reach site $T = 2^{n-1}$ with *n* pebbles, leaving only one pebble in play. Let $F_1(k)$ denote placing a pebble at site *k*, and $F_1(k)^{-1}$ denote removing a pebble from site *k*. Then

$$F_2(1,2) = F_1(1)F_1(2)F_1(1)^{-1}, (6.47)$$

leaves a pebble at site k = 2, using a maximum of two pebbles at intermediate

³We make the conservative assumption that we are not clever enough to know ahead of time what portion of the output from step k-1 might be needed later on. So we store a complete record of the configuration of the machine after step k-1, which is not to be erased until an updated record has been stored after the completion of a subsequent step.

⁴as pointed out by Bennett. For a recent discussion, see M. Li and P. Vitanyi, quant-ph/9703022.

stages. Similarly

$$F_3(1,4) = F_2(1,2)F_2(3,4)F_2(1,2)^{-1}, (6.48)$$

reaches site k = 4 using a maximum of three pebbles, and

$$F_4(1,8) = F_3(1,4)F_3(5,8)F_3(1,4)^{-1}, (6.49)$$

reaches k = 8 using four pebbles. Evidently we can construct $F_n(1, 2^{n-1})$ which uses a maximum of n pebbles and leaves a single pebble in play. (The routine

$$F_n(1, 2^{n-1})F_{n-1}(2^{n-1}+1, 2^{n-1}+2^{n-2})\dots F_1(2^n-1),$$
(6.50)

leaves all n pebbles in play, with the maximal pebble at site $k = 2^n - 1$.)

Interpreted as a routine for executing $T = 2^{n-1}$ steps of a computation, this strategy for playing the pebble game represents a simulation requiring space scaling like $n \sim \log T$. How long does the simulation take? At each level of the recursive procedure described above, two steps forward are replaced by two steps forward and one step back. Therefore, an irreversible computation with $T_{irr} = 2^n$ steps is simulated in $T_{rev} = 3^n$ steps, or

$$T_{\rm rev} = (T_{\rm irr})^{\log 3/\log 2}, = (T_{\rm irr})^{1.58},$$
 (6.51)

a modest power law slowdown.

In fact, we can improve the slowdown to

$$T_{\rm rev} \sim (T_{\rm irr})^{1+\varepsilon},$$
 (6.52)

for any $\varepsilon > 0$. Instead of replacing two steps forward with two forward and one back, we replace ℓ forward with ℓ forward and $\ell - 1$ back. A recursive procedure with *n* levels reaches site ℓ^n using a maximum of $n(\ell - 1) + 1$ pebbles. Now we have $T_{irr} = \ell^n$ and $T_{rev} = (2\ell - 1)^n$, so that

$$T_{\rm rev} = (T_{\rm irr})^{\log(2\ell-1)/\log\ell};$$
 (6.53)

the power characterizing the slowdown is

$$\frac{\log(2\ell-1)}{\log\ell} = \frac{\log 2\ell + \log\left(1 - \frac{1}{2\ell}\right)}{\log\ell} \simeq 1 + \frac{\log 2}{\log\ell},\tag{6.54}$$

and the space requirement scales as

$$S \simeq n\ell \simeq \ell \frac{\log T}{\log \ell}.$$
 (6.55)

Thus, for any fixed $\varepsilon > 0$, we can attain S scaling like log T, and a slowdown no worse than $(T_{\rm irr})^{1+\varepsilon}$. (This is not the optimal way to play the Pebble game if our objective is to get as far as we can with as few pebbles as possible. We use more pebbles to get to step T, but we get there faster.)

We have now seen that a reversible circuit can simulate a circuit composed of irreversible gates efficiently — without requiring unreasonable memory resources or causing an unreasonable slowdown. Why is this important? You might worry that, because reversible computation is "harder" than irreversible computation, the classification of complexity depends on whether we compute reversibly or irreversibly. But this is not the case, because a reversible computer can simulate an irreversible computer pretty easily.

6.2 Quantum Circuits

Now we are ready to formulate a mathematical model of a quantum computer. We will generalize the circuit model of classical computation to the quantum circuit model of quantum computation.

A classical computer processes bits. It is equipped with a finite set of gates that can be applied to sets of bits. A quantum computer processes qubits. We will assume that it too is equipped with a discrete set of fundamental components, called *quantum qates*. Each quantum gate is a unitary transformation that acts on a fixed number of qubits. In a quantum computation, a finite number n of qubits are initially set to the value $|00...0\rangle$. A circuit is executed that is constructed from a finite number of quantum gates acting on these qubits. Finally, a Von Neumann measurement of all the qubits (or a subset of the qubits) is performed, projecting each onto the basis $\{|0\rangle, |1\rangle\}$. The outcome of this measurement is the result of the computation.

Several features of this model require comment:

(1) It is implicit but important that the Hilbert space of the device has a preferred decomposition into a tensor product of low-dimensional spaces, in this case the two-dimensional spaces of the qubits. Of course, we could have considered a tensor product of, say, gutrits instead. But anyway we assume there is a natural decomposition into subsystems that is respected by the quantum gates — which act on only a few subsystems at a time. Mathematically, this feature of the gates is crucial for establishing a clearly defined notion of quantum complexity. Physically, the fundamental reason for a natural decomposition into subsystems is *locality*; feasible quantum gates must act in a bounded spatial region, so the computer decomposes into subsystems that interact only with their neighbors.

- (2) Since unitary transformations form a continuum, it may seem unnecessarily restrictive to postulate that the machine can execute only those quantum gates chosen from a discrete set. We nevertheless accept such a restriction, because we do not want to invent a new physical implementation each time we are faced with a new computation to perform.
- (3) We might have allowed our quantum gates to be superoperators, and our final measurement to be a POVM. But since we can easily simulate a superoperator by performing a unitary transformation on an extended system, or a POVM by performing a Von Neumann measurement on an extended system, the model as formulated is of sufficient generality.
- (4) We might allow the final measurement to be a collective measurement, or a projection into a different basis. But any such measurement can be implemented by performing a suitable unitary transformation followed by a projection onto the standard basis $\{|0\rangle, |1\rangle\}^n$. Of course, complicated collective measurements can be transformed into measurements in the standard basis only with some difficulty, but it is appropriate to take into account this difficulty when characterizing the complexity of an algorithm.
- (5) We might have allowed measurements at intermediate stages of the computation, with the subsequent choice of quantum gates conditioned on the outcome of those measurements. But in fact the same result can always be achieved by a quantum circuit with all measurements postponed until the end. (While we can postpone the measurements in principle, it might be very useful in practice to perform measurements at intermediate stages of a quantum algorithm.)

A quantum gate, being a unitary transformation, is reversible. In fact, a classical reversible computer is a special case of a quantum computer. A

classical reversible gate

$$x^{(n)} \to y^{(n)} = f(x^{(n)}),$$
 (6.56)

implementing a permutation of *n*-bit strings, can be regarded as a unitary transformation that acts on the "computational basis $\{|x_i\rangle\}$ according to

$$U:|x_i\rangle \to |y_i\rangle. \tag{6.57}$$

This action is unitary because the 2^n strings $|y_i\rangle$ are all mutually orthogonal. A quantum computation constructed from such classical gates takes $|0...0\rangle$ to one of the computational basis states, so that the final measurement is deterministic.

There are three main issues concerning our model that we would like to address. The first issue is *universality*. The most general unitary transformation that can be performed on n qubits is an element of $U(2^n)$. Our model would seem incomplete if there were transformations in $U(2^n)$ that we were unable to reach. In fact, we will see that there are many ways to choose a discrete set of *universal quantum gates*. Using a universal gate set we can construct circuits that compute a unitary transformation that comes as close as we please to any element in $U(2^n)$.

Thanks to universality, there is also a machine independent notion of *quantum complexity*. We may define a new complexity class BQP — the class of decision problems that can be solved, with high probability, by polynomial-size quantum circuits. Since one universal quantum computer can simulate another efficiently, the class does not depend on the details of our hardware (on the universal gate set that we have chosen).

Notice that a quantum computer can easily simulate a probabilistic classical computer: it can prepare $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and then project to $\{|0\rangle, |1\rangle\}$, generating a random bit. Therefore BQP certainly contains the class BPP. But as we discussed in Chapter 1, it seems to be quite reasonable to expect that BQP is actually larger than BPP, because a probabilistic classical computer cannot easily simulate a quantum computer. The fundamental difficulty is that the Hilbert space of n qubits is huge, of dimension 2^n , and hence the mathematical description of a typical vector in the space is exceedingly complex. Our second issue is to better characterize the resources needed to simulate a quantum computer on a classical computer. We will see that, despite the vastness of Hilbert space, a classical computer can simulate an n-qubit quantum computer even if limited to an amount of memory space that is polynomial in n. This means the BQP is contained in the complexity class PSPACE, the decision problems that can be solved with polynomial space, but may require exponential time. (We know that NP is also contained in PSPACE, since checking if $C(x^{(n)}, y^{(m)}) = 1$ for each $y^{(m)}$ can be accomplished with polynomial space.⁵

The third important issue we should address is *accuracy*. The class BQP is defined formally under the idealized assumption that quantum gates can be executed with perfect precision. Clearly, it is crucial to relax this assumption in any realistic implementation of quantum computation. A polynomial size quantum circuit family that solves a hard problem would not be of much interest if the quantum gates in the circuit were required to have exponential accuracy. In fact, we will show that this is not the case. An idealized *T*-gate quantum circuit can be simulated with acceptable accuracy by noisy gates, provided that the error probability per gate scales like 1/T.

We see that quantum computers pose a serious challenge to the strong Church–Turing thesis, which contends that any physically reasonable model of computation can be simulated by probabilistic classical circuits with at worst a polynomial slowdown. But so far there is no firm proof that

$$BPP \neq BQP. \tag{6.58}$$

Nor is such a proof necessarily soon to be expected.⁶ Indeed, a corollary would be

$$BPP \neq PSPACE,$$
 (6.59)

which would settle one of the long-standing and pivotal open questions in complexity theory.

It might be less unrealistic to hope for a proof that $BPP \neq BQP$ follows from another standard conjecture of complexity theory such as $P \neq NP$. So far no such proof has been found. But while we are not yet able to prove that quantum computers have capabilities far beyond those of conventional computers, we nevertheless might uncover evidence suggesting that $BPP \neq$ BQP. We will see that there are problems that seem to be hard (in classical computation) yet can be efficiently solved by quantum circuits.

⁵Actually there is another rung of the complexity hierarchy that may separate BQP and PSPACE; we can show that $BQP \subseteq P^{\#P} \subseteq PSPACE$, but we won't consider $P^{\#P}$ any further here.

⁶That is, we ought not to expect a "nonrelativized proof." A separation between BPP and BQP "relative to an oracle" will be established later in the chapter.

Thus it seems likely that the classification of complexity will be different depending on whether we use a classical computer or a quantum computer to solve a problem. If such a separation really holds, it is the quantum classification that should be regarded as the more fundamental, for it is better founded on the physical laws that govern the universe.

6.2.1 Accuracy

Let's discuss the issue of accuracy. We imagine that we wish to implement a computation in which the quantum gates U_1, U_2, \ldots, U_T are applied sequentially to the initial state $|\varphi_0\rangle$. The state prepared by our ideal quantum circuit is

$$|\varphi_T\rangle = \boldsymbol{U}_T \boldsymbol{U}_{T-1} \dots \boldsymbol{U}_2 \boldsymbol{U}_1 |\varphi_0\rangle. \tag{6.60}$$

But in fact our gates do not have perfect accuracy. When we attempt to apply the unitary transformation U_t , we instead apply some "nearby" unitary transformation \tilde{U}_t . (Of course, this is not the most general type of error that we might contemplate – the unitary U_t might be replaced by a *superoperator*. Considerations similar to those below would apply in that case, but for now we confine our attention to "unitary errors.")

The errors cause the actual state of the computer to wander away from the ideal state. How far does it wander? Let $|\varphi_t\rangle$ denote the ideal state after t quantum gates are applied, so that

$$|\varphi_t\rangle = \boldsymbol{U}_t |\varphi_{t-1}\rangle. \tag{6.61}$$

But if we apply the actual transformation \tilde{U}_t , then

$$\tilde{\boldsymbol{U}}_t |\varphi_{t-1}\rangle = |\varphi_t\rangle + |E_t\rangle, \qquad (6.62)$$

where

$$|E_t\rangle = (\tilde{\boldsymbol{U}}_t - \boldsymbol{U}_t)|\varphi_{t-1}\rangle, \qquad (6.63)$$

is an unnormalized vector. If $|\tilde{\varphi}_t\rangle$ denotes the actual state after t steps, then we have

$$\begin{split} |\tilde{\varphi}_1\rangle &= |\varphi_1\rangle + |E_1\rangle, \\ |\tilde{\varphi}_2\rangle &= \tilde{\boldsymbol{U}}_2 |\tilde{\varphi}_1\rangle = |\varphi_2\rangle + |E_2\rangle + \tilde{\boldsymbol{U}}_2 |E_1\rangle, \end{split}$$
(6.64)

and so forth; we ultimately obtain

$$|\tilde{\varphi}_T\rangle = |\varphi_T\rangle + |E_T\rangle + \tilde{\boldsymbol{U}}_T |E_{T-1}\rangle + \tilde{\boldsymbol{U}}_T \tilde{\boldsymbol{U}}_{T-1} |E_{T-2}\rangle + \ldots + \tilde{\boldsymbol{U}}_T \tilde{\boldsymbol{U}}_{T-1} \ldots \tilde{\boldsymbol{U}}_2 |E_1\rangle.$$
 (6.65)

Thus we have expressed the difference between $|\tilde{\varphi}_T\rangle$ and $|\varphi_T\rangle$ as a sum of T remainder terms. The worst case yielding the largest deviation of $|\tilde{\varphi}_T\rangle$ from $|\varphi_T\rangle$ occurs if all remainder terms line up in the same direction, so that the errors interfere constructively. Therefore, we conclude that

$$\| |\tilde{\varphi}_T \rangle - |\varphi_T \rangle \| \leq \| |E_T \rangle \| + \| |E_{T-1} \rangle \|$$

$$+ \ldots + \| |E_2 \rangle \| + \| |E_1 \rangle \|,$$
 (6.66)

where we have used the property $|| U | E_i \rangle || = || |E_i \rangle ||$ for any unitary U.

Let $\| A \|_{sup}$ denote the sup norm of the operator A — that is, the maximum modulus of an eigenvalue of A. We then have

$$|| |E_t\rangle || = || \left(\tilde{\boldsymbol{U}}_t - \boldsymbol{U}_t \right) |\varphi_{t-1}\rangle || \leq || \tilde{\boldsymbol{U}}_t - \boldsymbol{U}_t ||_{\sup}$$
(6.67)

(since $|\varphi_{t-1}\rangle$ is normalized). Now suppose that, for each value of t, the error in our quantum gate is bounded by

$$\| \boldsymbol{U}_t - \boldsymbol{U}_t \|_{\sup} < \varepsilon.$$
 (6.68)

Then after T quantum gates are applied, we have

$$\| \left| \tilde{\varphi}_T \right\rangle - \left| \varphi_T \right\rangle \| < T\varepsilon; \tag{6.69}$$

in this sense, the accumulated error in the state grows linearly with the length of the computation.

The distance bounded in eq. (6.68) can equivalently be expressed as $\| \mathbf{W}_t - \mathbf{1} \|_{\sup}$, where $\mathbf{W}_t = \tilde{\mathbf{U}}_t \mathbf{U}_t^{\dagger}$. Since \mathbf{W}_t is unitary, each of its eigenvalues is a phase $e^{i\theta}$, and the corresponding eigenvalue of $\mathbf{W}_t - 1$ has modulus

$$|e^{i\theta} - 1| = (2 - 2\cos\theta)^{1/2}, \tag{6.70}$$

so that eq. (6.68) is the requirement that each eigenvalue satisfies

$$\cos\theta > 1 - \varepsilon^2/2, \tag{6.71}$$

(or $|\theta| \lesssim \varepsilon$, for ε small). The origin of eq. (6.69) is clear. In each time step, $|\tilde{\varphi}\rangle$ rotates relative to $|\varphi\rangle$ by (at worst) an angle of order ε , and the distance between the vectors increases by at most of order ε .

How much accuracy is good enough? In the final step of our computation, we perform an orthogonal measurement, and the probability of outcome a, in the ideal case, is

$$P(a) = |\langle a | \varphi_T \rangle|^2. \tag{6.72}$$

Because of the errors, the actual probability is

$$\tilde{P}(a) = |\langle a | \tilde{\varphi}_T \rangle|^2.$$
 (6.73)

If the actual vector is close to the ideal vector, then the probability distributions are close, too. If we sum over an orthonormal basis $\{|a\rangle\}$, we have

$$\sum_{a} |\tilde{P}(a) - P(a)| \le 2 \parallel |\tilde{\varphi}_T\rangle - |\varphi_T\rangle \parallel, \tag{6.74}$$

as you will show in a homework exercise. Therefore, if we keep $T\varepsilon$ fixed (and small) as T gets large, the error in the probability distribution also remains fixed. In particular, if we have designed a quantum algorithm that solves a decision problem correctly with probability greater $\frac{1}{2} + \delta$ (in the ideal case), then we can achieve success probability greater than $\frac{1}{2}$ with our noisy gates, if we can perform the gates with an accuracy $T\varepsilon < O(\delta)$. A quantum circuit family in the BQP class can really solve hard problems, as long as we can improve the accuracy of the gates linearly with the computation size T.

6.2.2 BQP \subseteq **PSPACE**

Of course a classical computer can simulate any quantum circuit. But how much memory does the classical computer require? Naively, since the simulation of an *n*-qubit circuit involves manipulating matrices of size 2^n , it may seem that an amount of memory space exponential in *n* is needed. But we will now show that the simulation can be done to acceptable accuracy (albeit very slowly!) in polynomial space. This means that the quantum complexity class BQP is contained in the class PSPACE of problems that can be solved with polynomial space.

The object of the classical simulation is to compute the probability for each possible outcome a of the final measurement

$$\operatorname{Prob}(a) = |\langle a | \boldsymbol{U}_T | 0 \rangle|^2, \qquad (6.75)$$

where

$$\boldsymbol{U}_T = \boldsymbol{U}_T \boldsymbol{U}_{T-1} \dots \boldsymbol{U}_2 \boldsymbol{U}_1, \qquad (6.76)$$

is a product of T quantum gates. Each U_t , acting on the n qubits, can be represented by a $2^n \times 2^n$ unitary matrix, characterized by the complex matrix elements

$$\langle y|\boldsymbol{U}_t|x\rangle,$$
 (6.77)

where $x, y \in \{0, 1, ..., 2^n - 1\}$. Writing out the matrix multiplication explicitly, we have

$$\langle a | \boldsymbol{U}_T | 0 \rangle = \sum_{\{x_t\}} \langle a | \boldsymbol{U}_T | x_{T-1} \rangle \langle x_{T-1} | \boldsymbol{U}_{T-1} | x_{T-2} \rangle \dots$$
$$\dots \langle x_2 | \boldsymbol{U}_2 | x_1 \rangle \langle x_1 | \boldsymbol{U}_1 | 0 \rangle. \tag{6.78}$$

Eq. (6.78) is a sort of "path integral" representation of the quantum computation – the probability amplitude for the final outcome a is expressed as a coherent sum of amplitudes for each of a vast number $(2^{n(T-1)})$ of possible computational paths that begin at 0 and terminate at a after T steps.

Our classical simulator is to add up the $2^{n(T-1)}$ complex numbers in eq. (6.78) to compute $\langle a | \boldsymbol{U}_T | 0 \rangle$. The first problem we face is that finite size classical circuits do integer arithmetic, while the matrix elements $\langle y | \boldsymbol{U}_t | x \rangle$ need not be rational numbers. The classical simulator must therefore settle for an approximate calculation to reasonable accuracy. Each term in the sum is a product of T complex factors, and there are $2^{n(T-1)}$ terms in the sum. The accumulated errors are sure to be small if we express the matrix elements to m bits of accuracy, with m large compared to n(T-1). Therefore, we can replace each complex matrix element by pairs of signed integers, taking values in $\{0, 1, 2, \ldots, 2^{m-1}\}$. These integers give the binary expansion of the real and imaginary part of the matrix element, expressed to precision 2^{-m} .

Our simulator will need to compute each term in the sum eq. (6.78) and accumulate a total of all the terms. But each addition requires only a modest amount of scratch space, and furthermore, since only the accumulated subtotal need be stored for the next addition, not much space is needed to sum all the terms, even though there are exponentially many.

So it only remains to consider the evaluation of a typical term in the sum, a product of T matrix elements. We will require a classical circuit that

evaluates

$$\langle y|\boldsymbol{U}_t|x\rangle;$$
 (6.79)

this circuit accepts the 2n bit input (x, y), and outputs the 2m-bit value of the (complex) matrix element. Given a circuit that performs this function, it will be easy to build a circuit that multiplies the complex numbers together without using much space.

Finally, at this point, we appeal to the properties we have demanded of our quantum gate set — the gates from a discrete set, and each gate acts on a bounded number of qubits. Because there are a fixed (and finite) number of gates, there are only a fixed number of gate subroutines that our simulator needs to be able to call. And because the gate acts on only a few qubits, nearly all of its matrix elements vanish (when n is large), and the value $\langle y|\mathbf{U}|x\rangle$ can be determined (to the required accuracy) by a simple circuit requiring little memory.

For example, in the case of a single-qubit gate acting on the first qubit, we have

$$\langle y_1 y_2 \dots y_n | \boldsymbol{U} | x_1 x_2 \dots x_n \rangle = 0 \text{ if } x_2 x_3 \dots x_n \neq y_2 y_3 \dots y_n.$$
(6.80)

A simple circuit can compare x_2 with y_2, x_3 with y_3 , *etc.*, and output zero if the equality is not satisfied. In the event of equality, the circuit outputs one of the four complex numbers

$$\langle y_1 | \boldsymbol{U} | x_1 \rangle,$$
 (6.81)

to m bits of precision. A simple circuit can encode the 8m bits of this 2×2 complex-valued matrix. Similarly, a simple circuit, requiring only space polynomial in n and m, can evaluate the matrix elements of any gate of fixed size.

We conclude that a classical computer with space bounded above by poly(n) can simulate an *n*-qubit universal quantum computer, and therefore that BQP \subseteq PSPACE. Of course, it is also evident that the simulation we have described requires exponential time, because we need to evaluate the sum of $2^{n(T-1)}$ complex numbers. (Actually, most of the terms vanish, but there are still an exponentially large number of nonvanishing terms.)

6.2.3 Universal quantum gates

We must address one more fundamental question about quantum computation; how do we construct an adequate set of quantum gates? In other words, what constitutes a universal quantum computer?

We will find a pleasing answer. Any generic two-qubit gate suffices for universal quantum computation. That is, for all but a set of measure zero of 4×4 unitary matrices, if we can apply that matrix to any pair of qubits, then we can construct an *n*-qubit circuit that computes a transformation that comes as close as we please to any element of $U(2^n)$.

Mathematically, this is not a particularly deep result, but physically it is very interesting. It means that, in the quantum world, as long as we can devise a generic interaction between two qubits, and we can implement that interaction accurately between any two qubits, we can compute anything, no matter how complex. Nontrivial computation is ubiquitous in quantum theory.

Aside from this general result, it is also of some interest to exhibit particular universal gate sets that might be particularly easy to implement physically. We will discuss a few examples.

There are a few basic elements that enter the analysis of any universal quantum gate set.

(1) Powers of a generic gate

Consider a "generic" k-qubit gate. This is a $2^k \times 2^k$ unitary matrix U with eigenvalues $e^{i\theta_1}, e^{i\theta_2}, \ldots e^{i\theta_{2^k}}$. For all but a set of measure zero of such matrices, each θ_i is an irrational multiple of π , and all the θ_i 's are incommensurate (each θ_i/θ_j is also irrational). The positive integer power U^n of U has eigenvalues

$$e^{in\theta_1}, e^{in\theta_2}, \dots, e^{in\theta_{2^k}}.$$
(6.82)

Each such list of eigenvalues defines a point in a 2^k -dimensional torus (the product of 2^k circles). As *n* ranges over positive integer values, these points densely fill the whole torus, if \boldsymbol{U} is generic. If $\boldsymbol{U} = e^{iA}$, positive integer powers of \boldsymbol{U} come as close as we please to $\boldsymbol{U}(\lambda) = e^{i\lambda A}$, for any real λ . We say that any $\boldsymbol{U}(\lambda)$ is *reachable* by positive integer powers of \boldsymbol{U} .

(2) Switching the leads

There are a few (classical) transformations that we can implement just by switching the labels on k qubits, or in other words, by applying the gate U to the qubits in a different order. Of the $(2^k)!$ permutations of the length-k strings, k! can be realized by swapping qubits. If a gate applied to k qubits with a standard ordering is U, and P is a permutation implemented by swapping qubits, then we can construct the gate

$$\boldsymbol{U}' = P\boldsymbol{U}P^{-1},\tag{6.83}$$

just by switching the leads on the gate. For example, swapping two qubits implements the transposition

$$P:|01\rangle \leftrightarrow |10\rangle, \tag{6.84}$$

or

$$P = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$
(6.85)

acting on basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$. By switching leads, we obtain a gate

$$\Box U' = \Box P = U = P^{-1} =$$

We can also construct any positive integer power of U', $(PUP^{-1})^n = PU^nP^{-1}$.

(3) Completing the Lie algebra

We already remarked that if $U = e^{iA}$ is generic, then powers of U are dense in the torus $\{e^{i\lambda A}\}$. We can further argue that if $U = e^{iA}$ and $U' = e^{iB}$ are generic gates, we can compose them to come arbitrarily close to

$$e^{i(\alpha A+\beta B)}$$
 or $e^{-\gamma[A,B]}$, (6.86)

6.2. QUANTUM CIRCUITS

for any real α, β, γ . Thus, the "reachable" transformations have a closed *Lie algebra*. We say that $U = e^{iA}$ is generated by *A*; then if *A* and *B* are both generic generators of reachable transformations, so are real linear combinations of *A* and *B*, and (*i* times) the commutator of *A* and *B*.

We first note that

$$\lim_{n \to \infty} (e^{i\alpha A/n} e^{i\beta B/n})^n = \lim_{n \to \infty} \left(1 + \frac{i}{n} (\alpha A + \beta B) \right)^n = e^{i(\alpha A + \beta B)}.$$
(6.87)

Therefore, any $e^{i(\alpha A+\beta B)}$ is reachable if each $e^{i\alpha A/n}$ and $e^{i\beta B/n}$ is. Furthermore

$$\lim_{n \to \infty} \left(e^{iA/\sqrt{n}} e^{iB/\sqrt{n}} e^{-iA/\sqrt{n}} e^{-iB/\sqrt{n}} \right)^n = \lim_{n \to \infty} \left[1 - \frac{1}{n} (AB - BA) \right]^n = e^{-[A,B]},$$
(6.88)

so $e^{-[A,B]}$ is also reachable.

By invoking the observations (1), (2), and (3) above, we will be able to show that a generic two-qubit gate is universal.

Deutsch gate. It was David Deutsch (1989) who first pointed out the existence of a universal quantum gate. Deutsch's three-bit universal gate is a quantum cousin of the Toffoli gate. It is the controlled-controlled-R transformation



that applies \boldsymbol{R} to the third qubit if the first two qubits have the value 1; otherwise it acts trivially. Here

$$\boldsymbol{R} = -i\boldsymbol{R}_x(\theta) = (-i)\exp\left(i\frac{\theta}{2}\boldsymbol{\sigma}_x\right) = (-i)\left(\cos\frac{\theta}{2} + i\boldsymbol{\sigma}_x\sin\frac{\theta}{2}\right)$$
(6.89)

is, up to a phase, a rotation by θ about the x-axis, where θ is a particular angle incommensurate with π .

The *n*th power of the Deutsch gate is the controlled-controlled- \mathbf{R}^n . In particular, $\mathbf{R}^4 = \mathbf{R}_x(4\theta)$, so that all one-qubit transformations generated by $\boldsymbol{\sigma}_x$ are reachable by integer powers of \mathbf{R} . Furthermore the (4n + 1)st power is

$$(-i)\left[\cos\frac{(4n+1)\theta}{2} + i\boldsymbol{\sigma}_x\sin\frac{(4n+1)\theta}{2}\right],\tag{6.90}$$

which comes as close as we please to σ_x . Therefore, the Toffoli gate is reachable with integer powers of the Deutsch gate, and the Deutsch gate is universal for classical computation.

Acting on the three-qubit computational basis

$$\{|000\rangle, |001\rangle, |010\rangle, |011\rangle, |100\rangle, |101\rangle, |110\rangle, |111\rangle\}, \tag{6.91}$$

the generator of the Deutsch gate transposes the last two elements

$$|110\rangle \leftrightarrow |111\rangle. \tag{6.92}$$

We denote this 8×8 matrix as

$$(\boldsymbol{\sigma}_x)_{67} = \begin{pmatrix} 0 & 0 \\ \hline \\ 0 & \sigma_x \end{pmatrix}.$$
 (6.93)

With Toffoli gates, we can perform any permutation of these eight elements, in particular

$$P = (6m)(7n), (6.94)$$

for any m and n. So we can also reach any transformation generated by

$$P(\boldsymbol{\sigma}_x)_{67}P = (\boldsymbol{\sigma}_x)_{mn}.$$
(6.95)

Furthermore,

$$[(\boldsymbol{\sigma}_x)_{56}, (\boldsymbol{\sigma}_x)_{67}] = \begin{bmatrix} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \end{bmatrix} = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix} = i(\boldsymbol{\sigma}_y)_{57},$$
(6.96)

and similarly, we can reach any unitary generated by $(\sigma_y)_{mn}$. Finally

$$[(\boldsymbol{\sigma}_x)_{mn}, (\boldsymbol{\sigma}_y)_{mn}] = i(\boldsymbol{\sigma}_z)_{mn}, \qquad (6.97)$$

So we can reach any transformation generated by a linear combination of the $(\sigma_{x,y,z})_{mn}$'s. These span the whole SU(8) Lie Algebra, so we can generate any three-qubit unitary (aside from an irrelevant overall phase).

Now recall that we have already found that we can construct the n-bit Toffoli gate by composing three-bit Toffoli gates. The circuit



uses one scratch bit to construct a four-bit Deutsch gate ((controlled)³- \mathbf{R}) from the three-bit Deutsch gate and two three-bit Toffoli gates, and a similar circuit constructs the *n*-bit Deutsch gate from a three-bit Deutsch gate and two (n - 1)-bit Toffoli gates. Once we have an *n*-bit Deutsch gate, and universal classical computation, exactly the same argument as above shows that we can reach any transformation in $SU(2^n)$.

Universal two-qubit gates. For reversible classical computation, we saw that three-bit gates are needed for universality. But in quantum computation, two-bit gates turn out to be adequate. Since we already know that the Deutsch gate is universal, we can establish this by showing that the Deutsch gate can be constructed by composing two-qubit gates.

In fact, if



denotes the controlled-U gate (the 2 × 2 unitary U is applied to the second qubit if the first qubit is 1; otherwise the gate acts trivially) then a controlledcontrolled- U^2 gate is obtained from the circuit



the power of \boldsymbol{U} applied to the third qubit is

$$y - (x \oplus y) + x = x + y - (x + y - 2xy) = 2xy.$$
(6.98)

Therefore, we can construct Deutsch's gate from the controlled-U, controlled U^{-1} and controlled-NOT gates, where

$$\boldsymbol{U}^2 = -i\boldsymbol{R}_x(\theta); \tag{6.99}$$

we may choose

$$\boldsymbol{U} = e^{-i\frac{\pi}{4}} \boldsymbol{R}_x \left(\frac{\theta}{2}\right). \tag{6.100}$$

Positive powers of U came as close as we please to σ_x and U^{-1} , so from the controlled-U alone we can construct the Deutsch gate. Therefore, the controlled- $\left(e^{-i\frac{\pi}{4}} R_x\left(\frac{\theta}{2}\right)\right)$ is itself a universal gate, for θ/π irrational.

(Note that the above construction shows that, while we cannot construct the Toffoli gate from two-bit reversible classical gates, we *can* construct it from a controlled "square root of NOT" — a controlled-U with $U^2 = \sigma_x$.)

Generic two-bit gates. Now we have found particular two-bit gates (controlled rotations) that are universal gates. Therefore, for universality, it is surely sufficient if we can construct transformations that are dense in the U(4) acting on a pair of qubits.

In fact, though, any generic two-qubit gate is sufficient to generate all of U(4). As we have seen, if $e^{i\boldsymbol{A}}$ is a generic element of U(4), we can reach any transformation generated by \boldsymbol{A} . Furthermore, we can reach any transformations generated by an element of the minimal Lie algebra containing \boldsymbol{A} and

$$\boldsymbol{B} = P\boldsymbol{A}P^{-1} \tag{6.101}$$

where P is the permutation $(|01\rangle \leftrightarrow |10\rangle)$ obtained by switching the leads.

Now consider a general A, (expanded in terms of a basis for the Lie algebra of U(4)), and consider a particular scheme for constructing 16 elements of the algebra by successive commutations, starting from A and B. The elements so constructed are linearly independent (and it follows that any transformation in U(4) is reachable) if the determinant of a particular 16×16 matrix vanishes. Unless this vanishes identically, its zeros occur only on a submanifold of vanishing measure. But in fact, we can choose, say

$$\boldsymbol{A} = (\alpha I + \beta \sigma_x + \gamma \sigma_y)_{23}, \tag{6.102}$$

(for incommensurate α, β, γ), and show by explicit computation that the entire 16-dimension Lie Algebra is actually generated by successive commutations, starting with \boldsymbol{A} and \boldsymbol{B} . Hence we conclude that failure to generate the entire U(4) algebra is nongeneric, and find that almost all two-qubit gates are universal.

Other adequate sets of gates. One can also see that universal quantum computation can be realized with a gate set consisting of *classical* multiqubit gates and quantum single-qubit gates. For example, we can see that the XOR gate, combined with one-qubit gates, form a universal set. Consider the circuit



which applies ABC to the second qubit if x = 0, and $A\sigma_x B\sigma_x C$ to the second qubit if x = 1. If we can find A, B, C such that

$$ABC = 1$$
$$A\sigma_x B\sigma_x C = U, \qquad (6.103)$$

then this circuit functions as a controlled-U gate. In fact unitary 2 × 2 A, B, C with this property exist for any unitary U with determinant one (as you'll show in an exercise). Therefore, the XOR plus arbitrary one-qubit transformations form a universal set. Of course, two generic (noncommuting) one-qubit transformations are sufficient to reach all. In fact, with an XOR and a *single* generic one-qubit rotation, we can construct a second one-qubit rotation that does not commute with the first. Hence, an XOR together with just one generic single-qubit gate constitutes a universal gate set.

If we are able to perform a Toffoli gate, then even certain nongeneric one-qubit transformations suffice for universal computation. For example (another exercise) the Toffoli gate, together with $\pi/2$ rotations about the x and z axes, are a universal set.

Precision. Our discussion of universality has focused on *reachability* without any regard for *complexity*. We have only established that we can construct a quantum circuit that comes as close as we please to a desired element of $U(2^n)$, and we have not considered the size of the circuit that we need. But from the perspective of quantum complexity theory, universality is quite significant because it implies that one quantum computer can simulate another to reasonable accuracy without an unreasonable slowdown.

Actually, we have not been very precise up until now about what it means for one unitary transformation to be "close" to another; we should define a topology. One possibility is to use the sup norm as in our previous discussion of accuracy — the distance between matrices \boldsymbol{U} and \boldsymbol{W} is then $\| \boldsymbol{U} - \boldsymbol{W} \|_{sup}$. Another natural topology is associated with the inner product

$$\langle \boldsymbol{W} | \boldsymbol{U} \rangle \equiv \operatorname{tr} \boldsymbol{W}^{\dagger} \boldsymbol{U}$$
 (6.104)

(if U and W are $N \times N$ matrices, this is just the usual inner product on C^{N^2} , where we regard U_{ij} as a vector with N^2 components). Then we may define the distance squared between matrices as

$$\|\boldsymbol{U} - \boldsymbol{W}\|^2 \equiv \langle \boldsymbol{U} - \boldsymbol{W} | \boldsymbol{U} - \boldsymbol{W} \rangle.$$
(6.105)

For the purpose of analyzing complexity, just about any reasonable topology will do.

The crucial point is that given any universal gate set, we can reach within distance ε of any desired unitary transformation that acts on a fixed number of qubits, using a quantum circuit whose size is bounded above by a polynomial in ε^{-1} . Therefore, one universal quantum computer can simulate another, to accuracy ε , with a slowdown no worse than a factor that is polynomial in ε^{-1} . Now we have already seen that to have a high probability of getting the right answer when we perform a quantum circuit of size T, we should implement each quantum gate to an accuracy that scales like T^{-1} . Therefore, if you have a quantum circuit family of polynomial size that runs

6.3. SOME QUANTUM ALGORITHMS

on your quantum computer, I can devise a polynomial size circuit family that runs on my machine, and that emulates your machine to acceptable accuracy.

Why can a poly (ε^{-1}) -size circuit reach a given k-qubit U to within distance ε ? We know for example that the positive integer powers of a generic k-qubit $e^{i\mathbf{A}}$ are dense in the 2^k -torus $\{e^{i\lambda\mathbf{A}}\}$. The region of the torus within distance ε of any given point has volume of order ε^{2^k} , so (asymptotically for ε sufficiently small) we can reach any $\{e^{i\lambda\mathbf{A}}\}$ to within distance ε with $(e^{i\lambda\mathbf{A}})^n$, for some integer n of order ε^{-2^k} . We also know that we can obtain transformations $\{e^{i\mathbf{A}_a}\}$ where the \mathbf{A}_a 's span the full $U(2^k)$ Lie algebra, using circuits of fixed size (independent of ε). We may then approach any $\exp(i\sum_a \alpha_a \mathbf{A}_a)$ as in eq. (6.87), also with polynomial convergence.

In principle, we should be able to do much better, reaching a desired k-qubit unitary within distance ε using just $poly(log(\varepsilon^{-1}))$ quantum gates. Since the number of size-T circuits that we can construct acting on k qubits is exponential in T, and the circuits fill $U(2^k)$ roughly uniformly, there should be a size-T circuit reaching within a distance of order e^{-T} of any point in $U(2^k)$. However, it might be a computationally hard problem *classically* to work out the circuit that comes exponentially close to the unitary we are trying to reach. Therefore, it would be dishonest to rely on this more efficient construction in an asymptotic analysis of quantum complexity.

6.3 Some Quantum Algorithms

While we are not yet able to show that $BPP \neq BQP$, there are three approaches that we can pursue to study the differences between the capabilities of classical and quantum computers:

- (1) Nonexponential speedup. We can find quantum algorithms that are demonstrably faster than the best classical algorithm, but not *exponentially* faster. These algorithms shed no light on the conventional classification of complexity. But they do demonstrate a type of separation between tasks that classical and quantum computers can perform. Example: Grover's quantum speedup of the search of an unsorted data base.
- (2) "Relativized" exponential speedup. We can consider the problem of analyzing the contents of a "quantum black box." The box performs an

a priori unknown) unitary transformation. We can prepare an input for the box, and we can measure its output; our task is to find out what the box does. It is possible to prove that quantum black boxes (computer scientists call them oracles⁷) exist with this property: By feeding quantum superpositions to the box, we can learn what is inside with an *exponential* speedup, compared to how long it would take if we were only allowed classical inputs. A computer scientist would say that $BPP \neq BQP$ "relative to the oracle." Example: Simon's exponential quantum speedup for finding the period of a 2 to 1 function.

(3) Exponential speedup for "apparently" hard problems. We can exhibit a quantum algorithm that solves a problem in polynomial time, where the problem appears to be hard classically, so that it is strongly suspected (though not proved) that the problem is not in *BPP*. Example: Shor's factoring algorithm.

Deutsch's problem. We will discuss examples from all three approaches. But first, we'll warm up by recalling an example of a simple quantum algorithm that was previously discussed in §1.5: Deutsch's algorithm for distinguishing between constant and balanced functions $f : \{0, 1\} \rightarrow \{0, 1\}$. We are presented with a quantum black box that computes f(x); that is, it enacts the two-qubit unitary transformation

$$U_f: |x\rangle|y\rangle \to |x\rangle|y \oplus f(x)\rangle,$$
 (6.106)

which flips the second qubit iff f(first qubit) = 1. Our assignment is to determine whether f(0) = f(1). If we are restricted to the "classical" inputs $|0\rangle$ and $|1\rangle$, we need to access the box twice (x = 0 and x = 1) to get the answer. But if we are allowed to input a coherent superposition of these "classical" states, then once is enough.

The quantum circuit that solves the problem (discussed in $\S1.5$) is:



⁷The term "oracle" signifies that the box responds to a query *immediately*; that is, the time it takes the box to operate is not included in the complexity analysis.

Here H denotes the Hadamard transform

$$\boldsymbol{H}:|x\rangle \to \frac{1}{\sqrt{2}} \sum_{y} (-1)^{xy} |y\rangle, \qquad (6.107)$$

or

$$H: |0\rangle \to \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$$
$$|1\rangle \to \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle); \qquad (6.108)$$

that is, \boldsymbol{H} is the 2 \times 2 matrix

$$\boldsymbol{H}: \left(\begin{array}{cc} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{array}\right). \tag{6.109}$$

The circuit takes the input $|0\rangle|1\rangle$ to

$$\begin{aligned} |0\rangle|1\rangle &\to \frac{1}{2}(|0\rangle + |1\rangle)(|0\rangle - |1\rangle) \\ &\to \frac{1}{2}\left((-1)^{f(0)}|0\rangle + (-1)^{f(1)}|1\rangle\right)(|0\rangle - |1\rangle) \\ &\to \frac{1}{2}\left[\left((-1)^{f(0)} + (-1)^{f(1)}\right)|0\rangle \\ &+ \left((-1)^{f(0)} - (-1)^{f(1)}\right)|1\rangle\right] \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle). \end{aligned}$$
(6.110)

Then when we measure the first qubit, we find the outcome $|0\rangle$ with probability one if f(0) = f(1) (constant function) and the outcome $|1\rangle$ with probability one if $f(0) \neq f(1)$ (balanced function).

A quantum computer enjoys an advantage over a classical computer because it can invoke *quantum parallelism*. Because we input a superposition of $|0\rangle$ and $|1\rangle$, the output is sensitive to both the values of f(0) and f(1), even though we ran the box just once.

Deutsch–Jozsa problem. Now we'll consider some generalizations of Deutsch's problem. We will continue to assume that we are to analyze a quantum black box ("quantum oracle"). But in the hope of learning something about complexity, we will imagine that we have a family of black boxes,

with variable input size. We are interested in how the time needed to find out what is inside the box scales with the size of the input (where "time" is measured by how many times we query the box).

In the *Deutsch–Jozsa problem*, we are presented with a quantum black box that computes a function taking n bits to 1,

$$f: \{0,1\}^n \to \{0,1\},$$
 (6.111)

and we have it on good authority that f is either constant (f(x) = c for all x) or balanced (f(x) = 0 for exactly $\frac{1}{2}$ of the possible input values). We are to solve the decision problem: Is f constant or balanced?

In fact, we can solve this problem, too, accessing the box only once, using the same circuit as for Deutsch's problem (but with x expanded from one bit to n bits). We note that if we apply n Hadamard gates in parallel to n-qubits.

$$\boldsymbol{H}^{(n)} = \boldsymbol{H} \otimes \boldsymbol{H} \otimes \ldots \otimes \boldsymbol{H}, \tag{6.112}$$

then the n-qubit state transforms as

$$\boldsymbol{H}^{(n)}:|x\rangle \to \prod_{i=1}^{n} \left(\frac{1}{\sqrt{2}} \sum_{y_i=\{0,1\}} (-1)^{x_i y_i} |y_i\rangle\right) \equiv \frac{1}{2^{n/2}} \sum_{y=0}^{2^{n-1}} (-1)^{x \cdot y} |y\rangle, \tag{6.113}$$

where x, y represent *n*-bit strings, and $x \cdot y$ denotes the *bitwise* AND (or mod 2 scalar product)

$$x \cdot y = (x_1 \wedge y_1) \oplus (x_2 \wedge y_2) \oplus \ldots \oplus (x_n \wedge y_n).$$
(6.114)

Acting on the input $(|0\rangle)^n |1\rangle$, the action of the circuit is

$$(|0\rangle)^{n}|1\rangle \rightarrow \left(\frac{1}{2^{n/2}}\sum_{x=0}^{2^{n}-1}|x\rangle\right)\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \rightarrow \left(\frac{1}{2^{n/2}}\sum_{x=0}^{2^{n}-1}(-1)^{f(x)}|x\rangle\right)\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \rightarrow \left(\frac{1}{2^{n}}\sum_{x=0}^{2^{n}-1}\sum_{y=0}^{2^{n}-1}(-1)^{f(x)}(-1)^{x\cdot y}|y\rangle\right)\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$$
(6.115)

Now let us evaluate the sum

$$\frac{1}{2^n} \sum_{x=0}^{2^n-1} (-1)^{f(x)} (-1)^{x \cdot y}.$$
(6.116)

6.3. SOME QUANTUM ALGORITHMS

If f is a constant function, the sum is

$$(-1)^{f(x)} \left(\frac{1}{2^n} \sum_{x=0}^{2^n - 1} (-1)^{x \cdot y} \right) = (-1)^{f(x)} \delta_{y,0}; \tag{6.117}$$

it vanishes unless y = 0. Hence, when we measure the *n*-bit register, we obtain the result $|y = 0\rangle \equiv (|0\rangle)^n$ with probability one. But if the function is balanced, then for y = 0, the sum becomes

$$\frac{1}{2^n} \sum_{x=0}^{2^{n-1}} (-1)^{f(x)} = 0, \qquad (6.118)$$

(because half of the terms are (+1) and half are (-1)). Therefore, the probability of obtaining the measurement outcome $|y = 0\rangle$ is zero.

We conclude that one query of the quantum oracle suffices to distinguish constant and balanced function with 100% confidence. The measurement result y = 0 means constant, any other result means balanced.

So quantum computation solves this problem neatly, but is the problem really hard classically? If we are restricted to classical input states $|x\rangle$, we can query the oracle repeatedly, choosing the input x at random (without replacement) each time. Once we obtain distinct outputs for two different queries, we have determined that the function is balanced (not constant). But if the function is in fact constant, we will not be *certain* it is constant until we have submitted $2^{n-1}+1$ queries and have obtained the same response every time. In contrast, the quantum computation gives a definite response in only one go. So in this sense (if we demand absolute certainty) the classical calculation requires a number of queries exponential in n, while the quantum computation does not, and we might therefore claim an exponential quantum speedup.

But perhaps it is not reasonable to demand absolute certainty of the classical computation (particularly since any real quantum computer will be susceptible to errors, so that the quantum computer will also be unable to attain absolute certainty.) Suppose we are satisfied to guess balanced or constant, with a probability of success

$$P(\text{success}) > 1 - \varepsilon. \tag{6.119}$$

If the function is actually balanced, then if we make k queries, the probability of getting the same response every time is $p = 2^{-(k-1)}$. If after receiving the

same response k consecutive times we guess that the function is balanced, then a quick Bayesian analysis shows that the probability that our guess is wrong is $\frac{1}{2^{k-1}+1}$ (assuming that balanced and constant are a priori equally probable). So if we guess after k queries, the probability of a wrong guess is

$$1 - P(\text{success}) = \frac{1}{2^{k-1}(2^{k-1}+1)}.$$
(6.120)

Therefore, we can achieve success probability $1-\varepsilon$ for $\varepsilon^{-1} = 2^{k-1}(2^{k-1}+1)$ or $k \sim \frac{1}{2} \log \left(\frac{1}{\varepsilon}\right)$. Since we can reach an exponentially good success probability with a polynomial number of trials, it is not really fair to say that the problem is hard.

Bernstein–Vazirani problem. Exactly the same circuit can be used to solve another variation on the Deutsch–Jozsa problem. Let's suppose that our quantum black box computes one of the functions f_a , where

$$f_a(x) = a \cdot x, \tag{6.121}$$

and a is an n-bit string. Our job is to determine a.

The quantum algorithm can solve this problem with certainty, given just one (n-qubit) quantum query. For this particular function, the quantum state in eq. (6.115) becomes

$$\frac{1}{2^n} \sum_{x=0}^{2^n-1} \sum_{y=0}^{2^n-1} (-1)^{a \cdot x} (-1)^{x \cdot y} |y\rangle.$$
(6.122)

But in fact

$$\frac{1}{2^n} \sum_{x=0}^{2^n-1} (-1)^{a \cdot x} (-1)^{x \cdot y} = \delta_{a,y}, \qquad (6.123)$$

so this state is $|a\rangle$. We can execute the circuit once and measure the *n*-qubit register, finding the *n*-bit string *a* with probability one.

If only classical queries are allowed, we acquire only one bit of information from each query, and it takes n queries to determine the value of a. Therefore, we have a clear separation between the quantum and classical difficulty of the problem. Even so, this example does not probe the relation of BPPto BQP, because the classical problem is not hard. The number of queries required classically is only linear in the input size, not exponential.

6.3. SOME QUANTUM ALGORITHMS

Simon's problem. Bernstein and Vazirani managed to formulate a variation on the above problem that *is* hard classically, and so establish for the first time a "relativized" separation between quantum and classical complexity. We will find it more instructive to consider a simpler example proposed somewhat later by Daniel Simon.

Once again we are presented with a quantum black box, and this time we are assured that the box computes a function

$$f: \{0,1\}^n \to \{0,1\}^n,$$
 (6.124)

that is 2-to-1. Furthermore, the function has a "period" given by the n-bit string a; that is

$$f(x) = f(y) \qquad \text{iff} \qquad y = x \oplus a, \tag{6.125}$$

where here \oplus denotes the bitwise XOR operation. (So *a* is the period if we regard *x* as taking values in $(Z_2)^n$ rather than Z_{2^n} .) This is all we know about *f*. Our job is to determine the value of *a*.

Classically this problem is *hard*. We need to query the oracle an exponentially large number of times to have any reasonable probability of finding a. We don't learn anything until we are fortunate enough to choose two queries x and y that happen to satisfy $x \oplus y = a$. Suppose, for example, that we choose $2^{n/4}$ queries. The number of pairs of queries is less than $(2^{n/4})^2$, and for each pair $\{x, y\}$, the probability that $x \oplus y = a$ is 2^{-n} . Therefore, the probability of successfully finding a is less than

$$2^{-n}(2^{n/4})^2 = 2^{-n/2}; (6.126)$$

even with exponentially many queries, the success probability is exponentially small.

If we wish, we can frame the question as a decision problem: Either f is a 1-1 function, or it is 2-to-1 with some randomly chosen period a, each occurring with an a priori probability $\frac{1}{2}$. We are to determine whether the function is 1-to-1 or 2-to-1. Then, after $2^{n/4}$ classical queries, our probability of making a correct guess is

$$P(\text{success}) < \frac{1}{2} + \frac{1}{2^{n/2}},$$
 (6.127)

which does not remain bounded away from $\frac{1}{2}$ as n gets large.

But with quantum queries the problem is easy! The circuit we use is essentially the same as above, but now *both* registers are expanded to nqubits. We prepare the equally weighted superposition of all n-bit strings (by acting on $|0\rangle$ with $\mathbf{H}^{(n)}$), and then we query the oracle:

$$U_f:\left(\sum_{x=0}^{2^n-1}|x\rangle\right)|0\rangle \to \sum_{x=0}^{2^n-1}|x\rangle|f(x)\rangle.$$
(6.128)

Now we measure the second register. (This step is not actually necessary, but I include it here for the sake of pedagogical clarity.) The measurement outcome is selected at random from the 2^{n-1} possible values of f(x), each occurring equiprobably. Suppose the outcome is $f(x_0)$. Then because both x_0 and $x_0 \oplus a$, and only these values, are mapped by f to $f(x_0)$, we have prepared the state

$$\frac{1}{\sqrt{2}}(|x_0\rangle + |x_0 \oplus a\rangle) \tag{6.129}$$

in the first register.

Now we want to extract some information about a. Clearly it would do us no good to measure the register (in the computational basis) at this point. We would obtain either the outcome x_0 or $x_0 \oplus a$, each occurring with probability $\frac{1}{2}$, but neither outcome would reveal anything about the value of a.

But suppose we apply the Hadamard transform $\boldsymbol{H}^{(n)}$ to the register before we measure:

$$\boldsymbol{H}^{(n)} : \frac{1}{\sqrt{2}} (|x_0\rangle + |x_0 + a\rangle)
\rightarrow \frac{1}{2^{(n+1)/2}} \sum_{y=0}^{2^n - 1} \left[(-1)^{x_0 \cdot y} + (-1)^{(x_0 \oplus a) \cdot y} \right] |y\rangle
= \frac{1}{2^{(n-1)/2}} \sum_{a \cdot y = 0} (-1)^{x_0 \cdot y} |y\rangle.$$
(6.130)

If $a \cdot y = 1$, then the terms in the coefficient of $|y\rangle$ interfere destructively. Hence only states $|y\rangle$ with $a \cdot y = 0$ survive in the sum over y. The measurement outcome, then, is selected at random from all possible values of y such that $a \cdot y = 0$, each occurring with probability $2^{-(n-1)}$.

6.4. QUANTUM DATABASE SEARCH

We run this algorithm repeatedly, each time obtaining another value of y satisfying $y \cdot a = 0$. Once we have found n such linearly independent values $\{y_1, y_2, y_3 \dots y_n\}$ (that is, linearly independent over $(Z_2)^n$), we can solve the equations

$$y_{1} \cdot a = 0$$

$$y_{2} \cdot a = 0$$

$$\vdots$$

$$y_{n} \cdot a = 0,$$

(6.131)

to determine a unique value of a, and our problem is solved. It is easy to see that with O(n) repetitions, we can attain a success probability that is exponentially close to 1.

So we finally have found an example where, given a particular type of quantum oracle, we can solve a problem in polynomial time by exploiting quantum superpositions, while exponential time is required if we are limited to classical queries. As a computer scientist might put it:

There exists an oracle relative to which $BQP \neq BPP$.

Note that whenever we compare classical and quantum complexity relative to an oracle, we are considering a quantum oracle (queries and replies are states in Hilbert space), but with a preferred orthonormal basis. If we submit a classical query (an element of the preferred basis) we always receive a classical response (another basis element). The issue is whether we can achieve a significant speedup by choosing more general quantum queries.

6.4 Quantum Database Search

The next algorithm we will study also exhibits, like Simon's algorithm, a speedup with respect to what can be achieved with a classical algorithm. But in this case the speedup is merely quadratic (the quantum time scales like the square root of the classical time), in contrast to the exponential speedup in the solution to Simon's problem. Nevertheless, the result (discovered by Lov Grover) is extremely interesting, because of the broad utility of the algorithm.

Heuristically, the problem we will address is: we are confronted by a very large unsorted database containing $N \gg 1$ items, and we are to locate one particular item, to find a needle in the haystack. Mathematically, the database is represented by a table, or a function f(x), with $x \in \{0, 1, 2, \ldots, N-1\}$. We have been assured that the entry a occurs in the table exactly once; that is, that f(x) = a for only one value of x. The problem is, given a, to find this value of x.

If the database has been properly *sorted*, searching for x is easy. Perhaps someone has been kind enough to list the values of a in ascending order. Then we can find x by looking up only $\log_2 N$ entries in the table. Let's suppose $N \equiv 2^n$ is a power of 2. First we look up f(x) for $x = 2^{n-1} - 1$, and check if f(x) is greater than a. If so, we next look up f at $x = 2^{n-2} - 1$, *etc.* With each table lookup, we reduce the number of candidate values of x by a factor of 2, so that n lookups suffice to sift through all 2^n sorted items. You can use this algorithm to look up a number in the Los Angeles phone book, because the names are listed in lexicographic order.

But now suppose that you know someone's phone number, and you want to look up her *name*. Unless you are fortunate enough to have access to a reverse directory, this is a tedious procedure. Chances are you will need to check quite a few entries in the phone book before you come across her number.

In fact, if the N numbers are listed in a random order, you will need to look up $\frac{1}{2}N$ numbers before the probability is $P = \frac{1}{2}$ that you have found her number (and hence her name). What Grover discovered is that, if you have a quantum phone book, you can learn her name with high probability by consulting the phone book only about \sqrt{N} times.

This problem, too, can be formulated as an oracle or "black box" problem. In this case, the oracle is the phone book, or lookup table. We can input a name (a value of x) and the oracle outputs either 0, if $f(x) \neq a$, or 1, if f(x) = a. Our task is to find, as quickly as possible, the value of x with

$$f(x) = a. \tag{6.132}$$

Why is this problem important? You may have never tried to find in the phone book the name that matches a given number, but if it weren't so hard you might try it more often! More broadly, a rapid method for searching an unsorted database could be invoked to solve any problem in NP. Our oracle could be a subroutine that interrogates every potential "witness" y that could

potentially testify to certify a solution to the problem. For example, if we are confronted by a graph and need to know if it admits a Hamiltonian path, we could submit a path to the "oracle," and it could quickly answer whether the path is Hamiltonian or not. If we knew a fast way to query the oracle about all the possible paths, we would be able to find a Hamiltonian path efficiently (if one exists).

6.4.1 The oracle

So "oracle" could be shorthand for a subroutine that quickly evaluates a function to check a proposed solution to a decision problem, but let us continue to regard the oracle abstractly, as a black box. The oracle "knows" that of the 2^n possible strings of length n, one (the "marked" string or "solution" ω) is special. We submit a query x to the oracle, and it tells us whether $x = \omega$ or not. It returns, in other words, the value of a function $f_{\omega}(x)$, with

$$f_{\omega}(x) = 0, \qquad x \neq \omega,$$

$$f_{\omega}(x) = 1, \qquad x = \omega.$$
(6.133)

But furthermore, it is a *quantum* oracle, so it can respond to queries that are superpositions of strings. The oracle is a quantum black box that implements the unitary transformation

$$U_{f_{\omega}}: |x\rangle|y\rangle \to |x\rangle|y \oplus f_{\omega}(x)\rangle,$$
 (6.134)

where $|x\rangle$ is an *n*-qubit state, and $|y\rangle$ is a single-qubit state.

As we have previously seen in other contexts, we may choose the state of the single-qubit register to be $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, so that the oracle acts as

$$\boldsymbol{U}_{f_{\omega}} : |x\rangle \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle)$$

$$\rightarrow (-1)^{f_{\omega}(x)} |x\rangle \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle). \qquad (6.135)$$

We may now ignore the second register, and obtain

$$\boldsymbol{U}_{\omega}:|\boldsymbol{x}\rangle\to(-1)^{f_{\omega}(\boldsymbol{x})}|\boldsymbol{x}\rangle,\tag{6.136}$$

or

$$\boldsymbol{U}_{\omega} = 1 - 2|\omega\rangle\langle\omega|. \tag{6.137}$$

The oracle flips the sign of the state $|\omega\rangle$, but acts trivially on any state orthogonal to $|\omega\rangle$. This transformation has a simple geometrical interpretation. Acting on any vector in the 2ⁿ-dimensional Hilbert space, U_{ω} reflects the vector about the hyperplane orthogonal to $|\omega\rangle$ (it preserves the component in the hyperplane, and flips the component along $|\omega\rangle$).

We know that the oracle performs this reflection for some particular computational basis state $|\omega\rangle$, but we know nothing *a priori* about the value of the string ω . Our job is to determine ω , with high probability, consulting the oracle a minimal number of times.

6.4.2 The Grover iteration

As a first step, we prepare the state

$$|s\rangle = \frac{1}{\sqrt{N}} \left(\sum_{x=0}^{N-1} |x\rangle \right), \qquad (6.138)$$

The equally weighted superposition of all computational basis states – this can be done easily by applying the Hadamard transformation to each qubit of the initial state $|x = 0\rangle$. Although we do not know the value of ω , we do know that $|\omega\rangle$ is a computational basis state, so that

$$|\langle \omega | s \rangle| = \frac{1}{\sqrt{N}},\tag{6.139}$$

irrespective of the value of ω . Were we to measure the state $|s\rangle$ by projecting onto the computational basis, the probability that we would "find" the marked state $|\omega\rangle$ is only $\frac{1}{N}$. But following Grover, we can repeatedly iterate a transformation that enhances the probability amplitude of the unknown state $|\omega\rangle$ that we are seeking, while suppressing the amplitude of all of the undesirable states $|x \neq \omega\rangle$. We construct this Grover iteration by combining the unknown reflection U_{ω} performed by the oracle with a known reflection that we can perform ourselves. This known reflection is

$$\boldsymbol{U}_s = 2|s\rangle\langle s| - 1, \qquad (6.140)$$

which preserves $|s\rangle$, but flips the sign of any vector orthogonal to $|s\rangle$. Geometrically, acting on an arbitrary vector, it preserves the component along $|s\rangle$ and flips the component in the hyperplane orthogonal to $|s\rangle$.

6.4. QUANTUM DATABASE SEARCH

We'll return below to the issue of constructing a quantum circuit that implements U_s ; for now let's just assume that we can perform U_s efficiently.

One Grover iteration is the unitary transformation

$$\boldsymbol{R}_{\text{grov}} = \boldsymbol{U}_s \boldsymbol{U}_{\omega}, \qquad (6.141)$$

one oracle query followed by our reflection. Let's consider how \mathbf{R}_{grov} acts in the plane spanned by $|\omega\rangle$ and $|s\rangle$. This action is easiest to understand if we visualize it geometrically. Recall that

$$|\langle s|\omega\rangle| = \frac{1}{\sqrt{N}} \equiv \sin\theta, \qquad (6.142)$$

so that $|s\rangle$ is rotated by θ from the axis $|\omega^{\perp}\rangle$ normal to $|\omega\rangle$ in the plane. U_{ω} reflects a vector in the plane about the axis $|\omega^{\perp}\rangle$, and U_s reflects a vector about the axis $|s\rangle$. Together, the two reflections rotate the vector by 2θ :

The Grover iteration, then, is nothing but a rotation by 2θ in the plane determined by $|s\rangle$ and $|\omega\rangle$.

6.4.3 Finding 1 out of 4

Let's suppose, for example, that there are N = 4 items in the database, with one marked item. With classical queries, the marked item could be found in the 1st, 2nd, 3rd, or 4th query; on the average $2\frac{1}{2}$ queries will be needed before we are successful and four are needed in the worst case.⁸ But since $\sin \theta = \frac{1}{\sqrt{N}} = \frac{1}{2}$, we have $\theta = 30^{\circ}$ and $2\theta = 60^{\circ}$. After one Grover iteration, then, we rotate $|s\rangle$ to a 90° angle with $|\omega^{\perp}\rangle$; that is, it lines up with $|\omega\rangle$. When we measure by projecting onto the computational basis, we obtain the result $|\omega\rangle$ with certainty. Just one quantum query suffices to find the marked state, a notable improvement over the classical case.

⁸Of course, if we know there is one marked state, the 4th query is actually superfluous, so it might be more accurate to say that at most three queries are needed, and $2\frac{1}{4}$ queries are required on the average.

There is an alternative way to visualize the Grover iteration that is sometimes useful, as an "inversion about the average." If we expand a state $|\psi\rangle$ in the computational basis

$$|\psi\rangle = \sum_{x} a_x |x\rangle, \tag{6.143}$$

then its inner product with $|s\rangle = \frac{1}{\sqrt{N}} \sum_{x} |x\rangle$ is

$$\langle s|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{x} a_x = \sqrt{N} \langle a \rangle,$$
 (6.144)

where

$$\langle a \rangle = \frac{1}{N} \sum_{x} a_x, \tag{6.145}$$

is the mean of the amplitude. Then if we apply $U_s = 2|s\rangle\langle s| - 1$ to $|\psi\rangle$, we obtain

$$\boldsymbol{U}_{s}|\psi\rangle = \sum_{x} (2\langle a\rangle - a_{x})|x\rangle; \qquad (6.146)$$

the amplitudes are transformed as

$$\boldsymbol{U}_s: a_x - \langle a \rangle \to \langle a \rangle - a_x, \tag{6.147}$$

that is the coefficient of $|x\rangle$ is inverted about the mean value of the amplitude.

If we consider again the case N = 4, then in the state $|s\rangle$ each amplitude is $\frac{1}{2}$. One query of the oracle flips the sign of the amplitude of marked state, and so reduces the mean amplitude to $\frac{1}{4}$. Inverting about the mean then brings the amplitudes of all unmarked states from $\frac{1}{2}$ to zero, and raises the amplitude of the marked state from $-\frac{1}{2}$ to 1. So we recover our conclusion that one query suffices to find the marked state with certainty.

We can also easily see that one query is sufficient to find a marked state if there are N entries in the database, and exactly $\frac{1}{4}$ of them are marked. Then, as above, one query reduces the mean amplitude from $\frac{1}{\sqrt{N}}$ to $\frac{1}{2\sqrt{N}}$, and inversion about the mean then reduces the amplitude of each unmarked state to zero.

(When we make this comparison between the number of times we need to consult the oracle if the queries can be quantum rather than classical, it

6.4. QUANTUM DATABASE SEARCH

may be a bit unfair to say that only one query is needed in the quantum case. If the oracle is running a routine that computes a function, then some scratch space will be filled with garbage during the computation. We will need to erase the garbage by running the computation backwards in order to maintain quantum coherence. If the classical computation is irreversible there is no need to run the oracle backwards. In this sense, one query of the quantum oracle may be roughly equivalent, in terms of complexity, to two queries of a classical oracle.)

6.4.4 Finding 1 out of N

Let's return now to the case in which the database contains N items, and exactly one item is marked. Each Grover iteration rotates the quantum state in the plane determined by $|s\rangle$ and $|\omega\rangle$; after T iterations, the state is rotated by $\theta + 2T\theta$ from the $|\omega^{\perp}\rangle$ axis. To optimize the probability of finding the marked state when we finally perform the measurement, we will iterate until this angle is close to 90°, or

$$(2T+1)\theta \simeq \frac{\pi}{2} \Rightarrow 2T+1 \simeq \frac{\pi}{2\theta},$$
 (6.148)

we recall that $\sin \theta = \frac{1}{\sqrt{N}}$, or

$$\theta \simeq \frac{1}{\sqrt{N}},\tag{6.149}$$

for N large; if we choose

$$T = \frac{\pi}{4}\sqrt{N}(1 + O(N^{-1/2})), \qquad (6.150)$$

then the probability of obtaining the measurement result $|\omega\rangle$ will be

$$Prob(\omega) = \sin^2((2T+1)\theta) = 1 - O\left(\frac{1}{N}\right).$$
 (6.151)

We conclude that only about $\frac{\pi}{4}\sqrt{N}$ queries are needed to determine ω with high probability, a quadratic speedup relative to the classical result.
6.4.5 Multiple solutions

If there are r > 1 marked states, and r is known, we can modify the number of iterations so that the probability of finding one of the marked states is still very close to 1. The analysis is just as above, except that the oracle induces a reflection in the hyperplane orthogonal to the vector

$$\left|\tilde{\omega}\right\rangle = \frac{1}{\sqrt{r}} \left(\sum_{i=1}^{r} \left|\omega_{i}\right\rangle\right),\tag{6.152}$$

the equally weighted superposition of the marked computational basis states $|\omega_i\rangle$. Now

$$\langle s|\tilde{\omega}\rangle = \sqrt{\frac{r}{N}} \equiv \sin\theta,$$
 (6.153)

and a Grover iteration rotates a vector by 2θ in the plane spanned by $|s\rangle$ and $|\tilde{\omega}\rangle$; we again conclude that the state is close to $|\tilde{\omega}\rangle$ after a number of iterations

$$T \simeq \frac{\pi}{4\theta} = \frac{\pi}{4} \sqrt{\frac{N}{r}}.$$
(6.154)

If we then measure by projecting onto the computational basis, we will find one of the marked states (each occurring equiprobably) with probability close to one. (As the number of solutions increases, the time needed to find one of them declines like $r^{-1/2}$, as opposed to r^{-1} in the classical case.)

Note that if we continue to perform further Grover iterations, the vector continues to rotate, and so the probability of finding a marked state (when we finally measure) begins to decline. The Grover algorithm is like baking a soufflé – if we leave it in the oven for too long, it starts to fall. Therefore, if we don't know anything about the number of marked states, we might fail to find one of them. For example, $T \sim \frac{\pi}{4}\sqrt{N}$ iterations is optimal for r = 1, but for r = 4, the probability of finding a marked state after this many iterations is quite close to zero.

But even if we don't know r a priori, we can still find a solution with a quadratic speed up over classical algorithms (for $r \ll N$). For example, we might choose the number of iterations to be random in the range 0 to $\frac{\pi}{4}\sqrt{N}$. Then the expected probability of finding a marked state is close to 1/2 for each r, so we are unlikely to fail to find a marked state after several repetitions. And each time we measure, we can submit the state we find to the oracle as a classical query to confirm whether that state is really marked.

In particular, if we don't find a solution after several attempts, there probably is no solution. Hence with high probability we can correctly answer the yes/no question, "Is there a marked state?" Therefore, we can adopt the Grover algorithm to solve any NP problem, where the oracle checks a proposed solution, with a quadratic speedup over a classical exhaustive search.

6.4.6 Implementing the reflection

To perform a Grover iteration, we need (aside from the oracle query) a unitary transformation

$$\boldsymbol{U}_s = 2|s\rangle\langle s| - 1, \qquad (6.155)$$

that reflects a vector about the axis defined by the vector $|s\rangle$. How do we build this transformation efficiently from quantum gates? Since $|s\rangle = \mathbf{H}^{(n)}|0\rangle$, where $\mathbf{H}^{(n)}$ is the bitwise Hadamard transformation, we may write

$$\boldsymbol{U}_{s} = \boldsymbol{H}^{(n)}(2|0\rangle\langle 0| - 1)\boldsymbol{H}^{(n)}, \qquad (6.156)$$

so it will suffice to construct a reflection about the axis $|0\rangle$. We can easily build this reflection from an *n*-bit Toffoli gate $\theta^{(n)}$.

Recall that

$$\boldsymbol{H}\boldsymbol{\sigma}_{x}\boldsymbol{H} = \boldsymbol{\sigma}_{z}; \tag{6.157}$$

a bit flip in the Hadamard rotated basis is equivalent to a flip of the relative phase of $|0\rangle$ and $|1\rangle$. Therefore:



after conjugating the last bit by $\boldsymbol{H}, \theta^{(n)}$ becomes controlled⁽ⁿ⁻¹⁾- $\boldsymbol{\sigma}_z$, which flips the phase of $|11...|1\rangle$ and acts trivially on all other computational basis states. Conjugating by NOT⁽ⁿ⁾, we obtain \boldsymbol{U}_s , aside from an irrelevant overall minus sign.

You will show in an exercise that the *n*-bit Toffoli gate $\theta^{(n)}$ can be constructed from 2n - 5 3-bit Toffoli gates $\theta^{(3)}$ (if sufficient scratch space is available). Therefore, the circuit that constructs U_s has a size *linear* in $n = \log N$. Grover's database search (assuming the oracle answers a query instantaneously) takes a time of order $\sqrt{N} \log N$. If we regard the oracle as a subroutine that performs a function evaluation in polylog time, then the search takes time of order $\sqrt{N} \log N$).

6.5 The Grover Algorithm Is Optimal

Grover's quadratic quantum speedup of the database search is already interesting and potentially important, but surely with more cleverness we can do better, can't we? No, it turns out that we can't. Grover's algorithm provides the fastest possible quantum search of an unsorted database, if "time" is measured according to the number of queries of the oracle.

Considering the case of a single marked state $|\omega\rangle$, let $U(\omega, T)$ denote a quantum circuit that calls the oracle T times. We place no restriction on the circuit aside from specifying the number of queries; in particular, we place no limit on the number of quantum gates. This circuit is applied to an initial

state $|\psi(0)\rangle$, producing a final state

$$|\psi_{\omega}(t)\rangle = \boldsymbol{U}(\omega, T)|\psi(0)\rangle. \tag{6.158}$$

Now we are to perform a measurement designed to distinguish among the N possible values of ω . If we are to be able to perfectly distinguish among the possible values, the states $|\psi_{\omega}(t)\rangle$ must all be mutually orthogonal, and if we are to distinguish correctly with high probability, they must be nearly orthogonal.

Now, if the states $\{|\psi_{\omega}\rangle$ are an orthonormal basis, then, for any fixed normalized vector $|\varphi\rangle$,

$$\sum_{\omega=0}^{N-1} \| |\psi_{\omega}\rangle - |\varphi\rangle \|^2 \ge 2N - 2\sqrt{N}.$$
(6.159)

(The sum is minimized if $|\varphi\rangle$ is the equally weighted superposition of all the basis elements, $|\varphi\rangle = \frac{1}{\sqrt{N}} \sum_{\omega} |\psi_{\omega}\rangle$, as you can show by invoking a Lagrange multiplier to perform the constrained extremization.) Our strategy will be to choose the state $|\varphi\rangle$ suitably so that we can use this inequality to learn something about the number T of oracle calls.

Our circuit with T queries builds a unitary transformation

$$\boldsymbol{U}(\omega,T) = \boldsymbol{U}_{\omega}\boldsymbol{U}_{T}\boldsymbol{U}_{\omega}\boldsymbol{U}_{T-1}\dots\boldsymbol{U}_{\omega}\boldsymbol{U}_{1}, \qquad (6.160)$$

where U_{ω} is the oracle transformation, and the U_t 's are arbitrary non-oracle transformations. For our state $|\varphi(T)\rangle$ we will choose the result of applying $U(\omega, T)$ to $|\psi(0)\rangle$, except with each U_{ω} replaced by 1; that is, the same circuit, but with all queries submitted to the "empty oracle." Hence,

$$|\varphi(T)\rangle = \boldsymbol{U}_T \boldsymbol{U}_{T-1} \dots \boldsymbol{U}_2 \boldsymbol{U}_1 |\psi(0)\rangle, \qquad (6.161)$$

while

$$|\psi_{\omega}(T)\rangle = \boldsymbol{U}_{\omega}\boldsymbol{U}_{T}\boldsymbol{U}_{\omega}\boldsymbol{U}_{T-1}\dots\boldsymbol{U}_{\omega}\boldsymbol{U}_{1}|\psi(0)\rangle.$$
(6.162)

To compare $|\varphi(T)\rangle$ and $|\psi_{\omega}(T)\rangle$, we appeal to our previous analysis of the effect of errors on the accuracy of a circuit, regarding the ω oracle as an "erroneous" implementation of the empty oracle. The error vector in the *t*-th step (cf. eq. (6.63)) is

$$||E(\omega,t)\rangle|| = ||(\mathbf{U}_{\omega} - 1)|\varphi(t)\rangle||$$

= 2|\langle \overline \vert \v

since $U_{\omega} = 1 - 2|\omega\rangle\langle\omega|$. After T queries we have (cf. eq. (6.66))

$$\| |\psi_{\omega}(T)\rangle - |\varphi(T)\rangle \| \le 2\sum_{t=1}^{T} |\langle \omega|\varphi(t)\rangle|.$$
(6.164)

From the identity

$$\left(\sum_{t=1}^{T} c_t\right)^2 + \frac{1}{2} \sum_{s,t=1}^{T} (c_s - c_t)^2$$
$$= \sum_{s,t=1}^{T} \left(c_t c_s + \frac{1}{2} c_s^2 - c_t c_s + \frac{1}{2} c_s^2\right) = T \sum_{t=1}^{T} c_t^2, \qquad (6.165)$$

we obtain the inequality

$$\left(\sum_{t=1}^{T} c_t\right)^2 \le T \sum_{t=1}^{T} c_t^2, \tag{6.166}$$

which applied to eq. (6.164) yields

$$\| |\psi_{\omega}(T)\rangle - |\varphi(T)\rangle \|^{2} \leq 4T \left(\sum_{t=1}^{T} |\langle \omega | \varphi(t) \rangle|^{2}\right).$$
(6.167)

Summing over ω we find

$$\sum_{\omega} \| |\psi_{\omega}(T)\rangle - |\varphi(T)\rangle \|^{2} \leq 4T \sum_{t=1}^{T} \langle \varphi(t) |\varphi(t)\rangle = 4T^{2}.$$
(6.168)

Invoking eq. (6.159) we conclude that

$$4T^2 \ge 2N - 2\sqrt{N}, \tag{6.169}$$

if the states $|\psi_{\omega}(T)\rangle$ are mutually orthogonal. We have, therefore, found that any quantum algorithm that can distinguish all the possible values of the marked state must query the oracle T times where

$$T \ge \sqrt{\frac{N}{2}},\tag{6.170}$$

(ignoring the small correction as $N \to \infty$). Grover's algorithm finds ω in $\frac{\pi}{4}\sqrt{N}$ queries, which exceeds this bound by only about 11%. In fact, it is

possible to refine the argument to improve the bound to $T \geq \frac{\pi}{4}\sqrt{N}(1-\varepsilon)$, which is asymptotically saturated by the Grover algorithm.⁹ Furthermore, we can show that Grover's circuit attains the optimal success probability in $T \leq \frac{\pi}{4}\sqrt{N}$ queries.

One feels a twinge of disappointment (as well as a surge of admiration for Grover) at the realization that the database search algorithm cannot be improved. What are the implications for quantum complexity?

For many optimization problems in the NP class, there is no better method known than exhaustive search of all the possible solutions. By exploiting quantum parallelism, we can achieve a quadratic speedup of exhaustive search. Now we have learned that the quadratic speedup is the best possible if we rely on the power of sheer quantum parallelism, if we don't design our quantum algorithm to exploit the specific structure of the problem we wish to solve. Still, we might do better if we are sufficiently clever.

The optimality of the Grover algorithm might be construed as evidence that $BQP \not\supseteq NP$. At least, if it turns out that $NP \subseteq BQP$ and $P \neq NP$, then the NP problems must share a deeply hidden structure (for which there is currently no evidence) that is well-matched to the peculiar capabilities of quantum circuits.

Even the quadratic speedup may prove useful for a variety of NP-complete optimization problems. But a quadratic speedup, unlike an exponential one, does not really move the frontier between solvability and intractability. Quantum computers may someday outperform classical computers in performing exhaustive search, but only if the clock speed of quantum devices does not lag too far behind that of their classical counterparts.

6.6 Generalized Search and Structured Search

In the Grover iteration, we perform the transformation $U_s = 2|s\rangle\langle s| - 1$, the reflection in the axis defined by $|s\rangle = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle$. Why this axis? The advantage of the state $|s\rangle$ is that it has the same overlap with each and every computational basis state. Therefore, the overlap of any marked state $|\omega\rangle$ with $|s\rangle$ is guaranteed to be $|\langle \omega | s \rangle| = 1/\sqrt{N}$. Hence, if we know the number of marked states, we can determine how many iterations are required to find a marked state with high probability – the number of iterations needed does

⁹C. Zalka, "Grover's Quantum Searching Algorithm is Optimal," quant-ph/9711070.

not depend on which states are marked.

But of course, we could choose to reflect about a different axis. If we can build the unitary U (with reasonable efficiency) then we can construct

$$\boldsymbol{U}(2|0\rangle\langle 0|-1)\boldsymbol{U}^{\dagger} = 2\boldsymbol{U}|0\rangle\langle 0|\boldsymbol{U}^{\dagger}-1, \qquad (6.171)$$

which reflects in the axis $U|0\rangle$.

Suppose that

$$|\langle \omega | \boldsymbol{U} | 0 \rangle| = \sin \theta, \qquad (6.172)$$

where $|\omega\rangle$ is the marked state. Then if we replace U_s in the Grover iteration by the reflection eq. (6.171), one iteration performs a rotation by 2θ in the plane determined by $|\omega\rangle$ and $U|0\rangle$ (by the same argument we used for U_s). Thus, after T iterations, with $(2T + I)\theta \cong \pi/2$, a measurement in the computational basis will find $|\omega\rangle$ with high probability. Therefore, we can still search a database if we replace $H^{(n)}$ by U in Grover's quantum circuit, as long as $U|0\rangle$ is not orthogonal to the marked state.¹⁰ But if we have no *a priori* information about which state is marked, then $H^{(n)}$ is the best choice, not only because $|s\rangle$ has a known overlap with each marked state, but also because it has the largest *average* overlap with all the possible marked states.

But sometimes when we are searching a database, we *do* have some information about where to look, and in that case, the generalized search strategy described above may prove useful.¹¹

As an example of a problem with some auxiliary structure, suppose that f(x, y) is a one-bit-valued function of the two *n*-bit strings x and y, and we are to find the unique solution to f(x, y) = 1. With Grover's algorithm, we can search through the N^2 possible values $(N = 2^n)$ of (x, y) and find the solution (x_0, y_0) with high probability after $\frac{\pi}{4}N$ iterations, a quadratic speedup with respect to classical search.

But further suppose that g(x) is a function of x only, and that it is known that g(x) = 1 for exactly M values of x, where $1 \ll M \ll N$. And furthermore, it is known that $g(x_0) = 1$. Therefore, we can use g to help us find the solution (x_0, y_0) .

¹⁰L.K. Grover "Quantum Computers Can Search Rapidly By Using Almost Any Transformation," quant-ph/9712011.

¹¹E. Farhi and S. Gutmann, "Quantum-Mechanical Square Root Speedup in a Structured Search Problem," quant-ph/9711035; L.K. Grover, "Quantum Search On Structured Problems," quant-ph/9802035.

Now we have two oracles to consult, one that returns the value of f(x, y), and the other returning the value of g(x). Our task is to find (x_0, y_0) with a minimal number of queries.

Classically, we need of order NM queries to find the solution with reasonable probability. We first evaluate g(x) for each x; then we restrict our search for a solution to f(x, y) = 1 to only those M values of x such that g(x) = 1. It is natural to wonder whether there is a way to perform a quantum search in a time of order the square root of the classical time. Exhaustive search that queries only the f oracle requires time $N \gg \sqrt{NM}$, and so does not do the job. We need to revise our method of quantum search to take advantage of the structure provided by g.

A better method is to first apply Grover's algorithm to g(x). In about $\frac{\pi}{4}\sqrt{\frac{N}{M}}$ iterations, we prepare a state that is close to the equally weighted superposition of the M solutions to g(x) = 1. In particular, the state $|x_0\rangle$ appears with amplitude $\frac{1}{\sqrt{M}}$. Then we apply Grover's algorithm to f(x, y) with x fixed. In about $\frac{\pi}{4}\sqrt{N}$ iterations, the state $|x_0\rangle|s\rangle$ evolves to a state quite close to $|x_0\rangle|y_0\rangle$. Therefore $|x_0, y_0\rangle$ appears with amplitude $\frac{1}{\sqrt{M}}$.

The unitary transformation we have constructed so far, in about $\frac{\pi}{4}\sqrt{N}$ queries, can be regarded as the transformation \boldsymbol{U} that defines a generalized search. Furthermore, we know that

$$\langle x_0, y_0 | U | 0, 0 \rangle \cong \frac{1}{\sqrt{M}}.$$
 (6.173)

Therefore, if we iterate the generalized search about $\frac{\pi}{4}\sqrt{M}$ times, we will have prepared a state that is quite close to $|x_0, y_0\rangle$. With altogether about

$$\left(\frac{\pi}{4}\right)^2 \sqrt{NM},\tag{6.174}$$

queries, then, we can find the solution with high probability. This is indeed a quadratic speedup with respect to the classical search.

6.7 Some Problems Admit No Speedup

The example of structured search illustrates that quadratic quantum speedups over classical algorithms can be attained for a variety of problems, not just for an exhaustive search of a structureless database. One might even dare to hope that quantum parallelism enables us to significantly speedup any classical algorithm. This hope will now be dashed – for many problems, no quantum speedup is possible.

We continue to consider problems with a quantum black box, an oracle, that computes a function f taking n bits to 1. But we will modify our notation a little. The function f can be represented as a string of $N = 2^n$ bits

$$X = X_{N-1} X_{N-2} \dots X_1 X_0, \tag{6.175}$$

where X_i denotes f(i). Our problem is to evaluate some one-bit-valued function of X, that is, to answer a yes/no question about the properties of the oracle. What we will show is that for some functions of X, we can't evaluate the function with low error probability using a quantum algorithm, unless the algorithm queries the oracle as many times (or nearly as many times) as required with a classical algorithm.¹²

The key idea is that any Boolean function of the X_i 's can be represented as a polynomial in the X_i 's. Furthermore, the probability distribution for a quantum measurement can be expressed as a polynomial in X, where the degree of the polynomial is 2T, if the measurement follows T queries of the oracle. The issue, then, is whether a polynomial of degree 2T can provide a reasonable approximation to the Boolean function of interest.

The action of the oracle can be represented as

$$\boldsymbol{U}_O: |i, y; z\rangle \to |i, y \oplus X_i; z\rangle, \qquad (6.176)$$

where *i* takes values in $\{0, 1, ..., N-1\}, y \in \{0, 1\}$, and *z* denotes the state of auxiliary qubits not acted upon by the oracle. Therefore, in each 2×2 block spanned by $|i, 0, z\rangle$ and $|i, 1, z\rangle, U_O$ is the 2×2 matrix

$$\begin{pmatrix}
1 - X_i & X_i \\
X_i & 1 - X_i
\end{pmatrix}.$$
(6.177)

Quantum gates other than oracle queries have no dependence on X. Therefore after a circuit with T queries acts on any initial state, the resulting state $|\psi\rangle$ has amplitudes that are (at most) Tth-degree polynomials in X. If we perform a POVM on $|\psi\rangle$, then the probability $\langle \psi | \boldsymbol{F} | \psi \rangle$ of the outcome associated with the positive operator \boldsymbol{F} can be expressed as a polynomial in X of degree at most 2T.

¹²E. Farhi, et al., quant-ph/9802045; R. Beals, et al., quant-ph/9802049.

Now any Boolean function of the X_i 's can be expressed (uniquely) as a polynomial of degree $\leq N$ in the X_i 's. For example, consider the OR function of the $N X_i$'s; it is

$$OR(X) = 1 - (1 - X_0)(1 - X_1) \cdots (1 - X_{N-1}), \qquad (6.178)$$

a polynomial of degree N.

Suppose that we would like our quantum circuit to evaluate the OR function *with certainty*. Then we must be able to perform a measurement with two outcomes, 0 and 1, where

$$Prob(0) = 1 - OR(X),$$

$$Prob(1) = OR(X).$$
(6.179)

But these expressions are polynomials of degree N, which can arise only if the circuit queries the oracle at least T times, where

$$T \ge \frac{N}{2}.\tag{6.180}$$

We conclude that no quantum circuit with fewer than N/2 oracle calls can compute OR exactly. In fact, for this function (or any function that takes the value 0 for just one of its N possible arguments), there is a stronger conclusion (exercise): we require $T \ge N$ to evaluate OR with certainty.

On the other hand, evaluating the OR function (answering the yes/no question, "Is there a marked state?") is just what the Grover algorithm can achieve in a number of queries of order \sqrt{N} . Thus, while the conclusion is correct that N queries are needed to evaluate OR with certainty, this result is a bit misleading. We can evaluate OR probabilistically with far fewer queries. Apparently, the Grover algorithm can construct a polynomial in X that, though only of degree $O(\sqrt{N})$, provides a very adequate approximation to the N-th degree polynomial OR(X).

But OR, which takes the value 1 for every value of X except $X = \vec{0}$, is a very simple Boolean function. We should consider other functions that might pose a more serious challenge for the quantum computer.

One that comes to mind is the PARITY function: PARITY(X) takes the value 0 if the string X contains an even number of 1's, and the value 1 if the string contains an odd number of 1's. Obviously, a classical algorithm must query the oracle N times to determine the parity. How much better

can we do by submitting quantum queries? In fact, we can't do much better at all – at least N/2 quantum queries are needed to find the correct value of PARITY(X), with probability of success greater than $\frac{1}{2} + \delta$.

In discussing PARITY it is convenient to use new variables

$$\ddot{X}_i = 1 - 2X_i, \tag{6.181}$$

that take values ± 1 , so that

$$PARITY(\tilde{X}) = \prod_{i=0}^{N-1} \tilde{X}_i, \qquad (6.182)$$

also takes values ± 1 . Now, after we execute a quantum circuit with altogether T queries of the oracle, we are to perform a POVM with two possible outcomes \mathbf{F}_{even} and \mathbf{F}_{odd} ; the outcome will be our estimate of PARITY(\tilde{X}). As we have already noted, the probability of obtaining the outcome even (say) can be expressed as a polynomial $P_{\text{even}}^{(2T)}$ of degree (at most) 2T in \tilde{X} ,

$$\langle \boldsymbol{F}_{\text{even}} \rangle = P_{\text{even}}^{(2T)}(\tilde{X}).$$
 (6.183)

How often is our guess correct? Consider the sum

$$\sum_{\{\tilde{X}\}} P_{\text{even}}^{(2T)}(\tilde{X}) \cdot \text{PARITY}(\tilde{X})$$
$$= \sum_{\{\tilde{X}\}} P_{\text{even}}^{(2T)}(\tilde{X}) \prod_{i=0}^{N-1} \tilde{X}_i.$$
(6.184)

Since each term in the polynomial $P_{\text{even}}^{(2T)}(\tilde{X})$ contains at most 2T of the \tilde{X}_i 's, we can invoke the identity

$$\sum_{\tilde{X}_i \in \{0,1\}} \tilde{X}_i = 0, \tag{6.185}$$

to see that the sum in eq. (6.184) must vanish if N > 2T. We conclude that

$$\sum_{\text{par}(\tilde{X})=1} P_{\text{even}}^{(2T)}(\tilde{X}) = \sum_{\text{par}(\tilde{X})=-1} P_{\text{even}}^{(2T)}(\tilde{X});$$
(6.186)

hence, for T < N/2, we are just as likely to guess "even" when the actual $PARITY(\tilde{X})$ is odd as when it is even (on average). Our quantum algorithm

fails to tell us anything about the value of PARITY(\tilde{X}); that is, averaged over the (a priori equally likely) possible values of X_i , we are just as likely to be right as wrong.

We can also show, by exhibiting an explicit algorithm (exercise), that N/2 queries (assuming N even) are *sufficient* to determine PARITY (either probabilistically or deterministically.) In a sense, then, we can achieve a factor of 2 speedup compared to classical queries. But that is the best we can do.

6.8 Distributed database search

We will find it instructive to view the quantum database search algorithm from a fresh perspective. We imagine that two parties, Alice and Bob, need to arrange to meet on a mutually agreeable day. Alice has a calendar that lists $N = 2^n$ days, with each day marked by either a 0, if she is unavailable that day, or a 1, if she is available. Bob has a similar calendar. Their task is to find a day when they will both be available.

Alice and Bob both have quantum computers, but they are very far apart from one another. (Alice is on earth, and Bob has traveled to the Andromeda galaxy). Therefore, it is very expensive for them to communicate. They urgently need to arrange their date, but they must economize on the amount of information that they send back and forth.

Even if there exists a day when both are available, it might not be easy to find it. If Alice and Bob communicate by sending classical bits back and forth, then in the worst case they will need to exchange of order $N = 2^n$ calendar entries to have a reasonable chance of successfully arranging their date.. We will ask: can they do better by exchanging qubits instead?¹³ (The quantum

¹³In an earlier version of these notes, I proposed a different scenario, in which Alice and Bob had nearly identical tables, but with a single mismatched entry; their task was to find the location of the mismatched bit. However, that example was poorly chosen, because the task can be accomplished with only $\log N$ bits of classical communication. (Thanks to Richard Cleve for pointing out this blunder.) We want Alice to learn the address (a binary string of length n) of the one entry where her table differs from Bob's. So Bob computes the parity of the N/2 entries in his table with a label that takes the value 0 in its least significant bit, and he sends that one parity bit to Alice. Alice compares to the parity of the same entries in her table, and she infers one bit (the least significant bit) of the address of the mismatched entry. Then they do the same for each of the remaining n-1 bits, until Alice knows the complete address of the "error". Altogether just n bits

information highway from earth to Andromeda was carefully designed and constructed, so it does not cost much more to send qubits instead of bits.)

To someone familiar with the basics of quantum information theory, this sounds like a foolish question. Holevo's theorem told us once and for all that a single qubit can convey no more than one bit of classical information. On further reflection, though, we see that Holevo's theorem does not really settle the issue. While it bounds the mutual information of a state preparation with a measurement outcome, it does not assure us (at least not directly) that Alice and Bob need to exchange as many qubits as bits to compare their calendars. Even so, it comes as a refreshing surprise¹⁴ to learn that Alice and Bob can do the job by exchanging $O(\sqrt{N} \log N)$ qubits, as compared to O(N) classical bits.

To achieve this Alice and Bob must work in concert, implementing a distributed version of the database search. Alice has access to an oracle (her calendar) that computes a function $f_A(x)$, and Bob has an oracle (his calendar) that computes $f_B(x)$. Together, they can query the oracle

$$f_{AB}(x) = f_A(x) \wedge f_B(x)$$
 (6.187)

Either one of them can implement the reflection U_s , so they can perform a complete Grover iteration, and can carry out exhaustive search for a suitable day x such that $f_{AB}(x) = 1$ (when Alice and Bob are both available). If a mutually agreeable day really exists, they will succeed in finding it after of order \sqrt{N} queries.

How do Alice and Bob query f_{AB} ? We'll describe how they do it acting on any one of the computational basis states $|x\rangle$. First Alice performs

$$|x\rangle|0\rangle \to |x\rangle|f_A(x)\rangle,$$
 (6.188)

and then she sends the n + 1 qubits to Bob. Bob performs

$$|x\rangle|f_A(x)\rangle \to (-1)^{f_A(x)\wedge f_B(x)}|x\rangle|f_A(x)\rangle.$$
(6.189)

This transformation is evidently unitary, and you can easily verify that Bob can implement it by querying his oracle. Now the phase multiplying $|x\rangle$ is $(-1)^{f_{AB}(x)}$ as desired, but $|f_A(x)\rangle$ remains stored in the other register, which

are sent (and all from Bob to Alice).

¹⁴H. Burhman, *et al.*, "Quantum vs. Classical Communication and Computation," quant-ph/9802040.

would spoil the coherence of a superposition of x values. Bob cannot erase that register, but Alice can. So Bob sends the n + 1 qubits back to Alice, and she consults her oracle once more to perform

$$(-1)^{f_A(x)\wedge f_B(x)}|x\rangle|f_A(x)\rangle \to (-1)^{f_A(x)\wedge f_B(x)}|x\rangle|0\rangle.$$
(6.190)

By exchanging 2(n+1) qubits, the have accomplished one query of the f_{AB} oracle, and so can execute one Grover iteration.

Suppose, for example, that Alice and Bob know that there is only one mutually agreeable date, but they have no *a priori* information about which date it is. After about $\frac{\pi}{4}\sqrt{N}$ iterations, requiring altogether

$$Q \cong \frac{\pi}{4}\sqrt{N} \cdot 2(\log N + 1), \tag{6.191}$$

qubit exchanges, Alice measures, obtaining the good date with probability quite close to 1.

Thus, at least in this special context, exchanging $O(\sqrt{N} \log N)$ qubits is as good as exchanging O(N) classical bits. Apparently, we have to be cautious in interpreting the Holevo bound, which ostensibly tells us that a qubit has no more information-carrying capacity than a bit!

If Alice and Bob don't know in advance how many good dates there are, they can still perform the Grover search (as we noted in §6.4.5), and will find a solution with reasonable probability. With $2 \cdot \log N$ bits of classical communication, they can verify whether the date that they found is really mutually agreeable.

6.8.1 Quantum communication complexity

More generally, we may imagine that several parties each possess an *n*-bit input, and they are to evaluate a function of all the inputs, with one party eventually learning the value of the function. What is the minimum amount of communication needed to compute the function (either deterministically or probabilistically)? The well-studied branch of classical complexity theory that addresses this question is called *communication complexity*. What we established above is a quadratic separation between quantum and classical communication complexity, for a particular class of two-party functions. Aside from replacing the exchange of classical bits by the exchange of qubits, there are other interesting ways to generalize classical communication complexity. One is to suppose that the parties share some preexisting entangled state (either Bell pairs or multipartite entanglement), and that they may exploit that entanglement along with classical communication to perform the function evaluation. Again, it is not immediately clear that the shared entanglement will make things any easier, since entanglement alone doesn't permit the parties to exchange classical messages. But it turns out that the entanglement *does* help, at least a little bit.¹⁵

The analysis of communication complexity is a popular past time among complexity theorists, but this discipline does not yet seem to have assumed a prominent position in practical communications engineering. Perhaps this is surprising, considering the importance of efficiently distributing the computational load in parallelized computing, which has become commonplace. Furthermore, it seems that nearly all communication in real life can be regarded as a form of remote computation. I don't really need to receive all the bits that reach me over the telephone line, especially since I will probably retain only a few bits of information pertaining to the call tomorrow (the movie we decided to go to). As a less prosaic example, we on earth may need to communicate with a robot in deep space, to instruct it whether to enter and orbit around a distant star system. Since bandwidth is extremely limited, we would like it to compute the correct answer to the Yes/No question "Enter orbit?" with minimal exchange of information between earth and robot.

Perhaps a future civilization will exploit the known quadratic separation between classical and quantum communication complexity, by exchanging qubits rather than bits with its flotilla of spacecraft. And perhaps an exponential separation will be found, at least in certain contexts, which would significantly boost the incentive to develop the required quantum communications technology.

6.9 Periodicity

So far, the one case for which we have found an exponential separation between the speed of a quantum algorithm and the speed of the corresponding

¹⁵R. Cleve, et al., "Quantum Entanglement and the Communication Complexity of the Inner Product Function," quant-ph/9708019; W. van Dam, et al., "Multiparty Quantum Communication Complexity," quant-ph/9710054.

classical algorithm is the case of Simon's problem. Simon's algorithm exploits quantum parallelism to speed up the search for the period of a function. Its success encourages us to seek other quantum algorithms designed for other kinds of period finding.

Simon studied periodic functions taking values in $(Z_2)^n$. For that purpose the *n*-bit Hadamard transform $\mathbf{H}^{(n)}$ was a powerful tool. If we wish instead to study periodic functions taking values in Z_{2^n} , the (discrete) Fourier transform will be a tool of comparable power.

The moral of Simon's problem is that, while finding needles in a haystack may be difficult, finding *periodically* spaced needles in a haystack can be far easier. For example, if we scatter a photon off of a periodic array of needles, the photon is likely to be scattered in one of a set of preferred directions, where the Bragg scattering condition is satisfied. These preferred directions depend on the spacing between the needles, so by scattering just one photon, we can already collect some useful information about the spacing. We should further explore the implications of this metaphor for the construction of efficient quantum algorithms.

So imagine a quantum oracle that computes a function

$$f: \{0,1\}^n \to \{0,1\}^m,$$
 (6.192)

that has an unknown period r, where r is a positive integer satisfying

$$1 \ll r \ll 2^n. \tag{6.193}$$

That is,

$$f(x) = f(x + mr),$$
 (6.194)

where *m* is any integer such that *x* and x + mr lie in $\{0, 1, 2, ..., 2^n - 1\}$. We are to find the period *r*. Classically, this problem is *hard*. If *r* is, say, of order $2^{n/2}$, we will need to query the oracle of order $2^{n/4}$ times before we are likely to find two values of *x* that are mapped to the same value of f(x), and hence learn something about *r*. But we will see that there is a quantum algorithm that finds *r* in time poly (n).

Even if we know how to compute efficiently the function f(x), it may be a hard problem to determine its period. Our quantum algorithm can be applied to finding, in poly(n) time, the period of any function that we can compute in poly(n) time. Efficient period finding allows us to efficiently solve a variety of (apparently) hard problems, such as factoring an integer, or evaluating a discrete logarithm.

The key idea underlying quantum period finding is that the Fourier transform can be evaluated by an efficient quantum circuit (as discovered by Peter Shor). The quantum Fourier transform (QFT) exploits the power of quantum parallelism to achieve an exponential speedup of the well-known (classical) fast Fourier transform (FFT). Since the FFT has such a wide variety of applications, perhaps the QFT will also come into widespread use someday.

6.9.1 Finding the period

The QFT is the unitary transformation that acts on the computational basis according to

$$QFT:|x\rangle \to \frac{1}{\sqrt{N}} \sum_{y=0}^{N-1} e^{2\pi i x y/N} |y\rangle, \qquad (6.195)$$

where $N = 2^n$. For now let's suppose that we can perform the QFT efficiently, and see how it enables us to extract the period of f(x).

Emulating Simon's algorithm, we first query the oracle with the input $\frac{1}{\sqrt{N}}\sum_{x} |x\rangle$ (easily prepared by applying $\boldsymbol{H}^{(n)}$ to $|0\rangle$), and so prepare the state

$$\frac{1}{\sqrt{N}}\sum_{x=0}^{N-1}|x\rangle|f(x)\rangle.$$
(6.196)

Then we measure the output register, obtaining the result $|f(x_0)\rangle$ for some $0 \le x_0 < r$. This measurement prepares in the input register the coherent superposition of the A values of x that are mapped to $f(x_0)$:

$$\frac{1}{\sqrt{A}} \sum_{j=0}^{A-1} |x_0 + jr\rangle, \tag{6.197}$$

where

$$N - r \le x_0 + (A - 1)r < N, \tag{6.198}$$

or

$$A - 1 < \frac{N}{r} < A + 1. \tag{6.199}$$

6.9. PERIODICITY

Actually, the measurement of the output register is unnecessary. If it is omitted, the state of the input register is an incoherent superposition (summed over $x_0 \in \{0, 1, \ldots r - 1\}$) of states of the form eq. (6.197). The rest of the algorithm works just as well acting on this initial state.

Now our task is to extract the value of r from the state eq. (6.197) that we have prepared. Were we to measure the input register by projecting onto the computational basis at this point, we would learn nothing about r. Instead (cf. Simon's algorithm), we should Fourier transform first and *then* measure.

By applying the QFT to the state eq. (6.197) we obtain

$$\frac{1}{\sqrt{NA}} \sum_{y=0}^{N-1} e^{2\pi i x_0 y} \sum_{j=0}^{A-1} e^{2\pi i j r y/N} |y\rangle.$$
(6.200)

If we now measure in the computational basis, the probability of obtaining the outcome y is

$$\operatorname{Prob}(y) = \frac{A}{N} \left| \frac{1}{A} \sum_{j=0}^{A-1} e^{2\pi i j r y/N} \right|^2.$$
(6.201)

This distribution strongly favors values of y such that yr/N is close to an integer. For example, if N/r happened to be an integer (and therefore equal to A), we would have

$$\operatorname{Prob}(y) = \frac{1}{r} \left| \frac{1}{A} \sum_{j=0}^{A-1} e^{2\pi i j y/A} \right| = \begin{cases} \frac{1}{r} & y = A \cdot (\operatorname{integer}) \\ 0 & \operatorname{otherwise.} \end{cases}$$
(6.202)

More generally, we may sum the geometric series

$$\sum_{j=0}^{A-1} e^{i\theta j} = \frac{e^{iA\theta} - 1}{e^{i\theta} - 1},$$
(6.203)

where

$$\theta_y = \frac{2\pi yr \pmod{N}}{N}.$$
(6.204)

There are precisely r values of y in $\{0, 1, \dots, N-1\}$ that satisfy

$$-\frac{r}{2} \le yr \pmod{N} \le \frac{r}{2}.$$
(6.205)

(To see this, imagine marking the multiples of r and N on a number line ranging from 0 to rN - 1. For each multiple of N, there is a multiple of r no more than distance r/2 away.) For each of these values, the corresponding θ_y satisfies.

$$-\pi \frac{r}{N} \le \theta_y \le \pi \frac{r}{N}.$$
(6.206)

Now, since $A - 1 < \frac{N}{r}$, for these values of θ_y all of the terms in the sum over j in eq. (6.203) lie in the same half-plane, so that the terms interfere constructively and the sum is substantial.

We know that

$$|1 - e^{i\theta}| \le |\theta|, \tag{6.207}$$

because the straight-line distance from the origin is less than the arc length along the circle, and for $A|\theta| \leq \pi$, we know that

$$|1 - e^{iA\theta}| \ge \frac{2A|\theta|}{\pi},\tag{6.208}$$

because we can see (either graphically or by evaluating its derivative) that this distance is a convex function. We actually have $A < \frac{N}{r} + 1$, and hence $A\theta_y < \pi \left(1 + \frac{r}{N}\right)$, but by applying the above bound to

$$\left|\frac{e^{i(A-1)\theta} - 1}{e^{i\theta} - 1} + e^{i(A-1)\theta}\right| \ge \left|\frac{e^{i(A-1)\theta} - 1}{e^{i\theta} - 1}\right| - 1,$$
(6.209)

we can still conclude that

$$\left|\frac{e^{iA\theta} - 1}{e^{i\theta} - 1}\right| \ge \frac{2(A-1)|\theta|}{\pi|\theta|} - 1 = \frac{2}{\pi}A - \left(1 + \frac{2}{\pi}\right).$$
(6.210)

Ignoring a possible correction of order 2/A, then, we find

$$\operatorname{Prob}(y) \ge \left(\frac{4}{\pi^2}\right) \frac{1}{r},\tag{6.211}$$

for each of the r values of y that satisfy eq. (6.205). Therefore, with a probability of at least $4/\pi^2$, the measured value of y will satisfy

$$k\frac{N}{r} - \frac{1}{2} \le y \le k\frac{N}{r} + \frac{1}{2},\tag{6.212}$$

or

$$\frac{k}{r} - \frac{1}{2N} \le \frac{y}{N} \le \frac{k}{r} + \frac{1}{2N},$$
(6.213)

where k is an integer chosen from $\{0, 1, \ldots, r-1\}$. The output of the computation is reasonable likely to be within distance 1/2 of an integer multiple of N/r.

Suppose that we know that $r < M \ll N$. Thus N/r is a rational number with a denominator less than M. Two distinct rational numbers, each with denominator less than M, can be no closer together than $1/M^2$, since $\frac{a}{b} - \frac{c}{d} = \frac{ad-bc}{bd}$. If the measurement outcome y satisfies eq. (6.212), then there is a unique value of k/r (with r < M) determined by y/N, provided that $N \ge M^2$. This value of k/r can be efficiently extracted from the measured y/N, by the continued fraction method.

Now, with probability exceeding $4/\pi^2$, we have found a value of k/r where k is selected (roughly equiprobably) from $\{0, 1, 2, \ldots, r-1\}$. It is reasonably likely that k and r are relatively prime (have no common factor), so that we have succeeded in finding r. With a query of the oracle, we may check whether f(x) = f(x+r). But if $\text{GCD}(k,r) \neq 1$, we have found only a factor (r_1) of r.

If we did not succeed, we could test some nearby values of y (the measured value might have been close to the range $-r/2 \leq yr \pmod{N} \leq r/2$ without actually lying inside), or we could try a few multiples of r (the value of GCD(k, r), if not 1, is probably not large). That failing, we resort to a repetition of the quantum circuit, this time (with probability at least $4/\pi^2$) obtaining a value k'/r. Now k', too, may have a common factor with r, in which case our procedure again determines a factor (r_2) of r. But it is reasonably likely that GCD(k, k') = 1, in which case $r = LCM, (r_1, r_2)$. Indeed, we can estimate the probability that randomly selected k and k' are relatively prime as follows: Since a prime number p divides a fraction 1/p of all numbers, the probability that p divides both k and k' is $1/p^2$. And k and k' are coprime if and only if there is no prime p that divides both. Therefore,

Prob
$$(k, k' \text{ coprime}) = \prod_{\text{prime } p} \left(1 - \frac{1}{p^2}\right) = \frac{1}{\zeta(2)} = \frac{6}{\pi^2} \simeq .607$$
 (6.214)

(where $\zeta(z)$ denotes the Riemann zeta function). Therefore, we are likely to succeed in finding the period r after some constant number (independent of N) of repetitions of the algorithm.

6.9.2 From FFT to QFT

Now let's consider the implementation of the quantum Fourier transform. The Fourier transform

$$\sum_{x} f(x)|x\rangle \to \sum_{y} \left(\frac{1}{\sqrt{N}} \sum_{x} e^{2\pi i x y/N} f(x)\right) |y\rangle, \tag{6.215}$$

is multiplication by an $N \times N$ unitary matrix, where the (x, y) matrix element is $(e^{2\pi i/N})^{xy}$. Naively, this transform requires $O(N^2)$ elementary operations. But there is a well-known and very useful (classical) procedure that reduces the number of operations to $O(N \log N)$. Assuming $N = 2^n$, we express xand y as binary expansions

$$x = x_{n-1} \cdot 2^{n-1} + x_{n-2} \cdot 2^{n-2} + \dots + x_1 \cdot 2 + x_0$$

$$y = y_{n-1} \cdot 2^{n-1} + y_{n-2} \cdot 2^{n-2} + \dots + y_1 \cdot 2 + y_0.$$
 (6.216)

In the product of x and y, we may discard any terms containing n or more powers of 2, as these make no contribution to $e^{2\pi i x y}/2^n$. Hence

$$\frac{xy}{2^n} \equiv y_{n-1}(.x_0) + y_{n-2}(.x_1x_0) + y_{n-3}(.x_2x_1x_0) + \dots + y_1(.x_{n-2}x_{n-3}\dots x_0) + y_0(.x_{n-1}x_{n-2}\dots x_0),$$
(6.217)

where the factors in parentheses are binary expansions; e.g.,

$$x_2 x_1 x_0 = \frac{x_2}{2} + \frac{x_1}{2^2} + \frac{x_0}{2^3}.$$
 (6.218)

We can now evaluate

$$\tilde{f}(x) = \frac{1}{\sqrt{N}} \sum_{y} e^{2\pi i x y/N} f(y),$$
(6.219)

for each of the N values of x. But the sum over y factors into n sums over $y_k = 0, 1$, which can be done sequentially in a time of order n.

With quantum parallelism, we can do far better. From eq. (6.217) we obtain

$$QFT : |x\rangle \to \frac{1}{\sqrt{N}} \sum_{y} e^{2\pi i x y/N} |y\rangle$$
$$= \frac{1}{\sqrt{2^{n}}} \left(|0\rangle + e^{2\pi i (.x_{0})} |1\rangle \right) \left(|0\rangle + e^{2\pi i (.x_{1}x_{0})} |1\rangle \right)$$
$$\dots \left(|0\rangle + e^{2\pi i (.x_{n-1}x_{n-2}...x_{0})} |1\rangle \right).$$
(6.220)

The QFT takes each computational basis state to an *unentangled* state of n qubits; thus we anticipate that it can be efficiently implemented. Indeed, let's consider the case n = 3. We can readily see that the circuit



does the job (but note that the order of the bits has been reversed in the output). Each Hadamard gate acts as

$$\boldsymbol{H}: |x_k\rangle \to \frac{1}{\sqrt{2}} \left(|0\rangle + e^{2\pi i (.x_k)} |1\rangle \right).$$
(6.221)

The other contributions to the relative phase of $|0\rangle$ and $|1\rangle$ in the kth qubit are provided by the two-qubit conditional rotations, where

$$\boldsymbol{R}_d = \begin{pmatrix} 1 & 0\\ 0 & e^{i\pi/2^d} \end{pmatrix}, \qquad (6.222)$$

and d = (k - j) is the "distance" between the qubits.

In the case n = 3, the QFT is constructed from three \boldsymbol{H} gates and three controlled- \boldsymbol{R} gates. For general n, the obvious generalization of this circuit requires $n \boldsymbol{H}$ gates and $\binom{n}{2} = \frac{1}{2}n(n-1)$ controlled R's. A two qubit gate is applied to each pair of qubits, again with controlled relative phase $\pi/2^d$, where d is the "distance" between the qubits. Thus the circuit family that implements QFT has a size of order $(\log N)^2$.

We can reduce the circuit complexity to linear in $\log N$ if we are willing to settle for an implementation of fixed accuracy, because the two-qubit gates acting on distantly separated qubits contribute only exponentially small phases. If we drop the gates acting on pairs with distance greater than m, than each term in eq. (6.217) is replaced by an approximation to m bits of accuracy; the total error in $xy/2^n$ is certainly no worse than $n2^{-m}$, so we can achieve accuracy ε in $xy/2^n$ with $m \geq \log n/\varepsilon$. If we retain only the gates acting on qubit pairs with distance m or less, then the circuit size is $mn \sim n \log n/\varepsilon$. In fact, if we are going to measure in the computational basis immediately after implementing the QFT (or its inverse), a further simplification is possible – no two-qubit gates are needed at all! We first remark that the controlled – \mathbf{R}_d gate acts symmetrically on the two qubits – it acts trivially on $|00\rangle, |01\rangle$, and $|10\rangle$, and modifies the phase of $|11\rangle$ by $e^{i\theta_d}$. Thus, we can interchange the "control" and "target" bits without modifying the gate. With this change, our circuit for the 3-qubit QFT can be redrawn as:



Once we have measured $|y_0\rangle$, we know the value of the control bit in the controlled- \mathbf{R}_1 gate that acted on the first two qubits. Therefore, we will obtain the same probability distribution of measurement outcomes if, instead of applying controlled- \mathbf{R}_1 and then measuring, we instead measure y_0 first, and then apply $(\mathbf{R}_1)^{y_0}$ to the next qubit, conditioned on the outcome of the measurement of the first qubit. Similarly, we can replace the controlled- \mathbf{R}_1 and controlled- \mathbf{R}_2 gates acting on the third qubit by the single qubit rotation

$$(\mathbf{R}_2)^{y_0} (\mathbf{R}_1)^{y_1},$$
 (6.223)

(that is, a rotation with relative phase $\pi(.y_1y_0)$) after the values of y_1 and y_0 have been measured.

Altogether then, if we are going to measure after performing the QFT, only n Hadamard gates and n-1 single-qubit rotations are needed to implement it. The QFT is remarkably simple!

6.10 Factoring

6.10.1 Factoring as period finding

What does the factoring problem (finding the prime factors of a large composite positive integer) have to do with periodicity? There is a well-known

6.10. FACTORING

(randomized) reduction of factoring to determining the period of a function. Although this reduction is not directly related to quantum computing, we will discuss it here for completeness, and because the prospect of using a quantum computer as a factoring engine has generated so much excitement.

Suppose we want to find a factor of the *n*-bit number N. Select pseudorandomly a < N, and compute the greatest common divisor GCD(a, N), which can be done efficiently (in a time of order $(\log N)^3$) using the Euclidean algorithm. If $\text{GCD}(a, N) \neq 1$ then the GCD is a nontrivial factor of N, and we are done. So suppose GCD(a, N) = 1.

[Aside: The Euclidean algorithm. To compute $\text{GCD}(N_1, N_2)$ (for $N_1 > N_2$) first divide N_1 by N_2 obtaining remainder R_1 . Then divide N_2 by R_1 , obtaining remainder R_2 . Divide R_1 by R_2 , etc. until the remainder is 0. The last nonzero remainder is $R = \text{GCD}(N_1, N_2)$. To see that the algorithm works, just note that (1) R divides all previous remainders and hence also N_1 and N_2 , and (2) any number that divides N_1 and N_2 will also divide all remainders, including R. A number that divides both N_1 and N_2 , and also is divided by any number that divides both N_1 and N_2 must be $\text{GCD}(N_1, N_2)$. To see how long the Euclidean algorithm takes, note that

$$R_j = qR_{j+1} + R_{j+2}, (6.224)$$

where $q \ge 1$ and $R_{j+2} < R_{j+1}$; therefore $R_{j+2} < \frac{1}{2}R_j$. Two divisions reduce the remainder by at least a factor of 2, so no more than 2 log N_1 divisions are required, with each division using $O((\log N)^2)$ elementary operations; the total number of operations is $O((\log N)^3)$.]

The numbers a < N coprime to N (having no common factor with N) form a finite group under multiplication mod N. [Why? We need to establish that each element a has an inverse. But for given a < N coprime to N, each $ab \pmod{N}$ is distinct, as b ranges over all b < N coprime to N.¹⁶ Therefore, for some b, we must have $ab \equiv 1 \pmod{N}$; hence the inverse of a exists.] Each element a of this finite group has a finite order r, the smallest positive integer such that

$$a^r \equiv 1 \pmod{N}.\tag{6.225}$$

¹⁶If N divides ab - ab', it must divide b - b'.

The order of $a \mod N$ is the period of the function

$$f_{N,a}(x) = a^x \pmod{N}.$$
 (6.226)

We know there is an efficient quantum algorithm that can find the period of a function; therefore, if we can compute $f_{N,a}$ efficiently, we can find the order of a efficiently.

Computing $f_{N,a}$ may look difficult at first, since the exponent x can be very large. But if $x < 2^m$ and we express x as a binary expansion

$$x = x_{m-1} \cdot 2^{m-1} + x_{m-2} \cdot 2^{m-2} + \ldots + x_0, \tag{6.227}$$

we have

$$a^{x} (\text{mod } N) = (a^{2^{m-1}})^{x_{m-1}} (a^{2^{m-2}})^{x_{m-2}} \dots (a)^{x_0} \pmod{N}.$$

(6.228)

Each $a^{2^{j}}$ has a large exponent, but can be computed efficiently by a *classical* computer, using repeated squaring

$$a^{2^{j}} \pmod{N} = (a^{2^{j-1}})^2 \pmod{N}.$$
 (6.229)

So only m-1 (classical) mod N multiplications are needed to assemble a table of all a^{2^j} 's.

The computation of $a^x \pmod{N}$ is carried out by executing a routine:

INPUT 1

For
$$j = 0$$
 to $m - 1$, if $x_i = 1$, MULTIPLY a^{2^j} .

This routine requires at most $m \mod N$ multiplications, each requiring of order $(\log N)^2$ elementary operations.¹⁷ Since r < N, we will have a reasonable chance of success at extracting the period if we choose $m \sim 2 \log N$. Hence, the computation of $f_{N,a}$ can be carried out by a circuit family of size $O((\log N)^3)$. Schematically, the circuit has the structure:

76

¹⁷Using tricks for performing efficient multiplication of very large numbers, the number of elementary operations can be reduced to $O(\log N \log \log \log \log \log N)$; thus, asymptotically for large N, a circuit family with size $O(\log^2 N \log \log \log \log \log N)$ can compute $f_{N,a}$.



Multiplication by a^{2^j} is performed if the control qubit x_j has the value 1.

Suppose we have found the period r of $a \mod N$. Then *if* r is even, we have

N divides
$$\left(a^{\frac{r}{2}}+1\right)\left(a^{\frac{r}{2}}-1\right)$$
. (6.230)

We know that N does not divide $a^{r/2} - 1$; if it did, the order of a would be $\leq r/2$. Thus, *if* it is also the case that N does not divide $a^{r/2} + 1$, or

$$a^{r/2} \neq -1 \pmod{N},$$
 (6.231)

then N must have a nontrivial common factor with each of $a^{r/2} \pm 1$. Therefore, $\text{GCD}(N, a^{r/2} + 1) \neq 1$ is a factor (that we can find efficiently by a classical computation), and we are done.

We see that, once we have found r, we succeed in factoring N unless either (1) r is odd or (2) r is even and $a^{r/2} \equiv -1 \pmod{N}$. How likely is success?

Let's suppose that N is a product of two prime factors $p_1 \neq p_2$,

$$N = p_1 p_2$$
 (6.232)

(this is actually the least favorable case). For each $a < p_1p_2$, there exist unique $a_1 < p_1$ and $a_2 < p_2$ such that

$$a \equiv a_1 \pmod{p_1}$$

$$a \equiv a_2 \pmod{p_2}.$$
 (6.233)

Choosing a random a < N is, therefore, equivalent to choosing random $a, < p_1$ and $a_2 < p_2$.

[Aside: We're using the Chinese Remainder Theorem. The a solving eq. (6.233) is unique because if a and b are both solutions, then both

 p_1 and p_2 must divide a-b. The solution exists because every $a < p_1p_2$ solves eq. (6.233) for *some* a_1 and a_2 . Since there are exactly p_1p_2 ways to choose a_1 and a_2 , and exactly p_1p_2 ways to choose a, uniqueness implies that there is an a corresponding to each pair a_1, a_2 .]

Now let r_1 denote the order of $a_1 \mod p_1$ and r_2 denote the order of $a_2 \mod p_2$. The Chinese remainder theorem tells us that $a^r \equiv 1 \pmod{p_1 p_2}$ is equivalent to

$$a_1^r \equiv 1 \pmod{p_1}$$

$$a_2^r \equiv 1 \pmod{p_2}.$$
(6.234)

Therefore $r = \text{LCM}(r_1, r_2)$. If r_1 and r_2 are both odd, then so is r, and we lose.

But if either r_1 or r_2 is even, then so is r, and we are still in the game. If

$$a^{r/2} \equiv -1 \pmod{p_1}$$

 $a^{r/2} \equiv -1 \pmod{p_2}.$ (6.235)

Then we have $a^{r/2} \equiv -1 \pmod{p_1 p_2}$ and we still lose. But if either

$$a^{r/2} \equiv -1 \pmod{p_1}$$

 $a^{r/2} \equiv 1 \pmod{p_2},$ (6.236)

or

$$a^{r/2} \equiv 1 \pmod{p_1}$$

$$a^{r/2} \equiv -1 \pmod{p_2}, \tag{6.237}$$

then $a^{r/2} \not\equiv -1 \pmod{p_1 p_2}$ and we win. (Of course, $a^{r/2} \equiv 1 \pmod{p_1}$ and $a^{r/2} \equiv 1 \pmod{p_2}$ is not possible, for that would imply $a^{r/2} \equiv 1 \pmod{p_1 p_2}$, and r could not be the order of a.)

Suppose that

$$r_1 = 2^{c_1} \cdot \text{odd}$$

$$r_2 = 2^{c_2} \cdot \text{odd},$$
(6.238)

where $c_1 > c_2$. Then $r = \text{LCM}(r_1, r_2) = 2r_2 \cdot \text{integer}$, so that $a^{r/2} \equiv 1 \pmod{p_2}$ and eq. (6.236) is satisfied – we win! Similarly $c_2 > c_1$ implies eq. (6.237) – again we win. But for $c_1 = c_2$, $r = r_1 \cdot (\text{odd}) = r_2 \cdot (\text{odd}')$ so that eq. (6.235) is satisfied – in that case we lose.

6.10. FACTORING

Okay – it comes down to: for $c_1 = c_2$ we lose, for $c_1 \neq c_2$ we win. How likely is $c_1 \neq c_2$?

It helps to know that the multiplicative group mod p is cyclic – it contains a primitive element of order p - 1, so that all elements are powers of the primitive element. [Why? The integers mod p are a finite *field*. If the group were not cyclic, the maximum order of the elements would be q , so $that <math>x^q \equiv 1 \pmod{p}$ would have p - 1 solutions. But that can't be: in a finite field there are no more than q qth roots of unity.]

Suppose that $p - 1 = 2^k \cdot s$, where s is odd, and consider the orders of all the elements of the cyclic group of order p - 1. For brevity, we'll discuss only the case k = 1, which is the least favorable case for us. Then if b is a primitive (order 2s) element, the even powers of b have odd order, and the odd powers of b have order 2· (odd). In this case, then, $r = 2^c$ · (odd) where $c \in \{0, 1\}$, each occuring equiprobably. Therefore, if p_1 and p_2 are both of this (unfavorable) type, and a_1, a_2 are chosen randomly, the probability that $c_1 \neq c_2$ is $\frac{1}{2}$. Hence, once we have found r, our probability of successfully finding a factor is at least $\frac{1}{2}$, if N is a product of two distinct primes. If N has more than two distinct prime factors, our odds are even better. The method fails if N is a prime power, $N = p^{\alpha}$, but prime powers can be efficiently factored by other methods.

6.10.2 RSA

Does anyone care whether factoring is easy or hard? Well, yes, some people do.

The presumed difficulty of factoring is the basis of the security of the widely used RSA¹⁸ scheme for public key cryptography, which you may have used yourself if you have ever sent your credit card number over the internet.

The idea behind public key cryptography is to avoid the need to exchange a secret key (which might be intercepted and copied) between the parties that want to communicate. The enciphering key is public knowledge. But using the enciphering key to infer the deciphering key involves a prohibitively difficult computation. Therefore, Bob can send the enciphering key to Alice and everyone else, but only Bob will be able to decode the message that Alice (or anyone else) encodes using the key. Encoding is a "one-way function" that is easy to compute but very hard to invert.

¹⁸For Rivest, Shamir, and Adleman

(Of course, Alice and Bob could have avoided the need to exchange the public key if they had decided on a private key in their previous clandestine meeting. For example, they could have agreed to use a long random string as a one-time pad for encoding and decoding. But perhaps Alice and Bob never anticipated that they would someday need to communicate privately. Or perhaps they did agree in advance to use a one-time pad, but they have now used up their private key, and they are loath to reuse it for fear that an eavesdropper might then be able to break their code. Now they are two far apart to safely exchange a new private key; public key cryptography appears to be their most secure option.)

To construct the public key Bob chooses two large prime numbers p and q. But he does not publicly reveal their values. Instead he computes the product

$$N = pq. \tag{6.239}$$

Since Bob knows the prime factorization of N, he also knows the value of the Euler function $\varphi(N)$ – the number of number less than N that are coprime with N. In the case of a product of two primes it is

$$\varphi(N) = N - p - q + 1 = (p - 1)(q - 1), \tag{6.240}$$

(only multiples of p and q share a factor with N). It is easy to find $\varphi(N)$ if you know the prime factorization of N, but it is hard if you know only N.

Bob then pseudo-randomly selects $e < \varphi(N)$ that is coprime with $\varphi(N)$. He reveals to Alice (and anyone else who is listening) the value of N and e, but nothing else.

Alice converts her message to ASCII, a number a < N. She encodes the message by computing

$$b = f(a) = a^e \pmod{N},$$
 (6.241)

which she can do quickly by repeated squaring. How does Bob decode the message?

Suppose that a is coprime to N (which is overwhelmingly likely if p and q are very large – anyway Alice can check before she encodes). Then

$$a^{\varphi(N)} \equiv 1 \pmod{N} \tag{6.242}$$

(Euler's theorem). This is so because the numbers less than N and coprime to N form a group (of order $\varphi(N)$) under mod N multiplication. The order of any group element must divide the order of the group (the powers of a form a subgroup). Since $\text{GCD}(e, \varphi(N) = 1$, we know that e has a multiplicative inverse $d = e^{-1} \mod \varphi(N)$:

$$ed \equiv 1 \pmod{\varphi(N)}.$$
(6.243)

The value of d is Bob's closely guarded secret; he uses it to decode by computing:

$$f^{-1}(b) = b^d \pmod{N}$$

= $a^{ed} \pmod{N}$
= $a \cdot (a^{\varphi(N)})^{\text{integer}} \pmod{N}$
= $a \pmod{N}$. (6.244)

[Aside: How does Bob compute $d = e^{-1}$? The multiplicative inverse is a byproduct of carrying out the Euclidean algorithm to compute $\text{GCD}(e, \varphi(N)) = 1$. Tracing the chain of remainders from the bottom up, starting with $R_n = 1$:

$$1 = R_n = R_{n-2} - q_{n-1}R_{n-1}$$

$$R_{n-1} = R_{n-3} - q_{n-2}R_{n-2}$$

$$R_{n-2} = R_{n-4} - q_{n-3}R_{n-3}$$

$$etc...$$
(6.245)

(where the q_j 's are the quotients), so that

$$1 = (1 + q_{n-1}q_{n-2})R_{n-2} - q_{n-1}R_{n-3}$$

$$1 = (-q_{n-1} - q_{n-3}(1 + q_{n-1}q_{n-2}))R_{n-3}$$

$$+ (1 + q_{n-1}q_{n-2})R_{n-4},$$

etc.... (6.246)

Continuing, we can express 1 as a linear combination of any two successive remainders; eventually we work our way up to

$$1 = d \cdot e + q \cdot \varphi(N), \tag{6.247}$$

and identify d as $e^{-1} \pmod{\varphi(N)}$.

Of course, if Eve has a superfast factoring engine, the RSA scheme is insecure. She factors N, finds $\varphi(N)$, and quickly computes d. In fact, she does not really need to factor N; it is sufficient to compute the order modulo N of the encoded message $a^e \pmod{N}$. Since e is coprime with $\varphi(N)$, the order of $a^e \pmod{N}$ is the same as the order of a (both elements generate the same *orbit*, or cyclic subgroup). Once the order $\operatorname{Ord}(a)$ is known, Eve computes \tilde{d} such that

$$de \equiv 1 \pmod{\operatorname{Ord}(a)} \tag{6.248}$$

so that

$$(a^e)^{\tilde{d}} \equiv a \cdot (a^{\operatorname{Ord}(a)})^{\operatorname{integer}} \pmod{N} \equiv a \pmod{N},$$
(6.249)

and Eve can decipher the message. If our only concern is to defeat RSA, we run the Shor algorithm to find $r = \operatorname{Ord}(a^e)$, and we needn't worry about whether we can use r to extract a factor of N or not.

How important are such prospective cryptographic applications of quantum computing? When fast quantum computers are readily available, concerned parties can stop using RSA, or can use longer keys to stay a step ahead of contemporary technology. However, people with secrets sometimes want their messages to remain confidential for a while (30 years?). They may not be satisfied by longer keys if they are not confident about the pace of future technological advances.

And if they shun RSA, what will they use instead? Not so many suitable one-way functions are known, and others besides RSA are (or may be) vulnerable to a quantum attack. So there really is a lot at stake. If fast large scale quantum computers become available, the cryptographic implications may be far reaching.

But while quantum theory taketh away, quantum theory also giveth; quantum computers may compromise public key schemes, but also offer an alternative: secure quantum key distribution, as discussed in Chapter 4.

6.11 Phase Estimation

There is an alternative way to view the factoring algorithm (due to Kitaev) that deepens our insight into how it works: we can factor because we can

measure efficiently and accurately the eigenvalue of a certain unitary operator.

Consider a < N coprime to N, let x take values in $\{0, 1, 2, ..., N-1\}$, and let U_a denote the unitary operator

$$\boldsymbol{U}_a:|x\rangle \to |ax \;(\text{mod }N)\rangle. \tag{6.250}$$

This operator is unitary (a permutation of the computational basis) because multiplication by $a \mod N$ is invertible.

If the order of $a \mod N$ is r, then

$$U_a^r = 1.$$
 (6.251)

It follows that all eigenvalues of U_a are rth roots of unity:

$$\lambda_k = e^{2\pi i k/r}, \quad k \in \{0, 1, 2, \dots, r-1\}.$$
 (6.252)

The corresponding eigenstates are

$$|\lambda_k\rangle = \frac{1}{\sqrt{r}} \sum_{j=0}^{r-1} e^{-2\pi i k j/r} |a^j x_0 \pmod{N}\rangle;$$
 (6.253)

associated with each orbit of length r generated by multiplication by a, there are r mutually orthogonal eigenstates.

 U_a is not hermitian, but its *phase* (the Hermitian operator that generates U_a) is an observable quantity. Suppose that we can perform a measurement that projects onto the basis of U_a eigenstates, and determines a value λ_k selected equiprobably from the possible eigenvalues. Hence the measurement determines a value of k/r, as does Shor's procedure, and we can proceed to factor N with a reasonably high success probability. But how do we measure the eigenvalues of a unitary operator?

Suppose that we can execute the unitary U conditioned on a control bit, and consider the circuit:



Here $|\lambda\rangle$ denotes an eigenstate of U with eigenvalue λ ($U|\lambda\rangle = \lambda |\lambda\rangle$). Then the action of the circuit on the control bit is

$$|0\rangle \rightarrow \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \rightarrow \frac{1}{\sqrt{2}}(|0\rangle + \lambda|1\rangle)$$
$$\rightarrow \frac{1}{2}(1+\lambda)|0\rangle + \frac{1}{2}(1-\lambda)|1\rangle.$$
(6.254)

Then the outcome of the measurement of the control qubit has probability distribution

$$Prob(0) = \left|\frac{1}{2}(1+\lambda)\right|^2 = \cos^2(\pi\phi)$$
$$Prob(1) = \left|\frac{1}{2}(1-\lambda)\right|^2 = \sin^2(\pi\phi),$$
(6.255)

where $\lambda = e^{2\pi i \phi}$.

As we have discussed previously (for example in connection with Deutsch's problem), this procedure distinguishes with certainty between the eigenvalues $\lambda = 1$ ($\phi = 0$) and $\lambda = -1$ ($\phi = 1/2$). But other possible values of λ can also be distinguished, albeit with less statistical confidence. For example, suppose the state on which U acts is a superposition of U eigenstates

$$\alpha_1 |\lambda_1\rangle + \alpha_2 |\lambda_2\rangle. \tag{6.256}$$

And suppose we execute the above circuit n times, with n distinct control bits. We thus prepare the state

$$\alpha_{1}|\lambda_{1}\rangle \left(\frac{1+\lambda_{1}}{2}|0\rangle + \frac{1-\lambda_{1}}{2}|1\rangle\right)^{\otimes n} + \alpha_{2}|\lambda_{2}\rangle \left(\frac{1+\lambda_{2}}{2}|0\rangle + \frac{1-\lambda_{2}}{2}|1\rangle\right)^{\otimes n}.$$
(6.257)

If $\lambda_1 \neq \lambda_2$, the overlap between the two states of the *n* control bits is exponentially small for large *n*; by measuring the control bits, we can perform the orthogonal projection onto the $\{|\lambda_1\rangle, |\lambda_2\rangle\}$ basis, at least to an excellent approximation.

If we use enough control bits, we have a large enough sample to measure Prob $(0) = \frac{1}{2}(1 + \cos 2\pi\phi)$ with reasonable statistical confidence. By executing a controlled- $(i\mathbf{U})$, we can also measure $\frac{1}{2}(1 + \sin 2\pi\phi)$ which suffices to determine ϕ modulo an integer. However, in the factoring algorithm, we need to measure the phase of $e^{2\pi i k/r}$ to exponential accuracy, which seems to require an exponential number of trials. Suppose, though, that we can efficiently compute high powers of U (as is the case for U_a) such as

$$\boldsymbol{U}^{2^{j}}.$$
 (6.258)

By applying the above procedure to measurement of $U^{2^{j}}$, we determine

$$\exp(2\pi i 2^j \phi),\tag{6.259}$$

where $e^{2\pi i\phi}$ is an eigenvalue of U. Hence, measuring U^{2^j} to one bit of accuracy is equivalent to measuring the *j*th bit of the eigenvalue of U.

We can use this phase estimation procedure for order finding, and hence factorization. We invert eq. (6.253) to obtain

$$x_0 \rangle = \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} |\lambda_k\rangle; \qquad (6.260)$$

each computational basis state (for $x_0 \neq 0$) is an equally weighted superposition of r eigenstates of U_a .

Measuring the eigenvalue, we obtain $\lambda_k = e^{2\pi i k/r}$, with k selected from $\{0, 1, \ldots, r-1\}$ equiprobably. If $r < 2^n$, we measure to 2n bits of precision to determine k/r. In principle, we can carry out this procedure in a computer that stores fewer qubits than we would need to evaluate the QFT, because we can attack just one bit of k/r at a time.

But it is instructive to imagine that we incorporate the QFT into this phase estimation procedure. Suppose the circuit



acts on the eigenstate $|\lambda\rangle$ of the unitary transformation U. The conditional U prepares $\frac{1}{\sqrt{2}}(|0\rangle + \lambda|1\rangle)$, the conditional U^2 prepares $\frac{1}{\sqrt{2}}(|0\rangle + \lambda^2|1\rangle)$, the conditional U^4 prepares $\frac{1}{\sqrt{2}}(|0\rangle + \lambda^4|1\rangle)$, and so on. We could perform a Hadamard and measure each of these qubits to sample the probability distribution governed by the *j*th bit of ϕ , where $\lambda = e^{2\pi i \phi}$. But a more efficient method is to note that the state prepared by the circuit is

$$\frac{1}{\sqrt{2^m}} \sum_{y=0}^{2^m-1} e^{2\pi i \phi y} |y\rangle.$$
(6.261)

A better way to learn the value of ϕ is to perform the QFT^(m), not the Hadamard $\mathbf{H}^{(m)}$, before we measure.

Considering the case m = 3 for clarity, the circuit that prepares and then Fourier analyzes the state

$$\frac{1}{\sqrt{8}} \sum_{y=0}^{7} e^{2\pi i \phi y} |y\rangle \tag{6.262}$$



This circuit very nearly carries out our strategy for phase estimation outlined above, but with a significant modification. Before we execute the final Hadamard transformation and measurement of \tilde{y}_1 and \tilde{y}_2 , some conditional phase rotations are performed. It is those phase rotations that distinguish the QFT⁽³⁾ from Hadamard transform $\boldsymbol{H}^{(3)}$, and they strongly enhance the reliability with which we can extract the value of ϕ .

We can understand better what the conditional rotations are doing if we suppose that $\phi = k/8$, for $k \in \{0, 1, 2..., 7\}$; in that case, we know that the Fourier transform will generate the output $\tilde{y} = k$ with probability one. We may express k as the binary expansion

$$k = k_2 k_1 k_0 \equiv k_2 \cdot 4 + k_1 \cdot 2 + k_0. \tag{6.263}$$

is

In fact, the circuit for the least significant bit \tilde{y}_0 of the Fourier transform is precisely Kitaev's measurement circuit applied to the unitary U^4 , whose eigenvalue is

$$(e^{2\pi i\phi})^4 = e^{i\pi k} = e^{i\pi k_0} = \pm 1.$$
(6.264)

The measurement circuit distinguishes eigenvalues ± 1 perfectly, so that $\tilde{y}_0 = k_0$.

The circuit for the next bit \tilde{y}_1 is almost the measurement circuit for U^2 , with eigenvalue

$$(e^{2\pi i\phi})^2 = e^{i\pi k/2} = e^{i\pi(k_1 \cdot k_0)}.$$
(6.265)

Except that the conditional phase rotation has been inserted, which multiplies the phase by $\exp[i\pi(\cdot k_0)]$, resulting in $e^{i\pi k_1}$. Again, applying a Hadamard followed by measurement, we obtain the outcome $\tilde{y}_1 = k_1$ with certainty. Similarly, the circuit for \tilde{y}_2 measures the eigenvalue

$$e^{2\pi i\phi} = e^{i\pi k/4} = e^{i\pi(k_2 \cdot k_1 k_0)},\tag{6.266}$$

except that the conditional rotation removes $e^{i\pi(\cdot k_1k_0)}$, so that the outcome is $\tilde{y}_2 = k_2$ with certainty.

Thus, the QFT implements the phase estimation routine with maximal cleverness. We measure the less significant bits of ϕ first, and we exploit the information gained in the measurements to improve the reliability of our estimate of the more significant bits. Keeping this interpretation in mind, you will find it easy to remember the circuit for the QFT⁽ⁿ⁾!

6.12 Discrete Log

Sorry, I didn't have time for this.

6.13 Simulation of Quantum Systems

Ditto.
6.14 Summary

Classical circuits. The complexity of a problem can be characterized by the size of a uniform family of logic circuits that solve the problem: The problem is hard if the size of the circuit is a superpolynomial function of the size of the input. One classical universal computer can simulate another efficiently, so the classification of complexity is machine independent. The 3-bit Toffoli gate is universal for classical reversible computation. A reversible computer can simulate an irreversible computer without a significant slowdown and without unreasonable memory resources.

Quantum Circuits. Although there is no proof, it seems likely that polynomial-size quantum circuits cannot be simulated by polynomial-size probabilistic classical circuits ($BQP \neq BPP$); however, polynomial space is sufficient ($BQP \subseteq PSPACE$). A noisy quantum circuit can simulate an ideal quantum circuit of size T to acceptable accuracy if each quantum gate has an accuracy of order 1/T. One universal quantum computer can simulate another efficiently, so that the complexity class BQP is machine independent. A generic two-qubit quantum gate, if it can act on any two qubits in a device, is adequate for universal quantum computation. A controlled-NOT gate plus a generic one-qubit gate is also adequate.

Fast Quantum Searching. Exhaustive search for a marked item in an unsorted database of N items can be carried out by a quantum computer in a time of order \sqrt{N} , but no faster. Quadratic quantum speedups can be achieved for some structured search problems, too, but some oracle problems admit no significant quantum speedup. Two parties, each in possession of a table with N entries, can locate a "collision" between their tables by exchanging $O(\sqrt{N})$ qubits, in apparent violation of the spirit (but not the letter) of the Holevo bound.

Period Finding. Exploiting quantum parallelism, the Quantum Fourier Transform in an N-dimensional space can be computed in time of order $(\log N)^2$ (compared to time $N \log N$ for the classical fast Fourier transform); if we are to measure immediately afterward, one qubit gates are sufficient to compute the QFT. Thus quantum computers can efficiently solve certain problems with a periodic structure, such as factoring and the discrete log problem.

6.15 Exercises

6.1 Linear simulation of Toffoli gate.

In class we constructed the *n*-bit Toffoli gate $(\theta^{(n)})$ from 3-bit Toffoli gates $(\theta^{(3)}, s)$. The circuit required only one bit of scratch space, but the number of gates was exponential in *n*. With more scratch, we can substantially reduce the number of gates.

- a) Find a circuit family with $2n 5 \quad \theta^{(3)}$'s that evaluates $\theta^{(n)}$. (Here n 3 scratch bits are used, which are set to 0 at the beginning of the computation and return to the value 0 at the end.)
- b) Find a circuit family with $4n 12 \quad \theta^{(3)}$'s that evaluates $\theta^{(n)}$, which works irrespective of the initial values of the scratch bits. (Again the n 3 scratch bits return to their initial values, but they don't need to be set to zero at the beginning.)

6.2 A universal quantum gate set.

The purpose of this exercise is to complete the demonstration that the controlled-NOT and arbitrary one-qubit gates constitute a universal set.

a) If U is any unitary 2×2 matrix with determinant one, find unitary A, B, and C such that

$$ABC = 1 \tag{6.267}$$

$$A\sigma_x B\sigma_x C = U. \tag{6.268}$$

Hint: From the Euler angle construction, we know that

$$\boldsymbol{U} = \boldsymbol{R}_{z}(\psi)\boldsymbol{R}_{y}(\theta)\boldsymbol{R}_{z}(\phi), \qquad (6.269)$$

where, e.g., $\mathbf{R}_{z}(\phi)$ denotes a rotation about the z-axis by the angle ϕ . We also know that, e.g.,

$$\boldsymbol{\sigma}_{x}\boldsymbol{R}_{z}(\phi)\boldsymbol{\sigma}_{x} = \boldsymbol{R}_{z}(-\phi). \tag{6.270}$$

b) Consider a two-qubit controlled phase gate: it applies $U = e^{i\alpha} 1$ to the second qubit if the first qubit has value $|1\rangle$, and acts trivially otherwise. Show that it is actually a one-qubit gate.

c) Draw a circuit using controlled-NOT gates and single-qubit gates that implements controlled-U, where U is an arbitrary 2×2 unitary transformation.

6.3 Precision.

The purpose of this exercise is to connect the accuracy of a quantum state with the accuracy of the corresponding probability distribution.

a) Let $\|A\|_{sup}$ denote the sup norm of the operator A, and let

$$\|\boldsymbol{A}\|_{tr} = tr\left[(\boldsymbol{A}^{\dagger}\boldsymbol{A})^{1/2}\right], \qquad (6.271)$$

denote its *trace norm*. Show that

$$\|\boldsymbol{A}\boldsymbol{B}\|_{\mathrm{tr}} \leq \|\boldsymbol{B}\|_{\mathrm{sup}} \cdot \|\boldsymbol{A}\|_{\mathrm{tr}} \text{ and } |\operatorname{tr}\boldsymbol{A}| \leq \|\boldsymbol{A}\|_{\mathrm{tr}} .$$
(6.272)

b) Suppose ρ and $\tilde{\rho}$ are two density matrices, and $\{|a\rangle\}$ is a complete orthonormal basis, so that

$$P_{a} = \langle a | \boldsymbol{\rho} | a \rangle,$$

$$\tilde{P}_{a} = \langle a | \tilde{\boldsymbol{\rho}} | a \rangle,$$
(6.273)

are the corresponding probability distributions. Use (a) to show that

$$\sum_{a} |P_a - \tilde{P}_a| \leq \| \boldsymbol{\rho} - \tilde{\boldsymbol{\rho}} \|_{\mathrm{tr}} .$$
 (6.274)

c) Suppose that $\rho = |\psi\rangle\langle\psi|$ and $\tilde{\rho} = |\tilde{\psi}\rangle\langle\tilde{\psi}|$ are pure states. Use (b) to show that

$$\sum_{a} |P_a - \tilde{P}_a| \leq 2 || |\psi\rangle - |\tilde{\psi}\rangle ||. \qquad (6.275)$$

6.4 Continuous-time database search

A quantum system with an *n*-qubit Hilbert space has the Hamiltonian

$$\boldsymbol{H}_{\omega} = E|\omega\rangle\langle\omega|, \qquad (6.276)$$

where $|\omega\rangle$ is an unknown computational-basis state. You are to find the value of ω by the following procedure. Turn on a time-independent perturbation \boldsymbol{H}' of the Hamiltonian, so that the total Hamiltonian becomes

$$\boldsymbol{H} = \boldsymbol{H}_{\omega} + \boldsymbol{H}'. \tag{6.277}$$

Prepare an initial state $|\psi_0\rangle$, and allow the state to evolve, as governed by \boldsymbol{H} , for a time T. Then measure the state. From the measurement result you are to infer ω .

a) Suppose the initial state is chosen to be

$$|s\rangle = \frac{1}{2^{n/2}} \sum_{x=0}^{2^n - 1} |x\rangle, \qquad (6.278)$$

and the perturbation is

$$\mathbf{H}' = E|s\rangle\langle s|. \tag{6.279}$$

Solve the time-independent Schrödinger equation

$$i\frac{d}{dt}|\psi\rangle = \boldsymbol{H}|\psi\rangle \tag{6.280}$$

to find the state at time T. How should T be chosen to optimize the likelihood of successfully determining ω ?

b) Now suppose that we may choose $|\psi_0\rangle$ and H' however we please, but we demand that the state of the system after time T is $|\omega\rangle$, so that the measurement determines ω with success probability one. Derive a lower bound that T must satisfy, and compare to your result in (a). (Hint: As in our analysis in class, compare evolution governed by Hwith evolution governed by H' (the case of the "empty oracle"), and use the Schrödinger equation to bound how rapidly the state evolving according to H deviates from the state evolving according to H'.)

Chapter 7

Quantum Error Correction

7.1 A Quantum Error-Correcting Code

In our study of quantum algorithms, we have found persuasive evidence that a quantum computer would have extraordinary power. But will quantum computers really work? Will we ever be able to build and operate them?

To do so, we must rise to the challenge of protecting quantum information from errors. As we have already noted in Chapter 1, there are several aspects to this challenge. A quantum computer will inevitably interact with its surroundings, resulting in decoherence and hence in the decay of the quantum information stored in the device. Unless we can successfully combat decoherence, our computer is sure to fail. And even if we were able to prevent decoherence by perfectly isolating the computer from the environment, errors would still pose grave difficulties. Quantum gates (in contrast to classical gates) are unitary transformations chosen from a continuum of possible values. Thus quantum gates cannot be implemented with perfect accuracy; the effects of small imperfections in the gates will accumulate, eventually leading to a serious failure in the computation. Any effective strategem to prevent errors in a quantum computer must protect against small unitary errors in a quantum circuit, as well as against decoherence.

In this and the next chapter we will see how clever encoding of quantum information can protect against errors (in principle). This chapter will present the theory of quantum error-correcting codes. We will learn that quantum information, suitably encoded, can be deposited in a quantum memory, exposed to the ravages of a noisy environment, and recovered without damage (if the noise is not too severe). Then in Chapter 8, we will extend the theory in two important ways. We will see that the recovery procedure can work effectively even if occasional errors occur during recovery. And we will learn how to *process* encoded information, so that a quantum *computation* can be executed successfully despite the debilitating effects of decoherence and faulty quantum gates.

A quantum error-correcting code (QECC) can be viewed as a mapping of k qubits (a Hilbert space of dimension 2^k) into n qubits (a Hilbert space of dimension 2^n), where n > k. The k qubits are the "logical qubits" or "encoded qubits" that we wish to protect from error. The additional n - kqubits allow us to store the k logical qubits in a redundant fashion, so that the encoded information is not easily damaged.

We can better understand the concept of a QECC by revisiting an example that was introduced in Chapter 1, Shor's code with n = 9 and k = 1. We can characterize the code by specifying two basis states for the code subspace; we will refer to these basis states as $|\bar{0}\rangle$, the "logical zero" and $|\bar{1}\rangle$, the "logical one." They are

$$\begin{split} |\bar{0}\rangle &= \left[\frac{1}{\sqrt{2}}(|000\rangle + |111\rangle)\right]^{\otimes 3},\\ |\bar{1}\rangle &= \left[\frac{1}{\sqrt{2}}(|000\rangle - |111\rangle)\right]^{\otimes 3}; \end{split}$$
(7.1)

each basis state is a 3-qubit cat state, repeated three times. As you will recall from the discussion of cat states in Chapter 4, the two basis states can be distinguished by the 3-qubit observable $\sigma_x^{(1)} \otimes \sigma_x^{(2)} \otimes \sigma_x^{(3)}$ (where $\sigma_x^{(i)}$ denotes the Pauli matrix σ_x acting on the ith qubit); we will use the notation $X_1 X_2 X_3$ for this operator. (There is an implicit $I \otimes I \otimes \cdots \otimes I$ acting on the remaining qubits that is suppressed in this notation.) The states $|\bar{0}\rangle$ and $|\bar{1}\rangle$ are eigenstates of $X_1 X_2 X_3$ with eigenvalues +1 and -1 respectively. But there is no way to distinguish $|\bar{0}\rangle$ from $|\bar{1}\rangle$ (to gather any information about the value of the logical qubit) by observing any one or two of the qubits in the block of nine. In this sense, the logical qubit is encoded *nonlocally*; it is written in the nature of the entanglement among the qubits in the block. This nonlocal property of the encoded information provides protection against noise, if we assume that the noise is local (that it acts independently, or nearly so, on the different qubits in the block).

Suppose that an unknown quantum state has been prepared and encoded as $a|\bar{0}\rangle + b|\bar{1}\rangle$. Now an error occurs; we are to diagnose the error and reverse it. How do we proceed? Let us suppose, to begin with, that a single bit flip occurs acting on one of the first three qubits. Then, as discussed in Chapter 1, the location of the bit flip can be determined by measuring the two-qubit operators

$$\boldsymbol{Z}_1 \boldsymbol{Z}_2 , \quad \boldsymbol{Z}_2 \boldsymbol{Z}_3. \tag{7.2}$$

The logical basis states $|\bar{0}\rangle$ and $|\bar{1}\rangle$ are eigenstates of these operators with eigenvalue 1. But flipping any of the three qubits changes these eigenvalues. For example, if $Z_1Z_2 = -1$ and $Z_2Z_3 = 1$, then we infer that the first qubit has flipped relative to the other two. We may recover from the error by flipping that qubit back.

It is crucial that our measurement to diagnose the bit flip is a collective measurement on two qubits at once — we learn the value of Z_1Z_2 , but we must not find out about the separate values of Z_1 and Z_2 , for to do so would damage the encoded state. How can such a collective measurement be performed? In fact we can carry out collective measurements if we have a quantum computer that can execute controlled-NOT gates. We first introduce an additional "ancilla" qubit prepared in the state $|0\rangle$, then execute the quantum circuit

and finally measure the ancilla qubit. If the qubits 1 and 2 are in a state with $\mathbf{Z}_1\mathbf{Z}_2 = -1$ (either $|0\rangle_1|1\rangle_2$ or $|1\rangle_1|0\rangle_2$), then the ancilla qubit will flip once and the measurement outcome will be $|1\rangle$. But if qubits 1 and 2 are in a state with $\mathbf{Z}_1\mathbf{Z}_2 = 1$ (either $|0\rangle_1|0\rangle_2$ or $|1\rangle_1|1\rangle_2$), then the ancilla qubit will flip either twice or not at all, and the measurement outcome will be $|0\rangle$. Similarly, the two-qubit operators

can be measured to diagnose bit flip errors in the other two clusters of three qubits.

A three-qubit code would suffice to protect against a single bit flip. The reason the 3-qubit clusters are repeated three times is to protect against phase errors as well. Suppose now that a phase error

$$|\psi\rangle \to \mathbf{Z}|\psi\rangle$$
 (7.4)

occurs acting on one of the nine qubits. We can diagnose in which cluster the phase error occurred by measuring the two six-qubit observables

$$X_1 X_2 X_3 X_4 X_5 X_6,$$

 $X_4 X_5 X_6 X_7 X_8 X_9.$ (7.5)

The logical basis states $|\bar{0}\rangle$ and $|\bar{1}\rangle$ are both eigenstates with eigenvalue one of these observables. A phase error acting on any one of the qubits in a particular cluster will change the value of XXX in that cluster relative to the other two; the location of the change can be identified by measuring the observables in eq. (7.5). Once the affected cluster is identified, we can reverse the error by applying Z to one of the qubits in that cluster.

How do we measure the six-qubit observable $X_1X_2X_3X_4X_5X_6$? Notice that if its control qubit is initially in the state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, and its target is an eigenstate of X (that is, NOT) then a controlled-NOT acts according to

$$\text{CNOT}: \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |x\rangle \to \frac{1}{\sqrt{2}}(|0\rangle + (-1)^{x}|1\rangle) \otimes |x\rangle;$$
(7.6)

it acts trivially if the target is the X = 1 (x = 0) state, and it flips the control if the target is the X = -1 (x = 1) state. To measure a product of X's, then, we execute the circuit

and then measure the ancilla in the $\frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$ basis.

We see that a single error acting on any one of the nine qubits in the block will cause no irrevocable damage. But if two bit flips occur in a single cluster of three qubits, then the encoded information *will* be damaged. For example, if the first two qubits in a cluster both flip, we will misdiagnose the error and attempt to recover by flipping the third. In all, the errors, together with our mistaken recovery attempt, apply the operator $X_1 X_2 X_3$ to the code block. Since $|\bar{0}\rangle$ and $|\bar{1}\rangle$ are eigenstates of $X_1 X_2 X_3$ with distinct eigenvalues, the effect of two bit flips in a single cluster is a *phase error* in the encoded qubit:

$$\boldsymbol{X}_1 \boldsymbol{X}_2 \boldsymbol{X}_3 : a|\bar{0}\rangle + b|\bar{1}\rangle \to a|\bar{0}\rangle - b|\bar{1}\rangle .$$
 (7.7)

The encoded information will also be damaged if phase errors occur in two different clusters. Then we will introduce a phase error into the third cluster in our misguided attempt at recovery, so that altogether $Z_1Z_4Z_7$ will have been applied, which flips the encoded qubit:

$$\mathbf{Z}_1 \mathbf{Z}_4 \mathbf{Z}_7 : a|\bar{0}\rangle + b|\bar{1}\rangle \to a|\bar{1}\rangle + b|\bar{0}\rangle .$$
 (7.8)

If the likelihood of an error is small enough, and if the errors acting on distinct qubits are not strongly correlated, then using the nine-qubit code will allow us to preserve our unknown qubit more reliably than if we had not bothered to encode it at all. Suppose, for example, that the environment acts on each of the nine qubits, independently subjecting it to the depolarizing channel described in Chapter 3, with error probability p. Then a bit flip occurs with probability $\frac{2}{3}p$, and a phase flip with probability $\frac{2}{3}p$. (The probability that both occur is $\frac{1}{3}p$). We can see that the probability of a phase error affecting the logical qubit is bounded above by $4p^2$, and the probability of a bit flip error is bounded above by $12p^2$. The total error probability is no worse than $16p^2$; this is an improvement over the error probability p for an unprotected qubit, provided that p < 1/16.

Of course, in this analysis we have implicitly assumed that encoding, decoding, error syndrome measurement, and recovery are all performed flaw-lessly. In Chapter 8 we will examine the more realistic case in which errors occur during these operations.

7.2 Criteria for Quantum Error Correction

In our discussion of error recovery using the nine-qubit code, we have assumed that each qubit undergoes either a bit-flip error or a phase-flip error (or both). This is not a realistic model for the errors, and we must understand how to implement quantum error correction under more general conditions.

To begin with, consider a single qubit, initially in a pure state, that interacts with its environment in an arbitrary manner. We know from Chapter 3 that there is no loss or generality (we may still represent the most general superoperator acting on our qubit) if we assume that the initial state of the environment is a pure state, which we will denote as $|0\rangle_E$. Then the evolution of the qubit and its environment can be described by a unitary transformation

$$U: |0\rangle \otimes |0\rangle_E \to |0\rangle \otimes |e_{00}\rangle_E + |1\rangle \otimes |e_{01}\rangle_E , |1\rangle \otimes |0\rangle_E \to |0\rangle \otimes |e_{10}\rangle_E + |1\rangle \otimes |e_{11}\rangle_E ;$$
(7.9)

here the four $|e_{ij}\rangle_E$ are states of the environment that need not be normalized or mutually orthogonal (though they do satisfy some constraints that follow from the unitarity of U). Under U, an arbitrary state $|\psi\rangle = a|0\rangle + b|1\rangle$ of the qubit evolves as

$$U: (a|0\rangle + b|1\rangle)|0\rangle_E \to a(|0\rangle|e_{00}\rangle_E + |1\rangle|e_{01}\rangle_E) + b(|0\rangle|e_{10}\rangle_E + |1\rangle|e_{11}\rangle_E)$$

$$= (a|0\rangle + b|1\rangle) \otimes \frac{1}{2}(|e_{00}\rangle_E + |e_{11}\rangle_E)$$
$$+ (a|0\rangle - b|1\rangle) \otimes \frac{1}{2}(|e_{00}\rangle_E - |e_{11}\rangle_E)$$
$$+ (a|1\rangle + b|0\rangle) \otimes \frac{1}{2}(|e_{01}\rangle_E + |e_{10}\rangle_E)$$
$$+ (a|1\rangle - b|0\rangle) \otimes \frac{1}{2}(|e_{01}\rangle_E - |e_{10}\rangle_E)$$

$$\equiv \boldsymbol{I}|\psi\rangle \otimes |e_I\rangle_E + \boldsymbol{X}|\psi\rangle \otimes |e_X\rangle_E + \boldsymbol{Y}|\psi\rangle \otimes |e_Y\rangle_E + \boldsymbol{Z}|\psi\rangle \otimes |e_Z\rangle_E.$$
(7.10)

The action of U can be expanded in terms of the (unitary) Pauli operators $\{I, X, Y, Z\}$, simply because these are a basis for the vector space of 2×2 matrices. Heuristically, we might interpret this expansion by saying that one of four possible things happens to the qubit: nothing (I), a bit flip (X), a phase flip (Z), or both (Y = iXZ). However, this classification should not be taken literally, because unless the states $\{|e_I\rangle, |e_X\rangle, |e_Y\rangle, |e_Z\rangle\}$ of the environment are all mutually orthogonal, there is no conceivable measurement that could perfectly distinguish among the four alternatives.

Similarly, an arbitrary $2^n \times 2^n$ matrix acting on an *n*-qubit Hilbert space can be expanded in terms of the 2^{2n} operators

$$\{\boldsymbol{I}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}\}^{\otimes n}; \tag{7.11}$$

that is, each such operator can be expressed as a tensor-product "string" of single-qubit operators, with each operator in the string chosen from among the identity and the three Pauli matrices X, Y, and Z. Thus, the action of an arbitrary unitary operator on n qubits plus their environment can be expanded as

$$|\psi\rangle \otimes |0\rangle_E \to \sum_a \boldsymbol{E}_a |\psi\rangle \otimes |e_a\rangle_E;$$
 (7.12)

here the index *a* ranges over 2^{2n} values. The $\{E_a\}$ are the linearly independent Pauli operators acting on the *n* qubits, and the $\{|e_a\rangle_E\}$ are the corresponding states of the environment (which are *not* assumed to be normalized or mutually orthogonal). A crucial feature of this expansion for what follows is that each E_a is a unitary operator.

Eq. (7.12) provides the conceptual foundation of quantum error correction. In devising a quantum error-correcting code, we identify a subset \mathcal{E} of all the Pauli operators,

$$\mathcal{E} \subseteq \{ \boldsymbol{E}_a \} \equiv \{ \boldsymbol{I}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \}^{\otimes n} ; \qquad (7.13)$$

these are the errors that we wish to be able to correct. Our aim will be to perform a collective measurement of the *n* qubits in the code block that will enable us to diagnose which error $E_a \in \mathcal{E}$ occurred. If $|\psi\rangle$ is a state in the code subspace, then for some (but not all) codes this measurement will prepare a state $E_a |\psi\rangle \otimes |e_a\rangle_E$, where the value of *a* is known from the measurement outcome. Since E_a is unitary, we may proceed to apply $E_a^{\dagger} (=$ $E_a)$ to the code block, thus recovering the undamaged state $|\psi\rangle$.

Each Pauli operator can be assigned a *weight*, an integer t with $0 \le t \le n$; the weight is the number of qubits acted on by a nontrivial Pauli matrix $(\mathbf{X}, \mathbf{Y}, \text{ or } \mathbf{Z})$. Heuristically, then, we can interpret a term in the expansion eq. (7.12) where \mathbf{E}_a has weight t as an event in which errors occur on t qubits (but again we cannot take this interpretation too literally if the states $\{|e_a\rangle_E\}$ are not mutually orthogonal). Typically, we will take \mathcal{E} to be the set of all Pauli operators of weight up to and including t; then if we can recover from any error superoperator with support on the set \mathcal{E} , we will say that the code can correct t errors. In adopting such an error set, we are implicitly assuming that the errors afflicting different qubits are only weakly correlated with one another, so that the amplitude for more than t errors on the n qubits is relatively small.

Given the set \mathcal{E} of errors that are to be corrected, what are the necessary and sufficient conditions to be satisfied by the code subspace in order that recovery is possible? Let us denote by $\{ |\bar{i}\rangle \}$ an orthonormal basis for the code subspace. (We will refer to these basis elements as "codewords".) It will clearly be *necessary* that

$$\langle \bar{j} | \boldsymbol{E}_b^{\dagger} \boldsymbol{E}_a | \bar{i} \rangle = 0, \quad i \neq j,$$
(7.14)

where $E_{a,b} \in \mathcal{E}$. If this condition were not satisfied for some $i \neq j$, then errors would be able to destroy the perfect distinguishability of orthogonal codewords, and encoded quantum information could surely be damaged. (A more explicit derivation of this necessary condition will be presented below.) We can also easily see that a *sufficient* condition is

$$\langle \bar{j} | \boldsymbol{E}_b^{\dagger} \boldsymbol{E}_a | \bar{i} \rangle = \delta_{ab} \delta_{ij}. \tag{7.15}$$

In this case the E_a 's take the code subspace to a set of mutually orthogonal "error subspaces"

$$\mathcal{H}_a = \boldsymbol{E}_a \mathcal{H}_{code}.\tag{7.16}$$

Suppose, then that an arbitrary state $|\psi\rangle$ in the code subspace is prepared, and subjected to an error. The resulting state of code block and environment is

$$\sum_{\boldsymbol{E}_{a}\in\mathcal{E}}\boldsymbol{E}_{a}|\psi\rangle\otimes|e_{a}\rangle_{E},$$
(7.17)

where the sum is restricted to the errors in the set \mathcal{E} . We may then perform an orthogonal measurement that projects the code block onto one of the spaces \mathcal{H}_a , so that the state becomes

$$\boldsymbol{E}_a|\psi\rangle\otimes|\boldsymbol{e}_a\rangle_E.\tag{7.18}$$

We finally apply the unitary operator E_a^{\dagger} to the code block to complete the recovery procedure.

7.2. CRITERIA FOR QUANTUM ERROR CORRECTION

A code that satisfies the condition eq. (7.15) is called a *nondegenerate* code. This terminology signifies that there is a measurement that can unambiguously diagnose the error $E_a \in \mathcal{E}$ that occurred. But the example of the nine-qubit code has already taught us that more general codes are possible. The nine-qubit code is *degenerate*, because phase errors acting on different qubits in the same cluster of three affect the code subspace in precisely the same way $(e.g., \mathbb{Z}_1 | \psi \rangle = \mathbb{Z}_2 | \psi \rangle$). Though no measurement can determine which qubit suffered the error, this need not pose an obstacle to successful recovery.

The necessary and sufficient condition for recovery to be possible is easily stated:

$$\langle \bar{j} | \boldsymbol{E}_b^{\dagger} \boldsymbol{E}_a | \bar{i} \rangle = C_{ba} \delta_{ij}, \qquad (7.19)$$

where $\mathbf{E}_{a,b} \in \mathcal{E}$, and $C_{ba} = \langle \bar{i} | \mathbf{E}_b^{\dagger} \mathbf{E}_a | \bar{i} \rangle$ is an arbitrary Hermitian matrix. The nontrivial content of this condition that goes beyond the weaker necessary condition eq. (7.14) is that $\langle \bar{i} | \mathbf{E}_b^{\dagger} \mathbf{E}_a | \bar{i} \rangle$ does not depend on i. The origin of this condition is readily understood — were it otherwise, in identifying an error subspace \mathcal{H}_a we would acquire some information about the encoded state, and so would inevitably disturb that state.

To prove that the condition eq. (7.19) is necessary and sufficient, we invoke the theory of superoperators developed in Chapter 3. Errors acting on the code block are described by a superoperator, and the issue is whether another superoperator (the recovery procedure) can be constructed that will reverse the effect of the error. In fact, we learned in Chapter 3 that the only superoperators that can be inverted are unitary operators. But now we are demanding a bit less. We are not required to be able to reverse the action of the error superoperator on any state in the *n*-qubit code block; rather, it is enough to be able to reverse the errors when the initial state resides in the *k*-qubit encoded subspace.

An alternative way to express the action of an error on one of the code basis states $|\bar{i}\rangle$ (and the environment) is

$$|\bar{i}\rangle \otimes |0\rangle_E \to \sum_{\mu} \boldsymbol{M}_{\mu} |\bar{i}\rangle \otimes |\mu\rangle_E,$$
 (7.20)

where now the states $|\mu\rangle_E$ are elements of an *orthonormal basis* for the environment, and the matrices M_{μ} are linear combinations of the Pauli operators

 E_a contained in \mathcal{E} , satisfying the operator-sum normalization condition

$$\sum_{\mu} \boldsymbol{M}_{\mu}^{\dagger} \boldsymbol{M}_{\mu} = \boldsymbol{I}.$$
 (7.21)

The error can be reversed by a recovery superoperator if there exist operators \mathbf{R}_{ν} such that

$$\sum_{\nu} \boldsymbol{R}_{\nu}^{\dagger} \boldsymbol{R}_{\nu} = \boldsymbol{I}, \qquad (7.22)$$

and

$$\sum_{\mu,\nu} \mathbf{R}_{\nu} \mathbf{M}_{\mu} |\bar{i}\rangle \otimes |\mu\rangle_{E} \otimes |\nu\rangle_{A}$$
$$= |\bar{i}\rangle \otimes |\text{stuff}\rangle_{EA}; \qquad (7.23)$$

here the $|\nu\rangle_A$'s are elements of an orthonormal basis for the Hilbert space of the *ancilla* that is employed to implement the recovery operation, and the state $|\text{stuff}\rangle_{EA}$ of environment and ancilla must not depend on *i*. It follows that

$$\boldsymbol{R}_{\nu}\boldsymbol{M}_{\mu}|\bar{i}\rangle = \lambda_{\nu\mu}|\bar{i}\rangle; \qquad (7.24)$$

for each μ and ν ; the product $\mathbf{R}_{\nu}\mathbf{M}_{\mu}$ acting on the code subspace is a multiple of the identity. Using the normalization condition satisfied by the \mathbf{R}_{ν} 's, we infer that

$$\boldsymbol{M}_{\delta}^{\dagger}\boldsymbol{M}_{\mu}|\bar{i}\rangle = \boldsymbol{M}_{\delta}^{\dagger}\left(\sum_{\nu}\boldsymbol{R}_{\nu}^{\dagger}\boldsymbol{R}_{\nu}\right)\boldsymbol{M}_{\mu}|\bar{i}\rangle = \sum_{\nu}\lambda_{\nu\delta}^{*}\lambda_{\nu\mu}|\bar{i}\rangle, \qquad (7.25)$$

so that $M_{\delta}^{\dagger}M_{\mu}$ is likewise a multiple of the identity acting on the code subspace. In other words

$$\langle \bar{j} | \boldsymbol{M}_{\delta}^{\dagger} \boldsymbol{M}_{\mu} | \bar{i} \rangle = C_{\delta \mu} \delta_{ij};$$
 (7.26)

since each E_a in \mathcal{E} is a linear combination of M_{μ} 's, eq. (7.19) then follows.

Another instructive way to understand why eq. (7.26) is a necessary condition for error recovery is to note that if the code block is prepared in the state $|\psi\rangle$, and an error acts according to eq. (7.20), then the density matrix for the environment that we obtain by tracing over the code block is

$$\rho_E = \sum_{\mu,\nu} |\mu\rangle_E \langle \psi | \boldsymbol{M}_{\nu}^{\dagger} \boldsymbol{M}_{\mu} | \psi \rangle_E \langle \nu |.$$
(7.27)

Error recovery can proceed successfully only if there is no way to acquire any information about the state $|\psi\rangle$ by performing a measurement on the environment. Therefore, we require that ρ_E be independent of $|\psi\rangle$, if $|\psi\rangle$ is any state in the code subspace; eq. (7.26) then follows.

To see that eq. (7.26) is sufficient for recovery as well as necessary, we can explicitly construct the superoperator that reverses the error. For this purpose it is convenient to choose our basis $\{|\mu\rangle_E\}$ for the environment so that the matrix $C_{\delta\mu}$ in eq. (7.26) is diagonalized:

$$\langle \bar{j} | \boldsymbol{M}_{\delta}^{\dagger} \boldsymbol{M}_{\mu} | \bar{i} \rangle = C_{\mu} \delta_{\delta \mu} \delta_{ij} , \qquad (7.28)$$

where $\sum_{\mu} C_{\mu} = 1$ follows from the operator-sum normalization condition. For each ν with $C_{\nu} \neq 0$, let

$$\boldsymbol{R}_{\nu} = \frac{1}{\sqrt{C_{\nu}}} \sum_{i} |\bar{i}\rangle \langle \bar{i} | \boldsymbol{M}_{\nu}^{\dagger}, \qquad (7.29)$$

so that \boldsymbol{R}_{ν} acts according to

$$\boldsymbol{R}_{\nu}: \boldsymbol{M}_{\mu} | \bar{i} \rangle \to \sqrt{C_{\nu}} \delta_{\mu\nu} | \bar{i} \rangle.$$
 (7.30)

Then we easily see that

$$\sum_{\mu,\nu} \mathbf{R}_{\nu} \mathbf{M}_{\mu} |\bar{i}\rangle \otimes |\mu\rangle_{E} \otimes |\nu\rangle_{A}$$
$$= |\bar{i}\rangle \otimes (\sum_{\nu} \sqrt{C_{\nu}} |\nu\rangle_{E} \otimes |\nu\rangle_{A}); \qquad (7.31)$$

the superoperator defined by the \mathbf{R}_{ν} 's does indeed reverse the error. It only remains to check that the \mathbf{R}_{ν} 's satisfy the normalization condition. We have

$$\sum_{\nu} \boldsymbol{R}_{\nu}^{\dagger} \boldsymbol{R}_{\nu} = \sum_{\nu,i} \frac{1}{C_{\nu}} \sum_{\nu} \boldsymbol{M}_{\nu} |\bar{i}\rangle \langle \bar{i} | \boldsymbol{M}_{\nu}^{\dagger} , \qquad (7.32)$$

which is the orthogonal projection onto the space of states that can be reached by errors acting on codewords. Thus we can complete the specification of the recovery superoperator by adding one more element to the operator sum — the projection onto the complementary subspace.

In brief, eq. (7.19) is a sufficient condition for error recovery because it is possible to choose a basis for the error operators (not necessarily the Pauli operator basis) that diagonalizes the matrix C_{ab} , and in this basis we can unambiguously diagnose the error by performing a suitable orthogonal measurement. (The eigenmodes of C_{ab} with eigenvalue zero, like $\mathbf{Z}_1 - \mathbf{Z}_2$ in the case of the 9-qubit code, correspond to errors that occur with probability zero.) We see that, once the set \mathcal{E} of possible errors is specified, the recovery operation is determined. In particular, no information is needed about the states $|e_a\rangle_E$ of the environment that are associated with the errors \mathbf{E}_a . Therefore, the code works equally effectively to control unitary errors or decoherence errors (as long as the amplitude for errors outside of the set \mathcal{E} is negligible). Of course, in the case of a nondegenerate code, C_{ab} is already diagonal in the Pauli basis, and we can express the recovery basis as

$$\boldsymbol{R}_{a} = \sum_{i} |\bar{i}\rangle \langle \bar{i}| \boldsymbol{E}_{a}^{\dagger} ; \qquad (7.33)$$

there is an \mathbf{R}_a corresponding to each \mathbf{E}_a in \mathcal{E} .

We have described error correction as a two step procedure: first a collective measurement is conducted to diagnose the error, and secondly, based on the measurement outcome, a unitary transformation is applied to reverse the error. This point of view has many virtues. In particular, it is the quantum measurement procedure that seems to enable us to tame a continuum of possible errors, as the measurement projects the damaged state into one of a discrete set of outcomes, for each of which there is a prescription for recovery. But in fact measurement is not an essential ingredient of quantum error correction. The recovery superoperator of eq. (7.31) may of course be viewed as a unitary transformation acting on the code block and an ancilla. This superoperator can describe a measurement followed by a unitary operator if we imagine that the ancilla is subjected to an orthogonal measurement, but the measurement is not necessary.

If there is no measurement, we are led to a different perspective on the reversal of decoherence achieved in the recovery step. When the code block interacts with its environment, it becomes entangled with the environment, and the Von Neumann entropy of the environment increases (as does the entropy of the code block). If we are unable to control the environment, that increase in its entropy can never be reversed; how then, is quantum error correction possible? The answer provided by eq. (7.31) is that we may apply a unitary transformation to the data and to an ancilla that we *do* control. If the criteria for quantum error correction are satisfied, this unitary can be chosen to transform the entanglement of the data with the environment into

entanglement of ancilla with environment, restoring the purity of the data in the process, as in:

While measurement is not a necessary part of error correction, the ancilla is absolutely essential. The ancilla serves as a depository for the entropy inserted into the code block by the errors — it "heats" as the data "cools." If we are to continue to protect quantum information stored in quantum memory for a long time, a continuous supply of ancilla qubits should be provided that can be discarded after use. Alternatively, if the ancilla is to be recycled, it must first be erased. As discussed in Chapter 1, the erasure is dissipative and requires the expenditure of power. Thus principles of thermodynamics dictate that we cannot implement (quantum) error correction for free. Errors cause entropy to seep into the data. This entropy can be transferred to the ancilla by means of a reversible process, but work is needed to pump entropy from the ancilla back to the environment.

7.3 Some General Properties of QECC's

7.3.1 Distance

A quantum code is said to be *binary* if it can be represented in terms of qubits. In a binary code, a code subspace of dimension 2^k is embedded in a space of dimension 2^n , where k and n > k are integers. There is actually no need to require that the dimensions of these spaces be powers of two (see the exercises); nevertheless we will mostly confine our attention here to binary coding, which is the simplest case.

In addition to the block size n and the number of encoded qubits k, another important parameter characterizing a code is its *distance d*. The distance d is the minimum weight of a Pauli operator \boldsymbol{E} such that

$$\langle \bar{i} | \boldsymbol{E}_a | \bar{j} \rangle \neq C_a \delta_{ij}.$$
 (7.34)

We will describe a quantum code with block size n, k encoded qubits, and distance d as an "[[n, k, d]] quantum code." We use the double-bracket no-

tation for quantum codes, to distinguish from the [n, k, d] notation used for classical codes.

We say that an QECC can correct t errors if the set \mathcal{E} of \mathbf{E}_a 's that allow recovery includes all Pauli operators of weigh t or less. Our definition of distance implies that the criterion for error correction

$$\langle \bar{i} | \boldsymbol{E}_{a}^{\dagger} \boldsymbol{E}_{b} | \bar{j} \rangle = C_{ab} \delta_{ij}, \qquad (7.35)$$

will be satisfied by all Pauli operators $E_{a,b}$ of weight t or less, provided that $d \ge 2t + 1$. Therefore, a QECC with distance d = 2t + 1 can correct t errors.

7.3.2 Located errors

A distance d = 2t + 1 code can correct t errors, irrespective of the location of the errors in the code block. But in some cases we may know that particular qubits are especially likely to have suffered errors. Perhaps we saw a hammer strike those qubits. Or perhaps you sent a block of n qubits to me, but t < nof the qubits were lost and never received. I am confident that the n - tqubits that did arrive were well packaged and were received undamaged. But I replace the t missing qubits with the (arbitrarily chosen) state $|00...0\rangle$, realizing full well that these qubits are likely to be in error.

A given code can protect against more errors if the errors occur at known locations instead of unknown locations. In fact, a QECC with distance d = t + 1 can correct t errors at known locations. In this case, the set \mathcal{E} of errors to be corrected is the set of all Pauli operators with *support* at the t specified locations (each \mathbf{E}_a acts trivially on the other n-t qubits). But then, for each \mathbf{E}_a and \mathbf{E}_b in \mathcal{E} , the product $\mathbf{E}_a^{\dagger}\mathbf{E}_b$ also has weight at most t. Therefore, the error correction criterion is satisfied for all $\mathbf{E}_{a,b} \in \mathcal{E}$, provided the code has distance at least t + 1.

In particular, a QECC that corrects t errors in arbitrary locations can correct 2t errors in known locations.

7.3.3 Error detection

In some cases we may be satisfied to detect whether an error has occurred, even if we are unable to fully diagnose or reverse the error. A measurement designed for error detection has two possible outcomes: "good" and "bad." If the good outcome occurs, we are assured that the quantum state is undamaged. If the bad outcome occurs, damage has been sustained, and the state should be discarded.

If the error superoperator has its support on the set \mathcal{E} of all Pauli operators of weight up to t, and it is possible to make a measurement that correctly diagnoses *whether* an error has occurred, then it is said that we can detect t errors. Error detection is easier than error correction, so a given code can detect more errors than it can correct. In fact, a QECC with distance d = t + 1 can detect t errors.

Such a code has the property that

$$\langle \bar{i} | \boldsymbol{E}_a | \bar{j} \rangle = C_a \delta_{ij} \tag{7.36}$$

for every Pauli operator E_a of weight t or less, or

$$\boldsymbol{E}_{a}|\bar{i}\rangle = C_{a}|\bar{i}\rangle + |\varphi_{ai}^{\perp}\rangle , \qquad (7.37)$$

where $|\varphi_{ai}^{\perp}\rangle$ is an unnormalized vector orthogonal to the code subspace. Therefore, the action on a state $|\psi\rangle$ in the code subspace of an error superoperator with support on \mathcal{E} is

$$|\psi\rangle \otimes |0\rangle_E \to \sum_{\boldsymbol{E}_a \in \mathcal{E}} \boldsymbol{E}_a |\psi\rangle \otimes |e_a\rangle_E = |\psi\rangle \otimes \left(\sum_{\boldsymbol{E}_a \in \mathcal{E}} C_a |e_a\rangle_E\right) + |\text{orthog}\rangle,$$
(7.38)

where $|orthog\rangle$ denotes a vector orthogonal to the code subspace.

Now we can perform a "fuzzy" orthogonal measurement on the data, with two outcomes: the state is projected onto either the code subspace or the complementary subspace. If the first outcome is obtained, the undamaged state $|\psi\rangle$ is recovered. If the second outcome is found, an error has been detected. We conclude that our QECC with distance d can detect d - 1errors. In particular, then, a QECC that can correct t errors can detect 2terrors.

7.3.4 Quantum codes and entanglement

A QECC protects quantum information from error by encoding it *nonlo-cally*, that is, by sharing it among many qubits in a block. Thus a quantum codeword is a highly entangled state.

In fact, a distance d = t+1 nondegenerate code has the following property: Choose any state $|\psi\rangle$ in the code subspace and any t qubits in the block. Trace over the remaining n - t qubits to obtain

$$\boldsymbol{\rho}^{(t)} = \operatorname{tr}_{(n-t)} |\psi\rangle \langle \psi| , \qquad (7.39)$$

the density matrix of the t qubits. Then this density matrix is totally random:

$$\boldsymbol{\rho}^{(t)} = \frac{1}{2^t} \boldsymbol{I}; \tag{7.40}$$

(In any distance-(t + 1) code, we cannot acquire any information about the encoded data by observing any t qubits in the block; that is, $\rho^{(t)}$ is a constant, independent of the codeword. But only if the code is nondegenerate will the density matrix of the t qubits be a multiple of the identity.)

To verify the property eq. (7.40), we note that for a nondegenerate distance-(t+1) code,

$$\langle \bar{i} | \boldsymbol{E}_a | \bar{j} \rangle = 0 \tag{7.41}$$

for any \boldsymbol{E}_a of nonzero weight up to t, so that

$$\operatorname{tr}(\boldsymbol{\rho}^{(t)}\boldsymbol{E}_a) = 0, \tag{7.42}$$

for any t-qubit Pauli operator E_a other than the identity. Now $\rho^{(t)}$, like any Hermitian $2^t \times 2^t$ matrix, can be expanded in terms of Pauli operators:

$$\boldsymbol{\rho}^{(t)} = \left(\frac{1}{2^t}\right) \boldsymbol{I} + \sum_{\boldsymbol{E}_a \neq \boldsymbol{I}} \rho_a \boldsymbol{E}_a . \qquad (7.43)$$

Since the E_a 's satisfy

$$\left(\frac{1}{2^t}\right)\operatorname{tr}(\boldsymbol{E}_a\boldsymbol{E}_b) = \delta_{ab} , \qquad (7.44)$$

we find that each $\rho_a = 0$, and we conclude that $\boldsymbol{\rho}^{(t)}$ is a multiple of the identity.

7.4 Probability of Failure

7.4.1 Fidelity bound

If the support of the error superoperator contains only the Pauli operators in the set \mathcal{E} that we know how to correct, then we can recover the encoded quantum information with perfect fidelity. But in a realistic error model, there will be a small but nonzero amplitude for errors that are not in \mathcal{E} , so that the recovered state will not be perfect. What can we say about the fidelity of the recovered state?

The Pauli operator expansion of the error superoperator can be divided into a sum over the "good" operators (those in \mathcal{E}), and the "bad" ones (those not in \mathcal{E}), so that it acts on a state $|\psi\rangle$ in the code subspace according to

$$\begin{split} |\psi\rangle \otimes |0\rangle_E &\to \sum_a \boldsymbol{E}_a |\psi\rangle \otimes |e_a\rangle_E \\ \equiv & \sum_{\boldsymbol{E}_a \in \mathcal{E}} \boldsymbol{E}_a |\psi\rangle \otimes |e_a\rangle_E + \sum_{\boldsymbol{E}_b \notin \mathcal{E}} \boldsymbol{E}_b |\psi\rangle \otimes |e_b\rangle_E \\ \equiv & |\text{GOOD}\rangle + |\text{BAD}\rangle . \end{split}$$
(7.45)

The recovery operation (a unitary acting on the data and the ancilla) then maps $|\text{GOOD}\rangle$ to a state $|\text{GOOD}'\rangle$ of data, environment, and ancilla, and $|\text{BAD}\rangle$ to a state $|\text{BAD}'\rangle$, so that after recovery we obtain the state

$$|\text{GOOD}'\rangle + |\text{BAD}'\rangle;$$
 (7.46)

here (since recovery works perfectly acting on the good state)

$$|\text{GOOD}'\rangle = |\psi\rangle \otimes |s\rangle_{EA}$$
, (7.47)

where $|s\rangle_{EA}$ is some state of the environment and ancilla.

Suppose that the states $|\text{GOOD}\rangle$ and $|\text{BAD}\rangle$ are orthogonal. This would hold if, in particular, all of the "good" states of the environment are orthogonal to all of the "bad" states; that is, if

$$\langle e_a | e_b \rangle = 0 \quad \text{for} \quad \boldsymbol{E}_a \in \mathcal{E}, \ \boldsymbol{E}_b \notin \mathcal{E}.$$
 (7.48)

Let $\rho_{\rm rec}$ denote the density matrix of the recovered state, obtained by tracing out the environment and ancilla, and let

$$F = \langle \psi | \boldsymbol{\rho}_{\text{rec}} | \psi \rangle \tag{7.49}$$

be its fidelity. Now, since $|BAD'\rangle$ is orthogonal to $|GOOD'\rangle$ (that is, $|BAD'\rangle$ has no component along $|\psi\rangle|s\rangle_{EA}$), the fidelity will be

$$F = \langle \psi | \boldsymbol{\rho}_{\text{GOOD}'} | \psi \rangle + \langle \psi | \boldsymbol{\rho}_{\text{BAD}'} | \psi \rangle , \qquad (7.50)$$

where

$$\boldsymbol{\rho}_{\text{GOOD}'} = \operatorname{tr}_{EA}(|\text{GOOD}'\rangle\langle \text{GOOD}'|) , \quad \boldsymbol{\rho}_{\text{BAD}'} = \operatorname{tr}_{EA}(|\text{BAD}'\rangle\langle \text{BAD}'|) .$$
(7.51)

The fidelity of the recovered state therefore satisfies

$$F \ge \langle \psi | \boldsymbol{\rho}_{\text{GOOD}'} | \psi \rangle = || |s\rangle_{EA} ||^2 = || |\text{GOOD}' \rangle ||^2 .$$
(7.52)

Furthermore, since the recovery operation is unitary, we have $|| |GOOD' \rangle || = || |GOOD \rangle ||$, and hence

$$F \ge \| |\text{GOOD}\rangle \|^2 = \| \sum_{\boldsymbol{E}_a \in \mathcal{E}} \boldsymbol{E}_a |\psi\rangle \otimes |e_a\rangle_E \|^2 \quad . \tag{7.53}$$

In general, though, $|BAD\rangle$ need not be orthogonal to $|GOOD\rangle$, so that $|BAD'\rangle$ need not be orthogonal to $|GOOD'\rangle$. Then $|BAD'\rangle$ might have a component along $|GOOD'\rangle$ that interferes destructively with $|GOOD'\rangle$ and so reduces the fidelity. We can still obtain a lower bound on the fidelity in this more general case by resolving $|BAD'\rangle$ into a component along $|GOOD'\rangle$ and an orthogonal component, as

$$|BAD'\rangle = |BAD'_{\parallel}\rangle + |BAD'_{\perp}\rangle \tag{7.54}$$

Then reasoning just as above we obtain

$$F \ge \| |\text{GOOD}'\rangle + |\text{BAD}'_{\parallel}\rangle \|^2 \tag{7.55}$$

Of course, since both the error operation and the recovery operation are unitary acting on data, environment, and ancilla, the complete state $|\text{GOOD}'\rangle + |\text{BAD}'\rangle$ is normalized, or

$$\| |\text{GOOD}'\rangle + |\text{BAD}'_{\parallel}\rangle \|^2 + \| |\text{BAD}'_{\perp}\rangle \|^2 = 1$$
, (7.56)

and eq. (7.55) becomes

$$F \ge 1 - \| |\text{BAD}'_{\perp} \rangle \|^2$$
 . (7.57)

Finally, the norm of $|BAD'_{\perp}\rangle$ cannot exceed the norm of $|BAD'\rangle$, and we conclude that

$$1 - F \leq \| |\text{BAD}'\rangle \|^2 = \| |\text{BAD}\rangle \|^2 \equiv \| \sum_{\boldsymbol{E}_b \notin \mathcal{E}} \boldsymbol{E}_b |\psi\rangle \otimes |e_b\rangle_E \|^2 .$$
(7.58)

This is our general bound on the "failure probability" of the recovery operation. The result eq. (7.53) then follows in the special case where $|\text{GOOD}\rangle$ and $|\text{BAD}\rangle$ are orthogonal states.

7.4.2 Uncorrelated errors

Let's now consider some implications of these results for the case where errors acting on distinct qubits are completely uncorrelated. In that case, the error superoperator is a tensor product of single-qubit superoperators. If in fact the errors act on all the qubits in the same way, we can express the n-qubit superoperator as

$$\$_{\text{error}}^{(n)} = \left[\$_{\text{error}}^{(1)}\right]^{\otimes n} , \qquad (7.59)$$

where $\$_{\text{error}}^{(1)}$ is a one-qubit superoperator whose action (in its unitary representation) has the form

$$\begin{aligned} |\psi\rangle \otimes |0\rangle_E \to |\psi\rangle \otimes |e_I\rangle_E + \boldsymbol{X}|\psi\rangle \otimes |e_X\rangle_E + \boldsymbol{Y}|\psi\rangle \otimes |e_Y\rangle_E \\ + \boldsymbol{Z}|\psi\rangle \otimes |e_Z\rangle_E . \end{aligned}$$
(7.60)

The effect of the errors on encoded information is especially easy to analyze if we suppose further that each of the three states of the environment $|e_{X,Y,Z}\rangle$ is orthogonal to the state $|e_I\rangle$. In that case, a record of whether or not an error occurred for each qubit is permanently imprinted on the environment, and it is sensible to speak of a probability of error p_{error} for each qubit, where

$$\langle e_I | e_I \rangle = 1 - p_{\text{error}} . \tag{7.61}$$

If our quantum code can correct t errors, then the "good" Pauli operators have weight up to t, and the "bad" Pauli operators have weight greater than t; recovery is certain to succeed unless at least t + 1 qubits are subjected to errors. It follows that the fidelity obeys the bound

$$1 - F \leq \sum_{s=t+1}^{n} {n \choose s} p_{\text{error}}^{s} \left(1 - p_{\text{error}}\right)^{n-s} \leq {n \choose t+1} p_{\text{error}}^{t+1}.$$
(7.62)

(For each of the $\binom{n}{t+1}$ ways of choosing t+1 locations, the probability that errors occurs at every one of those locations is p_{error}^{t+1} , where we disregard whether additional errors occur at the remaining n-t-1 locations. Therefore, the final expression in eq. (7.62) is an upper bound on the probability that at least t+1 errors occur in the block of n qubits.) For p_{error} small and tlarge, the fidelity of the encoded data is a substantial improvement over the fidelity F = 1 - O(p) maintained by an unprotected qubit.

For a general error superoperator acting on a single qubit, there is no clear notion of an "error probability;" the state of the qubit and its environment obtained when the Pauli operator \boldsymbol{I} acts is not orthogonal to (and so cannot be perfectly distinguished from) the state obtained when the Pauli operators $\boldsymbol{X}, \boldsymbol{Y}$, and \boldsymbol{Z} act. In the extreme case there is no decoherence at all — the "errors" arise because unknown unitary transformations act on the qubits. (If the unitary transformation \boldsymbol{U} acting on a qubit were known, we could recover from the "error" simply by applying \boldsymbol{U}^{\dagger} .)

Consider uncorrelated unitary errors acting on the n qubits in the code block, each of the form (up to an irrelevant phase)

$$\boldsymbol{U}^{(1)} = \sqrt{1-p} + i\sqrt{p} \ \boldsymbol{W},\tag{7.63}$$

where \boldsymbol{W} is a (traceless, Hermitian) linear combination of \boldsymbol{X} , \boldsymbol{Y} , and \boldsymbol{Z} , satisfying $\boldsymbol{W}^2 = \boldsymbol{I}$. If the state $|\psi\rangle$ of the qubit is prepared, and then the unitary error eq. (7.63) occurs, the fidelity of the resulting state is

$$F = \left| \langle \psi | \boldsymbol{U}^{(1)} | \psi \rangle \right|^2 = 1 - p \left(1 - \left(\langle \psi | \boldsymbol{W} | \psi \rangle \right)^2 \right) \geq 1 - p .$$
(7.64)

If a unitary error of the form eq. (7.63) acts on each of the *n* qubits in the code block, and the resulting state is expanded in terms of Pauli operators as in eq. (7.45), then the state $|BAD\rangle$ (which arises from terms in which W acts on at least t + 1 qubits) has a norm of order $(\sqrt{p})^{t+1}$, and eq. (7.58) becomes

$$1 - F = O(p^{t+1}) . (7.65)$$

We see that coding provides an improvement in fidelity of the same order irrespective of whether the uncorrelated errors are due to decoherence or due to unknown unitary transformations.

7.5. CLASSICAL LINEAR CODES

To avoid confusion, let us emphasize the meaning of "uncorrelated" for the purpose of the above discussion. We consider a unitary error acting on n qubits to be "uncorrelated" if it is a tensor product of single-qubit unitary transformations, irrespective of how the unitaries acting on distinct qubits might be related to one another. For example, an "error" whereby all qubits rotate by an angle θ about a common axis is effectively dealt with by quantum error correction; after recovery the fidelity will be $F = 1 - O(\theta^{2(t+1)})$, if the code can protect against t uncorrelated errors. In contrast, a unitary error that would cause more trouble is one of the form $U^{(n)} \sim 1 + i\theta E_{\text{bad}}^{(n)}$, where $E_{\text{bad}}^{(n)}$ is an n-qubit Pauli operator whose weight is greater than t. Then $|\text{BAD}\rangle$ has a norm of order θ , and the typical fidelity after recovery will be $F = 1 - O(\theta^2)$.

7.5 Classical Linear Codes

Quantum error-correcting codes were first invented less than four years ago, but classical error-correcting codes have a much longer history. Over the past fifty years, a remarkably beautiful and powerful theory of classical coding has been erected. Much of this theory can be exploited in the construction of QECC's. Here we will quickly review just a few elements of the classical theory, confining our attention to binary linear codes.

In a binary code, k bits are encoded in a binary string of length n. That is, from among the 2^n strings of length n, we designate a subset containing 2^k strings – the codewords. A k-bit message is encoded by selecting one of these 2^k codewords.

In the special case of a binary linear code, the codewords form a kdimensional closed linear subspace C of the binary vector space F_2^n . That is, the bitwise XOR of two codewords is another codeword. The space C of the code is spanned by a basis of k vectors v_1, v_2, \ldots, v_k ; an arbitrary codeword may be expressed as a linear combination of these basis vectors:

$$v(\alpha_1, \dots, \alpha_k) = \sum_i \alpha_i v_i , \qquad (7.66)$$

where each $\alpha_i \in \{0, 1\}$, and addition is modulo 2. We may say that the length-*n* vector $v(\alpha_1 \dots \alpha_k)$ encodes the *k*-bit message $\alpha = (\alpha_1, \dots, \alpha_k)$.

The k basis vectors $v_1, \ldots v_k$ may be assembled into a $k \times n$ matrix

$$G = \begin{pmatrix} v_1 \\ \vdots \\ v_k \end{pmatrix}, \tag{7.67}$$

called the *generator matrix* of the code. Then in matrix notation, eq. (7.66) can be rewritten as

$$v(\alpha) = \alpha G ; \qquad (7.68)$$

the matrix G, acting to the left, encodes the message α .

An alternative way to characterize the k-dimensional code subspace of F_2^n is to specify n - k linear constraints. There is an $(n - k) \times n$ matrix H such that

$$Hv = 0 \tag{7.69}$$

for all those and only those vectors v in the code C. This matrix H is called the parity check matrix of the code C. The rows of H are n - k linearly independent vectors, and the code space is the space of vectors that are *orthogonal* to all of these vectors. Orthogonality is defined with respect to the mod 2 bitwise inner product; two length-n binary strings are orthogonal is they "collide" (both take the value 1) at an even number of locations. Note that

$$HG^T = 0 ;$$
 (7.70)

where G^T is the transpose of G; the rows of G are orthogonal to the rows of H.

For a classical bit, the only kind of error is a bit flip. An error occurring in an *n*-bit string can be characterized by an *n*-component vector e, where the 1's in e mark the locations where errors occur. When afflicted by the error e, the string v becomes

$$v \to v + e \ . \tag{7.71}$$

Errors can be detected by applying the parity check matrix. If v is a codeword, then

$$H(v+e) = Hv + He = He$$
. (7.72)

7.5. CLASSICAL LINEAR CODES

He is called the syndrome of the error e. Denote by \mathcal{E} the set of errors $\{e_i\}$ that we wish to be able to correct. Error recovery will be possible if and only if all errors e_i have distinct syndromes. If this is the case, we can unambiguously diagnose the error given the syndrome He, and we may then recover by flipping the bits specified by e as in

$$v + e \to (v + e) + e = v$$
. (7.73)

On the other hand, if $He_1 = He_2$ for $e_1 \neq e_2$ then we may misinterpret an e_1 error as an e_2 error; our attempt at recovery then has the effect

$$v + e_1 \to v + (e_1 + e_2) \neq v.$$
 (7.74)

The recovered message $v + e_1 + e_2$ lies in the code, but it differs from the intended message v; the encoded information has been damaged.

The distance d of a code C is the minimum weight of any vector $v \in C$, where the weight is the number of 1's in the string v. A linear code with distance d = 2t + 1 can correct t errors; the code assigns a distinct syndrome to each $e \in \mathcal{E}$, where \mathcal{E} contains all vectors of weight t or less. This is so because, if $He_1 = He_2$, then

$$0 = He_1 + He_2 = H(e_1 + e_2) , \qquad (7.75)$$

and therefore $e_1 + e_2 \in C$. But if e_1 and e_2 are unequal and each has weight no larger than t, then the weight of $e_1 + e_2$ is greater than zero and no larger than 2t. Since d = 2t + 1, there is no such vector in C. Hence He_1 and He_2 cannot be equal.

A useful concept in classical coding theory is that of the dual code. We have seen that the $k \times n$ generator matrix G and the $(n-k) \times n$ parity check matrix H of a code C are related by $HG^T = 0$. Taking the transpose, it follows that $GH^T = 0$. Thus we may regard H^T as the generator and G as the parity check of an (n-k)-dimensional code, which is denoted C^{\perp} and called the dual of C. In other words, C^{\perp} is the orthogonal complement of C in F_2^n . A vector is self-orthogonal if it has even weight, so it is possible for C and C^{\perp} to intersect. A code contains its dual if all of its codewords have even weight and are mutually orthogonal. If n = 2k it is possible that $C = C^{\perp}$, in which case C is said to be self-dual.

An identity relating the code C and its dual C^{\perp} will prove useful in the

following section:

$$\sum_{v \in C} (-1)^{v \cdot u} = \begin{cases} 2^k & u \in C^{\perp} \\ 0 & u \notin C^{\perp} \end{cases}.$$
 (7.76)

The nontrivial content of the identity is the statement that the sum vanishes for $u \notin C^{\perp}$. This readily follows from the familiar identity

$$\sum_{v \in \{0,1\}^k} (-1)^{v \cdot w} = 0, \ w \neq 0, \tag{7.77}$$

where v and w are strings of length k. We can express $v \in G$ as

$$v = \alpha G, \tag{7.78}$$

where α is a k-vector. Then

$$\sum_{v \in C} (-1)^{v \cdot u} = \sum_{\alpha \in \{0,1\}^k} (-1)^{\alpha \cdot Gu} = 0,$$
(7.79)

for $Gu \neq 0$. Since G, the generator matrix of C, is the parity check matrix for C^{\perp} , we conclude that the sum vanishes for $u \notin C^{\perp}$.

7.6 CSS Codes

Principles from the theory of classical linear codes can be adapted to the construction of quantum error-correcting codes. We will describe here a family of QECC's, the Calderbank–Shor–Steane (or CSS) codes, that exploit the concept of a dual code.

Let C_1 be a classical linear code with $(n-k_1) \times n$ parity check matrix H_1 , and let C_2 be a *subcode* of C_1 , with $(n-k_2) \times n$ parity check H_2 , where $k_2 < k_1$. The first $n - k_1$ rows of H_2 coincide with those of H_1 , but there are $k_1 - k_2$ additional linearly independent rows; thus each word in C_2 is contained in C_1 , but the words in C_2 also obey some additional linear constraints.

The subcode C_2 defines an equivalence relation in C_1 ; we say that $u, v \in C_1$ are equivalent $(u \equiv v)$ if and only if there is a w in C_2 such that u = v + w. The equivalence classes are the *cosets* of C_2 in C_1 .

7.6. CSS CODES

A CSS code is a $k = k_1 - k_2$ quantum code that associates a codeword with each equivalence class. Each element of a basis for the code subspace can be expressed as

$$\left|\bar{w}\right\rangle = \frac{1}{\sqrt{2^{k_2}}} \sum_{v \in C_2} \left|v + w\right\rangle \,,\tag{7.80}$$

an equally weighted superposition of all the words in the coset represented by w. There are $2^{k_1-k_2}$ cosets, and hence $2^{k_1-k_2}$ linearly independent codewords. The states $|\bar{w}\rangle$ are evidently normalized and mutually orthogonal; that is, $\langle \bar{w} | \bar{w}' \rangle = 0$ if w and w' belong to different cosets.

Now consider what happens to the codeword $|\bar{w}\rangle$ if we apply the bitwise Hadamard transform $\boldsymbol{H}^{(n)}$:

$$\boldsymbol{H}^{(n)}: \quad |\bar{w}\rangle_{F} \equiv \frac{1}{\sqrt{2^{k_{2}}}} \sum_{v \in C_{2}} |v+w\rangle$$

$$\rightarrow \quad |\bar{w}\rangle_{P} \equiv \frac{1}{\sqrt{2^{n}}} \sum_{u} \frac{1}{\sqrt{2^{k_{2}}}} \sum_{v \in C_{2}} (-1)^{u \cdot v} (-1)^{u \cdot w} |u\rangle$$

$$= \frac{1}{\sqrt{2^{n-k_{2}}}} \sum_{u \in C_{2}^{\perp}} (-1)^{u \cdot w} |u\rangle ; \qquad (7.81)$$

we obtain a coherent superposition, weighted by phases, of words in the dual code C_2^{\perp} (in the last step we have used the identity eq. (7.76)). It is again manifest in this last expression that the codeword depends only on the C_2 coset that w represents — shifting w by an element of C_2 has no effect on $(-1)^{w \cdot w}$ if u is in the code dual to C_2 .

Now suppose that the code C_1 has distance d_1 and the code C_2^{\perp} has distance d_2^{\perp} , such that

$$d_1 \ge 2t_F + 1$$
,
 $d_2^{\perp} \ge 2t_P + 1$. (7.82)

Then we can see that the corresponding CSS code can correct t_F bit flips and t_P phase flips. If e is a binary string of length n, let $\boldsymbol{E}_e^{(\text{flip})}$ denote the Pauli operator with an \boldsymbol{X} acting at each location i where $e_i = 1$; it acts on the state $|v\rangle$ according to

$$\boldsymbol{E}_{e}^{(\text{flip})}: |v\rangle \to |v+e\rangle$$
 . (7.83)

And let $E_e^{(\text{phase})}$ denote the Pauli operator with a Z acting where $e_i = 1$; its action is

$$\boldsymbol{E}_{e}^{(\text{phase})}:|v\rangle \to (-1)^{v.e}|v\rangle , \qquad (7.84)$$

which in the Hadamard rotated basis becomes

$$E_e^{(\text{phase})} : |u\rangle \to |u+e\rangle .$$
 (7.85)

Now, in the original basis (the F or "flip" basis), each basis state $|\bar{w}\rangle_F$ of the CSS code is a superposition of words in the code C_1 . To diagnose bit flip error, we perform on data and ancilla the unitary transformation

$$|v\rangle \otimes |0\rangle_A \to |v\rangle \otimes |H_1v\rangle_A , \qquad (7.86)$$

and then measure the ancilla. The measurement result H_1e_F is the *bit flip* syndrome. If the number of flips is t_F or fewer, we may correctly infer from this syndrome that bit flips have occurred at the locations labeled by e_F . We recover by applying X to the qubits at those locations.

To correct phase errors, we first perform the bitwise Hadamard transformation to rotate from the F basis to the P ("phase") basis. In the P basis, each basis state $|\bar{w}\rangle_P$ of the CSS code is a superposition of words in the code C_2^{\perp} . To diagnose phase errors, we perform a unitary transformation

$$|v\rangle \otimes |0\rangle_A \to |v\rangle \otimes |G_2 v\rangle_A , \qquad (7.87)$$

and measure the ancilla $(G_2$, the generator matrix of C_2 , is also the parity check matrix of C_2^{\perp}). The measurement result G_2e_P is the *phase error syn*drome. If the number of phase errors is t_P or fewer, we may correctly infer from this syndrome that phase errors have occurred at locations labeled by e_P . We recover by applying \boldsymbol{X} (in the P basis) to the qubits at those locations. Finally, we apply the bitwise Hadamard transformation once more to rotate the codewords back to the original basis. (Equivalently, we may recover from the phase errors by applying \boldsymbol{Z} to the affected qubits after the rotation back to the F basis.)

If e_F has weight less than d_1 and e_P has weight less than d_2^{\perp} , then

$$\langle \bar{w} | \boldsymbol{E}_{e_P}^{(\text{phase})} \boldsymbol{E}_{e_F}^{(\text{flip})} | \bar{w}' \rangle = 0$$
(7.88)

(unless $e_F = e_P = 0$). Any Pauli operator can be expressed as a product of a phase operator and a flip operator — a \boldsymbol{Y} error is merely a bit flip and

phase error both afflicting the same qubit. So the distance d of a CSS code satisfies

$$d \geq \min(d_1, d_2^{\perp}) . \tag{7.89}$$

CSS codes have the special property (not shared by more general QECC's) that the recovery procedure can be divided into two separate operations, one to correct the bit flips and the other to correct the phase errors.

The unitary transformation eq. (7.86) (or eq. (7.87)) can be implemented by executing a simple quantum circuit. Associated with each of the $n - k_1$ rows of the parity check matrix H_1 is a bit of the syndrome to be extracted. To find the *a*th bit of the syndrome, we prepare an ancilla bit in the state $|0\rangle_{A,a}$, and for each value of λ with $(H_1)_{a\lambda} = 1$, we execute a controlled-NOT gate with the ancilla bit as the target and qubit λ in the data block as the control. When measured, the ancilla qubit reveals the value of the parity check bit $\sum_{\lambda} (H_1)_{a\lambda} v_{\lambda}$.

Schematically, the full error correction circuit for a CSS code has the form:

Separate syndromes are measured to diagnose the bit flip errors and the phase errors. An important special case of the CSS construction arises when a code C contains its dual C^{\perp} . Then we may choose $C_1 = C$ and $C_2 = C^{\perp} \subseteq C$; the C parity check is computed in both the F basis and the P basis to determine the two syndromes.

7.7 The 7-Qubit Code

The simplest of the CSS codes is the [[n, k, d]] = [7, 1, 3] quantum code first formulated by Andrew Steane. It is constructed from the classical 7-bit Hamming code.

The Hamming code is an [n, k, d] = [7, 4, 3] classical code with the 3×7

parity check matrix

$$H = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}.$$
 (7.90)

To see that the distance of the code is d = 3, first note that the weight-3 string (1110000) passes the parity check and is, therefore, in the code. Now we need to show that there are no vectors of weight 1 or 2 in the code. If e_1 has weight 1, then He_1 is one of the columns of H. But no column of H is trivial (all zeros), so e_1 cannot be in the code. Any vector of weight 2 can be expressed as $e_1 + e_2$, where e_1 and e_2 are distinct vectors of weight 1. But

$$H(e_1 + e_2) = He_1 + He_2 \neq 0, \tag{7.91}$$

because all columns of H are distinct. Therefore $e_1 + e_2$ cannot be in the code.

The rows of H themselves pass the parity check, and so are also in the code. (Contrary to one's usual linear algebra intuition, a nonzero vector over the finite field F_2 can be orthogonal to itself.) The generator matrix G of the Hamming code can be written as

$$G = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 \end{pmatrix} ;$$
(7.92)

the first three rows coincide with the rows of H, and the weight-3 codeword (1110000) is appended as the fourth row.

The dual of the Hamming code is the [7,3,4] code generated by H. In this case the dual of the code is actually contained in the code — in fact, it is the *even subcode* of the Hamming code, containing all those and only those Hamming codewords that have even weight. The odd codeword (1110000) is a representative of the nontrivial coset of the even subcode. For the CSS construction, we will choose C_1 to be the Hamming code, and C_2 to be its dual, the even subcode. Therefore, $C_2^{\perp} = C_1$ is again the Hamming code; we will use the Hamming parity check both to detect bit flips in the F basis and to detect phase flips in the P basis.

7.7. THE 7-QUBIT CODE

In the F basis, the two orthonormal codewords of this CSS code, each associated with a distinct cos t of the even subcode, can be expressed as

$$|\bar{0}\rangle_{F} = \frac{1}{\sqrt{8}} \sum_{\substack{\text{even } v \\ \in \text{ Hamming}}} |v\rangle ,$$

$$|\bar{1}\rangle_{F} = \frac{1}{\sqrt{8}} \sum_{\substack{\text{odd } v \\ \in \text{ Hamming}}} |v\rangle .$$
(7.93)

Since both $|\bar{0}\rangle$ and $|\bar{1}\rangle$ are superpositions of Hamming codewords, bit flips can be diagnosed in this basis by performing an H parity check. In the Hadamard rotated basis, these codewords become

$$\boldsymbol{H}^{(7)}: |\bar{0}\rangle_{F} \to |\bar{0}\rangle_{P} \equiv \left(\frac{1}{4}\right) \sum_{v \in \text{ Hamming}} |v\rangle = \frac{1}{\sqrt{2}} (|\bar{0}\rangle_{F} + |\bar{1}\rangle_{F}) |\bar{1}\rangle_{F} \to |\bar{1}\rangle_{P} \equiv \left(\frac{1}{4}\right) \sum_{v \in \text{ Hamming}} (-1)^{\text{wt}(v)} |v\rangle = \frac{1}{\sqrt{2}} (|\bar{0}\rangle_{F} - |\bar{1}\rangle_{F}).$$
(7.94)

In this basis as well, the states are superpositions of Hamming codewords, so that bit flips in the P basis (phase flips in the original basis) can again be diagnosed with an H parity check. (We note in passing that for this code, performing the bitwise Hadamard transformation also implements a Hadamard rotation on the encoded data, a point that will be relevant to our discussion of fault-tolerant quantum computation in the next chapter.)

Steane's quantum code can correct a single bit flip and a single phase flip on any one of the seven qubits in the block. But recovery will fail if two different qubits both undergo either bit flips or phase flips. If e_1 and e_2 are two distinct weight-one strings then $He_1 + He_2$ is a sum of two distinct columns of H, and hence a third column of H (all seven of the nontrivial strings of length 3 appear as columns of H.) Therefore, there is another weight-one string e_3 such that $He_1 + He_2 = He_3$, or

$$H(e_1 + e_2 + e_3) = 0 ; (7.95)$$

thus $e_1 + e_2 + e_3$ is a weight-3 word in the Hamming code. We will interpret the syndrome He_3 as an indication that the error $v \to v + e_3$ has arisen, and we will attempt to recover by applying the operation $v \to v + e_3$. Altogether then, the effect of the two bit flip errors and our faulty attempt at recovery will be to add $e_1 + e_2 + e_3$ (an odd-weight Hamming codeword) to the data, which will induce a flip of the *encoded* qubit

$$|\bar{0}\rangle_F \leftrightarrow |\bar{1}\rangle_F.$$
 (7.96)

Similarly, two phase flips in the F basis are two bit flips in the P basis, which (after the botched recovery) induce on the encoded qubit

$$|\bar{0}\rangle_P \leftrightarrow |\bar{1}\rangle_P,$$
 (7.97)

or equivalently

$$\begin{split} |\bar{0}\rangle_F &\to |\bar{0}\rangle_F \\ |\bar{1}\rangle_F &\to -|\bar{1}\rangle_F, \end{split} \tag{7.98}$$

a phase flip of the encoded qubit in the F basis. If there is one bit flip and one phase flip (either on the same qubit or different qubits) then recovery will be successful.

7.8 Some Constraints on Code Parameters

Shor's code protects one encoded qubit from an error in any single one of nine qubits in a block, and Steane's code reduces the block size from nine to seven. Can we do better still?

7.8.1 The Quantum Hamming bound

To understand how much better we might do, let's see if we can derive any bounds on the distance d = 2t + 1 of an [[n, k, d]] quantum code, for given n and k. At first, suppose we limit our attention to *nondegenerate* codes, which assign a distinct syndrome to each possible error. On a given qubit, there are three possible linearly independent errors $\boldsymbol{X}, \boldsymbol{Y}$, or \boldsymbol{Z} . In a block of n qubits, there are $\binom{n}{j}$ ways to choose j qubits that are affected by errors, and three possible errors for each of these qubits; therefore the total number of possible errors of weight up to t is

$$N(t) = \sum_{j=0}^{t} 3^{j} \binom{n}{j}.$$
 (7.99)

If there are k encoded qubits, then there are 2^k linearly independent codewords. If all $\mathbf{E}_a|\bar{j}\rangle$'s are linearly independent, where \mathbf{E}_a is any error of weight up to t and $|\bar{i}\rangle$ is any element of a basis for the codewords, then the dimension 2^n of the Hilbert space of n qubits must be large enough to accommodate $N(t) \cdot 2^k$ independent vectors; hence

$$N(t) = \sum_{j=0}^{t} 3^{j} \binom{n}{j} \le 2^{n-k}.$$
(7.100)

This result is called the quantum Hamming bound. An analogous bound applies to classical block codes, but without the factor of 3^{j} , since there is only one type of error (a flip) that can affect a classical bit. We also emphasize that the quantum Hamming bound applies only in the case of nondegenerate coding, while the classical Hamming bound applies in general. However, no degenerate quantum codes that violate the quantum Hamming code have yet been constructed (as of January, 1999).

In the special case of a code with one encoded qubit (k = 1) that corrects one error (t = 1), the quantum Hamming bound becomes

$$1 + 3n \le 2^{n-1},\tag{7.101}$$

which is satisfied for $n \ge 5$. In fact, the case n = 5 saturates the inequality (1 + 15 = 16). A nondegenerate [[5, 1, 3]] quantum code, if it exists, is *perfect*: The entire 32-dimensional Hilbert space of the five qubits is needed to accommodate all possible one-qubit errors acting on all codewords — there is no wasted space.

7.8.2 The no-cloning bound

We could still wonder, though, if there is a degenerate n = 4 code that can correct one error. In fact, it is easy to see that no such code can exist. We already know that a code that corrects t errors at arbitrary locations can also be used to correct 2t errors at known locations. Suppose that we have a [[4,1,3]] quantum code. Then we could encode a single qubit in the fourqubit block, and split the block into two sub-blocks, each containing two qubits.

– Figure –

If we append $|00\rangle$ to each of those two sub-blocks, then the original block has spawned two offspring, each with two located errors. If we were able to correct the two located errors in each of the offspring, we would obtain two identical copies of the parent block — we would have cloned an unknown quantum state, which is impossible. Therefore, no [[4, 1, 3]] quantum code can exist. We conclude that n = 5 is the minimal block size of a quantum code that corrects one error, whether the code is degenerate or not.

The same reasoning shows that an $[[n, k \ge 1, d]]$ code can exist only for

$$n > 2(d-1) . (7.102)$$

7.8.3 The quantum Singleton bound

We will now see that this result eq. (7.102) can be strengthened to

$$n - k \ge 2(d - 1). \tag{7.103}$$

Eq. (7.103) resembles the Singleton bound on classical code parameters,

$$n-k \ge d-1,$$
 (7.104)

and so has been called the "quantum Singleton bound." For a classical *linear* code, the Singleton bound is a near triviality: the code can have distance d only if any d-1 columns of the parity check matrix are linearly independent. Since the columns have length n-k, at most n-k columns can be linearly independent; therefore d-1 cannot exceed n-k. The Singleton bound also applies to nonlinear codes.

An elegant proof of the quantum Singleton bound can be found that exploits the subadditivity of the Von Neumann entropy discussed in §5.2. We begin by introducing a k-qubit ancilla, and constructing a pure state that maximally entangles the ancilla with the 2^k codewords of the QECC:

$$|\Psi\rangle_{AQ} = \frac{1}{\sqrt{2^k}} \sum |x\rangle_A |\bar{x}\rangle_Q , \qquad (7.105)$$

where $\{|x\rangle_A\}$ denotes an orthonormal basis for the 2^k -dimensional Hilbert space of the ancilla, and $\{|\bar{x}\rangle_Q\}$ denotes an orthonormal basis for the 2^k dimensional code subspace. If we trace over the length-*n* code block Q, the density matrix ρ_A of the ancilla is $\frac{1}{2^k}\mathbf{1}$, which has entropy

$$S(A) = k = S(Q).$$
 (7.106)
Now, if the code has distance d, then d-1 located errors can be corrected; or, as we have seen, no observable acting on d-1 of the n qubits can reveal any information about the encoded state. Equivalently, the observable can reveal nothing about the state of the ancilla in the entangled state $|\Psi\rangle$.

Now, since we already know that n > 2(d-1) (if $k \ge 1$), let us imagine dividing the code block Q into three disjoint parts: a set of d-1 qubits $Q_{d-1}^{(1)}$, another disjoint set of d-1 qubits $Q_{d-1}^{(2)}$, and the remaining qubits $Q_{n-2(d-1)}^{(3)}$. If we trace out $Q^{(2)}$ and $Q^{(3)}$, the density matrix we obtain must contain no correlations between $Q^{(1)}$ and the ancilla A. This means that the entropy of system $AQ^{(1)}$ is additive:

$$S(Q^{(2)}Q^{(3)}) = S(AQ^{(1)}) = S(A) + S(Q^{(1)}).$$
(7.107)

Similarly,

$$S(Q^{(1)}Q^{(3)}) = S(AQ^{(2)}) = S(A) + S(Q^{(2)}).$$
(7.108)

Furthermore, in general, Von Neumann entropy is subadditive, so that

$$S(Q^{(1)}Q^{(3)}) \leq S(Q^{(1)}) + S(Q^{(3)})$$

$$S(Q^{(2)}Q^{(3)}) \leq S(Q^{(2)}) + S(Q^{(3)})$$
(7.109)

Combining these inequalities with the equalities above, we find

$$S(A) + S(Q^{(2)}) \leq S(Q^{(1)}) + S(Q^{(3)})$$

$$S(A) + S(Q^{(1)}) \leq S(Q^{(2)}) + S(Q^{(3)}).$$
(7.110)

Both of these inequalities can be simultaneously satisfied only if

$$S(A) \le S(Q^{(3)})$$
 (7.111)

Now $Q^{(3)}$ has dimension n - 2(d - 1), and its entropy is bounded above by its dimension so that

$$S(A) = k \le n - 2(d - 1), \tag{7.112}$$

which is the quantum Singleton bound.

The [[5, 1, 3]] code saturates this bound, but for most values of n and k the bound is not tight. Rains has obtained the stronger result that an [[n, k, 2t + 1]] code with $k \ge 1$ must satisfy

$$t \le \left[\frac{n+1}{6}\right],\tag{7.113}$$

(where [x] = "floor x" is the greatest integer greater than or equal to x. Thus, the minimal length of a k = 1 code that can correct t = 1, 2, 3, 4, 5 errors is n = 5, 11, 17, 23, 29 respectively. Codes with all of these parameters have actually been constructed, except for the [[23, 1, 9]] code.

7.9 Stabilizer Codes

7.9.1 General formulation

We will be able to construct a (nondegenerate) [[5, 1, 3]] quantum code, but to do so, we will need a more powerful procedure for constructing quantum codes than the CSS procedure.

Recall that to establish a criterion for when error recovery is possible, we found it quite useful to expand an error superoperator in terms of the n-qubit Pauli operators. But up until now we have not exploited the group structure of these operators (a product of Pauli operators is a Pauli operator). In fact, we will see that group theory is a powerful tool for constructing QECC's.

For a single qubit, we will find it more convenient now to choose all of the Pauli operators to be represented by real matrices, so I will now use a notation in which \boldsymbol{Y} denotes the anti-hermitian matrix

$$\boldsymbol{Y} = \boldsymbol{Z}\boldsymbol{X} = i\boldsymbol{\sigma}\boldsymbol{y} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \qquad (7.114)$$

satisfying $Y^2 = -I$. Then the operators

$$\{\pm \mathbf{I}, \pm \mathbf{X}, \pm \mathbf{Y}, \pm \mathbf{Z}\} \equiv \pm \{\mathbf{I}, \mathbf{X}, \mathbf{Y}, \mathbf{Z}\},$$
(7.115)

are the elements of a group of order $8.^1$ The *n*-fold tensor products of singlequbit Pauli operators also form a group

$$G_n = \pm \{ \boldsymbol{I}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \}^{\oplus n}, \qquad (7.116)$$

of order $|G_n| = 2^{2n+1}$ (since there are 4^n possible tensor products, and another factor of 2 for the \pm sign) we will refer to G_n as the *n*-qubit Pauli group. (In fact, we will use the term "Pauli group" both to refer to the abstract

¹It is not the quaternionic group but the *other* non-abelian group of order 8 — the symmetry group of the square. The element \mathbf{Y} , of order 4, can be regarded as the 90° rotation of the plane, while \mathbf{X} and \mathbf{Z} are reflections about two orthogonal axes.

group G_n , and to its dimension- 2^n faithful unitary representation by tensor products of 2×2 matrices; its only irreducible representation of dimension greater than 1.) Note that G_n has the two element center $Z_2 = \{\pm I^{\otimes n}\}$. If we quotient out its center, we obtain the group $\bar{G}_n \equiv G_n/Z_2$; this group can also be regarded as a binary vector space of dimension 2^{2n} , a property that we will exploit below.

The (2ⁿ-dimensional representation of the) Pauli group G_n evidently has these properties:

- (i) Each $M \in G_n$ is unitary, $M^{-1} = M^{\dagger}$.
- (ii) For each element $M \in G_n$, $M^2 = \pm I \equiv \pm I^{\otimes n}$. Furthermore, $M^2 = I$ if the number of Y's in the tensor product is even, and $M^2 = -I$ if the number of Y's is odd.
- (iii) If $M^2 = I$, then M is hermitian $(M = M^{\dagger})$; if $M^2 = -I$, then M is anti-hermitian $(M = -M^{\dagger})$.
- (iv) Any two elements $M, N \in G_n$ either commute or anti-commute: $MN = \pm NM$.

We will use the Pauli group to characterize a QECC in the following way: Let S denote an abelian subgroup of the *n*-qubit Pauli group G_n . Thus all elements of S acting on \mathcal{H}_{2^n} can be simultaneously diagonalized. Then the stabilizer code $\mathcal{H}_S \subseteq \mathcal{H}_{2^n}$ associated with S is the simultaneous eigenspace with eigenvalue 1 of all elements of S. That is,

$$|\psi\rangle \in \mathcal{H}_S \quad \text{iff} \quad \boldsymbol{M}|\psi\rangle = |\psi\rangle \text{ for all } \boldsymbol{M} \in S.$$
 (7.117)

The group S is called the *stabilizer* of the code, since it preserves all of the codewords.

The group S can be characterized by its generators. These are elements $\{M_i\}$ that are *independent* (no one can be expressed as a product of others) and such that each element of S can be expressed as a product of elements of $\{M_i\}$. If S has n-k generators, we can show that the code space \mathcal{H}_S has dimension 2^k — there are k encoded qubits.

To verify this, first note that each $M \in S$ must satisfy $M^2 = I$; if $M^2 = -I$, then M cannot have the eigenvalue +1. Furthermore, for each $M \neq \pm I$ in G_n that squares to one, the eigenvalues +1 and -1 have equal

degeneracy. This is because for each $M \neq \pm I$, there is an $N \in G_n$ that anti-commutes with M,

$$\boldsymbol{N}\boldsymbol{M} = -\boldsymbol{M}\boldsymbol{N} \; ; \tag{7.118}$$

therefore, $\boldsymbol{M}|\psi\rangle = |\psi\rangle$ if and only if $\boldsymbol{M}(\boldsymbol{N}|\psi\rangle) = -\boldsymbol{N}|\psi\rangle$, and the action of the unitary \boldsymbol{N} establishes a 1 - 1 correspondence between the +1 eigenstates of \boldsymbol{M} and the -1 eigenstates. Hence there are $\frac{1}{2}(2^n) = 2^{n-1}$ mutually orthogonal states that satisfy

$$\boldsymbol{M}_1|\psi\rangle = |\psi\rangle , \qquad (7.119)$$

where M_1 is one of the generators of S.

Now let M_2 be another element of G_n that commutes with M_1 such that $M_2 \neq \pm I, \pm M_1$. We can find an $N \in G_n$ that commutes with M_1 but anti-commutes with M_2 ; therefore N preserves the +1 eigenspace of M_1 , but within this space, it interchanges the +1 and -1 eigenstates of M_2 . It follows that the space satisfying

$$\boldsymbol{M}_1|\psi\rangle = \boldsymbol{M}_2|\psi\rangle = |\psi\rangle, \qquad (7.120)$$

has dimension 2^{n-2} .

Continuing in this way, we note that if M_j is independent of $\{M_1, M_2, \ldots, M_{j-1}\}$, then there is an N that commutes with M_1, \ldots, M_{j-1} , but anti-commutes with M_j (we'll discuss in more detail below how such an N can be found). Therefore, restricted to the space with $M_1 = M_2 = \ldots = M_{j-1} = 1, M_j$ has as many +1 eigenvectors as -1 eigenvectors. So adding another generator always cuts the dimension of the simultaneous eigenspace in half. With n - k generators, the dimension of the remaining space is $2^n (1/2)^{n-k} = 2^k$.

The stabilizer language is useful because it provides a simple way to characterize the errors that the code can detect and correct. We may think of the n - k stabilizer generators M_1, \ldots, M_{n-k} , as the *check operators* of the code, the collective observables that we measure to diagnose the errors. If the encoded information is undamaged, then we will find $M_i = 1$ for each of the generators; but if $M_i = -1$ for some i, then the data is orthogonal to the code subspace and an error has been detected.

Recall that the error superoperator can be expanded in terms of elements E_a of the Pauli group. A particular E_a either commutes or anti-commutes with a particular stabilizer generator M. If E_a and M commute, then

$$\boldsymbol{M}\boldsymbol{E}_{a}|\psi\rangle = \boldsymbol{E}_{a}\boldsymbol{M}|\psi\rangle = \boldsymbol{E}_{a}|\psi\rangle,$$
 (7.121)

for $|\psi\rangle \in \mathcal{H}_S$, so the error preserves the value M = 1. But if E_a and M anti-commute, then

$$\boldsymbol{M}\boldsymbol{E}_{a}|\psi\rangle = -\boldsymbol{E}_{a}\boldsymbol{M}|\psi\rangle = -\boldsymbol{E}_{a}|\psi\rangle, \qquad (7.122)$$

so that the error flips the value of M, and the error can be detected by measuring M.

For stabilizer generators M_i and errors E_a , we may write

$$M_i E_a = (-1)^{s_{ia}} E_a M_i.$$
 (7.123)

The s_{ia} 's, i = 1, ..., n - k constitute a *syndrome* for the error \mathbf{E}_a , as $(-1)^{s_{ia}}$ will be the result of measuring \mathbf{M}_i if the error \mathbf{E}_a occurs. In the case of a nondegenerate code, the s_{ia} 's will be distinct for all $\mathbf{E}_a \in \mathcal{E}$, so that measuring the n - k stabilizer generators will diagnose the error completely.

More generally, let us find a condition to be satisfied by the stabilizer that is sufficient to ensure that error recovery is possible. Recall that it is sufficient that, for each $E_a, E_b \in \mathcal{E}$, and normalized $|\psi\rangle$ in the code subspace, we have

$$\langle \psi | \boldsymbol{E}_{a}^{\dagger} \boldsymbol{E}_{b} | \psi \rangle = C_{ab}, \qquad (7.124)$$

where C_{ab} is independent of $|\psi\rangle$. We can see that this condition is satisfied provided that, for each $E_a, E_b \in \mathcal{E}$, one of the following holds:

- 1) $E_a^{\dagger} E_b \in S$,
- 2) There is an $M \in S$ that anti-commutes with $E_a^{\dagger} E_b$.
- **Proof:** In case (1) $\langle \psi | \boldsymbol{E}_{a}^{\dagger} \boldsymbol{E}_{b} | \psi \rangle = \langle \psi | \psi \rangle = 1$, for $|\psi \rangle \in \mathcal{H}_{S}$. In case (2), suppose $\boldsymbol{M} \in S$ and $\boldsymbol{M} \boldsymbol{E}_{a}^{\dagger} \boldsymbol{E}_{b} = -\boldsymbol{E}_{a}^{\dagger} \boldsymbol{E}_{b} \boldsymbol{M}$. Then

$$\langle \psi | oldsymbol{E}_a^\dagger oldsymbol{E}_b | \psi
angle = \langle \psi | oldsymbol{E}_a^\dagger oldsymbol{E}_b oldsymbol{M} | \psi
angle$$

$$= -\langle \psi | \boldsymbol{M} \boldsymbol{E}_{a}^{\dagger} \boldsymbol{E}_{b} | \psi \rangle = -\langle \psi | \boldsymbol{E}_{a}^{\dagger} \boldsymbol{E}_{b} | \psi \rangle, \qquad (7.125)$$

and therefore $\langle \psi | \boldsymbol{E}_a^{\dagger} \boldsymbol{E}_b | \psi \rangle = 0.$

Thus, a stabilizer code that corrects $\{\mathcal{E}\}$ is a space \mathcal{H}_S fixed by an abelian subgroup S of the Pauli group, where either (1) or (2) is satisfied by each $\mathbf{E}_a^{\dagger} \mathbf{E}_b$ with $\mathbf{E}_{a,b} \in \mathcal{E}$. The code is nondegenerate if condition (1) is not satisfied for any $\mathbf{E}_a^{\dagger} \mathbf{E}_b$.

Evidently we could also just as well choose the code subspace to be any one of the 2^{n-k} simultaneous eigenspaces of n-k independent commuting elements of G_n . But in fact all of these codes are equivalent. We may regard two stabilizer codes as *equivalent* if they differ only according to how the qubits are labeled, and how the basis for each single-qubit Hilbert space is chosen – that is the stabilizer of one code is transformed to the stabilizer of the other by a permutation of the qubits together with a tensor product of single-qubit transformations. If we partition the stabilizer generators into two sets $\{M_1, \ldots, M_j\}$ and $\{M_{j+1}, \ldots, M_{n-k}\}$, then there exists an $N \in G_n$ that commutes with each member of the first set and anti-commutes with each member of the second set. Applying N to $|\psi\rangle \in \mathcal{H}_s$ preserves the eigenvalues of the first set while flipping the eigenvalues of the second set. Since N is just a tensor product of single-qubit unitary transformations, there is no loss of generality (up to equivalence) in choosing all of the eigenvalues to be one. Furthermore, since minus signs don't really matter when the stabilizer is specified, we may just as well say that two codes are equivalent if, up to phases, the stabilizers differ by a permutation of the n qubits, and permutations on each individual qubits of the operators X, Y, Z.

Recovery may fail if there is an $E_a^{\dagger} E_b$ that *commutes* with the stabilizer but does not lie in the stabilizer. This is an operator that preserves the code subspace \mathcal{H}_S but may act nontrivially in that space; thus it can modify encoded information. Since $E_a |\psi\rangle$ and $E_b |\psi\rangle$ have the same syndrome, we might mistakenly interpret an E_a error as an E_b error; the effect of the error together with the attempt at recovery is that $E_b^{\dagger} E_a$ gets applied to the data, which can cause damage.

A stabilizer code with distance d has the property that each $E \in G_n$ of weight less than d either lies in the stabilizer or anti-commutes with some element of the stabilizer. The code is nondegenerate if the stabilizer contains no elements of weight less than d. A distance d = 2t + 1 code can correct t errors, and a distance s + 1 code can detect s errors or correct s errors at known locations.

7.9.2 Symplectic Notation

Properties of stabilizer codes are often best explained and expressed using the language of linear algebra. The stabilizer S of the code, an order 2^{n-k} abelian subgroup of the Pauli group with all elements squaring to the identity, can equivalently be regarded as a dimension n-k closed linear subspace of F_2^{2n} , self orthogonal with respect to a certain (symplectic) inner product.

The group $\bar{G}_n = G_n/Z_2$ is isomorphic to the binary vector space F_2^{2n} . We establish this by observing that, since $\boldsymbol{Y} = \boldsymbol{Z}\boldsymbol{X}$, any element \boldsymbol{M} of the Pauli group (up to the \pm sign) can be expressed as a product of \boldsymbol{Z} 's and \boldsymbol{X} 's; we may write

$$\boldsymbol{M} = \boldsymbol{Z}_M \cdot \boldsymbol{X}_M \tag{7.126}$$

where Z_M is a tensor product of Z's and X_M is a tensor product of X's. More explicitly, a Pauli operator may be written as

$$(\alpha|\beta) \equiv \boldsymbol{Z}(\alpha)\boldsymbol{X}(\beta) = \bigotimes_{i=1}^{n} \boldsymbol{Z}^{\alpha_{i}} \cdot \bigotimes_{i=1}^{n} \boldsymbol{X}^{\beta_{i}}, \qquad (7.127)$$

where α and β are binary strings of length n. (Then \mathbf{Y} acts at the locations where α and β "collide.") Multiplication in \overline{G}_n maps to addition in F_2^{2n} :

$$(\alpha|\beta)(\alpha'|\beta') = (-1)^{\alpha'\cdot\beta}(\alpha + \alpha'|\beta + \beta') ; \qquad (7.128)$$

the phase arises because $\alpha' \cdot \beta$ counts the number of times a Z is interchanged with a X as the product is rearranged into the standard form of eq. (7.127).

It follows from eq. (7.128) that the commutation properties of the Pauli operators can be expressed in the form

$$(\alpha|\beta)(\alpha'|\beta') = (-1)^{\alpha \cdot \beta' + \alpha' \cdot \beta}(\alpha'|\beta')(\alpha|\beta)$$
(7.129)

Thus two Pauli operators commute if and only if the corresponding vectors are orthogonal with respect to the "symplectic" inner product

$$\alpha \cdot \beta' + \alpha' \cdot \beta \ . \tag{7.130}$$

We also note that the square of a Pauli operator is

$$(\alpha|\beta)^2 = (-1)^{\alpha \cdot \beta} \boldsymbol{I} , \qquad (7.131)$$

since $\alpha \cdot \beta$ counts the number of **Y**'s in the operator; it squares to the identity if and only if

$$\alpha \cdot \beta = 0 \ . \tag{7.132}$$

Note that a closed subspace, where each element has this property, is automatically self-orthogonal, since

$$\alpha \cdot \beta' + \alpha' \cdot \beta = (\alpha + \alpha') \cdot (\beta + \beta') - \alpha \cdot \beta - \alpha' \cdot \beta' = 0 ;$$
(7.133)

in the group language, that is, a subgroup of G_n with each element squaring to I is automatically abelian.

Using the linear algebra language, some of the statements made earlier about the Pauli group can be easily verified by counting linear constraints. Elements are independent if the corresponding vectors are linearly independent over F_2^{2n} , so we may think of the n-k generators of the stabilizer as a basis for a linear subspace of dimension n - k. We will use the notation S to denote both the linear space and the corresponding abelian group. Then S^{\perp} denotes the dimension-n + k space of vectors that are orthogonal to each vector in S (with respect to the symplectic inner product). Note that S^{\perp} contains S, since all vectors in S are mutually orthogonal. In the group language, corresponding to S^{\perp} is the normalizer (or centralizer) group $N(S) \ (\equiv S^{\perp})$ of S in G_n — the subgroup of G_n containing all elements that commute with each element of S. Since S is abelian, it is contained in its own normalizer, which also contains other elements (to be further discussed below). The stabilizer of a distance d code has the property that each $(\alpha|\beta)$ whose weight $\sum_{i} (\alpha_i \vee \beta_i)$ is less than d either lies in the stabilizer subspace S or lies *outside* the orthogonal space S^{\perp} .

A code can be characterized by its stabilizer, a stabilizer by its generators, and the n - k generators can be represented by an $(n - k) \times 2n$ matrix

$$H = (H_Z | H_X). (7.134)$$

Here each row is a Pauli operator, expressed in the $(\alpha|\beta)$ notation. The syndrome of an error $E_a = (\alpha_a|\beta_a)$ is determined by its commutation properties with the generators $M_i = (\alpha'_i|\beta'_i)$; that is

$$s_{ia} = (\alpha_a | \beta_a) \cdot (\alpha'_i | \beta'_i) = \alpha_a \cdot \beta'_i + \alpha'_i \cdot \beta_a.$$
(7.135)

In the case of a nondegenerate code, each error has a distinct syndrome. If the code is degenerate, there may be several errors with the same syndrome, but we may apply any one of the E_a^{\dagger} corresponding to the observed syndrome in order to recover.

7.9.3 Some examples of stabilizer codes

(a) The nine-qubit code. This [[9, 1, 3]] code has eight stabilizer generators that can be expressed as

$$Z_1Z_2, \quad Z_2Z_3 \quad Z_4Z_5 \quad Z_5Z_6, \quad Z_7Z_8 \quad Z_8Z_9$$

 $X_1X_2X_3X_4X_5X_6, \quad X_4X_5X_6X_7X_8X_9.$ (7.136)

In the notation of eq. (7.134) these become

$\left(\begin{array}{rrrr} 1 & 1 & 0 \\ 0 & 1 & 1 \end{array}\right)$	0	0	
0	$\begin{array}{cccc} 1 & 1 & 0 \\ 0 & 1 & 1 \end{array}$	0	0
0	0	$\begin{array}{cccc} 1 & 1 & 0 \\ 0 & 1 & 1 \end{array}$	
	0		$\begin{array}{cccccccccccccccccccccccccccccccccccc$

(b) The seven-qubit code. This [[7,1,3]] code has six stabilizer generators, which can be expressed as

$$\tilde{H} = \begin{pmatrix} H_{\text{ham}} & 0\\ 0 & H_{\text{ham}} \end{pmatrix}, \qquad (7.137)$$

where H_{ham} is the 3 × 7 parity-check matrix of the classical [7,4,3] Hamming code. The three check operators

detect the bit flips, and the three check operators

$$M_{4} = X_{1}X_{3}X_{5}X_{7}$$

$$M_{5} = X_{2}X_{3}X_{6}X_{7}$$

$$M_{6} = X_{4}X_{5}X_{6}X_{7},$$
(7.139)

detect the phase errors. The space with $M_1 = M_2 = M_3 = 1$ is spanned by the codewords that satisfy the Hamming parity check. Recalling that a Hadamard change of basis interchanges Z and X, we see that the space with $M_4 = M_5 = M_6$ is spanned by codewords that satisfy the Hamming parity check in the Hadamard-rotated basis. Indeed, we constructed the seven-qubit code by demanding that the Hamming parity check be satisfied in both bases. The generators commute because the Hamming code contains its dual code; *i.e.*, each row of H_{ham} satisfies the Hamming parity check.

(c) CSS codes. Recall whenever an [n, k, d] classical code C contains its dual code C^{\perp} , we can perform the CSS construction to obtain an [[n, 2k - n, d]] quantum code. The stabilizer of this code can be written as

$$\tilde{H} = \begin{pmatrix} H & 0 \\ 0 & H \end{pmatrix}$$
(7.140)

where H is the $(n-k) \times n$ parity check matrix of C. As for the sevenqubit code, the stabilizers commute because C contains C^{\perp} , and the code subspace is spanned by states that satisfy the H parity check in both the F-basis and the P-basis. Equivalently, codewords obey the Hparity check and are invariant under

$$|v\rangle \to |v+w\rangle,$$
 (7.141)

where $w \in C^{\perp}$.

(d) More general CSS codes. Consider, more generally, a stabilizer whose generators can each be chosen to be either a product of Z's $(\alpha|0)$ or a product of X's $(0|\beta)$. Then the generators have the form

$$\tilde{H} = \begin{pmatrix} H_Z & 0\\ 0 & H_X \end{pmatrix}.$$
(7.142)

7.9. STABILIZER CODES

Now, what condition must H_X and H_Z satisfy if the Z-generators and X-generators are to commute? Since Z's must collide with X's an even number of times, we have

$$H_X H_Z^T = H_Z H_X^T = 0 . (7.143)$$

But this is just the requirement that the dual C_X^{\perp} of the code whose parity check is H_X be contained in the code C_Z whose parity check is H_Z . In other words, this QECC fits into the CSS framework, with

$$C_2 = C_X^\perp \subseteq C_1 = C_Z. \tag{7.144}$$

So we may characterize CSS codes as those and only those for which the stabilizer has generators of the form eq. (7.142).

However there is a caveat. The code defined by eq. (7.142) will be nondegenerate if errors are restricted to weight less than $d = \min(d_Z, d_X)$ (where d_Z is the distance of C_Z , and d_X the distance of C_X). But the true distance of the QECC could exceed d. For example, the 9-qubit code is in this generalized sense a CSS code. But in that case the classical code C_X is distance 1, reflecting that, *e.g.*, Z_1Z_2 is contained in the stabilizer. Nevertheless, the distance of the CSS code is d = 3, since no weight-2 Pauli operator lies in $S^{\perp} \setminus S$.

7.9.4 Encoded qubits

We have seen that the troublesome errors are those in $S^{\perp} \setminus S$ — those that commute with the stabilizer, but lie outside of it. These Pauli operators are also of interest for another reason: they can be regarded as the "logical" operations that act on the encoded data that is protected by the code.

Appealing to the "linear algebra" viewpoint, we can see that the normalizer S^{\perp} of the stabilizer contains n + k independent generators – in the 2n-dimensional space of the $(\alpha|\beta)$'s, the subspace containing the vectors that are orthogonal to each of n - k linearly independent vectors has dimension 2n - (n - k) = n + k. Of the n + k vectors that span this space, n - kcan be chosen to be the generators of the stabilizer itself. The remaining 2k generators preserve the code subspace because they commute with the stabilizer, but act nontrivially on the k encoded qubits.

In fact, these 2k operations can be chosen to be the single-qubit operators $\bar{Z}_i, \bar{X}_i, i = 1, 2, ..., k$, where \bar{Z}_i, \bar{X}_i are the Pauli operators Z and X acting

on the encoded qubit labeled by *i*. First, note that we can extend the n-k stabilizer generators to a maximal set of *n* commuting operators. The *k* operators that we add to the set may be denoted $\bar{Z}_1, \ldots, \bar{Z}_k$. We can then regard the simultaneous eigenstates of $\bar{Z}_1 \ldots, \bar{Z}_k$ (in the code subspace \mathcal{H}_S) as the logical basis states $|\bar{z}_1, \ldots, \bar{z}_k\rangle$, with $\bar{z}_j = 0$ corresponding to $\bar{Z}_j = 1$ and $\bar{z}_j = 1$ corresponding to $\bar{Z}_j = -1$.

The remaining k generators of the normalizer may be chosen to be mutually commuting and to commute with the stabilizer, but then they will not commute with any of the \bar{Z}_i 's. By invoking a Gram-Schmidt orthonormalization procedure, we can choose these generators, denoted \bar{X}_i , to diagonalize the symplectic form, so that

$$\bar{\boldsymbol{Z}}_i \bar{\boldsymbol{X}}_j = (-1)^{\delta_{ij}} \bar{\boldsymbol{X}}_j \bar{\boldsymbol{Z}}_i. \tag{7.145}$$

Thus, each \bar{X}_j flips the eigenvalue of the corresponding \bar{Z}_j , and it can so be regarded as the Pauli operator X acting on encoded qubit i

(a) The 9-qubit Code. As we have discussed previously, the logical operators can be chosen to be

$$\bar{\boldsymbol{Z}} = \boldsymbol{X}_1 \boldsymbol{X}_2 \boldsymbol{X}_3 , \bar{\boldsymbol{X}} = \boldsymbol{Z}_1 \boldsymbol{Z}_4 \boldsymbol{Z}_7 .$$
 (7.146)

These anti-commute with one another (an X and a Z collide at position 1), commute with the stabilizer generators, and are independent of the generators (no element of the stabilizer contains three X's or three Z's).

(b) The 7-qubit code. We have seen that

$$\bar{\boldsymbol{X}} = \boldsymbol{X}_1 \boldsymbol{X}_2 \boldsymbol{X}_3 , \bar{\boldsymbol{Z}} = \boldsymbol{Z}_1 \boldsymbol{Z}_2 \boldsymbol{Z}_3 ;$$
 (7.147)

then \bar{X} adds an odd Hamming codeword and \bar{Z} flips the phase of an odd Hamming codeword. These operations implement a bit flip and phase flip respectively in the basis $\{|0\rangle_F, |1\rangle_F\}$ defined in eq. (7.93).

7.10 The 5-Qubit Code

All of the QECC's that we have considered so far are of the CSS type — each stabilizer generator is either a product of \mathbf{Z} 's or a product of \mathbf{X} 's. But not all stabilizer codes have this property. An example of a non-CSS stabilizer code is the perfect nondegenerate [[5,1,3]] code.

Its four stabilizer generators can be expressed

$$M_1 = XZZXI,$$

$$M_2 = IXZZX,$$

$$M_3 = XIXZZ,$$

$$M_4 = ZXIXZ,$$

(7.148)

 $M_{2,3,4}$ are obtained from M_1 by performing a cyclic permutation of the qubits. (The fifth operator obtained by a cyclic permutation of the qubits, $M_5 = ZZXIX = M_1M_2M_3M_4$ is not independent of the other four.) Since a cyclic permutation of a generator is another generator, the code itself is cyclic — a cyclic permutation of a codeword is a codeword.

Clearly each M_i contains no Y's and so squares to I. For each pair of generators, there are two collisions between an X and a Z, so that the generators commute. One can quickly check that each Pauli operator of weight 1 or weight 2 anti-commutes with at least one generator, so that the distance of the code is 3.

Consider, for example, whether there are error operators with support on the first two qubits that commute with all four generators. The weight-2 operator, to commute with the IX in M_2 and the XI in M_3 , must be XX. But XX anti-commutes with the XZ in M_1 and the ZX in M_4 .

In the symplectic notation, the stabilizer may be represented as

$$\tilde{H} = \begin{pmatrix} 01100 & 10010\\ 00110 & 01001\\ 00011 & 10100\\ 10001 & 01010 \end{pmatrix}$$
(7.149)

This matrix has a nice interpretation, as each of its columns can be regarded as the *syndrome* of a single-qubit error. For example, the single-qubit bit flip operator X_j , commutes with M_i if M_i has an I or X in position j, and anti-commutes if M_i has a Z in position j. Thus the table

	$oldsymbol{X}_1$	$oldsymbol{X}_2$	$oldsymbol{X}_3$	$oldsymbol{X}_4$	$oldsymbol{X}_5$
$oldsymbol{M}_1$	0	1	1	0	0
$oldsymbol{M}_2$	0	0	1	1	0
$oldsymbol{M}_3$	0	0	0	1	1
$oldsymbol{M}_4$	1	0	0	0	1

lists the outcome of measuring $M_{1,2,3,4}$ in the event of a bit flip. (For example, if the first bit flips, the measurement outcomes $M_1 = M_2 = M_3 = 1, M_4 = -1$, diagnose the error.) Similarly, the right half of \tilde{H} can be regarded as the syndrome table for the phase errors.

	$oldsymbol{Z}_1$	$oldsymbol{Z}_2$	$oldsymbol{Z}_3$	$oldsymbol{Z}_4$	$oldsymbol{Z}_5$
$oldsymbol{M}_1$	1	0	0	1	0
$oldsymbol{M}_2$	0	1	0	0	1
$oldsymbol{M}_3$	1	0	1	0	0
$oldsymbol{M}_4$	0	1	0	1	0

Since \boldsymbol{Y} anti-commutes with both \boldsymbol{X} and \boldsymbol{Z} , we obtain the syndrome for the error \boldsymbol{Y}_i by summing the *i*th columns of the \boldsymbol{X} and \boldsymbol{Z} tables:

	$oldsymbol{Y}_1$	\boldsymbol{Y}_2	$oldsymbol{Y}_3$	\boldsymbol{Y}_4	$oldsymbol{Y}_5$
$oldsymbol{M}_1$	1	1	1	1	0
$oldsymbol{M}_2$	0	1	1	1	1
$oldsymbol{M}_3$	1	0	1	1	1
$oldsymbol{M}_4$	1	1	0	1	1

We find by inspection that the 15 columns of the X, Y, and Z syndrome tables are all distinct, and so we verify again that our code is a nondegenerate code that corrects one error. Indeed, the code is perfect — each of the 15 nontrivial binary strings of length 4 appears as a column in one of the tables.

Because of the cyclic property of the code, we can easily characterize all 15 nontrivial elements of its stabilizer. Aside from $M_1 = XZZXI$ and the four operators obtained from it by cyclic permutations of the qubit, the stabilizer also contains

$$\boldsymbol{M}_{3}\boldsymbol{M}_{4} = -\boldsymbol{Y}\boldsymbol{X}\boldsymbol{X}\boldsymbol{Y}\boldsymbol{I}, \tag{7.150}$$

plus its cyclic permutations, and

$$\boldsymbol{M}_{2}\boldsymbol{M}_{5} = -\boldsymbol{Z}\boldsymbol{Y}\boldsymbol{Y}\boldsymbol{Z}\boldsymbol{I}, \tag{7.151}$$

and its cyclic permutations. Evidently, all elements of the stabilizer are weight-4 Pauli operators.

For our logical operators, we may choose

$$\bar{Z} = ZZZZZ,$$

 $\bar{X} = XXXXX;$ (7.152)

these commute with $M_{1,2,3,4}$, square to I, and anti-commute with one another. Being weight 5, they are not themselves contained in the stabilizer. Therefore if we don't mind destroying the encoded state, we can determine the value of \bar{Z} for the encoded qubit by measuring Z of each qubit and evaluating the parity of the outcomes. In fact, since the code is distance three, there are elements of $S^{\perp} \setminus S$ of weight-three; alternate expressions for \bar{Z} and \bar{X} can be obtained by multiplying by elements of the stabilizer. For example we can choose

$$\bar{Z} = (ZZZZZ) \cdot (-ZYYZI) = -IXXIZ, \quad (7.153)$$

(or one of its cyclic permutations), and

$$\bar{\boldsymbol{X}} = (\boldsymbol{X}\boldsymbol{X}\boldsymbol{X}\boldsymbol{X}\boldsymbol{X}\boldsymbol{X}) \cdot (-\boldsymbol{Y}\boldsymbol{X}\boldsymbol{X}\boldsymbol{Y}\boldsymbol{I}) = -\boldsymbol{Z}\boldsymbol{I}\boldsymbol{I}\boldsymbol{Z}\boldsymbol{X},$$
(7.154)

(or one of its cyclic permutations). So it is possible to ascertain the value of \bar{X} or \bar{Z} by measuring X or Z of only three of the five qubits in the block, and evaluating the parity of the outcomes.

If we wish, we can construct an orthonormal basis for the code subspace, as follows. Starting from any state $|\psi_0\rangle$, we can obtain

$$|\Psi_0\rangle = \sum_{\boldsymbol{M}\in S} \boldsymbol{M} |\psi_0\rangle.$$
(7.155)

This (unnormalized) state obeys $\mathbf{M}'|\Psi_0\rangle = |\Psi_0\rangle$ for each $\mathbf{M}' \in S$, since multiplication by an element of the stabilizer merely permutes the terms in the sum. To obtain the $\bar{\mathbf{Z}} = 1$ encoded state $|\bar{0}\rangle$, we may start with the state $|00000\rangle$, which is also a $\bar{\mathbf{Z}} = 1$ eigenstate, but not in the stabilizer; we find (up to normalization)

$$\begin{split} \bar{0} \rangle &= \sum_{\boldsymbol{M} \in S} |00000\rangle \\ &= |00000\rangle + (\boldsymbol{M}_1 + \text{cyclic perms}) |00000\rangle \\ &+ (\boldsymbol{M}_3 \boldsymbol{M}_4 + \text{cyclic perms}) |00000\rangle + (\boldsymbol{M}_2 \boldsymbol{M}_5 + \text{cyclic perms}) |00000\rangle \\ &= |00000\rangle + (110010\rangle + \text{cyclic perms}) \\ &- (|11110\rangle + \text{cyclic perms}) \\ &- (|01100\rangle + \text{cyclic perms}). \end{split}$$
(7.156)

We may then find $|\bar{1}\rangle$ by applying \bar{X} to $|\bar{0}\rangle$, that is by flipping all 5 qubits:

$$|\bar{1}\rangle = \bar{\boldsymbol{X}}|\bar{0}\rangle = |11111\rangle + (|01101\rangle + \text{ cyclic perms}) - (|00001\rangle + \text{ cyclic perms}) - (|10011\rangle + \text{ cyclic perms}) .$$
(7.157)

How is the syndrome measured? A circuit that can be executed to measure $M_1 = XZZXI$ is:

The Hadamard rotations on the first and fourth qubits rotate M_1 to the tensor product of Z's ZZZZI, and the CNOT's then imprint the value of this operator on the ancilla. The final Hadamard rotations return the encoded block to the standard code subspace. Circuits for measuring $M_{2,3,4}$ are obtained from the above by cyclically permuting the five qubits in the code block.

What about encoding? We want to construct a unitary transformation

$$\boldsymbol{U}_{\text{encode}} : |0000\rangle \otimes (a|0\rangle + b|1\rangle) \to a|\bar{0}\rangle + b|\bar{1}\rangle.$$
 (7.158)

We have already seen that $|00000\rangle$ is a $\bar{Z} = 1$ eigenstate, and that $|00001\rangle$ is a $\bar{Z} = -1$ eigenstate. Therefore (up to normalization)

$$a|\bar{0}\rangle + b|\bar{1}\rangle = \left(\sum_{\boldsymbol{M}\in S} \boldsymbol{M}\right)|0000\rangle \otimes (a|0\rangle + b|1\rangle).$$
 (7.159)

So we need to figure out how to construct a circuit that applies $(\sum M)$ to an initial state.

Since the generators are independent, each element of the stabilizer can be expressed as a product of generators as a unique way, and we may therefore rewrite the sum as

$$\sum_{\boldsymbol{M}\in S} \boldsymbol{M} = (\boldsymbol{I} + \boldsymbol{M}_4)(\boldsymbol{I} + \boldsymbol{M}_3)(\boldsymbol{I} + \boldsymbol{M}_2)(\boldsymbol{I} + \boldsymbol{M}_1) .$$
(7.160)

Now to proceed further it is convenient to express the stabilizer in an alternative form. Note that we have the freedom to replace the generator M_i by M_iM_j without changing the stabilizer. This replacement is equivalent to adding the *j*th row to the *i*th row in the matrix \tilde{H} . With such row operations, we can perform a Gaussian elimination on the 4×5 matrix H_X , and so obtain the new presentation for the stabilizer

$$\tilde{H}' = \begin{pmatrix} 11011 & 10001\\ 00110 & 01001\\ 11000 & 00101\\ 10111 & 00011 \end{pmatrix} , \qquad (7.161)$$

or

$$M_{1} = YZIZY$$

$$M_{2} = IXZZX$$

$$M_{3} = ZZXIX$$

$$M_{4} = ZIZYY$$
(7.162)

In this form M_i applies an X (flip) only to qubits *i* and 5 in the block.

Adopting this form for the stabilizer, we can apply $\frac{1}{\sqrt{2}}(I + M_1)$ to a state $|0, z_2, z_3, z_4, z_5\rangle$ by executing the circuit

– Figure –

The Hadamard prepares $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle$. If the first qubit is $|0\rangle$, the other operations don't do anything, so I is applied. But if the first qubit is $|1\rangle$, then X has been applied to this qubit, and the other gates in the circuit apply

ZZIZY, conditioned on the first qubit being $|1\rangle$. Hence, $YZIZY = M_1$ has been applied. Similar circuits can be constructed that apply $\frac{1}{\sqrt{2}}(I + M_2)$ to $|z_1, 0, z_3, z_4, z_5\rangle$, and so forth. Apart from the Hadamard gates each of these circuits applies only Z's and conditional Z's to qubits 1 through 4; these qubits never flip. (It was to ensure thus that we performed the Gaussian elimination on H_X .) Therefore, we can construct our encoding circuit as

– Figure –

Furthermore, each Z gate acting on $|0\rangle$ can be replaced by the identity, so we may simplify the circuit by eliminating all such gates, obtaining

– Figure –

This procedure can be generalized to construct an encoding circuit for any stabilizer code.

Since the encoding transformation is unitary, we can use its adjoint to decode. And since each gate squares to $\pm I$, the decoding circuit is just the encoding circuit run in reverse.

7.11 Quantum secret sharing

The [[5, 1, 3]] code provides a nice illustration of a possible application of QECC's.²

Suppose that some top secret information is to be entrusted to n parties. Because none is entirely trusted, the secret is divided into n shares, so that each party, with access to his share alone, can learn nothing at all about the secret. But if enough parties get together and pool their shares, they can decipher the secret or some part of it.

In particular, an (m, n) threshold scheme has the property that m shares are sufficient to reconstruct all of the secret information. But from m - 1

 $^{^2\}mathrm{R.}$ Cleve, D. Gottesman, and H.-K. Lo, "How to Share a Quantum Secret," quant-ph/9901025.

7.11. QUANTUM SECRET SHARING

shares, no information at all can be extracted. (This is called a *threshold* scheme because as shares 1, 2, 3..., m-1 are collected one by one, nothing is learned, but the next share crosses the threshold and reveals everything.)

We should distinguish too kinds of secrets: a classical secret is an *a priori* unknown bit string, while a quantum secret is an *a priori* unknown quantum state. Either type of secret can be shared. In particular, we can distribute a classical secret among several parties by selecting one from an ensemble of mutually orthogonal (entangled) quantum states, and dividing the state among the parties.

We can see, for example, that the [[5,1,3]] code may be employed in a (3,5) threshold scheme, where the shared information is classical. One classical bit is encoded by preparing one of the two orthogonal states $|\bar{0}\rangle$ or $|\bar{1}\rangle$ and then the five qubits are distributed to five parties. We have seen that (since the code is nondegenerate) if any two parties get together, then the density matrix ρ their two qubits is

$$\boldsymbol{\rho}^{(2)} = \frac{1}{4} \mathbf{1} \ . \tag{7.163}$$

Hence, they learn nothing about the quantum state from any measurement of their two qubits. But we have also seen that the code can correct two located errors or two erasures. When any three parties get together, they may correct the two errors (the two missing qubits) and perfectly reconstruct the encoded state $|\bar{0}\rangle$ or $|\bar{1}\rangle$.

It is also clear that by a similar procedure a single qubit of quantum information can be shared – the [[5, 1, 3]] code is also the basis of a ((3, 5)) quantum threshold scheme (we use the ((m, n)) notation if the shared information is quantum information, and the (m, n) notation if the shared information is classical). How does this quantum-secret-sharing scenario generalize to more qubits? Suppose we prepare a pure state $|\psi\rangle$ of n qubits — can it be employed in an ((m, n)) threshold scheme?

We know that m qubits must be sufficient to reconstruct the state; hence n-m erasures can be corrected. It follows from our general error correction criterion that the expectation value of any weight-(n-m) observable must be independent of the state $|\psi\rangle$

$$\langle \psi | \boldsymbol{E} | \psi \rangle$$
 independent of $| \psi \rangle$, $\operatorname{wt}(\boldsymbol{E}) \leq n - m$. (7.164)

Thus, if m parties have all the information, the other n - m parties have no information at all. That makes sense, since quantum information cannot be cloned.

On the other hand, we know that m-1 shares reveal nothing, or that

$$\langle \psi | \boldsymbol{E} | \psi \rangle$$
 independent of $| \psi \rangle$, wt $(\boldsymbol{E}) \le m - 1$. (7.165)

It then follows that m-1 erasures can be corrected, or that the other n-m+1 parties have all the information.

From these two observations we obtain the two inequalities

$$n - m < m \quad \Rightarrow \quad n < 2m ,$$

$$m - 1 < n - m + 1 \quad \Rightarrow \quad n > 2m - 2 .$$
 (7.166)

It follows that

$$n = 2m - 1$$
, (7.167)

in an ((m, n)) pure state quantum threshold scheme, where each party has a single qubit. In other words, the threshold is reached as the number of qubits in hand crosses over from the minority to the majority of all n qubits.

We see that if each share is a qubit, a quantum pure state threshold scheme is a [[2m-1, k, m]] quantum code with $k \ge 1$. But in fact the [[3, 1, 2]] and [[7, 1, 4]] codes do not exist, and it follows from the Rains bound that the m > 3 codes do not exist. In a sense, then, the [[5, 1, 3]] code is the unique quantum threshold scheme.

There are a number of caveats — the restriction n = 2m - 1 continues to apply if each share is a q-dimensional system rather than a qubit, but various

$$[[2m-1,1,k]]_q (7.168)$$

codes can be constructed for q > 2. (See the exercises for an example.)

Also, we might allow the shared information to be a mixed state (that encodes a pure state). For example, if we discard one qubit of the five qubit block, we have a ((3, 4)) scheme. Again, once we have three qubits, we can correct two erasures, one arising because the fourth share is in the hands of another party, the other arising because a qubit has been thrown away.

Finally, we have assumed that the shared information is quantum information. But if we are only sharing classical information instead, then the conditions for correcting erasures are less stringent. For example, a Bell pair may be regarded as a kind of (2, 2) threshold scheme for two bits of classical information, where the classical information is encoded by choosing one of the four mutually orthogonal states $|\phi^{\pm}\rangle$, $|\psi^{\pm}\rangle$. A party in possession of one of the two qubits is unable to access any of this classical information. But this is not a scheme for sharing a quantum secret, since linear combinations of these Bell states do *not* have the property that $\rho = \frac{1}{2}\mathbf{1}$ if we trace out one of the two qubits.

7.12 Some Other Stabilizer Codes

7.12.1 The [[6,0,4]] code

A k = 0 quantum code has a one-dimensional code subspace; that is, there is only one encoded state. The code cannot be used to store unknown quantum information, but even so, k = 0 codes can have interesting properties. Since they can detect and diagnose errors, they might be useful for a study of the correlations in decoherence induced by interactions with the environment.

If k = 0, then S and S^{\perp} coincide – a Pauli operator that commutes with all elements of the stabilizer must lie in the stabilizer. In this case, the distance d is defined as the minimum weight of any Pauli operator in the stabilizer. Thus a distance-d code can "detect d - 1 errors;" that is, if any Pauli operator of weight less than d acts on the code state, the result is orthogonal to that state.

Associated with the [[5, 1, 3]] code is a [[6, 0, 4]] code, whose encoded state can be expressed as

$$|0\rangle \otimes |\bar{0}\rangle + |1\rangle \otimes |\bar{1}\rangle, \qquad (7.169)$$

where $|\bar{0}\rangle$ and $|\bar{1}\rangle$ are the \bar{Z} eigenstates of the [[5, 1, 3]] code. You can verify that this code has distance d = 4 (an exercise).

The [[6, 0, 4]] code is interesting because its code state is maximally entangled. We may choose any three qubits from among the six. The density matrix $\rho^{(3)}$ of those three, obtained by tracing over the other three, is totally random, $\rho^{(3)} = \frac{1}{8}I$. In this sense, the [[6, 0, 4]] state is a natural multiparticle analog of the two-qubit Bell states. It is far "more entangled" than the six-qubit cat state $\frac{1}{\sqrt{2}}(|000000\rangle + |111111\rangle)$. If we measure any one of the six qubits in the cat state, in the $\{|0\rangle, |1\rangle\}$ basis, we know everything about the state we have prepared of the remaining five qubits. But we may measure any observable we please acting on any *three* qubits in the [[6, 0, 4]] state, and we learn *nothing* about the remaining three qubits, which are still described by $\rho^{(3)} = \frac{1}{8}I$.

Our [[6, 0, 4]] state is all the more interesting in that it turns out (but is not so simple to prove) that its generalizations to more qubits do not exist. That is, there are no [[2n, 0, n + 1]] binary quantum codes for n > 3. You'll see in the exercises, though, that there are other, nonbinary, maximally entangled states that can be constructed.

7.12.2 The [[2m, 2m - 2, 2]] error-detecting codes

The Bell state $|\phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ is a [[2, 0, 2]] code with stabilizer generators

$$\begin{array}{l} \boldsymbol{Z}\boldsymbol{Z} \\ \boldsymbol{X}\boldsymbol{X} \\ \boldsymbol{X} \end{array} , \qquad (7.170)$$

The code has distance two because no weight-one Pauli operator commutes with both generators (none of X, Y, Z commute with both X and Z). Correspondingly, a bit flip (X) or a phase flip (Z), or both (Y) acting on either qubit in $|\phi^+\rangle$, takes it to an orthogonal state (one of the other Bell states $|\phi^-\rangle, |\psi^+\rangle, |\psi^-\rangle$).

One way to generalize the Bell states to more qubits is to consider the n = 4, k = 2 code with stabilizer generators

$$\begin{array}{l} \boldsymbol{Z}\boldsymbol{Z}\boldsymbol{Z}\boldsymbol{Z} \ , \\ \boldsymbol{X}\boldsymbol{X}\boldsymbol{X}\boldsymbol{X} \ . \end{array} \tag{7.171}$$

This is a distance d = 2 code for the same reason as before. The code subspace is spanned by states of even parity (ZZZZ) that are invariant under a simultaneous flip of all four qubits (XXX). A basis is:

$$\begin{array}{l} |0000\rangle + |1111\rangle , \\ |0011\rangle + |1100\rangle , \\ |0101\rangle + |1010\rangle , \\ |0110\rangle + |1001\rangle . \end{array}$$
(7.172)

Evidently, an X or a Z acting on any qubit takes each of these states to a state orthogonal to the code subspace; thus any single-qubit error can be detected.

A further generalization is the [[2m, 2m - 2, 2]] code with stabilizer generators

(the length is required to be even so that the generators will commute. The code subspace is spanned by our familiar friends the 2^{n-2} cat states

$$\frac{1}{\sqrt{2}}(|x\rangle + |\neg x\rangle), \tag{7.174}$$

where x is an even-weight string of length n = 2m.

7.12.3 The [[8,3,3]] code

As already noted in our discussion of the [[5, 1, 3]] code, a stabilizer code with generators

$$\tilde{H} = (H_Z | H_X), \tag{7.175}$$

can correct one error if: (1) the columns of \tilde{H} are distinct (a distinct syndrome for each X and Z error) and (2) each sum of a column of H_Z with the corresponding column of H_X is distinct from each column of \tilde{H} and distinct from all other such sums (each Y error can be distinguished from all other one-qubit errors).

We can readily construct a 5×16 matrix \tilde{H} with this property, and so derive the stabilizer of an [[8, 3, 3]] code; we choose

$$\tilde{H} = \begin{pmatrix} H & H^{\sigma} \\ 11111111 & 0000000 \\ 00000000 & 1111111 \end{pmatrix} .$$
(7.176)

Here H is the 3×8 matrix

$$H = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 \end{pmatrix}$$
(7.177)

whose columns are all the distinct binary strings of length 3, and H^{σ} is obtained from H by performing a suitable permutation of the columns. This

permutation is chosen so that the eight sums of columns of H with corresponding columns of H^{σ} are all distinct. We may see by inspection that a suitable choice is

$$H^{\sigma} = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \end{pmatrix}$$
(7.178)

as the column sums are then

The last two rows of \tilde{H} serve to distinguish each X syndrome from each Y syndrome or Z syndrome, and the above mentioned property of H^{σ} ensures that all Y syndromes are distinct. Therefore, we have constructed a length-8 code with k = 8-5 = 3 that can correct one error. It is actually the simplest in an infinite class of $[[2^m, 2^m - m - 2, 3]]$ codes constructed by Gottesman, with $m \geq 3$.

The [[8, 3, 3]] quantum code that we have just described is a close cousin of the "extended Hamming code," the self-dual [8,4,4] classical code that is obtained from the [7,3,4] dual of the Hamming code by adding an extra parity bit. Its parity check matrix (which is also its generator matrix) is

$$H_{\rm EH} = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix}$$
(7.180)

This matrix $H_{\rm EH}$ has the property that, not only are its eight columns distinct, but also each *sum* of two columns is distinct from all columns; since the sum of two columns has 0, not 1, as its fourth bit.

7.13 Codes Over GF(4)

We constructed the [[5, 1, 3]] code by guessing the stabilizer generators, and checking that d = 3. Is there a more systematic method?

In fact, there is. Our suspicion that the [[5, 1, 3]] code might exist was aroused by the observation that its parameters saturate the quantum sphere-packing inequality for t = 1 codes:

$$1 + 3n = 2^{n-k}, (7.181)$$

(16 = 16 for n = 5 and k = 1). To a coding theorist, this equation might look familiar.

Aside from the binary codes we have focused on up to now, classical codes can also be constructed from length-*n* strings of symbols that take values, not in $\{0, 1\}$, but in the finite field with *q* elements GF(q). Such finite fields exist for any $q = p^m$, where *p* is prime. (*GF* is short for "Galois Field," in honor of their discoverer.)

For such nonbinary codes, we may model error as addition by an element of the field, a cyclic shift of the q symbols. Then there are q-1 nontrivial errors. The weight of a vector in $GF(q)^n$ is the number of its nonzero elements, and the distance between two vectors is the weight of their difference (the number of elements that disagree). An $[n, k, d]_q$ classical code consists of q^k codewords in $GF(q)^n$, where the minimal distance between a pair is d. The sphere packing bound that must be satisfied for an $[n, k, d]_q$ code to exist becomes, for d = 3,

$$1 + (q-1)n \le q^{n-k}.$$
(7.182)

In fact, the perfect binary Hamming codes that saturate this bound for q = 2 with parameters

$$n = 2^m - 1, \ k = n - m, \tag{7.183}$$

admit a generalization to any GF(q); perfect Hamming codes over GF(q) can be constructed with

$$n = \frac{q^m - 1}{q - 1}, \ k = n - m.$$
 (7.184)

The [[5, 1, 3]] quantum code is descended from the classical $[5, 3, 3]_4$ Hamming code (the case q = 4 and m = 2).

What do the classical GF(4) codes have to do with binary quantum stabilizer codes? The connection arises because the stabilizer can be associated with a set of vectors over GF(4) closed under addition. The field GF(4) has four elements that may be denoted $0, 1, \omega, \bar{\omega}$, where

$$1 + 1 = \omega + \omega = \bar{\omega} + \bar{\omega} = 0,$$

$$1 + \omega = \bar{\omega},$$
(7.185)

and $\omega^2 = \bar{\omega}$, $\omega \bar{\omega} = 1$. Thus, the additive structure of GF(4) echos the multiplicative structure of the Pauli operators $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}$. Indeed, the length-2n binary string $(\alpha|\beta)$ that we have used to denote an element of the Pauli group can equivalently be regarded as a length-n vector in $GF(4)^n$

$$(\alpha|\beta) \leftrightarrow \alpha + \beta\omega. \tag{7.186}$$

The stabilizer, with 2^{n-k} elements, can be regarded as a subcode of GF(4), closed under addition and containing 2^{n-k} codewords.

Note that the code need not be a vector space over GF(4), as it is not required to be closed under multiplication by a scalar $\in GF(4)$. In the special case where the code is a vector space, it is called a *linear* code.

Much is known about codes over GF(4), so this connection opened the door for the (classical) coding theorists to construct many QECC's.³ However, not every subcode of $GF(4)^n$ is associated with a quantum code; we have not yet imposed the requirement that the stabilizer is abelian – the $(\alpha|\beta)$'s that span the code must be mutually orthogonal in the symplectic inner product

$$\alpha \cdot \beta' + \alpha' \cdot \beta . \tag{7.187}$$

This orthogonality condition might look strange to a coding theorist, who is more accustomed to defining the inner product of two vectors in $GF(4)^n$ as an element of GF(4) given by

$$v * u = \bar{v}_1 u_1 + \dots + \bar{v}_n u_n ,$$
 (7.188)

where conjugation, denoted by a bar, interchanges ω and $\bar{\omega}$. If this "hermitian" inner product * of two vectors v and u is

$$v * u = a + b\omega \in GF(4) , \qquad (7.189)$$

 $^{^3 {\}rm Calderbank},$ Rains, Shor, and Sloane, "Quantum error correction via codes over GF(4)," quant-ph/9608006.

then our symplectic inner product is

$$v \cdot u = b \ . \tag{7.190}$$

Therefore, vanishing of the symplectic inner product is a weaker condition than vanishing of the hermitian inner product. In fact, though, in the special case of a *linear* code, self-orthogonality with respect to the hermitian inner product is actually equivalent to self-orthogonality with respect to the symplectic inner product. We observe that if $v * u = a + b\omega$, orthogonality in the symplectic inner product requires b = 0. But if u is in a linear code, then so is $\bar{\omega}u$ where

$$v * (\bar{\omega}u) = b + a\bar{\omega} \tag{7.191}$$

so that

$$v \cdot (\bar{\omega}u) = a . \tag{7.192}$$

We see that if v and u belong to a linear GF(4) code and are orthogonal with respect to the symplectic inner product, then they are also orthogonal with respect to the hermitian inner product. We conclude then, that a linear GF(4) code defines a quantum stabilizer code if and only if the code is self-orthogonal in the hermitian inner product. Classical codes with these properties have been much studied.

In particular, consider again the $[5, 3, 3]_4$ Hamming code. Its parity check matrix (in an unconventional presentation) can be expressed as

$$H = \begin{pmatrix} 1 & \omega & \omega & 1 & 0 \\ 0 & 1 & \omega & \omega & 1 \end{pmatrix},$$
(7.193)

which is also the generator matrix of its dual, a linear self-orthogonal $[5, 2, 4]_4$ code. In fact, this $[5, 2, 4]_4$ code, with $4^2 = 16$ codewords, is precisely the stabilizer of the [[5, 1, 3]] quantum code. By identifying $1 \equiv \mathbf{X}, \omega \equiv \mathbf{Z}$, we recognize the two rows of H as the stabilizer generators $\mathbf{M}_1, \mathbf{M}_2$. The dual of the Hamming code is a linear code, so linear combinations of the rows are contained in the code. Adding the rows and multiplying by ω we obtain

$$\omega(1,\bar{\omega},0,\bar{\omega},1) = (\omega,1,0,1,\omega), \tag{7.194}$$

which is M_4 . And if we add M_4 to M_2 and multiply by $\bar{\omega}$, we find

$$\bar{\omega}(\omega, 0, \omega, \bar{\omega}, \bar{\omega}) = (1, 0, 1, \omega, \omega), \qquad (7.195)$$

which is M_3 .

The [[5, 1, 3]] code is just one example of a quite general construction. Consider a subcode C of $GF(4)^n$ that is additive (closed under addition), and self-orthogonal (contained in its dual) with respect to the symplectic inner product. This GF(4) code can be identified with the stabilizer of a binary QECC with length n. If the GF(4) code contains 2^{n-k} codewords, then the QECC has k encoded qubits. The distance d of the QECC is the minimum weight of a vector in $C^{\perp} \setminus C$.

Another example of a self-orthogonal linear GF(4) code is the dual of the m = 3 Hamming code with

$$n = \frac{1}{3}(4^3 - 1) = 21. \tag{7.196}$$

The Hamming code has 4^{n-m} codewords, and its dual has $4^m = 2^6$ codewords. We immediately obtain a QECC with parameters

$$[[21, 15, 3]], (7.197)$$

that can correct one error.

7.14 Good Quantum Codes

A family of [[n, k, d]] codes is good if it contains codes whose "rate" R = k/nand "error probability" p = t/n (where (t = (d - 1)/2) both approach a nonzero limit as $n \to \infty$. We can use the stabilizer formalism to prove a "quantum Gilbert-Varshamov" bound that demonstrates the existence of good quantum codes. In fact, good codes can be chosen to be nondegenerate.

We will only sketch the argument, without carrying out the requisite counting precisely. Let $\mathcal{E} = \{\mathbf{E}_a\}$ be a set of errors to be corrected, and denote by $\mathcal{E}^{(2)} = \{\mathbf{E}_a^{\dagger} \mathbf{E}_b\}$, the products of pairs of elements of \mathcal{E} . Then to construct a nondegenerate code that can correct the errors in \mathcal{E} , we must find a set of stabilizer generators such that some generator anti-commutes with each element of $\mathcal{E}^{(2)}$.

To see if a code with length n and k qubits can do the job, begin with the set $\mathcal{S}^{(n-k)}$ of all abelian subgroups of the Pauli group with n-k generators. We will gradually pare away the subgroups that are unsuitable stabilizers for correcting the errors in \mathcal{E} , and then see if any are left.

Each nontrivial error E_a commutes with a fraction $\sim 1/2^{n-k}$ of all groups contained in $\mathcal{S}^{(n-k)}$, since it is required to commute with each of the n-kgenerators of the group. (There is a small correction to this fraction that we may ignore for large n.) Each time we add another element to $\mathcal{E}^{(2)}$, a fraction 2^{k-n} of all stabilizer candidates must be rejected. When $\mathcal{E}^{(2)}$ has been fully assembled, we have rejected at worst a fraction

$$|\mathcal{E}^{(2)}| \cdot 2^{k-n},$$
 (7.198)

of all the subgroups contained in $\mathcal{S}^{(n-k)}$ (where $|\mathcal{E}^{(2)}|$ is the number of elements of $\mathcal{E}^{(2)}$.) As long as this fraction is less than one, a stabilizer that does the job will exist for large n.

If we want to correct t = pn errors, then $\mathcal{E}^{(2)}$ contains operators of weight at most 2t and we may estimate

$$\log_2 |\mathcal{E}^{(2)}| \lesssim \log_2 \left[\binom{n}{2pn} 3^{2pn} \right] \sim n \left[H_2(2p) + 2p \log_2 3 \right].$$
(7.199)

Therefore, nondegenerate quantum stabilizer codes that correct pn errors exist, with asymptotic vote R = k/n given by

$$\log_2 |\mathcal{E}^{(2)}| + k - n < 0, \quad \text{or} \quad R < 1 - H_2(2p) - 2p \log_2 3.$$

(7.200)

Thus is the (asymptotic form of the) quantum Gilbert–Varshamov bound.

We conclude that codes with a nonzero rate must exist that protect against errors that occur with any error probability $p < p_{\rm GV} \simeq .0946$. The maximum error probability allowed by the Rains bound is p = 1/6, for a code that can protect against every error operator of weight $\leq pn$.

Though good quantum codes exist, the explicit construction of families of good codes is quite another matter. Indeed, no such constructions are known.

7.15 Some Codes that Correct Multiple Errors

7.15.1 Concatenated codes

Up until now, all of the QECC's that we have explicitly constructed have d = 3 (or d = 2), and so can correct one error (at best). Now we will

describe some examples of codes that have higher distance.

A particularly simple way to construct codes that can correct more errors is to concatenate codes that can correct one error. A concatenated code is a code within a code. Suppose we have two k = 1 QECC's, an $[[n_1, 1, d_1]]$ code C_1 code and an $[[n_2, 1, d_2]]$ code C_2 . Imagine constructing a length n_2 codeword of C_2 , and expanding the codeword as a coherent superposition of product states, in which each qubit is in one of the states $|0\rangle$ or $|1\rangle$. Now replace each qubit by a length- n_1 encoded state using the code C_1 ; that is replace $|0\rangle$ by $|\bar{0}\rangle$ and $|1\rangle$ by $|\bar{1}\rangle$ of C_1 . The result is a code with length $n = n_1n_2, k = 1$, and distance no less than $d = d_1d_2$. We will call C_2 the "outer" code and C_1 the "inner" code.

In fact, we have already discussed one example of this construction: Shor's 9-qubit code. In that case, the inner code is the three-qubit repetition code with stabilizer generators

$$\boldsymbol{ZZI}, \quad \boldsymbol{IZZ}, \quad (7.201)$$

and the outer code is the three-qubit "phase code" with stabilizer generators

$$XXI, IXX (7.202)$$

(the Hadamard rotated repetition code). We construct the stabilizer of the concatenated code as follows: Acting on each of the three qubits contained in the block of the outer code, we include the two generators Z_1Z_2, Z_2Z_3 of the inner code (six generators altogether). Then we add the two generators of the outer code, but with X, Z replaced by the *encoded* operations of the inner code; in this case, these are the two generators

$$\bar{\boldsymbol{X}}\bar{\boldsymbol{X}}\bar{\boldsymbol{I}},\ \bar{\boldsymbol{I}}\bar{\boldsymbol{X}}\bar{\boldsymbol{X}},$$
 (7.203)

where $\overline{I} = III$ and $\overline{X} = XXX$. You will recognize these as the eight stabilizer generators of Shor's code that we have described earlier. In this case, the inner and outer codes both have distance 1 (*e.g.*, ZII commutes with the stabilizer of the inner code), yet the concatenated code has distance $3 > d_1d_2 = 1$. This happens because the code has been cleverly constructed so that the weight 1 and 2 encoded operations of the inner code do not commute with the stabilizer of the outer code. (It would have been different if we had concatenated the repetition code with itself rather than with the phase code!) We can obtain a distance 9 code (capable of correcting four errors) by concatenating the [[5, 1, 3]] code with itself. The length n = 25 is the smallest for any known code with k = 1 and d = 9. (An [[n, 1, 9]] code with n = 23, 24 would be consistent with the Rains bound, but it is unknown whether such a code really exists.)

The stabilizer of the [[25, 1, 9]] concatenated code has 24 generators. Of these, 20 are obtained as the four generators $M_{1,2,3,4}$ acting on each of the five subblocks of the outer code, and the remaining four are the *encoded* operators $\bar{M}_{1,2,3,4}$ of the outer code. Notice that the stabilizer contains elements of weight 4 (the stabilizer elements acting on each of the five inner codes); therefore, the code is degenerate. This is typical of concatenated codes.

There is no need to stop at two levels of concatenation; from L QECC's with parameters $[[n_1, 1, d_1]], \ldots, [[n_L, 1, d_L]]$, we can construct a hierarchical code with altogether L levels of codes within codes; it has length

$$n = n_1 n_2 \dots n_L, \tag{7.204}$$

and distance

$$d \ge d_1 d_2 \dots d_L. \tag{7.205}$$

In particular, by concatenating the [[5, 1, 3]] code L times, we may construct a code with parameters

$$[[5^L, 1, 3^L]]. (7.206)$$

Strictly speaking, this family of codes cannot protect against a number of errors that scales linearly with the length. Rather the ratio of the number t of errors that can be corrected to the length n is

$$\frac{t}{n} \sim \frac{1}{2} \left(\frac{3}{5}\right)^L \,, \tag{7.207}$$

which tends to zero for large L. But the distance d may be a deceptive measure of how well the code performs — it is all right if recovery fails for *some* ways of choosing $t \ll pn$ errors, so long as recovery will be successful for the *typical* ways of choosing pn faulty qubits. In fact, concatenated codes *can* correct pn typical errors, for n large and p > 0.

Actually, the way concatenated codes are usually used does not fully exploit their power to correct errors. To be concrete, consider the [[5, 1, 3]] code in the case where each of the five qubits is independently subjected to the depolarizing channel with error probability p (that is X, Y, Z errors each occur with probability p/3). Recovery is sure to succeed if fewer than two errors occur in the block. Therefore, as in §7.4.2, we can bound the failure probability by

$$p_{\text{fail}} \equiv p^{(1)} \le {5 \choose 2} p^2 = 10p^2.$$
 (7.208)

Now consider the performance of the concatenated [[25, 1, 9]] code. To keep life easy, we will perform recovery in a simple (but nonoptimal) way: First we perform recovery on each of the five subblocks, measuring $M_{1,2,3,4}$ to obtain an error syndrome for each subblock. After correcting the subblocks, we then measure the stabilizer generators $\bar{M}_{1,2,3,4}$ of the outer code, to obtains its syndrome, and apply an encoded \bar{X} , \bar{Y} , or \bar{Z} to one of the subblocks if the syndrome reveals an error.

For the outer code, recovery will succeed if at most one of the subblocks is damaged, and the probability $p^{(1)}$ of damage to a subblock is bounded as in eq. (7.208); we conclude that the probability of a botched recovery for the [[25, 1, 9]] code is bounded above by

$$p^{(2)} \le 10(p^{(1)})^2 \le 10(10p^2)^2 = 1000p^4.$$
 (7.209)

Our recovery procedure is clearly not the best possible, because four errors can induce failure if there are two each in two different subblocks. Since the code has distance nine, there is a better procedure that would always recover successfully from four errors, so that $p^{(2)}$ would be of order p^5 rather than p^4 . Still, the suboptimal procedure has the advantage that it is very easily generalized, (and analyzed) if there are many levels of concatenation.

Indeed, if there are L levels of concatenation, we begin recovery at the innermost level and work our way up. Solving the recursion

$$p^{(\ell)} \le C[p^{(\ell-1)}]^2,$$
 (7.210)

starting with $p^{(0)} = p$, we conclude that

$$p^{(L)} \le \frac{1}{C} (Cp)^{2^L},$$
 (7.211)

(where here C = 10). We see that as long as p < 1/10, we can make the failure probability as small as we please by adding enough levels to the code.

We may write

$$p^{(L)} \le p_o \left(\frac{p}{p_o}\right)^{2^L},\tag{7.212}$$

where $p_o = \frac{1}{10}$ is an estimate of the *threshold* error probability that can be tolerated (we will obtain better codes and better estimates of this threshold below). Note that to obtain

$$p^{(L)} < \varepsilon, \tag{7.213}$$

we may choose the block size $n = 5^L$ so that

$$n \le \left[\frac{\log(p_o/\varepsilon)}{\log(p_o/p)}\right]^{\log_2 5}.$$
(7.214)

In principle, the concatenated code at a high level could fail with many fewer than n/10 errors, but these would have to be distributed in a highly conspiratorial fashion that is quite unlikely for n large.

The concatenated encoding of an unknown quantum state can be carried out level by level. For example to encode $a|0\rangle + b|1\rangle$ in the [[25, 1, 9]] block, we could first prepare the state $a|\bar{0}\rangle + b|\bar{1}\rangle$ in the five qubit block, using the encoding circuit described earlier, and also prepare four five-qubit blocks in the state $|\bar{0}\rangle$. The $a|\bar{0}\rangle + |\bar{1}\rangle$ can be encoded at the next level by executing the encoded circuit yet again, but this time with all gates replaced by encoded gates acting on five-qubit blocks. We will see in the next chapter how these encoded gates are constructed.

7.15.2 Toric codes

The toric codes are another family of codes that, like concatenated codes, offer much better performance than would be expected on the basis of their distance. They'll be described by Professor Kitaev (who discovered them).

7.15.3 Reed–Muller codes

Another way to construct codes that can correct many errors is to invoke the CSS construction. Recall, in particular, the special case of that construction that applies to a classical code C that is contained in its dual code (we

then say that C is "weakly self-dual"). In the CSS construction, there is a codeword associated with each coset of C in C^{\perp} . Thus we obtain an [[n, k, d]] quantum code, where n is the length of C, d is (at least) the distance of C^{\perp} , and $k = \dim C^{\perp} - \dim C$. Therefore, for the construction of CSS codes that correct many errors, we seek weakly self-dual classical codes with a large minimum distance.

One class of weakly self-dual classical codes are the Reed-Muller codes. Though these are not especially efficient, they are very convenient, because they are easy to encode, recovery is simple, and it is not difficult to explain their mathematical structure.⁴

To prepare for the construction of Reed-Muller codes, consider Boolean functions on m bits,

$$f: \{0,1\}^m \to \{0,1\}$$
. (7.215)

There are 2^{2^m} such functions forming what we may regard as a binary vector space of dimension 2^m . It will be useful to have a basis for this space. Recall (§6.1), that any Boolean function has a disjunctive normal form. Since the NOT of a bit x is 1 - x, and the OR of two bits x and y can be expressed as

$$x \lor y = x + y - xy$$
, (7.216)

any of the Boolean functions can be expanded as a polynomial in the *m* binary variables $x_{m-1}, x_{m-2}, \ldots, x_1, x_0$. A basis for the vector space of polynomials consists of the 2^m functions

1,
$$x_i, x_i x_j, x_i x_j x_k, \dots,$$
 (7.217)

(where, since $x^2 = x$, we may choose the factors of each monomial to be distinct). Each such function f can be represented by a binary string of length 2^m , whose value in the position labeled by the binary string $x_{m-1}x_{m-2}\ldots x_1x_0$

⁴See, *e.g.*, MacWilliams and Sloane, Chapter 13.

is $f(x_{m-1}, x_{m-2}, ..., x_1, x_0)$. For example, for m = 3,

$$1 = (1111111)$$

$$x_{0} = (10101010)$$

$$x_{1} = (11001100)$$

$$x_{2} = (11110000)$$

$$x_{0}x_{1} = (10001000)$$

$$x_{0}x_{2} = (10100000)$$

$$x_{1}x_{2} = (11000000)$$

$$x_{0}x_{1}x_{2} = (10000000)$$
. (7.218)

A subspace of this vector space is obtained if we restrict the degree of the polynomial to r or less. This subspace is the Reed-Muller (or RM) code, denoted R(r, m). Its length is $n = 2^m$ and its dimension is

$$k = 1 + \binom{m}{1} + \binom{m}{2} + \ldots + \binom{m}{r}.$$
 (7.219)

Some special cases of interest are:

- R(0,m) is the length- 2^m repetition code.
- R(m-1,m) is the dual of the repetition code, the space on all length- 2^m even-weight strings.
- R(1,3) is the n = 8, k = 4 code spanned by $1, x_0, x_1, x_2$; it is in fact the [8, 4, 4] extended Hamming code that we have already discussed.
- More generally, R(m-2,m) is a d = 4 extended Hamming code for each $m \ge 3$. If we puncture this code (remove the last bit from all codewords) we obtain the $[n = 2^m - 1, k = n - m, d = 3]$ perfect Hamming code.
- R(1,m) has $d = 2^{m-1} = \frac{1}{2}n$ and k = m. It is the dual of the extended Hamming code, and is known as a "first-order" Reed-Muller code. It is of considerable practical interest in its own right, both because of its large distance and because it is especially easy to decode.

We can compute the distance of the code R(r, m) by invoking induction on m. First we must determine how R(m + 1, r) is related to R(m, r). A function of x_m, \ldots, x_0 can be expressed as

$$f(x_m, \dots, x_0) = g(x_{m-1}, \dots, x_0) + x_m h(x_{m-1}, \dots, x_0) ,$$
(7.220)

and if f has degree r, then g must be of degree r and h of degree r-1. Regarding f as a vector of length 2^{m+1} , we have

$$f = (g|g) + (h|0) \tag{7.221}$$

where g, h are vectors of length 2^m . Consider the distance between f and

$$f' = (g'|g') + (h'|0) . (7.222)$$

For h = h' and $f \neq f'$ this distance is $\operatorname{wt}(f - f') = 2 \cdot \operatorname{wt}(g - g') \geq 2 \cdot \operatorname{dist}(R(r,m))$; for $h \neq h'$ it is at least $\operatorname{wt}(h - h') \geq \operatorname{dist}(R(r-1,m))$. If d(r,m) denotes the distance of R(r,m), then we see that

$$d(r, m+1) = \min\left(2 \ d(r, m), d(r-1, m)\right) \ . \tag{7.223}$$

Now we can show that $d(r, m) = 2^{m-r}$ by induction on m. To start with, we check that $d(r, m = 1) = 2^{1-r}$ for r = 0, 1; R(1, 1) is the space of all length 2 strings, and R(0, 1) is the length-2 repetition code. Next suppose that $d = 2^{m-r}$ for all $m \leq M$ and $0 \leq r \leq m$. Then we infer that

$$d(r, m+1) = \min(2^{m-r+1}, 2^{m-r+1}) = 2^{m-r+1}, \qquad (7.224)$$

for each $1 \leq r \leq m$. It is also clear that d(m + 1, m + 1) = 1, since R(m + 1, m + 1) is the space of all binary strings of length 2^{m+1} , and that $d(0, m + 1) = 2^{m+1}$, since R(0, m + 1) is the length- 2^{m+1} repetition code. This completes the inductive step, and proves $d(r, m) = 2^{m-r}$.

It follows, in particular, that R(m-1,m) has distance 2, and therefore that the dual of R(r,m) is R(m-r-1,m). First we notice that the binomial coefficients $\binom{m}{j}$ sum to 2^m , so that R(m-r-1) has the right dimension to be $R(r,m)^{\perp}$. It suffices, then, to show that R(m-r-1) is contained in R(r,m). But if $f \in R(r,m)$ and $g \in R(m-r-1,m)$, their product is a polynomial of degree at most m-1, and is therefore in R(m-1,m). Each
vector in R(m-1,m) has even weight, so the inner product $f \cdot g$ vanishes; hence g is in the dual $R(v,m)^{\perp}$. This shows that

$$R(r,m)^{\perp} = R(m-r-1,m).$$
(7.225)

It is because of this nice duality property that Reed–Muller codes are wellsuited for the CSS construction of quantum codes.

In particular, the Reed–Muller code is weakly self-dual for $r \leq m-r-1$, or $2r \geq m-1$, and self-dual for 2r = m-1. In the self-dual case, the distance is

$$d = 2^{m-r} = 2^{\frac{1}{2}(m+1)} = \sqrt{2n} , \qquad (7.226)$$

and the number of encoded bits is

$$k = \frac{1}{2}n = 2^{m-1} . (7.227)$$

These self-dual codes, for m = 3, 5, 7, have parameters

 $[8,4,4], \quad [32,16,8], \quad [128,64,16] . \tag{7.228}$

(The [8, 4, 4] code is the extended Hamming code as we have already noted.) Associated with these self-dual codes are the k = 0 quantum codes with parameters

$$[[8,0,4]], \quad [[32,0,8]], \quad [[128,0,16]], \quad (7.229)$$

and so forth.

One way to obtain a k = 1 quantum code is to *puncture* the self-dual Reed-Muller code, that is, to delete one of the $n = 2^m$ bits from the code. (It turns out not to matter *which* bit we delete.) The classical punctured code has parameters $n = 2^m - 1$, $d = 2^{\frac{1}{2}(m-1)} - 1 = \sqrt{2(n+1)} - 1$, and $k = \frac{1}{2}(n+1)$. Furthermore, the dual of the punctured code is its even subcode. (The even subcode consists of those RM codewords for which the bit removed by the puncture is zero, and it follows from the self-duality of the RM code that these are orthogonal to all the words (both odd and even weight) of the punctured code.) From these punctured codes, we obtain, via the CSS construction, k = 1 quantum codes with parameters

$$[[7,1,3]], \quad [[31,1,7]], \quad [[127,1,15]], \quad (7.230)$$

and so forth. The [7, 4, 3] Hamming code is obtained by puncturing the [8, 4, 4] RM code, and the corresponding [7, 1, 3] QECC is of course Steane's code. These QECC's have a distance that increases like the square root of their length.

These k = 1 codes are not among the most efficient of the known QECC's. Nevertheless they are of special interest, since their properties are especially conducive to implementing fault-tolerant quantum gates on the encoded data, as we will see in Chapter 8. In particular, one useful property of the self-dual RM codes is that they are "doubly even" — all codewords have a weight that is an integral multiple of four.

Of course, we can also construct quantum codes with k > 1 by applying the CSS construction to the RM codes. For example R(3, 6), with parameters

$$n = 2^{m} = 64$$

$$d = 2^{m-r} = 8$$

$$k = 1 + 6 + {6 \choose 2} + {6 \choose 3} = 1 + 6 + 15 + 20 = 42 , \qquad (7.231)$$

is dual to R(2,6), with parameters

$$n = 2^{m} = 64$$

$$d = 2^{m-r} = 16$$

$$k = 1 + 6 + \binom{6}{2} = 1 + 6 + 15 = 22 , \qquad (7.232)$$

and so the CSS construction yields a QECC with parameters

$$[[64, 20, 8]] . (7.233)$$

Many other weakly self-dual codes are known and can likewise be employed.

7.15.4 The Golay Code

From the perspective of pure mathematics, the most important error-correcting code (classical or quantum) ever discovered is also one of the first ever described in a published article — the Golay code. Here we will briefly describe the Golay code, as it too can be transformed into a nice QECC via the CSS construction. (Perhaps this QECC is not really important enough to deserve a section of this chapter; still, I have included it just for fun.)

The (extended) Golay code is a self-dual [24, 12, 8] classical code. If we puncture it (remove any one of its 24 bits), we obtain the [23, 12, 7] Golay code, which can correct three errors. This code is actually perfect, as it saturates the sphere-packing bound:

$$1 + \binom{23}{1} + \binom{23}{2} + \binom{23}{3} = 2^{11} = 2^{23-12}.$$
 (7.234)

In fact, perfect codes that correct more than one error are extremely rare. It can be shown⁵ that the *only* perfect codes (linear or nonlinear) over *any* finite field that can correct more than one error are the [23, 12, 7] code and one other binary code discovered by Golay, with parameters [11, 6, 5].

The [24, 12, 8] Golay code has a very intricate symmetry. The symmetry is characterized by its automorphism group — the group of permutations of the 24 bits that take codewords to codewords. This is the Mathieu group M_{24} , a sporadic simple group of order 244,823,040 that was discovered in the 19th century.

The $2^{12} = 4096$ codewords have the weight distribution (in an obvious notation)

$$0^{1}8^{759}12^{2576}16^{759}24^{1} (7.235)$$

Note in particular that each weight is a multiple of 4 (the code is doubly even). What is the significance of the number 759 (= 3.11.23)? In fact it is

$$\binom{24}{5} / \binom{8}{5} = 759, \tag{7.236}$$

and it arises for this combination reason: with each weight-8 codeword we associate the eight-element set ("octad") where the codeword has its support. Each 5-element subset of the 24 bits is contained in exactly one octad (a reflection of the code's large symmetry).

What makes the Golay code important in mathematics? Its discovery in 1949 set in motion a sequence of events that led, by around 1980, to a complete classification of the finite simple groups. This classification is one of the greatest achievements of 20th century mathematics.

(A group is simple if it contains no nontrivial normal subgroup. The finite simple groups may be regarded as the building blocks of all finite groups in

⁵MacWilliams and Sloane §6.10.

the sense that for any finite group G there is a unique decomposition of the form

$$G \equiv G_0 \supseteq G_1 \supseteq G_2 \ge \ldots \supseteq G_n, \tag{7.237}$$

where each G_{j+1} is a normal subgroup of G_j , and each quotient group G_j/G_{j+1} is simple. The finite simple groups can be classified into various infinite families, plus 26 additional "sporadic" simple groups that resist classification.)

The Golay code led Leech, in 1964, to discover an extraordinarily close packing of spheres in 24 dimensions, known as the *Leech Lattice* Λ . The lattice points (the centers of the spheres) are 24-component integer-valued vectors with these properties: to determine if $\vec{x} = (x_1, x_2, \ldots, x_{24})$ is contained in Λ , write each component x_j in binary notation,

$$x_j = \dots x_{j3} x_{j2} x_{j1} x_{j0} . (7.238)$$

Then $\vec{x} \in \Lambda$ if

- (i) The x_{i0} 's are either all 0's or all 1's.
- (ii) The x_{j2} 's are an even parity 24-bit string if the x_{j0} 's are 0, and an odd parity 24-bit string if the x_{j0} 's are 1.
- (iii) The x_{j1} 's are a 24-bit string contained in the Golay code.

When these rules are applied, a negative number is represented by its binary complement, *e.g.*

$$\begin{array}{l}
-1 = \dots 11111 , \\
-2 = \dots 11110 , \\
-3 = \dots 11101 , \\
& \text{etc.} \\
\end{array} (7.239)$$

We can easily check that Λ is a lattice; that is, it is closed under addition. (Bits other than the last three in the binary expansion of the x_j 's are unrestricted).

We can now count the number of nearest neighbors to the origin (or the number of spheres that touch any given sphere). These points are all $(distance)^2 = 32$ away from the origin:

$$\begin{array}{rcccc} (\pm 2)^8 & : & 2^7 \cdot 759 \\ (\pm 3)(\mp 1)^{23} & : & 2^{12} \cdot 24 \\ (\pm 4)^2 & : & 2^2 \cdot \binom{24}{2} \end{array} .$$
 (7.240)

That is, there are $759 \cdot 2^7$ neighbors that have eight components with the values ± 2 — their support is on one of the 759 weight-8 Golay codewords, and the number of – signs must be even. There are $2^{12} \cdot 24$ neighbors that have one component with value ± 3 (this component can be chosen in 24 ways) and the remaining 23 components have the value (∓ 1). If, say, +3 is chosen, then the position of the +3, together with the position of the -1's, can be any of the 2^{11} Golay codewords with value 1 at the position of the +3. There are $2^2 \cdot \binom{24}{2}$ neighbors with two components each taking the value ± 4 (the signs are unrestricted). Altogether, the coordination number of the lattice is 196, 560.

The Leech lattice has an extraordinary automorphism group discovered by Conway in 1968. This is the finite subgroup of the 24-dimensional rotation group SO(24) that preserves the lattice. The order of this finite group (known as $\cdot 0$, or "dot oh") is

$$2^{22} \cdot 3^9 \cdot 5^4 \cdot 7^2 \cdot 11 \cdot 13 \cdot 23 = 8,315,553,613,086,720,000 \simeq 8.3 \times 10^{18}.$$
(7.241)

If its two element center is modded out, the sporadic simple group $\cdot 1$ is obtained. At the time of its discovery, $\cdot 1$ was the largest of the sporadic simple groups that had been constructed.

The Leech lattice and its automorphism group eventually (by a route that won't be explained here) led Griess in 1982 to the construction of the most amazing sporadic simple group of all (whose existence had been inferred earlier by Fischer and Griess). It is a finite subgroup of the rotation group in 196,883 dimensions, whose order is approximately 8.08×10^{53} . This behemoth known as F_1 has earned the nickname "the monster" (though Griess prefers to call it "the friendly giant".) It is the largest of the sporadic simple groups, and the last to be discovered.

Thus the classification of the finite simple groups owes much to (classical) coding theory, and to the Golay code in particular. Perhaps the theory of

QECC's can also bequeath to mathematics something of value and broad interest!

Anyway, since the (extended) [24, 12, 8] Golay code is self-dual, the [23, 12, 7] code obtained by puncturing it is weakly self dual; its [23, 11, 8] dual is its even subcode. From it, a [23, 1, 7] QECC can be constructed by the CSS method. This code is not the most efficient quantum code that can correct three errors (there is a [17, 1, 7] code that saturates the Rains bound), but it has especially nice properties that are conducive to fault-tolerant quantum computation, as we will see in Chapter 8.

7.16 The Quantum Channel Capacity

As we have formulated it up until now, our goal in constructing quantum error correcting codes has been to maximize the distance d of the code, given its length n and the number k of encoded qubits. Larger distance provides better protection against errors, as a distance d code can correct d-1 erasures, or (d-1)/2 errors at unknown locations. We have observed that "good" codes can be constructed, that maintain a finite rate k/n for n large, and correct a number of errors pn that scales linearly with n.

Now we will address a related but rather different question about the asymptotic performance of QECC's. Consider a superoperator \$ that acts on density operators in a Hilbert space \mathcal{H} . Now consider \$ acting independently each copy of \mathcal{H} contained in the *n*-fold tensor product

$$\mathcal{H}^{(n)} = \mathcal{H} \otimes \ldots \otimes \mathcal{H}. \tag{7.242}$$

We would like to select a code subspace $\mathcal{H}_{code}^{(n)}$ of $\mathcal{H}^{(n)}$ such that quantum information residing in $\mathcal{H}_{code}^{(n)}$ can be subjected to the superoperator

$$\$^{(n)} = \$ \otimes \ldots \otimes \$, \tag{7.243}$$

and yet can still be decoded with high fidelity.

The rate of a code is defined as

$$R = \frac{\log \mathcal{H}_{\text{code}}^{(n)}}{\log \mathcal{H}^{(n)}}; \qquad (7.244)$$

this is the number of qubits employed to carry one qubit of encoded information. The quantum channel capacity Q(\$) of the superoperator \$ is the

7.16. THE QUANTUM CHANNEL CAPACITY

maximum asymptotic rate at which quantum information can be sent over the channel with arbitrarily good fidelity. That is, Q(\$) is the largest number such that for any R < Q(\$) and any $\varepsilon > 0$, there is a code $\mathcal{H}_{code}^{(n)}$ with rate at least R, such that for any $|\psi\rangle \in \mathcal{H}_{code}^{(n)}$, the state ρ recovered after $|\psi\rangle$ passes through $\$^{(n)}$ has fidelity

$$F = \langle \psi | \boldsymbol{\rho} | \psi \rangle > 1 - \varepsilon. \tag{7.245}$$

Thus, Q(\$) is a quantum version of the capacity defined by Shannon for a classical noisy channel. As we have already seen in Chapter 5, this Q(\$) is not the only sort of capacity that can be associated with a quantum channel. It is also of considerable interest to ask about C(\$), the maximum rate at which *classical* information can be transmitted through a quantum channel with arbitrarily small probability of error. A formal answer to this question was formulated in §5.4, but only for a restricted class of possible encoding schemes; the general answer is still unknown. The quantum channel capacity Q(\$) is even less well understood than the classical capacity C(\$) of a quantum channel. Note that Q(\$) is not the same thing as the maximum asymptotic rate k/n that can be achieved by "good" [[n, k, d]] QECC's with positive d/n. In the case of the quantum channel capacity we need not insist that the code correct any possible distribution of pn errors, as long as the errors that cannot be corrected become highly atypical for n large.

Here we will mostly limit the discussion to two interesting examples of quantum channels acting on a single qubit — the quantum erasure channel (for which Q is exactly known), and the depolarizing channel (for which Q is still unknown, but useful upper and lower bounds can be derived).

What are these channels? In the case of the quantum erasure channel, a qubit transmitted through the channel either arrives intact, or (with probability p) becomes lost and is never received. We can find a unitary representation of this channel by embedding the qubit in the three-dimensional Hilbert space of a qubit with orthonormal basis $\{|0\rangle, |1\rangle, |2\rangle\}$. The channel acts according to

$$|0\rangle \otimes |0\rangle_{E} \rightarrow \sqrt{1-p}|0\rangle \otimes |0\rangle_{E} + \sqrt{p}|2\rangle \otimes |1\rangle_{E}, |1\rangle \otimes |0\rangle_{E} \rightarrow \sqrt{1-p}|1\rangle \otimes |0\rangle_{E} + \sqrt{p}|2\rangle \otimes |2\rangle_{E},$$
(7.246)

where $\{|0\rangle_E, |1\rangle_E, |2\rangle_E\}$ are mutually orthogonal states of the environment. The receiver can measure the observable $|2\rangle\langle 2|$ to determined whether the qubit is undamaged or has been "erased." The depolarizing channel (with error probability p) was discussed at length in §3.4.1. We see that, for $p \leq 3/4$, we may describe the fate of a qubit transmitted through the channel this way: with probability 1 - q(where q = 4/3p), the qubit arrives undamaged, and with probability q it is *destroyed*, in which case it is described by the random density matrix $\frac{1}{2}\mathbf{1}$.

Both the erasure channel and the depolarizing channel destroy a qubit with a specified probability. The crucial difference between the two channels is that in the case of the erasure channel, the receiver knows which qubits have been destroyed; in the case of the depolarizing channel, the damaged qubits carry no identifying marks, which makes recovery more challenging. Of course, for both channels, the sender has no way to know ahead of time which qubits will be obliterated.

7.16.1 Erasure channel

The quantum channel capacity of the erasure channel can be precisely determined. First we will derive an upper bound on Q, and then we will show that codes exist that achieve high fidelity and attain a rate arbitrarily close to the upper bound.

As the first step in the derivation of an upper bound on the capacity, we show that Q = 0 for $p > \frac{1}{2}$.

– Figure –

We observe that the erasure channel can be realized if Alice sends a qubit to Bob, and a third party Charlie decides at random to either *steal* the qubit (with probability p) or allow the qubit to pass unscathed to Bob (with probability 1-p).

If Alice sends a large number n of qubits, then about (1-p)n reach Bob, and pn are intercepted by Charlie. Hence for $p > \frac{1}{2}$, Charlie winds up in possession of more qubits than Bob, and if Bob can recover the quantum information encoded by Alice, then certainly Charlie can as well. Therefore, if Q(p) > 0 for $p > \frac{1}{2}$, Bob and Charlie can clone the unknown encoded quantum states sent by Alice, which is impossible. (Strictly speaking, they can clone with fidelity $F = 1 - \varepsilon$, for any $\varepsilon > 0$.) We conclude that Q(p) = 0for $p > \frac{1}{2}$.

7.16. THE QUANTUM CHANNEL CAPACITY

To obtain a bound on Q(p) in the case $p < \frac{1}{2}$, we will appeal to the following lemma. Suppose that Alice and Bob are connected by both a perfect noiseless channel and a noisy channel with capacity Q > 0. And suppose that Alice sends m qubits over the perfect channel and n qubits over the noisy channel. Then the number r of encoded qubits that Bob may recover with arbitrarily high fidelity must satisfy

$$r \le m + Qn. \tag{7.247}$$

We derive this inequality by noting that Alice and Bob can simulate the m qubits sent over the perfect channel by sending m/Q over the noisy channel and so achieve a rate

$$R = \frac{r}{m/Q+n} = \left(\frac{r}{m+Qn}\right)Q,\tag{7.248}$$

over the noisy channel. Were r to exceed m + Qn, this rate R would exceed the capacity, a contradiction. Therefore eq. (7.247) is satisfied.

How consider the erasure channel with error probability p_1 , and suppose $Q(p_1) > 0$. Then we can bound $Q(p_2)$ for $p_2 \leq p_1$ by

$$Q(p_2) \le 1 - \frac{p_2}{p_1} + \frac{p_2}{p_1}Q(p_1).$$
 (7.249)

(In other words, if we plot Q(p) in the (p, Q) plane, and we draw a straight line segment from any point (p_1, Q_1) on the plot to the point (p = 0, Q = 1), then the curve Q(p) must lie on or below the segment in the interval $0 \le p \le p_1$; if Q(p) is twice differentiable, then its second derivative cannot be positive.) To obtain this bound, imagine that Alice sends n qubits to Bob, knowing ahead of time that $n(1 - p_2/p_1)$ specified qubits will arrive safely. The remaining $n(p_2/p_1)$ qubits are erased with probability p_1 . Therefore, Alice and Bob are using both a perfect channel and an erasure channel with erasure probability p_1 ; eq. (7.247) holds, and the rate R they can attain is bounded by

$$R \le 1 - \frac{p_2}{p_1} + \frac{p_2}{p_1}Q(p_1). \tag{7.250}$$

On the other hand, for n large, altogether about np_2 qubits are erased, and $(1 - p_2)n$ arrive safely. Thus Alice and Bob have an erasure channel with erasure probability p_2 , except that they have the additional advantage of

knowing ahead of time that some of the qubits that Alice sends are invulnerable to erasure. With this information, they can be no worse off than without it; eq. (7.249) then follows. The same bound applies to the depolarizing channel as well.

Now, the result Q(p) = 0 for p > 1/2 can be combined with eq. (7.249). We conclude that the curve Q(p) must be on or below the straight line connecting the points (p = 0, Q = 1) and (p = 1/2, Q = 0), or

$$Q(p) \le 1 - 2p, \quad 0 \le p \le \frac{1}{2}.$$
 (7.251)

In fact, there are stabilizer codes that actually attain the rate 1 - 2p for $0 \le p \le 1/2$. We can see this by borrowing an idea from Claude Shannon, and averaging over random stabilizer codes. Imagine choosing, in succession, altogether n - k stabilizer generators. Each is selected from among the 4^n Pauli operators, where all have equal a priori probability, except that each generator is required to commute with all generators chosen in previous rounds.

Now Alice uses this stabilizer code to encode an arbitrary quantum state in the 2^k -dimensional code subspace, and sends the *n* qubits to Bob over an erasure channel with erasure probability *p*. Will Bob be able to recover the state sent by Alice?

Bob replaces each erased qubit by a qubit in the state $|0\rangle$, and then proceeds to measure all n - k stabilizer generators. From this syndrome measurement, he hopes to infer the Pauli operator E acting on the replaced qubits. Once E is known, we can apply E^{\dagger} to recover a perfect duplicate of the state sent by Alice. For n large, the number of qubits that Bob must replace is about pn, and he will recover successfully if there is a unique Pauli operator E that can produce the syndrome that he finds. If more than one Pauli operator acting on the replaced qubits has this same syndrome, then recovery may fail.

How likely is failure? Since there are about pn replaced qubits, there are about 4^{pn} Pauli operators with support on these qubits. Furthermore, for any particular Pauli operator \boldsymbol{E} , a random stabilizer code generates a random syndrome — each stabilizer generator has probability 1/2 of commuting with \boldsymbol{E} , and probability 1/2 of anti-commuting with \boldsymbol{E} . Therefore, the probability that two Pauli operators have the same syndrome is $(1/2)^{n-k}$.

There is at least one particular Pauli operator acting on the replaced qubits that has the syndrome found by Bob. But the probability that another Pauli operator has this same syndrome (and hence the probability of a recovery failure) is no worse than

$$P_{\text{fail}} \le 4^{pn} \left(\frac{1}{2}\right)^{n-k} = 2^{-n(1-2p-R)}.$$
 (7.252)

where R = k/n is the rate. Eq. (7.252) bounds the failure probability if we *average* over all stabilizer codes with rate R; it follows that at least one particular stabilizer code must exist whose failure probability also satisfies the bound.

For that particular code, P_{fail} gets arbitrarily small as $n \to \infty$, for any rate $R = 1-2p-\delta$ strictly less than 1-2p. Therefore R = 1-2p is asymptotically attainable; combining this result with the inequality eq. (7.251) we obtain the capacity of the quantum erasure channel:

$$Q(p) = 1 - 2p, \quad 0 \le p \le \frac{1}{2}$$
 (7.253)

If we wanted assurance that a distinct syndrome could be assigned to all ways of damaging pn erased qubits, then we would require an [[n, k, d]]quantum code with distance d > pn. Our Gilbert–Varshamov bound of §7.14 guarantees the existence of such a code for

$$R < 1 - H_2(p) - p \log_2 3. \tag{7.254}$$

This rate can be achieved by a code that recovers from any of the possible ways of erasing up to pn qubits. It lies strictly below the capacity for p > 0, because to achieve high average fidelity, it suffices to be able to correct the *typical* erasures, rather than all possible erasures.

7.16.2 Depolarizing channel

The capacity of the depolarizing channel is still not precisely known, but we can obtain some interesting upper and lower bounds.

As for the erasure channel, we can find an upper bound on the capacity by invoking the no-cloning theorem. Recall that for the depolarizing channel with error probability p < 3/4, each qubit either passes safely with probability 1 - 4/3p, or is randomized (replaced by the maximally mixed state $\rho = \frac{1}{2}\mathbf{1}$) with probability q = 4/3p. An eavesdropper Charlie, then, can simulate the channel by intercepting qubits with probability q, and replacing each stolen qubit with a maximally mixed qubit. For q > 1/2, Charlie steals more than half the qubits and is in a better position than Bob to decode the state sent by Alice. Therefore, to disallow cloning, the rate at which quantum information is sent from Alice to Bob must be strictly zero for q > 1/2or p > 3/8:

$$Q(p) = 0, \quad p > \frac{3}{8}.$$
 (7.255)

In fact we can obtain a stronger bound by noting that Charlie can choose a better eavesdropping strategy – he can employ the optimal *approximate* cloner that you studied in a homework problem. This device, applied to each qubit sent by Alice, replaces it by two qubits that each approximate the original with fidelity F = 5/6, or

$$|\psi\rangle\langle\psi| \rightarrow \left[(1-q)|\psi\rangle\langle\psi| + q\frac{1}{2}\mathbf{1}\right]^{\otimes 2},$$
 (7.256)

where F = 5/6 = 1 - 1/2q. By operating the cloner, both Charlie and Bob can receive Alice's state transmitted through the q = 1/3 depolarizing channel. Therefore, the attainable rate must vanish; otherwise, by combining the approximate cloner with quantum error correction, Bob and Charlie would be able to clone Alice's unknown state *exactly*. We conclude that the capacity vanishes for q > 1/3 or p > 1/4:

$$Q(p) = 0, \quad p > \frac{1}{4}.$$
 (7.257)

Invoking the bound eq. (7.249) we infer that

$$Q(p) \le 1 - 4p, \quad 0 \le p \le \frac{1}{4}.$$
 (7.258)

This result actually coincides with our bound on the rate of [[n, k, d]] codes with $k \ge 1$ and $d \ge 2pn + 1$ found in §7.8. A bound on the capacity is *not* the same thing as a bound on the allowable error probability for an [[n, k, d]]code (and in the latter case the Rains bound is tighter). Still, the similarity of the two results bound may not be a complete surprise, as both bounds are derived from the no-cloning theorem.

We can obtain a lower bound on the capacity by estimating the rate that can be attained through random stabilizer coding, as we did for the erasure channel. Now, when Bob measures the n-k (randomly chosen, commuting) stabilizer generators, he hopes to obtain a syndrome that points to a unique one among the typical Pauli error operators that can arise with nonnegligible probability when the depolarizing channel acts on the n qubits sent by Alice. The number $N_{\rm typ}$ of typical Pauli operators with total probability $1-\varepsilon$ can be bounded by

$$N_{\rm typ} \le 2^{n(H_2(p) + p\log_2 3 + \delta)},\tag{7.259}$$

for any $\delta, \varepsilon > 0$ and *n* sufficiently large. Bob's attempt at recovery can fail if another among these typical Pauli operators has the same syndrome as the actual error operator. Since a random code assigns a random (n - k)-bit syndrome to each Pauli operator, the failure probability can be bounded as

$$P_{\text{fail}} \le 2^{n(H_2(p) + p \log_2 3 + \delta)} 2^{k-n} + \varepsilon .$$
(7.260)

Here the second term bounds the probability of an atypical error, and the first bounds the probability of an ambiguous syndrome in the case of a typical error. We see that the failure probability, averaged over random stabilizer codes, becomes arbitrarily small for large n, for any $\delta' < 0$ and rate R such that

$$R \equiv \frac{k}{n} < 1 - H_2(p) - p \log_2 3 - \delta'.$$
(7.261)

If the failure probability, averaged over codes, is small, there is a particular code with small failure probability, and we conclude that the rate R is attainable; the capacity of the depolarizing channel is bounded below as

$$Q(p) \geq 1 - H_2(p) - p \log_2 3$$
. (7.262)

Not coincidentally, the rate attainable by random coding agrees with the asymptotic form of the quantum Hamming upper bound on the rate of nondegenerate [[n, k, d]] codes with d > 2pn; we arrive at both results by assigning a distinct syndrome to each of the typical errors. Of course, the Gilbert– Varshamov lower bound on the rate of [[n, k, d]] codes lies below Q(p), as it is obtained by demanding that the code can correct *all* the errors of weight pn or less, not just the typical ones.

This random coding argument can also be applied to a somewhat more general channel, in which X, Y, and Z errors occur at different rates. (We'll

call this a "Pauli channel.") If an X error occurs with probability p_X , a Y error with probability p_Y , a Z error with probability p_Z , and no error with probability $p_I \equiv 1 - p_X - p_Y - p_Z$, then the number of typical errors on n qubits is

$$\frac{n!}{(p_X n)!(p_Y n)!(p_Z n)!(p_I n)!} \sim 2^{nH(p_I, p_X, p_Y, p_Z)},$$
(7.263)

where

$$H \equiv H(p_I, p_X, p_Y, p_Z) = -p_I \log_2 p_I - p_X \log_2 p_X - p_Y \log_2 p_Y - p_Z \log_2 p_Z,$$
(7.264)

is the Shannon entropy of the probability distribution $\{p_I, p_X, p_Y, p_Z\}$. Now we find

$$Q(p_I, p_X, p_Y, p_Z) \geq 1 - H(p_I, p_X, p_Y, p_Z); \qquad (7.265)$$

if the rate R satisfies R < 1 - H, then again it is highly unlikely that a single syndrome of a random stabilizer code will point to more than one typical error operator.

7.16.3 Degeneracy and capacity

Our derivation of a lower bound on the capacity of the depolarizing channel closely resembles the argument in §5.1.3 for a lower bound on the capacity of the classical binary symmetric channel. In the classical case, there was a matching upper bound. If the rate were larger, then there would not be enough syndromes to attach to all of the typical errors.

In the quantum case, the derivation of the matching upper bound does not carry through, because a quantum code can be degenerate. We may not need a distinct syndrome for each typical error, as some of the possible errors could act trivially on the code subspace. Indeed, not only does the derivation fail; the matching upper bound is actually false – rates exceeding $1 - H_2(p) - p \log_2 3$ actually can be attained.⁶

Shor and Smolin investigated the rate that can be achieved by concatenated codes, where the outer code is a random stabilizer code, and the inner

⁶P.W. Shor and J.A. Smolin, "Quantum Error-Correcting Codes Need Not Completely Reveal the Error Syndrome" quant-ph/9604006; D.P. DiVincen, P.W. Shor, and J.A. Smolin, "Quantum Channel Capacity of Very Noisy Channels," quant-ph/9706061.

code is a degenerate code with a relatively small block size. Their idea is that the degeneracy of the inner code will allow enough typical errors to act trivially in the code space that a higher rate can be attained than through random coding alone.

To investigate this scheme, imagine that encoding and decoding are each performed in two stages. In the first stage, using the (random) outer code that she and Bob have agreed on, Alice encodes the state that she has selected in a large *n*-qubit block. In the second stage, Alice encodes each of these *n*-qubits in a block of *m* qubits, using the inner code. Similarly, when Bob receives the nm qubits, he first decodes each inner block of *m*, and then subsequently decodes the block of *n*.

We can evidently describe this procedure in an alternative language — Alice and Bob are using just the outer code, but the qubits are being transmitted through a composite channel.

This modified channel consists (as shown) of: first the inner encoder, then propagation through the original noisy channel, and finally inner decoding and inner recovery. The rate that can be attained through the original channel, via concatenated coding, is the same as the rate that can be attained through the modified channel, via random coding.

Specifically, suppose that the inner code is an m-qubit repetition code, with stabilizer

$$Z_1 Z_2, \ Z_1 Z_3, \ Z_1 Z_4, \dots, Z_1 Z_m.$$
 (7.266)

This is not much of a quantum code; it has distance 1, since it is insensitive to phase errors — each \mathbf{Z}_j commutes with the stabilizer. But in the present context its important feature is it high degeneracy, all \mathbf{Z}_i errors are equivalent.

The encoding (and decoding) circuit for the repetition code consists of just m-1 CNOT's, so our composite channel looks like (in the case m=3)

– Figure –

where denotes the original noisy channel. (We have also suppressed the final recovery step of the decoding; *e.g.*, if the measured qubits both read 1, we should flip the data qubit. In fact, to simplify the analysis of the composite channel, we will dispense with this step.)

Since we recall that a CNOT propagates bit flips forward (from control to target) and phase flips backward (from target to control), we see that for each possible measurement outcome of the auxiliary qubits, the composite channel is a Pauli channel. If we imagine that this measurement of the m-1 inner block qubits is performed for each of the n qubits of the outer block, then Pauli channels act independently on each of the n qubits, but the channels acting on different qubits have different parameters (error probabilities $p_I^{(i)}, p_X^{(i)}, p_Y^{(i)}, p_Z^{(i)}$ for the *i*th qubit). Now the number of typical error operators acting on the n qubits is

$$2^{\sum_{i=1}^{n} H_i}$$
 (7.267)

where

$$H_i = H(p_I^{(i)}, p_X^{(i)}, p_Y^{(i)}, p_Z^{(i)}),$$
(7.268)

is the Shannon entropy of the Pauli channel acting on the ith qubit. By the law of large numbers, we will have

$$\sum_{i=1}^{n} H_i = n \langle H \rangle, \tag{7.269}$$

for large n, where $\langle H \rangle$ is the Shannon entropy, averaged over the 2^{m-1} possible classical outcomes of the measurement of the extra qubits of the inner code. Therefore, the rate that can be attained by the random outer code is

$$R = \frac{1 - \langle H \rangle}{m},\tag{7.270}$$

(we divide by m, because the concatenated code has a length m times longer than the random code).

Shor and Smolin discovered that there are repetition codes (values of m) for which, in a suitable range of p, $1-\langle H \rangle$ is positive while $1-H_2(p)-p \log_2 3$ is negative. In this range, then, the capacity Q(p) is nonzero, showing that the lower bound eq. (7.262) is not tight.

7.16. THE QUANTUM CHANNEL CAPACITY

A nonvanishing asymptotic rate is attainable through random coding for $1 - H_2(p) - p \log_2 3 > 0$, or $p < p_{\max} \simeq .18929$. If a random outer code is concatenated with a 5-qubit inner repetition code (m = 5 turns out to be the optimal choice), then $1 - \langle H \rangle > 0$ for $p < p'_{\max} \simeq .19036$; the maximum error probability for which a nonzero rate is attainable increases by about 0.6%. It is not obvious that the concatenated code should outperform the random code in this range of error probability, though as we have indicated, it might have been expected because of the (phase) degeneracy of the repetition code. Nor is it obvious that m = 5 should be the best choice, but this can be verified by an explicit calculation of $\langle H \rangle$.⁷

The depolarizing channel is one of the very simplest of quantum channels. Yet even for this case, the problem of characterizing and calculating the capacity is largely unsolved. This example illustrates that, due to the possibility of degenerate coding, the capacity problem is considerably more subtle for quantum channels than for classical channels.

We have seen that (if the errors are well described by the depolarizing channel), quantum information can be recovered from a quantum memory with arbitrarily high fidelity, as long as the probability of error per qubit is less than 19%. This is an improvement relative to the 10% error rate that we found could be handled by concatenation of the [[5, 1, 3]] code. In fact [[n, k, d]] codes that can recover from any distribution of up to pn errors do not exist for p > 1/6, according to the Rains bound. Nonzero capacity is possible for error rates between 16.7% and 19% because it is sufficient for the QECC to be able to correct the typical errors rather than all possible errors.

However, the claim that recovery is possible even if 19% of the qubits sustain damage is highly misleading in an important respect. This result applies if encoding, decoding, and recovery can be executed flawlessly. But these operations are actually very intricate quantum computations that in practice will certainly be susceptible to error. We will not fully understand how well coding can protect quantum information from harm until we have learned to design an error recovery protocol that is robust even if the execution of the protocol is flawed. Such *fault-tolerant* protocols will be developed in Chapter 8.

⁷In fact a very slight further improvement can be achieved by concatenating a random code with the 25-qubit generalized Shor code described in the exercises – then a nonzero rate is attainable for $p < p''_{\text{max}} \simeq .19056$ (another 0.1% better than the maximum tolerable error probability with repetition coding).

7.17 Summary

Quantum error-correcting codes: Quantum error correction can protect quantum information from both decoherence and "unitary errors" due to imperfect implementations of quantum gates. In a (binary) quantum errorcorrecting code (QECC), the 2^k-dimensional Hilbert space \mathcal{H}_{code} of k encoded qubits is embedded in the 2ⁿ-dimensional Hilbert space of n qubits. Errors acting on the n qubits are reversible provided that $\langle \psi | \mathbf{M}_{\nu}^{\dagger} \mathbf{M}_{\mu} | \psi \rangle / \langle \psi | \psi \rangle$ is independent of $|\psi\rangle$ for any $|\psi\rangle \in \mathcal{H}_{code}$ and any two Kraus operators $\mathbf{M}_{\mu,\nu}$ occuring in the expansion of the error superoperator. The recovery superoperator transforms entanglement of the environment with the code block into entanglement of the environment with an ancilla that can then be discarded.

Quantum stabilizer codes: Most QECC's that have been constructed are stabilizer codes. A binary stabilizer code is characterized by its stabilizer S, an abelian subgroup of the n-qubit Pauli group $G_n = \{I, X, Y, Z\}^{\otimes n}$ (where X, Y, Z are the single-qubit Pauli operators). The code subspace is the simultaneous eigenspace with eigenvalue one of all elements of S; if S has n-k independent generators, then there are k encoded qubits. A stabilizer code can correct each error in a subset \mathcal{E} of G_n if for each $\mathbf{E}_a, \mathbf{E}_b \in \mathcal{E}$, $E_a^{\dagger}E_b$ either lies in the stabilizer S or outside of the normalizer S^{\perp} of the stabilizer. If some $E_a^{\dagger} E_b$ is in S for $E_{a,b} \in \mathcal{E}$ the code is *degenerate*; otherwise it is nondegenerate. Operators in $S^{\perp} \setminus S$ are "logical" operators that act on encoded quantum information. The stabilizer S can be associated with an additive code over the finite field GF(4) that is self-orthogonal with respect to a symplectic inner product. The *weight* of a Pauli operator is the number of qubits on which its action is nontrivial, and the distance d of a stabilizer code is the minimum weight of an element of $S^{\perp} \setminus S$. A code with length n, k encoded qubits, and distance d is called an [[n, k, d]] quantum code. If the code enables recovery from any error superoperator with support on Pauli operators of weight t or less, we say that the code "can correct t errors." A code with distance d can correct $\left[\frac{d-1}{2}\right]$ in unknown locations or d-1 errors in known locations. "Good" families of stabilizer codes can be constructed in which d/n and k/n remain bounded away from zero as $n \to \infty$.

Examples: The code of minimal length that can correct one error is a [[5, 1, 3,]] quantum code associated with a classical GF(4) Hamming code. Given a classical linear code C_1 and subcode $C_2 \subseteq C_1$, a Calderbank-Shor-Steane (CSS) quantum code can be constructed with $k = \dim(C_1) - \dim(C_2)$ encoded qubits. The distance d of the CSS code satisfies $d \geq \min(d_1, d_2^{\perp})$, where d_1 is the distance of C_1 and d_2^{\perp} is the distance of C_2^{\perp} , the dual of C_2 . The simplest CSS code is a [[7, 1, 3]] quantum code constructed from the [7, 4, 3] classical Hamming code and its even subcode. An [[n_1 , 1, d_1]] quantum code can be *concatenated* with an [[n_2 , 1, d_2]] code to obtain a degenerate [[n_1n_2 , 1, d_1]] code with $d \ge d_1d_2$.

Quantum channel capacity: The quantum channel capacity of a superoperator (noisy quantum channel) is the maximum rate at which quantum information can be transmitted over the channel and decoded with arbitrarily good fidelity. The capacity of the binary quantum erasure channel with erasure probability p is Q(p) = 1 - 2p, for $0 \le p \le 1/2$. The capacity of the binary depolarizing channel is no yet known. The problem of calculating the capacity is subtle because the optimal code may be degenerate; in particular, random codes do not attain an asymptotically optimal rate over a quantum channel.

7.18 Exercises

7.1 Phase error-correcting code

- a) Construct stabilizer generators for an n = 3, k = 1 code that can correct a single bit flip; that is, ensure that recovery is possible for any of the errors in the set $\mathcal{E} = \{III, XII, IXI, IIX\}$. Find an orthonormal basis for the two-dimensional code subspace.
- b) Construct stabilizer generators for an n = 3, k = 1 code that can correct a single phase error; that is, ensure that recovery is possible for any of the errors in the set $\mathcal{E} = \{III, ZII, IIZ\}$. Find an orthonormal basis for the two-dimensional code subspace.

7.2 Error-detecting codes

- a) Construct stabilizer generators for an [[n, k, d]] = [[3, 0, 2]] quantum code. With this code, we can detect any single-qubit error. Find the encoded state. (Does it look familiar?)
- b) Two QECC's C_1 and C_2 (with the same length n) are equivalent if a permutation of qubits, combined with single-qubit unitary transformations, transforms the code subspace of C_1 to that of C_2 . Are all [[3, 0, 2]] stabilizer codes equivalent?

c) Does a [[3, 1, 2]] stabilizer code exist?

7.3 Maximal entanglement

Consider the [[5,1,3]] quantum code, whose stabilizer generators are $M_1 = \mathbf{X}\mathbf{Z}\mathbf{Z}\mathbf{X}\mathbf{I}$, and $M_{2,3,4}$ obtained by cyclic permutations of M_1 , and choose the encoded operation $\bar{\mathbf{Z}}$ to be $\bar{\mathbf{Z}} = \mathbf{Z}\mathbf{Z}\mathbf{Z}\mathbf{Z}\mathbf{Z}$. From the encoded states $|\bar{0}\rangle$ with $\bar{\mathbf{Z}}|\bar{0}\rangle = |\bar{0}\rangle$ and $|\bar{1}\rangle$ with $\bar{\mathbf{Z}}|\bar{1}\rangle = -|\bar{1}\rangle$, construct the n = 6, k = 0 code whose encoded state is

$$\frac{1}{\sqrt{2}} \left(|0\rangle \otimes |\bar{0}\rangle + |1\rangle \otimes |\bar{1}\rangle \right) . \tag{7.271}$$

- a) Construct a set of stabilizer generators for this n = 6, k = 0 code.
- b) Find the distance of this code. (Recall that for a k = 0 code, the distance is defined as the minimum weight of any element of the stabilizer.)
- c) Find $\rho^{(3)}$, the density matrix that is obtained if three qubits are selected and the remaining three are traced out.

7.4 Codewords and nonlocality

For the [[5,1,3]] code with stabilizer generators and logical operators as in the preceding problem,

- a) Express \bar{Z} as a weight-3 Pauli operator, a tensor product of I's, X's, and Z's (no Y's). Note that because the code is cyclic, all cyclic permutations of your expression are equivalent ways to represent \bar{Z} .
- b) Use the Einstein locality assumption (local hidden variables) to predict a relation between the five (cyclically related) observables found in (a) and the observable ZZZZZ. Is this relation among observables satisfied for the state $|\bar{0}\rangle$?
- c) What would Einstein say?

7.5 Generalized Shor code

For integer $m \ge 2$, consider the $n = m^2$, k = 1 generalization of Shor's nine-qubit code, with code subspace spanned by the two states:

$$|\bar{0}\rangle = (|000\dots0\rangle + |111\dots1\rangle)^{\otimes m} , |\bar{1}\rangle = (|000\dots0\rangle - |111\dots1\rangle)^{\otimes m} .$$
 (7.272)

a) Construct stabilizer generators for this code, and construct the logical operations \bar{Z} and \bar{X} such that

$$\bar{\boldsymbol{Z}}|\bar{0}\rangle = |\bar{0}\rangle , \qquad \bar{\boldsymbol{X}}|\bar{0}\rangle = |\bar{1}\rangle , \bar{\boldsymbol{Z}}|\bar{1}\rangle = -|\bar{1}\rangle , \qquad \bar{\boldsymbol{X}}|\bar{1}\rangle = |\bar{0}\rangle .$$
 (7.273)

- b) What is the distance of this code?
- c) Suppose that m is odd, and suppose that each of the $n = m^2$ qubits is subjected to the depolarizing channel with error probability p. How well does this code protect the encoded qubit? Specifically, (i) estimate the probability, to leading nontrivial order in p, of a logical bit-flip error $|\bar{0}\rangle \leftrightarrow |\bar{1}\rangle$, and (ii) estimate the probability, to leading nontrivial order in p, of a logical phase error $|\bar{0}\rangle \rightarrow |\bar{0}\rangle$, $|\bar{1}\rangle \rightarrow -|\bar{1}\rangle$.
- d) Consider the asymptotic behavior of your answer to (c) for m large. What condition on p should be satisfied for the code to provide good protection against (i) bit flips and (ii) phase errors, in the $n \to \infty$ limit?

7.6 Encoding circuits

For an [[n,k,d]] quantum code, an encoding transformation is a unitary \boldsymbol{U} that acts as

$$\boldsymbol{U}: |\psi\rangle \otimes |0\rangle^{\otimes (n-k)} \to |\bar{\psi}\rangle , \qquad (7.274)$$

where $|\psi\rangle$ is an arbitrary k-qubit state, and $|\bar{\psi}\rangle$ is the corresponding encoded state. Design a quantum circuit that implements the encoding transformation for

- a) Shor's [[9,1,3]] code.
- b) Steane's [[7,1,3]] code.

7.7 Shortening a quantum code

a) Consider a binary [[n, k, d]] stabilizer code. Show that it is possible to choose the n - k stabilizer generators so that at most two act nontrivially on the last qubit. (That is, the remaining n - k - 2generators apply I to the last qubit.)

- b) These n-k-2 stabilizer generators that apply I to the last qubit will still commute and are still independent if we drop the last qubit. Hence they are the generators for a code with length n-1 and k+1encoded qubits. Show that if the original code is nondegenerate, then the distance of the shortened code is at least d-1. (Hint: First show that if there is a weight-t element of the (n-1)-qubit Pauli group that commutes with the stabilizer of the shortened code, then there is an element of the n-qubit Pauli group of weight at most t+1 that commutes with the stabilizer of the original code.)
- c) Apply the code-shortening procedure of (a) and (b) to the [[5, 1, 3]] QECC. Do you recognize the code that results? (Hint: It may be helpful to exploit the freedom to perform a change of basis on some of the qubits.)

7.8 Codes for qudits

A qudit is a d-dimensional quantum system. The Pauli operators I, X, Y, Z acting on qubits can be generalized to qudits as follows. Let $\{|0\rangle, |1\rangle, \ldots, |d-1\rangle\}$ denote an orthonormal basis for the Hilbert space of a single qudit. Define the operators:

$$\begin{aligned} \boldsymbol{X} : & |j\rangle \to |j+1 \pmod{d}\rangle , \\ \boldsymbol{Z} : & |j\rangle \to \omega^j |j\rangle , \end{aligned}$$
(7.275)

where $\omega = \exp(2\pi i/d)$. Then the $d \times d$ Pauli operators $E_{r,s}$ are

$$E_{r,s} \equiv X^r Z^s$$
, $r, s = 0, 1, \dots, d-1$ (7.276)

- a) Are the $\boldsymbol{E}_{r,s}$'s a basis for the space of operators acting on a qudit? Are they unitary? Evaluate $\operatorname{tr}(\boldsymbol{E}_{r,s}^{\dagger}\boldsymbol{E}_{t,u})$.
- b) The Pauli operators obey

$$\boldsymbol{E}_{r,s}\boldsymbol{E}_{t,u} = (\eta_{r,s;t,u})\boldsymbol{E}_{t,u}\boldsymbol{E}_{r,s} , \qquad (7.277)$$

where $\eta_{r,s;t,u}$ is a phase. Evaluate this phase.

The *n*-fold tensor products of these qudit Pauli operators form a group $G_n^{(d)}$ of order d^{2n+1} (and if we mod out its *d*-element center, we obtain

the group $\bar{G}_n^{(d)}$ of order d^{2n}). To construct a stabilizer code for qudits, we choose an abelian subgroup of $G_n^{(d)}$ with n - k generators; the code is the simultaneous eigenstate with eigenvalue one of these generators. If d is prime, then the code subspace has dimension d^k : k logical qudits are encoded in a block of n qudits.

c) Explain how the dimension might be different if d is not prime. Hint: Consider the case d = 4 and n = 1.)

7.9 Syndrome measurement for qudits

Errors on qudits are diagnosed by measuring the stabilizer generators. For this purpose, we may invoke the two-qudit gate SUM (which generalizes the controlled-NOT), acting as

$$\text{SUM}: |j\rangle \otimes |k\rangle \to |j\rangle \otimes |k+j \pmod{d} \ . \tag{7.278}$$

a) Describe a quantum circuit containing SUM gates that can be executed to measure an *n*-qudit observable of the form

$$\bigotimes_{a} \boldsymbol{Z}_{a}^{s_{a}} . \tag{7.279}$$

If d is prime, then for each r, s = 0, 1, 2, ..., d-1, there is a single-qudit unitary operator $U_{r,s}$ such that

$$\boldsymbol{U}_{r,s}\boldsymbol{E}_{r,s}\boldsymbol{U}_{r,s}^{\dagger} = \boldsymbol{Z} . \qquad (7.280)$$

b) Describe a quantum circuit containing SUM gates and $\boldsymbol{U}_{r,s}$ gates that can be executed to measure an arbitrary element of $G_n^{(d)}$ of the form

$$\bigotimes_{a} \boldsymbol{E}_{r_a, s_a} \ . \tag{7.281}$$

7.10 Error-detecting codes for qudits

A qudit with d = 3 is called a *qutrit*. Consider a qutrit stabilizer code with length n = 3 and k = 1 encoded qutrit defined by the two stabilizer generators

$$ZZZ$$
, XXX . (7.282)

- a) Do the generators commute?
- b) Find the distance of this code.
- c) In terms of the orthonormal basis $\{|0\rangle, |1\rangle, |2\rangle\}$ for the qutrit, write out explicitly an orthonormal basis for the three-dimensional code subspace.
- d) Construct the stabilizer generators for an n = 3m qutrit code (where m is any positive integer), with k = n 2, that can detect one error.
- e) Construct the stabilizer generators for a qudit code that detects one error, with parameters n = d, k = d 2.

7.11 Error-correcting code for qudits

Consider an n = 5, k = 1 qudit stabilizer code with stabilizer generators

(the second, third, and fourth generators are obtained from the first by a cyclic permutation of the qudits).

- a) Find the order of each generator. Are the generators really independent? Do they commute? Is the fifth cyclic permutation $Z Z^{-1} X^{-1} I X$ independent of the rest?
- b) Find the distance of this code. Is the code nondegenerate?
- c) Construct the encoded operations \bar{X} and \bar{Z} , each expressed as an operator of weight 3. (Be sure to check that these operators obey the right commutation relations for any value of d.)

92

Lecture Notes for Physics 219: Quantum Computation

John Preskill California Institute of Technology

14 June 2004

Contents

9	Topological quantum computation	4
9.1	Anyons, anyone?	4
9.2	Flux-charge composites	7
9.3	Spin and statistics	9
9.4	Combining anyons	11
9.5	Unitary representations of the braid group	13
9.6	Topological degeneracy	16
9.7	Toric code revisited	20
9.8	The nonabelian Aharonov-Bohm effect	21
9.9	Braiding of nonabelian fluxons	24
9.10	Superselection sectors of a nonabelian superconductor	29
9.11	Quantum computing with nonabelian fluxons	32
9.12	Anyon models generalized	40
	9.12.1 Labels	40
	9.12.2 Fusion spaces	41
	9.12.3 Braiding: the <i>R</i> -matrix	44
	9.12.4 Associativity of fusion: the <i>F</i> -matrix	45
	9.12.5 Many anyons: the standard basis	46
	9.12.6 Braiding in the standard basis: the <i>B</i> -matrix	47
9.13	Simulating anyons with a quantum circuit	49
9.14	Fibonacci anyons	52
9.15	Quantum dimension	53
9.16	Pentagon and hexagon equations	58
9.17	Simulating a quantum circuit with Fibonacci anyons	61
9.18	Epilogue	63
	9.18.1 Chern-Simons theory	63
	9.18.2 <i>S</i> -matrix	64
	9.18.3 Edge excitations	65
9.19	Bibliographical notes	65

Contents

References

67

3

Topological quantum computation

9

9.1 Anyons, anyone?

A central theme of quantum theory is the concept of *indistinguishable* particles (also called *identical particles*). For example, all electrons in the world are exactly alike. Therefore, for a system with many electrons, an operation that exchanges two of the electrons (swaps their positions) is a symmetry — it leaves the physics unchanged. This symmetry is represented by a unitary transformation acting on the many-electron wave function.

For the indistinguishable particles in three-dimensional space that we normally talk about in physics, particle exchanges are represented in one of two distinct ways. If the particles are bosons (like, for example, ⁴He atoms in a superfluid), then an exchange of two particles is represented by the identity operator: the wave function is invariant, and we say the particles obey Bose statistics. If the particles are fermions (like, for example, electrons in a metal), than an exchange is represented by multiplication by (-1): the wave function changes sign, and we say that the particles obey Fermi statistics.

The concept of identical-particle statistics becomes ambiguous in one spatial dimension. The reason is that for two particles to swap positions in one dimension, the particles need to pass through one another. If the wave function changes sign when two identical particles are exchanged, we could say that the particles are noninteracting fermions, but we could just as well say that the particles are interacting bosons, such that the sign change is induced by the interaction as the particles pass one another. More generally, the exchange could modify the wavefunction by a multiplicative phase $e^{i\theta}$ that could take values other than +1 or -1, but we could account for this phase change by describing the particles as either bosons or fermions.

Thus, identical-particle statistics is a rather tame concept in three (or more) spatial dimensions and also in one dimension. But in between these two dull cases, in two dimensions, a remarkably rich variety of types of particle statistics are possible, so rich that we have far to go before we can give a useful classification of all of the possibilities.

Indistinguishable particles in two dimensions that are neither bosons nor fermions are called *anyons*. Anyons are a fascinating theoretical construct, but do they have anything to do with the physics of real systems that can be studied in the laboratory? The remarkable answer is: "Yes!" Even in our three-dimensional world, a two-dimensional gas of electrons can be realized by trapping the electrons in a thin layer between two slabs of semiconductor, so that at low energies, electron motion in the direction orthogonal to the layer is frozen out. In a sufficiently strong magnetic field and at sufficiently low temperature, and if the electrons in the material are sufficiently mobile, the two-dimensional electron gas attains a profoundly entangled ground state that is separated from all excited states by a nonzero energy gap. Furthermore, the low-energy particle excitations in the systems do not have the quantum numbers of electrons; rather they are anyons, and carry electric charges that are fractions of the electron charge. The anyons have spectacular effects on the transport properties of the sample, manifested as the fractional quantum Hall effect.

Anyons will be our next topic. But why? True, I have already said enough to justify that anyons are a deep and fascinating subject. But this is not a course about the unusual behavior of exotic phases attainable in condensed matter systems. It is a course about quantum computation.

In fact, there is a connection, first appreciated by Alexei Kitaev in 1997: anyons provide an unusual, exciting and perhaps promising means of realizing fault-tolerant quantum computation.

So that sounds like something we should be interested in. After all, I have already given 12 lectures on the theory of quantum error correction and fault-tolerant quantum computing. It is a beautiful theory; I have enjoyed telling you about it and I hope you enjoyed hearing about it. But it is also daunting. We've seen that an ideal quantum circuit can be simulated faithfully by a circuit with noisy gates, provided the noisy gates are not *too* noisy, and we've seen that the overhead in circuit size and depth required for the simulation to succeed is reasonable. These observations greatly boost our confidence that large scale quantum computers will really be built and operated someday. Still, for fault tolerance to be effective, quantum gates need to have quite high fidelity (by the current standards of experimental physics), and the overhead cost of achieving fault tolerance is substantial. Even though reliable quantum computation with noisy gates is possible in principle, there always will

9 Topological quantum computation

be a strong incentive to improve the fidelity of our computation by improving the hardware rather than by compensating for the deficiencies of the hardware through clever circuit design. By using anyons, we might achieve fault tolerance by designing hardware with an intrinsic resistance to decoherence and other errors, significantly reducing the size and depth blowups of our circuit simulations. Clearly, then, we have ample motivation for learning about anyons. Besides, it will be fun!

In some circles, this subject has a reputation (not fully deserved in my view) for being abstruse and inaccessible. I intend to start with the basics, and not to clutter the discussion with details that are mainly irrelevant to our central goals. That way, I hope to keep the presentation clear without really dumbing it down.

What are these goals? I will *not* be explaining how the theory of anyons connects with observed phenomena in fractional quantum Hall systems. In particular, *abelian* anyons arise in most of these applications. From a quantum information viewpoint, abelian anyons are relevant to robust *storage* of quantum information (and we have already gotten a whiff of that connection in our study of toric quantum codes). We will discuss abelian anyons here, but our main interest will be in *nonabelian* anyons, which as we will see can be endowed with surprising computational power.

Kitaev (quant-ph/9707021) pointed out that a system of nonabelian anyons with suitable properties can efficiently simulate a quantum circuit; this idea was elaborated by Ogburn and me (quant-ph/9712048), and generalized by Mochon (quant-ph/0206128, quant-ph/0306063). In Kitaev's original scheme, measurements were required to simulate some quantum gates. Freedman, Larsen and Wang (quant-ph/000110) observed that if we use the right kind of anyons, all measurements can be postponed until the readout of the final result of the computation. Freedman, Kitaev, and Wang (quant-ph/0001071) also showed that a system of anyons can be simulated efficiently by a quantum circuit; thus the anyon quantum computer and the quantum circuit model have equivalent computational power. The aim of these lectures is to explain these important results.

We will focus on the applications of anyons to quantum computing, not on the equally important issue of how systems of anyons with desirable properties can be realized in practice.^{*} It will be left to you to figure that out!

^{*} Two interesting approaches to realizing nonabelian anyons — using superconducting junction arrays and using cold atoms trapped in optical lattices — have been discussed in the recent literature.

9.2 Flux-charge composites

For those of us who are put off by abstract mathematical constructions, it will be helpful to begin our exploration of the theory of anyons by thinking about a concrete model. So let's start by recalling a more familiar concept, the *Aharonov-Bohm effect*.

Imagine electromagnetism in a two-dimensional world, where a "flux tube" is a localized "pointlike" object (in three dimensions, you may envision a plane intersecting a magnetic solenoid directed perpendicular to the plane). The flux might be enclosed behind an impenetrable wall, so that an object outside can never visit the region where the magnetic field is nonzero. But even so, the magnetic field has a measurable influence on charged particles outside the flux tube. If an electric charge q is adiabatically transported (counterclockwise) around a flux Φ , the wave function of the charge acquires a *topological phase* $e^{iq\Phi}$ (where we use units with $\hbar = c = 1$). Here the world "topological" means that the Aharonov-Bohm phase is robust when we deform the trajectory of the charged particle all that matters is the "winding number" of the charge about the flux.

The concept of topological invariance arises naturally in the study of fault tolerance. Topological properties are those that remain invariant when we smoothly deform a system, and a fault-tolerant quantum gate is one whose action on protected information remains invariant (or nearly so) when we deform the implementation of the gate by adding noise. The topological invariance of the Aharonov-Bohm phenomenon is the essential property that we hope to exploit in the design of quantum gates that are intrinsically robust.

We usually regard the Aharonov-Bohm effect as a phenomenon that occurs in quantum electrodynamics, where the photon is exactly massless. But it is useful to recognize that Aharonov-Bohm phenomena can also occur in massive theories. For example, we might consider a "superconducting" system composed of charge e particles, such that composite objects with charge ne form a condensate (where n is an integer). In this superconductor, there is a quantum of flux $\Phi_0 = 2\pi/ne$, the minimal nonzero flux such that a charge-(ne) particle in the condensate, when transported around the flux, acquires a *trivial* Aharonov-Bohm phase. An isolated region that contains a flux quantum is an island of normal material surrounded by the superconducting condensate, prevented from spreading because the magnetic flux cannot penetrate into the superconductor. That is, it is a stable particle, what we could call a "fluxon." When one of the charge-e particles is transported around a fluxon, its wave function acquires the nontrivial topological phase $e^{ie\Phi_0} = e^{2\pi i/n}$. But in the superconductor, the photon acquires a mass via the Higgs mechanism, and there are no massless particles. That topological phases

9 Topological quantum computation

are compatible with massive theories is important, because massless particles are easily excited, a potentially copious source of decoherence.

Now, let's imagine that, in our two-dimensional world, flux and electric charge are permanently bound together (for some reason). A fluxon can be envisioned as flux Φ confined inside an impenetrable circular wall, and an electric charge q is stuck to the *outside* of the wall. What is the angular momentum of this flux-charge composite? Suppose that we carefully rotate the object counterclockwise by angle 2π , returning it to its original orientation. In doing so, we have transported the charge qabout the flux Φ , generating a topological phase $e^{iq\Phi}$. This rotation by 2π is represented in Hilbert space by the unitary transformation

$$\boldsymbol{U}(2\pi) = e^{-i2\pi\boldsymbol{J}} = e^{iq\Phi} , \qquad (9.1)$$

where J is the angular momentum. We conclude, then, that the possible eigenvalues of angular momentum are

$$J = m - \frac{q\Phi}{2\pi} \quad (m = \text{integer}) . \tag{9.2}$$

We can characterize this spectrum by an angular variable $\theta \in [0, 2\pi)$, defined by $\theta = q\Phi \pmod{2\pi}$, and say that the eigenvalues are shifted away from integer values by $-\theta/2\pi$. We will refer to the phase $e^{i\theta}$ that represents a counterclockwise rotation by 2π as the *topological spin* of the composite object.

But shouldn't a rotation by 2π act trivially on a physical system (isn't it the same as doing nothing)? No, we know better than that, from our experience with *spinors* in three dimensions. For a system with fermion number F, we have

$$e^{-2\pi i J} = (-1)^F$$
; (9.3)

if the fermion number is odd, the eigenvalues of J are shifted by 1/2from the integers. This shift is physically acceptable because there is a $(-1)^F$ superselection rule: no observable local operator can change the value of $(-1)^F$ (there is no physical process that can create or destroy an isolated fermion). Acting on a coherent superposition of states with different values of $(-1)^F$, the effect of $e^{-2\pi i J}$ is

$$e^{-i2\pi J} (a | \text{ even } F \rangle + b | \text{ odd } F \rangle) = a | \text{ even } F \rangle - b | \text{ odd } F \rangle .$$
(9.4)

The relative sign in the superposition flips, but this has no detectable physical effects, since all observables are block diagonal in the $(-1)^F$ basis.

Similarly, in two dimensions, the shift in the angular momentum spectrum $e^{-2\pi i J} = e^{i\theta}$ has no unacceptable physical consequences if there is

a θ superselection rule, ensuring that the relative phase in a superposition of states with different values of θ is physically inaccessible (not just in practice but even in principle). As for fermions, there is no allowed physical process that can create of destroy an isolated anyon.

In three dimensions, only $\theta = 0$, π are allowed, because (as you probably know) of a topological property of the three-dimensional rotation group SO(3): a closed path in SO(3) beginning at the identity and ending at a rotation by 4π can be smoothly contracted to a trivial path. It follows that a rotation by 4π really is represented by the identity, and therefore that the eigenvalues of a rotation by 2π are +1 and -1. But the twodimensional rotation group SO(2) does not have this topological property, so that any value of θ is possible in principle.

Note that the angular momentum J changes sign under time reversal (T) and also under parity (P). Unless $\theta = 0$ or π , the spectrum of J is asymmetric about zero, and therefore a theory of anyons typically will not be T or P invariant. In our flux-charge composite model the origin of this symmetry breaking is not mysterious — it arises from the nonzero magnetic field. But in a system with no intrinsic breaking of T and P, if anyons occur then either these symmetries must be broken spontaneously, or else the particle spectrum must be "doubled" so that for each anyon with exchange phase $e^{i\theta}$ there also exists an otherwise identical particle with exchange phase $e^{-i\theta}$.

9.3 Spin and statistics

For identical particles in three dimensions, there is a well known connection between spin and statistics: indistinguishable particles with integer spin are bosons, and those with half-odd-integer spin are fermions. In two dimensions, the spin can be any real number. What does this new possibility of "fractional spin" imply about statistics? The answer is that statistics, too, can be "fractionalized"!

What happens if we perform an exchange of two of our flux-charge composite objects, in a counterclockwise sense? Each charge q is adiabatically transported *half way* around the flux Φ of the other object. We can anticipate, then, that each charge will acquire an Aharonov-Bohm phase that is half of the phase generated by a complete revolution of the charge about the flux. Adding together the phases arising from the transport of both charges, we find that the exchange of the two flux-charge composites changes their wave function by the phase

$$\exp\left[i\left(\frac{1}{2}q\Phi + \frac{1}{2}q\Phi\right)\right] = e^{iq\Phi} = e^{i\theta} = e^{-2\pi i \boldsymbol{J}} .$$
(9.5)

The phase generated when the two objects are exchanged matches the

phase generated when one of the two objects is rotated by 2π . Thus the connection between spin and statistics continues to hold, in a form that is a natural generalization of the connection that applies to bosons and fermions.

The origin of this connection is fairly clear in our flux-charge composite model, but in fact it holds much more generally. Why? Reading textbooks on relativistic quantum field theory, one can easily get the impression that the spin-statistics connection is founded on Lorentz invariance, and has something to do with the properties of the complexified Lorentz group. Actually, this impression is quite misleading. All that is essential for a spin-statistics connection to hold is *the existence of antiparticles*. Special relativity is not an essential ingredient.

Consider an anyon, characterized by the phase θ , and suppose that this particle has a corresponding antiparticle. This means that the particle and its antiparticle, when combined, have trivial quantum numbers (in particular, zero angular momentum) and therefore that there are physical processes in which particle-antiparticle pairs can be created and annihilated. Draw a world line in spacetime that represents a process in which two particle-antiparticle pairs are created (one pair on the left and the other pair on the right), the particle from the pair on the right is exchanged in a counterclockwise sense with the particle from the pair on the left, and then both pairs reannihilate. (The world line has an orientation; if directed forward in time it represents a particle, and if directed backward in time it represents an antiparticle.) Turning our diagram 90° , we obtain a depiction of a process in which a single particle-antiparticle pair is created, the particle and antiparticle are exchanged in a *clock*wise sense, and then the pair reannihilates. Turning it 90° yet again, we have a process in which two pairs are created and the *antiparticle* from the pair on the right is exchanged, in a counterclockwise sense, with the antiparticle from the pair on the left, before reannihilation.



What do we conclude from these manipulations? Denote by R_{ab} the unitary operator that represents a counterclockwise exchange of particles of types a and b (so that the inverse operator R_{ab}^{-1} represents a clockwise exchange), and denote by \bar{a} the antiparticle of a. We have found that

$$R_{aa} = R_{a\bar{a}}^{-1} = R_{\bar{a}\bar{a}} \ . \tag{9.6}$$

If a is an anyon with exchange phase $e^{i\theta}$, then its antiparticle \bar{a} also has the same exchange phase. Furthermore, when a and \bar{a} are exchanged counterclockwise, the phase acquired is $e^{-i\theta}$.

These conclusions are unsurprising when we interpret them from the perspective of our flux-charge composite model of anyons. The antiparticle of the object with flux Φ and charge q has flux $-\Phi$ and charge -q. Hence, when we exchange two antiparticles, the minus signs cancel and the effect is the same as though the particles were exchanged. But if we exchange a particle and an antiparticle, then the relative sign of charge and flux results in the exchange phase $e^{-iq\Phi} = e^{-i\theta}$.

But what is the connection between these observations about statistics and the spin? Continuing to contemplate the same spacetime diagram, let us consider its implications regarding the *orientation* of the particles. For keeping track of the orientation, it is convenient to envision the particle world line not as a thread but as a *ribbon* in spacetime. I claim that our process can be smoothly deformed to one in which a particle-antiparticle pair is created, the particle is rotated counterclockwise by 2π , and then the pair reannihilates. A convenient way to verify this assertion is to take off your belt (or borrow a friend's). The buckle at one end specifies an orientation; point your thumb toward the buckle, and following the righthand rule, twist the belt by 2π before rebuckling it. You should be able to check that you can lay out the belt to match the spacetime diagram for any of the exchange processes described earlier, and also for the process in which the particle rotates by 2π .

Thus, in a topological sense, rotating a particle counterclockwise by 2π is really the same thing as exchanging two particles in a counterclockwise sense (or exchanging particle and antiparticle in a clockwise sense), which provides a satisfying explanation for a general spin-statistics connection.[†] I emphasize again that this argument invokes processes in which particle-antiparticle pairs are created and annihilated, and therefore the existence of antiparticles is an essential prerequisite for it to apply.

9.4 Combining anyons

We know that a composite object composed of two fermions is a boson. What happens when we build a composite object by combining two anyons?

[†] Actually, this discussion has been oversimplified. Though it is adequate for abelian anyons, we will see that it must be amended for nonabelian anyons, because R_{ab} has more than one eigenvalue in the nonabelian case. Similarly, the discussion in the next section of "combining anyons" will need to be elaborated because, in the nonabelian case, more than one kind of composite anyon can be obtained when two anyons are fused together.

9 Topological quantum computation

Suppose that a is an anyon with exchange phase $e^{i\theta}$, and that we build a "molecule" from n of these a anyons. What phase is acquired under a counterclockwise exchange of the two molecules?

The answer is clear in our flux-charge composite model. Each of the n charges in one molecule acquires a phase $e^{i\theta/2}$ when transported half way around each of the n fluxes in the other molecule. Altogether then, $2n^2$ factors of the phase $e^{i\theta/2}$ are generated, resulting in the total phase

$$e^{i\theta_n} = e^{in^2\theta} . (9.7)$$

Said another way, the phase $e^{i\theta}$ occurs altogether n^2 times because in effect n anyons in one molecule are being exchanged with n anyons in the other molecule. Contrary to what we might have naively expected, if we split a fermion (say) into two identical constituents, the constituents have, not an exchange phase of $\sqrt{-1} = i$, but rather $(e^{i\pi})^{1/4} = e^{i\pi/4}$.

This behavior is compatible with the spin-statistics connection: the angular momentum J of the n-anyon molecule satisfies

$$e^{-2\pi i J_n} = e^{-2\pi i n^2 J} = e^{i n^2 \theta} . (9.8)$$

For example, consider a molecule of two anyons, and imagine rotating the molecule counterclockwise by 2π . Not only does each anyon in the molecule rotate by 2π ; in addition one of the anyons revolves around the other. One revolution is equivalent to two successive exchanges, so that the phase generated by the revolution is $e^{i2\theta}$. The total effect of the two rotations and the revolution is the phase

$$\exp\left[i\left(\theta + \theta + 2\theta\right)\right] = e^{i4\theta} . \tag{9.9}$$

Another way to understand why the angular momenta of the anyons in the molecule do not combine additively is to note that the total angular momentum of the molecule consists of two parts — the spin angular momentum S of each of the two anyons (which *is* additive) and the *orbital* angular momentum L of the anyon pair. Because the counterclockwise transport of one anyon around the other generates the nontrivial phase $e^{i2\theta}$, the dependence of the two-anyon wavefunction ψ on the relative azimuthal angle φ is not single-valued; instead,

$$\psi(\varphi + 2\pi) = e^{-i2\theta}\psi(\varphi) . \qquad (9.10)$$

This means that the spectrum of the orbital angular momentum L is shifted away from integer values:

$$e^{-i2\pi \boldsymbol{L}} = e^{2i\theta} , \qquad (9.11)$$
and this orbital angular momentum combines additively with the spin ${\cal S}$ to produce the total angular momentum

$$-2\pi J = -2\pi L - 2\pi S = 2\theta + 2\theta + 2\pi (\text{integer}) = 4\theta + 2\pi (\text{integer}) . \quad (9.12)$$

What if, on the other hand, we build a molecule $\bar{a}a$ from an anyon a and its antiparticle \bar{a} ? Then, as we've seen, the spin S has the same value as for the aa molecule. But the exchange phase has the opposite value, so that the noninteger part of the orbital angular momentum is $-2\pi L = -2\theta$ instead of $-2\pi L = 2\theta$, and the total angular momentum J = L + S is an integer. This property is necessary, of course, if the $\bar{a}a$ pair is to be able to annihilate without leaving behind an object that carries nontrivial angular momentum.

9.5 Unitary representations of the braid group

We have already noted that the angular momentum spectrum has different properties in two spatial dimensions than in three dimensions because SO(2) has different topological properties than SO(3) (SO(3) has a compact simply connected covering group SU(2), but SO(2) does not). This observation provides one way to see why anyons are possible in two dimensions but not in three. It is also instructive to observe that particle *exchanges* have different topological properties in two spatial dimensions than in three dimensions.

As we have found in our discussion of the relation between the statistics of particles and of antiparticles, it is useful to envision exchanges of particles as processes taking place in spacetime. In particular, it is convenient to imagine that we are computing the quantum transition amplitude for a time-dependent process involving n particles by evaluating a sum over particle histories (though for our purposes it will not actually be necessary to calculate any path integrals).

Consider a system of n indistinguishable pointlike particles confined to a two-dimensional spatial surface (which for now we may assume is the plane), and suppose that no two particles are permitted to occupy coincident positions. We may think of a configuration of the particles at a fixed time as a plane with n "punctures" at specified locations — that is, we associate with each particle a hole in the surface with infinitesimal radius. The condition that the particles are forbidden to coincide is enforced by demanding that there are exactly n punctures in the plane at any time. Furthermore, just as the particles are indistinguishable, each puncture is the same as any other. Thus if we were to perform a permutation of the n punctures, this would have no physical effect; all the punctures are the same anyway, so it makes no difference which one is which. All that matters is the n distinct particle positions in the plane.

9 Topological quantum computation

To evaluate the quantum amplitude for a configuration of n particles at specified initial positions at time t = 0 to evolve to a configuration of n particles at specified final positions at time t = T, we are to sum over all classical histories for the n particles that interpolate between the fixed initial configuration and the fixed final configuration, weighted by the phase e^{iS} , where S is the classical action of the history. If we envision each particle world line as a thread, each history for the n particles becomes a *braid*, where each particle on the initial (t = 0) time slice can be connected by a thread to any one of the particles on the final (t = T) time slice. Furthermore, since the particle world lines are forbidden to cross, the braids fall into distinct topological classes that cannot be smoothly deformed one to another, and the path integral can be decomposed as a sum of contributions, with each contribution arising from a different topological class of histories.

Nontrivial exchange operations acting on the particles on the final time slice change the topological class of the braid. Thus we see that the elements of the symmetry group generated by exchanges are in one-to-one correspondence with the topological classes. This (infinite) group is called B_n , the braid group on *n* strands; the group composition law corresponds to concatenation of braids (that is, following one braid with another). In the quantum theory, the quantum state of the *n* indistinguishable particles belongs to a Hilbert space that transforms as a unitary representation of the braid group B_n .

The group can be presented as a set of generators that obey particular defining relations. To understand the defining relations, we may imagine that the *n* particles occupy *n* ordered positions (labeled 1, 2, 3, ..., *n*) arranged on a line. Let σ_1 denote a counterclockwise exchange of the particles that initially occupy positions 1 and 2, let σ_2 denote a counterclockwise exchange of the particles that initially occupy positions 2 and 3, and so on. Any braid can be constructed as a succession of exchanges of neighboring particles; hence $\sigma_1, \sigma_2, \ldots, \sigma_{n-1}$ are the group generators.

The defining relations satisfied by these generators are of two types. The first type is

$$\sigma_j \sigma_k = \sigma_k \sigma_j , \quad |j - k| \ge 2 , \qquad (9.13)$$

which just says that exchanges of disjoint pairs of particles commute. The second, slightly more subtle, type of relation is

$$\sigma_j \sigma_{j+1} \sigma_j = \sigma_{j+1} \sigma_j \sigma_{j+1} , \quad j = 1, 2, \dots, n-2 ,$$
 (9.14)

which is sometimes called the Yang-Baxter relation. You can verify the Yang-Baxter relation by drawing the two braids $\sigma_1\sigma_2\sigma_1$ and $\sigma_2\sigma_1\sigma_2$ on a piece of paper, and observing that both describe a process in which the particles initially in positions 1 and 3 are exchanged counterclockwise

about the particle labeled 2, which stays fixed — *i.e.*, these are topologically equivalent braids.



Since the braid group is infinite, it has an infinite number of unitary irreducible representations, and in fact there are an infinite number of onedimensional representations. Indistinguishable particles that transform as a one-dimensional representation of the braid group are said to be abelian anyons. In the one-dimensional representations, each generator σ_j of B_n is represented by a phase $\sigma_j = e^{i\theta_j}$. Furthermore, the Yang-Baxter relation becomes $e^{i\theta_j}e^{i\theta_{j+1}}e^{i\theta_j} = e^{i\theta_{j+1}}e^{i\theta_j}e^{i\theta_{j+1}}$, which implies $e^{i\theta_j} = e^{i\theta_{j+1}} \equiv e^{i\theta}$ — all exchanges are represented by the same phase. Of course, that makes sense; if the particles are really indistinguishable, the exchange phase ought not to depend on which pair is exchanged. For $\theta = 0$ we obtain bosons, and for $\theta = \pi$, fermions

The braid group also has many nonabelian representations that are of dimension greater than one; indistinguishable particles that transform as such representations are said to be *nonabelian anyons* (or, sometimes, *nonabelions*). To understand the physical properties of nonabelian anyons we will need to understand the mathematical structure of some of these representations. In these lectures, I hope to convey some intuition about nonabelian anyons by discussing some examples in detail.

For now, though, we can already anticipate the main goal we hope to fulfill. For nonabelian anyons, the irreducible representation of B_n realized by n anyons acts on a "topological vector space" V_n whose dimension D_n increases exponentially with n. And for anyons with suitable properties, the image of the representation may be *dense* in $SU(D_n)$. Then braiding of anyons can simulate a quantum computation — any (special) unitary transformation acting on the exponentially large vector space V_n can be realized with arbitrarily good fidelity by executing a suitably chosen braid.

Thus we are keenly interested in the nonabelian representations of the braid group. But we should also emphasize (and will discuss at greater

9 Topological quantum computation

length later on) that there is more to a model of anyons than a mere representation of the braid group. In our flux tube model of abelian anyons, we were able to describe not only the effects of an exchange of anyons, but also the types of particles that can be obtained when two or more anyons are combined together. Likewise, in a general anyon model, the anyons are of various types, and the model incorporates "fusion rules" that specify what types can be obtained when two anyons of particular types are combined. Nontrivial consistency conditions arise because fusion is associate (fusing a with b and then fusing the result with c is equivalent to fusing b with c and then fusing the result with a), and because the fusion rules must be consistent with the braiding rules. Though these consistency conditions are realizable in principle.

9.6 Topological degeneracy

But before moving on to nonabelian anyons, there is another important idea concerning abelian anyons that we should discuss. In any model of anyons (indeed, in any local quantum system with a mass gap), there is a ground state or *vacuum state*, the state in which no particles are present. On the plane the ground state is unique, but for a two-dimensional surface with nontrivial topology, the ground state is degenerate, with the degree of degeneracy depending on the topology. We have already encountered this phenomenon of "topological degeneracy" in the model of abelian anyons that arose in our study of a particular quantum error-correcting code, Kitaev's toric code. Now we will observe that topological degeneracy is a general feature of any model of (abelian) anyons.

We can arrive at the concept of topological degeneracy by examining the representations of a simple operator algebra. Consider the case of the torus, represented as a square with opposite sides identified, and consider the two fundamental 1-cycles of the torus: C_1 , which winds around the square in the x_1 direction, and C_2 which winds around in the x_2 direction. A unitary operator T_1 can be constructed that describes a process in which an anyon-antianyon pair is created, the anyon propagates around C_1 , and then the pair reannihilates. Similarly a unitary operator T_2 can be constructed that describes a process in which the pair is created, and the anyon propagates around the cycle C_2 before the pair reannihilates. Each of the operators T_1 and T_2 preserves the ground state of the system (the state with no particles); indeed, each commutes with the Hamiltonian H of the system and so either can be simultaneously diagonalized with H (T_1 and T_2 are both symmetries).

However, T_1 and T_2 do not commute with one another. If our torus has infinite spatial volume, and there is a mass gap (so that the only

interactions among distantly separated anyons are due to the Aharonov-Bohm effect), then the commutator of T_1 and T_2 is

$$T_2^{-1}T_1^{-1}T_2T_1 = e^{-i2\theta}I , \qquad (9.15)$$

where $e^{i\theta}$ is the anyon's exchange phase. The nontrivial commutator arises because the process in which (1) an anyon winds around C_1 , (2) an anyon winds around C_2 (3) an anyon winds around C_1 in the reverse direction, and (4) an anyon winds around C_2 in the reverse direction, is topologically equivalent to a process in which one anyon winds clockwise around another. To verify this claim, view the action of $T_2^{-1}T_1^{-1}T_2T_1$ as a process in spacetime. First note that the process described by the operator $T_1^{-1}T_1$, in which an anyon world line first sweeps though C_1 and then immediately traverses C_1 in the reverse order, can be deformed to a process in which the anyon world line traverses a topologically trivial loop that can be smoothly shrunk to a point (in keeping with the property that $T_1^{-1}T_1$ is really the identity operator). In similar fashion, the process described by the operator $T_2^{-1}T_1^{-1}T_2T_1$ can be deformed to one where the anyon world lines traverse two closed loops, but such that the world lines *link* once with one another; furthermore, one loop pierces the surface bounded by the other loop in a direction opposite to the orientation inherited by the surface via the right-hand rule from its bounding loop. This process can be smoothly deformed to one in which two pairs are created, one anyon winds clockwise around the other, and then both pairs annihilate. The clockwise winding is equivalent to two successive clockwise exchanges, represented in our one-dimensional representation of the braid group by the phase $e^{-i2\theta}$. We conclude that T_1 and T_2 are noncommuting, except in the cases $\theta = 0$ (bosons) and $\theta = \pi$ (fermions).



9 Topological quantum computation

Since T_1 and T_2 both commute with the Hamiltonian H, both preserve the eigenspaces of H, but since T_1 and T_2 do not commute with one another, they cannot be simultaneously diagonalized. Since T_1 is unitary, its eigenvalues are phases; let us use the angular variable $\alpha \in [0, 2\pi)$ to label an eigenstate of T_1 with eigenvalue $e^{i\alpha}$:

$$T_1|\alpha\rangle = e^{i\alpha}|\alpha\rangle \ . \tag{9.16}$$

Then applying T_2 to the T_1 eigenstate advances the value of α by 2θ :

$$T_1(T_2|\alpha\rangle) = e^{i2\theta}T_2T_1|\alpha\rangle = e^{i2\theta}e^{i\alpha}(T_2|\alpha\rangle) \quad . \tag{9.17}$$

Suppose that θ is a *rational* multiple of 2π , which we may express as

$$\theta = \pi p/q , \qquad (9.18)$$

where q and p (p < 2q) are positive integers with no common factor. Then we conclude that T_1 must have at least q distinct eigenvalues; T_1 acting on α generates an orbit with q distinct values:

$$\alpha + \left(\frac{2\pi p}{q}\right) k \pmod{2\pi}, \quad k = 0, 1, 2, \dots, q-1.$$
(9.19)

Since T_1 commutes with H, on the torus the ground state of our anyonic system (indeed, any energy eigenstate) must have a degeneracy that is an integer multiple of q. Indeed, generically (barring further symmetries or accidental degeneracies), the degeneracy is expected to be exactly q.

For a two-dimensional surface with genus g (a sphere with g "handles"), the degree of this topological degeneracy becomes q^g , because there are operators analogous to T_1 and T_2 associated with each of the g handles, and all of the T_1 -like operators can be simultaneously diagonalized. Furthermore, we can apply a similar argument to a finite planar medium if single anyons can be created and destroyed at the edges of the system. For example, consider an annulus in which anyons can appear or disappear at the inner and outer edges. Then we could define the unitary operator T_1 as describing a process in which an anyon winds counterclockwise around the annulus, and a unitary operator T_2 as describing a process in which an anyon appears at the outer edge, propagates to the inner edge, and disappears. These operators T_1 and T_2 have the same commutator as the corresponding operators defined on the torus, and so we conclude as before that the ground state on the annulus is q-fold degenerate for $\theta = \pi p/q$. For a disc with h holes, there is an operator analogous to T_1 that winds an anyon counterclockwise around each of the holes, and an operator analogous to T_2 that propagates an anyon from the outer boundary of the disk to the edge of the hole; thus the degeneracy is q^h .

What we have described here is a robust *topological quantum memory*. The phase $e^{i2\theta} = e^{i2\pi p/q} \equiv \omega$ acquired when one anyon winds counterclockwise around another is a primitive qth root of unity, and in the case of a planar system with holes, the operator T_1 can be regarded as the encoded Pauli operator \overline{Z} acting on a q-dimension system associated with a particular hole. Physically, the eigenvalue ω^s of \overline{Z} just counts the number s of anyons that are "stuck" inside the hole. The operator T_2 can be regarded as the complementary Pauli operator \bar{X} that increments the value of s by carrying one anyon from the boundary of the system and depositing it in the hole. Since the quantum information is encoded in a nonlocal property of the system, it is well protected from environmental decoherence. By the same token depositing a quantum state in the memory, and reading it out, might be challenging for this system, though in principle Z could be measured by, say, performing an interference experiment in which an anyon projectile scatters off of a hole. We will see later that by using nonabelian anyons we will be able to simplify the readout; in addition, with nonabelian anyons we can use topological properties to process quantum information as well as to store it.

Just how robust is this quantum memory? We need to worry about errors due to thermal fluctuations and due to quantum fluctuations. Thermal fluctuations might excite the creation of anyons, and thermal anyons might diffuse around one of the holes in the sample, or from one boundary to another, causing an encoded error. Thermal errors are heavily suppressed by the Boltzman factor $e^{-\Delta/T}$, if the temperature T is sufficiently small compared to the energy gap Δ (the minimal energy cost of creating a single anyon at the edge of the sample, or a pair of anyons in the bulk). The harmful quantum fluctuations are tunneling processes in which a virtual anyon-antianyon pair appears and the anyon propagates around a hole before reannihilating, or a virtual anyon appears at the edge of a hole and propagates to another boundary before disappearing. These errors due to quantum tunneling are heavily suppressed if the holes are sufficiently large and sufficiently well separated from one another and from the outer boundary.[‡]

Note that our conclusion that the topological degeneracy is finite hinged on the assumption that the angle θ is a rational multiple of π . We may say that a theory of anyons is *rational* if the topological degeneracy is finite for any surface of finite genus (and, for nonabelian anyons, if the

[‡] If you are familiar with Euclidean path integral methods, you'll find it easy to verify that in the leading semiclassical approximation the amplitude A for such a tunneling process in which the anyon propagates a distance L has the form $A = Ce^{-L/L_0}$, where C is a constant and $L_0 = \hbar (2m^* \Delta)^{-1/2}$; here \hbar is Planck's constant and m^* is the effective mass of the anyon, defined so that the kinetic energy of an anyon traveling at speed v is $\frac{1}{2}m^*v^2$.

topological vector space V_n is finite-dimensional for any finite number of anyons n). We may anticipate that the anyons that arise in any physically reasonable system will be rational in this sense, and therefore should be expected to have exchange phases that are roots of unity.

9.7 Toric code revisited

If these observations about topological degeneracy seem hauntingly familiar, it may be because we used quite similar arguments in our discussion of the toric code.

The toric code can be regarded as the (degenerate) ground state of a system of qubits that occupy the links of a square lattice on the torus, with Hamiltonian

$$H = -\frac{1}{4}\Delta\left(\sum_{P} Z_{P} + \sum_{S} X_{S}\right) , \qquad (9.20)$$

where the plaquette operator $Z_P = \bigotimes_{\ell \in P} Z_\ell$ is the tensor product of Z's acting on the four qubits associated with the links contained in plaquette P, and the site operator $X_S \bigotimes_{\ell \ni S} X_\ell$ is the tensor product of X's acting on the four qubits associated with the links that meet at the site S. These plaquette and site operators are just the (commuting) stabilizer generators for the toric code. The ground state is the simultaneous eigenstate with eigenvalue 1 of all the stabilizer generators.

This model has two types of localized particle excitations — plaquette excitations where $Z_P = -1$, which we might think of as magnetic fluxons, and site excitations where $X_S = -1$, which we might think of as electric charges. A Z error acting on a link creates a pair of charges on the two site joined by the link, and an X error acting on a link creates a pair of fluxons on the two plaquettes that share the link. The energy gap Δ is the cost of creating a pair of either type.

The charges are bosons relative to one another (they have a trivial exchange phase $e^{i\theta} = 0$), and the fluxons are also bosons relative to one another. Since the fluxons are distinguishable from the charges, it does not make sense to exchange a charge with a flux. But what makes this an anyon model is that a phase (-1) is acquired when a charge is carried around a flux. The degeneracy of the ground state (the dimension of the code space) can be understood as a consequence of this property of the particles.

For this model on the torus, because there are two types of particles, there are two types of T_1 operators: $T_{1,S}$, which propagates a charge (site defect) around the 1-cycle C_1 , and $T_{1,P}$, which propagates a fluxon (plaquette defect) around C_1 . Similarly there are two types of T_2 operators, $T_{2,S}$ and $T_{2,P}$. The nontrivial commutators are

$$T_{2,P}^{-1}T_{1,S}^{-1}T_{2,P}T_{1,S} = -1 = T_{2,S}^{-1}T_{1,P}^{-1}T_{2,S}T_{1,P} , \qquad (9.21)$$

both arising from processes in which world lines of charges and fluxon link once with one another. Thus $T_{1,S}$ and $T_{2,S}$ can be diagonalized simultaneously, and can be regarded as the encoded Pauli operators \bar{Z}_1 and \bar{Z}_2 acting on two protected qubits. The operator $T_{2,P}$, which commutes with \bar{Z}_1 and anticommutes with \bar{Z}_2 , can be regarded as the encoded \bar{X}_1 , and similarly $T_{1,P}$ is the encoded \bar{X}_2 .

On the torus, the degeneracy of the four ground states is exact for the ideal Hamiltonian we constructed (the particles have infinite effective masses). Weak local perturbations will break the degeneracy, but only by an amount that gets exponentially small as the linear size L of the torus increases. To be concrete, suppose the perturbation is a uniform "magnetic field" pointing in the \hat{z} direction, coupling to the magnetic moments of the qubits:

$$H' = -h \sum_{\ell} Z_{\ell} . \qquad (9.22)$$

Because of the nonzero energy gap, for the purpose of computing in perturbation theory the leading contribution to the splitting of the degeneracy, it suffices to consider the effect of the perturbation in the fourdimensional subspace spanned by the ground states of the unperturbed system. In the toric code, the operators with nontrivial matrix elements in this subspace are those such that Z_{ℓ} 's act on links that form a closed loop that wraps around the torus (or X_{ℓ} 's act on links whose dual links form a closed loop that wraps around the torus). For an $L \times L$ lattice on the torus, the minimal length of such a closed loop is L; therefore nonvanishing matrix elements do not arise in perturbation theory until the Lth order, and are suppressed by h^L . Thus, for small h and large L, memory errors due to quantum fluctuations occur only with exponentially small amplitude.

9.8 The nonabelian Aharonov-Bohm effect

There is a beautiful abstract theory of nonabelian anyons, and in due course we will delve into that theory a bit. But I would prefer to launch our study of the subject by describing a more concrete model.

With that goal in mind, let us recall some properties of *chromodynamics*, the theory of the quarks and gluons contained within atomic nuclei and other strongly interacting particles. In the real world, quarks are permanently bound together and can never be isolated, but for our discussion let us imagine a fictitious world in which the forces between quarks are weak, so that the characteristic distance scale of quark confinement is very large.

Quarks carry a degree of freedom known metaphorically as *color*. That is, there are three kinds of quarks, which in keeping with the metaphor we call red (R), yellow (Y), and blue (B). Quarks of all three colors are physical identical, except that when we bring two quarks together, we can tell whether their colors are the same (the Coulombic interaction between like colors is repulsive), or different (distinct colors attract). There is nothing to prevent me from establishing a *quark bureau of standards* in my laboratory, where colored quarks are sorted into three bins; all the quarks in the same bin have the same color, and quarks in different bins have different colors. We may attach (arbitrary) labels to the three bins -R, Y, and B.

If while taking a hike outside by lab, I discover a previously unseen quark, I may at first be unsure of its color. But I can find out. I capture the quark and carry it back to my lab, being very careful not to disturb its color along the way (in chromodynamics, there is a notion of *parallel transport* of color). Once back at the quark bureau of standards, I can compare this new quark to the previously calibrated quarks in the bins, and so determine whether the new quark should be labeled R, Y, or B.

It sounds simple but there is a catch: in chromodynamics, the parallel transport of color is *path dependent* due to an Aharonov-Bohm phenomenon that affects color. Suppose that at the quark bureau of standards a quark is prepared whose color is described by the quantum state

$$|\psi_q\rangle = q_R|R\rangle + q_Y|Y\rangle + q_B|B\rangle ; \qquad (9.23)$$

it is a coherent superposition with amplitudes q_R, q_Y, q_B for the red, yellow, and blue states. The quark is carried along a path that winds around a *color magnetic flux tube* and is returned to the quark bureau of standards where its color can be recalibrated. Upon its return the color state has been rotated:

$$\begin{pmatrix} q'_R \\ q'_Y \\ q'_B \end{pmatrix} = U \begin{pmatrix} q_R \\ q_Y \\ q_B \end{pmatrix} , \qquad (9.24)$$

where U is a (special) unitary 3×3 matrix. Similarly, when a newly discovered quark is carried back to the bureau of standards, the outcome of a measurement of its color will depend on whether it passed to the left or the right of the flux tube during its voyage.

This path dependence of the parallel transport of color is closely analogous to the path dependence of the parallel transport of a tangent vector on a curved Riemannian manifold. In chromodynamics, a magnetic field is the "curvature" whose strength determines the amount of path dependence. In general, the SU(3) matrix U that describes the effect of parallel transport of color about a closed path depends on the *basepoint* x_0 where the path begins and ends, as well as on the closed loop C traversed by the path — when it is important to specify the loop and basepoint we will use the notation $U(C, x_0)$. The eigenvalues of the matrix U have an invariant "geometrical" meaning characterizing the parallel transport, but U itself depends on the conventions we have established at the basepoint. You might prefer to choose a different orthonormal basis for the color space at the basepoint x_0 than the basis I chose, so that your standard colors R, Y, and B differ from mine by the action of an SU(3) matrix $V(x_0)$. Then, while I characterize the effect or parallel transport around the loop

$$V(x_0)U(C, x_0)V(x_0)^{-1}$$
, (9.25)

that differs from mine by conjugation by $V(x_0)$. Physicists sometimes speak of this freedom to redefine conventions as a choice of *gauge*, and say that U itself is *gauge dependent* while its eigenvalues are *gauge invariant*.

C with the matrix U, you characterize it with another matrix

Chromodynamics, on the distance scales we consider here (much smaller than the characteristic distance scale of quark confinement), is a theory like electrodynamics with long-range Coulombic interactions among quarks, mediated by "gluon" fields. We will prefer to consider a theory that retains some of the features of chromodynamics (in particular the path dependence of color transport), but without the easily excited light gluons. In the case of electrodynamics, we eliminated the light photon by considering a "superconductor" in which charged particles form a condensate, magnetic fields are expelled, and the magnetic flux of an isolated object is quantized. Let us appeal to the same idea here. We consider a nonabelian superconductor in two spatial dimensions. This world contains particles that carry "magnetic flux" (similar to the color magnetic flux in chromodynamics) and particles that carry charge (similar to the colored quarks of chromodynamics). The flux takes values in a *nonabelian finite* group G, and the charges are unitary irreducible representations of the group G. In this setting, we can formulate some interesting models of nonabelian anyons.

Let R denote a particular irreducible representation of G, whose dimension is denoted |R|. We may establish a "charge bureau of standards," and define there an arbitrarily chosen orthonormal basis for the |R|-dimensional vector space acted upon by R:

$$|R,i\rangle$$
, $i = 1, 2, \dots |R|$. (9.26)

When a charge R is transported around a closed path that encloses a flux $a \in G$, there is a nontrivial Aharonov-Bohm effect — the basis for R is

rotated by a unitary matrix $D^{R}(a)$ that represents a:

$$|R,j\rangle \mapsto \sum_{i=1}^{|R|} |R,i\rangle D_{ij}^R(a) . \qquad (9.27)$$

The matrix elements $D_{ij}^R(a)$ are measurable in principle, for example by conducting interference experiments in which a beam of calibrated charges can pass on either side of the flux. (The phase of the complex number $D_{ij}^R(a)$ determines the magnitude of the *shift* of the interference fringes, and the modulus of $D_{ij}^R(a)$ determines the *visibility* of the fringes.) Thus once we have chosen a standard basis for the charges, we can use the charges to attach labels (elements of G) to all fluxes. The flux labels are unambiguous as long as the representation R is faithful, and barring any group automorphisms (which create ambiguities that we are free to resolve however we please).

However, the group elements that we attach to the fluxes depend on our conventions. Suppose I am presented with k fluxons (particles that carry flux), and that I use my standard charges to measure the flux of each particle. I assign group elements $a_1, a_2, \ldots, a_k \in G$ to the k fluxons. You are then asked to measure the flux, to verify my assignments. But your standard charges differ from mine, because they have been surreptitiously transported around another flux (one that I would label with $g \in G$). Therefore you will assign the group elements $ga_1g^{-1}, ga_2g^{-1}, \ldots, ga_kg^{-1}$ to the k fluxons; our assignments differ by an overall conjugation by g.

The moral of this story is that the assignment of group elements to fluxons is inherently ambiguous and has no invariant meaning. But because the valid assignments of group elements to fluxons differ only by conjugation by some element $g \in G$, the *conjugacy class* of the flux in G does have an invariant meaning on which all observers will agree. Indeed, even if we fix our conventions at the charge bureau of standards, the group element that we assign to a particular fluxon may change if that fluxon takes part in a physical process in which *it* braids with other fluxons. For that reason, the fluxons belonging to the same conjugacy class should all be regarded as indistinguishable particles, even though they come in many varieties (one for each representative of the class) that can be distinguished when we make measurements at a particular time and place: The fluxons are *nonabelian anyons*.

9.9 Braiding of nonabelian fluxons

We will see that, for a nonabelian superconductor with suitable properties, it is possible to operate a fault-tolerant universal quantum computer by manipulating the fluxons. The key thing to understand is what happens when two fluxons are exchanged with one another.

For this purpose, imagine that we carefully calibrate two fluxons, and label them with elements of the group G. The labels are assigned by establishing a standard basis for the charged particles at a basepoint x_0 . Then a standard path, designated α , is chosen that begins at x_0 , winds counterclockwise around the fluxon on the *left*, and returns to x_0 . Finally, charged particles are carried around the closed path α , and it is observed that under this parallel transport, the particles are acted upon by D(a), where D is the representation of G according to which the charged particles transform, and $a \in G$ is the particular group element that we assign to the fluxon. Similarly, another standard path, designated β , is chosen that begins at x_0 , winds counterclockwise around the fluxon on the *right*, and returns to x_0 ; the effect of parallel transport around β is found to be D(b), and so the fluxon on the right is labeled with $b \in G$.

Now imagine that a counterclockwise exchange of the two fluxons is performed, after which the calibration procedure is repeated. How will the fluxons be labeled now?

To find the answer, consider the path $\alpha\beta\alpha^{-1}$; here we use α^{-1} to denote the path α traversed in reverse order, and we have adopted the convention that $\alpha\beta\alpha^{-1}$ denotes the path in which α^{-1} is traversed first, followed by β and then α . Now observe that if, as the two fluxons are exchanged counterclockwise, we deform the paths so that they are never crossed by the fluxons, then the path $\alpha\beta\alpha^{-1}$ is deformed to the path α , while the path α is deformed to β :

$$\alpha\beta\alpha^{-1}\mapsto\alpha\,,\quad\alpha\mapsto\beta\,.$$
 (9.28)



It follows that the effect of transporting a charge around the path α , after the exchange, is equivalent to the effect of transport around the path $\alpha\beta\alpha^{-1}$, before the exchange; similarly, the effect of transport around β , after the exchange, is the same as the effect of transport around α before. We conclude that the braid operator R representing a counterclockwise

exchange acts on the fluxes according to

$$R: |a,b\rangle \to |aba^{-1},a\rangle . \tag{9.29}$$

Of course, if the fluxes a and b are commuting elements of G, all the braiding does is swap the positions of the two labels. But if a and b do not commute, the effect of the exchange is more subtle and interesting. The asymmetric form of the action of R is a consequence of our conventions and of the (counterclockwise) sense of the exchange; the inverse operator R^{-1} representing a clockwise exchange acts as

$$R^{-1}: |a,b\rangle \to |b,b^{-1}ab\rangle . \tag{9.30}$$

Note that the total flux of the pair of fluxons can be detected by a charged particle that traverses the path $\alpha\beta$ that encloses both members of the pair. Since in principle the charge detecting this total flux could be far, far away, the exchange ought not to alter the total flux; indeed, we find that the product flux *ab* is preserved by *R* and by R^{-1} .

The effect of two successive counterclockwise exchanges is the "monodromy" operator R^2 , representing the counterclockwise winding of one fluxon about the other, whose action is

$$R^{2}: |a,b\rangle \mapsto |(ab)a(ab)^{-1}, (ab)b(ab)^{-1}\rangle ; \qquad (9.31)$$

both fluxes are conjugated by the total flux ab. That is, winding a counterclockwise about b conjugates b by a (and similarly, winding b clockwise about a conjugates a by b^{-1}). The nontrivial monodromy means that if many fluxons are distributed in the plane, and one of these fluxons is to be brought to my laboratory for analysis, the group element I assign to the fluxon may depend on the path the flux follows as it travels to my lab. If for one choice of path the flux is labeled by $a \in G$, then for other paths any other element bab^{-1} might in principle be assigned. Thus, the conjugacy class in G represented by the fluxon is invariant, but the particular representative of that class is ambiguous.

For example, suppose the group is $G = S_3$, the permutation group on three objects. One of the conjugacy classes contains all of the twocycle permutations (transpositions of two objects), the three elements $\{(12), (23), (31)\}$. When two such two-cycles fluxes are combined, there are three possibilities for the total flux — the trivial flux e, or one of the three-cycle fluxes (123) or (132). If the total flux is trivial, the braiding of the two fluxes is also trivial (a and $b = a^{-1}$ commute). But if the total flux is nontrivial, then the braid operator R has orbits of length three:

$$R: |(12), (23)\rangle \mapsto |(31), (12)\rangle \mapsto |(23), (31)\rangle \mapsto |(12), (23)\rangle ,$$

$$R: |(23), (12)\rangle \mapsto |(31), (23)\rangle \mapsto |(12), (31)\rangle \mapsto |(23), (12)\rangle ,$$

(9.32)

Thus, if the two fluxons are exchanged three times, they swap positions (the number of exchanges is odd), yet the labeling of the state is unmodified. This observation means that there can be quantum interference between the "direct" and "exchange" scattering of two fluxons that carry distinct labels in the same conjugacy class, reinforcing the notion that fluxes carrying conjugate labels ought to be regarded as indistinguishable particles.

Since the braid operator acting on pairs of two-cycle fluxes satisfies $R^3 = I$, its eigenvalues are third roots of unity. For example, by taking linear combinations of the three states with total flux (123), we obtain the R eigenstates

$$R = 1: |(12), (23)\rangle + |(31), (12)\rangle + |(23), (31)\rangle,$$

$$R = \omega: |(12), (23)\rangle + \bar{\omega}|(31), (12)\rangle + \omega|(23), (31)\rangle,$$

$$R = \bar{\omega}: |(12), (23)\rangle + \omega|(31), (12)\rangle + \bar{\omega}|(23), (31)\rangle, \quad (9.33)$$

where $\omega = e^{2\pi i/3}$.

Although a pair of fluxes $|a, a^{-1}\rangle$ with trivial total flux has trivial braiding properties, it is interesting for another reason — it carries *charge*. The way to detect the charge of an object is to carry a flux *b* around the object (counterclockwise); this modifies the object by the action of $D^R(b)$ for some representation *R* of *G*. If the charge is zero then the representation is trivial — D(b) = I for all $b \in G$. But if we carry flux *b* counterclockwise around the state $|a, a^{-1}\rangle$, the state transforms as

$$|a, a^{-1}\rangle \mapsto |bab^{-1}, ba^{-1}b^{-1}\rangle$$
, (9.34)

a nontrivial action (for at least some b) if a belongs to a conjugacy class with more than one element. In fact, for each conjugacy class α , there is a unique state $|0; \alpha\rangle$ with zero charge, the uniform superposition of the class representatives:

$$|0;\alpha\rangle = \frac{1}{\sqrt{|\alpha|}} \sum_{a \in \alpha} |a, a^{-1}\rangle , \qquad (9.35)$$

where $|\alpha|$ denotes the order of α . A pair of fluxons in the class α that can be created in a local process must not carry any conserved charges and therefore must be in the state $|0; \alpha\rangle$. Other linear combinations orthogonal to $|0; \alpha\rangle$ carry nonzero charge. This charge carried by a pair of fluxons can be detected by other fluxons, yet oddly the charge cannot be localized on the core of either particle in the pair. Rather it is a collective property of the pair. If two fluxons with a nonzero total charge are brought together, complete annihilation of the pair will be forbidden by charge conservation, even though the total flux is zero.

9 Topological quantum computation

In the case of a pair of fluxons from the two-cycle class of $G = S_3$, for example, there is a two-dimensional subspace with trivial total flux and nontrivial charge, for which we may choose the basis

$$\begin{aligned} |0\rangle &= |(12), (12)\rangle + \bar{\omega}|(23), (23)\rangle + \omega|(31), (31)\rangle ,\\ |1\rangle &= |(12), (12)\rangle + \omega|(23), (23)\rangle + \bar{\omega}|(31), (31)\rangle . \end{aligned}$$
(9.36)

If a flux b is carried around the pair, both fluxes are conjugated by b; therefore the action (by conjugation) of S_3 on these states is

$$D(12) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad D(23) = \begin{pmatrix} 0 & \bar{\omega} \\ \omega & 0 \end{pmatrix}, \quad D(31) = \begin{pmatrix} 0 & \omega \\ \bar{\omega} & 0 \end{pmatrix},$$
$$D(123) = \begin{pmatrix} \omega & 0 \\ 0 & \bar{\omega} \end{pmatrix}, \quad D(132) = \begin{pmatrix} \bar{\omega} & 0 \\ 0 & \omega \end{pmatrix}.$$
(9.37)

This action is just the two-dimensional irreducible representation R = [2] of S_3 , and so we conclude that the charge of the pair of fluxons is [2].

Furthermore, under braiding this charge carried by a pair of fluxons can be transferred to other particles. For example, consider a pair of particles, each of which carries charge but no flux (I will refer to such particles as *chargeons*), such that the total charge of the pair is trivial. If one of the chargeons transforms as the unitary irreducible representation R of G, there is a unique conjugate representation \bar{R} that can be combined with R to give the trivial representation; if $\{|R, i\rangle\}$ is a basis for R, then a basis $\{|\bar{R}, i\rangle\}$ can be chosen for \bar{R} , such that the chargeon pair with trivial charge can be expressed as

$$|0;R\rangle = \frac{1}{\sqrt{|R|}} \sum_{i} |R,i\rangle \otimes |\bar{R},i\rangle . \qquad (9.38)$$

Imagine that we create a pair of fluxons in the state $|0; \alpha\rangle$ and also create a pair of chargeons in the state $|0; R\rangle$. Then we wind the chargeon with charge R counterclockwise around the fluxon with flux in class α , and bring the two chargeons together again to see if they will annihilate. What happens?

For a fixed value $a \in \alpha$ of the flux, the effect of the winding on the state of the two chargeons is

$$|0;R\rangle \mapsto \frac{1}{\sqrt{|R|}} \sum_{i,j} |R,j\rangle \otimes |\bar{R},i\rangle D_{ji}^{R}(a) ; \qquad (9.39)$$

if the charge of the pair were now measured, the probability that zero total charge would be found is the square of the overlap of this state with $|0; R\rangle$, which is

$$\operatorname{Prob}(0) = \left|\frac{\chi^R(a)}{|R|}\right|^2 , \qquad (9.40)$$

where

$$\chi^{R}(a) = \sum_{i} D_{ii}^{R}(a) = \text{tr } D^{R}(a)$$
(9.41)

is the *character* of the representation R, evaluated at a. In fact, the character (a trace) is unchanged by conjugation — it takes the same value for all $a \in \alpha$. Therefore, eq. (9.40) is also the probability that the pair of chargeons has zero total charge when one chargeon (initially a member of a pair in the state $|0; R\rangle$ winds around one fluxon (initially a member of a pair in the state $|0;\alpha\rangle$). Of course, since the total charge of all four particles is zero and charge is conserved, after the winding the two pairs have opposite charges — if the pair of chargeons has total charge R', then the pair of fluxons must have total charge \bar{R}' , combined with R' to give trivial total charge. A pair of particles with zero total charge and flux can annihilate, leaving no stable particle behind, while a pair with nonzero charge will be unable to annihilate completely. We conclude, then, that if the world lines of a fluxon pair and a chargeon pair link once, the probability that both pairs will be able to annihilate is given by eq. (9.40). This probability is less than one, provided that the representation of Ris not one dimensional and the class α is not represented trivially. Thus the linking of the world lines induces an exchange of charge between the two pairs.

For example, in the case where α is the two-cycle class of $G = S_3$ and R = [2] (the two-dimensional irreducible representation of S_3), we see from eq. (9.37) that $\chi^{[2]}(\alpha) = 0$. Therefore, charge is transferred with certainty; after the winding, both the fluxon pair and the chargeon pair transform as R' = [2].

9.10 Superselection sectors of a nonabelian superconductor

In our discussion so far of the nonabelian superconductor, we have been considering two kinds of particles: *fluxons*, which carry flux but no charge, and *chargeons*, which carry charge but no flux. These are not the most general possible particles. It will be instructive to consider what happens when we build a composite particle by combining a fluxon with a chargeon. In particular, what is the charge of the composite? This question is surprisingly subtle; to answer cogently, we should think carefully about how the charge can be measured.

In principle, charge can be measured in an Aharonov-Bohm interference experiment. We could hide the object whose charge is to be found behind a screen in between two slits, shoot a beam of carefully calibrated fluxons at the screen, and detect the fluxons on the other side. From the shift and visibility of the interference pattern revealed by the detected positions of the fluxons, we can determine $D^R(b)$ for each $b \in G$, and so deduce R.

9 Topological quantum computation

However, there is a catch if the object being analyzed carries a nontrivial flux $a \in G$ as well as charge. Since carrying a flux b around the flux a changes a to bab^{-1} , the two possible paths followed by the b flux do *not* interfere, if a and b do not commute. After the b flux is detected, we could check whether the a flux has been modified, and determine whether the b flux passed through the slit on the left or the slit on the right. Since the flux $(a \text{ or } bab^{-1})$ is correlated with the "which way" information (left or right slit), the interference is destroyed.

Therefore, the experiment reveals information about the charge only if a and b commute. Hence the charge attached to a flux a is not described as an irreducible representation of G; instead it is an irreducible representation of a subgroup of G, the normalizer N(a) of a in G, which is defined as

$$N(a) = \{ b \in G | ab = ba \} .$$
(9.42)

The normalizers N(a) and $N(bab^{-1})$ are isomorphic, so we may associate the normalizer with a conjugacy class α of G rather than with a particular element, and denote it as $N(\alpha)$. Therefore, each type of particle that can occur in our nonabelian superconductor really has two labels: a conjugacy class α describing the flux, and an irreducible representation $R^{(\alpha)}$ of $N(\alpha)$ describing the charge. We say that α and $R^{(\alpha)}$ label the superselection sectors of the theory, as these are the properties of a localized object that must be conserved in all local physical processes. For particles that carry the labels $(\alpha, R^{(\alpha)})$, it is possible to establish a "bureau of standards" where altogether $|\alpha| \cdot |R^{(\alpha)}| \equiv d_{(\alpha,R^{(\alpha)})}$ different particle species can be distinguished at a particular time and place this number is called the *dimension* of the sector. But if these particles are braided with other particles the species may change, while the labels $(\alpha, R^{(\alpha)})$ remain invariant.

In any theory of anyons, a dimension can be assigned to each particle type, although as we will see, in general the dimension need not be an integer, and may have no direct interpretation in terms the counting of distinct species of the same type. The *total dimension* \mathcal{D} can be defined by summing over all types; in the case of a nonabelian superconductor we have

$$\mathcal{D}^2 = \sum_{\alpha} \sum_{R^{(\alpha)}} d^2_{(\alpha, R^{(\alpha)})} = \sum_{\alpha} |\alpha|^2 \sum_{R^{(\alpha)}} |R^{(\alpha)}|^2 .$$
(9.43)

Since the sum over the dimension squared for all irreducible representations of a finite group is the order of the group, and the order of the normalizer $N(\alpha)$ is $|G|/|\alpha|$, we obtain

$$\mathcal{D}^2 = \sum_{\alpha} |\alpha| \cdot |G| = |G|^2 ; \qquad (9.44)$$

the total dimension is $\mathcal{D} = |G|$.

For the case $G = S_3$ there are 8 particle types, listed here:

Type	Flux	Charge	Dim
Α	e	[+]	1
В	e	[-]	1
С	e	[2]	2
D	(12)	[+]	3
E	(12)	[-]	3
\mathbf{F}	(123)	[1]	2
G	(123)	$[\omega]$	2
Н	(123)	$[\bar{\omega}]$	2

If the flux is trivial (e), then the charge can be any one of the three irreducible representations of S_3 — the trivial one-dimensional representation [+], the nontrivial one-dimensional representation [-], or the twodimensional representation [2]. If the flux is a two-cycle, then the normalizer group is Z_2 , and the charge can be either the trivial representation [+] or the nontrivial representation [-]. And if the flux is a three-cycle, then the normalizer group is Z_3 , and the charge can be either the trivial representation [1], the nontrivial representation [ω], or its conjugate representation [$\bar{\omega}$]. You can verify that the total dimension is $\mathcal{D} = |S_3| = 6$, as expected.

Note that since a commutes with all elements of N(a) by definition, the matrix $D^{R^{(a)}}(a)$ that represents a in the irreducible representation $R^{(a)}$ commutes with all matrices in the representation; therefore by Schur's lemma it is a multiple of the identity:

$$D^{R^{(a)}}(a) = \exp\left(i\theta_{R^{(a)}}\right)I .$$
(9.45)

To appreciate the significance of the phase $\exp(i\theta_{R^{(a)}})$, consider a fluxcharge composite in which a chargeon in representation $R^{(a)}$ is bound to the flux a, and imagine rotating the composite object counterclockwise by 2π . This rotation carries the charge around the flux, generating the phase

$$e^{-2\pi i J} = e^{i\theta_{R(a)}}$$
; (9.46)

therefore each superselection sector has a definite value of the *topological* spin, determined by $\theta_{R^{(a)}}$.

When two different particle types are fused together, the composite object can be of various types, and the *fusion rules* of the theory specify which types are possible. The flux of the composite can belong to any of the conjugacy classes that can be obtained as a product of representatives of the classes that label the two constituents. Finding the charge of the composite is especially tricky, as we must decompose a tensor product of representations of two different normalizer groups as a sum of representations of the normalizer of the product flux. In the case $G = S_3$, the rule governing the fusion of two particles of type D, for example, is

$$D \times D = A + C + F + G + H \tag{9.47}$$

We have already noted that the fusion of two two-cycle fluxes can yield either a trivial total flux or a three-cycle flux, and that the charge of the composite with trivial total flux can be either [+] or [2]. If the total flux is a three-cycle, then the charge eigenstates are just the braid operator eigenstates that we constructed in eq. (9.33).

For a system of two anyons, why should the eigenstates of the total charge also be eigenstates of the braid operator? We can understand this connection more generally by thinking about the angular momentum of the two-anyon composite object. The monodromy operator R^2 captures the effect of winding one particle counterclockwise around another. This winding is almost the same thing as rotating the composite system counterclockwise by 2π , except that the rotation of the composite system also rotates both of the constituents. We can compensate for the rotation of the constituents by following the counterclockwise rotation of the composite by a clockwise rotation of the constituents. Therefore, the monodromy operator can be expressed as

$$(R_{ab}^c)^2 = e^{-2\pi i J_c} e^{2\pi i J_a} e^{2\pi i J_b} = e^{i(\theta_c - \theta_a - \theta_b)} .$$
(9.48)

Here R_{ab}^c denotes the braid operator for a counterclockwise exchange of particles of types a and b that are combined together into a composite of type c, and we are using a more succinct notation than before, in which a, b, c are complete labels for the superselection sectors (specifying, in the nonabelian superconductor model, both the flux and the charge). Since each superselection sector has a definite topological spin, and the monodromy operator is diagonal in the topological spin basis, we see that eigenstates of the braid operator coincide with charge eigenstates. Note that eq. (9.48) generalizes our earlier observations about abelian anyons — that a composite of two identical anyons has topological spin $e^{i4\theta}$, and that the exchange phase of an anyon-antianyon pair (with trivial total spin) is $e^{-i\theta}$.

9.11 Quantum computing with nonabelian fluxons

A model of anyons is characterized by the answers to two basic questions: (1) What happens when two anyons are combined together (what are the *fusion rules*)? (2) What happens when two anyons are exchanged (what are the *braiding rules*)? We have discussed how these questions are answered in the special case of a nonabelian superconductor model associated with a nonabelian finite group G, and now we wish to see how these fusion and braiding rules can be invoked in a simulation of a quantum circuit.

In formulating the simulation, we will assume these physical capabilities:

- Pair creation and identification. We can create pairs of particles, and for each pair we can identify the particle type (the conjugacy class α of the flux of each particle in the pair, and the particles's charge — an irreducible representation $R^{(\alpha)}$ of the flux's normalizer group $N(\alpha)$). This assumption is reasonable because there is no symmetry relating particles of different types; they have distinguishable physical properties — for example, different energy gaps and effective masses. In practice, the only particle types that will be needed are fluxons that carry no charge and chargeons that carry no flux.
- Pair annihilation. We can bring two particles together, and observe whether the pair annihilates completely. Thus we obtain the answer to the question: Does this pair of particles have trivial flux and charge, or not? This assumption is reasonable, because if the pair carries a nontrivial value of some conserved quantity, a localized excitation must be left behind when the pair fuses, and this leftover particle is detectable in principle.
- *Braiding.* We can guide the particles along specified trajectories, and so perform exchanges of the particles. Quantum gates will be simulated by choosing particles world lines that realize particular braids.

These primitive capabilities allow us to realize some further derived capabilities that will be used repeatedly. First, we can use the chargeons to calibrate the fluxons and assemble a flux bureau of standards. Suppose that we are presented with two pairs of fluxons in the states $|a, a^{-1}\rangle$ and $|b, b^{-1}\rangle$, and we wish to determine whether the fluxes a and b match or not. We create a chargeon-antichargeon pair, where the charge of the chargeon is the irreducible representation R of G. Then we carry the chargeon around a closed path that encloses the first member of the first fluxon pair and the second member of the second fluxon pair, we reunite the chargeon and antichargeon, and observed whether the chargeon pair annihilates or not. Since the total flux enclosed by the chargeon's path is ab^{-1} , the chargeon pair annihilates with probability

$$Prob(0) = \left|\frac{\chi^R(ab^{-1})}{|R|}\right|^2 , \qquad (9.49)$$

9 Topological quantum computation

which is less than one if the flux ab^{-1} is not the identity (assuming that the representation R is not one-dimensional and represents ab^{-1} nontrivially). Thus, if annihilation of the chargeon pair does not occur, we know for sure that a and b are distinct fluxes, and each time annihilation does occur, it becomes increasingly likely that a and b are equal. By repeating this procedure a modest number of times, we can draw a conclusion about whether a and b are the same, with high statistical confidence.

This procedure allows us to sort the fluxon pairs into bins, where each pair in a bin has the same flux. If a bin contains n pairs, its state is, in general, a mixture of states of the form

$$\sum_{a \in G} \psi_a |a, a^{-1}\rangle^{\otimes n} . \tag{9.50}$$

By discarding just one pair in the bin, each such state becomes a mixture

$$\sum_{a \in g} \rho_a \left(|a\rangle \langle a| \right)^{\otimes (n-1)} ; \qquad (9.51)$$

we may regard each bin as containing (n-1) pairs, all with the same definite flux, but where that flux is as yet unknown.

Which bin is which? We want to label the bins with elements of G. To arrive at a consistent labeling, we withdraw fluxon pairs from three different bins. Suppose the three pairs are $|a, a^{-1}\rangle$, $|b, b^{-1}\rangle$, and $|c, c^{-1}\rangle$, and that we want to check whether c = ab. We create a chargeon-antichargeon pair, carry the chargeon around a closed path that encloses the first member of the first fluxon pair, the first member of the second fluxon pair, and second member of the third fluxon pair, and observe whether the reunited chargeon pair annihilates or not. Since the total flux enclosed by the chargeon's path is abc^{-1} , by repeating this procedure we can determine with high statistical confidence whether ab and c are the same. Such observations allow us to label the bins in some manner that is consistent with the group composition rule. This labeling is unique apart from group automorphisms (and ambiguities arising from any automorphisms may be resolved arbitrarily).

Once the flux bureau of standards is established, we can use it to measure the unknown flux of an unlabeled pair. If the state of the pair to be measured is $|d, d^{-1}\rangle$, we can withdraw the labeled pair $|a, a^{-1}\rangle$ from a bin, and use chargeon pairs to measure the flux ad^{-1} . By repeating this procedure with other labeled fluxes, we can eventually determine the value of the flux d, realizing a projective measurement of the flux.

For a simulation of a quantum circuit using fluxons, we will need to perform logic gates that act upon the value of the flux. The basic gate we will use is realized by winding counterclockwise a fluxon pair with state $|a, a^{-1}\rangle$ around the first member of another fluxon pair with state $|b, b^{-1}\rangle$. Since the $|a, a^{-1}\rangle$ pair has trivial total flux, the $|b, b^{-1}\rangle$ pair is unaffected by this procedure. But since in effect the flux *b* travels counterclockwise about both members of the pair whose initial state was $|a, a^{-1}\rangle$, this pair is transformed as

$$|a, a^{-1}\rangle \mapsto |bab^{-1}, ba^{-1}b^{-1}\rangle$$
 . (9.52)

We will refer to this operation as the *conjugation gate* acting on the fluxon pair.

To summarize what has been said so far, our primitive and derived capabilities allow us to: (1) Perform a projective flux measurement, (2) perform a destructive measurement that determines whether or not the flux and charge of a pair is trivial, and (3) execute a conjugation gate. Now we must discuss how to simulate a quantum circuit using these capabilities.

The next step is to decide how to encode qubits using fluxons. Appropriate encodings can be chosen in many ways; we will stick to one particular choice that illustrates the key ideas — namely we will encode a qubit by using a pair of fluxons, where the total flux of the pair is trivial. We select two noncommuting elements $a, b \in G$, where $b^2 = e$, and choose a computational basis for the qubit

$$|\bar{0}\rangle = |a, a^{-1}\rangle , \quad |\bar{1}\rangle = |bab^{-1}, ba^{-1}b^{-1}\rangle .$$
 (9.53)

The crucial point is that a single isolated fluxon with flux *a* looks identical to a fluxon with the conjugate flux bab^{-1} . Therefore, if the two fluxons in a pair are kept far apart from one another, local interactions with the environment will not cause a superposition of the states $|\bar{0}\rangle$ and $|\bar{1}\rangle$ to decohere. The quantum information is protected from damage because it is stored nonlocally, by exploiting a topological degeneracy of the states where the fluxon and antifluxon are pinned to fixed and distantly separated positions.

However, in contrast with the topological degeneracy that arises in systems with abelian anyons, this protected qubit can be measured relatively easily, without resorting to delicate interferometric procedures that extract Aharonov-Bohm phases. We have already described how to measure flux using previously calibrated fluxons; therefore we can perform a projective measurement of the encoded Pauli operator \bar{Z} (a projection onto the basis $\{|\bar{0}\rangle, |\bar{1}\rangle\}$). We can also measure the complementary Pauli operator \bar{X} , albeit destructively and imperfectly. The \bar{X} eigenstates are

$$|\pm\rangle = \frac{1}{\sqrt{2}} \left(|\bar{0}\rangle \pm |\bar{1}\rangle \right) \equiv \frac{1}{\sqrt{2}} \left(|a, a^{-1}\rangle \pm |bab^{-1}, ba^{-1}b^{-1}\rangle \right) ; \qquad (9.54)$$

therefore the state $|-\rangle$ is *orthogonal* to the zero-charge state

$$|0;\alpha\rangle = \frac{1}{\sqrt{|\alpha|}} \left(\sum_{c \in \alpha} |c, c^{-1}\rangle \right) , \qquad (9.55)$$

where α is the conjugacy class that contains *a*. On the other hand, the state $|+\rangle$ has a nonzero overlap with $|0;\alpha\rangle$

$$\langle +|0;\alpha\rangle = \sqrt{2/|\alpha|} ; \qquad (9.56)$$

Therefore, if the two members of the fluxon pair are brought together, complete annihilation is impossible if the state of the pair is $|-\rangle$, and occurs with probability $\operatorname{Prob}(0) = 2/|\alpha|$ if the state is $|+\rangle$.

Note that it is also possible to prepare a fluxon pair in the state $|+\rangle$. One way to do that is to create a pair in the state $|0; \alpha\rangle$. If α contains only the two elements a and bab^{-1} we are done. Otherwise, we compare the newly created pair with calibrated pairs in each of the states $|c, c^{-1}\rangle$, where $c \in \alpha$ and c is distinct from both a and bab^{-1} . If the pair fails to match any of these $|c, c^{-1}\rangle$ pairs, its state must be $|+\rangle$.

To go further, we need to characterize the computational power of the conjugation gate. Let us use a more compact notation, in which the state $|x, x^{-1}\rangle$ of a fluxon pair is simply denoted $|x\rangle$, and consider the transformations of the state $|x, y, z\rangle$ that can be built from conjugation gates. By winding the third pair through the first, either counterclockwise or clockwise, we can execute the gates

$$|x, y, z\rangle \mapsto |x, y, xzx^{-1}\rangle$$
, $|x, y, z\rangle \mapsto |x, y, x^{-1}zx\rangle$, (9.57)

and by winding the third pair through the second, either counterclockwise or clockwise, we can execute

$$|x, y, z\rangle \mapsto |x, y, yzy^{-1}\rangle$$
, $|x, y, z\rangle \mapsto |x, y, y^{-1}zy\rangle$; (9.58)

furthermore, by borrowing a pair with flux $|c\rangle$ from the bureau of standards, we can execute

$$|x, y, z\rangle \mapsto |x, y, czc^{-1}\rangle \tag{9.59}$$

for any constant $c \in G$. Composing these elementary operations, we can execute any gate of the form

$$|x, y, z\rangle \mapsto |x, y, fzf^{-1}\rangle$$
, (9.60)

where the function f(x, y) can be expressed in *product form* — that is, as a finite product of group elements, where the elements appearing in

the product may be the inputs x and y, their inverses x^{-1} and y^{-1} , or constant elements of G, each of which may appear in the product any number of times.

What are the functions f(x, y) that can be expressed in this form? The answer depends on the structure of the group G, but the following characterization will suffice for our purposes. Recall that a subgroup Hof a finite group G is normal if for any $h \in H$ and any $g \in G$, $ghg^{-1} \in H$, and recall that a finite group G is said to be simple if G has no normal subgroups other than G itself and the trivial group $\{e\}$. It turns out that if G is a simple nonabelian finite group, then any function f(x, y) can be expressed in product form. In the computer science literature, a closely related result is often called Barrington's theorem.

In particular, then, if the group G is a nonabelian simple group, there is a function f realizable in product form such that

$$f(a,a) = f(a,bab^{-1}) = f(bab^{-1},a) = e$$
, $f(bab^{-1},bab^{-1}) = b$. (9.61)

Thus for $x, y, z \in \{a, bab^{-1}\}$, the action eq. (9.60) causes the flux of the third pair to "flip" if and only if $x = y = bab^{-1}$; we have constructed from our elementary operations a Toffoli gate in the computational basis. Therefore, conjugation gates suffice for universal reversible *classical* computation acting on the standard basis states.

The nonabelian simple group of minimal order is A_5 , the group of even permutations of five objects, with $|A_5| = 60$. Therefore, one concrete realization of universal classical computation using conjugation gates is obtained by choosing a to be the three-cycle element $a = (345) \in A_5$, and b to be the product of two-cycles $b = (12)(34) \in A_5$, so that $bab^{-1} = (435)$.

With this judicious choice of the group G, we achieve a topological realization of universal classical computation, but how can be go still further, to realize universal quantum computation? We have the ability to prepare computational basis states, to measure in the computational basis, and to execute Toffoli gates, but these tools are entirely classical. The only nonclassical tricks at our disposal are the ability to prepare $\bar{X} = 1$ eigenstates, and the ability to perform an imperfect destructive measurement of \bar{X} . Fortunately, these additional capabilities are sufficient.

In our previous discussions of quantum fault tolerance, we have noted that if we can do the classical gates Toffoli and CNOT, it suffices for universal quantum computation to be able to apply each of the Pauli operators X, Y, and Z, and to be able to perform projective measurements of each of X, Y, and Z. We already know how to apply the classical gate X and to measure Z (that is, project onto the computational basis). Projective measurement of X and Y, and execution of Z, are still missing from our repertoire. (Of course, if we can apply X and Z, we can also apply their product ZX = iY.) Next, let's see how to elevate our imperfect destructive measurement of X to a reliable projective measurement of X. Recall the action by conjugation of a CNOT on Pauli operators:

$$CNOT: XI \mapsto XX , \qquad (9.62)$$

where the first qubit is the control and the second qubit is the target of the CNOT. Therefore, CNOT gates, together with the ability to prepare X = 1 eigenstates and to perform destructive measurements of X, suffice to realize projective measurements of X. We can prepare an ancilla qubit in the X = 1 eigenstate, perform a CNOT with the ancilla as control and the data to be measured as target, and then measure the ancilla destructively. The measurement prepares the data in an eigenstate of X, whose eigenvalue matches the outcome of the measurement of the ancilla. In our case, the destructive measurement is not fully reliable, but we can repeat the measurement multiple times. Each time we prepare and measure a fresh ancilla, and after a few repetitions, we have acceptable statistical confidence in the inferred outcome of the measurement.

Now that we can measure X projectively, we can prepare X = -1 eigenstates as well as X = 1 eigenstates (for example, we follow a Z measurement with an X measurement until we eventually obtain the outcome X = -1). Then, by performing a CNOT gate whose target is an X = -1 eigenstate, we can realize the Pauli operator Z acting on the control qubit. It only remains to show that a measurement of Y can be realized.

Measurement of Y seems problematic at first, since our physical capabilities have not provided any means to distinguish between Y = 1 and Y = -1 eigenstates (that is, between a state ψ and its complex conjugate ψ^*). However, this ambiguity actually poses no serious difficulty, because it makes no difference how the ambiguity is resolved. Were we to replace measurement of Y by measurement of -Y in our simulation of a unitary transformation U, the effect would be that U^* is simulated instead; this replacement would not alter the probability distributions of outcomes for measurements in the standard computational basis.

To be explicit, we can formulate a protocol for measuring Y by noting first that applying a Toffoli gate whose target qubit is an X = -1 eigenstate realizes the controlled-phase gate $\Lambda(Z)$ acting on the two control qubits. By composing this gate with the CNOT gate $\Lambda(X)$, we obtain the gate $\Lambda(iY)$ acting as

$$\begin{split} \Lambda(iY): & |X = +1\rangle \otimes |Y = +1\rangle \mapsto |Y = +1\rangle \otimes |Y = +1\rangle ,\\ & |X = +1\rangle \otimes |Y = -1\rangle \mapsto |Y = -1\rangle \otimes |Y = -1\rangle ,\\ & |X = -1\rangle \otimes |Y = +1\rangle \mapsto |Y = -1\rangle \otimes |Y = +1\rangle ,\\ & |X = -1\rangle \otimes |Y = -1\rangle \mapsto |Y = +1\rangle \otimes |Y = -1\rangle , \end{split}$$

where the first qubit is the control and the second is the target. Now suppose that my trusted friend gives me just one qubit that he assures me has been prepared in the state $|Y = 1\rangle$. I know how to prepare $|X = 1\rangle$ states myself and I can execute $\Lambda(iY)$ gates; therefore since a $\Lambda(iY)$ gate with $|Y = 1\rangle$ as its target transforms $|X = 1\rangle$ to $|Y = 1\rangle$, I can make many copies of the $|Y = 1\rangle$ state I obtained from my friend. When I wish to measure Y, I apply the inverse of $\Lambda(iY)$, whose target is the qubit to be measured, and whose control is one of my Y = 1 states; then I perform an X measurement of the ancilla to read out the result of the Y measurement of the other qubit.

What if my friend lies to me, and gives me a copy of the state $|Y = -1\rangle$ instead? Then I'll make many copies of the $|Y = -1\rangle$ state, and I will be measuring -Y when I think I am measuring Y. My simulation will work just the same as before; I'll actually be simulating the complex conjugate of the ideal circuit, but that won't change the final outcome of the quantum computation. If my friend flipped a coin to decide whether to give me the $|Y = 1\rangle$ state or the $|Y = -1\rangle$, this too would have no effect on the fidelity of my simulation. Therefore, it turns out I don't need by friend's help at all — instead of using the $|Y = 1\rangle$ state I would have received from him, I may use the random state $\rho = I/2$ (an equally weighted mixture of $|Y = 1\rangle$ and $|Y = -1\rangle$, which I know how to prepare myself).

This completes the demonstration that we can simulate a quantum circuit efficiently and fault tolerantly using the fluxons and chargeons of a nonabelian superconductor, at least in the case where G is a simple nonabelian finite group.[§] Viewed as a whole, including all state preparation and calibration of fluxes, the simulation can be described this way: Many pairs of anyons (fluxons and chargeons) are prepared, the anyon world lines follow a particular braid, and pairs of anyons are fused to see whether they will annihilate. The simulation is nondeterministic in the sense that the actual braid executed by the anyons depends on the outcomes of measurements performed (via fusion) during the course of the simulation. It is robust if the temperature is low compared to the energy gap, and if particles are kept sufficiently far apart from one another (except when pairs are being created and fused), to suppress the exchange of virtual anyons. Small deformations in the world lines of the particles have no effect on the outcome of the computation, as long as the braiding of the particles is in the correct topological class.

[§] Mochon has shown that universal quantum computation is possible for a larger class of groups.

9.12 Anyon models generalized

Our discussion of the nonabelian superconductor model provides an existence proof for fault-tolerant quantum computation using anyons. But the model certainly has drawbacks. The scheme we described lacks beauty, elegance, or simplicity.

I have discussed this model in such detail because it is rather concrete and so helps us to build intuition about the properties of nonabelian anyons. But now that we understand better the key concepts of braiding and fusing in anyon models, we are ready to start thinking about anyons in a more general and abstract way. Our new perspective will lead us to new models, including some that are far simpler than those we have considered so far. We will be able to jettison much of the excess baggage that burdened the nonabelian superconductor model, such as the distinction between fluxons and chargeons, the calibration of fluxes, and the measurements required to simulate nonclassical gates. The simpler models we will now encounter are more naturally conducive to fault-tolerant computing, and more plausibly realizable in reasonable physical systems.

A model of anyons is a theory of particles on a two-dimensional surface (which we will assume to be the plane), where the particles carry locally conserved *charges*. We also assume that the theory has a mass gap, so that there are no long-range interactions between particles mediated by massless particles. The model has three defining properties:

- 1. A list of particle *types*. The types are labels that specify the possible values of the conserved charge that a particle can carry.
- 2. Rules for *fusing* and *splitting*, which specify the possible values of the charge that can be obtained when two particles of known charge are combined together, and the possible ways in which the charge carried by a single particle can be split into two parts.
- **3.** Rules for *braiding*, which specify what happens when two particles are exchanged (or when one particle is rotated by 2π).

Let's now discuss each of these properties in more detail.

9.12.1 Labels

I will use Latin letters $\{a, b, c, ...\}$ for the labels that distinguish different types of particles. (For the case of the nonabelian superconductor, the label was $(\alpha, R^{(\alpha)})$, specifying a conjugacy class and an irreducible representation of the normalizer of the class, but now our notation will be more compact). We will assume that the set of possible labels is finite. The symbol *a* represents the value of the conserved charge carried by the particle. Sometimes we say that this label specifies a *superselection sector* of the theory. This term just means that the label *a* is a property of a localized object that cannot be changed by any local physical process. That is, if one particle is at all times well isolated from other particles, its label will never change. In particular, local interactions between the particle and its environment may jostle the particle, but will not alter the label. This local conservation of charge is the essential reason that anyons are amenable to fault-tolerant quantum information processing.

There is one special label — the identity label 1. A particle with the label 1 is really the same thing as no particle at all. Furthermore, for each particle label a there is a conjugate label \bar{a} , and there is a *charge conjugation* operation C (where $C^2 = I$) acting on the labels that maps a label to its conjugate:

$$C: a \mapsto \bar{a} \mapsto a \ . \tag{9.64}$$

It is possible for a label to be self-conjugate, so that $\bar{a} = a$. For example, $\bar{1} = 1$.

We will want to consider states of n particles, where the particles have a specified order. Therefore, it is convenient to imagine that the particles are arranged on a particular line (such as the real axis) from left to right in consecutive order. The n particles are labeled $(a_1, a_2, a_3 \dots, a_n)$, where a_1 is attached to the particle furthest to the left, a_n to the particle furthest to the right.

9.12.2 Fusion spaces

When two particles are combined together, the composite object also has a charge. The *fusion rules* of the model specify the possible values of the total charge c when the constituents have charges a and b. These can be written

$$a \times b = \sum_{c} N_{ab}^{c} c , \qquad (9.65)$$

where each N_{ab}^c is a nonnegative integer and the sum is over the complete set of labels. Note that a, b and c are labels, not vector spaces; the product on the left-hand side is not a tensor product and the sum on the right-hand side is not a direct sum. Rather, the fusion rules can be regarded as an abstract relation on the label set that maps the ordered triple (a, b; c) to N_{ab}^c . This relation is symmetric in a and b $(a \times b = b \times a)$ — the possible charges of the composite do not depend on whether a is on the left or the right. Read backwards, the fusion rules specify the possible ways for the charge c to split into two parts with charges a and b.

If $N_{ab}^c = 0$, then charge c cannot be obtained when we combine a and b. If $N_{ab}^c = 1$, then c can be obtained — in a unique way. If $N_{ab}^c > 1$,

then c can be obtained in N_{ab}^c distinguishable ways. The notion that fusing two charges can yield a third charge in more than one possible way should be familiar from group representation theory. For example, the rule governing the fusion of two octet representations of SU(3) is

$$8 \times 8 = 1 + 8 + 8 + 10 + \overline{10} + 27 , \qquad (9.66)$$

so that $N_{88}^8 = 2$. We emphasize again, however, that while the fusion rules for group representations can be interpreted as a decomposition of a tensor product of vector spaces as a direct sum of vector spaces, in general the fusion rules in an anyon model have no such interpretation.

The N_{ab}^c distinguishable ways that c can arise by fusing a and b can be regarded as the orthonormal basis states of a Hilbert space V_{ab}^c . We call V_{ab}^c a *fusion space* and the states it contains *fusion states*. The basis elements for V_{ab}^c may be denoted

$$\{|ab; c, \mu\rangle, \quad \mu = 1, 2, \dots, N_{ab}^c\}.$$
 (9.67)

It is quite convenient to introduce a graphical notation for the fusion basis states:



The state $|ab; c, \mu\rangle$ is represented as a circle containing the symbol μ ; connected to the circle are lines labeled a and b with incoming arrows, representing the charges being fused, and a line labeled c with an outgoing arrow, representing the result of the fusion. There is a dual vector space V_c^{ab} describing the states that arise when charge c splits into charges aand b, and a dual basis with the sense of the arrow reversed (c coming in, a and b going out). The spaces V_{ab}^c with different values of c are mutually orthogonal, so that the fusion basis elements satisfy

$$\langle ab; c'\mu' | ab; c, \mu \rangle = \delta_c^{c'} \delta_\mu^{\mu'} , \qquad (9.68)$$

and the completeness of the fusion basis can be expressed as

$$\sum_{c,\mu} |ab;c,\mu\rangle\langle ab;c,\mu| = I_{ab} , \qquad (9.69)$$

where I_{ab} denotes the projector onto the space $\bigoplus_c V_{ab}^c$, the full Hilbert space for the anyon pair ab.



There are some natural isomorphisms among fusion spaces. First of all, $V_{ab}^c \cong V_{ba}^c$; these vector spaces are associated with different labelings of the two particles (if $a \neq b$) and so should be regarded as distinct, but they are isomorphic spaces because fusion is symmetric. We may also "raise and lower indices" of a fusion space by replacing a label by its conjugate, *e.g.*,

$$V_{ab}^c \cong V_{a\bar{c}}^{\bar{b}} \cong V_{ab\bar{c}}^1 \cong V_a^{\bar{b}c}, \cong V_{\bar{c}}^{\bar{a}\bar{b}} \cong \cdots; \qquad (9.70)$$

in the diagrammatic notation, we have the freedom to reverse the sense of a line while conjugating the line's label. The space $V_{ab\bar{c}}^1$, represented as a diagram with three incoming lines, is the space spanned by the distinguishable ways to obtain the trivial total charge 1 when fusing three particles with labels a, b, \bar{c} .

The charge 1 deserves its name because it fuses trivially with other particles:

$$a \times 1 = a \ . \tag{9.71}$$

Because of the isomorphism $V_{a1}^a \cong V_{a\bar{a}}^1$, we conclude that \bar{a} is the unique label that can fuse with a to yield 1, and that this fusion can occur in only one way. Similarly, $V_{a1}^a \cong V_1^{a\bar{a}}$ means that pairs of particles created out of the vacuum have conjugate charges.

An anyon model is *nonabelian* if

$$\dim\left(\bigoplus_{c} V_{ab}^{c}\right) = \sum_{c} N_{ab}^{c} \ge 2$$
(9.72)

for at least some pair of labels *ab*; otherwise the model is *abelian*. In an abelian model, any two particles fuse in a unique way, but in a nonabelian model, there are some pairs of particles that can fuse in more than one way, and there is a Hilbert space of two or more dimensions spanned by these distinguishable states. We will refer to this space as the "topological

Hilbert space" of the pair of anyons, to emphasize that this quantum information is encoded nonlocally — it is a collective property of the pair, not localized on either particle. Indeed, when the two particles with labels a and b are far apart, different states in the topological Hilbert space look identical locally. Therefore, this quantum information is well hidden, and invulnerable to decoherence due to local interactions with the environment.

It is for this reason that we propose to use nonabelian anyons in the operation of a quantum computer. Of course, nonlocally encoded information is not only hidden from the environment; we are unable to read it ourselves as well. However, with nonabelian anyons, we can have our cake and eat it too! At the conclusion of a quantum computation, when we are ready to perform the readout, we can bring the anyons together in pairs and observe the result of this fusion. In fact, it will suffice to distinguish the case where the charge of the composite is c = 1 from the case $c \neq 1$ — that is, to distinguish a residual particle (unable to decay because of its nontrivial conserved charge) from no particle at all.

Note that for each pair of anyons this topological Hilbert space is finitedimensional. An anyon model with this property is said to be *rational*. As in our discussion of the topologically degenerate ground state for an abelian model, anyons in rational nonabelian models always have topological spins that are roots of unity.

9.12.3 Braiding: the R-matrix

When two particles with labels a and b undergo a counterclockwise exchange, their total charge c is unchanged. Therefore, since the two particles swap positions on the line, the swap induces a natural isomorphism mapping the Hilbert space V_{ba}^c to V_{ab}^c ; this map is the braid operator

$$R: V_{ba}^c \to V_{ab}^c \ . \tag{9.73}$$

If we choose canonical bases $\{|ba; c, \mu\rangle\}$ and $\{|ab; c, \mu'\rangle\}$ for these two spaces, R can be expressed as the unitary matrix

$$R: |ba; c, \mu\rangle \mapsto \sum_{\mu'} |ab; c, \mu'\rangle \left(R^c_{ab}\right)^{\mu'}_{\mu} ; \qquad (9.74)$$

note that R may have a nontrivial action on the fusion states. When we represent the action of R diagrammatically, it is convenient to fix the positions of the labels a and b on the incoming lines, and twist the lines counterclockwise as they move toward the fusion vertex (μ) — the graph with twisted lines represents the state in V_{ab}^c obtained by applying R to $|ba; c, \mu\rangle$, which can be expanded in terms of the canonical basis for V_{ab}^c :



The *monodromy* operator

$$R^2: V_{ab}^c \to V_{ab}^c \tag{9.75}$$

is an isomorphism from V_{ab}^c to itself, representing the effect of winding *a* counterclockwise around *b*. As we already remarked in our discussion of the nonabelian superconductor, the monodromy operator is equivalent to rotating *c* by 2π while rotating *a* and *b* by -2π ; therefore, the eigenvalues of the monodromy operator are determined by the *topological spins* of the particles:

$$(R_{ab}^c)^2 = e^{-2\pi i J_c} e^{2\pi i J_a} e^{2\pi i J_b} \equiv e^{i(\theta_c - \theta_a - \theta_b)} .$$
(9.76)

Furthermore, as we argued for the case of abelian anyons, the topological spin is determined by the braid operator acting on a particle-antiparticle pair with trivial total charge:

$$e^{-i\theta_a} = R^1_{a\bar{a}} \tag{9.77}$$

(because creating a pair, exchanging, and annihilating is equivalent to rotating the particle by -2π).

9.12.4 Associativity of fusion: the F-matrix

Fusion is associative:

$$(a \times b) \times c = a \times (b \times c) . \tag{9.78}$$

Mathematically, this is an axiom satisfied by the fusion rules of an anyon model. Physically, it is imposed because the total charge of a system of three particles is an intrinsic property of the three particles, and ought not to depend on whether we first fuse a and b and then fuse the result with c, or first fuse b and c and then fuse the result with a.

Therefore, when three particles with charges a, b, c are fused to yield a total charge of d, there are two natural ways to decompose the topological Hilbert space in terms of the fusion spaces of pairs of particles:

$$V_{abc}^{d} \cong \bigoplus_{e} V_{ab}^{e} \otimes V_{eb}^{d} \cong \bigoplus_{e'} V_{ae'}^{d} \otimes V_{bc}^{e'} .$$

$$(9.79)$$

9 Topological quantum computation

Correspondingly, there are two natural orthonormal bases for V^d_{abc} , which we may denote

$$\begin{aligned} |(ab)c \to d; e\mu\nu\rangle &\equiv |ab; e, \mu\rangle \otimes |ec; d, \nu\rangle ,\\ |a(bc) \to d; e'\mu'\nu'\rangle &\equiv |ae'; d, \nu'\rangle \otimes |bc; e', \mu'\rangle , \end{aligned}$$
(9.80)

and which are related by a unitary transformation F:

$$|(ab)c \to d; e\mu\nu\rangle = \sum_{e'\mu'\nu'} |a(bc) \to d; e'\mu'\nu'\rangle \left(F^d_{abc}\right)^{e'\mu'\nu'}_{e\mu\nu} . \tag{9.81}$$



The unitary matrices F_{abc}^d are sometimes called *fusion matrices*; however, rather than risk causing confusion between F and the fusion *rules* N_{ab}^c , I will just call it the *F*-matrix.

9.12.5 Many anyons: the standard basis

In an anyonic quantum computer, we process the topological quantum state of n anyons by braiding the anyons. For describing this computation, it is convenient to adopt a standard basis for such a Hilbert space.

Suppose that n anyons with total charge c, arranged sequentially along a line, carry labels $a_1, a_2, a_3, \ldots, a_n$. Imagine fusing anyons 1 and 2, then fusing the result with anyon 3, then fusing the result with anyon 4, and so on. Associated with fusion in this order is a decomposition of the topological Hilbert space of the n anyons

$$V_{a_{1}a_{2}a_{3}\cdots a_{n}}^{c} \cong \bigoplus_{b_{1},b_{2},\dots,b_{n-2}} V_{a_{1}a_{2}}^{b_{1}} \otimes V_{b_{1}a_{3}}^{b_{2}} \otimes V_{b_{2}a_{4}}^{b_{3}} \otimes \cdots \otimes V_{b_{n-2}a_{n}}^{c} .$$
(9.82)

Note that this space does *not* have a natural decomposition as a tensor product of subsystems associated with the localized particles; rather, we have expressed it as a direct sum of many tensor products. For nonabelian anyons, its dimension

$$\dim \left(V_{a_1 a_2 a_3 \cdots a_n}^c \right) \equiv N_{a_1 a_2 a_3 \cdots a_n}^c$$

=
$$\sum_{b_1, b_2, b_3, \dots b_{n-2}} N_{a_1 a_2}^{b_1} N_{b_1 a_3}^{b_2} N_{b_2 a_4}^{b_3} \dots N_{b_{n-2} a_n}^c$$
(9.83)

46

is exponential in n; thus the topological Hilbert space is a suitable arena for powerful quantum information processing.

This decomposition of $V_{a_1a_2a_3\cdots a_n}^c$ suggests a standard basis whose elements are labeled by the intermediate charges $b_1, b_2, \ldots b_{n-2}$ and by the basis elements $\{|\mu_j\rangle\}$ for the fusion spaces $V_{b_{j-1},a_{j+1}}^{b_j}$:

$$\{|a_1a_2; b_1, \mu_1\rangle | b_1a_3; b_2, \mu_2\rangle \cdots | b_{n-3}a_{n-1}; b_{n-2}, \mu_{n-2}\rangle | b_{n-2}a_n; c, \mu_{n-1}\rangle\},$$
(9.84)

or in diagrammatic notation:



Of course, this basis is chosen arbitrarily. If we preferred, we could imagine fusing the particles in a different order, and would obtain a different basis that can be expressed in terms of our standard one with help from the F-matrix.

9.12.6 Braiding in the standard basis: the B-matrix

We would like to consider what happens to states of the topological vector space $V_{a_1a_2a_3\cdots a_n}^c$ of n anyons when the particles are exchanged with one another. Actually, since exchanges can swap the positions of particles with distinct labels, they may map one topological vector space to another by permuting the labels. Nevertheless, we can consider the direct sums of the vector spaces associated with all the possible permutations of the labels, which will provide a representation of the braid group B_n .

We would like to describe how this representation acts on the standard bases for these spaces. It suffices to say how exchanges of neighboring particles are represented; that is, to specify the action of the generators of the braid group. However, so far, we have discussed only the action of the braid group on a pair of particles with definite total charge (the R-matrix), which is not in itself enough to tell us its action on the standard bases.

The way out of this quandary is to observe that, by applying the F-matrix, we can move from the standard basis to the basis in which the R-matrix is block diagonal, apply R, and then apply F^{-1} to return to the standard basis:



The composition of these three operations, which expresses the effect of braiding in the standard basis, is denoted B and sometimes called the "braid matrix;" but to avoid confusion between B and R, I will just call it the *B*-matrix.

Consider exchanging the anyons in positions j and j+1 along the line. In our decomposition of $V_{a_1a_2a_3\cdots a_n}^c$, this exchange acts on the space

$$V_{b_{j-2,a_j,a_{j+1}}}^{b_j} = \bigoplus_{b_{j-1}} V_{b_{j-2,a_j}}^{b_{j-1}} \otimes V_{b_{j-1},a_{j+1}}^{b_j} .$$
(9.85)

To reduce the number of subscripts, we will call this space V^d_{acb} , which is transformed by the exchange as

$$B: V_{acb}^d \to V_{abc}^d . \tag{9.86}$$

Let us express the action of B in terms of the standard bases for the two spaces V^d_{acb} and V^d_{abc} .



To avoid cluttering the equations, I suppress the labels for the fusion space basis elements (it is obvious where they should go). Hence we write

$$B|(ac)b \to d; e\rangle = \sum_{f} B|a(cb) \to d; f\rangle \left(F_{acb}^{d}\right)_{e}^{f}$$

$$= \sum_{f} |a(bc) \to d; f\rangle R_{bc}^{f} \left(F_{acb}^{d}\right)_{e}^{f}$$

$$= \sum_{f,g} |(ab)c \to d; g\rangle \left[\left(F^{-1}\right)_{abc}^{d} \right]_{f}^{g} R_{bc}^{f} \left(F_{acb}^{d}\right)_{e}^{f},$$
(9.87)
or

$$B: |(ac)b \to d; e\rangle \mapsto \sum_{g} |(ab)c \to d; g\rangle \left(B^{d}_{abc}\right)^{g}_{e} , \qquad (9.88)$$

where

$$\left(B^d_{abc}\right)^g_e = \sum_f \left[\left(F^{-1}\right)^d_{abc} \right]^g_f R^f_{bc} \left(F^d_{acb}\right)^f_e \ . \tag{9.89}$$

We have expressed the action of the B-matrix in the standard basis in terms of the F-matrix and R-matrix, as desired.

Thus, the representation of the braid group realized by n anyons is completely characterized by the *F*-matrix and the *R*-matrix. Furthermore, we have seen that the *R* matrix also determines the topological spins of the anyons, so that we have actually constructed a representation of a larger group whose generators include both the exchanges of neighboring particles and 2π rotations of the particles. A good name for this group would be the *ribbon* group, as its elements are in one-to-one correspondence with the topological classes of braided ribbons (which can be twisted) rather than braided strings; however, mathematicians have already named it "the mapping class group for the sphere with *n* punctures."

And with that observation we have completed our description of an anyon model in this general setting. The model is specified by: (1) a label set, (2) the fusion rules, (3) the *R*-matrix, and (4) the *F* matrix.

The mathematical object we have constructed is called a *unitary topological modular functor*, and it is closely related to two other objects that have been much studied: *topological quantum field theories* in 2+1 space-time dimensions, and *conformal field theories* in 1+1 spacetime dimensions. However, we will just call it an *anyon model*.

9.13 Simulating anyons with a quantum circuit

A topological quantum computation is executed in three steps:

- **1.** Initialization: Particle-antiparticle pairs $c_1\bar{c}_1, c_2\bar{c}_2, c_3\bar{c}_3, \ldots, c_m\bar{c}_m$ are created. Each pair is of a specified type and has trivial total charge.
- **2.** Processing. The n = 2m particles are guided along trajectories, their world lines following a specified braid.
- **3.** *Readout.* Pairs of neighboring particles are fused together, and it is recorded whether each pair annihilates fully or not. This record is the output of the computation.

(In the case of the nonabelian superconductor model of computation, we allowed the braiding to be conditioned on the outcome of fusing carried out during the processing stage. But now we are considering a model in which all measurements are delayed until the final readout.)

How powerful is this model of computation? I claim that this topological quantum computer can be simulated efficiently by a quantum circuit. Since the topological Hilbert space of n anyons does not have a simple and natural decomposition as a tensor product of small subsystems, this claim may not be immediately obvious. To show it we must explain:

1. How to encode the topological Hilbert space using ordinary qubits.

- 2. How to represent braiding efficiently using quantum gates.
- **3.** How to simulate the fusion of an anyon pair.

Encoding. Since each pair produced during initialization has trivial total charge, the initial state of the n anyons also has trivial total charge. Therefore, the topological Hilbert space is

$$V_{a_1 a_2 a_3 \cdots a_n}^1 \cong \bigoplus_{b_1, b_2, \dots, b_{n-3}} V_{a_1 a_2}^{b_1} \otimes V_{b_1 a_3}^{b_2} \otimes \dots \otimes V_{b_{n-3} a_{n-1}}^{\bar{a}_n} , \qquad (9.90)$$

for some choice of the labels $a_1, a_2, a_3, \ldots a_n$; there are n-3 intermediate charges and n-2 fusion spaces appearing in each summand. Exchanges of the particles swap the labels, but after each exchange the vector space still has the form eq. (9.90) with labels given by some permutation of the original labels.

Although each n-anyon topological Hilbert spaces is not itself a tensor products of subsystems, all of these spaces are contained in

$$\left(\mathcal{H}_d\right)^{\otimes (n-2)} , \qquad (9.91)$$

where

$$\mathcal{H}_d = \bigoplus_{a,b,c} V_{abc}^1 \ . \tag{9.92}$$

Here, a, b, c are summed over the complete label set of the model (which we have assumed is finite), so that \mathcal{H}_d contains all the possible fusion states of three particles, and the dimension d of \mathcal{H}_d is

$$d = \sum_{a,b,c} N_{abc}^1 . (9.93)$$

Thus the state of n anyons can be encoded in the Hilbert space of n-2qudits for some constant d (which depends on the anyon model but is independent of n). The basis states of this qudit can be chosen to be $\{|a, b, c; \mu\rangle\}$, where μ labels an element of the basis for the fusion space V_{abc}^1 .

Braiding. In the topological quantum computer, a braid is executed by performing a sequence of exchanges, each acting on a pair of neighboring particles. The effect of each exchange in the standard basis is described by the *B*-matrix. How is *B* represented acting on our encoding of the topological vector space (using qudits)? Suppressing fusion states, our basis for two-qudit states can be denoted $|a, b, c\rangle |d, e, \bar{f}\rangle$. But in the topological quantum computer, the labels *d* and \bar{c} always match, and therefore to perform our simulation of braiding we need only consider two-qudit states whose labels match in this sense:

$$a \xrightarrow{e}_{\overline{d}} \xrightarrow{b}_{\overline{f}} = \sum_{g} \left(B_{aeb}^{f} \right)_{d}^{g} \xrightarrow{e}_{\overline{g}} \xrightarrow{b}_{\overline{g}} \xrightarrow{f}_{\overline{f}}$$

Then the action of the B-matrix on these basis states is

$$B: |a, b, \bar{d}\rangle |d, e, \bar{f}\rangle \mapsto \sum_{g} |a, e, \bar{g}\rangle |g, b, \bar{f}\rangle \left(B^{f}_{aeb}\right)^{g}_{d} .$$

$$(9.94)$$

As desired, we have represented the B as a $d^2 \times d^2$ matrix acting on a pair of neighboring qudits.

Fusion. Fusion of a pair of anyons can be simulated by a two-qudit measurement, which can be reduced to a single-qudit measurement with a little help from the F-matrix:



Consider a basis state $|a, b, \overline{d}\rangle | d, e, \overline{f} \rangle$ for a pair of neighboring qudits; what is the amplitude for the anyon pair (be) to have trivial total charge? Using an *F*-move, the state can be expanded as

$$F: \qquad |a, b, \bar{d}\rangle | d, e, \bar{f}\rangle \mapsto \sum_{g} |a, g, \bar{f}\rangle | b, \bar{g}, e\rangle \left(F_{abe}^{f}\right)_{d}^{g}$$
$$= |a, 1, \bar{f}\rangle | b, 1, e\rangle \left(F_{abe}^{f}\right)_{d}^{1} + \sum_{g \neq 1} |a, g, \bar{f}\rangle | b, \bar{g}, e\rangle \left(F_{abe}^{f}\right)_{d}^{g} ; (9.95)$$

we have separated the sum over g into the component for which (be) fuses to 1, plus the remainder. After the F-move which (is just a particular two-qudit unitary gate), we can sample the probability that (be) fuses to 1 by performing a projective measurement of the second qudit in the basis $\{|b, \bar{g}, e\rangle\}$, and recording whether g = 1.

This completes our demonstration that a quantum circuit can simulate efficiently a topological quantum computer.

9.14 Fibonacci anyons

Now we have established that topological quantum computation is no more powerful than the quantum circuit model — any problem that can be solved efficiently by braiding nonabelian anyons can also be solved efficiently with a quantum circuit. But is it *as* powerful? Can we simulate a universal quantum computer by braiding anyons? The answer depends on the specific properties of the anyons: some nonabelian anyon models are universal, others are not. To find the answer for a particular anyon model, we need to understand the properties of the representations of the braid group that are determined by the F-matrix and R-matrix.

Rather than give a general discussion, we will study one especially simple nonabelian anyon model, and demonstrate its computational universality. This model is the very simplest nonabelian model — conformal field theorists call it the "Yang-Lee model," but I will call it the "Fibonacci model" for reasons that will soon be clear.

In the Fibonacci model there are only two labels — the trivial label, which I will now denote 0, and a single nontrivial label that I will call 1, where $\overline{1} = 1$. And there is only one nontrivial fusion rule:

$$1 \times 1 = 0 + 1$$
; (9.96)

when two anyons are brought together they either annihilate, or fuse to become a single anyon. The model is nonabelian because two anyons can fuse in two distinguishable ways.

Consider the standard basis for the Hilbert space V_{1n}^b of n anyons, where each basis element describes a distinguishable way in which the n anyons could fuse to give total charge $b \in \{0, 1\}$. If the two anyons furthest to the left were fused first, the resulting charge could be 0 or 1; this charge could then fuse with the third anyon, yielding a total charge of 0 or 1, and so on. Finally, the last anyon fuses with the total charge of the first n-1 anyons to give the total charge b. Altogether n-2 intermediate charges $b_1, b_2, b_3, \ldots b_{n-2}$ appear in this description of the fusion process; thus the corresponding basis element can be designated with a binary string of length n-2. If the total charge is 0, the result of fusing the first n-1 anyons has to be 1, so the basis states are labeled by strings of length n-3.

However, not *all* binary strings are allowed — a 0 must always be followed by a 1. There cannot be two zeros in a row because when the charge 0 fuses with 1, a total charge of 1 is the only possible outcome. Otherwise, there is no restriction on the sequence. Therefore, the basis states are in one-to-one with the binary strings that do not contain two successive 0's.

Thus the dimensions $N_n^0 \equiv N_{1n}^0$ of the topological Hilbert spaces V_{1n}^0 obey a simple recursion relation. If the fusion of the first two particles yields trivial total charge, then the remaining n-2 particles can fuse in N_{n-2}^0 distinguishable ways, and if the fusion of the first two particles yields an anyon with nontrivial charge, then that anyon can fuse with the other n-2 anyons in N_{n-1}^0 ways; therefore,

$$N_n^0 = N_{n-1}^0 + N_{n-2}^0 . (9.97)$$

Since $N_1^0 = 0$ and $N_2^0 = 1$, the solution to this recursion relation is

$$n = 1 \ 2 \ 3 \ 4 \ 5 \ 6 \ 7 \ 8 \ 9 \dots$$

$$N_n^0 = 0 \ 1 \ 1 \ 2 \ 3 \ 5 \ 8 \ 13 \ 21 \dots$$
(9.98)

— the dimensions are Fibonacci numbers (which is why I am calling this model the "Fibonacci model").

The Fibonacci numbers grow with n at a rate $N_n^0 \approx C\phi^n$, where ϕ is the golden mean $\phi = \frac{1}{2} \left(1 + \sqrt{5}\right) \approx 1.618$. Because ϕ governs the rate at which the Hilbert space enlarges as anyons are added, we say that $d = \phi$ is the quantum dimension of the Fibonacci anyon. That this "dimension" is an irrational number illustrates vividly that the topological Hilbert space has no natural decomposition as a tensor product of subsystems — instead, the topologically encoded quantum information is a collective property of the n anyons.

9.15 Quantum dimension

We will return shortly to the properties of the Fibonacci model, but first let's explore more deeply the concept of quantum dimension. For a general anyon model, how should the dimension d_a of label a be defined? For this purpose, it is convenient to imagine a physical process in which two $a\bar{a}$ pairs are created (each with trivial total charge); then the particle a from the pair on the right fuses with the antiparticle \bar{a} from the pair on the left. Do these particles annihilate?

With suitable phase conventions, the *amplitude* for the annihilation to occur is a real number in the unit interval [0,1]. Let us define this

9 Topological quantum computation

number to be $1/d_a$, where d_a is the quantum dimension of a (and $1/d_a^2$ is the *probability* that annihilation occurs). Note that it is clear from this definition that $d_a = d_{\bar{a}}$. For the case in which the a is the label of an irreducible representation R_a of a group G, the dimension is just $d_a = |R_a|$, the dimension of the representation. This is easily understood pictorially:



If two pairs are created and then each pair annihilates immediately, the world lines of the pairs form two closed loops, and |R| counts the number of distinct "colors" that propagate around each loop. But if the particle from each pair annihilates the antiparticle from the other pair, there is only one closed loop and therefore one sum over colors; if we normalize the process on the left to unity, the amplitude for the process on the right is suppressed by a factor of 1/|R|. To say the same thing in an equation, the normalized state of an $R\bar{R}$ pair is

$$|R\bar{R}\rangle = \frac{1}{\sqrt{|R|}} \sum_{i} |i\rangle |\bar{i}\rangle , \qquad (9.99)$$

where $\{|i\rangle\}$ denotes an orthonormal basis for R and $\{|\bar{i}\rangle\}$ is a basis for \bar{R} . Suppose that two pairs $|R\bar{R}\rangle$ and $|R'\bar{R}'\rangle$ are created; if the pairs are fused after swapping partners, the amplitude for annihilation is

$$\langle R\bar{R}, R'\bar{R}'|R\bar{R}', R'\bar{R}\rangle = \frac{1}{|R|^2} \sum_{i,i',j,j'} \langle j\bar{j}, j'\bar{j}'|i\bar{i}', i'\bar{i}\rangle$$

= $\frac{1}{|R|^2} \sum_{i,i',j,j'} \delta_{ji} \delta_{ji'} \delta_{j'i'} \delta_{j'i} = \frac{1}{|R|^2} \sum_i \delta_{ii} = \frac{1}{|R|} .$ (9.100)

In general, though, the quantum dimension has no direct interpretation in terms of counting "colors," and there is no reason why it has to be an integer.

How are such quantum dimensions related to the dimensions of topological Hilbert spaces? To see the connection, if is very useful to alter our normalization conventions. Notice we can introduce many "zigzags" in the world line of a particle of type a by creating many $a\bar{a}$ pairs, and fusing the particle from each pair with the antiparticle from the neighboring pair. However, each zigzag reduces the amplitude by another factor of $1/d_a$. We can compensate for these factors of $1/d_a$ if we weight each pair creation or annihilation event by a factor of $\sqrt{d_a}$. With this new convention, we can bend the world line of a particle forward or backward in time without paying any penalty:



Now the weight assigned to a world line is a topological invariant (it is unchanged when we distort the line), and a world line of type a forming a closed loop is weighted by d_a .

With our new conventions, we can justify this sequence of manipulations:



Each diagram represents an inner product of two (unconventionally normalized) states. We have inserted a complete sum over the labels (c) and the corresponding fusion states (μ) that can arise when a and b fuse. Exploiting the topological invariance of the diagram, we have then turned it "inside out," then contracted the fusion states (acquiring the factor N_{ab}^c which counts the possible values of μ).

The equation that we have derived,

$$d_a d_b = \sum_c N_{ab}^c d_c \equiv \sum_c (N_a)_b^c d_c , \qquad (9.101)$$

says that the vector \vec{d} , whose components are the quantum dimensions, is an eigenvector with eigenvalue d_a of the matrix N_a that describes how the label *a* fuses with other labels:

$$N_a \vec{d} = d_a \vec{d} . \tag{9.102}$$

Furthermore, since N_a has nonnegative entries and all components of \vec{d} are positive, d_a is the largest eigenvalue of N_a and is nondegenerate. (This simple observation is sometimes called the *Perron-Frobenius theorem.*) For n anyons, each with label a, the topological Hilbert space $V_{aaa\cdots a}^b$ for the sector with total charge b has dimension

$$N^{b}_{aaa\cdots a} = \sum_{\{b_i\}} N^{b_1}_{aa} N^{b_2}_{ab_1} N^{b_3}_{ab_2} \dots N^{b}_{ab_{n-2}} = \langle b | (N_a)^{n-1} | a \rangle .$$
(9.103)

The matrix N_a can be diagonalized, and expressed as

$$N_a = |v\rangle d_a \langle v| + \cdots , \qquad (9.104)$$

where

$$|v\rangle = \frac{\vec{d}}{\mathcal{D}}, \quad \mathcal{D} = \sqrt{\sum_{c} d_{c}^{2}}, \qquad (9.105)$$

and subleading eigenvalues have been omitted; therefore

$$N^b_{aaa\cdots a} = d^n_a d_b / \mathcal{D}^2 + \cdots , \qquad (9.106)$$

where the ellipsis represents terms that are exponentially suppressed for large n. We see that the quantum dimension d_a controls the rate of growth of the *n*-particle Hilbert space for anyons of type a.

Because the label 0 with trivial charge fuses trivially, we have $d_0 = 1$. In the case of the Fibonacci model, it follows from the fusion rule $1 \times 1 = 0+1$ that $d_1^2 = 1 + d_1$, which is solved by $d_1 = \phi$ as we found earlier; therefore $\mathcal{D}^2 = d_0^2 + d_1^2 = 1 + \phi^2 = 2 + \phi$. Our formula becomes

$$N_{111\dots 1}^{0} = \left(\frac{1}{2+\phi}\right)\phi^{n} , \qquad (9.107)$$

which is an excellent approximation to the Fibonacci numbers even for modest values of n.

Suppose that an $a\bar{a}$ pair and a $b\bar{b}$ pair are both created. If the a and b particles are fused, with what probability $p(ab \rightarrow c)$ will their total charge be c? This question can be answered using the same kind of graphical manipulations:



Dividing by $d_a d_b$ to restore the proper renormalization of the inner product, we conclude that

$$p(ab \to c) = \frac{N_{ab}^c d_c}{d_a d_b} , \qquad (9.108)$$

which generalizes the formula $p(a\bar{a} \rightarrow 1) = 1/d_a^2$ that we used to define the quantum dimension, and satisfies the normalization condition

$$\sum_{c} p(ab \to c) = 1. \tag{9.109}$$

To arrive at another interpretation of the quantum dimension, imagine that a dense gas of anyons is created, which is then permitted to anneal for awhile — anyons collide and fuse, gradually reducing the population of particles. Eventually, but long before the thermal equilibrium is attained, the collision rate becomes so slow that the fusion process effectively turns off. By this stage, whatever the initial distribution of particles types, a steady state distribution is attained that is preserved by collisions. If in the steady state particles of type a appear with probability p_a , then

$$\sum_{ab} p_a p_b \ p(ab \to c) = p_c \ . \tag{9.110}$$

Using

$$\sum_{a} N_{ab}^{c} d_{a} = \sum_{a} N_{b\bar{c}}^{a} d_{\bar{a}} = d_{b}d_{\bar{c}} = d_{b}d_{c} , \qquad (9.111)$$

we can easily verify that this condition is satisfied by

$$p_a = \frac{d_a^2}{\mathcal{D}^2} \,. \tag{9.112}$$

We conclude that if anyons are created in a random process, those carrying labels with larger quantum dimension are more likely to be produced, in keeping with the property that anyons with larger dimension have more quantum states.

9.16 Pentagon and hexagon equations

To assess the computational power of an anyon model like the Fibonacci model, we need to know the braiding properties of the anyons, which are determined by the R and F matrices. We will see that the braiding rules are highly constrained by algebraic consistency conditions. For the Fibonacci model, these consistency conditions suffice to determine a unique braiding rule that is compatible with the fusion rules.

Consistency conditions arise because we can make a sequence of "F-moves" and "R-moves" to obtain an isomorphism relating two topological Hilbert spaces. The isomorphism can be regarded as a unitary matrix that relates the canonical orthonormal bases for two different spaces; this unitary transformation does not depend on the particular sequence of moves from which the isomorphism is constructed, only on the initial and final bases.

For example, there are five different ways to fuse four particles (without any particle exchanges), which are related by *F*-moves:



The basis shown furthest to the left in this pentagon diagram is the "left standard basis" { $|\text{left}; a, b\rangle$ }, in which particles 1 and 2 are fused first, the resulting charge a is fused with particle 3 to yield charge b, and then finally b is fused with particle 4 to yield the total charge 5. The basis shown furthest to the right is the "right standard basis" { $|\text{right}; c, d\rangle$ }, in which the particles are fused from right to left instead of left to right. Across the top of the pentagon, these two bases are related by two F-

moves, and we obtain

$$|\text{left}; a, b\rangle = \sum_{c,d} |\text{right}; c, d\rangle \left(F_{12c}^{5}\right)_{a}^{d} \left(F_{a34}^{5}\right)_{b}^{c}$$
 (9.113)

Across the bottom of the pentagon, the bases are related by three F-moves, and we find

$$|\text{left}; a, b\rangle = \sum_{c,d,e} |\text{right}; c, d\rangle \left(F_{234}^{d}\right)_{e}^{c} \left(F_{1e4}^{5}\right)_{b}^{d} \left(F_{123}^{b}\right)_{a}^{e} .$$
(9.114)

Equating our two expressions for $|left; a, b\rangle$, we obtain the *pentagon equation*:

$$\left(F_{12c}^{5}\right)_{a}^{d}\left(F_{a34}^{5}\right)_{b}^{c} = \left(F_{234}^{d}\right)_{e}^{c}\left(F_{1e4}^{5}\right)_{b}^{d}\left(F_{123}^{b}\right)_{a}^{e} . \tag{9.115}$$

Another nontrivial consistency condition is found by considering the various ways that three particles can fuse:



The basis { $|\text{left}; a\rangle$ } furthest to the left in this hexagon diagram is obtained if the particles are arranged in the order 123, and particles 1 and 2 are fused first, while the basis { $|\text{right}, c\rangle$ } furthest to the right is obtained if the particles are arranged in order 231, and particles 1 and 3 are fused first. Across the top of the hexagon, the two bases are related by the sequence of moves FRF:

$$|\text{left}, a\rangle = \sum_{b,c} |\text{right}; c\rangle \left(F_{231}^4\right)_b^c R_{1b}^4 \left(F_{123}^4\right)_a^b .$$
 (9.116)

Across the bottom of the hexagon, the bases are related by the sequence of moves RFR, and we find

$$|\text{left}, a\rangle = \sum_{c} |\text{right}; c\rangle R_{13}^{c} \left(F_{213}^{4}\right)_{a}^{c} R_{12}^{a}$$
. (9.117)

Equating our two expressions for $|left; a\rangle$, we obtain the *hexagon equation*:

$$R_{13}^c \left(F_{213}^4\right)_a^c R_{12}^a = \sum_b \left(F_{231}^4\right)_b^c R_{1b}^4 \left(F_{123}^4\right)_a^b .$$
(9.118)

A beautiful theorem, which I will not prove here, says that there are no further conditions that must be imposed to ensure the consistency of braiding and fusing. That is, for any choice of an initial and final basis for n anyons, all sequences of R-moves and F-moves that take the initial basis to the final basis yield the same isomorphism, provided that the pentagon equation and hexagon equation are satisfied. This theorem is an instance of the MacLane coherence theorem, a fundamental result in category theory. The pentagon and hexagon equations together are called the Moore-Seiberg polynomial equations — their relevance to physics was first appreciated in studies of (1+1)-dimensional conformal field theory during the 1980's.

A solution to the polynomial equations defines a viable anyon model. Therefore, there is a systematic procedure for constructing anyon models:

1. Choose a set of labels and assume a fusion rule.

2. Solve the polynomial equations for R and F.

If no solutions exist, then the hypothetical fusion rule is incompatible with the principles of local quantum physics and must be rejected. If there is more than one solution (not related to one another by any reshuffling of the labels, redefinition of bases, etc.), then each distinct solution defines a distinct model with the assumed fusion rule.

To illustrate the procedure, consider the polynomial equations for the Fibonacci fusion rule. There are only two F-matrices that arise, which we will denote as

$$F_{0111} \equiv F_0 , \quad F_{1111} \equiv F_1 .$$
 (9.119)

 F_0 is really the 1×1 matrix

$$(F_0)^b_a = \delta^1_a \delta^b_1 , \qquad (9.120)$$

while F_1 is a 2 × 2 matrix. The pentagon equation becomes

$$(F_c)^d_a (F_a)^c_b = \sum_e (F_d)^c_e (F_e)^d_b (F_b)^c_a . \qquad (9.121)$$

The general solution for $F \equiv F_1$ is

$$F = \begin{pmatrix} \tau & e^{i\phi}\sqrt{\tau} \\ e^{-i\phi}\sqrt{\tau} & -\tau \end{pmatrix} , \qquad (9.122)$$

where $e^{i\phi}$ is an arbitrary phase (which we can set to 1 with a suitable phase convention), and $\tau = (\sqrt{5} - 1)/2 = \phi - 1 \approx .618$, which satisfies

$$\tau^2 + \tau = 1 . (9.123)$$

The 2×2 *R*-matrix that describes a counterclockwise exchange of two Fibonacci anyons has two eigenvalues — R^0 for the case where the total charge of the pair of anyons is trivial, and R^1 for the case where the total charge is nontrivial. The hexagon equation becomes

$$R^{c}(F)_{a}^{c}R^{a} = (F)_{0}^{c}(F)_{a}^{0} + (F)_{1}^{c}R^{1}(F)_{a}^{1} .$$
(9.124)

Using the expression for F found by solving the pentagon equation, we can solve the hexagon equation for R, finding

$$R = \begin{pmatrix} e^{4\pi i/5} & 0\\ 0 & -e^{2\pi i/5} \end{pmatrix} , \quad F = \begin{pmatrix} \tau & \sqrt{\tau}\\ \sqrt{\tau} & -\tau \end{pmatrix} .$$
(9.125)

The only other solution is the complex conjugate of this one; this second solution really describes the same model, but with clockwise and counterclockwise braiding interchanged. Therefore, an anyon model with the Fibonacci fusion rule really *does* exist, and it is essentially unique.

9.17 Simulating a quantum circuit with Fibonacci anyons

Now we know enough to address whether a universal quantum computer can be simulated using Fibonacci anyons. We need to explain how qubits can be encoded with anyons, and how a universal set of quantum gates can be realized.

First we note that the Hilbert space $V_4^0 \equiv V_{1111}^0$ has dimension $N_4^0 = 2$; therefore a qubit can be encoded by four anyons with trivial total charge. The anyons are lined up in order 1234, numbered from left to right; in the standard basis state $|0\rangle$, anyons number 1 and number 2 fuse to yield total charge 0, while in the standard basis state $|1\rangle$, anyons 1 and 2 fuse to yield total charge 1. Acting on this standard basis, the braid group generator σ_1 (counterclockwise exchange of particles 1 and 2) is represented by

$$\sigma_1 \mapsto R = \begin{pmatrix} e^{4\pi i/5} & 0\\ 0 & -e^{2\pi i/5} \end{pmatrix} , \qquad (9.126)$$

while the generator σ_2 is represented by

$$\sigma_2 \mapsto B = F^{-1}RF$$
, $F = \begin{pmatrix} \tau & \sqrt{\tau} \\ \sqrt{\tau} & -\tau \end{pmatrix}$. (9.127)

These matrices generate a representation of the braid group B_3 on three strands whose image is dense in SU(2). Indeed, R and B generate Z_{10} subgroups of U(2), about two distinct axes, and there is no finite subgroup of U(2) that contains both of these subgroups — therefore, the representation closes on the group containing all elements of U(2) with determinant equal to a 10th root of unity. Similarly, for n anyons with trivial total charge, the image of the representation of the braid group is dense in SU(N_n^0).

To simulate a quantum circuit acting on n qubits, altogether 4n anyons are used. We have just seen that by braiding within each cluster of four anyons, arbitrary single-qubit gates can be realized. To complete a universal set, we will need two-qubit gates as well. But two neighboring qubits are encoded by eight anyons, and exchanges of these anyons generate a representation of B_8 whose image is dense in $SU(N_8^0) = SU(13)$, which of course includes the SU(4) that acts on the two encoded qubits. Therefore, each gate in a universal set can be simulated with arbitrary accuracy by some finite braid.

Since we can braid clockwise as well as counterclockwise, the inverse of each exchange gate is also in our repertoire. Therefore, we can apply the Solovay-Kitaev theorem to conclude that the universal gates of the circuit model can be simulated to accuracy ε with braids of length poly $(\log(1/\varepsilon))$. It follows that an ideal quantum circuit with L gates acting on all together n qubits can be simulated to fixed accuracy using 4n anyons and a braid of length $O(L \cdot \operatorname{poly}(\log(L)))$. As desired, we have shown that a universal quantum computer can be simulated efficiently with Fibonacci anyons. Note that, in contrast to the simulation using the nonabelian superconductor model, no intermediate measurements are needed to realize the universal gates.

In the analysis above, we have assumed that there are no errors in the simulation other than those limiting the accuracy of the Solovay-Kitaev approximation to the ideal gates. It is therefore implicit that the temperature is small enough compared to the energy gap of the model that thermally excited anyons are too rare to cause trouble, that the anyons are kept far enough apart from one another that uncontrolled exchange of charge can be neglected, and in general that errors in the topological quantum computation are unimportant. If the error rate is small but not completely negligible, then the standard theory of quantum fault tolerance can be invoked to boost the accuracy of the simulation as needed, at an additional overhead cost polylogarithmic in L. The faulttolerant procedure should include a method for controlling the "leakage" of the encoded qubits — that is, to prevent the drift of the clusters of four qubits from the two-dimensional computational space V_4^0 to its threedimensional orthogonal complement V_4^1 .

9.18 Epilogue

That is as far as I got in class. I will mention briefly here a few other topics that I might have covered if I had not run out of time.

9.18.1 Chern-Simons theory

We have discussed how anyon models can be constructed through a bruteforce solution to the polynomial equations. This method is foolproof, but in practice models are often constructed using other, more efficient methods. Indeed, most of the known anyon models have been found as instances of *Chern-Simons theory*.

The fusion rules of a Chern-Simons theory are a truncated version of the fusion rules for irreducible representations of a Lie group. For example, associated with the group SU(2) there is a tower of Chern-Simons theories indexed by a positive integer k. For SU(2), the irreducible representations carry labels $j = 0, \frac{1}{2}, 1, \frac{3}{2}, 2, \frac{5}{2}, \ldots$, and the fusion rules have the form

$$j_1 \times j_2 = \sum_{j=|j_2-j_1|}^{j_1+j_2} j$$
 (9.128)

In the Chern-Simons theory denoted $SU(2)_k$, the half-integer labels are limited to $j \leq k/2$, and the label j is contained in $j_1 \times j_2$ only if $j_1+j_2+j \leq k$.

For example, the $SU(2)_1$ model is abelian, and the nontrivial fusion rules of the $SU(2)_2$ model are

$$\frac{1}{2} \times \frac{1}{2} = 0 + 1 ,$$

$$\frac{1}{2} \times 1 = \frac{1}{2} ,$$

$$1 \times 1 = 0 .$$

$$(9.129)$$

Therefore, the label $\frac{1}{2}$ has quantum dimension $d_{1/2} = \sqrt{2}$, and the topological Hilbert space of 2m such anyons with total charge 0 has dimension

$$N^{0}_{\left(\frac{1}{2}\right)^{2m}} = 2^{m-1} . (9.130)$$

The polynomial equations for these fusion rules have multiple solutions (only one of which describes the braiding properties of the SU(2)₂ model), but none of the resulting models have computationally universal braiding rules. The space $V_{\frac{1}{2}\frac{1}{2}\frac{1}{2}\frac{1}{2}}^{0}$ is two-dimensional, and the 2 × 2 matrices $F \equiv F_{\frac{1}{2}\frac{1}{2}\frac{1}{2}\frac{1}{2}}$ and $R \equiv R_{\frac{1}{2}\frac{1}{2}}$ are, up to overall phases and complex conjugation,

$$F = H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1\\ 1 & -1 \end{pmatrix} , \quad R = P = \begin{pmatrix} 1 & 0\\ 0 & i \end{pmatrix} .$$
(9.131)

9 Topological quantum computation

There are Clifford-group quantum gates, inadequate for universality.

However, the $SU(2)_k$ models for $k \ge 3$ are computationally universal. The nontrivial fusion rules of $SU(2)_3$ are

$$\begin{array}{rcl} \frac{1}{2} \times \frac{1}{2} &= 0 + 1 \ , \\ \frac{1}{2} \times 1 &= \frac{1}{2} + \frac{3}{2} \ , \\ \frac{1}{2} \times \frac{3}{2} &= 1 \ , \\ 1 \times 1 &= 0 \ + 1 \ , \\ 1 \times \frac{3}{2} &= \frac{1}{2} \ , \\ \frac{3}{2} \times \frac{3}{2} &= 0 \ . \end{array}$$
(9.132)

The Fibonacci (Yang-Lee) model that we have studied is obtained by truncating $SU(2)_3$, further, eliminating the noninteger labels $\frac{1}{2}$ and $\frac{3}{2}$ (*i.e.*, this is the Chern-Simons theory $SO(3)_3$); then the only remaining nontrivial fusion rule is $1 \times 1 = 0 + 1$.

Wang (unpublished) has recently constructed all anyons models with no more than four labels, and has found that all of the models are closely related to the models that are found in Chern-Simons theory.

9.18.2 S-matrix

The *modular S-matrix* of an anyon model can be defined in terms of two anyon world lines that form a *Hopf link*:



Here \mathcal{D} is the total quantum dimension of the model, and we have used the normalization where unlinked loops would have the value $d_a d_b$; then the matrix S_a^b is symmetric and unitary. In abelian anyon models, the Hopf link arose in our discussion of topological degeneracy, where we characterized how the vacuum state of an anyon model on the torus is affected when an anyon is transported around one of the cycles of the torus. The *S*-matrix has a similar interpretation in the nonabelian case. By elementary reasoning, *S* can be related to the fusion rules:

$$(N_a)_b^c = \sum_d S_b^d \left(\frac{S_a^d}{S_1^d}\right) \left(S^{-1}\right)_d^c \;; \tag{9.133}$$

that is, the S-matrix simulaneously diagonalizes all the matrices $\{N_a\}$ (the Verlinde relation). Note that it follows from the definition that $S_1^a = d_a/\mathcal{D}$.

9.18.3 Edge excitations

In our formulation of anyon models, we have discussed the fusing and braiding of particles in the two-dimensional *bulk*. But there is another aspect of the physics of two-dimensional media that we have not yet discussed, the properties of the one-dimensional *edge* of the sample. Typically, if a two-dimensional system supports anyons in the bulk, there are also *chiral massless excitations* that propagate along the one-dimensional edge. At nonzero temperature T, there is an energy flux along the edge given by the expression

$$J = \frac{\pi}{12}c_{-}T^{2} ; \qquad (9.134)$$

here the constant c_{-} , called the *chiral central charge* of the edge, is a universal property that is unaffected by small changes in the underlying Hamiltonian of the system.

While this chiral central charge is an intrinsic property of the twodimensional medium, the properties of the anyons in the bulk do not determine it completely; rather we have

$$\frac{1}{\mathcal{D}}\sum_{a} d_a^2 e^{2\pi i J_a} = e^{(2\pi i/8)c_-} , \qquad (9.135)$$

where the sum is over the complete label set of the anyon model, and $e^{2\pi i J_a} = R^1_{a\bar{a}}$ is the topological spin of the label a. This expression relates the quantity c_- , characteristic of the edge theory, to the quantum dimensions and topological spins of the bulk theory, but determines c_- only modulo 8. Therefore, at least in principle, there can be multiple edge theories corresponding to a single theory of anyons in the bulk.

9.19 Bibliographical notes

Some of the pioneering papers on the theory of anyons are reprinted in [1].

What I have called the "nonabelian superconductor" model is often referred to in the literature as the "quantum double," and is studied using the representation theory of Hopf algebras. For a review see [2].

That nonabelian anyons can be used for fault-tolerant quantum computing was first suggested in [3]. This paper also discusses the toric code, and related lattice models that have nonabelian phases. A particular realization of universal quantum computation in a nonabelian superconductor was discussed in [4, 5]. My discussion of the universal gate set is based on [6], where more general models are also discussed. Other schemes, that make more extensive use of electric charges and that are universal for smaller groups (like S_3) are described in [7].

Diagrammatic methods, like those I used in the discussion of the quantum dimension, are extensively applied to derive properties of anyons in [8]. The role of the polynomial equations (pentagon and hexagon equations) in (1+1)-dimensional conformal field theory is discussed in [9].

Simulation of anyons using a quantum circuit is discussed in [10]. Simulation of a universal quantum computer using the anyons of the $SU(2)_{k=3}$ Chern-Simons theory is discussed in [11]. That the Yang-Lee model is also universal was pointed out in [12].

I did not discuss physical implementations in my lectures, but I list a few relevant references here anyway: Ideas about realizing abelian and nonabelian anyons using superconducting Josephson-junction arrays are discussed in [13]. A spin model with nearest-neighbor interactions that has nonabelian anyons (though not ones that are computationally universal) is proposed and solved in [14], and a proposal for realizing this model using cold atoms trapped in an optical lattice is described in [15]. Some ideas about realizing the (computationally universal) SU(2)_{k=3} model in a system of interacting electrons are discussed in [16].

Much of my understanding of the theory of computing with nonabelian anyons was derived from many helpful discussions with Alexei Kitaev.

References

- F. Wilczek, Fractional statistics and anyon superconductivity (World Scientific, Singapore, 1990).
- [2] M. de Wild Propitius and F. A. Bais, "Discrete gauge theories," arXiv: hep-th/9511201 (1995).
- [3] A. Yu. Kitaev, "Fault-tolerant quantum computation by anyons," Annals Phys. 303, 2-30 (2003), arXiv: quant-ph/9707021.
- [4] R. W. Ogburn and H. Preskill, "Topological quantum computation," Lect. Notes in Comp. Sci. 1509, 341-356 (1999).
- [5] J. Preskill, "Fault-tolerant quantum computation," arXiv: quantph/9712048 (1997).
- [6] C. Mochon, "Anyons from non-solvable finite groups are sufficient for universal quantum computation" Phys. Rev. A 67, 022315 (2003), quantph/0206128.
- [7] C. Mochon, "Anyon computers with smaller groups," Phys. Rev. A 69, 032306 (2004), arXiv: quant-ph/0306063.
- [8] J. Fröhlich and F. Gabbiani, "Braid statistics in local quantum theory," Rev. Math. Phys. 2:3, 251–353 (1990).
- G. Moore and N. Seiberg, "Classical and quantum conformal field theory," Comm. Math. Phys. 123, 171–254 (1989).
- [10] M. H. Freedman, A. Kitaev, and Z. Wang, "Simulation of topological field theories by quantum computers," Comm. Math. Phys. 227, 587-603 (2002), arXiv: quant-ph/0001071.
- [11] M. H. Freedman, M. Larsen, and Z. Wang, "A modular functor which is universal for quantum computation," arXiv: quant-ph/0001108 (2000).
- [12] G. Kuperberg, unpublished.
- [13] B. Doucot, L. B. Ioffe, and J. Vidal, "Discrete non-Abelian gauge theories in two-dimensional lattices and their realizations in Josephson-junction arrays," Phys. Rev. B 69, 214501 (2004), arXiv: cond-mat/0302104.

References

- [14] A. Kitaev, "Anyons in a spin model on the honeycomb lattice," unpublished.
- [15] L.-M. Duan, E. Demler, and M. D. Lukin, "Controlling spin exchange interactions of ultracold atoms in optical lattices," Phys. Rev. Lett. 91, 090402 (2003), arXiv: cond-mat/0210564.
- [16] M. Freedman, C. Nayak, K. Shtengel, K. Walker, and Zhenghan Wang, "A class of P, T-invariant topological phases of interacting electrons," condmat/0307511 (2003).

Quantum Information Chapter 10. Quantum Shannon Theory

John Preskill Institute for Quantum Information and Matter California Institute of Technology

Updated June 2016

For further updates and additional chapters, see: http://www.theory.caltech.edu/people/preskill/ph219/

Please send corrections to preskill@caltech.edu

Contents

			page v
	Prefe	ice	vi
10	Qua	ntum Shannon Theory	1
	10.1	Shannon for Dummies	1
		10.1.1 Shannon entropy and data compression	2
		10.1.2 Joint typicality, conditional entropy, and mutual information	. 4
		10.1.3 Distributed source coding	6
		10.1.4 The noisy channel coding theorem	7
	10.2	Von Neumann Entropy	12
		10.2.1 Mathematical properties of $H(\rho)$	14
		10.2.2 Mixing, measurement, and entropy	15
		10.2.3 Strong subadditivity	16
		10.2.4 Monotonicity of mutual information	18
		10.2.5 Entropy and thermodynamics	19
		10.2.6 Bekenstein's entropy bound.	20
		10.2.7 Entropic uncertainty relations	21
	10.3	Quantum Source Coding	23
		10.3.1 Quantum compression: an example	24
		10.3.2 Schumacher compression in general	27
	10.4	Entanglement Concentration and Dilution	30
	10.5	Quantifying Mixed-State Entanglement	35
		10.5.1 Asymptotic irreversibility under LOCC	35
		10.5.2 Squashed entanglement	37
		10.5.3 Entanglement monogamy	38
	10.6	Accessible Information	39
		10.6.1 How much can we learn from a measurement?	39
		10.6.2 Holevo bound	40
		10.6.3 Monotonicity of Holevo χ	41
		10.6.4 Improved distinguishability through coding: an example	42
		10.6.5 Classical capacity of a quantum channel	45
		10.6.6 Entanglement-breaking channels	48
	10.7	Quantum Channel Capacities and Decoupling	50
		10.7.1 Coherent information and the quantum channel capacity	50
		10.7.2 The decoupling principle	52
		10.7.3 Degradable channels	54

Contents

10.8 Quantum Protocols	56
10.8.1 Father: Entanglement-assisted quantum communication	56
10.8.2 Mother: Quantum state transfer	58
10.8.3 Operational meaning of strong subadditivity	61
10.8.4 Negative conditional entropy in thermodynamics	62
10.9 The Decoupling Inequality	64
10.9.1 Proof of the decoupling inequality	65
10.9.2 Proof of the mother inequality	66
10.9.3 Proof of the father inequality	68
10.9.4 Quantum channel capacity revisited	70
10.9.5 Black holes as mirrors	71
10.10 Summary	
10.11 Bibliographical Notes	
Exercises	
References	

iv

This article forms one chapter of *Quantum Information* which will be first published by Cambridge University Press.

© in the Work, John Preskill, 2016

NB: The copy of the Work, as displayed on this website, is a draft, pre-publication copy only. The final, published version of the Work can be purchased through Cambridge University Press and other standard distribution channels. This draft copy is made available for personal use only and must not be sold or re-distributed.

Preface

This is the 10th and final chapter of my book *Quantum Information*, based on the course I have been teaching at Caltech since 1997. An early version of this chapter (originally Chapter 5) has been available on the course website since 1998, but this version is substantially revised and expanded.

The level of detail is uneven, as I've aimed to provide a gentle introduction, but I've also tried to avoid statements that are incorrect or obscure. Generally speaking, I chose to include topics that are both useful to know and relatively easy to explain; I had to leave out a lot of good stuff, but on the other hand the chapter is already quite long.

My version of Quantum Shannon Theory is no substitute for the more careful treatment in Wilde's book [1], but it may be more suitable for beginners. This chapter contains occasional references to earlier chapters in my book, but I hope it will be intelligible when read independently of other chapters, including the chapter on quantum error-correcting codes.

This is a working draft of Chapter 10, which I will continue to update. See the URL on the title page for further updates and drafts of other chapters. Please send an email to preskill@caltech.edu if you notice errors.

Eventually, the complete book will be published by Cambridge University Press. I hesitate to predict the publication date — they have been far too patient with me.

Quantum Shannon Theory

Quantum information science is a synthesis of three great themes of 20th century thought: quantum physics, computer science, and information theory. Up until now, we have given short shrift to the information theory side of this trio, an oversight now to be remedied.

A suitable name for this chapter might have been *Quantum Information Theory*, but I prefer for that term to have a broader meaning, encompassing much that has already been presented in this book. Instead I call it *Quantum Shannon Theory*, to emphasize that we will mostly be occupied with generalizing and applying Claude Shannon's great (classical) contributions to a quantum setting. Quantum Shannon theory has several major thrusts:

- 1. Compressing quantum information.
- 2. Transmitting classical and quantum information through noisy quantum channels.
- 3. Quantifying, characterizing, transforming, and using quantum entanglement.

A recurring theme unites these topics — the properties, interpretation, and applications of Von Neumann entropy.

My goal is to introduce some of the main ideas and tools of quantum Shannon theory, but there is a lot we won't cover. For example, we will mostly consider information theory in an *asymptotic setting*, where the same quantum channel or state is used arbitrarily many times, thus focusing on issues of principle rather than more practical questions about devising efficient protocols.

10.1 Shannon for Dummies

Before we can understand Von Neumann entropy and its relevance to quantum information, we should discuss Shannon entropy and its relevance to classical information.

Claude Shannon established the two core results of classical information theory in his landmark 1948 paper. The two central problems that he solved were:

- 1. How much can a message be *compressed*; *i.e.*, how redundant is the information? This question is answered by the "source coding theorem," also called the "noiseless coding theorem."
- 2. At what *rate* can we communicate reliably over a noisy channel; *i.e.*, how much redundancy must be incorporated into a message to protect against errors? This question is answered by the "noisy channel coding theorem."

Quantum Shannon Theory

Both questions concern redundancy – how unexpected is the next letter of the message, on the average. One of Shannon's key insights was that entropy provides a suitable way to quantify redundancy.

I call this section "Shannon for Dummies" because I will try to explain Shannon's ideas quickly, minimizing distracting details. That way, I can compress classical information theory to about 14 pages.

10.1.1 Shannon entropy and data compression

A message is a string of letters, where each letter is chosen from an alphabet of k possible letters. We'll consider an idealized setting in which the message is produced by an "information source" which picks each letter by sampling from a probability distribution

$$X := \{x, p(x)\}; \tag{10.1}$$

that is, the letter has the value

$$x \in \{0, 1, 2, \dots k-1\} \tag{10.2}$$

with probability p(x). If the source emits an *n*-letter message the particular string $x = x_1x_2...x_n$ occurs with probability

$$p(x_1 x_2 \dots x_n) = \prod_{i=1}^n p(x_i).$$
 (10.3)

Since the letters are statistically independent, and each is produced by consulting the same probability distribution X, we say that the letters are *independent and identically distributed*, abbreviated *i.i.d.* We'll use X^n to denote the ensemble of *n*-letter messages in which each letter is generated independently by sampling from X, and $\vec{x} = (x_1 x_2 \dots x_n)$ to denote a string of bits.

Now consider long *n*-letter messages, $n \gg 1$. We ask: is it possible to compress the message to a shorter string of letters that conveys essentially the same information? The answer is: Yes, it's possible, unless the distribution X is uniformly random.

If the alphabet is binary, then each letter is either 0 with probability 1 - p or 1 with probability p, where $0 \le p \le 1$. For n very large, the law of large numbers tells us that typical strings will contain about n(1-p) 0's and about np 1's. The number of distinct strings of this form is of order the binomial coefficient $\binom{n}{np}$, and from the Stirling approximation $\log n! = n \log n - n + O(\log n)$ we obtain

$$\log \binom{n}{np} = \log \left(\frac{n!}{(np)! (n(1-p))!} \right)$$

$$\approx n \log n - n - (np \log np - np + n(1-p) \log n(1-p) - n(1-p))$$

$$= nH(p),$$
(10.4)

where

$$H(p) = -p\log p - (1-p)\log(1-p)$$
(10.5)

is the *entropy* function.

In this derivation we used the Stirling approximation in the appropriate form for natural logarithms. But from now on we will prefer to use logarithms with base 2, which is more convenient for expressing a quantity of information in bits; thus if no base is indicated, it will be understood that the base is 2 unless otherwise stated. Adopting this convention in the expression for H(p), the number of typical strings is of order $2^{nH(p)}$.

To convey essentially all the information carried by a string of n bits, it suffices to choose a block code that assigns a nonnegative integer to each of the typical strings. This block code needs to distinguish about $2^{nH(p)}$ messages (all occurring with nearly equal *a priori* probability), so we may specify any one of the messages using a binary string with length only slightly longer than nH(p). Since $0 \leq H(p) \leq 1$ for $0 \leq p \leq 1$, and H(p) = 1 only for $p = \frac{1}{2}$, the block code shortens the message for any $p \neq \frac{1}{2}$ (whenever 0 and 1 are not equally probable). This is Shannon's result. The key idea is that we do not need a codeword for every sequence of letters, only for the *typical* sequences. The probability that the actual message is atypical becomes negligible asymptotically, *i.e.*, in the limit $n \to \infty$.

Similar reasoning applies to the case where X samples from a k-letter alphabet. In a string of n letters, x typically occurs about np(x) times, and the number of typical strings is of order

$$\frac{n!}{\prod_{x} (np(x))!} \simeq 2^{-nH(X)},$$
(10.6)

where we have again invoked the Stirling approximation and now

$$H(X) = -\sum_{x} p(x) \log_2 p(x).$$
 (10.7)

is the Shannon entropy (or simply entropy) of the ensemble $X = \{x, p(x)\}$. Adopting a block code that assigns integers to the typical sequences, the information in a string of n letters can be compressed to about nH(X) bits. In this sense a letter x chosen from the ensemble carries, on the average, H(X) bits of information.

It is useful to restate this reasoning more carefully using the *strong law of large* numbers, which asserts that a sample average for a random variable almost certainly converges to its expected value in the limit of many trials. If we sample from the distribution $Y = \{y, p(y)\}$ n times, let $y_i, i \in \{1, 2, ..., n\}$ denote the *i*th sample, and let

$$\mu[Y] = \langle y \rangle = \sum_{y} y \ p(y) \tag{10.8}$$

denote the expected value of y. Then for any positive ε and δ there is a positive integer N such that

$$\left|\frac{1}{n}\sum_{i=1}^{n}y_{i}-\mu[Y]\right| \le \delta \tag{10.9}$$

with probability at least $1 - \varepsilon$ for all $n \ge N$. We can apply this statement to the random variable $\log_2 p(x)$. Let us say that a sequence of n letters is δ -typical if

$$H(X) - \delta \le -\frac{1}{n} \log_2 p(x_1 x_2 \dots x_n) \le H(X) + \delta;$$
 (10.10)

then the strong law of large numbers says that for any $\varepsilon, \delta > 0$ and *n* sufficiently large, an *n*-letter sequence will be δ -typical with probability $\geq 1 - \varepsilon$.

Since each δ -typical *n*-letter sequence \vec{x} occurs with probability $p(\vec{x})$ satisfying

$$p_{\min} = 2^{-n(H+\delta)} \le p(\vec{x}) \le 2^{-n(H-\delta)} = p_{\max},$$
 (10.11)

we may infer upper and lower bounds on the number $N_{typ}(\varepsilon, \delta, n)$ of typical sequences:

$$N_{\text{typ}} p_{\min} \le \sum_{\text{typical } x} p(x) \le 1, \quad N_{\text{typ}} p_{\max} \ge \sum_{\text{typical } x} p(x) \ge 1 - \varepsilon,$$
 (10.12)

implies

$$2^{n(H+\delta)} \ge N_{\text{typ}}(\varepsilon, \delta, n) \ge (1-\varepsilon)2^{n(H-\delta)}.$$
(10.13)

Therefore, we can encode all typical sequences using a block code with length $n(H + \delta)$ bits. That way, any message emitted by the source can be compressed and decoded successfully as long as the message is typical; the compression procedure achieves a success probability $p_{\text{success}} \geq 1 - \varepsilon$, no matter how the atypical sequences are decoded.

What if we try to compress the message even further, say to $H(X) - \delta'$ bits per letter, where δ' is a constant independent of the message length n? Then we'll run into trouble, because there won't be enough codewords to cover all the typical messages, and we won't be able to decode the compressed message with negligible probability of error. The probability p_{success} of successfully decoding the message will be bounded above by

$$p_{\text{success}} \le 2^{n(H-\delta')} 2^{-n(H-\delta)} + \varepsilon = 2^{-n(\delta'-\delta)} + \varepsilon; \qquad (10.14)$$

we can correctly decode only $2^{n(H-\delta')}$ typical messages, each occurring with probability no higher than $2^{-n(H-\delta)}$; we add ε , an upper bound on the probability of an atypical message, allowing optimistically for the possibility that we somehow manage to decode the atypical messages correctly. Since we may choose ε and δ as small as we please, this success probability becomes small as $n \to \infty$, if δ' is a positive constant.

The number of bits per letter encoding the compressed message is called the *rate* of the compression code, and we say a rate R is *achievable* asymptotically (as $n \to \infty$) if there is a sequence of codes with rate at least R and error probability approaching zero in the limit of large n. To summarize our conclusion, we have found that

Compression Rate =
$$H(X) + o(1)$$
 is achievable,
Compression Rate = $H(X) - \Omega(1)$ is not achievable, (10.15)

where o(1) denotes a positive quantity which may be chosen as small as we please, and $\Omega(1)$ denotes a positive constant. This is Shannon's source coding theorem.

We have not discussed at all the details of the compression code. We might imagine a huge lookup table which assigns a unique codeword to each message and vice versa, but because such a table has size exponential in n it is quite impractical for compressing and decompressing long messages. It is fascinating to study how to make the coding and decoding efficient while preserving a near optimal rate of compression, and quite important, too, if we really want to compress something. But this practical aspect of classical compression theory is beyond the scope of this book.

10.1.2 Joint typicality, conditional entropy, and mutual information

The Shannon entropy quantifies my *ignorance* per letter about the output of an information source. If the source X produces an n-letter message, then n(H(X) + o(1)) bits suffice to convey the content of the message, while $n(H(X) - \Omega(1))$ bits do not suffice.

Two information sources X and Y can be correlated. Letters drawn from the sources are governed by a joint distribution $XY = \{(x, y), p(x, y)\}$, in which a pair of letters (x, y) appears with probability p(x, y). The sources are independent if p(x, y) = p(x)p(y),

but correlated otherwise. If XY is a joint distribution, we use X to denote the marginal distribution, defined as

$$X = \left\{ x, p(x) = \sum_{y} p(x, y) \right\},$$
 (10.16)

and similarly for Y. If X and Y are correlated, then by reading a message generated by Y^n I reduce my ignorance about a message generated by X^n , which should make it possible to compress the output of X further than if I did not have access to Y.

To make this idea more precise, we use the concept of *jointly typical sequences*. Sampling from the distribution X^nY^n , that is, sampling n times from the joint distribution XY, produces a message $(\vec{x}, \vec{y}) = (x_1x_2 \dots x_n, y_1y_2 \dots y_n)$ with probability

$$p(\vec{x}, \vec{y}) = p(x_1, y_1) p(x_2, y_2) \dots p(x_n, y_n).$$
(10.17)

Let us say that (\vec{x}, \vec{y}) drawn from $X^n Y^n$ is jointly δ -typical if

$$2^{-n(H(X)+\delta)} \le p(\vec{x}) \le 2^{-n(H(X)-\delta)},$$

$$2^{-n(H(Y)+\delta)} \le p(\vec{y}) \le 2^{-n(H(Y)-\delta)},$$

$$2^{-n(H(XY)+\delta)} \le p(\vec{x}, \vec{y}) \le 2^{-n(H(XY)-\delta)}.$$
(10.18)

Then, applying the strong law of large numbers simultaneously to the three distributions X^n , Y^n , and X^nY^n , we infer that for $\varepsilon, \delta > 0$ and n sufficiently large, a sequence drawn from X^nY^n will be δ -typical with probability $\geq 1 - \varepsilon$. Using Bayes' rule, we can then obtain upper and lower bounds on the *conditional* probability $p(\vec{x}|\vec{y})$ for jointly typical sequences:

$$p(\vec{x}|\vec{y}) = \frac{p(\vec{x},\vec{y})}{p(\vec{y})} \ge \frac{2^{-n(H(XY)+\delta)}}{2^{-n(H(Y)-\delta)}} = 2^{-n(H(X|Y)+2\delta)},$$

$$p(\vec{x}|\vec{y}) = \frac{p(\vec{x},\vec{y})}{p(\vec{y})} \le \frac{2^{-n(H(XY)-\delta)}}{2^{-n(H(Y)+\delta)}} = 2^{-n(H(X|Y)-2\delta)}.$$
 (10.19)

Here we have introduced the quantity

$$H(X|Y) = H(XY) - H(Y) = \langle -\log p(x,y) + \log p(y) \rangle = \langle -\log p(x|y) \rangle, \quad (10.20)$$

which is called the *conditional entropy* of X given Y.

The conditional entropy quantifies my remaining ignorance about x once I know y. From eq.(10.19) we see that if (\vec{x}, \vec{y}) is jointly typical (as is the case with high probability for n large), then the number of possible values for \vec{x} compatible with the known value of \vec{y} is no more than $2^{n(H(X|Y)+2\delta)}$; hence we can convey \vec{x} with a high success probability using only H(X|Y) + o(1) bits per letter. On the other hand we can't do much better, because if we use only $2^{n(H(X|Y)-\delta')}$ codewords, we are limited to conveying reliably no more than a fraction $2^{-n(\delta'-2\delta)}$ of all the jointly typical messages. To summarize, H(X|Y) is the number of additional bits per letter needed to specify both \vec{x} and \vec{y} once \vec{y} is known. Similarly, H(Y|X) is the number of additional bits per letter needed to specify both \vec{x} and \vec{y} when \vec{x} is known.

The information about X that I gain when I learn Y is quantified by how much the

number of bits per letter needed to specify X is *reduced* when Y is known. Thus is

$$I(X;Y) \equiv H(X) - H(X|Y) = H(X) + H(Y) - H(XY) = H(Y) - H(Y|X),$$
(10.21)

which is called the *mutual information*. The mutual information I(X;Y) quantifies how X and Y are correlated, and is symmetric under interchange of X and Y: I find out as much about X by learning Y as about Y by learning X. Learning Y never *reduces* my knowledge of X, so I(X;Y) is obviously nonnegative, and indeed the inequality $H(X) \ge H(X|Y) \ge 0$ follows easily from the convexity of the log function.

Of course, if X and Y are completely uncorrelated, we have p(x, y) = p(x)p(y), and

$$I(X;Y) \equiv \left\langle \log \frac{p(x,y)}{p(x)p(y)} \right\rangle = 0; \qquad (10.22)$$

we don't find out anything about X by learning Y if there is no correlation between X and Y.

10.1.3 Distributed source coding

To sharpen our understanding of the operational meaning of conditional entropy, consider this situation: Suppose that the joint distribution XY is sampled n times, where Alice receives the n-letter message \vec{x} and Bob receives the n-letter message \vec{y} . Now Alice is to send a message to Bob which will enable Bob to determine \vec{x} with high success probability, and Alice wants to send as few bits to Bob as possible. This task is harder than in the scenario considered in §10.1.2, where we assumed that the encoder and the decoder share full knowledge of \vec{y} , and can choose their code for compressing \vec{x} accordingly. It turns out, though, that even in this more challenging setting Alice can compress the message she sends to Bob down to n(H(X|Y) + o(1)) bits, using a method called *Slepian-Wolf coding*.

Before receiving (\vec{x}, \vec{y}) , Alice and Bob agree to sort all the possible *n*-letter messages that Alice might receive into 2^{nR} possible bins of equal size, where the choice of bins is known to both Alice and Bob. When Alice receives \vec{x} , she sends nR bits to Bob, identifying the bin that contains \vec{x} . After Bob receives this message, he knows both \vec{y} and the bin containing \vec{x} . If there is a unique message in that bin which is jointly typical with \vec{y} , Bob decodes accordingly. Otherwise, he decodes arbitrarily. This procedure can fail either because \vec{x} and \vec{y} are not jointly typical, or because there is more than one message in the bin which is jointly typical with \vec{x} . Otherwise, Bob is sure to decode correctly.

Since \vec{x} and \vec{y} are jointly typical with high probability, the compression scheme works if it is unlikely for a bin to contain an incorrect message which is jointly typical with \vec{y} . If \vec{y} is typical, what can we say about the number $N_{\text{typ}|\vec{y}}$ of messages \vec{x} that are jointly typical with \vec{y} ? Using eq.(10.19), we have

$$1 \ge \sum_{\text{typical } \vec{x} | \vec{y}} p(\vec{x} | \vec{y}) \ge N_{\text{typ} | \vec{y}} \ 2^{-n(H(X|Y) + 2\delta)}, \tag{10.23}$$

and thus

$$N_{\text{typ}|\vec{y}} \le 2^{n(H(X|Y)+2\delta)}.$$
 (10.24)

Now, to estimate the probability of a decoding error, we need to specify how the bins are chosen. Let's assume the bins are chosen uniformly at random, or equivalently, let's consider averaging uniformly over all codes that divide the length-*n* strings into 2^{nR} bins of equal size. Then the probability that a particular bin contains a message jointly typical with a specified \vec{y} purely by accident is bounded above by

$$2^{-nR} N_{\text{typ}|\vec{y}} \ge 2^{-n(R-H(X|Y)-2\delta)}.$$
(10.25)

We conclude that if Alice sends R bits to Bob per each letter of the message x, where

$$R = H(X|Y) + o(1), (10.26)$$

then the probability of a decoding error vanishes in the limit $n \to \infty$, at least when we average over uniformly all codes. Surely, then, there must exist a particular sequence of codes Alice and Bob can use to achieve the rate R = H(X|Y) + o(1), as we wanted to show.

In this scenario, Alice and Bob jointly know (x, y), but initially neither Alice nor Bob has access to all their shared information. The goal is to merge all the information on Bob's side with minimal communication from Alice to Bob, and we have found that H(X|Y) + o(1) bits of communication per letter suffice for this purpose. Similarly, the information can be merged on Alice's side using H(Y|X) + o(1) bits of communication per letter from Bob to Alice.

10.1.4 The noisy channel coding theorem

Suppose Alice wants to send a message to Bob, but the communication channel linking Alice and Bob is noisy. Each time they use the channel, Bob receives the letter y with probability p(y|x) if Alice sends the letter x. Using the channel $n \gg 1$ times, Alice hopes to transmit a long message to Bob.

Alice and Bob realize that to communicate reliably despite the noise they should use some kind of code. For example, Alice might try sending the same bit k times, with Bob using a majority vote of the k noisy bits he receives to decode what Alice sent. One wonders: for a given channel, is it possible to ensure perfect transmission asymptotically, *i.e.*, in the limit where the number of channel uses $n \to \infty$? And what can be said about the *rate* of the code; that is, how many bits must be sent per letter of the transmitted message?

Shannon answered these questions. He showed that *any* channel can be used for perfectly reliable communication at an asymptotic nonzero rate, as long as there is *some* correlation between the channel's input and its output. Furthermore, he found a useful formula for the optimal rate that can be achieved. These results are the content of the *noisy channel coding theorem*.

Capacity of the binary symmetric channel.

To be concrete, suppose we use the binary alphabet $\{0, 1\}$, and the *binary symmetric* channel; this channel acts on each bit independently, flipping its value with probability p, and leaving it intact with probability 1 - p. Thus the conditional probabilities characterizing the channel are

$$p(0|0) = 1 - p, \quad p(0|1) = p, p(1|0) = p, \qquad p(1|1) = 1 - p.$$
(10.27)

Quantum Shannon Theory

We want to construct a family of codes with increasing block size n, such that the probability of a decoding error goes to zero as $n \to \infty$. For each n, the code contains 2^k codewords among the 2^n possible strings of length n. The rate R of the code, the number of encoded data bits transmitted per physical bit carried by the channel, is

$$R = \frac{k}{n}.\tag{10.28}$$

To protect against errors, we should choose the code so that the codewords are as "far apart" as possible. For given values of n and k, we want to maximize the number of bits that must be flipped to change one codeword to another, the Hamming distance between the two codewords. For any n-bit input message, we expect about np of the bits to flip — the input diffuses into one of about $2^{nH(p)}$ typical output strings, occupying an "error sphere" of "Hamming radius" np about the input string. To decode reliably, we want to choose our input codewords so that the error spheres of two different codewords do not overlap substantially. Otherwise, two different inputs will sometimes yield the same output, and decoding errors will inevitably occur. To avoid such decoding ambiguities, the total number of strings contained in all $2^k = 2^{nR}$ error spheres should not exceed the total number 2^n of bits in the output message; we therefore require

$$2^{nH(p)}2^{nR} \le 2^n \tag{10.29}$$

or

$$R \le 1 - H(p) := C(p). \tag{10.30}$$

If transmission is highly reliable, we cannot expect the rate of the code to exceed C(p). But is the rate R = C(p) actually *achievable* asymptotically?

In fact transmission with R = C - o(1) and negligible decoding error probability is possible. Perhaps Shannon's most ingenious idea was that this rate can be achieved by an average over "random codes." Though choosing a code at random does not seem like a clever strategy, rather surprisingly it turns out that random coding achieves as high a rate as any other coding scheme in the limit $n \to \infty$. Since C is the optimal rate for reliable transmission of data over the noisy channel it is called the *channel capacity*.

Suppose that X is the uniformly random ensemble for a single bit (either 0 with $p = \frac{1}{2}$) or 1 with $p = \frac{1}{2}$), and that we sample from X^n a total of 2^{nR} times to generate 2^{nR} "random codewords." The resulting code is known by both Alice and Bob. To send nR bits of information, Alice chooses one of the codewords and sends it to Bob by using the channel n times. To decode the n-bit message he receives, Bob draws a "Hamming sphere" with "radius" slightly large than np, containing

$$2^{n(H(p)+\delta)}$$
 (10.31)

strings. If this sphere contains a unique codeword, Bob decodes the message accordingly. If the sphere contains more than one codeword, or no codewords, Bob decodes arbitrarily.

How likely is a decoding error? For any positive δ , Bob's decoding sphere is large enough that it is very likely to contain the codeword sent by Alice when n is sufficiently large. Therefore, we need only worry that the sphere might contain another codeword just by accident. Since there are altogether 2^n possible strings, Bob's sphere contains a fraction

$$f = \frac{2^{n(H(p)+\delta)}}{2^n} = 2^{-n(C(p)-\delta)},$$
(10.32)

of all the strings. Because the codewords are uniformly random, the probability that Bob's sphere contains any particular codeword aside from the one sent by Alice is f, and the probability that the sphere contains any one of the $2^{nR} - 1$ invalid codewords is no more than

$$2^{nR}f = 2^{-n(C(p)-R-\delta)}.$$
(10.33)

Since δ may be as small as we please, we may choose R = C(p) - c where c is any positive constant, and the decoding error probability will approach zero as $n \to \infty$.

When we speak of codes chosen at random, we really mean that we are averaging over many possible codes. The argument so far has shown that the *average* probability of error is small, where we average over the choice of random code, and for each specified code we also average over all codewords. It follows that there must be a particular sequence of codes such that the average probability of error (when we average over the codewords) vanishes in the limit $n \to \infty$. We would like a stronger result – that the probability of error is small for *every* codeword.

To establish the stronger result, let p_i denote the probability of a decoding error when codeword *i* is sent. For any positive ε and sufficiently large *n*, we have demonstrated the existence of a code such that

$$\frac{1}{2^{nR}} \sum_{i=1}^{2^{nR}} p_i \le \varepsilon. \tag{10.34}$$

Let $N_{2\varepsilon}$ denote the number of codewords with $p_i \geq 2\varepsilon$. Then we infer that

$$\frac{1}{2^{nR}}(N_{2\varepsilon})2\varepsilon \le \varepsilon \text{ or } N_{2\varepsilon} \le 2^{nR-1};$$
(10.35)

we see that we can throw away at most half of the codewords, to achieve $p_i \leq 2\varepsilon$ for *every* codeword. The new code we have constructed has

$$Rate = R - \frac{1}{n}, \tag{10.36}$$

which approaches R as $n \to \infty$. We have seen, then, that the rate R = C(p) - o(1) is asymptotically achievable with negligible probability of error, where C(p) = 1 - H(p).

Mutual information as an achievable rate.

Now consider how to apply this random coding argument to more general alphabets and channels. The channel is characterized by p(y|x), the conditional probability that the letter y is received when the letter x is sent. We fix an ensemble $X = \{x, p(x)\}$ for the input letters, and generate the codewords for a length-n code with rate R by sampling 2^{nR} times from the distribution X^n ; the code is known by both the sender Alice and the receiver Bob. To convey an encoded nR-bit message, one of the 2^{nR} n-letter codewords is selected and sent by using the channel n times. The channel acts independently on the n letters, governed by the same conditional probability distribution p(y|x) each time it is used. The input ensemble X, together with the conditional probability characterizing the channel, determines the joint ensemble XY for each letter sent, and therefore the joint ensemble X^nY^n for the n uses of the channel.

To define a decoding procedure, we use the notion of joint typicality introduced in §10.1.2. When Bob receives the *n*-letter output message \vec{y} , he determines whether there is an *n*-letter input codeword \vec{x} jointly typical with \vec{y} . If such \vec{x} exists and is unique,

Bob decodes accordingly. If there is no \vec{x} jointly typical with \vec{y} , or more than one such \vec{x} , Bob decodes arbitrarily.

How likely is a decoding error? For any positive ε and δ , the (\vec{x}, \vec{y}) drawn from $X^n Y^n$ is jointly δ -typical with probability at least $1 - \varepsilon$ if n is sufficiently large. Therefore, we need only worry that there might more than one codeword jointly typical with \vec{y} .

Suppose that Alice samples X^n to generate a codeword \vec{x} , which she sends to Bob using the channel *n* times. Then Alice samples X^n a second time, producing another codeword \vec{x}' . With probability close to one, both \vec{y} and \vec{x}' are δ -typical. But what is the probability that \vec{x}' is *jointly* δ -typical with \vec{y} ?

Because the samples are independent, the probability of drawing these two codewords factorizes as $p(\vec{x}', \vec{x}) = p(\vec{x}')p(\vec{x})$, and likewise the channel output \vec{y} when the first codeword is sent is independent of the second channel input \vec{x}' , so $p(\vec{x}', \vec{y}) = p(\vec{x}')p(\vec{y})$. From eq.(10.18) we obtain an upper bound on the number $N_{i,t}$ of jointly δ -typical (\vec{x}, \vec{y}) :

$$1 \ge \sum_{j.t. \ (\vec{x}, \vec{y})} p(\vec{x}, \vec{y}) \ge N_{j.t.} \ 2^{-n(H(XY) + \delta)} \implies N_{j.t.} \le 2^{n(H(XY) + \delta)}.$$
(10.37)

We also know that each δ -typical \vec{x}' occurs with probability $p(\vec{x}') \leq 2^{-n(H(X)-\delta)}$ and that each δ -typical \vec{y} occurs with probability $p(\vec{y}) \leq 2^{-n(H(Y)-\delta)}$. Therefore, the probability that \vec{x}' and \vec{y} are jointly δ -typical is bounded above by

$$\sum_{j.t. \ (\vec{x}', \vec{y})} p(\vec{x}') p(\vec{y}) \leq N_{j.t.} \ 2^{-n(H(X)-\delta)} 2^{-n(H(Y)-\delta)}$$
$$\leq 2^{n(H(XY)+\delta)} 2^{-n(H(X)-\delta)} 2^{-n(H(Y)-\delta)}$$
$$= 2^{-n(I(X;Y)-3\delta)}.$$
(10.38)

If there are 2^{nR} codewords, all generated independently by sampling X^n , then the probability that *any* other codeword besides \vec{x} is jointly typical with \vec{y} is bounded above by

$$2^{nR}2^{-n(I(X;Y)-3\delta)} = 2^{n(R-I(X;Y)+3\delta)}.$$
(10.39)

Since ε and δ are as small as we please, we may choose R = I(X;Y) - c, where c is any positive constant, and the decoding error probability will approach zero as $n \to \infty$.

So far we have shown that the error probability is small when we average over codes and over codewords. To complete the argument we use the same reasoning as in our discussion of the capacity of the binary symmetric channel. There must exist a particular sequence of code with zero error probability in the limit $n \to \infty$, when we average over codewords. And by pruning the codewords, reducing the rate by a negligible amount, we can ensure that the error probability is small for *every* codeword. We conclude that the rate

$$R = I(X;Y) - o(1) \tag{10.40}$$

is asymptotically achievable with negligible probability of error. This result provides a concrete operational interpretation for the mutual information I(X; Y); it is the information per letter we can transmit over the channel, supporting the heuristic claim that I(X; Y) quantifies the information we gain about X when we have access to Y.

The mutual information I(X;Y) depends not only on the channel's conditional probability p(y|x) but also on the *a priori* probability p(x) defining the codeword ensemble X. The achievability argument for random coding applies for any choice of X, so we
have demonstrated that errorless transmission over the noisy channel is possible for any rate R strictly less than

$$C := \max_{X} I(X;Y). \tag{10.41}$$

This quantity C is called the *channel capacity*; it depends only on the conditional probabilities p(y|x) that define the channel.

Upper bound on the capacity.

We have now shown that any rate R < C is achievable, but can R exceed C with the error probability still approaching 0 for large n? To see that a rate for errorless transmission exceeding C is not possible, we reason as follows.

Consider any code with 2^{nR} codewords, and consider the uniform ensemble on the codewords, denoted \tilde{X}^n , in which each codeword occurs with probability 2^{-nR} . Evidently, then,

$$H(\tilde{X}^n) = nR. \tag{10.42}$$

Sending the codewords through n uses of the channel we obtain an ensemble \tilde{Y}^n of output states, and a joint ensemble $\tilde{X}^n \tilde{Y}^n$.

Because the channel acts on each letter independently, the conditional probability for n uses of the channel factorizes:

$$p(y_1y_2\cdots y_n|x_1x_2\cdots x_n) = p(y_1|x_1)p(y_2|x_2)\cdots p(y_n|x_n),$$
(10.43)

and it follows that the conditional entropy satisfies

$$H(\tilde{Y}^{n}|\tilde{X}^{n}) = \langle -\log p(\vec{y}|\vec{x}) \rangle = \sum_{i} \langle -\log p(y_{i}|x_{i}) \rangle$$
$$= \sum_{i} H(\tilde{Y}_{i}|\tilde{X}_{i}), \qquad (10.44)$$

where \tilde{X}_i and \tilde{Y}_i are the marginal probability distributions for the *i*th letter determined by our distribution on the codewords. Because Shannon entropy is subadditive, $H(XY) \leq H(X) + H(Y)$, we have

$$H(\tilde{Y}^n) \le \sum_i H(\tilde{Y}_i), \tag{10.45}$$

and therefore

$$I(\tilde{Y}^{n}; \tilde{X}^{n}) = H(\tilde{Y}^{n}) - H(\tilde{Y}^{n} | \tilde{X}^{n})$$

$$\leq \sum_{i} (H(\tilde{Y}_{i}) - H(\tilde{Y}_{i} | \tilde{X}_{i}))$$

$$= \sum_{i} I(\tilde{Y}_{i}; \tilde{X}_{i}) \leq nC.$$
(10.46)

The mutual information of the messages sent and received is bounded above by the sum of the mutual information per letter, and the mutual information for each letter is bounded above by the capacity, because C is defined as the maximum of I(X;Y) over all input ensembles.

Recalling the symmetry of mutual information, we have

$$I(\tilde{X}^n; \tilde{Y}^n) = H(\tilde{X}^n) - H(\tilde{X}^n | \tilde{Y}^n)$$

= $nR - H(\tilde{X}^n | \tilde{Y}^n) \le nC.$ (10.47)

Now, if we can decode reliably as $n \to \infty$, this means that the input codeword is completely determined by the signal received, or that the conditional entropy of the input (per letter) must get small

$$\frac{1}{n}H(\tilde{X}^n|\tilde{Y}^n) \to 0.$$
(10.48)

If errorless transmission is possible, then, eq. (10.47) becomes

$$R \le C + o(1),$$
 (10.49)

in the limit $n \to \infty$. The asymptotic rate cannot exceed the capacity. In Exercise 10.9, you will sharpen the statement eq.(10.48), showing that

$$\frac{1}{n}H(\tilde{X}^n|\tilde{Y}^n) \le \frac{1}{n}H_2(p_e) + p_e R,$$
(10.50)

where p_e denotes the decoding error probability, and $H_2(p_e) = -p_e \log_2 p_e - (1 - p_e) \log_2(1 - p_e)$.

We have now seen that the capacity C is the highest achievable rate of communication through the noisy channel, where the probability of error goes to zero as the number of letters in the message goes to infinity. This is Shannon's noisy channel coding theorem. What is particularly remarkable is that, although the capacity is achieved by messages that are many letters in length, we have obtained a *single-letter formula* for the capacity, expressed in terms of the optimal mutual information I(X;Y) for just a single use of the channel.

The method we used to show that R = C - o(1) is achievable, averaging over random codes, is not constructive. Since a random code has no structure or pattern, encoding and decoding are unwieldy, requiring an exponentially large code book. Nevertheless, the theorem is important and useful, because it tells us what is achievable, and not achievable, in principle. Furthermore, since I(X;Y) is a concave function of $X = \{x, p(x)\}$ (with $\{p(y|x)\}$ fixed), it has a unique local maximum, and C can often be computed (at least numerically) for channels of interest. Finding codes which can be efficiently encoded and decoded, and come close to achieving the capacity, is a very interesting pursuit, but beyond the scope of our lightning introduction to Shannon theory.

10.2 Von Neumann Entropy

In classical information theory, we often consider a source that prepares messages of n letters $(n \gg 1)$, where each letter is drawn independently from an ensemble $X = \{x, p(x)\}$. We have seen that the Shannon entropy H(X) is the number of incompressible bits of information carried per letter (asymptotically as $n \to \infty$).

We may also be interested in correlations among messages. The correlations between two ensembles of letters X and Y are characterized by conditional probabilities p(y|x). We have seen that the mutual information

$$I(X;Y) = H(X) - H(X|Y) = H(Y) - H(Y|X),$$
(10.51)

is the number of bits of information per letter about X that we can acquire by reading Y (or vice versa). If the p(y|x)'s characterize a noisy channel, then, I(X;Y) is the amount of information per letter that can be transmitted through the channel (given the *a priori* distribution X for the channel inputs).

We would like to generalize these considerations to quantum information. We may

imagine a source that prepares messages of n letters, but where each letter is chosen from an ensemble of quantum states. The signal alphabet consists of a set of quantum states $\{\rho(x)\}$, each occurring with a specified *a priori* probability p(x).

As we discussed at length in Chapter 2, the probability of any outcome of any measurement of a letter chosen from this ensemble, if the observer has no knowledge about which letter was prepared, can be completely characterized by the density operator

$$\boldsymbol{\rho} = \sum_{x} p(x)\boldsymbol{\rho}(x); \qquad (10.52)$$

for a POVM $\boldsymbol{E} = \{\boldsymbol{E}_a\}$, the probability of outcome *a* is

$$\operatorname{Prob}(a) = \operatorname{tr}(\boldsymbol{E}_a \boldsymbol{\rho}). \tag{10.53}$$

For this (or any) density operator, we may define the Von Neumann entropy

$$H(\boldsymbol{\rho}) = -\mathrm{tr}(\boldsymbol{\rho}\log\boldsymbol{\rho}). \tag{10.54}$$

Of course, we may choose an orthonormal basis $\{|a\rangle\}$ that diagonalizes ρ ,

$$\boldsymbol{\rho} = \sum_{a} \lambda_a |a\rangle \langle a|; \tag{10.55}$$

the vector of eigenvalues $\lambda(\boldsymbol{\rho})$ is a probability distribution, and the Von Neumann entropy of $\boldsymbol{\rho}$ is just the Shannon entropy of this distribution,

$$H(\boldsymbol{\rho}) = H(\lambda(\boldsymbol{\rho})). \tag{10.56}$$

If ρ_A is the density operator of system A, we will sometimes use the notation

$$H(A) := H(\boldsymbol{\rho}_A). \tag{10.57}$$

Our convention is to denote quantum systems with A, B, C, \ldots and classical probability distributions with X, Y, Z, \ldots

In the case where the signal alphabet $\{|\varphi(x)\rangle, p(x)\}$ consists of mutually orthogonal pure states, the quantum source reduces to a classical one; all of the signal states can be perfectly distinguished, and $H(\rho) = H(X)$, where X is the classical ensemble $\{x, p(x)\}$. The quantum source is more interesting when the signal states $\{\rho(x)\}$ are not mutually commuting. We will argue that the Von Neumann entropy quantifies the incompressible information content of the quantum source (in the case where the signal states are pure) much as the Shannon entropy quantifies the information content of a classical source.

Indeed, we will find that Von Neumann entropy plays multiple roles. It quantifies not only the *quantum* information content per letter of the pure-state ensemble (the minimum number of qubits per letter needed to reliably encode the information) but also its *classical* information content (the maximum amount of information per letter—in bits, not qubits—that we can gain about the preparation by making the best possible measurement). And we will see that Von Neumann information enters quantum information in yet other ways — for example, quantifying the entanglement of a bipartite pure state. Thus quantum information theory is largely concerned with the interpretation and uses of Von Neumann entropy, much as classical information theory is largely concerned with the interpretation and uses of Shannon entropy.

In fact, the mathematical machinery we need to develop quantum information theory is very similar to Shannon's mathematics (typical sequences, random coding, \ldots); so similar as to sometimes obscure that the conceptual context is really quite different.

The central issue in quantum information theory is that nonorthogonal quantum states cannot be perfectly distinguished, a feature with no classical analog.

10.2.1 Mathematical properties of $H(\rho)$

There are a handful of properties of the Von Neumann entropy $H(\rho)$ which are frequently useful, many of which are closely analogous to corresponding properties of the Shannon entropy H(X). Proofs of some of these are Exercises 10.1, 10.2, 10.3.

- 1. **Pure states**. A pure state $\rho = |\varphi\rangle\langle\varphi|$ has $H(\rho) = 0$.
- 2. Unitary invariance. The entropy is unchanged by a unitary change of basis,

$$H(\boldsymbol{U}\boldsymbol{\rho}\boldsymbol{U}^{-1}) = H(\boldsymbol{\rho}), \qquad (10.58)$$

because $H(\rho)$ depends only on the eigenvalues of ρ .

3. Maximum. If ρ has d nonvanishing eigenvalues, then

$$H(\boldsymbol{\rho}) \le \log d,\tag{10.59}$$

with equality when all the nonzero eigenvalues are equal. The entropy is maximized when the quantum state is maximally mixed.

4. Concavity. For $\lambda_1, \lambda_2, \dots, \lambda_n \ge 0$ and $\lambda_1 + \lambda_2 + \dots + \lambda_n = 1$,

$$H(\lambda_1 \boldsymbol{\rho}_1 + \dots + \lambda_n \boldsymbol{\rho}_n) \ge \lambda_1 H(\boldsymbol{\rho}_1) + \dots + \lambda_n H(\boldsymbol{\rho}_n).$$
(10.60)

The Von Neumann entropy is larger if we are *more ignorant* about how the state was prepared. This property is a consequence of the convexity of the log function.

5. Subadditivity. Consider a bipartite system AB in the state ρ_{AB} . Then

$$H(AB) \le H(A) + H(B) \tag{10.61}$$

(where $\rho_A = \operatorname{tr}_B(\rho_{AB})$ and $\rho_B = \operatorname{tr}_A(\rho_{AB})$), with equality for $\rho_{AB} = \rho_A \otimes \rho_B$. Thus, entropy is *additive* for uncorrelated systems, but otherwise the entropy of the whole is less than the sum of the entropy of the parts. This property is the quantum generalization of subadditivity of Shannon entropy:

$$H(XY) \le H(X) + H(Y).$$
 (10.62)

6. Bipartite pure states. If the state ρ_{AB} of the bipartite system AB is pure, then

$$H(A) = H(B),$$
 (10.63)

because ρ_A and ρ_B have the same nonzero eigenvalues.

7. Quantum mutual information. As in the classical case, we define the mutual information of two quantum systems as

$$I(A; B) = H(A) + H(B) - H(AB),$$
(10.64)

which is nonnegative because of the subadditivity of Von Neumann entropy, and zero only for a product state $\rho_{AB} = \rho_A \otimes \rho_B$.

8. Triangle inequality (Araki-Lieb inequality). For a bipartite system,

$$H(AB) \ge |H(A) - H(B)|.$$
 (10.65)

To derive the triangle inequality, consider the tripartite pure state $|\psi\rangle_{ABC}$ which purifies $\rho_{AB} = \operatorname{tr}_C(|\psi\rangle\langle\psi|)$. Since $|\psi\rangle$ is pure, H(A) = H(BC) and H(C) = H(AB); applying subadditivity to BC yields $H(A) \leq H(B) + H(C) = H(B) + H(AB)$. The same inequality applies with A and B interchanged, from which we obtain eq.(10.65).

The triangle inequality contrasts sharply with the analogous property of Shannon entropy,

$$H(XY) \ge H(X), H(Y). \tag{10.66}$$

The Shannon entropy of just part of a classical bipartite system cannot be greater than the Shannon entropy of the whole system. Not so for the Von Neumann entropy! For example, in the case of an entangled bipartite pure quantum state, we have H(A) = H(B) > 0, while H(AB) = 0. The entropy of the global system vanishes because our ignorance is minimal — we know as much about AB as the laws of quantum physics will allow. But we have incomplete knowledge of the parts A and B, with our ignorance quantified by H(A) = H(B). For a quantum system, but not for a classical one, information can be encoded in the correlations among the parts of the system, yet be invisible when we look at the parts one at a time.

Equivalently, a property that holds classically but not quantumly is

$$H(X|Y) = H(XY) - H(Y) \ge 0.$$
(10.67)

The Shannon conditional entropy H(X|Y) quantifies our remaining ignorance about X when we know Y, and equals zero when knowing Y makes us certain about X. On the other hand, the Von Neumann conditional entropy,

$$H(A|B) = H(AB) - H(B),$$
 (10.68)

can be negative; in particular we have H(A|B) = -H(A) = -H(B) < 0 if ρ_{AB} is an entangled pure state. How can it make sense that "knowing" the subsystem *B* makes us "more than certain" about the subsystem *A*? We'll return to this intriguing question in §10.8.2.

When X and Y are perfectly correlated, then H(XY) = H(X) = H(Y); the conditional entropy is H(X|Y) = H(Y|X) = 0 and the mutual information is I(X;Y) = H(X). In contrast, for a bipartite pure state of AB, the quantum state for which we may regard A and B as perfectly correlated, the mutual information is I(A;B) = 2H(A) = 2H(B). In this sense the quantum correlations are stronger than classical correlations.

10.2.2 Mixing, measurement, and entropy

The Shannon entropy also has a property called *Schur concavity*, which means that if $X = \{x, p(x)\}$ and $Y = \{y, q(y)\}$ are two ensembles such that $p \prec q$, then $H(X) \ge H(Y)$. In fact, any function on probability vectors is Schur concave if it is invariant under permutations of its arguments and also concave in each argument. Recall that $p \prec q$ (q majorizes p) means that "p is at least as random as q" in the sense that p = Dq for some doubly stochastic matrix D. Thus Schur concavity of H says that an ensemble with more randomness has higher entropy.

The Von Neumann entropy $H(\rho)$ of a density operator is the Shannon entropy of its vector of eigenvalues $\lambda(\rho)$. Furthermore, we showed in Exercise 2.6 that if the quantum state ensemble $\{|\varphi(x)\rangle, p(x)\}$ realizes ρ , then $p \prec \lambda(\rho)$; therefore $H(\rho) \leq H(X)$, where equality holds only for an ensemble of mutually orthogonal states. The decrease in

entropy $H(X)-H(\rho)$ quantifies how *distinguishability is lost* when we mix nonorthogonal pure states. As we will soon see, the amount of information we can gain by measuring ρ is no more than $H(\rho)$ bits, so some of the information about which state was prepared has been irretrievably lost if $H(\rho) < H(X)$.

If we perform an orthogonal measurement on ρ by projecting onto the basis $\{|y\rangle\}$, then outcome y occurs with probability

$$q(y) = \langle y | \boldsymbol{\rho} | y \rangle = \sum_{a} |\langle y | a \rangle|^2 \lambda_a, \quad \text{where} \quad \boldsymbol{\rho} = \sum_{a} \lambda_a |a\rangle \langle a| \tag{10.69}$$

and $\{|a\rangle\}$ is the basis in which ρ is diagonal. Since $D_{ya} = |\langle y|a\rangle|^2$ is a doubly stochastic matrix, $q \prec \lambda(\rho)$ and therefore $H(Y) \ge H(\rho)$, where equality holds only if the measurement is in the basis $\{|a\rangle\}$. Mathematically, the conclusion is that for a nondiagonal and nonnegative Hermitian matrix, the diagonal elements are more random than the eigenvalues. Speaking more physically, the outcome of an orthogonal measurement is easiest to predict if we measure an observable which commutes with the density operator, and becomes less predictable if we measure in a different basis.

This majorization property has a further consequence, which will be useful for our discussion of quantum compression. Suppose that ρ is a density operator of a *d*-dimensional system, with eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$ and that $\mathbf{E}' = \sum_{i=1}^{d'} |e_i\rangle\langle e_i|$ is a projector onto a subspace Λ of dimension $d' \leq d$ with orthonormal basis $\{|e_i\rangle\}$. Then

$$\operatorname{tr}\left(\boldsymbol{\rho}\boldsymbol{E}'\right) = \sum_{i=1}^{d'} \langle e_i | \boldsymbol{\rho} | e_i \rangle \leq \sum_{i=1}^{d'} \lambda_i, \qquad (10.70)$$

where the inequality follows because the diagonal elements of ρ in the basis $\{|e_i\rangle\}$ are majorized by the eigenvalues of ρ . In other words, if we perform a two-outcome orthogonal measurement, projecting onto either Λ or its orthogonal complement Λ^{\perp} , the probability of projecting onto Λ is no larger than the sum of the d' largest eigenvalues of ρ (the Ky Fan dominance principle).

10.2.3 Strong subadditivity

In addition to the subadditivity property $I(X;Y) \ge 0$, correlations of classical random variables obey a further property called *strong subadditivity*:

$$I(X;YZ) \ge I(X;Y). \tag{10.71}$$

This is the eminently reasonable statement that the correlations of X with YZ are at least as strong as the correlations of X with Y alone.

There is another useful way to think about (classical) strong subadditivity. Recalling the definition of mutual information we have

$$I(X;YZ) - I(X;Y) = -\left\langle \log \frac{p(x)p(y,z)}{p(x,y,z)} + \log \frac{p(x,y)}{p(x)p(y)} \right\rangle$$
$$= -\left\langle \log \frac{p(x,y)}{p(y)} \frac{p(y,z)}{p(y)} \frac{p(y)}{p(x,y,z)} \right\rangle$$
$$= -\left\langle \log \frac{p(x|y)p(z|y)}{p(x,z|y)} \right\rangle = \sum_{y} p(y)I(X;Z|y) \ge 0.$$
(10.72)

For each fixed y, p(x, z|y) is a normalized probability distribution with nonnegative

mutual information; hence I(X; YZ) - I(X; Y) is a convex combination of nonnegative terms and therefore nonnegative. The quantity I(X; Z|Y) := I(X; YZ) - I(X; Y) is called the *conditional mutual information*, because it quantifies how strongly X and Z are correlated when Y is known; strong subadditivity can be restated as the nonnegativity of conditional mutual information,

$$I(X;Z|Y) \ge 0.$$
 (10.73)

One might ask under what conditions strong subadditivity is satisfied as an equality; that is, when does the conditional mutual information vanish? Since I(X; Z|Y) is sum of nonnegative terms, each of these terms must vanish if I(X; Z|Y) = 0. Therefore for each y with p(y) > 0, we have I(X; Z|y) = 0. The mutual information vanishes only for a product distribution, therefore

$$p(x, z|y) = p(x|y)p(z|y) \implies p(x, y, z) = p(x|y)p(z|y)p(y).$$
 (10.74)

This means that the correlations between x and z arise solely from their shared correlation with y, in which case we say that x and z are *conditionally independent*.

Correlations of quantum systems also obey strong subadditivity:

$$I(A; BC) - I(A; B) := I(A; C|B) \ge 0.$$
(10.75)

But while the proof is elementary in the classical case, in the quantum setting strong subadditivity is a rather deep result with many important consequences. We will postpone the proof until §10.8.3, where we will be able to justify the quantum statement by giving it a clear operational meaning. We'll also see in Exercise 10.3 that strong subadditivity follows easily from another deep property, the monotonicity of relative entropy:

$$D(\boldsymbol{\rho}_A \| \boldsymbol{\sigma}_A) \le D(\boldsymbol{\rho}_{AB} \| \boldsymbol{\sigma}_{AB}), \qquad (10.76)$$

where

$$D(\boldsymbol{\rho} \| \boldsymbol{\sigma}) := \operatorname{tr} \, \boldsymbol{\rho} \left(\log \boldsymbol{\rho} - \log \boldsymbol{\sigma} \right) \,. \tag{10.77}$$

The relative entropy of two density operators on a system AB cannot be less than the induced relative entropy on the subsystem A. Insofar as we can regard the relative entropy as a measure of the "distance" between density operators, monotonicity is the reasonable statement that quantum states become no easier to distinguish when we look at the subsystem A than when we look at the full system AB. It also follows (Exercise 10.3), that the action of a quantum channel \mathcal{N} cannot increase relative entropy:

$$D(\mathcal{N}(\boldsymbol{\rho}) \| \mathcal{N}(\boldsymbol{\sigma})) \le D(\boldsymbol{\rho} \| \boldsymbol{\sigma})$$
(10.78)

There are a few other ways of formulating strong subadditivity which are helpful to keep in mind. By expressing the quantum mutual information in terms of the Von Neumann entropy we find

$$H(ABC) + H(B) \le H(AB) + H(BC).$$
 (10.79)

While A, B, C are three disjoint quantum systems, we may view AB and BC as overlapping systems with intersection B and union ABC; then strong subadditivity says that the sum of the entropies of two overlapping systems is at least as large as the sum of the

entropies of their union and their intersection. In terms of conditional entropy, strong subadditivity becomes

$$H(A|B) \ge H(A|BC); \tag{10.80}$$

loosely speaking, our ignorance about A when we know only B is no smaller than our ignorance about A when we know both B and C, but with the proviso that for quantum information "ignorance" can sometimes be negative!

As in the classical case, it is instructive to consider the condition for equality in strong subadditivity. What does it mean for systems to have *quantum conditional independence*, I(A; C|B) = 0? It is easy to formulate a sufficient condition. Suppose that system B has a decomposition as a direct sum of tensor products of Hilbert spaces

$$\mathcal{H}_B = \bigoplus_j \mathcal{H}_{B_j} = \bigoplus_j \mathcal{H}_{B_j^L} \otimes \mathcal{H}_{B_j^R}, \qquad (10.81)$$

and that the state of ABC has the block diagonal form

$$\boldsymbol{\rho}_{ABC} = \bigoplus_{j} p_{j} \ \boldsymbol{\rho}_{AB_{j}^{L}} \otimes \boldsymbol{\rho}_{B_{j}^{R}C}.$$
(10.82)

In each block labeled by j the state is a tensor product, with conditional mutual information

$$I(A;C|B_j) = I(A;B_jC) - I(A;B_j) = I(A;B_j^L) - I(A;B_j^L) = 0;$$
(10.83)

What is less obvious is that the converse is also true — any state with I(A; C|B) = 0 has a decomposition as in eq.(10.82). This is a useful fact, though we will not give the proof here.

10.2.4 Monotonicity of mutual information

Strong subadditivity implies another important property of quantum mutual information, its monotonicity — a quantum channel acting on system B cannot increase the mutual information of A and B. To derive monotonicity, suppose that a quantum channel $\mathcal{N}^{B\to B'}$ maps B to B'. Like any quantum channel, \mathcal{N} has an isometric extension, its Stinespring dilation $U^{B\to B'E}$, mapping B to B' and a suitable environment system E. Since the isometry U does not change the eigenvalues of the density operator, it preserves the entropy of B and of AB,

$$H(B) = H(B'E), \quad H(AB) = H(AB'E),$$
 (10.84)

which implies

$$I(A; B) = H(A) + H(B) - H(AB)$$

= $H(A) + H(B'E) - H(ABE') = I(A; B'E).$ (10.85)

From strong subadditivity, we obtain

$$I(A;B) = I(A;B'E) \ge I(A,B')$$
(10.86)

the desired statement of monotonicity.

10.2 Von Neumann Entropy

10.2.5 Entropy and thermodynamics

The concept of entropy first entered science through the study of thermodynamics, and the mathematical properties of entropy we have enumerated have many interesting thermodynamic implications. Here we will just mention a few ways in which the nonnegativity and monotonicity of quantum relative entropy relate to ideas encountered in thermodynamics.

There are two distinct ways to approach the foundations of quantum statistical physics. In one, we consider the evolution of an isolated closed quantum system, but ask what we will observe if we have access to only a portion of the full system. Even though the evolution of the full system is unitary, the evolution of a subsystem is not, and the subsystem may be accurately described by a thermal ensemble at late times. Information which is initially encoded locally in an out-of-equilibrium state becomes encoded more and more nonlocally as the system evolves, eventually becoming invisible to an observer confined to the subsystem.

In the other approach, we consider the evolution of an open system A, in contact with an unobserved environment E, and track the evolution of A only. From a fundamental perspective this second approach may be regarded as a special case of the first, since AE is closed, with A as a privileged subsystem. In practice, though, it is often more convenient to describe the evolution of an open system using a master equation as in Chapter 3, and to analyze evolution toward thermal equilibrium without explicit reference to the environment.

Free energy and the second law.

Tools of quantum Shannon theory can help us understand why the state of an open system with Hamiltonian H might be expected to be close to the thermal *Gibbs state*

$$\boldsymbol{\rho}_{\beta} = \frac{e^{-\beta \boldsymbol{H}}}{\operatorname{tr}\left(e^{-\beta \boldsymbol{H}}\right)},\tag{10.87}$$

where $kT = \beta^{-1}$ is the temperature. Here let's observe one noteworthy feature of this state. For an arbitrary density operator ρ , consider its *free energy*

$$F(\boldsymbol{\rho}) = E(\boldsymbol{\rho}) - \beta^{-1}S(\boldsymbol{\rho}) \tag{10.88}$$

where $E(\boldsymbol{\rho}) = \langle \boldsymbol{H} \rangle_{\boldsymbol{\rho}}$ denotes the expectation value of the Hamiltonian in this state; for this subsection we respect the conventions of thermodynamics by denoting Von Neumann entropy by $S(\boldsymbol{\rho})$ rather than $H(\boldsymbol{\rho})$ (lest H be confused with the Hamiltonian \boldsymbol{H}), and by using natural logarithms. Expressing $F(\boldsymbol{\rho})$ and the free energy $F(\boldsymbol{\rho}_{\beta})$ of the Gibbs state as

$$F(\boldsymbol{\rho}) = \operatorname{tr}(\boldsymbol{\rho}\boldsymbol{H}) - \beta^{-1}S(\boldsymbol{\rho}) = \beta^{-1}\operatorname{tr}\boldsymbol{\rho}(\ln\boldsymbol{\rho} + \beta\boldsymbol{H}),$$

$$F(\boldsymbol{\rho}_{\beta}) = -\beta^{-1}\ln\left(\operatorname{tr} e^{-\beta\boldsymbol{H}}\right),$$
(10.89)

we see that the relative entropy of ρ and ρ_{β} is

$$D(\boldsymbol{\rho} \| \boldsymbol{\rho}_{\beta}) = \operatorname{tr} (\boldsymbol{\rho} \ln \boldsymbol{\rho}) - \operatorname{tr} (\boldsymbol{\rho} \ln \boldsymbol{\rho}_{\beta}) = \beta \left(F(\boldsymbol{\rho}) - F(\boldsymbol{\rho}_{\beta}) \right) \ge 0,$$
(10.90)

with equality only for $\rho = \rho_{\beta}$. The nonnegativity of relative entropy implies that at a given temperature β^{-1} , the Gibbs state ρ_{β} has the lowest possible free energy. Our open

system, in contact with a thermal reservoir at temperature β^{-1} , will prefer the Gibbs state if it wishes to minimize its free energy.

What can we say about the *approach* to thermal equilibrium of an open system? We may anticipate that the joint unitary evolution of system and reservoir induces a quantum channel \mathcal{N} acting on the system alone, and we know that relative entropy is monotonic — if

$$\mathcal{N}: \boldsymbol{\rho} \mapsto \boldsymbol{\rho}', \quad \mathcal{N}: \boldsymbol{\sigma} \mapsto \boldsymbol{\sigma}',$$
 (10.91)

then

$$D(\boldsymbol{\rho}' \| \boldsymbol{\sigma}') \le D(\boldsymbol{\rho} \| \boldsymbol{\sigma}). \tag{10.92}$$

Furthermore, if the Gibbs state is an *equilibrium* state, we expect this channel to preserve the Gibbs state

$$\mathcal{N}: \boldsymbol{\rho}_{\beta} \mapsto \boldsymbol{\rho}_{\beta}; \tag{10.93}$$

therefore,

$$D(\boldsymbol{\rho}' \| \boldsymbol{\rho}_{\beta}) = \beta \left(F(\boldsymbol{\rho}') - F(\boldsymbol{\rho}_{\beta}) \right) \le \beta \left(F(\boldsymbol{\rho}) - F(\boldsymbol{\rho}_{\beta}) \right) = D(\boldsymbol{\rho} \| \boldsymbol{\rho}_{\beta}), \tag{10.94}$$

and hence

$$F(\boldsymbol{\rho}') \le F(\boldsymbol{\rho}). \tag{10.95}$$

Any channel that preserves the Gibbs state cannot increase the free energy; instead, free energy of an out-of-equilibrium state is monotonically decreasing under open-state evolution. This statement is a version of the second law of thermodynamics.

10.2.6 Bekenstein's entropy bound.

Similar ideas lead to Bekenstein's bound on entropy in quantum field theory. The fieldtheoretic details, though interesting, would lead us far afield. The gist is that Bekenstein proposed an inequality relating the energy and the entropy in a bounded spatial region. This bound was motivated by gravitational physics, but can be formulated without reference to gravitation, and follows from properties of relative entropy.

A subtlety is that entropy of a region is infinite in quantum field theory, because of contributions coming from arbitrarily short-wavelength quantum fluctuations near the boundary of the region. Therefore we have to make a subtraction to define a finite quantity. The natural way to do this is to subtract away the entropy of the same region in the vacuum state of the theory, as any finite energy state in a finite volume has the same structure as the vacuum at very short distances. Although the vacuum is a pure state, it, and any other reasonable state, has a marginal state in a finite region which is highly mixed, because of entanglement between the region and its complement.

For the purpose of our discussion here, we may designate any mixed state ρ_0 we choose supported in the bounded region as the "vacuum," and define a corresponding "modular Hamiltonian" K by

$$\boldsymbol{\rho}_0 = \frac{e^{-\boldsymbol{K}}}{\operatorname{tr}\left(e^{-\boldsymbol{K}}\right)}.\tag{10.96}$$

20

That is, we regard the state as the thermal mixed state of K, with the temperature arbitrarily set to unity (which is just a normalization convention for K). Then by rewriting eq.(10.90) we see that, for any state ρ , $D(\rho || \rho_0) \ge 0$ implies

$$S(\boldsymbol{\rho}) - S(\boldsymbol{\rho}_0) \le \operatorname{tr}(\boldsymbol{\rho} \boldsymbol{K}) - \operatorname{tr}(\boldsymbol{\rho}_0 \boldsymbol{K})$$
(10.97)

The left-hand side, the entropy with vacuum entropy subtracted, is not larger than the right-hand side, the (modular) energy with vacuum energy subtracted. This is one version of Bekenstein's bound. Here K, which is dimensionless, can be loosely interpreted as ER, where E is the energy contained in the region and R is its linear size.

While the bound follows easily from nonnegativity of relative entropy, the subtle part of the argument is recognizing that the (suitably subtracted) expectation value of the modular Hamiltonian is a reasonable way to define ER. The detailed justification for this involves properties of relativistic quantum field theory that we won't go into here. Suffice it to say that, because we constructed K by regarding the marginal state of the vacuum as the Gibbs state associated with the Hamiltonian K, we expect K to be linear in the energy, and dimensional analysis then requires inclusion of the factor of R(in units with $\hbar = c = 1$).

Bekenstein was led to conjecture such a bound by thinking about black hole thermodynamics. Leaving out numerical factors, just to get a feel for the orders of magnitude of things, the entropy of a black hole with circumference ~ R is $S \sim R^2/G$, and its mass (energy) is $E \sim R/G$, where G is Newton's gravitational constant; hence $S \sim ER$ for a black hole. Bekenstein realized that unless S = O(ER) for arbitrary states and regions, we could throw extra stuff into the region, making a black hole with lower entropy than the initial state, thus violating the (generalized) second law of thermodynamics. Though black holes provided the motivation, G drops out of the inequality, which holds even in nongravitational relativistic quantum field theories.

10.2.7 Entropic uncertainty relations

The uncertainty principle asserts that noncommuting observables cannot simultaneously have definite values. To translate this statement into mathematics, recall that a Hermitian observable \boldsymbol{A} has spectral representation

$$\boldsymbol{A} = \sum_{x} |x\rangle a(x)\langle x| \tag{10.98}$$

where $\{|x\rangle\}$ is the orthonormal basis of eigenvectors of A and $\{a(x)\}$ is the corresponding vector of eigenvalues; if A is measured in the state ρ , the outcome a(x) occurs with probability $p(x) = \langle x | \rho | x \rangle$. Thus A has expectation value $\operatorname{tr}(\rho A)$ and variance

$$(\Delta A)^2 = \operatorname{tr}\left(\boldsymbol{\rho}A^2\right) - (\operatorname{tr}\boldsymbol{\rho}A)^2.$$
(10.99)

Using the Cauchy-Schwarz inequality, we can show that if A and B are two Hermitian observables and $\rho = |\psi\rangle\langle\psi|$ is a pure state, then

$$\Delta A \Delta B \ge \frac{1}{2} |\langle \psi | [\boldsymbol{A}, \boldsymbol{B}] | \psi \rangle|.$$
(10.100)

Eq.(10.100) is a useful statement of the uncertainty principle, but has drawbacks. It depends on the state $|\psi\rangle$ and for that reason does not fully capture the incompatibility of the two observables. Furthermore, the variance does not characterize very well the

unpredictability of the measurement outcomes; entropy would be a more informative measure.

In fact there are *entropic uncertainty relations* which do not suffer from these deficiencies. If we measure a state ρ by projecting onto the orthonormal basis $\{|x\rangle\}$, the outcomes define a classical ensemble

$$X = \{x, p(x) = \langle x | \boldsymbol{\rho} | x \rangle\}; \tag{10.101}$$

that is, a probability vector whose entries are the diagonal elements of ρ in the *x*basis. The Shannon entropy H(X) quantifies how uncertain we are about the outcome before we perform the measurement. If $\{|z\rangle\}$ is another orthonormal basis, there is a corresponding classical ensemble Z describing the probability distribution of outcomes when we measure the same state ρ in the z-basis. If the two bases are incompatible, there is a tradeoff between our uncertainty about X and about Z, captured by the inequality

$$H(X) + H(Z) \ge \log\left(\frac{1}{c}\right) + H(\boldsymbol{\rho}), \qquad (10.102)$$

where

$$c = \max_{x,z} |\langle x|z \rangle|^2.$$
(10.103)

The second term on the right-hand side, which vanishes if ρ is a pure state, reminds us that our uncertainty increases when the state is mixed. Like many good things in quantum information theory, this entropic uncertainty relation follows from the monotonicity of the quantum relative entropy.

For each measurement there is a corresponding quantum channel, realized by performing the measurement and printing the outcome in a classical register,

$$\mathcal{M}_X : \boldsymbol{\rho} \mapsto \sum_x |x\rangle \langle x|\boldsymbol{\rho}|x\rangle \langle x| =: \boldsymbol{\rho}_X,$$

$$\mathcal{M}_Z : \boldsymbol{\rho} \mapsto \sum_z |z\rangle \langle z|\boldsymbol{\rho}|z\rangle \langle z| =: \boldsymbol{\rho}_Z.$$
 (10.104)

The Shannon entropy of the measurement outcome distribution is also the Von Neumann entropy of the corresponding channel's output state,

$$H(X) = H(\rho_X), \quad H(Z) = H(\rho_Z);$$
 (10.105)

the entropy of this output state can be expressed in terms of the relative entropy of input and output, and the entropy of the channel input, as in

$$H(X) = -\mathrm{tr}\boldsymbol{\rho}_X \log \boldsymbol{\rho}_X = -\mathrm{tr}\boldsymbol{\rho} \log \boldsymbol{\rho}_X = D(\boldsymbol{\rho} \| \boldsymbol{\rho}_X) + H(\boldsymbol{\rho}).$$
(10.106)

Using the monotonicity of relative entropy under the action of the channel \mathcal{M}_Z , we have

$$D(\boldsymbol{\rho} \| \boldsymbol{\rho}_X) \ge D(\boldsymbol{\rho}_Z \| \mathcal{M}_Z(\boldsymbol{\rho}_X)), \tag{10.107}$$

where

$$D(\boldsymbol{\rho}_{Z} \| \mathcal{M}_{Z}(\boldsymbol{\rho}_{X})) = -H(\boldsymbol{\rho}_{Z}) - \operatorname{tr} \boldsymbol{\rho}_{Z} \log \mathcal{M}_{Z}(\boldsymbol{\rho}_{X}), \qquad (10.108)$$

and

$$\mathcal{M}_{Z}(\boldsymbol{\rho}_{X}) = \sum_{x,z} |z\rangle \langle z|x\rangle \langle x|\boldsymbol{\rho}|x\rangle \langle x|z\rangle \langle z|.$$
(10.109)

Writing

$$\log \mathcal{M}_Z(\boldsymbol{\rho}_X) = \sum_{z} |z\rangle \log \left(\sum_{x} \langle z | x \rangle \langle x | \boldsymbol{\rho} | x \rangle \langle x | z \rangle \right) \langle z|, \qquad (10.110)$$

we see that

$$-\mathrm{tr}\boldsymbol{\rho}_{Z}\log\mathcal{M}_{Z}(\boldsymbol{\rho}_{X}) = -\sum_{z}\langle z|\boldsymbol{\rho}|z\rangle\log\left(\sum_{x}\langle z|x\rangle\langle x|\boldsymbol{\rho}|x\rangle\langle x|z\rangle\right).$$
 (10.111)

Now, because $-\log(\cdot)$ is a monotonically decreasing function, we have

$$-\log\left(\sum_{x}\langle z|x\rangle\langle x|\boldsymbol{\rho}|x\rangle\langle x|z\rangle\right) \ge -\log\left(\max_{x,z}|\langle x|z\rangle|^{2}\sum_{x}\langle x|\boldsymbol{\rho}|x\rangle\right)$$
$$=\log\left(\frac{1}{c}\right),$$
(10.112)

and therefore

$$-\mathrm{tr}\boldsymbol{\rho}_Z \log \mathcal{M}_Z(\boldsymbol{\rho}_X) \ge \log\left(\frac{1}{c}\right).$$
 (10.113)

Finally, putting together eq.(10.106), (10.107) (10.108), (10.113), we find

$$H(X) - H(\boldsymbol{\rho}) = D(\boldsymbol{\rho} \| \boldsymbol{\rho}_X) \ge D(\boldsymbol{\rho}_Z \| \mathcal{M}_Z(\boldsymbol{\rho}_X))$$

= $-H(Z) - \operatorname{tr} \boldsymbol{\rho}_Z \log \mathcal{M}_Z(\boldsymbol{\rho}_X) \ge -H(Z) + \log\left(\frac{1}{c}\right),$ (10.114)

which is equivalent to eq.(10.102).

We say that two different bases $\{|x\rangle\}$, $\{|z\rangle\}$ for a *d*-dimensional Hilbert space are *mutually unbiased* if for all x, z

$$|\langle x|z\rangle|^2 = \frac{1}{d}; \tag{10.115}$$

thus, if we measure any x-basis state $|x\rangle$ in the z-basis, all d outcomes are equally probable. For measurements in two mutually unbiased bases performed on a pure state, the entropic uncertainty relation becomes

$$H(X) + H(Z) \ge \log d. \tag{10.116}$$

Clearly this inequality is tight, as it is saturated by x-basis (or z-basis) states, for which H(X) = 0 and $H(Z) = \log d$.

10.3 Quantum Source Coding

What is the quantum analog of Shannon's source coding theorem?

Let's consider a long message consisting of n letters, where each letter is a pure quantum state chosen by sampling from the ensemble

$$\{|\varphi(x)\rangle, p(x)\}.$$
(10.117)

If the states of this ensemble are mutually orthogonal, then the message might as well be classical; the interesting quantum case is where the states are not orthogonal and

23

therefore not perfectly distinguishable. The density operator realized by this ensemble is

$$\boldsymbol{\rho} = \sum_{x} p(x) |\varphi(x)\rangle \langle \varphi(x)|, \qquad (10.118)$$

and the entire n-letter message has the density operator

$$\boldsymbol{\rho}^{\otimes n} = \boldsymbol{\rho} \otimes \cdots \otimes \boldsymbol{\rho}. \tag{10.119}$$

How redundant is the quantum information in this message? We would like to devise a quantum code allowing us to compress the message to a smaller Hilbert space, but without much compromising the fidelity of the message. Perhaps we have a quantum memory device, and we know the *statistical* properties of the recorded data; specifically, we know ρ . We want to conserve space on our (very expensive) quantum hard drive by compressing the data.

The optimal compression that can be achieved was found by Schumacher. As you might guess, the message can be compressed to a Hilbert space \mathcal{H} with

$$\dim \mathcal{H} = 2^{n(H(\boldsymbol{\rho}) + o(1))} \tag{10.120}$$

with negligible loss of fidelity as $n \to \infty$, while errorless compression to dimension $2^{n(H(\rho)-\Omega(1))}$ is not possible. In this sense, the Von Neumann entropy is the number of *qubits* of quantum information carried per letter of the message. Compression is always possible unless ρ is maximally mixed, just as we can always compress a classical message unless the information source is uniformly random. This result provides a precise operational interpretation for Von Neumann entropy.

Once Shannon's results are known and understood, the proof of Schumacher's compression theorem is not difficult, as the mathematical ideas needed are very similar to those used by Shannon. But conceptually quantum compression is very different from its classical counterpart, as the imperfect distinguishability of nonorthogonal quantum states is the central idea.

10.3.1 Quantum compression: an example

Before discussing Schumacher's quantum compression protocol in full generality, it is helpful to consider a simple example. Suppose that each letter is a single qubit drawn from the ensemble

$$|\uparrow_z\rangle = \begin{pmatrix} 1\\ 0 \end{pmatrix}, \quad p = \frac{1}{2},$$
 (10.121)

$$|\uparrow_x\rangle = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix}, \quad p = \frac{1}{2},$$
 (10.122)

so that the density operator of each letter is

$$\rho = \frac{1}{2} |\uparrow_z\rangle\langle\uparrow_z| + \frac{1}{2} |\uparrow_x\rangle\langle\uparrow_x|$$

= $\frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} = \begin{pmatrix} \frac{3}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} \end{pmatrix}.$ (10.123)

As is obvious from symmetry, the eigenstates of ρ are qubits oriented up and down along the axis $\hat{n} = \frac{1}{\sqrt{2}}(\hat{x} + \hat{z})$,

$$|0'\rangle \equiv |\uparrow_{\hat{n}}\rangle = \begin{pmatrix} \cos \frac{\pi}{8} \\ \sin \frac{\pi}{8} \end{pmatrix},$$

$$|1'\rangle \equiv |\downarrow_{\hat{n}}\rangle = \begin{pmatrix} \sin \frac{\pi}{8} \\ -\cos \frac{\pi}{8} \end{pmatrix};$$
 (10.124)

the eigenvalues are

$$\lambda(0') = \frac{1}{2} + \frac{1}{2\sqrt{2}} = \cos^2 \frac{\pi}{8},$$

$$\lambda(1') = \frac{1}{2} - \frac{1}{2\sqrt{2}} = \sin^2 \frac{\pi}{8};$$
 (10.125)

evidently $\lambda(0') + \lambda(1') = 1$ and $\lambda(0')\lambda(1') = \frac{1}{8} = \det \rho$. The eigenstate $|0'\rangle$ has equal (and relatively large) overlap with both signal states

$$|\langle 0'|\uparrow_z\rangle|^2 = |\langle 0'|\uparrow_x\rangle|^2 = \cos^2\frac{\pi}{8} = .8535,$$
 (10.126)

while $|1'\rangle$ has equal (and relatively small) overlap with both,

$$|\langle 1'|\uparrow_z\rangle|^2 = |\langle 1'|\uparrow_x\rangle|^2 = \sin^2\frac{\pi}{8} = .1465.$$
 (10.127)

Thus if we don't know whether $|\uparrow_z\rangle$ or $|\uparrow_x\rangle$ was sent, the best guess we can make is $|\psi\rangle = |0'\rangle$. This guess has the maximal *fidelity* with ρ

$$F = \frac{1}{2} |\langle \uparrow_z |\psi \rangle|^2 + \frac{1}{2} |\langle \uparrow_x |\psi \rangle|^2, \qquad (10.128)$$

among all possible single-qubit states $|\psi\rangle$ (F = .8535).

Now imagine that Alice needs to send three letters to Bob, but she can afford to send only two qubits. Still, she wants Bob to reconstruct her state with the highest possible fidelity. She could send Bob two of her three letters, and ask Bob to guess $|0'\rangle$ for the third. Then Bob receives two letters with perfect fidelity, and his guess has F = .8535for the third; hence F = .8535 overall. But is there a more clever procedure that achieves higher fidelity?

Yes, there is. By diagonalizing ρ , we decomposed the Hilbert space of a single qubit into a "likely" one-dimensional subspace (spanned by $|0'\rangle$) and an "unlikely" onedimensional subspace (spanned by $|1'\rangle$). In a similar way we can decompose the Hilbert space of three qubits into likely and unlikely subspaces. If $|\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \otimes |\psi_3\rangle$ is any signal state, where the state of each qubit is either $|\uparrow_z\rangle$ or $|\uparrow_x\rangle$, we have

$$\begin{aligned} |\langle 0'0'0'|\psi\rangle|^2 &= \cos^6\left(\frac{\pi}{8}\right) = .6219, \\ |\langle 0'0'1'|\psi\rangle|^2 &= |\langle 0'1'0'|\psi\rangle|^2 = |\langle 1'0'0'|\psi\rangle|^2 = \cos^4\left(\frac{\pi}{8}\right)\sin^2\left(\frac{\pi}{8}\right) = .1067, \\ |\langle 0'1'1'|\psi\rangle|^2 &= |\langle 1'0'1'|\psi\rangle|^2 = |\langle 1'1'0'|\psi\rangle|^2 = \cos^2\left(\frac{\pi}{8}\right)\sin^4\left(\frac{\pi}{8}\right) = .0183, \\ |\langle 1'1'1'|\psi\rangle|^2 &= \sin^6\left(\frac{\pi}{8}\right) = .0031. \end{aligned}$$
(10.129)

Thus, we may decompose the space into the likely subspace Λ spanned by $\{|0'0'0'\rangle, |0'0'1'\rangle, |0'1'0'\rangle, |1'0'0'\rangle\}$, and its orthogonal complement Λ^{\perp} . If we make an

incomplete orthogonal measurement that projects a signal state onto Λ or Λ^{\perp} , the probability of projecting onto the likely subspace Λ is

$$p_{\text{likely}} = .6219 + 3(.1067) = .9419, \tag{10.130}$$

while the probability of projecting onto the unlikely subspace is

$$p_{\text{unlikely}} = 3(.0183) + .0031 = .0581.$$
 (10.131)

To perform this measurement, Alice could, for example, first apply a unitary transformation U that rotates the four high-probability basis states to

$$|\cdot\rangle \otimes |\cdot\rangle \otimes |0\rangle, \tag{10.132}$$

and the four low-probability basis states to

$$|\cdot\rangle \otimes |\cdot\rangle \otimes |1\rangle; \tag{10.133}$$

then Alice measures the third qubit to perform the projection. If the outcome is $|0\rangle$, then Alice's input state has in effect been projected onto Λ . She sends the remaining two unmeasured qubits to Bob. When Bob receives this compressed two-qubit state $|\psi_{\text{comp}}\rangle$, he decompresses it by appending $|0\rangle$ and applying U^{-1} , obtaining

$$|\psi'\rangle = U^{-1}(|\psi_{\text{comp}}\rangle \otimes |0\rangle).$$
 (10.134)

If Alice's measurement of the third qubit yields $|1\rangle$, she has projected her input state onto the low-probability subspace Λ^{\perp} . In this event, the best thing she can do is send the state that Bob will decompress to the most likely state $|0'0'0'\rangle$ – that is, she sends the state $|\psi_{comp}\rangle$ such that

$$|\psi'\rangle = U^{-1}(|\psi_{\text{comp}}\rangle \otimes |0\rangle) = |0'0'0'\rangle.$$
(10.135)

Thus, if Alice encodes the three-qubit signal state $|\psi\rangle$, sends two qubits to Bob, and Bob decodes as just described, then Bob obtains the state ρ'

$$|\psi\rangle\langle\psi|\to\boldsymbol{\rho}'=\boldsymbol{E}|\psi\rangle\langle\psi|\boldsymbol{E}+|0'0'0'\rangle\langle\psi|(\boldsymbol{I}-\boldsymbol{E})|\psi\rangle\langle0'0'0'|,\qquad(10.136)$$

where E is the projection onto Λ . The fidelity achieved by this procedure is

$$F = \langle \psi | \boldsymbol{\rho}' | \psi \rangle = (\langle \psi | \boldsymbol{E} | \psi \rangle)^2 + (\langle \psi | (\boldsymbol{I} - \boldsymbol{E}) | \psi \rangle) (\langle \psi | 0' 0' 0' \rangle)^2$$

= (.9419)² + (.0581)(.6219) = .9234. (10.137)

This is indeed better than the naive procedure of sending two of the three qubits each with perfect fidelity.

As we consider longer messages with more letters, the fidelity of the compression improves, as long as we don't try to compress too much. The Von-Neumann entropy of the one-qubit ensemble is

$$H(\boldsymbol{\rho}) = H\left(\cos^2\frac{\pi}{8}\right) = .60088\dots$$
 (10.138)

Therefore, according to Schumacher's theorem, we can shorten a long message by the factor, say, .6009, and still achieve very good fidelity.

10.3 Quantum Source Coding

10.3.2 Schumacher compression in general

The key to Shannon's noiseless coding theorem is that we can code the typical sequences and ignore the rest, without much loss of fidelity. To quantify the compressibility of quantum information, we promote the notion of a typical *sequence* to that of a typical *subspace*. The key to Schumacher's noiseless quantum coding theorem is that we can code the typical subspace and ignore its orthogonal complement, without much loss of fidelity.

We consider a message of n letters where each letter is a pure quantum state drawn from the ensemble $\{|\varphi(x)\rangle, p(x)\}$, so that the density operator of a single letter is

$$\boldsymbol{\rho} = \sum_{x} p(x) |\varphi(x)\rangle \langle \varphi(x)|. \tag{10.139}$$

Since the letters are drawn independently, the density operator of the entire message is

$$\boldsymbol{\rho}^{\otimes n} \equiv \boldsymbol{\rho} \otimes \cdots \otimes \boldsymbol{\rho}. \tag{10.140}$$

We claim that, for *n* large, this density matrix has nearly all of its support on a subspace of the full Hilbert space of the messages, where the dimension of this subspace asymptotically approaches $2^{nH(\rho)}$.

This claim follows directly from the corresponding classical statement, for we may consider ρ to be realized by an ensemble of orthonormal pure states, its eigenstates, where the probability assigned to each eigenstate is the corresponding eigenvalue. In this basis our source of quantum information is effectively classical, producing messages which are tensor products of ρ eigenstates, each with a probability given by the product of the corresponding eigenvalues. For a specified n and δ , define the δ -typical subspace Λ as the space spanned by the eigenvectors of $\rho^{\otimes n}$ with eigenvalues λ satisfying

$$2^{-n(H-\delta)} > \lambda > 2^{-n(H+\delta)}.$$
(10.141)

Borrowing directly from Shannon's argument, we infer that for any $\delta, \varepsilon > 0$ and n sufficiently large, the sum of the eigenvalues of $\rho^{\otimes n}$ that obey this condition satisfies

$$\operatorname{tr}(\boldsymbol{\rho}^{\otimes n}\boldsymbol{E}) \ge 1 - \varepsilon, \tag{10.142}$$

where E denotes the projection onto the typical subspace Λ , and the number dim (Λ) of such eigenvalues satisfies

$$2^{n(H+\delta)} \ge \dim(\Lambda) \ge (1-\varepsilon)2^{n(H-\delta)}.$$
(10.143)

Our coding strategy is to send states in the typical subspace faithfully. We can make a measurement that projects the input message onto either Λ or Λ^{\perp} ; the outcome will be Λ with probability $p_{\Lambda} = \operatorname{tr}(\boldsymbol{\rho}^{\otimes n} \boldsymbol{E}) \geq 1 - \varepsilon$. In that event, the projected state is coded and sent. Asymptotically, the probability of the other outcome becomes negligible, so it matters little what we do in that case.

The coding of the projected state merely packages it so it can be carried by a minimal number of qubits. For example, we apply a unitary change of basis \boldsymbol{U} that takes each state $|\psi_{\text{typ}}\rangle$ in Λ to a state of the form

$$\boldsymbol{U}|\psi_{\rm typ}\rangle = |\psi_{\rm comp}\rangle \otimes |0_{\rm rest}\rangle, \qquad (10.144)$$

where $|\psi_{\text{comp}}\rangle$ is a state of $n(H + \delta)$ qubits, and $|0_{\text{rest}}\rangle$ denotes the state $|0\rangle \otimes \ldots \otimes |0\rangle$ of the remaining qubits. Alice sends $|\psi_{\text{comp}}\rangle$ to Bob, who decodes by appending $|0_{\text{rest}}\rangle$ and applying U^{-1} .

Quantum Shannon Theory

Suppose that

$$|\varphi(\vec{x})\rangle = |\varphi(x_1)\rangle \otimes \ldots \otimes |\varphi(x_n)\rangle, \qquad (10.145)$$

denotes any one of the n-letter pure state messages that might be sent. After coding, transmission, and decoding are carried out as just described, Bob has reconstructed a state

$$\begin{aligned} |\varphi(\vec{x})\rangle\langle\varphi(\vec{x})| &\mapsto \rho'(\vec{x}) = \boldsymbol{E}|\varphi(\vec{x})\rangle\langle\varphi(\vec{x})|\boldsymbol{E} \\ &+ \rho_{\text{Junk}}(\vec{x})\langle\varphi(\vec{x})|(\boldsymbol{I} - \boldsymbol{E})|\varphi(\vec{x})\rangle, \end{aligned} \tag{10.146}$$

where $\rho_{\text{Junk}}(\vec{x})$ is the state we choose to send if the measurement yields the outcome Λ^{\perp} . What can we say about the fidelity of this procedure?

The fidelity varies from message to message, so we consider the fidelity averaged over the ensemble of possible messages:

$$\bar{F} = \sum_{\vec{x}} p(\vec{x}) \langle \varphi(\vec{x}) | \boldsymbol{\rho}'(\vec{x}) | \varphi(\vec{x}) \rangle$$

$$= \sum_{\vec{x}} p(\vec{x}) \langle \varphi(\vec{x}) | \boldsymbol{E} | \varphi(\vec{x}) \rangle \langle \varphi(\vec{x}) | \boldsymbol{E} | \varphi(\vec{x}) \rangle$$

$$+ \sum_{\vec{x}} p(\vec{x}) \langle \varphi(\vec{x}) | \boldsymbol{\rho}_{\text{Junk}}(\vec{x}) | \varphi(\vec{x}) \rangle \langle \varphi(\vec{x}) | \boldsymbol{I} - \boldsymbol{E} | \varphi(\vec{x}) \rangle$$

$$\geq \sum_{\vec{x}} p(\vec{x}) \langle \varphi(\vec{x}) | \boldsymbol{E} | \varphi(\vec{x}) \rangle^{2}, \qquad (10.147)$$

where the last inequality holds because the "Junk" term is nonnegative. Since any real number z satisfies

$$(z-1)^2 \ge 0$$
, or $z^2 \ge 2z-1$, (10.148)

we have (setting $z = \langle \varphi(\vec{x}) | \boldsymbol{E} | \varphi(\vec{x}) \rangle$)

$$\langle \varphi(\vec{x}) | \boldsymbol{E} | \varphi(\vec{x}) \rangle^2 \ge 2 \langle \varphi(\vec{x}) | \boldsymbol{E} | \varphi(\vec{x}) \rangle - 1,$$
 (10.149)

and hence

$$\bar{F} \ge \sum_{\vec{x}} p(\vec{x}) (2\langle \varphi(\vec{x}) | \boldsymbol{E} | \varphi(\vec{x}) \rangle - 1)$$

= 2 tr($\boldsymbol{\rho}^{\otimes n} \boldsymbol{E}$) - 1 \ge 2(1 - \varepsilon) - 1 = 1 - 2\varepsilon. (10.150)

Since ε and δ can be as small as we please, we have shown that it is possible to compress the message to n(H + o(1)) qubits, while achieving an average fidelity that becomes arbitrarily good as n gets large.

Is further compression possible? Let us suppose that Bob will decode the message $\rho_{\text{comp}}(\vec{x})$ that he receives by appending qubits and applying a unitary transformation U^{-1} , obtaining

$$\boldsymbol{\rho}'(\vec{x}) = \boldsymbol{U}^{-1}(\boldsymbol{\rho}_{\text{comp}}(\vec{x}) \otimes |0\rangle \langle 0|) \boldsymbol{U}$$
(10.151)

("unitary decoding"), and suppose that $\rho_{\text{comp}}(\vec{x})$ has been compressed to $n(H - \delta')$ qubits. Then, no matter how the input messages have been encoded, the decoded messages are all contained in a subspace Λ' of Bob's Hilbert space with $\dim(\Lambda') = 2^{n(H-\delta')}$.

28

If the input message is $|\varphi(\vec{x})\rangle$, then the density operator reconstructed by Bob can be diagonalized as

$$\boldsymbol{\rho}'(\vec{x}) = \sum_{a_{\vec{x}}} |a_{\vec{x}}\rangle \lambda_{a_{\vec{x}}} \langle a_{\vec{x}}|, \qquad (10.152)$$

where the $|a_{\vec{x}}\rangle$'s are mutually orthogonal states in Λ' . The fidelity of the reconstructed message is

$$F(\vec{x}) = \langle \varphi(\vec{x}) | \boldsymbol{\rho}'(\vec{x}) | \varphi(\vec{x}) \rangle$$

= $\sum_{a_{\vec{x}}} \lambda_{a_{\vec{x}}} \langle \varphi(\vec{x}) | a_{\vec{x}} \rangle \langle a_{\vec{x}} | \varphi(\vec{x}) \rangle$
 $\leq \sum_{a_{\vec{x}}} \langle \varphi(\vec{x}) | a_{\vec{x}} \rangle \langle a_{\vec{x}} | \varphi(\vec{x}) \rangle \leq \langle \varphi(\vec{x}) | \boldsymbol{E}' | \varphi(\vec{x}) \rangle, \qquad (10.153)$

where E' denotes the orthogonal projection onto the subspace Λ' . The average fidelity therefore obeys

$$\bar{F} = \sum_{\vec{x}} p(\vec{x}) F(\vec{x}) \le \sum_{\vec{x}} p(\vec{x}) \langle \varphi(\vec{x}) | \boldsymbol{E}' | \varphi(\vec{x}) \rangle = \operatorname{tr}(\boldsymbol{\rho}^{\otimes n} \boldsymbol{E}').$$
(10.154)

But, according to the Ky Fan dominance principle discussed in §10.2.2, since E' projects onto a space of dimension $2^{n(H-\delta')}$, $\operatorname{tr}(\rho^{\otimes n}E')$ can be no larger than the sum of the $2^{n(H-\delta')}$ largest eigenvalues of $\rho^{\otimes n}$. The δ -typical eigenvalues of $\rho^{\otimes n}$ are no smaller than $2^{-n(H-\delta)}$, so the sum of the $2^{n(H-\delta')}$ largest eigenvalues can be bounded above:

$$\operatorname{tr}(\boldsymbol{\rho}^{\otimes n}\boldsymbol{E}') \le 2^{n(H-\delta')}2^{-n(H-\delta)} + \varepsilon = 2^{-n(\delta'-\delta)} + \varepsilon, \qquad (10.155)$$

where the $+ \varepsilon$ accounts for the contribution from the atypical eigenvalues. Since we may choose ε and δ as small as we please for sufficiently large n, we conclude that the average fidelity \overline{F} gets small as $n \to \infty$ if we compress to $H(\rho) - \Omega(1)$ qubits per letter. We find, then, that $H(\rho)$ qubits per letter is the optimal compression of the quantum information that can be achieved if we are to obtain good fidelity as n goes to infinity. This is Schumacher's quantum source coding theorem.

The above argument applies to any conceivable encoding scheme, but only to a restricted class of decoding schemes, unitary decodings. The extension of the argument to general decoding schemes is sketched in §10.6.3. The conclusion is the same. The point is that $n(H - \delta)$ qubits are too few to faithfully encode the typical subspace.

There is another useful way to think about Schumacher's quantum compression protocol. Suppose that Alice's density operator $\rho_A^{\otimes n}$ has a *purification* $|\psi\rangle_{RA}$ which Alice shares with Robert. Alice wants to convey her share of $|\psi\rangle_{RA}$ to Bob with high fidelity, sending as few qubits to Bob as possible. To accomplish this task, Alice can use the same procedure as described above, attempting to compress the state of A by projecting onto its typical subspace Λ . Alice's projection succeeds with probability

$$P(\boldsymbol{E}) = \langle \psi | \boldsymbol{I} \otimes \boldsymbol{E} | \psi \rangle = \operatorname{tr} \left(\boldsymbol{\rho}^{\otimes n} \boldsymbol{E} \right) \ge 1 - \varepsilon, \qquad (10.156)$$

where E projects onto Λ , and when successful prepares the state

$$\frac{(\boldsymbol{I} \otimes \boldsymbol{E}) |\psi\rangle}{\sqrt{P(\boldsymbol{E})}}.$$
(10.157)

Therefore, after Bob decompresses, the state he shares with Robert has fidelity F_e with

 $|\psi\rangle$ satisfying

$$F_e \ge \langle \psi | \mathbf{I} \otimes \mathbf{E} | \psi \rangle \langle \psi | \mathbf{I} \otimes \mathbf{E} | \psi \rangle = \left(\operatorname{tr} \left(\boldsymbol{\rho}^{\otimes n} \mathbf{E} \right) \right)^2 \ge (1 - \varepsilon)^2 \ge 1 - 2\varepsilon.$$
(10.158)

We conclude that Alice can transfer her share of the pure state $|\psi\rangle_{RA}$ to Bob by sending $nH(\rho) + o(n)$ qubits, achieving arbitrarily good entanglement fidelity F_e as $n \to \infty$.

To summarize, there is a close analogy between Shannon's classical source coding theorem and Schumacher's quantum source coding theorem. In the classical case, nearly all long messages are typical sequences, so we can code only these and still have a small probability of error. In the quantum case, nearly all long messages have nearly perfect overlap with the typical subspace, so we can code only the typical subspace and still achieve good fidelity.

Alternatively, Alice could send classical information to Bob, the string $x_1x_2 \cdots x_n$, and Bob could follow these classical instructions to reconstruct Alice's state $|\varphi(x_1)\rangle \otimes \ldots \otimes$ $|\varphi(x_n)\rangle$. By this means, they could achieve high-fidelity compression to H(X) + o(1)bits — or qubits — per letter, where X is the classical ensemble $\{x, p(x)\}$. But if $\{|\varphi(x)\rangle, p(x)\}$ is an ensemble of *nonorthogonal* pure states, this classically achievable amount of compression is not optimal; some of the classical information about the preparation of the state is redundant, because the nonorthogonal states cannot be perfectly distinguished. Schumacher coding goes further, achieving optimal compression to $H(\rho) + o(1)$ qubits per letter. Quantum compression packages the message more efficiently than classical compression, but at a price — Bob receives the quantum state Alice intended to send, but Bob doesn't know what he has. In contrast to the classical case, Bob can't fully decipher Alice's quantum message accurately. An attempt to read the message will unavoidably disturb it.

10.4 Entanglement Concentration and Dilution

Any bipartite pure state that is not a product state is entangled. But *how* entangled? Can we compare two states and say that one is more entangled than the other?

For example, consider the two bipartite states

$$\begin{aligned} |\phi^{+}\rangle &= \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle), \\ |\psi\rangle &= \frac{1}{\sqrt{2}} |00\rangle + \frac{1}{2} |11\rangle + \frac{1}{2} |22\rangle. \end{aligned}$$
(10.159)

 $|\phi^+\rangle$ is a maximally entangled state of two qubits, while $|\psi\rangle$ is a *partially* entangled state of two *qutrits*. Which is more entangled?

It is not immediately clear that the question has a meaningful answer. Why should it be possible to find an unambiguous way of ordering all bipartite pure states according to their degree of entanglement? Can we compare a pair of qutrits with a pair of qubits any more than we can compare apples and oranges?

A crucial feature of entanglement is that it cannot be created by local operations and classical communication (LOCC). In particular, if Alice and Bob share a bipartite pure state, its Schmidt number does not increase if Alice or Bob performs a unitary transformation on her/his share of the state, nor if Alice or Bob measures her/his share, even if Alice and Bob exchange classical messages about their actions and measurement outcomes. Therefore, any quantitative measure of entanglement should have the property that LOCC cannot increase it, and it should also vanish for an unentangled product

30

state. An obvious candidate is the Schmidt number, but on reflection it does not seem very satisfactory. Consider

$$|\psi_{\varepsilon}\rangle = \sqrt{1 - 2|\varepsilon|^2} |00\rangle + \varepsilon |11\rangle + \varepsilon |22\rangle, \qquad (10.160)$$

which has Schmidt number 3 for any $|\varepsilon| > 0$. Do we really want to say that $|\psi_{\varepsilon}\rangle$ is "more entangled" than $|\phi^+\rangle$? Entanglement, after all, can be regarded as a resource — we might plan to use it for teleportation, for example — and it seems clear that $|\psi_{\varepsilon}\rangle$ (for $|\varepsilon| \ll 1$) is a less valuable resource than $|\phi^+\rangle$.

It turns out, though, that there is a natural and useful way to quantify the entanglement of any bipartite pure state. To compare two states, we use LOCC to convert both states to a common currency that can be compared directly. The common currency is *maximal* entanglement, and the amount of shared entanglement can be expressed in units of Bell pairs (maximally entangled two-qubit states), also called *ebits* of entanglement.

To quantify the entanglement of a particular bipartite pure state, $|\psi\rangle_{AB}$, imagine preparing *n* identical copies of that state. Alice and Bob share a large supply of maximally entangled *Bell pairs*. Using LOCC, they are to convert *k* Bell pairs $(|\phi^+\rangle_{AB})^{\otimes k})$ to *n* high-fidelity copies of the desired state $(|\psi\rangle_{AB})^{\otimes n}$). What is the minimum number k_{\min} of Bell pairs with which they can perform this task?

To obtain a precise answer, we consider the *asymptotic* setting, requiring arbitrarily high-fidelity conversion in the limit of large n. We say that a rate R of conversion from $|\phi^+\rangle$ to $|\psi\rangle$ is asymptotically achievable if for any $\varepsilon, \delta > 0$, there is an LOCC protocol with

$$\frac{k}{n} \le R + \delta, \tag{10.161}$$

which prepares the target state $|\psi^+\rangle^{\otimes n}$ with fidelity $F \ge 1 - \varepsilon$. We define the *entangle*ment cost E_C of $|\psi\rangle$ as the infimum of achievable conversion rates:

$$E_C(|\psi\rangle) := \inf \{ \text{achievable rate for creating } |\psi\rangle \text{ from Bell pairs} \}.$$
 (10.162)

Asymptotically, we can create many copies of $|\psi\rangle$ by consuming E_C Bell pairs per copy.

Now imagine that *n* copies of $|\psi\rangle_{AB}$ are already shared by Alice and Bob. Using LOCC, Alice and Bob are to convert $(|\psi\rangle_{AB})^{\otimes n}$ back to the standard currency: k' Bell pairs $|\phi^+\rangle_{AB}^{\otimes k'}$. What is the maximum number k'_{max} of Bell pairs they can extract from $|\psi\rangle_{AB}^{\otimes n}$? In this case we say that a rate R' of conversion from $|\psi\rangle$ to $|\phi^+\rangle$ is asymptotically achievable if for any $\varepsilon, \delta > 0$, there is an LOCC protocol with

$$\frac{k'}{n} \ge R' - \delta,\tag{10.163}$$

which prepares the target state $|\phi^+\rangle^{\otimes k'}$ with fidelity $F \ge 1 - \varepsilon$. We define the *distillable* entanglement E_D of $|\psi\rangle$ as the supremum of achievable conversion rates:

 $E_D(|\psi\rangle) := \sup \{ \text{achievable rate for distilling Bell pairs from } |\psi\rangle \}.$ (10.164)

Asymptotically, we can convert many copies of $|\psi\rangle$ to Bell pairs, obtaining E_D Bell pairs per copy of $|\psi\rangle$ consumed.

Since it is an in inviolable principle that LOCC cannot create entanglement, it is certain that

$$E_D(|\psi\rangle) \le E_C(|\psi\rangle); \tag{10.165}$$

otherwise Alice and Bob could increase their number of shared Bell pairs by converting

them to copies of $|\psi\rangle$ and then back to Bell pairs. In fact the entanglement cost and distillable entanglement are *equal* for bipartite pure states. (The story is more complicated for bipartite mixed states; see §10.5.) Therefore, for pure states at least we may drop the subscript, using $E(|\psi\rangle)$ to denote the *entanglement* of $|\psi\rangle$. We don't need to distinguish between entanglement cost and distillable entanglement because conversion of entanglement from one form to another is an asymptotically *reversible* process. *E* quantifies both what we have to pay in Bell pairs to create $|\psi\rangle$, and value of $|\psi\rangle$ in Bell pairs for performing tasks like quantum teleportation which consume entanglement.

But what is the value of $E(|\psi\rangle_{AB})$? Perhaps you can guess — it is

$$E(|\psi\rangle_{AB}) = H(\boldsymbol{\rho}_A) = H(\boldsymbol{\rho}_B), \qquad (10.166)$$

the Von Neumann entropy of Alice's density operator ρ_A (or equivalently Bob's density operator ρ_B). This is clearly the right answer in the case where $|\psi\rangle_{AB}$ is a product of k Bell pairs. In that case ρ_A (or ρ_B) is $\frac{1}{2}I$ for each qubit in Alice's possession

$$\boldsymbol{\rho}_A = \left(\frac{1}{2}\boldsymbol{I}\right)^{\otimes k},\tag{10.167}$$

and

$$H(\boldsymbol{\rho}_A) = k \ H\left(\frac{1}{2}\boldsymbol{I}\right) = k. \tag{10.168}$$

How do we see that $E = H(\rho_A)$ is the right answer for any bipartite pure state?

Though it is perfectly fine to use Bell pairs as the common currency for comparing bipartite entangled states, in the asymptotic setting it is simpler and more natural to allow fractions of a Bell pair, which is what we'll do here. That is, we'll consider a maximally entangled state of two *d*-dimensional systems to be $\log_2 d$ Bell pairs, even if *d* is not a power of two. So our goal will be to show that Alice and Bob can use LOCC to convert shared maximal entanglement of systems with dimension $d = 2^{n(H(\boldsymbol{\rho}_A)+\delta)}$ into *n* copies of $|\psi\rangle$, for any positive δ and with arbitrarily good fidelity as $n \to \infty$, and conversely that Alice and Bob can use LOCC to convert *n* copies of $|\psi\rangle$ into a shared maximally entangled state of *d*-dimensional systems with arbitrarily good fidelity, where $d = 2^{n(H(\boldsymbol{\rho}_A)-\delta)}$. This suffices to demonstrate that $E_C(|\psi\rangle) = E_D(|\psi\rangle) = H(\boldsymbol{\rho}_A)$.

First let's see that if Alice and Bob share $k = n(H(\rho_A) + \delta)$ Bell pairs, then they can prepare $|\psi\rangle_{AB}^{\otimes n}$ with high fidelity using LOCC. They perform this task, called *entan*glement dilution, by combining quantum teleportation with Schumacher compression. To get started, Alice locally creates n copies of $|\psi\rangle_{AC}$, where A and C are systems she controls in her laboratory. Next she wishes to teleport the C^n share of these copies to Bob, but to minimize the consumption of Bell pairs, she should compress C^n before teleporting it.

If A and C are d-dimensional, then the bipartite state $|\psi\rangle_{AC}$ can be expressed in terms of its Schmidt basis as

$$|\psi\rangle_{AC} = \sqrt{p_0} |00\rangle + \sqrt{p_1} |11\rangle + \ldots + \sqrt{p_{d-1}} |d-1, d-1\rangle, \qquad (10.169)$$

and n copies of the state can be expressed as

$$\begin{aligned} |\psi\rangle_{AC}^{\otimes n} &= \sum_{x_1,\dots,x_n=0}^{d-1} \sqrt{p(x_1)\dots p(x_n)} \ |x_1x_2\dots x_n\rangle_{A^n} \otimes |x_1x_2\dots x_n\rangle_{C^n} \\ &= \sum_{\vec{x}} \sqrt{p(\vec{x})} \ |\vec{x}\rangle_{A^n} \otimes |\vec{x}\rangle_{C^n}, \end{aligned}$$
(10.170)

where $\sum_{\vec{x}} p(\vec{x}) = 1$. If Alice attempts to project onto the δ -typical subspace of C^n , she succeeds with high probability

$$P = \sum_{\delta - \text{typical } \vec{x}} p(\vec{x}) \ge 1 - \varepsilon$$
(10.171)

and when successful prepares the post-measurement state

$$|\Psi\rangle_{A^nC^n} = P^{-1/2} \sum_{\delta - \text{typical } \vec{x}} \sqrt{p(\vec{x})} \ |\vec{x}\rangle_{A^n} \otimes |\vec{x}\rangle_{C^n}, \tag{10.172}$$

such that

$$\langle \Psi | \psi^{\otimes n} \rangle = P^{-1/2} \sum_{\delta - \text{typical } \vec{x}} p(\vec{x}) = \sqrt{P} \ge \sqrt{1 - \varepsilon}.$$
 (10.173)

Since the typical subspace has dimension at most $2^{n(H(\boldsymbol{\rho})+\delta)}$, Alice can teleport the C^n half of $|\Psi\rangle$ to Bob with perfect fidelity using no more than $n(H(\boldsymbol{\rho}) + \delta)$ Bell pairs shared by Alice and Bob. The teleportation uses LOCC: Alice's entangled measurement, classical communication from Alice to Bob to convey the measurement outcome, and Bob's unitary transformation conditioned on the outcome. Finally, after the teleportation, Bob decompresses, so that Alice and Bob share a state which has high fidelity with $|\psi\rangle_{AB}^{\otimes n}$. This protocol demonstrates that the entanglement cost E_C of $|\psi\rangle$ is not more than $H(\boldsymbol{\rho}_A)$.

Now consider the distillable entanglement E_D . Suppose Alice and Bob share the state $|\psi\rangle_{AB}^{\otimes n}$. Since $|\psi\rangle_{AB}$ is, in general, a *partially* entangled state, the entanglement that Alice and Bob share is in a diluted form. They wish to *concentrate* their shared entanglement, squeezing it down to the smallest possible Hilbert space; that is, they want to convert it to maximally-entangled pairs. We will show that Alice and Bob can "distill" at least

$$k' = n(H(\boldsymbol{\rho}_A) - \delta) \tag{10.174}$$

Bell pairs from $|\psi\rangle_{AB}^{\otimes n},$ with high likelihood of success.

To illustrate the concentration of entanglement, imagine that Alice and Bob have n copies of the two-qubit state $|\psi\rangle$, which is

$$|\psi(p)\rangle = \sqrt{1-p} |00\rangle + \sqrt{p} |11\rangle, \qquad (10.175)$$

where $0 \le p \le 1$, when expressed in its Schmidt basis. That is, Alice and Bob share the state

$$|\psi(p)\rangle^{\otimes n} = (\sqrt{1-p} |00\rangle + \sqrt{p} |11\rangle)^{\otimes n}.$$
 (10.176)

When we expand this state in the $\{|0\rangle, |1\rangle\}$ basis, we find 2^n terms, in each of which Alice and Bob hold exactly the same binary string of length n.

Now suppose Alice (or Bob) performs a local measurement on her (his) n qubits, measuring the *total* spin along the *z*-axis

$$\boldsymbol{\sigma}_{3}^{(\text{total})} = \sum_{i=1}^{n} \boldsymbol{\sigma}_{3}^{(i)}.$$
(10.177)

Equivalently, the measurement determines the Hamming weight of Alice's n qubits, the number of $|1\rangle$'s in Alice's *n*-bit string; that is, the number of spins pointing up.

In the expansion of $|\psi(p)\rangle^{\otimes n}$ there are $\binom{n}{m}$ terms in which Alice's string has Hamming weight m, each occurring with the same amplitude: $(1-p)^{(n-m)/2} p^{m/2}$. Hence the probability that Alice's measurement finds Hamming weight m is

$$p(m) = \binom{n}{m} (1-p)^{n-m} p^m.$$
(10.178)

Furthermore, because Alice is careful not to acquire any additional information besides the Hamming weight when she conducts the measurement, by measuring the Hamming weight m she prepares a *uniform* superposition of all $\binom{n}{m}$ strings with m up spins. Because Alice and Bob have perfectly correlated strings, if Bob were to measure the Hamming weight of his qubits he would find the same outcome as Alice. Alternatively, Alice could report her outcome to Bob in a classical message, saving Bob the trouble of doing the measurement himself. Thus, Alice and Bob share a maximally entangled state

$$\sum_{i=1}^{D} |i\rangle_A \otimes |i\rangle_B, \tag{10.179}$$

where the sum runs over the $D = \binom{n}{m}$ strings with Hamming weight m.

For n large the binomial distribution $\{p(m)\}$ approaches a sharply peaked function of m with mean $\mu = np$ and variance $\sigma^2 = np(1-p)$. Hence the probability of a large deviation from the mean,

$$|m - np| = \Omega(n), \tag{10.180}$$

is $p = \exp(-\Omega(n))$. Using Stirling's approximation, it then follows that

$$2^{n(H(p)-o(1))} \le D \le 2^{n(H(p)+o(1))}.$$
(10.181)

with probability approaching one as $n \to \infty$, where $H(p) = -p \log_2 p - (1-p) \log_2(1-p)$ is the entropy function. Thus with high probability Alice and Bob share a maximally entangled state of Hilbert spaces \mathcal{H}_A and \mathcal{H}_B with $\dim(\mathcal{H}_A) = \dim(\mathcal{H}_B) = D$ and $\log_2 D \ge n(H(p) - \delta)$. In this sense Alice and Bob can distill $H(p) - \delta$ Bell pairs per copy of $|\psi\rangle_{AB}$.

Though the number m of up spins that Alice (or Bob) finds in her (his) measurement is typically close to np, it can fluctuate about this value. Sometimes Alice and Bob will be lucky, and then will manage to distill more than H(p) Bell pairs per copy of $|\psi(p)\rangle_{AB}$. But the probability of doing substantially better becomes negligible as $n \to \infty$.

The same idea applies to bipartite pure states in larger Hilbert spaces. If A and B are d-dimensional systems, then $|\psi\rangle_{AB}$ has the Schmidt decomposition

$$|\psi(X)\rangle_{AB} = \sum_{i=0}^{d-1} \sqrt{p(x)} |x\rangle_A \otimes |x\rangle_B, \qquad (10.182)$$

where X is the classical ensemble $\{x, p(x)\}$, and $H(\rho_A) = H(\rho_B) = H(X)$. The Schmidt decomposition of n copies of ψ is

$$\sum_{x_1, x_2, \dots, x_n = 0}^{d-1} \sqrt{p(x_1)p(x_2)\dots p(x_n)} |x_1 x_2 \dots x_n\rangle_{A^n} \otimes |x_1 x_2 \dots x_n\rangle_{B^n}.$$
 (10.183)

Now Alice (or Bob) can measure the total number of $|0\rangle$'s, the total number of $|1\rangle$'s, etc. in her (his) possession. If she finds $m_0|0\rangle$'s, $m_1|1\rangle$'s, etc., then her measurement prepares a maximally entangled state with Schmidt number

$$D(m_0, m_1, \dots, m_{d-1}) = \frac{n!}{m_0! m_1! \dots m_{d-1}!}$$
(10.184)

and this outcome occurs with probability

$$p(m) = D(m_0, m_1, \dots, m_{d-1}) p(0)^{m_0} p(1)^{m_1} \dots p(d-1)^{m_{d-1}}.$$
(10.185)

For n large, Alice will typically find $m_x \approx np(x)$, and again the probability of a large deviation is small, so that, from Stirling's approximation

$$2^{n(H(X)-o(1))} \le D \le 2^{n(H(X)+o(1))}$$
(10.186)

with high probability. Thus, asymptotically for $n \to \infty$, $n(H(\rho_A) - o(1))$ high-fidelity Bell pairs can be distilled from n copies of $|\psi\rangle$, establishing that $E_D(|\psi\rangle) \ge H(\rho_A)$, and therefore $E_D(|\psi\rangle) = E_C(|\psi\rangle) = E(|\psi\rangle)$.

This entanglement concentration protocol uses local operations but does not require any classical communication. When Alice and Bob do the same measurement they always get the same outcome, so there is no need for them to communicate. Classical communication really is necessary, though, to perform entanglement dilution. The protocol we described here, based on teleportation, requires two bits of classical one-way communication per Bell pair consumed; in a more clever protocol this can be reduced to $O(\sqrt{n})$ bits, but no further. Since the classical communication cost is sublinear in n, the number of bits of classical communication needed per copy of $|\psi\rangle$ becomes negligible in the limit $n \to \infty$.

10.5 Quantifying Mixed-State Entanglement

10.5.1 Asymptotic irreversibility under LOCC

The entanglement cost E_C and the distillable entanglement E_D are natural and operationally meaningful ways to quantify entanglement. It's quite satisfying to find that, because entanglement dilution and concentration are asymptotically reversible for pure states, these two measures of pure-state bipartite entanglement agree, and provide another operational role for the Von Neumann entropy of a marginal quantum state.

We can define E_C and E_D for bipartite mixed states just as we did for pure states, but the story is more complicated — when we prepare many copies of a mixed state shared by Alice and Bob, the dilution of Bell pairs is not in general reversible, even asymptotically, and the distillable entanglement can be strictly less than the entanglement cost, though it can never be larger. There are even bipartite mixed states with nonzero entanglement cost and zero distillable entanglement, a phenomenon called *bound entanglement*. This irreversibility is not shocking; any bipartite operation which maps many copies of the pure state $|\phi^+\rangle_{AB}$ to many copies of the mixed state ρ_{AB} necessarily discards some information to the environment, and we don't normally expect a process that forgets information to be reversible.

This separation between E_C and E_D raises the question, what is the preferred way to quantify the amount of entanglement when two parties share a mixed quantum state? The answer is, it depends. Many different measures of bipartite mixed-state entanglement have been proposed, each with its own distinctive advantages and disadvantages. Even though they do not always agree, both E_C and E_D are certainly valid measures. A further distinction can be made between the rate E_{D1} at which entanglement can be distilled with one-way communication between the parties, and the rate E_D with two-way communication. There are bipartite mixed states for which $E_D > E_{D1}$, and even states for which E_D is nonzero while E_{D1} is zero. In contrast to the pure-state case, we don't have nice formulas for the values of the various entanglement measures, though there are useful upper and lower bounds. We will derive a lower bound on E_{D1} in §10.8.2 (the hashing inequality).

There are certain properties that any reasonable measure of bipartite quantum entanglement should have. The most important is that it must not increase under local operations and classical communication, because quantum entanglement cannot be created by LOCC alone. A function on bipartite states that is nonincreasing under LOCC is called an *entanglement monotone*. Note that an entanglement monotone will also be *invariant* under local unitary operations $U_{AB} = U_A \otimes U_B$, for if U_{AB} can reduce the entanglement for any state, its inverse can increase entanglement.

A second important property is that a bipartite entanglement measure must *vanish* for separable states. Recall from Chapter 4 that a bipartite mixed state is separable if it can be expressed as a convex combination of product states,

$$\boldsymbol{\rho}_{AB} = \sum_{x} p(x) \ |\alpha(x)\rangle \langle \alpha(x)|_A \ \otimes \ |\beta(x)\rangle \langle \beta(x)|_B. \tag{10.187}$$

A separable state is not entangled, as it can be created using LOCC. Via classical communication, Alice and Bob can establish a shared source of randomness, the distribution $X = \{x, p(x)\}$. Then they may jointly sample from X; if the outcome is x, Alice prepares $|\alpha(x)\rangle$ while Bob prepares $|\beta(x)\rangle$.

A third desirable property for a bipartite entanglement measure is that it should agree with $E = E_C = E_D$ for bipartite pure states. Both the entanglement cost and the distillable entanglement respect all three of these properties.

We remark in passing that, despite the irreversibility of entanglement dilution under LOCC, there is a mathematically viable way to formulate a reversible theory of bipartite entanglement which applies even to mixed states. In this formulation, we allow Alice and Bob to perform arbitrary bipartite operations that are incapable of creating entanglement; these include LOCC as well as additional operations which cannot be realized using LOCC. In this framework, dilution and concentration of entanglement become asymptotically reversible even for mixed states, and a unique measure of entanglement can be formulated characterizing the optimal rate of conversion between copies of ρ_{AB} and Bell pairs using these non-entangling operations.

Irreversible bipartite entanglement theory under LOCC, and also the reversible theory under non-entangling bipartite operations, are both examples of *resource theories*. In the resource theory framework, one or more parties are able to perform some restricted class of operations, and they are capable of preparing a certain restricted class of states using these operations. In addition, the parties may also have access to *resource states*, which are outside the class they can prepare on their own. Using their restricted operations, they can transform resource states from one form to another, or consume resource states to perform operations beyond what they could achieve with their restricted operations alone. The name "resource state" conveys that such states are valuable because they may be consumed to do useful things.

In a two-party setting, where LOCC is allowed or more general non-entangling operations are allowed, bipartite entangled states may be regarded as a valuable resource. Resource theory also applies if the allowed operations are required to obey certain symmetries; then states breaking this symmetry become a resource. In thermodynamics, states deviating from thermal equilibrium are a resource. Entanglement theory, as a particularly well developed resource theory, provides guidance and tools which are broadly applicable to many different interesting situations.

10.5.2 Squashed entanglement

As an example of an alternative bipartite entanglement measure, consider the squashed entanglement E_{sq} , defined by

$$E_{\rm sq}(\boldsymbol{\rho}_{AB}) = \inf\left\{\frac{1}{2}I(A;B|C):\boldsymbol{\rho}_{AB} = \operatorname{tr}_{C}(\boldsymbol{\rho}_{ABC})\right\}$$
(10.188)

The squashed entanglement of ρ_{AB} is the greatest lower bound on the quantum conditional mutual information of all possible extensions of ρ_{AB} to a tripartite state ρ_{ABC} ; it can be shown to be an entanglement monotone. The locution "squashed" conveys that choosing an optimal conditioning system C squashes out the non-quantum correlations between A and B.

For pure states the extension is superfluous, so that

$$E_{\rm sq}(|\psi\rangle_{AB}) = \frac{1}{2}I(A;B) = H(A) = H(B) = E(|\psi\rangle_{AB}).$$
(10.189)

For a separable state, we may choose the extension

$$\boldsymbol{\rho}_{ABC} = \sum_{x} p(x) \ |\alpha(x)\rangle \langle \alpha(x)|_A \ \otimes \ |\beta(x)\rangle \langle \beta(x)|_B \ \otimes \ |x\rangle \langle x|_C.$$
(10.190)

where $\{|x\rangle_C\}$ is an orthonormal set; the state ρ_{ABC} has the block-diagonal form eq.(10.82) and hence I(A; B|C) = 0. Conversely, if ρ_{AB} has any extension ρ_{ABC} with I(A; B|C) = 0, then ρ_{ABC} has the form eq.(10.82) and therefore ρ_{AB} is separable.

 $E_{\rm sq}$ is difficult to compute, because the infimum is to be evaluated over all possible extensions, where the system C may have arbitrarily high dimension. This property also raises the logical possibility that there are nonseparable states for which the infimum vanishes; conceivably, though a nonseparable ρ_{AB} can have no finite-dimensional extension for which I(A; B|C) = 0, perhaps I(A; B|C) can approach zero as the dimension of C increases. Fortunately, though this is not easy to show, it turns out that $E_{\rm sq}$ is strictly positive for any nonseparable state. In this sense, then, it is a faithful entanglement measure, strictly positive if and only if the state is nonseparable.

One desirable property of E_{sq} , not shared by E_C and E_D , is its additivity on tensor products (Exercise 10.6),

$$E_{\rm sq}(\boldsymbol{\rho}_{AB} \otimes \boldsymbol{\rho}_{A'B'}) = E_{\rm sq}(\boldsymbol{\rho}_{AB}) + E_{\rm sq}(\boldsymbol{\rho}_{A'B'}). \tag{10.191}$$

Though, unlike E_C and E_D , squashed entanglement does not have an obvious operational

meaning, any additive entanglement monotone which matches E for bipartite pure states is bounded above and below by E_C and E_D respectively,

$$E_C \ge E_{\rm sq} \ge E_D. \tag{10.192}$$

10.5.3 Entanglement monogamy

Classical correlations are *polyamorous*; they can be shared among many parties. If Alice and Bob read the same newspaper, then they have information in common and become correlated. Nothing prevents Claire from reading the same newspaper; then Claire is just as strongly correlated with Alice and with Bob as Alice and Bob are with one another. Furthermore, David, Edith, and all their friends can read the newspaper and join the party as well.

Quantum correlations are not like that; they are harder to share. If Bob's state is pure, then the tripartite quantum state is a product $\rho_B \otimes \rho_{AC}$, and Bob is completely uncorrelated with Alice and Claire. If Bob's state is mixed, then he can be entangled with other parties. But if Bob is fully entangled with Alice (shares a pure state with Alice), then the state is a product $\rho_{AB} \otimes \rho_C$; Bob has used up all his ability to entangle by sharing with Alice, and Bob cannot be correlated with Claire at all. Conversely, if Bob shares a pure state with Claire, the state is $\rho_A \otimes \rho_{BC}$, and Bob is uncorrelated with Alice. Thus we say that quantum entanglement is *monogamous*.

Entanglement measures obey monogamy inequalities which reflect this tradeoff between Bob's entanglement with Alice and with Claire in a three-party state. Squashed entanglement, in particular, obeys a monogamy relation following easily from its definition, which was our primary motivation for introducing this quantity; we have

$$E_{\rm sq}(A;B) + E_{\rm sq}(A;C) \le E_{\rm sq}(A;BC).$$
 (10.193)

In particular, in the case of a pure tripartite state, $E_{sq} = H(A)$ is the (pure-state) entanglement shared between A and BC. The inequality is saturated if Alice's system is divided into subsystems A_1 and A_2 such that the tripartite pure state is

$$|\psi\rangle_{ABC} = |\psi_1\rangle_{A_1B} \otimes |\psi_2\rangle_{A_2C}.$$
(10.194)

In general, combining eq.(10.192) with eq.(10.193) yields

$$E_D(A; B) + E_D(A; C) \le E_C(A; BC);$$
 (10.195)

loosely speaking, the entanglement cost $E_C(A; BC)$ imposes a ceiling on Alice's ability to entangle with Bob and Claire individually, requiring her to trade in some distillable entanglement with Bob to increase her distillable entanglement with Claire.

To prove the monogamy relation eq.(10.193), we note that mutual information obeys a *chain rule* which is really just a restatement of the definition of conditional mutual information:

$$I(A; BC) = I(A; C) + I(A; B|C).$$
(10.196)

A similar equation follows directly from the definition if we condition on a fourth system D,

$$I(A; BC|D) = I(A; C|D) + I(A; B|CD).$$
(10.197)

Now, $E_{sq}(A; BC)$ is the infimum of I(A; BC|D) over all possible extensions of ρ_{ABC} to ρ_{ABCD} . But since ρ_{ABCD} is also an extension of ρ_{AB} and ρ_{AC} , we have

$$I(A; BC|D) \ge E_{sq}(A; C) + E_{sq}(A; B)$$
 (10.198)

for any such extension. Taking the infimum over all ρ_{ABCD} yields eq.(10.193).

A further aspect of monogamy arises when we consider extending a quantum state to more parties. We say that the bipartite state ρ_{AB} of systems A and B is k-extendable if there is a (k+1)-part state $\rho_{AB_1...B_k}$ whose marginal state on AB_j matches ρ_{AB} for each j = 1, 2, ...k, and such that $\rho_{AB_1...B_k}$ is invariant under permutations of the k systems $B_1, B_2...B_k$. Separable states are k-extendable for every k, and entangled pure states are not even 2-extendable. Every entangled mixed state fails to be k-extendable for some finite k, and we may regard the maximal value k_{max} for which such a symmetric extension exists as a rough measure of how entangled the state is — bipartite entangled states with larger and larger k_{max} are closer and closer to being separable.

10.6 Accessible Information

10.6.1 How much can we learn from a measurement?

Consider a game played by Alice and Bob. Alice prepares a quantum state drawn from the ensemble $\mathcal{E} = \{\rho(x), p(x)\}$ and sends the state to Bob. Bob knows this ensemble, but not the particular state that Alice chose to send. After receiving the state, Bob performs a POVM with elements $\{E(y)\} \equiv E$, hoping to find out as much as he can about what Alice sent. The conditional probability that Bob obtains outcome y if Alice sent $\rho(x)$ is $p(y|x) = \operatorname{tr}(E(y)\rho(x))$, and the joint distribution governing Alice's preparation and Bob's measurement is p(x, y) = p(y|x)p(x).

Before he measures, Bob's ignorance about Alice's state is quantified by H(X), the number of "bits per letter" needed to specify x; after he measures his ignorance is reduced to H(X|Y) = H(XY) - H(Y). The improvement in Bob's knowledge achieved by the measurement is Bob's *information gain*, the mutual information

$$I(X;Y) = H(X) - H(X|Y).$$
(10.199)

Bob's best strategy (his *optimal measurement*) maximizes this information gain. The best information gain Bob can achieve,

$$\operatorname{Acc}(\mathcal{E}) = \max_{\boldsymbol{E}} I(X;Y), \qquad (10.200)$$

is a property of the ensemble \mathcal{E} called the *accessible information* of \mathcal{E} .

If the states $\{\rho(x)\}$ are mutually orthogonal they are perfectly distinguishable. Bob can identify Alice's state with certainty by choosing E(x) to be the projector onto the support of $\rho(x)$; Then $p(y|x) = \delta_{x,y} = p(x|y)$, hence $H(X|Y) = \langle -\log p(x|y) \rangle = 0$ and $\operatorname{Acc}(\mathcal{E}) = H(X)$. Bob's task is more challenging if Alice's states are not orthogonal. Then no measurement will identify the state perfectly, so H(X|Y) is necessarily positive and $\operatorname{Acc}(\mathcal{E}) < H(X)$.

Though there is no simple general formula for the accessible information of an ensemble, we can derive a useful upper bound, called the *Holevo bound*. For the special case of an ensemble of pure states $\mathcal{E} = \{|\varphi(x)\rangle, p(x)\}$, the Holevo bound becomes

Acc(
$$\mathcal{E}$$
) $\leq H(\boldsymbol{\rho})$, where $\boldsymbol{\rho} = \sum_{x} p(x) |\varphi(x)\rangle \langle \varphi(x)|,$ (10.201)

Quantum Shannon Theory

and a sharper statement is possible for an ensemble of mixed states, as we will see. Since the entropy for a quantum system with dimension d can be no larger than $\log d$, the Holevo bound asserts that Alice, by sending n qubits to Bob $(d = 2^n)$ can convey no more than n bits of information. This is true even if Bob performs a sophisticated collective measurement on all the qubits at once, rather than measuring them one at a time.

Therefore, if Alice wants to convey classical information to Bob by sending qubits, she can do no better than treating the qubits as though they were classical, sending each qubit in one of the two orthogonal states $\{|0\rangle, |1\rangle\}$ to transmit one bit. This statement is not so obvious. Alice might try to stuff more classical information into a single qubit by sending a state chosen from a large alphabet of pure single-qubit signal states, distributed uniformly on the Bloch sphere. But the enlarged alphabet is to no avail, because as the number of possible signals increases the signals also become less distinguishable, and Bob is not able to extract the extra information Alice hoped to deposit in the qubit.

If we can send information more efficiently by using an alphabet of mutually orthogonal states, why should we be interested in the accessible information for an ensemble of non-orthogonal states? There are many possible reasons. Perhaps Alice finds it easier to send signals, like coherent states, which are imperfectly distinguishable rather than mutually orthogonal. Or perhaps Alice sends signals to Bob through a noisy channel, so that signals which are orthogonal when they enter the channel are imperfectly distinguishable by the time they reach Bob.

The accessible information game also arises when an experimental physicist tries to measure an unknown classical force using a quantum system as a probe. For example, to measure the z-component of a magnetic field, we may prepare a spin- $\frac{1}{2}$ particle pointing in the x-direction; the spin precesses for time t in the unknown field, producing an ensemble of possible final states (which will be an ensemble of mixed states if the initial preparation is imperfect, or if decoherence occurs during the experiment). The more information we can gain about the final state of the spin, the more accurately we can determine the value of the magnetic field.

10.6.2 Holevo bound

Recall that quantum mutual information obeys monotonicity — if a quantum channel maps B to B', then $I(A; B) \ge I(A; B')$. We derive the Holevo bound by applying monotonicity of mutual information to the accessible information game. We will suppose that Alice records her chosen state in a classical register X and Bob likewise records his measurement outcome in another register Y, so that Bob's information gain is the mutual information I(X; Y) of the two registers. After Alice's preparation of her system A, the joint state of XA is

$$\boldsymbol{\rho}_{XA} = \sum_{x} p(x) |x\rangle \langle x| \otimes \boldsymbol{\rho}(x).$$
(10.202)

Bob's measurement is a quantum channel mapping A to AY according to

$$\boldsymbol{\rho}(x) \mapsto \sum_{y} \boldsymbol{M}(y) \boldsymbol{\rho}(x) \boldsymbol{M}(y)^{\dagger} \otimes |y\rangle \langle y|, \qquad (10.203)$$

where $\boldsymbol{M}(y)^{\dagger}\boldsymbol{M}(y) = \boldsymbol{E}(y)$, yielding the state for XAY

$$\boldsymbol{\rho}_{XAY}' = \sum_{x} p(x) |x\rangle \langle x| \otimes \boldsymbol{M}(y) \boldsymbol{\rho}(x) \boldsymbol{M}(y)^{\dagger} \otimes |y\rangle \langle y|.$$
(10.204)

Now we have

$$I(X;Y)_{\rho'} \le I(X;AY)_{\rho'} \le I(X;A)_{\rho},$$
 (10.205)

where the subscript indicates the state in which the mutual information is evaluated; the first inequality uses strong subadditivity in the state ρ' , and the second uses monotonicity under the channel mapping ρ to ρ' .

The quantity I(X; A) is an intrinsic property of the ensemble \mathcal{E} ; it is denoted $\chi(\mathcal{E})$ and called the *Holevo chi* of the ensemble. We have shown that however Bob chooses his measurement his information gain is bounded above by the Holevo chi; therefore,

$$\operatorname{Acc}(\mathcal{E}) \le \chi(\mathcal{E}) := I(X; A)_{\rho}.$$
(10.206)

This is the Holevo bound.

Now let's calculate $I(X; A)_{\rho}$ explicitly. We note that

$$H(XA) = -\operatorname{tr}\left(\sum_{x} p(x)|x\rangle\langle x|\otimes\boldsymbol{\rho}(x)\log\left(\sum_{x'} p(x')|x'\rangle\langle x'|\otimes\boldsymbol{\rho}(x')\right)\right)$$
$$= -\sum_{x} \operatorname{tr} p(x)\boldsymbol{\rho}(x)\left(\log p(x) + \log \boldsymbol{\rho}(x)\right)$$
$$= H(X) + \sum_{x} p(x)H(\boldsymbol{\rho}(x)), \qquad (10.207)$$

and therefore

$$H(A|X) = H(XA) - H(X) = \sum_{x} p(x)H(\boldsymbol{\rho}(x)).$$
(10.208)

Using I(X; A) = H(A) - H(A|X), we then find

$$\chi(\mathcal{E}) = I(X; A) = H(\boldsymbol{\rho}_A) - \sum_x p(x) H(\boldsymbol{\rho}_A(x)) \equiv H(A)_{\mathcal{E}} - \langle H(A) \rangle_{\mathcal{E}}$$
(10.209)

For an ensemble of pure states, χ is just the entropy of the density operator arising from the ensemble, but for an ensemble \mathcal{E} of mixed states it is a strictly smaller quantity – the difference between the entropy $H(\rho_{\mathcal{E}})$ of the convex sum of signal states and the convex sum $\langle H \rangle_{\mathcal{E}}$ of the signal state entropies; this difference is always nonnegative because of the concavity of the entropy function (or because mutual information is nonnegative).

10.6.3 Monotonicity of Holevo χ

Since Holevo χ is the mutual information I(X; A) of the classical register X and the quantum system A, the monotonicity of mutual information also implies the monotonicity of χ . If $\mathcal{N} : A \to A'$ is a quantum channel, then $I(X; A') \leq I(X; A)$ and therefore

$$\chi(\mathcal{E}') \le \chi(\mathcal{E}),\tag{10.210}$$

where

$$\mathcal{E} = \{\boldsymbol{\rho}(x)\}, p(x)\} \quad \text{and} \quad \mathcal{E}' = \{\boldsymbol{\rho}'(x) = \mathcal{N}(\boldsymbol{\rho}(x)), p(x)\}.$$
(10.211)

Quantum Shannon Theory

A channel cannot increase the Holevo χ of an ensemble.

Its monotonicity provides a further indication that $\chi(\mathcal{E})$ is a useful measure of the information encoded in an ensemble of quantum states; the decoherence described by a quantum channel can reduce this quantity, but never increases it. In contrast, the Von Neumann entropy may either increase or decrease under the action of a channel. Mapping pure states to mixed states can increase H, but a channel might instead map the mixed states in an ensemble to a fixed pure state $|0\rangle\langle 0|$, decreasing H and improving the purity of each signal state, but without improving the distinguishability of the states.

We discussed the asymptotic limit $H(\rho)$ on quantum compression per letter in §10.3.2. There we considered unitary decoding; invoking the monotonicity of Holevo χ clarifies why more general decoders cannot do better. Suppose we compress and decompress the ensemble $\mathcal{E}^{\otimes n}$ using an encoder \mathcal{N}_e and a decoder \mathcal{N}_d , where both maps are quantum channels:

$$\mathcal{E}^{\otimes n} \xrightarrow{\mathcal{N}_e} \tilde{\mathcal{E}}^{(n)} \xrightarrow{\mathcal{N}_d} \tilde{\mathcal{E}}^{\prime(n)} \approx \mathcal{E}^{\otimes n}$$
(10.212)

The Holevo χ of the input pure-state product ensemble is additive, $\chi(\mathcal{E}^{\otimes n}) = H(\boldsymbol{\rho}^{\otimes n}) = nH(\boldsymbol{\rho})$, and χ of a *d*-dimensional system is no larger than $\log_2 d$; therefore if the ensemble $\tilde{\mathcal{E}}^{(n)}$ is compressed to q qubits per letter, then because of the monotonicity of χ the decompressed ensemble $\tilde{\mathcal{E}}'^{(n)}$ has Holevo chi per letter $\frac{1}{n}\chi(\tilde{\mathcal{E}}'^{(n)}) \leq q$. If the decompressed output ensemble has high fidelity with the input ensemble, its χ per letter should nearly match the χ per letter of the input ensemble, hence

$$q \ge \frac{1}{n}\chi(\tilde{\mathcal{E}}'^{(n)}) \ge H(\boldsymbol{\rho}) - \delta \tag{10.213}$$

for any positive δ and sufficiently large n. We conclude that high-fidelity compression to fewer than $H(\rho)$ qubits per letter is impossible asymptotically, even when the compression and decompression maps are arbitrary channels.

10.6.4 Improved distinguishability through coding: an example

To better acquaint ourselves with the concept of accessible information, let's consider a single-qubit example. Alice prepares one of the three possible pure states

$$\begin{aligned} |\varphi_1\rangle &= |\uparrow_{\hat{n}_1}\rangle = \begin{pmatrix} 1\\0 \end{pmatrix}, \\ |\varphi_2\rangle &= |\uparrow_{\hat{n}_2}\rangle = \begin{pmatrix} -\frac{1}{2}\\\frac{\sqrt{3}}{2} \end{pmatrix}, \\ |\varphi_3\rangle &= |\uparrow_{\hat{n}_3}\rangle = \begin{pmatrix} -\frac{1}{2}\\-\frac{\sqrt{3}}{2} \end{pmatrix}; \end{aligned}$$
(10.214)

a spin- $\frac{1}{2}$ object points in one of three directions that are symmetrically distributed in the *xz*-plane. Each state has *a priori* probability $\frac{1}{3}$. Evidently, Alice's signal states are nonorthogonal:

$$\langle \varphi_1 | \varphi_2 \rangle = \langle \varphi_1 | \varphi_3 \rangle = \langle \varphi_2 | \varphi_3 \rangle = -\frac{1}{2}.$$
 (10.215)

Bob's task is to find out as much as he can about what Alice prepared by making a

suitable measurement. The density matrix of Alice's ensemble is

$$\boldsymbol{\rho} = \frac{1}{3} (|\varphi_1\rangle\langle\varphi_1| + |\varphi_2\rangle\langle\varphi_3| + |\varphi_3\rangle\langle\varphi_3|) = \frac{1}{2} \boldsymbol{I}, \qquad (10.216)$$

which has $H(\rho) = 1$. Therefore, the Holevo bound tells us that the mutual information of Alice's preparation and Bob's measurement outcome cannot exceed 1 bit.

In fact, though, the accessible information is considerably less than the one bit allowed by the Holevo bound. In this case, Alice's ensemble has enough symmetry that it is not hard to guess the optimal measurement. Bob may choose a POVM with three outcomes, where

$$\boldsymbol{E}_{a} = \frac{2}{3} (\boldsymbol{I} - |\varphi_{a}\rangle\langle\varphi_{a}|), \quad a = 1, 2, 3;$$
(10.217)

we see that

$$p(a|b) = \langle \varphi_b | \boldsymbol{E}_a | \varphi_b \rangle = \begin{cases} 0 & a = b, \\ \frac{1}{2} & a \neq b. \end{cases}$$
(10.218)

The measurement outcome *a excludes* the possibility that Alice prepared *a*, but leaves equal *a posteriori* probabilities $\left(p = \frac{1}{2}\right)$ for the other two states. Bob's information gain is

$$I = H(X) - H(X|Y) = \log_2 3 - 1 = .58496.$$
(10.219)

To show that this measurement is really optimal, we may appeal to a variation on a theorem of Davies, which assures us that an optimal POVM can be chosen with three E_a 's that share the same three-fold symmetry as the three states in the input ensemble. This result restricts the possible POVM's enough so that we can check that eq. (10.217) is optimal with an explicit calculation. Hence we have found that the ensemble $\mathcal{E} = \{|\varphi_a\rangle, p_a = \frac{1}{3}\}$ has accessible information.

$$Acc(\mathcal{E}) = \log_2\left(\frac{3}{2}\right) = .58496...$$
 (10.220)

The Holevo bound is not saturated.

Now suppose that Alice has enough cash so that she can afford to send two qubits to Bob, where again each qubit is drawn from the ensemble \mathcal{E} . The obvious thing for Alice to do is prepare one of the *nine* states

$$|\varphi_a\rangle \otimes |\varphi_b\rangle, \quad a, b = 1, 2, 3,$$
 (10.221)

each with $p_{ab} = 1/9$. Then Bob's best strategy is to perform the POVM eq. (10.217) on each of the two qubits, achieving a mutual information of .58496 bits per qubit, as before.

But, determined to do better, Alice and Bob decide on a different strategy. Alice will prepare one of *three* two-qubit states

$$|\Phi_a\rangle = |\varphi_a\rangle \otimes |\varphi_a\rangle, \quad a = 1, 2, 3,$$
 (10.222)

each occurring with *a priori* probability $p_a = 1/3$. Considered one-qubit at a time, Alice's choice is governed by the ensemble \mathcal{E} , but now her two qubits have (classical) correlations – both are prepared the same way.

The three $|\Phi_a\rangle$'s are linearly independent, and so span a three-dimensional subspace

of the four-dimensional two-qubit Hilbert space. In Exercise 10.4, you will show that the density operator

$$\boldsymbol{\rho} = \frac{1}{3} \left(\sum_{a=1}^{3} |\Phi_a\rangle \langle \Phi_a| \right), \qquad (10.223)$$

has the nonzero eigenvalues 1/2, 1/4, 1/4, so that

$$H(\boldsymbol{\rho}) = -\frac{1}{2}\log_2\frac{1}{2} - 2\left(\frac{1}{4}\log_2\frac{1}{4}\right) = \frac{3}{2}.$$
 (10.224)

The Holevo bound requires that the accessible information *per qubit* is no more than 3/4 bit, which is at least consistent with the possibility that we can exceed the .58496 bits per qubit attained by the nine-state method.

Naively, it may seem that Alice won't be able to convey as much classical information to Bob, if she chooses to send one of only three possible states instead of nine. But on further reflection, this conclusion is not obvious. True, Alice has fewer signals to choose from, but the signals are *more distinguishable*; we have

$$\langle \Phi_a | \Phi_b \rangle = \frac{1}{4}, \quad a \neq b,$$
 (10.225)

instead of eq. (10.215). It is up to Bob to exploit this improved distinguishability in his choice of measurement. In particular, Bob will find it advantageous to perform *collective* measurements on the two qubits instead of measuring them one at a time.

It is no longer obvious what Bob's optimal measurement will be. But Bob can invoke a general procedure that, while not guaranteed optimal, is usually at least pretty good. We'll call the POVM constructed by this procedure a "pretty good measurement" (or PGM).

Consider some collection of vectors $|\tilde{\Phi}_a\rangle$ that are not assumed to be orthogonal or normalized. We want to devise a POVM that can distinguish these vectors reasonably well. Let us first construct

$$\boldsymbol{G} = \sum_{a} |\tilde{\Phi}_{a}\rangle \langle \tilde{\Phi}_{a}|; \qquad (10.226)$$

This is a positive operator on the space spanned by the $|\tilde{\Phi}_a\rangle$'s. Therefore, on that subspace, \boldsymbol{G} has an inverse, \boldsymbol{G}^{-1} and that inverse has a positive square root $\boldsymbol{G}^{-1/2}$. Now we define

$$\boldsymbol{E}_{a} = \boldsymbol{G}^{-1/2} |\tilde{\Phi}_{a}\rangle \langle \tilde{\Phi}_{a} | \boldsymbol{G}^{-1/2}, \qquad (10.227)$$

and we see that

$$\sum_{a} \boldsymbol{E}_{a} = \boldsymbol{G}^{-1/2} \left(\sum_{a} |\tilde{\Phi}_{a}\rangle \langle \tilde{\Phi}_{a}| \right) \boldsymbol{G}^{-1/2}$$
$$= \boldsymbol{G}^{-1/2} \boldsymbol{G} \boldsymbol{G}^{-1/2} = \boldsymbol{I}, \qquad (10.228)$$

on the span of the $|\tilde{\Phi}_a\rangle$'s. If necessary, we can augment these E_a 's with one more positive operator, the projection E_0 onto the orthogonal complement of the span of the $|\tilde{\Phi}_a\rangle$'s, and so construct a POVM. This POVM is the PGM associated with the vectors $|\tilde{\Phi}_a\rangle$.

In the special case where the $|\Phi_a\rangle$'s are orthogonal,

$$|\tilde{\Phi}_a\rangle = \sqrt{\lambda_a} |\phi_a\rangle,$$
 (10.229)

(where the $|\Phi_a\rangle$'s are orthonormal), we have

$$\boldsymbol{E}_{a} = \sum_{a,b,c} (|\phi_{b}\rangle \lambda_{b}^{-1/2} \langle \phi_{b}|) (|\phi_{a}\rangle \lambda_{a} \langle \phi_{a}|) (|\phi_{c}\rangle \lambda_{c}^{-1/2} \langle \phi_{c}|)$$
$$= |\phi_{a}\rangle \langle \phi_{a}|; \qquad (10.230)$$

this is the orthogonal measurement that perfectly distinguishes the $|\Phi_a\rangle$'s and so clearly is optimal. If the $|\tilde{\Phi}_a\rangle$'s are linearly independent but not orthogonal, then the PGM is again an orthogonal measurement (because *n* one-dimensional operators in an *n*dimensional space can constitute a POVM only if mutually orthogonal — see Exercise 3.11), but in that case the measurement may not be optimal.

In Exercise 10.4, you'll construct the PGM for the vectors $|\Phi_a\rangle$ in eq. (10.222), and you'll show that

$$p(a|a) = \langle \Phi_a | \mathbf{E}_a | \Phi_a \rangle = \frac{1}{3} \left(1 + \frac{1}{\sqrt{2}} \right)^2 = .971405$$
$$p(b|a) = \langle \Phi_a | \mathbf{E}_b | \Phi_a \rangle = \frac{1}{6} \left(1 - \frac{1}{\sqrt{2}} \right)^2 = .0142977, \quad (10.231)$$

(for $b \neq a$). It follows that the conditional entropy of the input is

$$H(X|Y) = .215893, (10.232)$$

and since $H(X) = \log_2 3 = 1.58496$, the information gain is

$$I(X;Y) = H(X) - H(X|Y) = 1.36907, (10.233)$$

a mutual information of .684535 bits per qubit. Thus, the improved distinguishability of Alice's signals has indeed paid off – we have exceeded the .58496 bits that can be extracted from a single qubit. We still didn't saturate the Holevo bound ($I \leq 1.5$ in this case), but we came a lot closer than before.

This example, first described by Peres and Wootters, teaches some useful lessons. First, Alice is able to convey more information to Bob by "pruning" her set of codewords. She is better off choosing among fewer signals that are more distinguishable than more signals that are less distinguishable. An alphabet of three letters encodes more than an alphabet of nine letters.

Second, Bob is able to read more of the information if he performs a collective measurement instead of measuring each qubit separately. His optimal orthogonal measurement projects Alice's signal onto a basis of *entangled* states.

10.6.5 Classical capacity of a quantum channel

This example illustrates how coding and collective measurement can enhance accessible information, but while using the code narrowed the gap between the accessible information and the Holevo chi of the ensemble, it did not close the gap completely. As is often the case in information theory, we can characterize the accessible information more precisely by considering an asymptotic i.i.d. setting. To be specific, we'll consider the task of sending classical information reliably through a noisy quantum channel $\mathcal{N}^{A\to B}$.

An ensemble of input signal states $\mathcal{E} = \{\rho(x), p(x)\}$ prepared by Alice is mapped by the channel to an ensemble of output signals $\mathcal{E}' = \{\mathcal{N}(\rho(x)), p(x)\}$. If Bob measures the output his information gain

$$\operatorname{Acc}(\mathcal{E}') \le I(X; B) = \chi(\mathcal{E}'). \tag{10.234}$$

is bounded above by the Holevo chi of the output ensemble \mathcal{E}' . To convey as much information through the channel as possible, Alice and Bob may choose the input ensemble \mathcal{E} that maximizes the Holevo chi of the output ensemble \mathcal{E}' . The maximum value

$$\chi(\mathcal{N}) := \max_{\mathcal{E}} \chi(\mathcal{E}') = \max_{\mathcal{E}} I(X; B), \qquad (10.235)$$

of $\chi(\mathcal{E}')$ is a property of the channel, which we will call the Holevo chi of \mathcal{N} .

As we've seen, Bob's actual optimal information gain in this *single-shot* setting may fall short of $\chi(\mathcal{E}')$ in general. But instead of using the channel just once, suppose that Alice and Bob use the channel $n \gg 1$ times, where Alice sends signal states chosen from a code, and Bob performs an optimal measurement to decode the signals he receives. Then an information gain of $\chi(\mathcal{N})$ bits per letter really can be achieved asymptotically as $n \to \infty$.

Let's denote Alice's ensemble of encoded *n*-letter signal states by $\tilde{\mathcal{E}}^{(n)}$, denote the ensemble of classical labels carried by the signals by \tilde{X}^n , and denote Bob's ensemble of measurement outcomes by \tilde{Y}^n . Let's say that the code has rate R if Alice may choose from among 2^{nR} possible signals to send. If classical information can be sent through the channel with rate R - o(1) such that Bob can decode the signal with negligible error probability as $n \to \infty$, then we say the rate R is *achievable*. The classical capacity $C(\mathcal{N})$ of the quantum channel $\mathcal{N}^{\mathcal{A}\to\mathcal{B}}$ is the supremum of all achievable rates.

Just as in our discussion of the capacity of a classical channel in §10.1.4, the conditional entropy per letter $\frac{1}{n}H(\tilde{X}^n|\tilde{Y}^n)$ approaches zero as $n \to \infty$ if the error probability is asymptotically negligible; therefore

$$R \leq \frac{1}{n} \left(I(\tilde{X}^n; \tilde{Y}^n) + o(1) \right)$$

$$\leq \frac{1}{n} \left(\max_{\mathcal{E}^{(n)}} I(X^n; B^n) + o(1) \right) = \frac{1}{n} \left(\chi(\mathcal{N}^{\otimes n}) + o(1) \right), \qquad (10.236)$$

where we obtain the first inequality as in eq. (10.47) and the second inequality by invoking the Holevo bound, optimized over all possible *n*-letter input ensembles. We therefore infer that

$$C(\mathcal{N}) \le \lim_{n \to \infty} \frac{1}{n} \chi\left(\mathcal{N}^{\otimes n}\right); \qquad (10.237)$$

the classical capacity is bounded above by the asymptotic Holevo χ per letter of the product channel $\mathcal{N}^{\otimes n}$.

In fact this upper bound is actually an achievable rate, and hence equal to the classical capacity $C(\mathcal{N})$. However, this formula for the classical capacity is not very useful as it stands, because it requires that we optimize the Holevo χ over message ensembles of arbitrary length; we say that the formula for capacity is *regularized* if, as in this case, it involves taking a limit in which the number of channel tends to infinity. It would be far preferable to reduce our expression for $C(\mathcal{N})$ to a *single-letter formula* involving just one use of the channel. In the case of a classical channel, the reduction of the regularized expression to a single-letter formula was possible, because the conditional entropy for n uses of the channel is additive as in eq.(10.44).
For quantum channels the situation is more complicated, as channels are known to exist such that the Holevo χ is strictly superadditive:

$$\chi\left(\mathcal{N}_1\otimes\mathcal{N}_2\right)>\chi\left(\mathcal{N}_1\right)+\chi\left(\mathcal{N}_2\right).\tag{10.238}$$

Therefore, at least for some channels, we are stuck with the not-very-useful regularized formula for the classical capacity. But we can obtain a single-letter formula for the optimal achievable communication rate if we put a restriction on the code used by Alice and Bob. In general, Alice is entitled to choose input codewords which are entangled across the many uses of the channel, and when such entangled codes are permitted the computation of the classical channel capacity may be difficult. But suppose we demand that all of Alice's codewords are product states. With that proviso the Holevo chi becomes subadditive, and we may express the optimal rate as

$$C_1(\mathcal{N}) = \chi(\mathcal{N}). \tag{10.239}$$

 $C_1(\mathcal{N})$ is called the *product-state capacity* of the channel.

Let's verify the subadditivity of χ for product-state codes. The product channel $\mathcal{N}^{\otimes n}$ maps product states to product states; hence if Alice's input signals are product states then so are Bob's output signals, and we can express Bob's *n*-letter ensemble as

$$\mathcal{E}^{(n)} = \{ \boldsymbol{\rho}(x_1) \otimes \boldsymbol{\rho}(x_2) \otimes \cdots \otimes \boldsymbol{\rho}(x_n), \ p(x_1 x_2 \dots x_n) \},$$
(10.240)

which has Holevo χ

$$\chi(\mathcal{E}^{(n)}) = I(X^n; B^n) = H(B^n) - H(B^n | X^n).$$
(10.241)

While the Von Neumann entropy is subadditive,

$$H(B^{n}) = \sum_{i=1}^{n} H(B_{i}); \qquad (10.242)$$

the (negated) conditional entropy

$$-H(B^{n}|X^{n}) = -\sum_{\vec{x}} p(\vec{x}) \ H(\boldsymbol{\rho}(\vec{x}))$$
(10.243)

(see eq.(10.209)) is not subadditive in general. But for the product-state ensemble eq.(10.240), since the entropy of a product is additive, we have

$$H(B^{n}|X^{n}) = \sum_{x_{1}, x_{2}, \dots, x_{n}} p(x_{1}x_{2}, \dots, x_{n}) \left(\sum_{i=1}^{n} H\left(\boldsymbol{\rho}(x_{i})\right)\right)$$
$$= \sum_{i=1}^{n} p_{i}(x_{i})H(\boldsymbol{\rho}(x_{i})) = \sum_{i=1}^{n} H(B_{i}|X_{i})$$
(10.244)

where $x_i = \{x_i, p_i(x_i)\}$ is the marginal probability distribution for the *i*th letter. Eq.(10.244) is a quantum analog of eq.(10.44), which holds for product-state ensembles but not in general for entangled ensembles. Combining eq.(10.241), (10.242), (10.244), we have

$$I(X^{n}; B^{n}) \leq \sum_{i=1}^{n} \left(H(B_{i}) - H(B_{i}|X_{i}) \right) = \sum_{i} I(X_{i}; B_{i}) \leq n\chi(\mathcal{N}).$$
(10.245)

Therefore the Holevo χ of a channel is subadditive when restricted to product-state codewords, as we wanted to show.

We won't give a careful argument here that $C_1(\mathcal{N})$ is an asymptotically achievable rate using product-state codewords; we'll just give a rough sketch of the idea. We demonstrate achievability with a random coding argument similar to Shannon's. Alice fixes an input ensemble $\mathcal{E} = \{ \boldsymbol{\rho}(x), p(x) \}$, and samples from the product ensemble $\mathcal{E}^{\otimes n}$ to generate a codeword; that is, the codeword

$$\boldsymbol{\rho}(\vec{x}) = \boldsymbol{\rho}(x_1) \otimes \boldsymbol{\rho}(x_2) \otimes \cdots \otimes \boldsymbol{\rho}(x_n)$$
(10.246)

is selected with probability $p(\vec{x}) = p(x_1)p(x_2) \dots p(x_n)$. (In fact Alice should choose each $\rho(\vec{x})$ to be pure to optimize the communication rate.) This codeword is sent via n uses of the channel \mathcal{N} , and Bob receives the product state

$$\mathcal{N}^{\otimes n}\left(\boldsymbol{\rho}(\vec{x})\right) = \mathcal{N}(\boldsymbol{\rho}(x_1)) \otimes \mathcal{N}(\boldsymbol{\rho}(x_2)) \otimes \cdots \otimes \mathcal{N}(\boldsymbol{\rho}(x_n)).$$
(10.247)

Averaged over codewords, the joint state of Alice's classical register X^n and Bob's system B^n is

$$\boldsymbol{\rho}_{X^n B^n} = \sum_{\vec{x}} p(\vec{x}) \ |\vec{x}\rangle \langle \vec{x}| \otimes \mathcal{N}^{\otimes n}(\boldsymbol{\rho}(\vec{x})).$$
(10.248)

To decode, Bob performs a POVM designed to distinguish the codewords effectively; a variant of the pretty good measurement described in §10.6.4 does the job well enough. The state Bob receives is mostly supported on a typical subspace with dimension $2^{n(H(B)+o(1))}$, and for each typical codeword that Alice sends, what Bob receives is mostly supported on a much smaller typical subspace with dimension $2^{n(H(B|X)+o(1))}$. The key point is that ratio of these spaces is exponential in the mutual information of X and B:

$$\frac{2^{n(H(B|X)+o(1))}}{2^{n(H(B)-o(1))}} = 2^{-n(I(X;B)-o(1))}$$
(10.249)

Each of Bob's POVM elements has support on the typical subspace arising from a particular one of Alice's codewords. The probability that any codeword is mapped purely by accident to the decoding subspace of a different codeword is suppressed by the ratio eq.(10.249). Therefore, the probability of a decoding error remains small even when there are 2^{nR} codewords to distinguish, for R = I(X; B) - o(1).

We complete the argument with standard Shannonisms. Since the probability of decoding error is small when we average over codes, it must also be small, averaged over codewords, for a particular sequence of codes. Then by pruning half of the codewords, reducing the rate by a negligible amount, we can ensure that the decoding errors are improbable for every codeword in the code. Therefore I(X; B) is an achievable rate for classical communication. Optimizing over all product-state input ensembles, we obtain eq.(10.239).

To turn this into an honest argument, we would need to specify Bob's decoding measurement more explicitly and do a careful error analysis. This gets a bit technical, so we'll skip the details. Somewhat surprisingly, though, it turns out to be easier to prove capacity theorems when quantum channels are used for other tasks besides sending classical information. We'll turn to that in §10.7.

10.6.6 Entanglement-breaking channels

Though Holevo chi is superadditive for some quantum channels, there are classes of channels for which chi is additive, and for any such channel \mathcal{N} the classical capacity

is $C = \chi(\mathcal{N})$ without any need for regularization. For example, consider *entanglement*breaking channels. We say that $\mathcal{N}^{A\to B}$ is entanglement breaking if for any input state ρ_{RA} , $I \otimes \mathcal{N}(\rho_{RA})$ is a separable state on RA — the action of \mathcal{N} on A always breaks its entanglement with R. We claim that if \mathcal{N}_1 is entanglement breaking, and \mathcal{N}_2 is an arbitrary channel, then

$$\chi\left(\mathcal{N}_1 \otimes \mathcal{N}_2\right) \le \chi(\mathcal{N}_1) + \chi(\mathcal{N}_2). \tag{10.250}$$

To bound the chi of the product channel, consider an input ensemble

$$\boldsymbol{\rho}_{XA_1A_2} = \sum_{x} p(x) |x\rangle \langle x| \otimes \boldsymbol{\rho}(x)_{A_1A_2}.$$
(10.251)

Because \mathcal{N}_1 is entanglement breaking, $\rho(x)_{A_1A_2}$ is mapped by the product channel to a separable state:

$$\mathcal{N}_1 \otimes \mathcal{N}_2 : \boldsymbol{\rho}(x)_{A_1 A_2} \mapsto \sum_y p(y|x) \ \boldsymbol{\sigma}(x, y)_{B_1} \otimes \boldsymbol{\tau}(x, y)_{B_2}.$$
(10.252)

Now $\chi(\mathcal{N}_1 \otimes \mathcal{N}_2)$ is the maximum of $I(X; B_1B_2)_{\rho'}$, evaluated in the state

$$\boldsymbol{\rho}_{XB_1B_2}' = \sum_{x,y} p(x)p(y|x)|x\rangle\langle x|\otimes\boldsymbol{\sigma}(x,y)_{B_1}\otimes\boldsymbol{\tau}(x,y)_{B_2}$$
(10.253)

which may be regarded as the marginal state (after tracing out Y) of

$$\tilde{\boldsymbol{\rho}}'_{XYB_1B_2} = \sum_{x,y} p(x,y) |x,y\rangle \langle x,y| \otimes \boldsymbol{\sigma}(x,y)_{B_1} \otimes \tilde{\boldsymbol{\tau}}(x,y)_{B_2}$$
(10.254)

Because $\tilde{\rho}'$ becomes a product state when conditioned on (x, y), it satisfies

$$H(B_1B_2|XY) = H(B_1|XY) + H(B_2|XY), \qquad (10.255)$$

and from the subadditivity and strong subadditivity of entropy we have

$$I(X; B_1B_2) \le I(XY; B_1B_2) = H(B_1B_2) - H(B_1B_2|XY)$$

$$\le H(B_1) + H(B_2) - H(B_1|XY) - H(B_2|XY)$$

$$= I(XY; B_1) + I(XY; B_2).$$
(10.256)

The right-hand side is bounded above by $\chi(\mathcal{N}_1) + \chi(\mathcal{N}_2)$, and maximizing the left-hand side yields eq.(10.250).

An example of an entanglement-breaking channel is a *classical-quantum channel*, also called a c-q *channel*, which acts according to

$$\mathcal{N}^{A \to B} : \boldsymbol{\rho}_A \mapsto \sum_x \langle x | \boldsymbol{\rho}_A | x \rangle \sigma(x)_B, \qquad (10.257)$$

where $\{|x\rangle\}$ is an orthonormal basis. In effect, the channel performs a complete orthogonal measurement on the input state and then prepares an output state conditioned on the measurement outcome. The measurement breaks the entanglement between system A and any other system with which it was initially entangled. Therefore, c-q channels are entanglement breaking and have additive Holevo chi.

10.7 Quantum Channel Capacities and Decoupling

10.7.1 Coherent information and the quantum channel capacity

As we have already emphasized, it's marvelous that the capacity for a classical channel can be expressed in terms of the optimal correlation between input and output for a *single use* of the channel,

$$C := \max_{X} I(X;Y).$$
 (10.258)

Another pleasing feature of this formula is its *robustness*. For example, the capacity does not increase if we allow the sender and receiver to share randomness, or if we allow feedback from receiver to sender. But for quantum channels the story is more complicated. We've seen already that no simple single-letter formula is known for the classical capacity of a quantum channel, if we allow entanglement among the channel inputs, and we'll soon see that the same is true for the quantum capacity. In addition, it turns out that entanglement shared between sender and receiver can boost the classical and quantum capacities of some channels, and so can "backward" communication from receiver to sender. There are a variety of different notions of capacity for quantum channels, all reasonably natural, and all with different achievable rates.

While Shannon's theory of classical communication over noisy classical channels is pristine and elegant, the same cannot be said for the theory of communication over noisy quantum channels, at least not in its current state. It's still a work in progress. Perhaps some day another genius like Shannon will construct a beautiful theory of quantum capacities. For now, at least there are a lot of interesting things we can say about achievable rates. Furthermore, the tools that have been developed to address questions about quantum capacities have other applications beyond communication theory.

The most direct analog of the classical capacity of a classical channel is the quantum capacity of a quantum channel, unassisted by shared entanglement or feedback. The quantum channel $\mathcal{N}^{A\to B}$ is a TPCP map from \mathcal{H}_A to \mathcal{H}_B , and Alice is to use the channel *n* times to convey a quantum state to Bob with high fidelity. She prepares her state $|\psi\rangle$ in a code subspace

$$\mathcal{H}^{(n)} \subseteq \mathcal{H}_A^{\otimes n} \tag{10.259}$$

and sends it to Bob, who applies a decoding map, attempting to recover $|\psi\rangle$. The rate R of the code is the number of encoded qubits sent per channel use,

$$R = \log_2 \dim \left(\mathcal{H}^{(n)} \right), \qquad (10.260)$$

We say that the rate R is *achievable* if there is a sequence of codes with increasing n such that for any $\varepsilon, \delta > 0$ and for sufficiently large n the rate is at least $R - \delta$ and Bob's recovered state ρ has fidelity $F = \langle \psi | \rho | \psi \rangle \geq 1 - \varepsilon$. The quantum channel capacity $Q(\mathcal{N})$ is the supremum of all achievable rates.

There is a regularized formula for $Q(\mathcal{N})$. To understand the formula we first need to recall that any channel $\mathcal{N}^{A\to B}$ has an isometric Stinespring dilation $U^{A\to BE}$ where E is the channel's "environment." Furthermore, any input density operator ρ_A has a purification; if we introduce a *reference system* R, for any ρ_A there is a pure state ψ_{RA} such that $\rho_A = \operatorname{tr}_R(|\psi\rangle\langle\psi|)$. (I will sometimes use ψ rather than the Dirac ket $|\psi\rangle$ to denote a pure state vector, when the context makes the meaning clear and the ket notation seems unnecessarily cumbersome.) Applying the channel's dilation to ψ_{RA} , we obtain an output pure state ϕ_{RBE} , which we represent graphically as:



We then define the one-shot quantum capacity of the channel \mathcal{N} by

$$Q_1(\mathcal{N}) := \max_A \left(-H(R|B)_{\phi_{RBE}} \right).$$
(10.261)

Here the maximum is taken over all possible input density operators $\{\rho_A\}$, and H(R|B) is the quantum conditional entropy

$$H(R|B) = H(RB) - H(B) = H(E) - H(B),$$
(10.262)

where in the last equality we used H(RB) = H(E) in a pure state of RBE. The quantity -H(R|B) has such a pivotal role in quantum communication theory that it deserves to have its own special name. We call it the *coherent information* from R to B and denote it

$$I_c(R\rangle B)_{\phi} = -H(R|B)_{\phi} = H(B)_{\phi} - H(E)_{\phi}.$$
(10.263)

This quantity does not depend on how the purification ϕ of the density operator ρ_A is chosen; any one purification can be obtained from any other by a unitary transformation acting on R alone, which does not alter H(B) or H(E). Indeed, since the expression H(B) - H(E) only depends on the marginal state of BE, for the purpose of computing this quantity we could just as well consider the input to the channel to be the mixed state ρ_A obtained from ψ_{RA} by tracing out the reference system R.

For a classical channel, H(R|B) is always nonnegative and the coherent information is never positive. In the quantum setting, $I_c(R \ge B)$ is positive if the reference system Ris more strongly correlated with the channel output B than with the environment E. Indeed, an alternative way to express the coherent information is

$$I_c(R \rangle B) = \frac{1}{2} \left(I(R; B) - I(R; E) \right) = H(B) - H(E), \qquad (10.264)$$

where we note that (because ϕ_{RBE} is pure)

$$I(R; B) = H(R) + H(B) - H(RB) = H(R) + H(B) - H(E),$$

$$I(R; E) = H(R) + H(E) - H(RE) = H(R) + H(E) - H(B).$$
(10.265)

Now we can state the regularized formula for the quantum channel capacity — it is the optimal asymptotic coherent information per letter

$$Q(\mathcal{N}^{A \to B}) = \lim_{n \to \infty} \max_{A^n} \frac{1}{n} I_c(R^n) B^n)_{\phi_{R^n B^n E^n}},$$
(10.266)

where the input density operator ρ_{A^n} is allowed to be entangled across the *n* channel uses. If coherent information were subadditive, we could reduce this expression to a single-letter quantity, the one-shot capacity $Q_1(\mathcal{N})$. But, unfortunately, for some channels the coherent information can be superadditive, in which case the regularized formula is not very informative. At least we can say that $Q_1(\mathcal{N})$ is an achievable rate, and therefore a lower bound on the capacity.

Quantum Shannon Theory

10.7.2 The decoupling principle

Before we address achievability, let's understand why eq.(10.266) is an upper bound on the capacity. First we note that the monotonicity of mutual information implies a corresponding monotonicity property for the coherent information. Suppose that the channel $\mathcal{N}_1^{A \to B}$ is followed by a channel $\mathcal{N}_2^{B \to C}$. Because mutual information is monotonic we have

$$I(R; A) \ge I(R; B) \ge I(R; C),$$
 (10.267)

which can also be expressed as

$$H(R) - H(R|A) \ge H(R) - H(R|B) \ge H(R) - H(R|C),$$
(10.268)

and hence

$$I_c(R \mid A) \ge I_c(R \mid B) \ge I_c(R \mid C).$$
(10.269)

A quantum channel cannot increase the coherent information, which has been called the *quantum data-processing inequality*.

Suppose now that ρ_A is a quantum code state, and that the two channels acting in succession are a noisy channel $\mathcal{N}^{A\to B}$ and the decoding map $\mathcal{D}^{B\to \hat{B}}$ applied by Bob to the channel output in order to recover the channel input. Consider the action of the dilation $U^{A\to BE}$ of \mathcal{N} followed by the dilation $V^{B\to \hat{B}B'}$ of \mathcal{D} on the input purification ψ_{RA} , under the assumption that Bob is able to recover *perfectly*:

$$\psi_{RA} \xrightarrow{\boldsymbol{U}} \phi_{RBE} \xrightarrow{\boldsymbol{V}} \tilde{\psi}_{R\hat{B}B'E} = \psi_{R\hat{B}} \otimes \chi_{B'E}.$$
(10.270)

If the decoding is perfect, then after decoding Bob holds in system \hat{B} the purification of the state of R, so that

$$H(R) = I_c(R\rangle A)_{\psi} = I_c(R\rangle B)_{\tilde{\psi}}.$$
(10.271)

Since the initial and final states have the same coherent information, the quantum data processing inequality implies that the same must be true for the intermediate state ϕ_{RBE} :

$$H(R) = I_c(R \mid B) = H(B) - H(E)$$

$$\implies H(B) = H(RE) = H(R) + H(E).$$
(10.272)

Thus the state of RE is a product state. We have found that if Bob is able to recover perfectly from the action of the channel dilation $U^{A\to BE}$ on the pure state ψ_{RA} , then, in the resulting channel output pure state ϕ_{RBE} , the marginal state ρ_{RE} must be the product $\rho_R \otimes \rho_E$.

Conversely, suppose that ψ_{RA} is an entangled pure state, and Alice wishes to transfer the purification of R to Bob by sending it through the noisy channel $U^{A\to BE}$. And suppose that in the resulting tripartite pure state ϕ_{RBE} , the marginal state of REfactorizes as $\rho_{RE} = \rho_R \otimes \rho_E$. Then B decomposes into subsystems $B = \hat{B}_1 B_2$ such that

$$\phi_{RBE} = \psi_{RB_1} \otimes \chi_{B_2E}. \tag{10.273}$$

Now Bob can construct an isometric decoder $V^{B_1 \to \hat{B}}$, which extracts the purification of R into Bob's preferred subsystem \hat{B} . Since all purifications of R differ by an isometry on Bob's side, Bob can choose his decoding map to output the state $\psi_{R\hat{B}}$; then the input

state of RA is successfully transmitted to $R\hat{B}$ as desired. Furthermore, we may choose the initial state to be a maximally entangled state Φ_{RA} of the reference system with the code space of a quantum code; if the marginal state of RE factorizes in the resulting output pure state ϕ_{RBE} , then by the relative state method of Chapter 3 we conclude that any state in the code space can be sent through the channel and decoded with perfect fidelity by Bob.

We have found that purified quantum information transmitted through the noisy channel is exactly correctable if and only if the reference system is completely uncorrelated with the channel's environment, or as we sometimes say, *decoupled* from the environment. This is the *decoupling principle*, a powerful notion underlying many of the key results in the theory of quantum channels.

So far we have shown that exact correctability corresponds to exact decoupling. But we can likewise see that approximate correctability corresponds to approximate decoupling. Suppose for example that the state of RE is close to a product state in the L^1 norm:

$$\|\boldsymbol{\rho}_{RE} - \boldsymbol{\rho}_R \otimes \boldsymbol{\rho}_E\|_1 \le \varepsilon. \tag{10.274}$$

As we learned in Chapter 2, if two density operators are close together in this norm, that means they also have fidelity close to one and hence purifications with a large overlap. Any purification of the product state $\rho_R \otimes \rho_E$ has the form

$$\phi_{RBE} = \psi_{RB_1} \otimes \chi_{B_2E}, \tag{10.275}$$

and since all purifications of ρ_{RE} can be transformed to one another by an isometry acting on the purifying system B, there is a way to choose the decomposition $B = B_1B_2$ such that

$$F(\boldsymbol{\rho}_{RE}, \boldsymbol{\rho}_{R} \otimes \boldsymbol{\rho}_{E}) = \left\| \langle \phi_{RBE} | \tilde{\phi}_{RBE} \rangle \right\|^{2} \ge 1 - \|\boldsymbol{\rho}_{RE} - \boldsymbol{\rho}_{R} \otimes \boldsymbol{\rho}_{E}\|_{1} \ge 1 - \varepsilon. \quad (10.276)$$

Furthermore, because fidelity is monotonic, both under tracing out E and under the action of Bob's decoding map, and because Bob can decode $\tilde{\phi}_{RBE}$ perfectly, we conclude that

$$F\left(\mathcal{D}^{B\to\hat{B}}\left(\boldsymbol{\rho}_{RB}\right),\psi_{R\hat{B}}\right) \ge 1-\varepsilon \tag{10.277}$$

if Bob chooses the proper decoding map \mathcal{D} . Thus approximate decoupling in the L^1 norm implies high-fidelity correctability. It is convenient to note that the argument still works the same way if ρ_{RE} is ε -close in the L^1 norm to $\tilde{\rho}_R \otimes \tilde{\rho}_E$, where $\tilde{\rho}_R$ is not necessarily tr_E (ρ_{RE}) and $\tilde{\rho}_E$ is not necessarily tr_R (ρ_{RE}). We'll use this form of the argument in what follows.

On the other hand, if (approximate) decoupling fails, the fidelity of Bob's decoded state will be seriously compromised. Suppose that in the state ϕ_{RBE} we have

$$H(R) + H(E) - H(RE) = \varepsilon > 0.$$
 (10.278)

Then the coherent information of ϕ is

$$I_c(R \rangle B)_{\phi} = H(B)_{\phi} - H(E)_{\phi} = H(RE)_{\phi} - H(E)_{\phi} = H(R)_{\phi} - \varepsilon.$$
(10.279)

By the quantum data processing inequality, we know that the coherent information of Bob's decoded state $\tilde{\psi}_{R\hat{B}}$ is no larger; hence

$$I_c(R)\hat{B})_{\tilde{\psi}} = H(R)_{\psi} - H(R\hat{B})_{\tilde{\psi}} \le H(R)_{\psi} - \varepsilon, \qquad (10.280)$$

Quantum Shannon Theory

and therefore

$$H(R\ddot{B})_{\tilde{\psi}} \ge \varepsilon \tag{10.281}$$

The deviation from perfect decoupling means that the decoded state of $R\hat{B}$ has some residual entanglement with the environment E, and is therefore impure.

Now we have the tools to derive an upper bound on the quantum channel capacity $Q(\mathcal{N})$. For *n* channel uses, let $\psi^{(n)}$ be a maximally entangled state of a reference system $\mathcal{H}_{R}^{(n)} \subseteq \mathcal{H}_{R}^{\otimes n}$ with a code space $\mathcal{H}_{A}^{(n)} \subseteq \mathcal{H}_{A}^{\otimes n}$, where dim $\mathcal{H}_{A}^{(n)} = 2^{nR}$, so that

$$I_c(R^n \rangle A^n)_{\psi^{(n)}} = H(R^n)_{\psi^{(n)}} = nR.$$
(10.282)

Now A^n is transmitted to B^n through $(U^{A\to BE})^{\otimes n}$, yielding the pure state $\phi^{(n)}$ of $R^n B^n E^n$. If Bob can decode with high fidelity, then his decoded state must have coherent information $H(R^n)_{\psi^{(n)}} - o(n)$, and the quantum data processing inequality then implies that

$$I_c(R^n \rangle B^n)_{\phi^{(n)}} = H(R^n)_{\psi^{(n)}} - o(n) = nR - o(n)$$
(10.283)

and hence

$$R = \frac{1}{n} I_c(R^n) B^n)_{\phi^{(n)}} + o(1).$$
(10.284)

Taking the limit $n \to \infty$ we see that the expression for $Q(\mathcal{N})$ in eq.(10.266) is an upper bound on the quantum channel capacity. In Exercise 10.10, you will sharpen the statement eq.(10.283), showing that

$$H(R^n) - I_c(R^n) B^n) \le 2H_2(\varepsilon) + 4\varepsilon nR.$$
(10.285)

To show that $Q(\mathcal{N})$ is an achievable rate, rather than just an upper bound, we will need to formulate a quantum version of Shannon's random coding argument. Our strategy (see §10.9.3) will be to demonstrate the existence of codes that achieve approximate decoupling of E^n from \mathbb{R}^n .

10.7.3 Degradable channels

Though coherent information can be superadditive in some cases, there are classes of channels for which the coherent information is additive, and therefore the quantum channel capacity matches the single-shot capacity, for which there is a single-letter formula. One such class is the class of *degradable channels*.

To understand what a degradable channel is, we first need the concept of a complementary channel. Any channel $\mathcal{N}^{A\to B}$ has a Stinespring dilation $U^{A\to BE}$, from which we obtain $\mathcal{N}^{A\to B}$ by tracing out the environment E. Alternatively we obtain the channel $\mathcal{N}_c^{A\to E}$ complementary to $\mathcal{N}^{A\to B}$ by tracing out B instead. Since we have the freedom to compose $U^{A\to BE}$ with an isometry $V^{E\to E}$ without changing $\mathcal{N}^{A\to B}$, the complementary channel is defined only up to an isometry acting on E. This lack of uniqueness need not trouble us, because the properties of interest for the complementary channel are invariant under such isometries.

We say that the channel $\mathcal{N}^{A\to B}$ is degradable if we can obtain its complementary channel by composing $\mathcal{N}^{A\to B}$ with a channel mapping B to E:

$$\mathcal{N}_{c}^{A \to E} = \mathcal{T}^{B \to E} \circ \mathcal{N}^{A \to B}. \tag{10.286}$$

54

In this sense, when Alice sends a state through the channel, Bob, who holds system B, receives a less noisy copy than Eve, who holds system E.

Now suppose that $U_1^{A_1 \to B_1 E_1}$ and $U_2^{A_2 \to B_2 E_2}$ are dilations of the degradable channels \mathcal{N}_1 and \mathcal{N}_2 . Alice introduces a reference system R and prepares an input pure state $\psi_{RA_1A_2}$, then sends the state to Bob via $\mathcal{N}_1 \otimes \mathcal{N}_2$, preparing the output pure state $\phi_{RB_1B_2E_1E_2}$. We would like to evaluate the coherent information $I_c(R)B_1B_2)_{\phi}$ in this state.

The key point is that because both channels are degradable, there is a product channel $\mathcal{T}_1 \otimes \mathcal{T}_2$ mapping B_1B_2 to E_1E_2 , and the monotonicity of mutual information therefore implies

$$I(B_1; B_2) \ge I(E_1; E_2). \tag{10.287}$$

Therefore, the coherent information satisfies

$$I_{c}(R \mid B_{1}B_{2}) = H(B_{1}B_{2}) - H(E_{1}E_{2})$$

= $H(B_{1}) + H(B_{2}) - I(B_{1}; B_{2}) - H(E_{1}) - H(E_{2}) + I(E_{1}; E_{2})$
 $\leq H(B_{1}) - H(E_{1}) + H(B_{2}) - H(E_{2}).$ (10.288)

These quantities are all evaluated in the state $\phi_{RB_1B_2E_1E_2}$. But notice that for the evaluation of $H(B_1)-H(E_1)$, the isometry $U_2^{A_2 \to B_2E_2}$ is irrelevant. This quantity is really the same as the coherent information $I_c(RA_2 \rangle B_1)$, where now we regard A_2 as part of the reference system for the input to channel \mathcal{N}_1 . Similarly $H(B_2) - H(E_2) = I_c(RA_1 \rangle B_2)$, and therefore,

$$I_c(R \mid B_1 \mid B_2) \le I_c(R \mid A_2 \mid B_1) + I_c(R \mid A_1 \mid B_2) \le Q_1(\mathcal{N}_1) + Q_1(\mathcal{N}_2), \quad (10.289)$$

where in the last inequality we use the definition of the one-shot capacity as coherent information maximized over all inputs. Since $Q_1(\mathcal{N}_1 \otimes \mathcal{N}_2)$ is likewise defined by maximizing the coherent information $I_c(R \otimes B_1 B_2)$, we find that

$$Q_1(\mathcal{N}_1 \otimes \mathcal{N}_2) \le Q_1(\mathcal{N}_1) + Q_1(\mathcal{N}_2) \tag{10.290}$$

if \mathcal{N}_1 and \mathcal{N}_2 are degradable.

The regularized formula for the capacity of \mathcal{N} is

$$Q(\mathcal{N}) = \lim_{n \to \infty} \frac{1}{n} Q_1(\mathcal{N}^{\otimes n}) \le Q_1(\mathcal{N}), \qquad (10.291)$$

where the last inequality follows from eq.(10.290) assuming that \mathcal{N} is degradable. We'll see that $Q_1(\mathcal{N})$ is actually an achievable rate, and therefore a single-letter formula for the quantum capacity of a degradable channel.

As a concrete example of a degradable channel, consider the *generalized dephasing* channel with dilation

$$U^{A \to BE} : |x\rangle_A \mapsto |x\rangle_B \otimes |\alpha_x\rangle_E, \tag{10.292}$$

where $\{|x\rangle_A\}$, $\{|x\rangle_B\}$ are orthonormal bases for \mathcal{H}_A , \mathcal{H}_B respectively, and the states $\{|\alpha_x\rangle_E\}$ of the environment are not necessarily orthogonal. The corresponding channel is

$$\mathcal{N}^{A \to B} : \boldsymbol{\rho} \mapsto \sum_{x, x'} |x\rangle \langle x| \boldsymbol{\rho} | x' \rangle \langle \alpha_{x'} | \alpha_x \rangle \langle x' |, \qquad (10.293)$$

which has the complementary channel

$$\mathcal{N}_{c}^{A \to E} : \boldsymbol{\rho} \mapsto \sum_{x} |\alpha_{x}\rangle \langle x | \boldsymbol{\rho} | x \rangle \langle \alpha_{x} |.$$
(10.294)

In the special case where the states $\{|\alpha_x\rangle_E = |x\rangle_E\}$ are orthonormal, we obtain the completely dephasing channel

$$\Delta^{A \to B} : \boldsymbol{\rho} \mapsto \sum_{x} |x\rangle \langle x| \boldsymbol{\rho} |x\rangle \langle x|, \qquad (10.295)$$

whose complement $\Delta^{A \to E}$ has the same form as $\Delta^{A \to B}$. We can easily check that

$$\mathcal{N}_c^{A \to E} = \mathcal{N}_c^{C \to E} \circ \Delta^{B \to C} \circ \mathcal{N}^{A \to B}; \qquad (10.296)$$

therefore $\mathcal{N}_c \circ \Delta$ degrades \mathcal{N} to \mathcal{N}_c . Thus \mathcal{N} is degradable and $Q(\mathcal{N}) = Q_1(\mathcal{N})$.

Further examples of degradable channels are discussed in Exercise 10.12.

10.8 Quantum Protocols

Using the decoupling principle in an i.i.d. setting, we can prove achievable rates for two fundamental quantum protocols. These are fondly known as the father and mother protocols, so named because each spawns a brood of interesting corollaries. We will formulate these protocols and discuss some of their "children" in this section, postponing the proofs until §10.9.

10.8.1 Father: Entanglement-assisted quantum communication

The father protocol is a scheme for entanglement-assisted quantum communication. Through many uses of a noisy quantum channel $\mathcal{N}^{A\to B}$, this protocol sends quantum information with high fidelity from Alice to Bob, while also consuming some previously prepared quantum entanglement shared by Alice and Bob. The task performed by the protocol is summarized by the *father resource inequality*

$$\left\langle \mathcal{N}^{A \to B} : \boldsymbol{\rho}_A \right\rangle + \frac{1}{2} I(R; E)[qq] \ge \frac{1}{2} I(R; B)[q \to q], \tag{10.297}$$

where the resources on the left-hand side can be used to achieve the result on the righthand side, in an asymptotic i.i.d. setting. That is, for any positive ε , the quantum channel \mathcal{N} may be used *n* times to transmit $\frac{n}{2}I(R;B) - o(n)$ qubits with fidelity $F \geq 1 - \varepsilon$, while consuming $\frac{n}{2}I(R;E) + o(n)$ ebits of entanglement shared between sender and receiver. These entropic quantities are evaluated in a tripartite pure state ϕ_{RBE} , obtained by applying the Stinespring dilation $U^{A \to BE}$ of $\mathcal{N}^{A \to B}$ to the purification ψ_{RA} of the input density operator ρ_A . Eq.(10.297) means that for any input density operator ρ_A , there exists a coding procedure that achieves the quantum communication at the specified rate by consuming entanglement at the specified rate.

To remember the father resource inequality, it helps to keep in mind that I(R; B) quantifies something good, the correlation with the reference system which survives transmission through the channel, while I(R; E) quantifies something bad, the correlation between the reference system R and the channel's environment E, which causes the transmitted information to decohere. The larger the good quantity I(R; B), the higher the rate of quantum communication. The larger the bad quantity I(R; E), the

56

more entanglement we need to consume to overcome the noise in the channel. To remember the factor of $\frac{1}{2}$ in front of I(R; B), consider the case of a noiseless quantum channel, where ψ_{RA} is maximally entangled; in that case there is no environment,

$$\phi_{RB} = \frac{1}{\sqrt{d}} \sum_{i=0}^{d-1} |i\rangle_R \otimes |i\rangle_B, \qquad (10.298)$$

and $\frac{1}{2}I(R; B) = H(R) = H(B) = \log_2 d$ is just the number of qubits in A. To remember the factor of $\frac{1}{2}$ in front of I(R; E), consider the case of a noiseless *classical* channel, where the quantum information completely decoheres in a preferred basis; in that case

$$\phi_{RBE} = \frac{1}{\sqrt{d}} \sum_{i=0}^{d-1} |i\rangle_R \otimes |i\rangle_B \otimes |i\rangle_E, \qquad (10.299)$$

and $I(R; B) = I(R; E) = H(R) = H(B) = \log_2 d$. Then the father inequality merely says that we can teleport $\frac{n}{2}$ qubits by consuming $\frac{n}{2}$ ebits and sending *n* classical bits.

Before proving the father resource inequality, we will first discuss a few of its interesting consequences.

Entanglement-assisted classical communication.

Suppose Alice wants to send classical information to Bob, rather than quantum information. Then we can use superdense coding to turn the quantum communication achieved by the father protocol into classical communication, at the cost of consuming some additional entanglement. By invoking the superdense coding resource inequality

$$SD: \quad [q \to q] + [qq] \ge 2[c \to c] \tag{10.300}$$

 $\frac{n}{2}I(R;B)$ times, and combining with the father resource inequality, we obtain I(R;B) bits of classical communication per use of the channel while consuming a number of ebits

$$\frac{1}{2}I(R;E) + \frac{1}{2}I(R;B) = H(R)$$
(10.301)

per channel use. Thus we obtain an achievable rate for entanglement-assisted classical communication through the noisy quantum channel:

$$\langle \mathcal{N}^{A \to B} : \boldsymbol{\rho}_A \rangle + H(R)[qq] \ge I(R;B)[c \to c].$$
 (10.302)

We may define the *entanglement-assisted classical capacity* $C_E(\mathcal{N})$ as the supremum over achievable rates of classical communication per channel use, assuming that an unlimited amount of entanglement is available at no cost. Then the resource inequality eq.(10.302) implies

$$C_E(\mathcal{N}) \ge \max_{A} I(R; B). \tag{10.303}$$

In this case there is a matching upper bound, so $C_E(\mathcal{N})$ is really an equality, and hence a single-letter formula for the entanglement-assisted classical capacity. Furthermore, eq.(10.302) tells us a rate of entanglement consumption which suffices to achieve the capacity. If we disregard the cost of entanglement, the father protocol shows that a rate can be achieved for entanglement-assisted quantum communication which is half the entanglement-assisted classical capacity $C_E(\mathcal{N})$ of the noisy channel \mathcal{N} . That's clearly true, since by consuming entanglement we can use teleportation to convert n bits of classical communication into n/2 qubits of quantum communication.

Quantum channel capacity.

It may be that Alice wants to send quantum information to Bob, but Alice and Bob are not so fortunate as to have pre-existing entanglement at their disposal. They can still make use of the father protocol, if we are willing to loan them some entanglement, which they are later required to repay. In this case we say that the entanglement *catalyzes* the quantum communication. Entanglement is needed to activate the process to begin with, but at the conclusion of the process no net entanglement has been consumed.

In this catalytic setting, Alice and Bob borrow $\frac{1}{2}I(R; E)$ ebits of entanglement per use of the channel to get started, execute the father protocol, and then sacrifice some of the quantum communication they have generated to replace the borrowed entanglement via the resource inequality

$$[q \to q] \ge [qq]. \tag{10.304}$$

After repaying their debt, Alice and Bob retain a number of qubits of quantum communication per channel use

$$\frac{1}{2}I(R;B) - \frac{1}{2}I(R;E) = H(B) - H(E) = I_c(R \mid B), \qquad (10.305)$$

the channel's coherent information from R to B. We therefore obtain the achievable rate for quantum communication

$$\langle \mathcal{N}^{A \to B} : \boldsymbol{\rho}_A \rangle \ge I_c(R \rangle B)[q \to q],$$
 (10.306)

albeit in the catalyzed setting. It can actually be shown that this same rate is achievable without invoking catalysis (see §10.9.4). As already discussed in §10.7.1, though, because of the superadditivity of coherent information this resource inequality does not yield a general single-letter formula for the quantum channel capacity $Q(\mathcal{N})$.

10.8.2 Mother: Quantum state transfer

In the mother protocol, Alice, Bob, and Eve initially share a tripartite pure state ϕ_{ABE} ; thus Alice and Bob together hold the purification of Eve's system E. Alice wants to send her share of this purification to Bob, using as few qubits of noiseless quantum communication as possible. Therefore, Alice divides her system A into two subsystems A_1 and A_2 , where A_1 is as small as possible and A_2 is uncorrelated with E. She keeps A_2 and sends A_1 to Bob. After receiving A_1 , Bob divides A_1B into two subsystems B_1 and B_2 , where B_1 purifies E and B_2 purifies A_2 . Thus, at the conclusion of the protocol, Bob holds the purification of E in B_1 , and in addition Alice and Bob share a bipartite pure state in A_2B_2 . The protocol is portrayed in the following diagram:



In the i.i.d. version of the mother protocol, the initial state is $\phi_{ABE}^{\otimes n}$, and the task achieved by the protocol is summarized by the mother resource inequality

$$\langle \phi_{ABE} \rangle + \frac{1}{2} I(A; E)[q \to q] \ge \frac{1}{2} I(A; B)[qq] + \langle \phi'_{\tilde{B}E} \rangle, \qquad (10.307)$$

where the resources on the left-hand side can be used to achieve the result on the righthand side, in an asymptotic i.i.d. setting, and the entropic quantities are evaluated in the state ϕ_{ABE} . That is, if $A_1^{(n)}$ denotes the state Alice sends and $A_2^{(n)}$ denotes the state she keeps, then for any positive ε , the state of $A_2^{(n)}E^n$ is ε -close in the L^1 norm to a product state, where $\log |A_1^{(n)}| = \frac{n}{2}I(A; E) + o(n)$, while $A_2^{(n)}B_2^{(n)}$ contains $\frac{n}{2}I(A; B) - o(n)$ shared ebits of entanglement. Eq.(10.307) means that for any input pure state ϕ_{ABE} there is a way to choose the subsystem $A_2^{(n)}$ of the specified dimension such that $A_2^{(n)}$ and E^n are nearly uncorrelated and the specified amount of entanglement is harvested in $A_2^{(n)}B_2^{(n)}$.

The mother protocol is in a sense *dual* to the father protocol. While the father protocol consumes entanglement to achieve quantum communication, the mother protocol consumes quantum communication and harvests entanglement. For the mother, I(A; B)quantifies the correlation between Alice and Bob at the beginning of the protocol (something *good*), and I(A; E) quantifies the noise in the initial shared entanglement (something *bad*). The mother protocol can also be viewed as a quantum generalization of the Slepian-Wolf distributed compression protocol discussed in §10.1.3. The mother protocol merges Alice's and Bob's shares of the purification of E by sending Alice's share to Bob, much as distributed source coding merges the classical correlations shared by Alice and Bob by sending Alice's classical information to Bob. For this reason the mother protocol has been called the *fully quantum Slepian-Wolf protocol*; the modifier "fully" will be clarified in §10.8.2, when we discuss a variant on quantum state transfer in which classical communication is assumed to be freely available.

We may also view the mother protocol as a generalization of the entanglement concentration protocol discussed in §10.4, extending that discussion in three ways:

- 1. The initial entangled state shared by Alice and Bob may be mixed rather than pure.
- 2. The communication from Alice to Bob is quantum rather than classical.
- 3. The amount of communication that suffices to execute the protocol is quantified by the resource inequality.

Also note that if the state of AE is pure (uncorrelated with B), then the mother protocol reduces to Schumacher compression. In that case $\frac{1}{2}I(A; E) = H(A)$, and the mother resource inequality states that the purification of A^n can be transferred to Bob with high fidelity using nH(A) + o(n) qubits of quantum communication.

Before proving the mother resource inequality, we will first discuss a few of its interesting consequences.

Hashing inequality.

Suppose Alice and Bob wish to distill entanglement from many copies of the state ϕ_{ABE} , using only local operations and classical communication (LOCC). In the catalytic setting, they can borrow some quantum communication, use the mother protocol to distill some shared entanglement, and then use classical communication and their harvested entanglement to repay their debt via quantum teleportation. Using the teleportation resource inequality

$$TP: \quad [qq] + 2[c \to c] \ge [q \to q] \tag{10.308}$$

 $\frac{n}{2}I(A; E)$ times, and combining with the mother resource inequality, we obtain

$$\langle \phi_{ABE} \rangle + I(A; E)[c \to c] \ge I_c(A \rangle B)[qq] + \langle \phi'_{\tilde{B}E} \rangle,$$
 (10.309)

since the net amount of distilled entanglement is $\frac{1}{2}I(A;B)$ per copy of ϕ achieved by the mother minus the $\frac{1}{2}I(A;E)$ per copy consumed by teleportation, and

$$\frac{1}{2}I(A;B) - \frac{1}{2}I(A;E) = H(B) - H(E) = I_c(A \mid B).$$
(10.310)

Eq.(10.309) is the hashing inequality, which quantifies an achievable rate for distilling ebits of entanglement shared by Alice and Bob from many copies of a mixed state ρ_{AB} , using one-way classical communication, assuming that $I_c(A \mid B) = -H(A \mid B)$ is positive. Furthermore, the hashing inequality tells us how much classical communication suffices for this purpose.

In the case where the state ρ_{AB} is pure, $I_c(A|B) = H(A) - H(AB) = H(A)$ and there is no environment E; thus we recover our earlier conclusion about concentration of pure-state bipartite entanglement — that H(A) Bell pairs can be extracted per copy, with a negligible classical communication cost.

State merging.

Suppose Alice and Bob share the purification of Eve's state, and Alice wants to transfer her share of the purification to Bob, where now unlimited *classical* communication from Alice to Bob is available at no cost. In contrast to the mother protocol, Alice wants to achieve the transfer with as little one-way quantum communication as possible, even if she needs to send more bits in order to send fewer qubits.

In the catalytic setting, Alice and Bob can borrow some quantum communication, perform the mother protocol, then use teleportation and the entanglement extracted by the mother protocol to repay some of the borrowed quantum communication. Combining teleportation of $\frac{n}{2}I(A; B)$ qubits with the mother resource inequality, we obtain

$$\langle \phi_{ABE} \rangle + H(A|B)[q \to q] + I(A;B)[c \to c] \ge \langle \phi'_{\tilde{B}E} \rangle,$$
 (10.311)

using

$$\frac{1}{2}I(A;E) - \frac{1}{2}I(A;B) = H(E) - H(B) = H(AB) - H(B) = H(A|B).$$
(10.312)

Eq.(10.311) is the *state-merging inequality*, expressing how much quantum and classical communication suffices to achieve the state transfer in an i.i.d. setting, assuming that H(A|B) is nonnegative.

Like the mother protocol, this state merging protocol can be viewed as a (partially) quantum version of the Slepian-Wolf protocol for merging classical correlations. In the classical setting, H(X|Y) quantifies Bob's remaining ignorance about Alice's information X when Bob knows only Y; correspondingly, Alice can reveal X to Bob by sending H(X|Y) bits per letter of X. Similarly, state merging provides an operational meaning to the quantum conditional information H(A|B), as the number of qubits per copy of ϕ that Alice sends to Bob to convey her share of the purification of E, assuming classical communication is free. In this sense we may regard H(A|B) as a measure of Bob's remaining "ignorance" about the shared purification of E when he holds only B.

Classically, H(X|Y) is nonnegative, and zero if and only if Bob is already certain about XY, but quantumly H(A|B) can be negative. How can Bob have "negative uncertainty" about the quantum state of AB? If H(A|B) < 0, or equivalently if I(A; E) < I(A; B), then the mother protocol yields more quantum entanglement than the amount of quantum communication it consumes. Therefore, when H(A|B) is negative (*i.e.* $I_c(A|B)$) is

positive), the mother resource inequality implies the Hashing inequality, asserting that classical communication from Alice to Bob not only achieves state transfer, but also distills -H(A|B) ebits of entanglement per copy of ϕ . These distilled ebits can be deposited in the entanglement bank, to be withdrawn as needed in future rounds of state merging, thus reducing the quantum communication cost of those future rounds. Bob's "negative uncertainty" today reduces the quantum communication cost of tasks to be performed tomorrow.

10.8.3 Operational meaning of strong subadditivity

The observation that H(A|B) is the quantum communication cost of state merging allows us to formulate a simple operational proof of the strong subadditivity of Von Neumann entropy, expressed in the form

$$H(A|BC) \le H(A|B), \text{ or } -H(A|B) \le -H(A|BC).$$
 (10.313)

When H(A|B) is positive, eq.(10.313) is the obvious statement that it is no harder to merge Alice's system with Bob's if Bob holds C as well as B. When H(A|B) is negative, eq.(10.313) is the obvious statement that Alice and Bob can distill no less entanglement using one-way classical communication if Bob holds C as well as B.

To complete this argument, we need to know that H(A|B) is not only *achievable* but also that it is the *optimal* quantum communication cost of state merging, and that -H(A|B) ebits is the optimal yield of hashing. The optimality follows from the principle that, for a bipartite pure state, k qubits of quantum communication cannot increase the shared entanglement of AB by more than k ebits.

If H(A|B) is negative, consider cutting the system ABE into the two parts AE and B, as in the following figure:



In the hashing protocol, applied to n copies of ϕ_{ABE} , the entanglement across this cut at the beginning of the protocol is nH(B). By the end of the protocol E^n has decoupled from $A_2^{(n)}$ and has entanglement nH(E) with $B_1^{(n)}$, ignoring o(n) corrections. If k ebits shared by Alice and Bob are distilled, the final entanglement across the AE-B cut is

$$nH(E) + k \le nH(B) \implies \frac{k}{n} \le H(B) - H(E) = -H(A|B).$$
(10.314)

This inequality holds because LOCC cannot increase the entanglement across the cut, and implies that no more than -H(A|B) ebits of entanglement per copy of ϕ_{ABE} can be distilled in the hashing protocol, asymptotically.

On the other hand, if H(A|B) is positive, at the conclusion of state merging $B_1^{(n)}$ is entangled with E^n , and the entanglement across the AE-B cut is at least nH(E). To achieve this increase in entanglement, the number of qubits sent from Alice to Bob must be at least

$$k \ge nH(E) - nH(B) \implies \frac{k}{n} \ge H(E) - H(B) = H(A|B)$$
(10.315)

This inequality holds because the entanglement across the cut cannot increase by more than the quantum communication across the cut, and implies that at least H(A|B)qubits must be sent per copy of ϕ_{ABE} to achieve state merging.

To summarize, we have proven strong subadditivity, not by the traditional route of sophisticated matrix analysis, but via a less direct method. This proof is built on two cornerstones of quantum information theory — the decoupling principle and the theory of typical subspaces — which are essential ingredients in the proof of the mother resource inequality.

10.8.4 Negative conditional entropy in thermodynamics

As a further application of the decoupling mother resource inequality, we now revisit Landauer's Principle, developing another perspective on the implications of negative quantum conditional entropy. Recall that erasure of a bit is a process which maps the bit to 0 irrespective of its initial value. This process is *irreversible* — knowing only the final state 0 after erasure, we cannot determine whether the initial state before erasure was 0 or 1. Irreversibility implies that erasure incurs an unavoidable thermodynamic cost. According to Landauer's Principle, erasing a bit at temperature T requires work no less than $W = kT \ln 2$.

A specific erasure procedure is analyzed in Exercise 10.16. Suppose a two-level quantum system has energy eigenstates $|0\rangle$, $|1\rangle$ with corresponding eigenvalues E_0 and E_1 , where $E = E_1 - E_0 \ge 0$. Initially the qubit is in an unknown mixture of these two states, and the energy splitting is E = 0. We erase the bit in three steps. In the first step, we bring the bit into contact with a heat bath at temperature T > 0, and wait for the bit to come to thermal equilibrium with the bath. In this step the bit "forgets" its initial value, but the bit is not yet erased because it has not been reset. In the second step, with the bit still in contact with the bath, we turn on a control field which slowly increases E_1 to a value much larger than kT while maintaining thermal equilibrium all the while, thus resetting the bit to $|0\rangle$. In the third step, we isolate the bit from the bath and turn off the control field, so the two states of the bit become degenerate again. As shown in Exercise 10.16, work $W = kT \ln 2$ is required to execute step 2, with the energy dissipated as heat flowing from bit to bath.

We can also run the last two steps backward, increasing E_1 while the bit is isolated from the bath, then decreasing E_1 with the bit in contact with the bath. This procedure maps the state $|0\rangle$ to the maximally mixed state of the bit, extracting work $W = kT \ln 2$ from the bath in the process.

Erasure is irreversible because the agent performing the erasure does not know the information being erased. (If a copy of the information were stored in her memory, survival of that copy would mean that the erasure had not succeeded). From an informationtheoretic perspective, the reduction in the thermodynamic entropy of the erased bit, and hence the work required to perform the erasure, arises because erasure reduces the agent's *ignorance* about the state of the bit, ignorance which is quantified by the Shannon entropy. But to be more precise, it is the *conditional* entropy of the system, given the state of the agent's memory, which captures the agent's ignorance before erasure and therefore also the thermodynamic cost of erasing. Thus the minimal work needed to erase system A should be expressed as

$$W(A|O) = H(A|O)kT\ln 2,$$
(10.316)

where O is the memory of the *observer* who performs the erasure, and H(A|O) quantifies that observer's ignorance about the state of A.

But what if A and O are quantum systems? We know that if A and O are entangled, then the conditional entropy H(A|O) can be negative. Does that mean we can erase A while *extracting* work rather than doing work?

Yes, we can! Suppose for example that A and O are qubits and their initial state is maximally entangled. By controlling the contact between AO and the heat bath, the observer can extract work $W = 2kT \log 2$ while transforming AO to a maximally mixed state, using the same work extraction protocol as described above. Then she can do work $W = kT \log 2$ to return A to the state $|0\rangle$. The net effect is to erase A while extracting work $W = kT \log 2$, satisfying the equality eq.(10.316).

To appreciate why this trick works, we should consider the joint state of AO rather than the state of A alone. Although the marginal state of A is mixed at the beginning of the protocol and pure at the end, the state of AO is pure at the beginning and mixed at the end. Positive work is extracted by sacrificing the purity of AO.

To generalize this idea, let's consider $n \gg 1$ copies of the state ρ_{AO} of system A and memory O. Our goal is to map the n copies of A to the erased state $|000...0\rangle$ while using or extracting the optimal amount of work. In fact, the optimal work per copy is given by eq.(10.316) in the $n \to \infty$ limit.

To achieve this asymptotic work per copy, the observer first projects A^n onto its typical subspace, succeeding with probability 1 - o(1). A unitary transformation then rotates the typical subspace to a subsystem \bar{A} containing n(H(A) + o(1)) qubits, while erasing the complementary qubits as in eq.(10.144). Now it only remains to erase \bar{A} .

The mother resource inequality ensures that we may decompose A into subsystems A_1A_2 such that A_2 contains $\frac{n}{2}(I(A; O) - o(1))$ qubits and is nearly maximally entangled with a subsystem of O^n . What is important for the erasure protocol is that we may identify a subsystem of $\overline{A}O^n$ containing n(I(A; O) - o(1)) qubits which is only distance o(1) away from a pure state. By controlling the contact between this subsystem and the heat bath, we may extract work $W = n(I(A; O) - o(1))kT \log 2$ while transforming the subsystem to a maximally mixed state. We then proceed to erase \overline{A} , expending work $kT \log |\overline{A}| = n(H(A) + o(1))kT \log 2$. The net work cost of the erasure, per copy of ρ_{AO} , is therefore

$$W = (H(A) - I(A; O) + o(1)) kT \log 2 = (H(A|O) + o(1)) kT \log 2,$$
(10.317)

and the erasure succeeds with probability 1 - o(1). A notable feature of the protocol is that only the subsystem of O^n which is entangled with A_2 is affected. Any correlation of the memory O with other systems remains intact, and can be exploited in the future to reduce the cost of erasure of those other systems.

As does the state merging protocol, this erasure protocol provides an operational interpretation of strong subadditivity. For positive H(A|O), $H(A|O) \ge H(A|OO')$ means that it is no harder to erase A if the observer has access to both O and O' than if she has access to O alone. For negative H(A|O), $-H(A|OO') \ge -H(A|O)$ means that we can extract at least as much work from AOO' as from its subsystem AO.

To carry out this protocol and extract the optimal amount of work while erasing A, we need to know which subsystem of O^n provides the purification of A_2 . The decoupling argument ensures that this subsystem exists, but does not provide a constructive method for finding it, and therefore no concrete protocol for erasing at optimal cost. This quandary is characteristic of Shannon theory; for example, Shannon's noisy channel coding theorem ensures the existence of a code that achieves the channel capacity, but does not provide any explicit code construction.

10.9 The Decoupling Inequality

Achievable rates for quantum protocols are derived by using random codes, much as in classical Shannon theory. But this similarity between classical and quantum Shannon theory is superficial — at a deeper conceptual level, quantum protocols differ substantially from classical ones. Indeed, the decoupling principle underlies many of the key findings of quantum Shannon theory, providing a unifying theme that ties together many different results. In particular, the mother and father resource inequalities, and hence all their descendants enumerated above, follow from an inequality that specifies a sufficient condition for decoupling.

This decoupling inequality addresses the following question: Suppose that Alice and Eve share a quantum state σ_{AE} , where A is an n-qubit system. This state may be mixed, but in general A and E are correlated; that is, I(A; E) > 0. Now Alice starts discarding qubits one at a time, where each qubit is a randomly selected two-dimensional subsystem of what Alice holds. Each time Alice discards a qubit, her correlation with E grows weaker. How many qubits should she discard so that the subsystem she retains has a negligible correlation with Eve's system E?

To make the question precise, we need to formalize what it means to discard a random qubit. More generally, suppose that A has dimension |A|, and Alice decomposes A into subsystems A_1 and A_2 , then discards A_1 and retains A_2 . We would like to consider many possible ways of choosing the discarded system with specified dimension $|A_1|$. Equivalently, we may consider a fixed decomposition $A = A_1A_2$, where we apply a unitary transformation U to A before discarding A_1 . Then discarding a random subsystem with dimension $|A_1|$ is the same thing as applying a random unitary U before discarding the fixed subsystem A_1 :



To analyze the consequences of discarding a random subsystem, then, we will need to be able to compute the expectation value of a function f(U) when we average Uuniformly over the group of unitary $|A| \times |A|$ matrices. We denote this expectation value as $\mathbb{E}_{U}[f(U)]$; to perform computations we will only need to know that \mathbb{E}_{U} is suitably normalized, and is invariant under left or right multiplication by any constant unitary matrix V:

$$\mathbb{E}_{\boldsymbol{U}}[\boldsymbol{I}] = 1, \quad \mathbb{E}_{\boldsymbol{U}}[f(\boldsymbol{U})] = \mathbb{E}_{\boldsymbol{U}}[f(\boldsymbol{V}\boldsymbol{U})] = \mathbb{E}_{\boldsymbol{U}}[f(\boldsymbol{U}\boldsymbol{V})]. \quad (10.318)$$

These conditions uniquely define $\mathbb{E}_{U}[f(U)]$, which is sometimes described as the integral over the unitary group using the *invariant measure* or *Haar measure* on the group.

If we apply the unitary transformation U to A, and then discard A_1 , the marginal

state of A_2E is

$$\boldsymbol{\sigma}_{A_{2}E}(\boldsymbol{U}) := \operatorname{tr}_{A_{1}}\left(\left(\boldsymbol{U}_{A} \otimes \boldsymbol{I}_{E}\right) \boldsymbol{\sigma}_{AE}\left(\boldsymbol{U}_{A}^{\dagger} \otimes \boldsymbol{I}_{E}\right)\right).$$
(10.319)

The decoupling inequality expresses how close (in the L^1 norm) σ_{A_2E} is to a product state when we average over U:

$$\left(\mathbb{E}_{\boldsymbol{U}}\left[\|\boldsymbol{\sigma}_{A_{2}E}(\boldsymbol{U})-\boldsymbol{\sigma}_{A_{2}}^{\max}\otimes\boldsymbol{\sigma}_{E}\|_{1}\right]\right)^{2} \leq \frac{|A_{2}|\cdot|E|}{|A_{1}|} \operatorname{tr}\left(\boldsymbol{\sigma}_{AE}^{2}\right), \quad (10.320)$$

where

$$\boldsymbol{\sigma}_{A_2}^{\max} := \frac{1}{|A_2|} \boldsymbol{I}$$
(10.321)

denotes the maximally mixed state on A_2 , and σ_E is the marginal state tr_A σ_{AE} .

This inequality has interesting consequences even in the case where there is no system E at all and σ_A is pure, where it becomes

$$\left(\mathbb{E}_{\boldsymbol{U}}\left[\|\boldsymbol{\sigma}_{A_{2}}(\boldsymbol{U})-\boldsymbol{\sigma}_{A_{2}}^{\max}\|_{1}\right]\right)^{2} \leq \frac{|A_{2}|}{|A_{1}|} \operatorname{tr}\left(\boldsymbol{\sigma}_{A}^{2}\right) = \frac{|A_{2}|}{|A_{1}|}.$$
(10.322)

Eq.(10.322) implies that, for a randomly chosen pure state of the bipartite system $A = A_1A_2$, where $|A_2|/|A_1| \ll 1$, the density operator on A_2 is very nearly maximally mixed with high probability. One can likewise show that the expectation value of the entanglement entropy of A_1A_2 is very close to the maximal value: $\mathbb{E}[H(A_2)] \ge \log_2 |A_2| - |A_2|/(2|A_1| \ln 2)$. Thus, if for example A_2 is 50 qubits and A_1 is 100 qubits, the typical entropy deviates from maximal by only about $2^{-50} \approx 10^{-15}$.

10.9.1 Proof of the decoupling inequality

To prove the decoupling inequality, we will first bound the distance between σ_{A_2E} and a product state in the L^2 norm, and then use the Cauchy-Schwarz inequality to obtain a bound on the L^1 distance. Eq.(10.320) follows from

$$\mathbb{E}_{\boldsymbol{U}}\left[\|\boldsymbol{\sigma}_{A_{2}E}(\boldsymbol{U}) - \boldsymbol{\sigma}_{A_{2}}^{\max} \otimes \boldsymbol{\sigma}_{E}\|_{2}^{2}\right] \leq \frac{1}{|A_{1}|} \operatorname{tr}\left(\boldsymbol{\sigma}_{AE}^{2}\right), \qquad (10.323)$$

combined with

$$(\mathbb{E}[f])^2 \le \mathbb{E}[f]^2$$
 and $||M||_1^2 \le d||M||_2^2$ (10.324)

(for nonnegative f), which implies

$$\left(\mathbb{E}\left[\|\cdot\|_{1}\right]\right)^{2} \leq \mathbb{E}\left[\|\cdot\|_{1}^{2}\right] \leq |A_{2}|\cdot|E|\cdot\mathbb{E}\left[\|\cdot\|_{2}^{2}\right].$$
(10.325)

We also note that

$$\|\boldsymbol{\sigma}_{A_{2}E} - \boldsymbol{\sigma}_{A_{2}}^{\max} \otimes \boldsymbol{\sigma}_{E}\|_{2}^{2} = \operatorname{tr} \left(\boldsymbol{\sigma}_{A_{2}E} - \boldsymbol{\sigma}_{A_{2}}^{\max} \otimes \boldsymbol{\sigma}_{E}\right)^{2}$$
$$= \operatorname{tr} \left(\boldsymbol{\sigma}_{A_{2}E}^{2}\right) - \frac{1}{|A_{2}|} \operatorname{tr} \left(\boldsymbol{\sigma}_{E}^{2}\right), \qquad (10.326)$$

because

$$\operatorname{tr}\left(\boldsymbol{\sigma}_{A_{2}}^{\max}\right)^{2} = \frac{1}{|A_{2}|};$$
 (10.327)

therefore, to prove eq.(10.323) it suffices to show

$$\mathbb{E}_{\boldsymbol{U}}\left[\operatorname{tr}\left(\boldsymbol{\sigma}_{A_{2}E}^{2}(\boldsymbol{U})\right)\right] \leq \frac{1}{|A_{2}|}\operatorname{tr}\left(\boldsymbol{\sigma}_{E}^{2}\right) + \frac{1}{|A_{1}|}\operatorname{tr}\left(\boldsymbol{\sigma}_{AE}^{2}\right).$$
(10.328)

We can facilitate the computation of $\mathbb{E}_{U}\left[\operatorname{tr}\left(\sigma_{A_{2}E}^{2}(U)\right)\right]$ using a clever trick. For any bipartite system BC, imagine introducing a second copy B'C' of the system. Then (Exercise 10.17)

$$\operatorname{tr}_{C}\left(\boldsymbol{\sigma}_{C}^{2}\right) = \operatorname{tr}_{BCB'C'}\left(\boldsymbol{I}_{BB'}\otimes\boldsymbol{S}_{CC'}\right)\left(\boldsymbol{\sigma}_{BC}\otimes\boldsymbol{\sigma}_{B'C'}\right),\qquad(10.329)$$

where $S_{CC'}$ denotes the swap operator, which acts as

$$\mathbf{S}_{CC'}:|i\rangle_C\otimes|j\rangle_{C'}\mapsto|j\rangle_C\otimes|i\rangle_{C'}.$$
(10.330)

In particular, then,

$$\operatorname{tr}_{A_{2}E}\left(\boldsymbol{\sigma}_{A_{2}E}^{2}(\boldsymbol{U})\right)$$

$$= \operatorname{tr}_{AEA'E'}\left(\boldsymbol{I}_{A_{1}A'_{1}}\otimes\boldsymbol{S}_{A_{2}A'_{2}}\otimes\boldsymbol{S}_{EE'}\right)\left(\boldsymbol{\sigma}_{AE}(\boldsymbol{U})\otimes\boldsymbol{\sigma}_{A'E'}(\boldsymbol{U})\right)$$

$$= \operatorname{tr}_{AEA'E'}\left(\boldsymbol{M}_{AA'}(\boldsymbol{U})\otimes\boldsymbol{S}_{EE'}\right)\left(\boldsymbol{\sigma}_{AE}\otimes\boldsymbol{\sigma}_{A'E'}\right), \qquad (10.331)$$

where

$$\boldsymbol{M}_{AA'}(\boldsymbol{U}) = \left(\boldsymbol{U}_A^{\dagger} \otimes \boldsymbol{U}_{A'}^{\dagger}\right) \left(\boldsymbol{I}_{A_1 A_1'} \otimes \boldsymbol{S}_{A_2 A_2'}\right) \left(\boldsymbol{U}_A \otimes \boldsymbol{U}_{A'}\right).$$
(10.332)

The expectation value of $M_{AA'}(U)$ is evaluated in Exercise 10.17; there we find

$$\mathbb{E}_{\boldsymbol{U}}[\boldsymbol{M}_{AA'}(\boldsymbol{U})] = c_{\boldsymbol{I}}\boldsymbol{I}_{AA'} + c_{\boldsymbol{S}}\boldsymbol{S}_{AA'}$$
(10.333)

where

$$c_{I} = \frac{1}{|A_{2}|} \left(\frac{1 - 1/|A_{1}|}{1 - 1/|A|} \right) \le \frac{1}{|A_{2}|},$$

$$c_{S} = \frac{1}{|A_{1}|} \left(\frac{1 - 1/|A_{2}|}{1 - 1/|A|} \right) \le \frac{1}{|A_{1}|}.$$
(10.334)

Plugging into eq.(10.331), we then obtain

$$\mathbb{E}_{\boldsymbol{U}}\left[\operatorname{tr}_{A_{2}E}\left(\boldsymbol{\sigma}_{A_{2}E}^{2}(\boldsymbol{U})\right)\right] \\
\leq \operatorname{tr}_{AEA'E'}\left(\left(\frac{1}{|A_{2}|}\boldsymbol{I}_{AA'}+\frac{1}{|A_{1}|}\boldsymbol{S}_{AA'}\right)\otimes\boldsymbol{S}_{EE'}\right)\left(\boldsymbol{\sigma}_{AE}\otimes\boldsymbol{\sigma}_{A'E'}\right) \\
= \frac{1}{|A_{2}|}\operatorname{tr}\left(\boldsymbol{\sigma}_{E}^{2}\right)+\frac{1}{|A_{1}|}\left(\boldsymbol{\sigma}_{AE}^{2}\right),$$
(10.335)

thus proving eq.(10.328) as desired.

10.9.2 Proof of the mother inequality

The mother inequality eq.(10.307) follows from the decoupling inequality eq.(10.320) in an i.i.d. setting. Suppose Alice, Bob, and Eve share the pure state $\phi_{ABE}^{\otimes n}$. Then there are jointly typical subspaces of A^n , B^n , and E^n , which we denote by \bar{A} , \bar{B} , \bar{E} , such that

$$\left|\bar{A}\right| = 2^{nH(A)+o(n)}, \quad \left|\bar{B}\right| = 2^{nH(B)+o(n)}, \quad \left|\bar{E}\right| = 2^{nH(E)+o(n)}.$$
 (10.336)

Furthermore, the normalized pure state $\phi'_{\bar{A}\bar{B}\bar{E}}$ obtained by projecting $\phi^{\otimes n}_{ABE}$ onto $\bar{A} \otimes \bar{B} \otimes \bar{E}$ deviates from $\phi^{\otimes n}_{ABE}$ by distance o(1) in the L^1 norm.

66

In order to transfer the purification of E^n to Bob, Alice first projects A^n onto its typical subspace, succeeding with probability 1 - o(1), and compresses the result. She then divides her compressed system \bar{A} into two parts $\bar{A}_1\bar{A}_2$, and applies a random unitary to \bar{A} before sending \bar{A}_1 to Bob. Quantum state transfer is achieved if \bar{A}_2 decouples from \bar{E} .

Because $\phi'_{\bar{A}\bar{B}\bar{E}}$ is close to $\phi^{\otimes n}_{ABE}$, we can analyze whether the protocol is successful by supposing the initial state is $\phi'_{\bar{A}\bar{B}\bar{E}}$ rather than $\phi^{\otimes n}_{ABE}$. According to the decoupling inequality

$$\left(\mathbb{E}_{\boldsymbol{U}} \left[\| \boldsymbol{\sigma}_{\bar{A}_{2}\bar{E}}(\boldsymbol{U}) - \boldsymbol{\sigma}_{\bar{A}_{2}}^{\max} \otimes \boldsymbol{\sigma}_{\bar{E}} \|_{1} \right] \right)^{2} \leq \frac{|A| \cdot |E|}{|\bar{A}_{1}|^{2}} \operatorname{tr} \left(\boldsymbol{\sigma}_{\bar{A}\bar{E}}^{2} \right)$$

$$= \frac{1}{|\bar{A}_{1}|^{2}} 2^{n(H(A) + H(E) + o(1))} \operatorname{tr} \left(\boldsymbol{\sigma}_{\bar{A}\bar{E}}^{2} \right) = \frac{1}{|\bar{A}_{1}|^{2}} 2^{n(H(A) + H(E) - H(B) + o(1))};$$
(10.337)

here we have used properties of typical subspaces in the second line, as well as the property that $\sigma_{\bar{A}\bar{E}}$ and $\sigma_{\bar{B}}$ have the same nonzero eigenvalues, because $\phi'_{\bar{A}\bar{B}\bar{E}}$ is pure. Eq.(10.337) bounds the L^1 distance of $\sigma_{\bar{A}_2\bar{E}}(U)$ from a product state when averaged

Eq.(10.337) bounds the L^1 distance of $\sigma_{\bar{A}_2\bar{E}}(U)$ from a product state when averaged over all unitaries, and therefore suffices to ensure the existence of at least one unitary transformation U such that the L^1 distance is bounded above by the right-hand side. Therefore, by choosing this U, Alice can decouple \bar{A}_2 from E^n to o(1) accuracy in the L^1 norm by sending to Bob

$$\log_2 |\bar{A}_1| = \frac{n}{2} \left(H(A) + H(E) - H(B) + o(1) \right) = \frac{n}{2} \left(I(A; E) + o(1) \right)$$
(10.338)

qubits, randomly chosen from the (compressed) typical subspace of A^n . Alice retains $nH(A) - \frac{n}{2}I(A; E) - o(n)$ qubits of her compressed system, which are nearly maximally mixed and uncorrelated with E^n ; hence at the end of the protocol she shares with Bob this many qubit pairs, which have high fidelity with a maximally entangled state. Since ϕ_{ABE} is pure, and therefore $H(A) = \frac{1}{2}(I(A; E) - I(A; B))$, we conclude that Alice and Bob distill $\frac{n}{2}I(A; B) - o(n)$ ebits of entanglement, thus proving the mother resource inequality.

We can check that this conclusion is plausible using a crude counting argument. Disregarding the o(n) corrections in the exponent, the state $\phi_{ABE}^{\otimes n}$ is nearly maximally mixed on a typical subspace of $A^n E^n$ with dimension $2^{nH(AE)}$, *i.e.* the marginal state on \overline{AE} can be realized as a nearly uniform ensemble of this many mutually orthogonal states. If \overline{A}_1 is randomly chosen and sufficiently small, we expect that, for each state in this ensemble, \overline{A}_1 is nearly maximally entangled with a subsystem of the much larger system $\overline{A}_2\overline{E}$, and that the marginal states on $\overline{A}_2\overline{E}$ arising from different states in the \overline{AE} ensemble have a small overlap. Therefore, we anticipate that tracing out \overline{A}_1 yields a state on $\overline{A}_2\overline{E}$ which is nearly maximally mixed on a subspace with dimension $|\overline{A}_1|2^{nH(AE)}$. Approximate decoupling occurs when this state attains full rank on $\overline{A}_2\overline{E}$, since in that case it is close to maximally mixed on $\overline{A}_2\overline{E}$ and therefore close to a product state on its support. The state transfer succeeds, therefore, provided

$$|\bar{A}_{1}|2^{nH(AE)} \approx |\bar{A}_{2}| \cdot |\bar{E}| = \frac{|\bar{A}| \cdot |\bar{E}|}{|\bar{A}_{1}|} \approx \frac{2^{n(H(A)+H(E))}}{|\bar{A}_{1}|} \implies |\bar{A}_{1}|^{2} \approx 2^{nI(A;E)},$$
(10.339)

as in eq.(10.338).

Our derivation of the mother resource inequality, based on random coding, does not

exhibit any concrete protocol that achieves the claimed rate, nor does it guarantee the existence of any protocol in which the required quantum processing can be executed efficiently. Concerning the latter point, it is notable that our derivation of the decoupling inequality applies not just to the expectation value averaged uniformly over the unitary group, but also to any average over unitary transformations which satisfies eq.(10.333). In fact, this identity is satisfied by a uniform average over the Clifford group, which means that there is some Clifford transformation on \overline{A} which achieves the rates specified in the mother resource inequality. Any Clifford transformation on n qubits can be reached by a circuit with $O(n^2)$ gates. Since it is also known that Schumacher compression can be achieved by a polynomial-time quantum computation, Alice's encoding operation can be carried out efficiently.

In fact, after compressing, Alice encodes the quantum information she sends to Bob using a stabilizer code (with Clifford encoder U), and Bob's task, after receiving \bar{A}_1 is to correct the erasure of \bar{A}_2 . Bob can replace each erased qubit by the standard state $|0\rangle$, and then measure the code's check operators. With high probability, there is a unique Pauli operator acting on the erased qubits that restores Bob's state to the code space, and the recovery operation can be efficiently computed using linear algebra. Hence, Bob's part of the mother protocol, like Alice's, can be executed efficiently.

10.9.3 Proof of the father inequality

One-shot version.

In the one-shot version of the father protocol, Alice and Bob share a pair of maximally entangled systems A_1B_1 , and in addition Alice holds input state ρ_{A_2} of system A_2 which she wants to convey to Bob. Alice encodes ρ_{A_2} by applying a unitary transformation Vto $A = A_1A_2$, then sends A to Bob via the noisy quantum channel $\mathcal{N}^{A \to B_2}$. Bob applies a decoding map $\mathcal{D}^{B_1B_2 \to \tilde{A}_2}$ jointly to the channel output and his half of the entangled state he shares with Alice, hoping to recover Alice's input state with high fidelity:



We would like to know how much shared entanglement suffices for Alice and Bob to succeed.

This question can be answered using the decoupling inequality. First we introduce a reference system R' which is maximally entangled with A_2 ; then Bob succeeds if his decoder can extract the purification of R'. Because the systems $R'B_1$ and A_1A_1 are maximally entangled, the encoding unitary V acting on A_1A_2 can be replaced by its transpose V^T acting on $R'B_1$. We may also replace \mathcal{N} by its Stinespring dilation $U^{A_1A_2 \to B_2E}$, so that the extended output state ϕ of $R'B_1B_2E$ is pure:



Finally we invoke the decoupling principle — if R' and E decouple, then R' is purified by a subsystem of B_1B_2 , which means that Bob can recover ρ_{A_2} with a suitable decoding map.

If we consider V, and hence also V^T , to be a random unitary, then we may describe the situation this way: We have a tripartite pure state ϕ_{RB_2E} , where $R = R'B_1$, and we would like to know whether the marginal state of R'E is close to a product state when the random subsystem B_1 is discarded from R. This is exactly the question addressed by the decoupling inequality, which in this case may be expressed as

$$\left(\mathbb{E}_{\boldsymbol{V}}\left[\|\boldsymbol{\sigma}_{R'E}(\boldsymbol{V}) - \boldsymbol{\sigma}_{R'}^{\max} \otimes \boldsymbol{\sigma}_{E}\|_{1}\right]\right)^{2} \leq \frac{|R| \cdot |E|}{|B_{1}|^{2}} \operatorname{tr}\left(\boldsymbol{\sigma}_{RE}^{2}\right), \quad (10.340)$$

Eq.(10.340) asserts that the L^1 distance from a product state is bounded above when averaged uniformly over all unitary V's; therefore there must be some particular encoding unitary V that satisfies the same bound. We conclude that near-perfect decoupling of R'E, and therefore high-fidelity decoding of B_2 , is achievable provided that

$$|A_1| = |B_1| \gg |R'| \cdot |E| \operatorname{tr} \left(\sigma_{RE}^2 \right) = |A_2| \cdot |E| \operatorname{tr} \left(\sigma_{B_2}^2 \right), \quad (10.341)$$

where to obtain the second equality we use the purity of ϕ_{RB_2E} and recall that the reference system R' is maximally entangled with A_2 .

i.i.d. version.

In the i.i.d. version of the father protocol, Alice and Bob achieve high fidelity entanglement-assisted quantum communication through n uses of the quantum channel $\mathcal{N}^{A\to B}$. The code they use for this purpose can be described in the following way: Consider an input density operator ρ_A of system A, which is purified by a reference system R. Sending the purified input state ψ_{RA} through $U^{A\to BE}$, the isometric dilation of $\mathcal{N}^{A\to B}$, generates the tripartite pure state ϕ_{RBE} . Evidently applying $(U^{A\to BE})^{\otimes n}$ to $\psi_{RA}^{\otimes n}$ produces $\phi_{RBE}^{\otimes n}$.

But now suppose that before transmitting the state to Bob, Alice projects A^n onto its typical subspace \bar{A} , succeeding with probability 1 - o(1) in preparing a state of $\bar{A}\bar{R}$ that is nearly maximally entangled, where \bar{R} is the typical subspace of R^n . Imagine dividing \bar{R} into a randomly chosen subsystem B_1 and its complementary subsystem R'; then there is a corresponding decomposition of $A = A_1A_2$ such that A_1 is very nearly maximally entangled with B_1 and A_2 is very nearly maximally entangled with R'.

If we interpret B_1 as Bob's half of an entangled state of A_1B_1 shared with Alice, this becomes the setting where the one-shot father protocol applies, if we ignore the small deviation from maximal entanglement in A_1B_1 and $R'A_2$. As for our analysis of the i.i.d. mother protocol, we apply the one-shot father inequality not to $\phi_{RBE}^{\otimes n}$, but rather to the nearby state $\phi'_{\bar{R}\bar{B}\bar{E}}$, where \bar{B} and \bar{E} are the typical subspaces of B^n and E^n respectively. Applying eq.(10.340), and using properties of typical subspaces, we can bound the square of the L^1 deviation of R'E from a product state, averaged over the choice of B_1 , by

$$\frac{|\bar{R}| \cdot |\bar{E}|}{|B_1|^2} \operatorname{tr} \left(\boldsymbol{\sigma}_{\bar{B}}^2 \right) = \frac{2^{n(H(R) + H(E) - H(B) + o(1))}}{|B_1|^2} = \frac{2^{n(I(R;E) + o(1))}}{|B_1|^2}; \quad (10.342)$$

hence the bound also applies for some particular way of choosing B_1 . This choice defines the code used by Alice and Bob in a protocol which consumes

$$\log_2 |B_1| = \frac{n}{2} I(R; E) + o(n) \tag{10.343}$$

ebits of entanglement, and conveys from Alice to Bob

$$nH(B) - \frac{n}{2}I(R;E) - o(n) = \frac{n}{2}I(R;B) - o(n)$$
(10.344)

high-fidelity qubits. This proves the father resource inequality.

10.9.4 Quantum channel capacity revisited

In §10.8.1 we showed that the coherent information is an achievable rate for quantum communication over a noisy quantum channel. That derivation, a corollary of the father resource inequality, applied to a catalytic setting, in which shared entanglement between sender and receiver can be borrowed and later repaid. It is useful to see that the same rate is achievable without catalysis, a result we can derive from an alternative version of the decoupling inequality.

This version applies to the setting depicted here:



A density operator ρ_A for system A, with purification ψ_{RA} , is transmitted through a channel $\mathcal{N}^{A\to B}$ which has the isometric dilation $U^{A\to BE}$. The reference system R has a decomposition into subsystems R_1R_2 . We apply a random unitary transformation V to R, then project R_1 onto a fixed vector $|0\rangle_{R_1}$, and renormalize the resulting state. In effect, then we are projecting R onto a subspace with dimension $|R_2|$, which purifies a corresponding code subspace of A. This procedure prepares a normalized pure state ϕ_{R_2BE} , and a corresponding normalized marginal state σ_{R_2E} of R_2E .

If R_2 decouples from E, then R_2 is purified by a subsystem of B, which means that the code subspace of A can be recovered by a decoder applied to B. A sufficient condition for approximate decoupling can be derived from the inequality

$$\left(\mathbb{E}_{\boldsymbol{V}}\left[\|\boldsymbol{\sigma}_{R_{2}E}(\boldsymbol{V})-\boldsymbol{\sigma}_{R_{2}}^{\max}\otimes\boldsymbol{\sigma}_{E}\|_{1}\right]\right)^{2} \leq |R_{2}|\cdot|E| \operatorname{tr}\left(\boldsymbol{\sigma}_{RE}^{2}\right).$$
(10.345)

Eq.(10.345) resembles eq.(10.320) and can be derived by a similar method. Note that the right-hand side of eq.(10.345) is enhanced by a factor of $|R_1|$ relative to the right-hand side of eq.(10.320). This factor arises because after projecting R_1 onto the fixed state $|0\rangle$ we need to renormalize the state by multiplying by $|R_1|$, while on the other hand the projection suppresses the expected distance squared from a product state by a factor $|R_1|$.

In the i.i.d. setting where the noisy channel is used n times, we consider $\phi_{BBE}^{\otimes n}$, and

project onto the jointly typical subspaces \overline{R} , \overline{B} , \overline{E} of \mathbb{R}^n , \mathbb{B}^n , \mathbb{E}^n respectively, succeeding with high probability. We choose a code by projecting \overline{R} onto a random subspace with dimension $|R_2|$. Then, the right-hand side of eq.(10.345) becomes

$$|R_2| \cdot 2^{n(H(E) - H(B) + o(1))}, \tag{10.346}$$

and since the inequality holds when we average uniformly over V, it surely holds for some particular V. That unitary defines a code which achieves decoupling and has the rate

$$\frac{1}{n}\log_2|R_2| = H(E) - H(B) - o(1) = I_c(R \mid B) - o(1).$$
(10.347)

Hence the coherent information is an achievable rate for high-fidelity quantum communication over the noisy channel.

10.9.5 Black holes as mirrors

As our final application of the decoupling inequality, we consider a highly idealized model of black hole dynamics. Suppose that Alice holds a k-qubit system A which she wants to conceal from Bob. To be safe, she discards her qubits by tossing them into a large black hole, where she knows Bob will not dare to follow. The black hole B is an (n-k)-qubit system, which grows to n qubits after merging with A, where n is much larger than k.

Black holes are not really completely black — they emit Hawking radiation. But qubits leak out of an evaporating black hole very slowly, at a rate per unit time which scales like $n^{-1/2}$. Correspondingly, it takes time $\Theta(n^{3/2})$ for the black hole to radiate away a significant fraction of its qubits. Because the black hole Hilbert space is so enormous, this is a very long time, about 10^{67} years for a solar mass black hole, for which $n \approx 10^{78}$. Though Alice's qubits might not remain secret forever, she is content knowing that they will be safe from Bob for 10^{67} years.

But in her haste, Alice fails to notice that her black hole is very, very old. It has been evaporating for so long that it has already radiated away more than half of its qubits. Let's assume that the joint state of the black hole and its emitted radiation is pure, and furthermore that the radiation is a Haar-random subsystem of the full system.

Because the black hole B is so old, |B| is much smaller than the dimension of the radiation subsystem; therefore, as in eq.(10.322), we expect the state of B to be very nearly maximally mixed with high probability. We denote by R_B the subsystem of the emitted radiation which purifies B; thus the state of BR_B is very nearly maximally entangled. We assume that R_B has been collected by Bob and is under his control.

To keep track of what happens to Alice's k qubits, we suppose that her k-qubit system A is maximally entangled with a reference system R_A . After A enters the black hole, Bob waits for a while, until the k'-qubit system A' is emitted in the black hole's Hawking radiation. After retrieving A', Bob hopes to recover the purification of R_A by applying a suitable decoding map to $A'R_B$. Can he succeed?

We've learned that Bob can succeed with high fidelity if the remaining black hole system B' decouples from Alice's reference system R_A . Let's suppose that the qubits emitted in the Hawking radiation are chosen randomly; that is, A' is a Haar-random k'-qubit subsystem of the *n*-qubit system AB, as depicted here:



The double lines indicate the very large systems B and B', and single lines the smaller systems A and A'. Because the radiated qubits are random, we can determine whether R_AB' decouples using the decoupling inequality, which for this case becomes

$$\mathbb{E}_{\boldsymbol{U}}\left[\|\boldsymbol{\sigma}_{B'R_{A}}(\boldsymbol{U})-\boldsymbol{\sigma}_{B'}^{\max}\otimes\boldsymbol{\sigma}_{R_{A}}\|_{1}\right] \leq \sqrt{\frac{|ABR_{A}|}{|A'|^{2}}} \operatorname{tr}\left(\boldsymbol{\sigma}_{ABR_{A}}^{2}\right).$$
(10.348)

Because the state of AR_A is pure, and B is maximally entangled with R_B , we have $\operatorname{tr}\left(\boldsymbol{\sigma}_{ABR_A}^2\right) = 1/|B|$, and therefore the Haar-averaged L^1 distance of $\boldsymbol{\sigma}_{B'R_A}$ from a product state is bounded above by

$$\sqrt{\frac{|AR_A|}{|A'|^2}} = \frac{|A|}{|A'|}.$$
(10.349)

Thus, if Bob waits for only k' = k + c qubits of Hawking radiation to be emitted after Alice tosses in her k qubits, Bob can decode her qubits with excellent fidelity $F \ge 1-2^{-c}$.

Alice made a serious mistake. Rather than waiting for $\Omega(n)$ qubits to emerge from the black hole, Bob can already decode Alice's secret quite well when he has collected just a few more than k qubits. And Bob is an excellent physicist, who knows enough about black hole dynamics to infer the encoding unitary transformation U, information he uses to find the right decoding map.

We could describe the conclusion, more prosaically, by saying that the random unitary U applied to AB encodes a good quantum error-correcting code, which achieves high-fidelity entanglement-assisted transmission of quantum information though an erasure channel with a high erasure probability. Of the n input qubits, only k' randomly selected qubits are received by Bob; the rest remain inside the black hole and hence are inaccessible. The input qubits, then, are erased with probability p = (n - k')/n, while nearly error-free qubits are recovered from the input qubits at a rate

$$R = \frac{k}{n} = 1 - p - \frac{k' - k}{n}; \tag{10.350}$$

in the limit $n \to \infty$ with c = k' - k fixed, this rate approaches 1 - p, the entanglementassisted quantum capacity of the erasure channel.

So far, we've assumed that the emitted system A' is a randomly selected subsystem of AB. That won't be true for a real black hole. However, it is believed that the internal dynamics of actually black holes mixes quantum information quite rapidly (the *fast scrambling conjecture*). For a black hole with temperature T, it takes time of order \hbar/kT for each qubit to be emitted in the Hawking radiation, and a time longer by only a factor of log n for the dynamics to mix the black hole degrees of freedom sufficiently for our decoupling estimate to hold with reasonable accuracy. For a solar mass black 10.10 Summary

hole, Alice's qubits are revealed just a few milliseconds after she deposits them, much faster than the 10^{67} years she had hoped for! Because Bob holds the system R_B which purifies B, and because he knows the right decoding map to apply to $A'R_B$, the black hole behaves like an information mirror — Alice's qubits bounce right back!

If Alice is more careful, she will dump her qubits into a young black hole instead. If we assume that the initial black hole B is in a pure state, then σ_{ABR_A} is also pure, and the Haar-averaged L^1 distance of $\sigma_{B'R_A}$ from a product state is bounded above by

$$\sqrt{\frac{|ABR_A|}{|A'|^2}} = \frac{2^{n+k}}{2^{2k'}} = \frac{1}{2^c}$$
(10.351)

after

$$k' = \frac{1}{2}(n+k+c) \tag{10.352}$$

qubits are emitted. In this case, Bob needs to wait a long time, until more than half of the qubits in AB are radiated away. Once Bob has acquired k + c more qubits than the number still residing in the black hole, he is empowered to decode Alice's k qubits with fidelity $F \ge 1 - 2^{-c}$. In fact, there is nothing special about Alice's subsystem A; by adjusting his decoding map appropriately, Bob can decode any k qubits he chooses from among the n qubits in the initial black hole AB.

There is far more to learn about quantum information processing by black holes, an active topic of current research (as of this writing in 2016), but we will not delve further into this fascinating topic here. We can be confident, though, that the tools and concepts of quantum information theory discussed in this book will be helpful for addressing the many unresolved mysteries of quantum gravity.

10.10 Summary

Shannon entropy and classical data compression. The Shannon entropy of an ensemble $X = \{x, p(x)\}$ is $H(X) \equiv \langle -\log p(x) \rangle$; it quantifies the compressibility of classical information. A message *n* letters long, where each letter is drawn independently from X, can be compressed to H(X) bits per letter (and no further), yet can still be decoded with arbitrarily good accuracy as $n \to \infty$.

Conditional entropy and information merging. The conditional entropy H(X|Y) = H(XY) - H(Y) quantifies how much the information source X can be compressed when Y is known. If n letters are drawn from XY, where Alice holds X and Bob holds Y, Alice can convey X to Bob by sending H(X|Y) bits per letter, asymptotically as $n \to \infty$.

Mutual information and classical channel capacity. The mutual information I(X;Y) = H(X) + H(Y) - H(XY) quantifies how information sources X and Y are correlated; when we learn the value of y we acquire (on the average) I(X;Y) bits of information about x, and vice versa. The capacity of a memoryless noisy classical communication channel is $C = \max_X I(X;Y)$. This is the highest number of bits per letter that can be transmitted through n uses of the channel, using the best possible code, with negligible error probability as $n \to \infty$.

Von Neumann entropy and quantum data compression. The Von Neumann entropy of a density operator ρ is

$$H(\boldsymbol{\rho}) = -\mathrm{tr}\boldsymbol{\rho}\log\boldsymbol{\rho};\tag{10.353}$$

it quantifies the compressibility of an ensemble of pure quantum states. A message *n* letters long, where each letter is drawn independently from the ensemble $\{|\varphi(x)\rangle, p(x)\}$, can be compressed to $H(\rho)$ qubits per letter (and no further) where $\rho = \sum_X p(x) |\varphi(x)\rangle \langle \varphi(x)|$, yet can still be decoded with arbitrarily good fidelity as $n \to \infty$.

Entanglement concentration and dilution. The entanglement E of a bipartite pure state $|\psi\rangle_{AB}$ is $E = H(\rho_A)$ where $\rho_A = \operatorname{tr}_B(|\psi\rangle\langle\psi|)$. With local operations and classical communication, we can prepare n copies of $|\psi\rangle_{AB}$ from nE Bell pairs (but not from fewer), and we can distill nE Bell pairs (but not more) from n copies of $|\psi\rangle_{AB}$, asymptotically as $n \to \infty$.

Accessible information. The *Holevo chi* of an ensemble $\mathcal{E} = \{ \boldsymbol{\rho}(x), p(x) \}$ of quantum states is

$$\chi(\mathcal{E}) = H\left(\sum_{x} p(x)\boldsymbol{\rho}(x)\right) - \sum_{x} p(x)H(\boldsymbol{\rho}(x)).$$
(10.354)

The accessible information of an ensemble \mathcal{E} of quantum states is the maximal number of bits of information that can be acquired about the preparation of the state (on the average) with the best possible measurement. The accessible information cannot exceed the Holevo chi of the ensemble. The product-state capacity of a quantum channel \mathcal{N} is

$$C_1(\mathcal{N}) = \max_{\mathcal{E}} \chi(\mathcal{N}(\mathcal{E})). \tag{10.355}$$

This is the highest number of classical bits per letter that can be transmitted through n uses of the quantum channel, with negligible error probability as $n \to \infty$, assuming that each codeword is a product state.

Decoupling and quantum communication. In a tripartite pure state ϕ_{RBE} , we say that systems R and E decouple if the marginal density operator of RE is a product state, in which case R is purified by a subsystem of B. A quantum state transmitted through a noisy quantum channel $\mathcal{N}^{A\to B}$ (with isometric dilation $U^{A\to BE}$) can be accurately decoded if a reference system R which purifies channel's input A nearly decouples from the channel's environment E.

Father and mother protocols. The *father and mother resource inequalities* specify achievable rates for entanglement-assisted quantum communication and quantum state transfer, respectively. Both follow from the *decoupling inequality*, which establishes a sufficient condition for approximate decoupling in a tripartite mixed state. By combining the father and mother protocols with superdense coding and teleportation, we can derive achievable rates for other protocols, including entanglement-assisted classical communication, quantum communication, entanglement distillation, and quantum state merging.

Homage to Ben Schumacher:

Ben. He rocks. I remember When He showed me how to fit A qubit In a small box.

I wonder how it feels

To be compressed. And then to pass A fidelity test.

Or does it feel At all, and if it does Would I squeal Or be just as I was?

If not undone I'd become as I'd begun And write a memorandum On being random. Had it felt like a belt Of rum?

And might it be predicted That I'd become addicted, Longing for my session Of compression?

I'd crawl To Ben again. And call, Put down your pen! Don't stall! Make me small!

10.11 Bibliographical Notes

Cover and Thomas [2] is an excellent textbook on classical information theory. Shannon's original paper [3] is still very much worth reading.

Nielsen and Chuang [4] provide a clear introduction to some aspects of quantum Shannon theory. Wilde [1] is a more up-to-date and very thorough account.

Properties of entropy are reviewed in [5]. Strong subadditivity of Von Neumann entropy was proven by Lieb and Ruskai [6], and the condition for equality was derived by Hayden *et al.* [7]. The connection between separability and majorization was pointed out by Nielsen and Kempe [8].

Bekenstein's entropy bound was formulated in [9] and derived by Casini [10]. Entropic uncertainty relations are reviewed in [11], and I follow their derivation. The original derivation, by Maassen and Uffink [12] uses different methods.

Schumacher compression was first discussed in [13, 14], and Bennett *et al.* [15] devised protocols for entanglement concentration and dilution. Measures of mixed-state entanglement are reviewed in [16]. The reversible theory of mixed-state entanglement was formulated by Brandão and Plenio [17]. Squashed entanglement was introduced by Christandl and Winter [18], and its monogamy discussed by Koashi and Winter [19]. Brandão, Christandl, and Yard [20] showed that squashed entanglement is positive for any nonseparable bipartite state. Doherty, Parrilo, and Spedalieri [21] showed that every nonseparable bipartite state fails to be k-extendable for some finite k.

The Holevo bound was derived in [22]. Peres-Wootters coding was discussed in [23]. The product-state capacity formula was derived by Holevo [24] and by Schumacher

and Westmoreland [25]. Hastings [26] showed that Holevo chi can be superadditive. Horodecki, Shor, and Ruskai [27] introduced entanglement-breaking channels, and additivity of Holevo chi for these channels was shown by Shor [28].

Necessary and sufficient conditions for quantum error correction were formulated in terms of the decoupling principle by Schumacher and Nielsen [29]; that (regularized) coherent information is an upper bound on quantum capacity was shown by Schumacher [30], Schumacher and Nielsen [29], and Barnum *et al.* [31]. That coherent information is an achievable rate for quantum communication was conjectured by Lloyd [32] and by Schumacher [30], then proven by Shor [33] and by Devetak [34]. Devetak and Winter [35] showed it is also an achievable rate for entanglement distillation. The quantum Fano inequality was derived by Schumacher [30].

Approximate decoupling was analyzed by Schumacher and Westmoreland [36], and used to prove capacity theorems by Devetak [34], by Horodecki *et al.* [37], by Hayden *et al.* [38], and by Abeyesinghe *et al.* [39]. The entropy of Haar-random subsystems had been discussed earlier, by Lubkin [40], Lloyd and Pagels [41], and Page [42]. Devetak, Harrow, and Winter [43, 44] introduced the mother and father protocols and their descendants. Devatak and Shor [45] introduced degradable quantum channels and proved that coherent information is additive for these channels. Bennett *et al.* [46, 47] found the single-letter formula for entanglement-assisted classical capacity. Superadditivity of coherent information was discovered by Shor and Smolin [48] and by DiVincenzo *et al.* [49]. Smith and Yard [50] found extreme examples of superadditivity, in which two zero-capacity channels have nonzero capacity when used jointly. The achievable rate for state merging was derived by Horodecki *et al.* [37], and used by them to prove strong subadditivity of Von Neumann entropy.

Decoupling was applied to Landuaer's principle by Renner *et al.* [51], and to black holes by Hayden and Preskill [52]. The fast scrambling conjecture was proposed by Sekino and Susskind [53].

Exercises

10.1 Positivity of quantum relative entropy

- a) Show that $\ln x \le x 1$ for all positive real x, with equality iff x = 1.
- b) The (classical) relative entropy of a probability distribution $\{p(x)\}$ relative to $\{q(x)\}$ is defined as

$$D(p \parallel q) \equiv \sum_{x} p(x) \left(\log p(x) - \log q(x) \right) .$$
 (10.356)

Show that

$$D(p \parallel q) \ge 0 , \qquad (10.357)$$

with equality iff the probability distributions are identical. **Hint**: Apply the inequality from (a) to $\ln(q(x)/p(x))$.

c) The quantum relative entropy of the density operator ρ with respect to σ is defined as

$$D(\boldsymbol{\rho} \parallel \boldsymbol{\sigma}) = \operatorname{tr} \boldsymbol{\rho} (\log \boldsymbol{\rho} - \log \boldsymbol{\sigma}) \quad . \tag{10.358}$$

Let $\{p_i\}$ denote the eigenvalues of ρ and $\{q_a\}$ denote the eigenvalues of σ .

Exercises

Show that

$$D(\boldsymbol{\rho} \parallel \boldsymbol{\sigma}) = \sum_{i} p_i \left(\log p_i - \sum_{a} D_{ia} \log q_a \right) , \qquad (10.359)$$

where D_{ia} is a doubly stochastic matrix. Express D_{ia} in terms of the eigenstates of ρ and σ . (A matrix is doubly stochastic if its entries are nonnegative real numbers, where each row and each column sums to one.)

d) Show that if D_{ia} is doubly stochastic, then (for each i)

$$\log\left(\sum_{a} D_{ia} q_{a}\right) \ge \sum_{a} D_{ia} \log q_{a} , \qquad (10.360)$$

with equality only if $D_{ia} = 1$ for some a.

e) Show that

$$D(\boldsymbol{\rho} \parallel \boldsymbol{\sigma}) \ge D(p \parallel r) , \qquad (10.361)$$

where $r_i = \sum_a D_{ia} q_a$.

f) Show that $D(\rho \parallel \sigma) \ge 0$, with equality iff $\rho = \sigma$.

10.2 Properties of Von Neumann entropy

a) Use nonnegativity of quantum relative entropy to prove the *subadditivity* of Von Neumann entropy

$$H(\boldsymbol{\rho}_{AB}) \le H(\boldsymbol{\rho}_A) + H(\boldsymbol{\rho}_B), \qquad (10.362)$$

with equality iff $\rho_{AB} = \rho_A \otimes \rho_B$. Hint: Consider the relative entropy of ρ_{AB} and $\rho_A \otimes \rho_B$.

b) Use subadditivity to prove the concavity of the Von Neumann entropy:

$$H(\sum_{x} p_x \boldsymbol{\rho}_x) \ge \sum_{x} p_x H(\boldsymbol{\rho}_x) . \qquad (10.363)$$

Hint: Consider

$$\boldsymbol{\rho}_{AB} = \sum_{x} p_x \left(\boldsymbol{\rho}_x \right)_A \otimes \left(|x\rangle \langle x| \right)_B \quad , \tag{10.364}$$

where the states $\{|x\rangle_B\}$ are mutually orthogonal.

c) Use the condition

$$H(\boldsymbol{\rho}_{AB}) = H(\boldsymbol{\rho}_A) + H(\boldsymbol{\rho}_B) \quad \text{iff} \quad \boldsymbol{\rho}_{AB} = \boldsymbol{\rho}_A \otimes \boldsymbol{\rho}_B \tag{10.365}$$

to show that, if all p_x 's are nonzero,

$$H\left(\sum_{x} p_{x} \boldsymbol{\rho}_{x}\right) = \sum_{x} p_{x} H(\boldsymbol{\rho}_{x})$$
(10.366)

iff all the ρ_x 's are identical.

10.3 Monotonicity of quantum relative entropy

Quantum relative entropy has a property called *monotonicity*:

$$D(\boldsymbol{\rho}_A \| \boldsymbol{\sigma}_A) \le D(\boldsymbol{\rho}_{AB} \| \boldsymbol{\sigma}_{AB}); \qquad (10.367)$$

The relative entropy of two density operators on a system AB cannot be less than the induced relative entropy on the subsystem A.

- a) Use monotonicity of quantum relative entropy to prove the strong subadditivity property of Von Neumann entropy. **Hint**: On a tripartite system *ABC*, consider the relative entropy of ρ_{ABC} and $\rho_A \otimes \rho_{BC}$.
- b) Use monotonicity of quantum relative entropy to show that the action of a quantum channel \mathcal{N} cannot increase relative entropy:

$$D(\mathcal{N}(\boldsymbol{\rho}) \| \mathcal{N}(\boldsymbol{\sigma}) \le D(\boldsymbol{\rho} \| \boldsymbol{\sigma}), \qquad (10.368)$$

Hint: Recall that any quantum channel has an isometric dilation.

10.4 The Peres–Wootters POVM.

Consider the Peres–Wootters information source described in §10.6.4 of the lecture notes. It prepares one of the three states

$$|\Phi_a\rangle = |\varphi_a\rangle \otimes |\varphi_a\rangle, \quad a = 1, 2, 3,$$
 (10.369)

each occurring with a priori probability $\frac{1}{3}$, where the $|\varphi_a\rangle$'s are defined in eq.(10.214).

a) Express the density matrix

$$\boldsymbol{\rho} = \frac{1}{3} \left(\sum_{a} |\Phi_a\rangle \langle \Phi_a| \right), \qquad (10.370)$$

in terms of the Bell basis of maximally entangled states $\{|\phi^{\pm}\rangle, |\psi^{\pm}\rangle\}$, and compute $H(\rho)$.

- b) For the three vectors $|\Phi_a\rangle$, a = 1, 2, 3, construct the "pretty good measurement" defined in eq.(10.227). (Again, expand the $|\Phi_a\rangle$'s in the Bell basis.) In this case, the PGM is an orthogonal measurement. Express the elements of the PGM basis in terms of the Bell basis.
- c) Compute the mutual information of the PGM outcome and the preparation.

10.5 Separability and majorization

The hallmark of entanglement is that in an entangled state the whole is less random than its parts. But in a separable state the correlations are essentially classical and so are expected to adhere to the classical principle that the parts are less disordered than the whole. The objective of this problem is to make this expectation precise by showing that if the bipartite (mixed) state ρ_{AB} is separable, then

$$\lambda(\boldsymbol{\rho}_{AB}) \prec \lambda(\boldsymbol{\rho}_{A}) , \quad \lambda(\boldsymbol{\rho}_{AB}) \prec \lambda(\boldsymbol{\rho}_{B}) .$$
 (10.371)

Here $\lambda(\rho)$ denotes the vector of eigenvalues of ρ , and \prec denotes majorization.

A separable state can be realized as an ensemble of pure product states, so that if ρ_{AB} is separable, it may be expressed as

$$\boldsymbol{\rho}_{AB} = \sum_{a} p_a |\psi_a\rangle \langle\psi_a| \otimes |\varphi_a\rangle \langle\varphi_a| . \qquad (10.372)$$

We can also diagonalize ρ_{AB} , expressing it as

$$\boldsymbol{\rho}_{AB} = \sum_{j} r_{j} |e_{j}\rangle \langle e_{j}| , \qquad (10.373)$$

where $\{|e_j\rangle\}$ denotes an orthonormal basis for AB; then by the HJW theorem, there is a unitary matrix V such that

$$\sqrt{r_j}|e_j\rangle = \sum_a V_{ja}\sqrt{p_a}|\psi_a\rangle \otimes |\varphi_a\rangle . \qquad (10.374)$$

Also note that ρ_A can be diagonalized, so that

$$\boldsymbol{\rho}_A = \sum_a p_a |\psi_a\rangle \langle \psi_a| = \sum_\mu s_\mu |f_\mu\rangle \langle f_\mu| ; \qquad (10.375)$$

here $\{|f_{\mu}\rangle\}$ denotes an orthonormal basis for A, and by the HJW theorem, there is a unitary matrix U such that

$$\sqrt{p_a}|\psi_a\rangle = \sum_{\mu} U_{a\mu}\sqrt{s_{\mu}}|f_{\mu}\rangle . \qquad (10.376)$$

Now show that there is a doubly stochastic matrix D such that

$$r_j = \sum_{\mu} D_{j\mu} s_{\mu} . (10.377)$$

That is, you must check that the entries of $D_{j\mu}$ are real and nonnegative, and that $\sum_j D_{j\mu} = 1 = \sum_{\mu} D_{j\mu}$. Thus we conclude that $\lambda(\rho_{AB}) \prec \lambda(\rho_A)$. Just by interchanging A and B, the same argument also shows that $\lambda(\rho_{AB}) \prec \lambda(\rho_B)$.

Remark: Note that it follows from the Schur concavity of Shannon entropy that, if ρ_{AB} is separable, then the von Neumann entropy has the properties $H(AB) \ge H(A)$ and $H(AB) \ge H(B)$. Thus, for separable states, conditional entropy is nonnegative: $H(A|B) = H(AB) - H(B) \ge 0$ and $H(B|A) = H(AB) - H(A) \ge 0$. In contrast, if H(A|B) is negative, then according to the hashing inequality the state of AB has positive distillable entanglement -H(A|B), and therefore is surely not separable.

10.6 Additivity of squashed entanglement

Suppose that Alice holds systems A, A' and Bob holds systems B, B'. How is the entanglement of AA' with BB' related to the entanglement of A with B and A' with B'? In this problem we will show that the squashed entanglement is superadditive,

$$E_{\rm sq}(\boldsymbol{\rho}_{ABA'B'}) \ge E_{\rm sq}(\boldsymbol{\rho}_{AB}) + E_{\rm sq}(\boldsymbol{\rho}_{A'B'}) \tag{10.378}$$

and is strictly additive for a tensor product,

$$E_{\rm sq}(\boldsymbol{\rho}_{AB} \otimes \boldsymbol{\rho}_{A'B'}) = E_{\rm sq}(\boldsymbol{\rho}_{AB}) + E_{\rm sq}(\boldsymbol{\rho}_{A'B'}). \tag{10.379}$$

a) Use the chain rule for mutual information eq.(10.196) and eq.(10.197) and the nonnegativity of quantum conditional mutual information to show that

$$I(AA'; BB'|C) \ge I(A; B|C) + I(A'; B'|AC),$$
(10.380)

and show that eq.(10.378) follows.

b) Show that for any extension $\rho_{ABC} \otimes \rho_{A'B'C'}$ of the product state $\rho_{AB} \otimes \rho_{A'B'}$, we have

$$I(AA'; BB'|CC') \le I(A; B|C) + I(A'; B'|C').$$
(10.381)

Conclude that

$$E_{\rm sq}(\boldsymbol{\rho}_{AB} \otimes \boldsymbol{\rho}_{A'B'}) \le E_{\rm sq}(\boldsymbol{\rho}_{AB}) + E_{\rm sq}(\boldsymbol{\rho}_{A'B'}), \qquad (10.382)$$

which, when combined with eq.(10.378), implies eq.(10.379).

10.7 The first law of Von Neumann entropy

Writing the density operator in terms of its modular Hamiltonian K as in §10.2.6,

$$\boldsymbol{\rho} = \frac{e^{-\boldsymbol{K}}}{\operatorname{tr}\left(e^{-\boldsymbol{K}}\right)},\tag{10.383}$$

consider how the entropy $S(\rho) = -\operatorname{tr}(\rho \ln \rho)$ changes when the density operator is perturbed slightly:

$$\boldsymbol{\rho} \to \boldsymbol{\rho}' = \boldsymbol{\rho} + \delta \boldsymbol{\rho}. \tag{10.384}$$

Since ρ and ρ' are both normalized density operators, we have tr $(\delta \rho) = 0$. Show that

$$S(\boldsymbol{\rho}') - S(\boldsymbol{\rho}) = \operatorname{tr}\left(\boldsymbol{\rho}'\boldsymbol{K}\right) - \operatorname{tr}\left(\boldsymbol{\rho}\boldsymbol{K}\right) + O\left(\left(\delta\boldsymbol{\rho}\right)^{2}\right); \qquad (10.385)$$

that is,

$$\delta S = \delta \langle \mathbf{K} \rangle \tag{10.386}$$

to first order in the small change in ρ . This statement generalizes the first law of thermodynamics; for the case of a thermal density operator with $\mathbf{K} = T^{-1}\mathbf{H}$ (where \mathbf{H} is the Hamiltonian and T is the temperature), it becomes the more familiar statement

$$\delta E = \delta \langle \boldsymbol{H} \rangle = T \delta S. \tag{10.387}$$

10.8 Information gain for a quantum state drawn from the uniform ensemble

Suppose Alice prepares a quantum state drawn from the ensemble $\{\rho(x), p(x)\}$ and Bob performs a measurement $\{E(y)\}$ yielding outcome y with probability $p(y|x) = \operatorname{tr} (E(y)\rho(x))$. As noted in §10.6.1, Bob's information gain about Alice's preparation is the mutual information I(X;Y) = H(X) - H(X|Y). If x is a continuous variable, while y is discrete, it is more convenient to use the symmetry of mutual information to write I(X;Y) = H(Y) - H(Y|X), where

$$H(Y|X) = \sum_{y} \int dx \cdot p(x) \cdot p(y|x) \cdot \log p(y|x); \qquad (10.388)$$

here p(x) is a probability density (that is, p(x)dx is the probability for x to lie in the interval [x, x + dx]).

For example, suppose that Alice prepares an arbitrary pure state $|\varphi\rangle$ chosen from the uniform ensemble in a *d*-dimensional Hilbert space, and Bob performs an orthogonal measurement projecting onto the basis $\{|e_y\rangle\}$, hoping to learn something about what Alice prepared. Then Bob obtains outcome *y* with probability

$$p(y|\theta) = |\langle e_y|\varphi\rangle|^2 \equiv \cos^2\theta \tag{10.389}$$

where θ is the angle between $|\varphi\rangle$ and $|e_y\rangle$. Because Alice's ensemble is uniform, Bob's outcomes are also uniformly distributed; hence $H(Y) = \log d$. Furthermore, the measurement outcome y reveals only information about θ ; Bob learns nothing Exercises

else about $|\varphi\rangle$. Therefore, eq.(10.388) implies that the information gain may be expressed as

$$I(X;Y) = \log d - d \int d\theta \cdot p(\theta) \cdot \cos^2 \theta \cdot \log \cos^2 \theta.$$
 (10.390)

Here $p(\theta)d\theta$ is the probability density for the vector $|\varphi\rangle$ to point in a direction making angle θ with the axis $|e_y\rangle$, where $0 \le \theta \le \pi/2$.

a) Show that

$$p(\theta) \cdot d\theta = -(d-1) \left[1 - \cos^2 \theta\right]^{d-2} \cdot d\cos^2 \theta.$$
 (10.391)

Hint: Choose a basis in which the fixed axis $|e_y\rangle$ is

$$|e_y\rangle = (1, \vec{0})$$
 (10.392)

and write

$$|\varphi\rangle = (e^{i\phi}\cos\theta, \psi^{\perp}), \qquad (10.393)$$

where $\theta \in [0, \pi/2]$, and $|\psi^{\perp}\rangle$ denotes a complex (d-1)-component vector with length $\sin \theta$. Now note that the phase ϕ resides on a circle of radius $\cos \theta$ (and hence circumference $2\pi \cos \theta$), while $|\psi^{\perp}\rangle$ lies on a sphere of radius $\sin \theta$ (thus the volume of the sphere, up to a multiplicative numerical constant, is $\sin^{2d-3} \theta$).

b) Now evaluate the integral eq. (10.390) to show that the information gain from the measurement, in *nats*, is

$$I(X;Y) = \ln d - \left(\frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{d}\right) .$$
 (10.394)

(Information is expressed in nats if logarithms are natural logarithms; I in nats is related to I in bits by $I_{\text{bits}} = I_{\text{nats}}/\ln 2$.) Hint: To evaluate the integral

$$\int_0^1 dx (1-x)^p x \ln x , \qquad (10.395)$$

observe that

$$x\ln x = \frac{d}{ds}x^s\Big|_{s=1}, \qquad (10.396)$$

and then calculate $\int_0^1 dx (1-x)^p x^s$ by integrating by parts repeatedly. c) Show that in the limit of large d, the information gain, in *bits*, approaches

$$I_{d=\infty} = \frac{1-\gamma}{\ln 2} = .60995\dots, \qquad (10.397)$$

where $\gamma = .57721...$ is Euler's constant.

Our computed value of H(Y|X) may be interpreted in another way: Suppose we fix an orthogonal measurement, choose a typical state, and perform the measurement repeatedly on that chosen state. Then the measurement outcomes will not be uniformly distributed. Instead the entropy of the outcomes will fall short of maximal by .60995 bits, in the limit of large Hilbert space dimension.

10.9 Fano's inequality

Suppose $X = \{x, p(x)\}$ is a probability distribution for a letter x drawn from an alphabet of d possible letters, and that XY is the joint distribution for x and another random variable y which is correlated with x. Upon receiving y we estimate the value of x by evaluating a function $\hat{x}(y)$. We may anticipate that if our estimate is usually correct, then the conditional entropy H(X|Y) must be small. In this problem we will confirm that expectation.

Let $e \in \{0,1\}$ denote a binary random variable which takes the value e = 0if $x = \hat{x}(y)$ and takes the value e = 1 if $x \neq \hat{x}(y)$, and let XYE denote the joint distribution for x, y, e. The error probability P_e is the probability that e = 1, averaged over this distribution. Our goal is to derive an upper bound on H(X|Y)depending on P_e .

a) Show that

$$H(X|Y) = H(X|YE) + H(E|Y) - H(E|XY).$$
(10.398)

Note that H(E|XY) = 0 because e is determined when x and y are know, and that $H(E|Y) \leq H(E)$ because mutual information is nonnegative. Therefore,

$$H(X|Y) \le H(X|YE) + H(E).$$
 (10.399)

b) Noting that

$$H(X|YE) = p(e=0)H(X|Y,e=0) + p(e=1)H(X|Y,e=1), \quad (10.400)$$

and that H(X|Y, e = 0) = 0 (because $x = \hat{x}(y)$ is determined by y when there is no error), show that

$$H(X|YE) \le P_e \log_2(d-1).$$
 (10.401)

c) Finally, show that

$$H(X|Y) \le H_2(P_e) + P_e \log_2(d-1),$$
 (10.402)

which is *Fano's inequality*.

d) Use Fano's inequality to derive eq.(10.50), hence completing the proof that the classical channel capacity C is an upper bound on achievable rates for communication over a noisy channel with negligible error probability.

10.10 A quantum version of Fano's inequality

a) In a d-dimensional system, suppose a density operator ρ approximates the pure state $|\psi\rangle$ with fidelity

$$F = \langle \psi | \boldsymbol{\rho} | \psi \rangle = 1 - \varepsilon. \tag{10.403}$$

Show that

$$H(\boldsymbol{\rho}) \le H_2(\varepsilon) + \varepsilon \log_2(d-1). \tag{10.404}$$

Hint: Recall that if a complete orthogonal measurement performed on the state ρ has distribution of outcomes X, then $H(\rho) \leq H(X)$, where H(X) is the Shannon entropy of X.

82
Exercises

b) As in §10.7.2, suppose that the noisy channel $\mathcal{N}^{A \to B}$ acts on the pure state ψ_{RA} , and is followed by the decoding map $\mathcal{D}^{B \to C}$. Show that

$$H(R)_{\rho} - I_c(R \mid B)_{\rho} \le 2H(RC)_{\sigma}, \qquad (10.405)$$

where

$$\boldsymbol{\rho}_{RB} = \mathcal{N}(\psi_{RA}), \quad \boldsymbol{\sigma}_{RC} = \mathcal{D} \circ \mathcal{N}(\psi_{RA}). \tag{10.406}$$

Therefore, if the decoder's output (the state of RC) is almost pure, then the coherent information of the channel \mathcal{N} comes close to matching its input entropy. **Hint**: Use the data processing inequality $I_c(R \ C)_{\sigma} \leq I_c(R \ B)_{\rho}$ and the subadditivity of von Neumann entropy. It is convenient to consider the joint pure state of the reference system, the output, and environments of the dilations of \mathcal{N} and \mathcal{D} .

c) Suppose that the decoding map recovers the channel input with high fidelity,

$$F(\mathcal{D} \circ \mathcal{N}(\psi_{RA}), \psi_{RC}) = 1 - \varepsilon.$$
(10.407)

Show that

$$H(R)_{\boldsymbol{\rho}} - I_c(R \mid B)_{\boldsymbol{\rho}} \le 2H_2(\varepsilon) + 2\varepsilon \log_2(d^2 - 1), \qquad (10.408)$$

assuming that R and C are d-dimensional. This is a quantum version of Fano's inequality, which we may use to derive an upper bound on the quantum channel capacity of \mathcal{N} .

10.11 Mother protocol for the GHZ state

The mother resource inequality expresses an asymptotic resource conversion that can be achieved if Alice, Bob, and Eve share n copies of the pure state ϕ_{ABE} : by sending $\frac{n}{2}I(A; E)$ qubits to Bob, Alice can destroy the correlations of her state with Eve's state, so that Bob alone holds the purification of Eve's state, and furthermore Alice and Bob share $\frac{n}{2}I(A; B)$ ebits of entanglement at the end of the protocol; here I(A; E) and I(A; B) denote quantum mutual informations evaluated in the state ϕ_{ABE} .

Normally, the resource conversion can be realized with arbitrarily good fidelity only in the limit $n \to \infty$. But in this problem we will see that the conversion can be perfect if Alice, Bob and Eve share only n = 2 copies of the three-qubit GHZ state

$$|\phi\rangle_{ABE} = \frac{1}{\sqrt{2}} \left(|000\rangle + |111\rangle\right) .$$
 (10.409)

The protocol achieving this perfect conversion uses the notion of *coherent classical* communication defined in Chapter 4.

a) Show that in the GHZ state $|\phi\rangle_{ABE}$, I(A; E) = I(A; B) = 1. Thus, for this state, the mother inequality becomes

$$2\langle \phi_{ABE} \rangle + [q \to q]_{AB} \ge [qq]_{AB} + 2\langle \phi'_{\tilde{B}E} \rangle . \tag{10.410}$$

b) Suppose that in the GHZ state Alice measures the Pauli operator X, gets the outcome +1 and broadcasts her outcome to Bob and Eve. What state do Bob and Eve then share? What if Alice gets the outcome -1 instead?

- c) Suppose that Alice, Bob, and Eve share just one copy of the GHZ state ϕ^{ABE} . Find a protocol such that, after one unit of *coherent classical communication* from Alice to Bob, the shared state becomes $|\phi^+\rangle_{AB} \otimes |\phi^+\rangle_{BE}$, where $|\phi^+\rangle = (|00\rangle + |11\rangle) / \sqrt{2}$ is a maximally entangled Bell pair.
- d) Now suppose that Alice, Bob, and Eve start out with two copies of the GHZ state, and suppose that Alice and Bob can borrow an ebit of entanglement, which will be repaid later, to catalyze the resource conversion. Use coherent superdense coding to construct a protocol that achieves the (catalytic) conversion eq. (10.410) perfectly.

10.12 Degradability of amplitude damping and erasure

The qubit amplitude damping channel $\mathcal{N}_{a.d.}^{A\to B}(p)$ discussed in §3.4.3 has the dilation $U^{A\to BE}$ such that

$$\begin{split} \boldsymbol{U} : & |0\rangle_A \mapsto |0\rangle_B \otimes |0\rangle_E, \\ & |1\rangle_A \mapsto \sqrt{1-p} \ |1\rangle_B \otimes |0\rangle_E + \sqrt{p} \ |0\rangle_B \otimes |1\rangle_E; \end{split}$$

a qubit in its "ground state" $|0\rangle_A$ is unaffected by the channel, while a qubit in the "excited state" $|1\rangle_A$ decays to the ground state with probability p, and the decay process excites the environment. Note that U is invariant under interchange of systems B and E accompanied by transformation $p \leftrightarrow (1-p)$. Thus the channel complementary to $\mathcal{N}_{\mathrm{a.d.}}^{A \to B}(p)$ is $\mathcal{N}_{\mathrm{a.d.}}^{A \to E}(1-p)$.

a) Show that $\mathcal{N}_{\text{a.d.}}^{A \to B}(p)$ is degradable for $p \leq 1/2$. Therefore, the quantum capacity of the amplitude damping channel is its optimized one-shot coherent information. **Hint**: It suffices to show that

$$\mathcal{N}_{\text{a.d.}}^{A \to E}(1-p) = \mathcal{N}_{\text{a.d.}}^{B \to E}(q) \circ \mathcal{N}_{\text{a.d.}}^{A \to B}(p), \qquad (10.411)$$

where $0 \le q \le 1$.

The erasure channel $\mathcal{N}_{\text{erase}}^{A \to B}(p)$ has the dilation $U^{A \to BE}$ such that

$$\boldsymbol{U}: |\psi\rangle_A \mapsto \sqrt{1-p} \; |\psi\rangle_B \otimes |e\rangle_E + \sqrt{p} \; |e\rangle_B \otimes |\psi\rangle_E; \tag{10.412}$$

Alice's system passes either to Bob (with probability 1-p) or to Eve (with probability p), while the other party receives the "erasure symbol" $|e\rangle$, which is orthogonal to Alice's Hilbert space. Because U is invariant under interchange of systems Band E accompanied by transformation $p \leftrightarrow (1-p)$, the channel complementary to $\mathcal{N}_{\text{erase}}^{A \to B}(p)$ is $\mathcal{N}_{\text{erase}}^{A \to E}(1-p)$.

b) Show that $\mathcal{N}_{\text{erase}}^{A \to B}(p)$ is degradable for $p \leq 1/2$. Therefore, the quantum capacity of the amplitude damping channel is its optimized one-shot coherent information. **Hint**: It suffices to show that

$$\mathcal{N}_{\text{erase}}^{A \to E}(1-p) = \mathcal{N}_{\text{erase}}^{B \to E}(q) \circ \mathcal{N}_{\text{erase}}^{A \to B}(p), \qquad (10.413)$$

where $0 \le q \le 1$.

c) Show that for $p \leq 1/2$ the quantum capacity of the erasure channel is

$$Q(\mathcal{N}_{\text{erase}}^{A \to B}(p)) = (1 - 2p)\log_2 d, \qquad (10.414)$$

where A is d-dimensional, and that the capacity vanishes for $1/2 \le p \le 1$.

Exercises

10.13 Quantum Singleton bound

As noted in chapter 7, an [[n, k, d]] quantum error-correcting code (k protected qudits in a block of n qudits, with code distance d) must obey the constraint

$$n - k \ge 2(d - 1),\tag{10.415}$$

the quantum Singleton bound. This bound is actually a corollary of a stronger statement which you will prove in this exercise.

Suppose that in the pure state ϕ_{RA} the reference system R is maximally entangled with a code subspace of A, and that E_1 and E_2 are two disjoint correctable subsystems of system A (erasure of either E_1 or E_2 can be corrected). You are to show that

$$\log|A| - \log|R| \ge \log|E_1| + \log|E_2|. \tag{10.416}$$

Let E^c denote the subsystem of A complementary to E_1E_2 , so that $A = E^c E_1E_2$.

a) Recalling the error correction conditions $\rho_{RE_1} = \rho_R \otimes \rho_{E_1}$ and $\rho_{RE_2} = \rho_R \otimes \rho_{E_2}$, show that $\phi_{RE^cE_1E_2}$ has the property

$$H(R) = H(E^c) - \frac{1}{2}I(E^c; E_1) - \frac{1}{2}I(E^c; E_2).$$
(10.417)

b) Show that eq.(10.417) implies eq.(10.416).

10.14 Capacities of the depolarizing channel

Consider the depolarizing channel $\mathcal{N}_{\text{depol.}}(p)$, which acts on a pure state $|\psi\rangle$ of a single qubit according to

$$\mathcal{N}_{\text{depol.}}(p) : |\psi\rangle\langle\psi| \mapsto \left(1 - \frac{4}{3}p\right)|\psi\rangle\langle\psi| + \frac{4}{3}p \cdot \frac{1}{2}\boldsymbol{I}.$$
 (10.418)

For this channel, compute the product-state classical capacity $C_1(p)$, the entanglement-assisted classical capacity $C_E(p)$, and the one-shot quantum capacity $Q_1(p)$. Plot the results as a function of p. For what value of p does Q_1 hit zero?

The depolarizing channel is not degradable, and in fact the quantum capacity Q(p) is larger than $Q_1(p)$ when the channel is sufficiently noisy. The function Q(p) is still unknown.

10.15 Noisy superdense coding and teleportation.

a) By converting the entanglement achieved by the mother protocol into classical communication, prove the noisy superdense coding resource inequality:

Noisy
$$SD: \langle \phi_{ABE} \rangle + H(A)[q \to q] \ge I(A;B)[c \to c].$$
 (10.419)

Verify that this matches the standard noiseless superdense coding resource inequality when ϕ is a maximally entangled state of AB.

b) By converting the entanglement achieved by the mother protocol into quantum communication, prove the noisy teleportation resource inequality:

Noisy
$$TP$$
: $\langle \phi_{ABE} \rangle + I(A; B)[c \to c] \ge I_c(A \rangle B)[q \to q].$ (10.420)

Verify that this matches the standard noiseless teleportation resource inequality when ϕ is a maximally entangled state of AB.

10.16 The cost of erasure

Erasure of a bit is a process in which the state of the bit is reset to 0. Erasure is *irreversible* — knowing only the final state 0 after erasure, we cannot determine whether the initial state before erasure was 0 or 1. This irreversibility implies that erasure incurs an unavoidable thermodynamic cost. According to *Landauer's Principle*, erasing a bit at temperature T requires work $W \ge kT \log 2$. In this problem you will verify that a particular procedure for achieving erasure adheres to Landauer's Principle.

Suppose that the two states of the bit both have zero energy. We erase the bit in two steps. In the first step, we bring the bit into contact with a reservoir at temperature T > 0, and wait for the bit to come to thermal equilibrium with the reservoir. In this step the bit "forgets" its initial value, but the bit is not yet erased because it has not been reset.

We reset the bit in the second step, by slowly turning on a control field λ which splits the degeneracy of the two states. For $\lambda \geq 0$, the state 0 has energy $E_0 = 0$ and the state 1 has energy $E_1 = \lambda$. After the bit thermalizes in step one, the value of λ increases gradually from the initial value $\lambda = 0$ to the final value $\lambda = \infty$; the increase in λ is slow enough that the qubit remains in thermal equilibrium with the reservoir at all times. As λ increases, the probability P(0) that the qubit is in the state 0 approaches unity — *i.e.*, the bit is reset to the state 0, which has zero energy.

- (a) For $\lambda \neq 0$, find the probability P(0) that the qubit is in the state 0 and the probability P(1) that the qubit is in the state 1.
- (b) How much work is required to increase the control field from λ to $\lambda + d\lambda$?
- (c) How much work is expended as λ increases slowly from $\lambda = 0$ to $\lambda = \infty$? (You will have to evaluate an integral, which can be done analytically.)

10.17 Proof of the decoupling inequality

In this problem we complete the derivation of the decoupling inequality sketched in §10.9.1.

a) Verify eq.(10.329).

To derive the expression for $\mathbb{E}_{\boldsymbol{U}}[\boldsymbol{M}_{AA'}(\boldsymbol{U})]$ in eq.(10.333), we first note that the invariance property eq.(10.318) implies that $\mathbb{E}_{\boldsymbol{U}}[\boldsymbol{M}_{AA'}(\boldsymbol{U})]$ commutes with $\boldsymbol{V} \otimes \boldsymbol{V}$ for any unitary \boldsymbol{V} . Therefore, by Schur's lemma, $\mathbb{E}_{\boldsymbol{U}}[\boldsymbol{M}_{AA'}(\boldsymbol{U})]$ is a weighted sum of projections onto irreducible representations of the unitary group. The tensor product of two fundamental representations of $\boldsymbol{U}(d)$ contains two irreducible representations — the symmetric and antisymmetric tensor representations. Therefore we may write

$$\mathbb{E}_{\boldsymbol{U}}\left[\boldsymbol{M}_{AA'}(\boldsymbol{U})\right] = c_{\text{sym}} \, \boldsymbol{\Pi}_{AA'}^{(\text{sym})} + c_{\text{anti}} \, \boldsymbol{\Pi}_{AA'}^{(\text{anti})}; \quad (10.421)$$

here $\Pi_{AA'}^{(\text{sym})}$ is the orthogonal projector onto the subspace of AA' symmetric under the interchange of A and A', $\Pi_{AA'}^{(\text{anti})}$ is the projector onto the antisymmetric subspace, and c_{sym} , c_{anti} are suitable constants. Note that

$$\boldsymbol{\Pi}_{AA'}^{(\text{sym})} = \frac{1}{2} \left(\boldsymbol{I}_{AA'} + \boldsymbol{S}_{AA'} \right),
\boldsymbol{\Pi}_{AA'}^{(\text{anti})} = \frac{1}{2} \left(\boldsymbol{I}_{AA'} - \boldsymbol{S}_{AA'} \right),$$
(10.422)

86

Exercises

where $S_{AA'}$ is the swap operator, and that the symmetric and antisymmetric subspaces have dimension $\frac{1}{2}|A|(|A|+1)$ and dimension $\frac{1}{2}|A|(|A|-1)$ respectively.

Even if you are not familiar with group representation theory, you might regard eq.(10.421) as obvious. We may write $M_{AA'}(U)$ as a sum of two terms, one symmetric and the other antisymmetric under the interchange of A and A'. The expectation of the symmetric part must be symmetric, and the expectation value of the antisymmetric part must be antisymmetric. Furthermore, averaging over the unitary group ensures that no symmetric state is preferred over any other.

b) To evaluate the constant c_{sym} , multiply both sides of eq.(10.421) by $\Pi_{AA'}^{(\text{sym})}$ and take the trace of both sides, thus finding

$$c_{\rm sym} = \frac{|A_1| + |A_2|}{|A| + 1}.$$
 (10.423)

c) To evaluate the constant c_{anti} , multiply both sides of eq.(10.421)) by $\Pi_{AA'}^{(\text{anti})}$ and take the trace of both sides, thus finding

$$c_{\text{anti}} = \frac{|A_1| - |A_2|}{|A| - 1}.$$
 (10.424)

d) Using

$$c_{I} = \frac{1}{2} (c_{\text{sym}} + c_{\text{anti}}), \quad c_{S} = \frac{1}{2} (c_{\text{sym}} - c_{\text{anti}})$$
 (10.425)

prove eq.(10.334).

References

- [1] M. M. Wilde, Quantum Information Theory (Cambridge, 2013).
- [2] T. M. Cover and J. A. Thomas, *Information Theory* (Wiley, 1991).
- [3] C. E Shannon and W. Weaver, The Mathematical Theory of Communication (Illinois, 1949).
- [4] M. A. Nielsen and I. L. Chuang, Quantum Computation and Quantum Information (Cambridge, 2000).
- [5] A. Wehrl, General properties of entropy, Rev. Mod. Phys. 50, 221 (1978).
- [6] E. H. Lieb and M. B. Ruskai, A fundamental property of quantum-mechanical entropy, Phys. Rev. Lett. 30, 434 (1973).
- [7] P. Hayden, R. Jozsa, D. Petz, and A. Winter, Structure of states which satisfy strong subadditivity with equality, Comm. Math. Phys. 246, 359-374 (2003).
- [8] M. A. Nielsen and J. Kempe, Separable states are more disordered globally than locally, Phys. Rev. Lett. 86, 5184 (2001).
- [9] J. Bekenstein, Universal upper bound on the entropy-to-energy ration of bounded systems, Phys. Rev. D 23, 287 (1981).
- [10] H. Casini, Relative entropy and the Bekenstein bound, Class. Quant. Grav. 25, 205021 (2008).
- [11] P. J. Coles, M. Berta, M. Tomamichel, S. Wehner, Entropic uncertainty relations and their applications, arXiv:1511.04857 (2015).
- [12] H. Maassen and J. Uffink, Phys. Rev. Lett. 60, 1103 (1988).
- [13] B. Schumacher, Quantum coding, Phys. Rev. A 51, 2738 (1995).
- [14] R. Jozsa and B. Schumacher, A new proof of the quantum noiseless coding theory, J. Mod. Optics 41, 2343-2349 (1994).
- [15] C. H. Bennett, H. J. Bernstein, S. Popescu, and B. Schumacher, Concentrating partial entanglement by local operations, Phys. Rev. A 53, 2046 (1996).
- [16] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, Quantum entanglement, Rev. Mod. Phys. 81, 865 (2009).
- [17] F. G. S. L. Brandão and M. B. Plenio, A reversible theory of entanglement and its relation to the second law, Comm. Math. Phys. 295, 829-851 (2010).
- [18] M. Christandl and A. Winter, "Squashed entanglement": an additive entanglement measure, J. Math. Phys. 45, 829 (2004).
- [19] M. Koashi and A. Winter, Monogamy of quantum entanglement and other correlations, Phys. Rev. A 69, 022309 (2004).
- [20] F. G. S. L. Brandão, M. Cristandl, and J. Yard, Faithful squashed entanglement, Comm. Math. Phys. 306, 805-830 (2011).
- [21] A. C. Doherty, P. A. Parrilo, and F. M. Spedalieri, Complete family of separability criteria, Phys. Rev. A 69, 022308 (2004).
- [22] A. S. Holevo, Bounds for the quantity of information transmitted by a quantum communication channel, Probl. Peredachi Inf. 9, 3-11 (1973).
- [23] A. Peres and W. K. Wootters, Optimal detection of quantum information, Phys. Rev. Lett 66, 1119 (1991).
- [24] A. S. Holevo, The capacity of the quantum channel with general signal states, arXiv: quantph/9611023.
- [25] B. Schumacher and M. D. Westmoreland, Sending classical information via noisy quantum channels, Phys. Rev. A 56, 131-138 (1997).

References

- [26] M. B. Hastings, Superadditivity of communication capacity using entangled inputs, Nature Physics 5, 255-257 (2009).
- [27] M. Horodecki, P. W. Shor, and M. B. Ruskai, Entanglement breaking channels, Rev. Math. Phys. 15, 629-641 (2003).
- [28] P. W. Shor, Additivity of the classical capacity for entanglement-breaking quantum channels, J. Math. Phys. 43, 4334 (2002).
- [29] B. Schumacher and M. A. Nielsen, Quantum data processing and error correction, Phys. Rev. A 54, 2629 (1996).
- [30] B. Schumacher, Sending entanglement through noisy quantum channels, Phys. Rev. A 54, 2614 (1996).
- [31] H. Barnum, E. Knill, and M. A. Nielsen, On quantum fidelities and channel capacities, IEEE Trans. Inf. Theory 46, 1317-1329 (2000).
- [32] S. Lloyd, Capacity of the noisy quantum channel, Phys. Rev. A 55, 1613 (1997).
- [33] P. W. Shor, unpublished (2002).
- [34] I. Devetak, The private classical capacity and quantum capacity of a quantum channel, IEEE Trans. Inf. Theory 51, 44-55 (2005).
- [35] I. Devetak and A. Winter, Distillation of secret key and entanglement from quantum states, Proc. Roy. Soc. A 461, 207-235 (2005).
- [36] B. Schumacher and M. D. Westmoreland, Approximate quantum error correction, Quant. Inf. Proc. 1, 5-12 (2002).
- [37] M. Horodecki, J. Oppenheim, and A. Winter, Quantum state merging and negative information, Comm. Math. Phys. 269, 107-136 (2007).
- [38] P. Hayden, M. Horodecki, A. Winter, and J. Yard, Open Syst. Inf. Dyn. 15, 7-19 (2008).
- [39] A. Abeyesinge, I. Devetak, P. Hayden, and A. Winter, Proc. Roy. Soc. A, 2537-2563 (2009).
- [40] E. Lubkin, Entropy of an n-system from its correlation with a k-reservoir, J. Math. Phys. 19, 1028 (1978).
- [41] S. Lloyd and H. Pagels, Complexity as thermodynamic depth, Ann. Phys. 188, 186-213 (1988)
- [42] D. N. Page, Average entropy of a subsystem, Phys. Rev. Lett. 71, 1291 (1993).
- [43] I. Devetak, A. W. Harrow, and A. Winter, A family of quantum protocols, Phys. Rev. Lett. 93, 230504 (2004).
- [44] I. Devetak, A. W. Harrow, and A. Winter, A resource framework for quantum Shannon theory, IEEE Trans. Inf. Theory 54, 4587-4618 (2008).
- [45] I. Devetak and P. W. Shor, The capacity of a quantum channel for simultaneous transmission of classical and quantum information, Comm. Math. Phys. 256, 287-303 (2005).
- [46] C. H. Bennett, P. W. Shor, J. A. Smolin, and A. V. Thapliyal, Entanglement-assisted classical capacity of noisy quantum channels, Phys. Rev. Lett. 83, 3081 (1999).
- [47] C. H. Bennett, P. W. Shor, J. A. Smolin, and A. V. Thapliyal, Entanglement-assisted classical capacity of a quantum channel and the reverse Shannon theorem, IEEE Trans. Inf. Theory 48, 2637-2655 (2002).
- [48] P. W. Shor and J. A. Smolin, Quantum error-correcting codes need not completely reveal the error syndrome, arXiv:quant-ph/9604006.
- [49] D. P. DiVincenzo, P. W. Shor, and J. A. Smolin, Quantum channel capacity of very noisy channels, Phys. Rev. A 57, 830 (1998).
- [50] G. Smith and J. Yard, Quantum communication with zero-capacity channels, Science 321, 1812-1815 (2008).
- [51] L. del Rio, J. Aberg, R. Renner, O. Dahlsten, and V. Vedral, The thermodynamic meaning of negative entropy, Nature 474, 61-63 (2011).
- [52] P. Hayden and J. Preskill, Black holes as mirrors: quantum information in random subsystems, JHEP 09, 120 (2007).
- [53] Y. Sekino and L. Susskind, Fast scramblers, JHEP 10, 065 (2008).