IEEE MILCOM

IEEE ComSoc
IEEE Communications Society

Biased Backpressure Routing for Multihop Wireless Networks with Heterogeneous Interfaces

Yujun Ming*, Zhongyuan Zhao*, Fikadu Dagefu‡, Justin Kong‡, Terrence Moore‡, Kevin Chan‡, Ananthram Swami‡, Santiago Segarra*

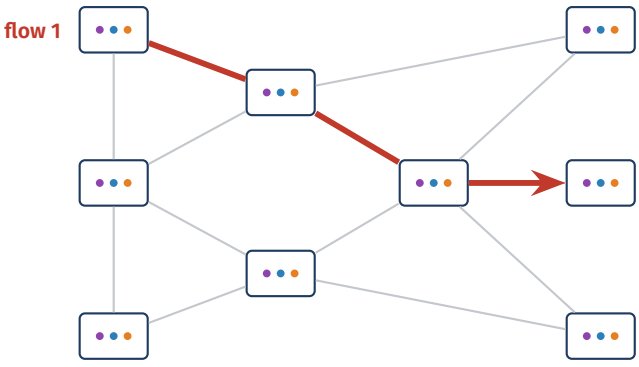
*Rice University, ‡US Army's DEVCOM Army Research Laboratory





12–16 October 2026 // National Capital Region, USA // IEEE Military Communications Conference

How can we route packets when every device has multiple wireless interfaces to choose from?

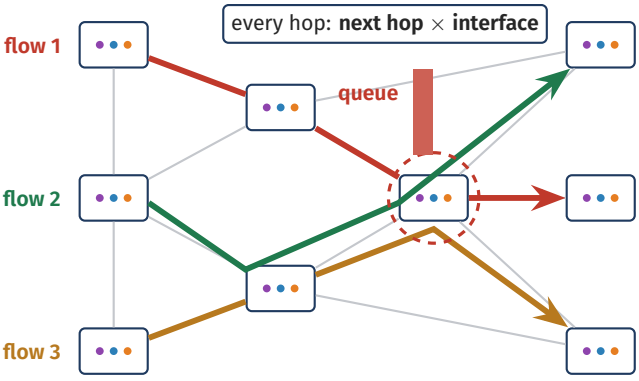


Each device carries **5 GHz**, **2.4 GHz** and **UHF**.

Standard approach:
find a good path for each flow

shortest-path / on-demand routing

How can we route packets when every device has multiple wireless interfaces to choose from?



Each device carries **5 GHz**, **2.4 GHz** and **UHF**.

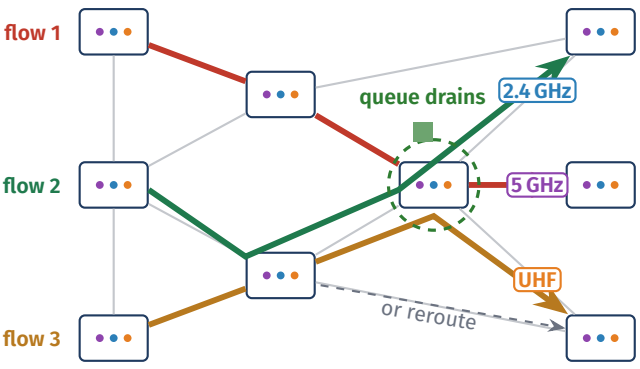
Standard approach:
find a good path for each flow

shortest-path / on-demand routing

Independent per-flow paths
→ shared bottleneck → congestion

With multiple interfaces:
– harder: next hop *and* interface

How can we route packets when every device has multiple wireless interfaces to choose from?



Each device carries **5 GHz**, **2.4 GHz** and **UHF**.

Standard approach:
find a good path for each flow

shortest-path / on-demand routing

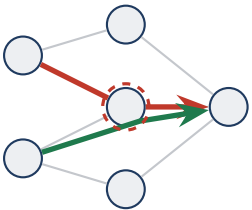
Independent per-flow paths
→ shared bottleneck → congestion

With multiple interfaces:
− harder: next hop *and* interface
+ new way out: not only **reroute**, also **switch interface**

More interfaces: more decisions, but also new ways out of congestion.

What if we optimize multiple flows together?

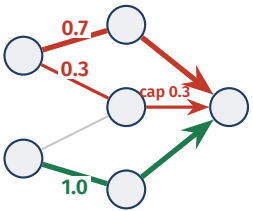
Independent path selection



flows overlap at a bottleneck

Minimum-Cost Flow (MCF) Baseline

offline, multi-commodity: $\min_f \sum_{e,c} w_e f_e^{(c)}$ s.t. link capacities



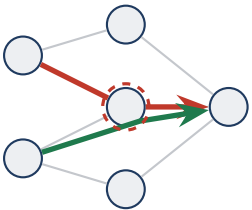
demand: 1 per flow
middle route: cheapest, but link capacity 0.3
$\Rightarrow$ optimal split **0.7 / 0.3**, **1.0**

✓ coordinates multiple flows ✓ uses link costs and capacities

flow split $\rightarrow$ forwarding probabilities (**MCF-Prob** ➕ in our results)

What if we optimize multiple flows together?

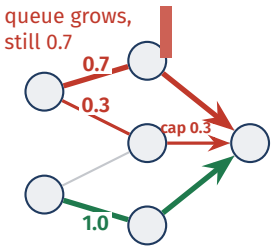
Independent path selection



flows overlap at a bottleneck

Minimum-Cost Flow (MCF) Baseline

offline, multi-commodity: $\min_f \sum_{e,c} w_e f_e^{(c)}$ s.t. link capacities



demand: 1 per flow
middle route: cheapest, but link capacity 0.3
$\Rightarrow$ optimal split **0.7 / 0.3**, **1.0**

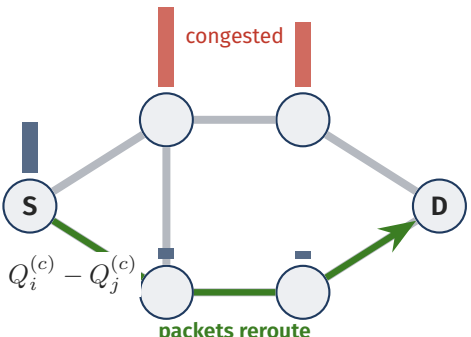
✓ coordinates multiple flows ✓ uses link costs and capacities

flow split $\rightarrow$ forwarding probabilities (**MCF-Prob** ➕ in our results)

But... **static probabilistic forwarding, no real-time queue state**

Can routing respond to congestion as it happens? $\rightarrow$ **Backpressure**

Routing with real-time congestion

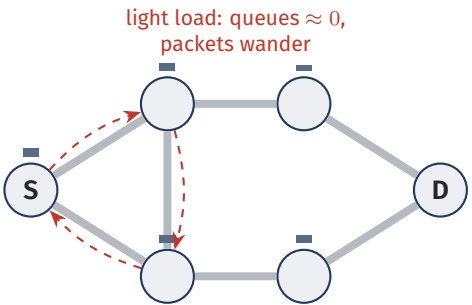


Backpressure

like water in pipes: packets flow from fuller queues to emptier ones; if one route jams, traffic flows around it on its own

✓ real-time queues
✓ distributed
✓ throughput-optimal / queue-stable

Routing with real-time congestion

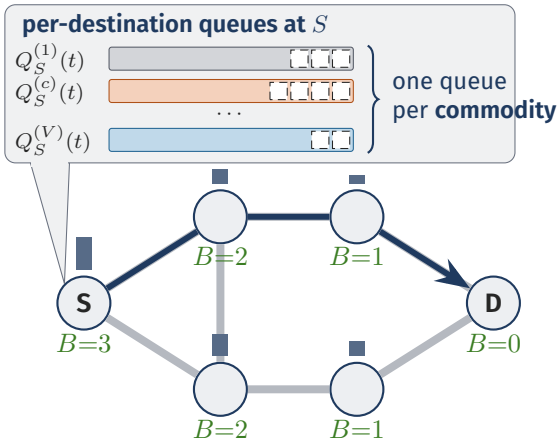


Backpressure

But: poor delay under light traffic

with small queues, pressure gives little direction

Routing with real-time congestion



Shortest-Path Biased Backpressure

$$P_{ij}^{(c)} = \underbrace{Q_i^{(c)} - Q_j^{(c)}}_{\text{queue pressure}} + \underbrace{B_i^{(c)} - B_j^{(c)}}_{\text{shortest-path bias}}$$

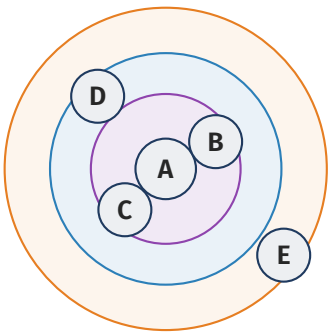
light traffic $\rightarrow$ shortest path
congestion $\rightarrow$ pressure reroutes

SP-BP handles multi-flow congestion. But there is still a second challenge.

But which interface should carry the packet?

wireless interface $\longleftrightarrow$ **modality**

physical device | our model



Modality 1 (mmWave): $A \to B, C$

high capacity, short range, blockage-sensitive

Modality 2 (5G): $A \to B, C, D$

medium range, medium capacity

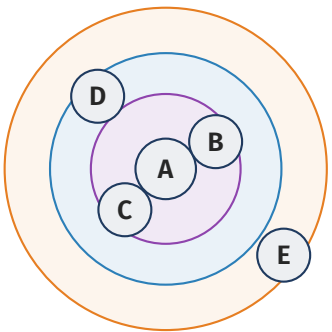
Modality 3 (Sub-GHz): $A \to B, C, D, E$

long reach; slow, low capacity

But which interface should carry the packet?

wireless interface ⟷ **modality**

physical device — our model



Modality 1 (mmWave): $A \to B, C$

high capacity, short range, blockage-sensitive

Modality 2 (5G): $A \to B, C, D$

medium range, medium capacity

Modality 3 (Sub-GHz): $A \to B, C, D, E$

long reach; slow, low capacity

Which modality? ⟷ Which next hop?

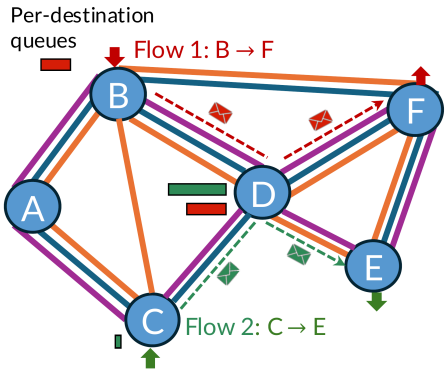
modality choice changes neighbors and rates

Can we turn these two different decisions into one?

Turn interface selection into routing

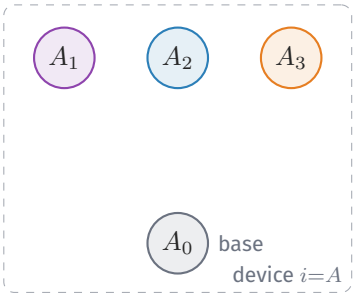


one physical device = one graph node



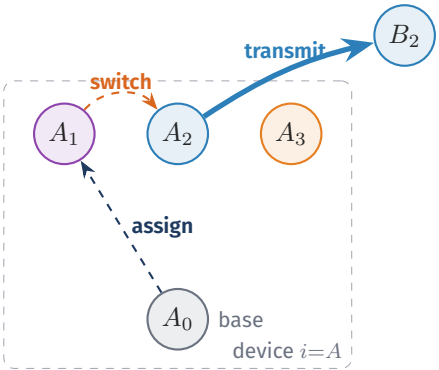
A multihop wireless network: three interfaces per device (link colors), per-destination queues.

Turn interface selection into routing



Base node i_0: where new packets enter
Modality nodes i_m: one per modality m

Turn interface selection into routing

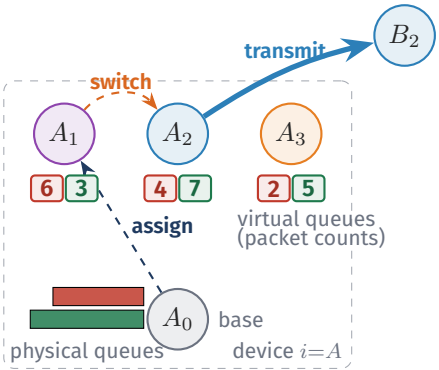


Base node i_0: where new packets enter
Modality nodes i_m: one per modality m

Virtual link $e^v = (i_n, i_m)$: inside device i
assign (i_0, i_m) **switch** (i_n, i_m)

Physical link $e^p = (i_m, j_m)$:
transmit to device j on modality m

Turn interface selection into routing



Base node i_0: where new packets enter
Modality nodes i_m: one per modality m

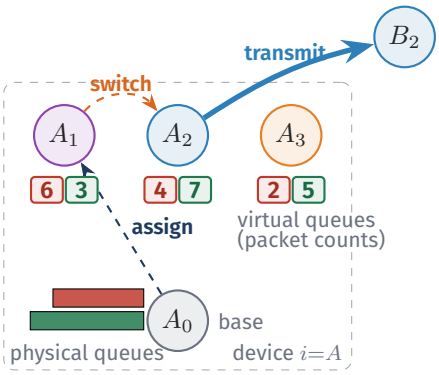
Virtual link $e^v = (i_n, i_m)$: inside device i
assign (i_0, i_m) **switch** (i_n, i_m)

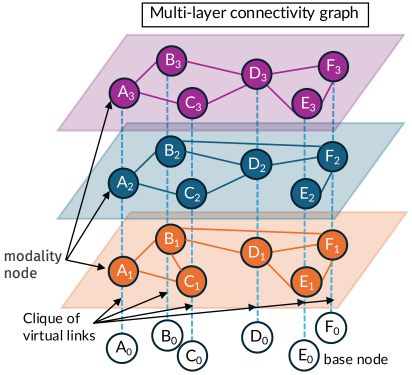
Physical link $e^p = (i_m, j_m)$:
transmit to device j on modality m

Physical queues at i_0: packets device i stores, one per flow

6 3 **Virtual queues** at i_m: packets assigned to modality m; they add up to the physical queue

Turn interface selection into routing





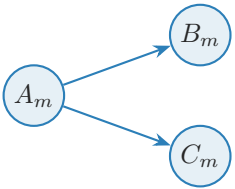
One layer per modality.

modality selection + next-hop selection $\rightarrow$ **one unified routing problem**

What does this representation give us?

Physical routing

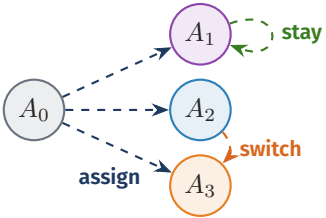
which next hop?



find a **physical** next hop within a layer

Modality selection

which modality?

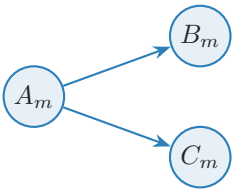


find a **virtual** next hop within the device

What does this representation give us?

Physical routing

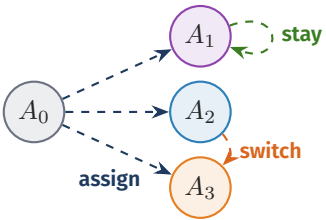
which next hop?



find a **physical** next hop within a layer

Modality selection

which modality?



find a **virtual** next hop within the device

Both become **next-hop selection** on the multi-layer graph

So how do we actually make these routing decisions?

How does SP-BP make a routing decision?

Slot-by-slot MaxWeight decision solved every slot t

$$\mu^\star(t) = \arg \max_{\mu \in \Pi_p(t)} \sum_{(i,j),\,c} \mu_{ij}^{(c)}\, P_{ij}^{(c)}$$

$\Pi_p(t)$: feasible physical actions (link rates, conflicts)

Decision $\mu_{ij}^{(c)}$

allocated transmission rate on link (i,j) for commodity c

Pressure $P_{ij}^{(c)}$ (biased)

$$P_{ij}^{(c)} = \underbrace{Q_i^{(c)} - Q_j^{(c)}}_{\text{queue differential}} + \underbrace{B_i^{(c)} - B_j^{(c)}}_{\text{shortest-path bias}}$$

Implemented by the **classical 4-step SP-BP** procedure (fully distributed)

Select optimal commodity → Link weighting → Conflict-aware MaxWeight scheduling → Assign link capacity

One routing objective

$$W_p = \sum_{e \in \mathcal{E}_p,\, c} \mu_e^{(c)} P_e^{(c)}$$

physical routing

$$W_v = \sum_{e \in \mathcal{E}_v,\, c} \nu_e^{(c)} P_e^{(c)}$$

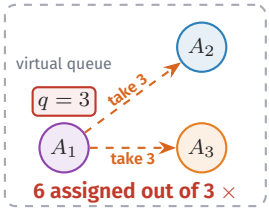
virtual routing / modality selection

$$(\mu^\star, \nu^\star) = \arg \max_{(\mu,\nu) \in \Pi(t)} \left(W_p + W_v\right)$$

physical and virtual routing under **one MaxWeight formulation**

The objective is the same. But the constraints are not.

Same objective. Different constraints.



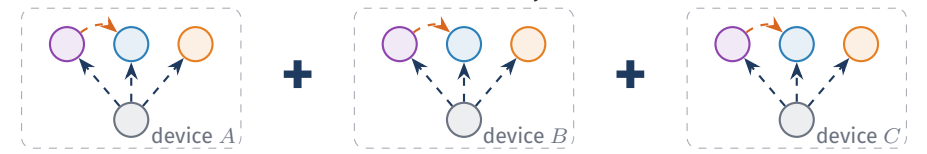
Same goal: maximize pressure

	Virtual routing	**Physical routing**
	$\arg\max_{\nu \in \Pi_v} W_v$	$\arg\max_{\mu \in \Pi_p} W_p$
where	inside a device	across devices
interference	none	yes: conflict graph
constraints	packet availability	link rates
solver	**?**	classical 4-step SP-BP

So how do we solve the virtual argmax?

How do we solve the virtual argmax?

$$\max_{\nu \in \Pi_v} W_v \;=\; \sum_{\text{device } i} \; \max_{\nu_i \in \Pi_v^i} \; \underbrace{\sum_{(i_n, i_m), c} \nu_{i_n i_m}^{(c)} P_{i_n i_m}^{(c)}}_{\text{only device } i\text{'s virtual links}}$$

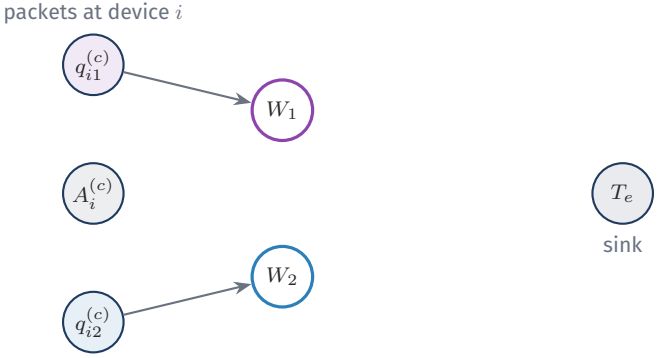


virtual links and packet-assignment constraints are **local to each device**

Each device solves its own minimum-cost flow, in parallel.

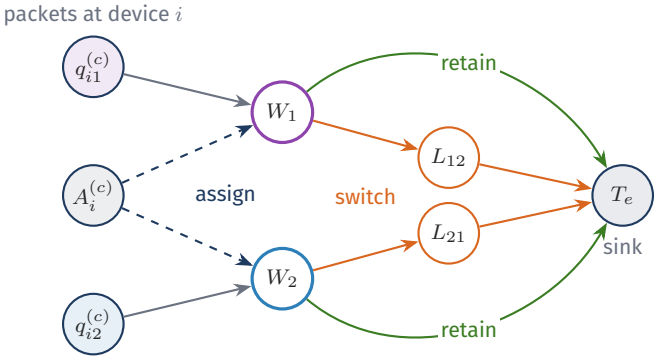
unlike the offline MCF baseline: solved **every slot**, with **current pressures** as costs ($-P$)

Per-Device Computational Graph for Minimum-Cost Flow



supply = available packets

Per-Device Computational Graph for Minimum-Cost Flow



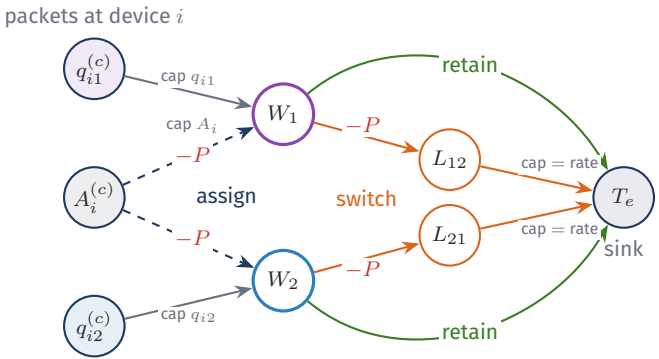
One commodity, two modalities

supply = available packets

retain: stay on the current modality

assign / **switch**: move to another modality

Per-Device Computational Graph for Minimum-Cost Flow



supply = available packets

retain: stay on the current modality

assign / **switch**: move to another modality

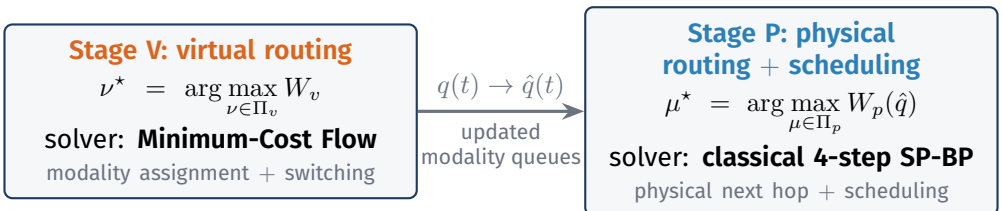
capacities = queued packets and virtual switching rate

costs = negative pressures

One commodity, two modalities

MCF solution = maximizer of the virtual pressure objective

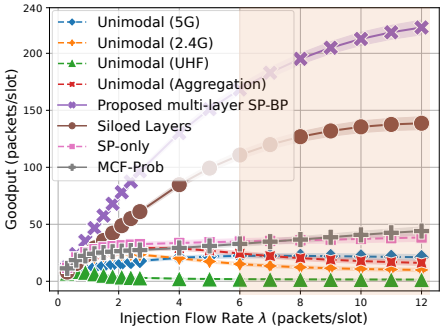
Putting everything together



Same pressure-based objective · different constraints · different solvers

So does this extra flexibility actually improve routing performance?

Dynamic modality switching matters most under congestion



Goodput vs. per-flow arrival rate λ; 100 nodes, streaming.



Aggregation one fat pipe



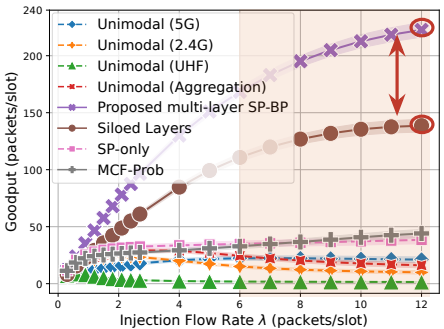
Siloed layers choose a modality at injection $\rightarrow$ **stay on that layer**



Proposed **switch if needed** $\rightarrow$ route

Higher load $\rightarrow$ widening gap

Dynamic modality switching matters most under congestion



Goodput vs. per-flow arrival rate λ; 100 nodes, streaming.



Aggregation
one fat pipe



Siloed layers
choose a modality at injection
$\rightarrow$ **stay on that layer**

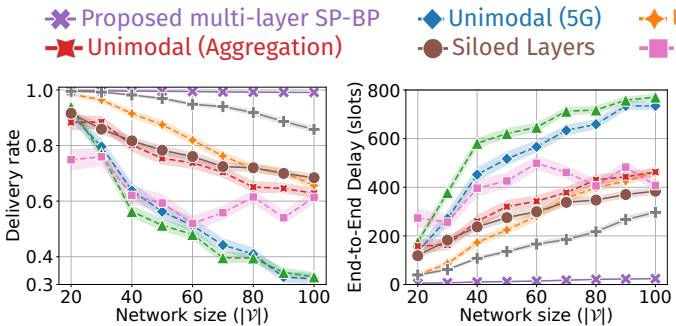


Proposed
switch if needed $\rightarrow$ route

same modalities + same SP-BP
key difference: dynamic cross-modality switching

The value of dynamic modality switching grows with congestion.

Multi-layer SP-BP remains robust across network sizes and traffic types



delivery ratio, higher is better

end-to-end delay, lower is better

Streaming flows

>**99% delivery** across tested sizes

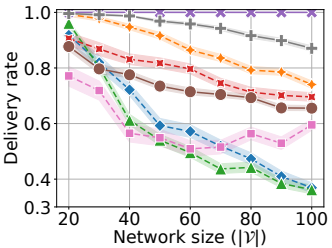
lowest end-to-end delay

stays best as the network grows from 20 to 100 nodes

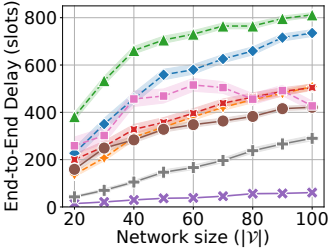
Multi-layer SP-BP remains robust across network sizes and traffic types

Dynamic multi-layer routing adapts across both network sizes and traffic patterns.

Proposed multi-layer SP-BP
Unimodal (5G)
Unimodal (2.4G)
Unimodal (UHF)
Unimodal (Aggregation)
Siloed Layers
SP-only
MCF-Prob



delivery ratio, higher is better



end-to-end delay, lower is better

Bursty flows

Hard for classical BP:
- direction only from queue build-up
- leftover packets drain slowly

Same ranking:
highest delivery, lowest delay at every network size

Takeaways



1. Interface selection becomes routing
our multi-layer model: one unified routing problem



2. Packets choose where to go *and* which modality
dynamically, as congestion changes



3. Interfaces inside the routing framework
not a separate control problem

Thank you!