ECE 598HH: Special Topics in Wireless Networks and Mobile Systems

Lecture 21: Opportunistic Routing

Haitham Hassanieh

*These slides are courtesy of Dina Katabi

Lecture Outline

- Single Path Routing
- Opportunistic Routing with ExOR
- Intra-Flow Network Coding with MORE

Traditional Single Path Routing

Represent the wireless network as a graph

- Two nodes have an edge if they can communicate (i.e., are within radio range)
- Each edge is labeled with a weight (where a smaller weight indicates a preferred edge)

Run shortest path algorithm on the graph (e.g., Dijkstra)

- Produce the minimum weight path between every pair of nodes

How do you pick the edge weights?

- i.e., what metric should shortest path minimize?

A straw-man route metric (1):

Assign all edges the same weight → Minimize number of hops

Reasoning:

- Links in route share radio spectrum
- Extra hops reduce throughput

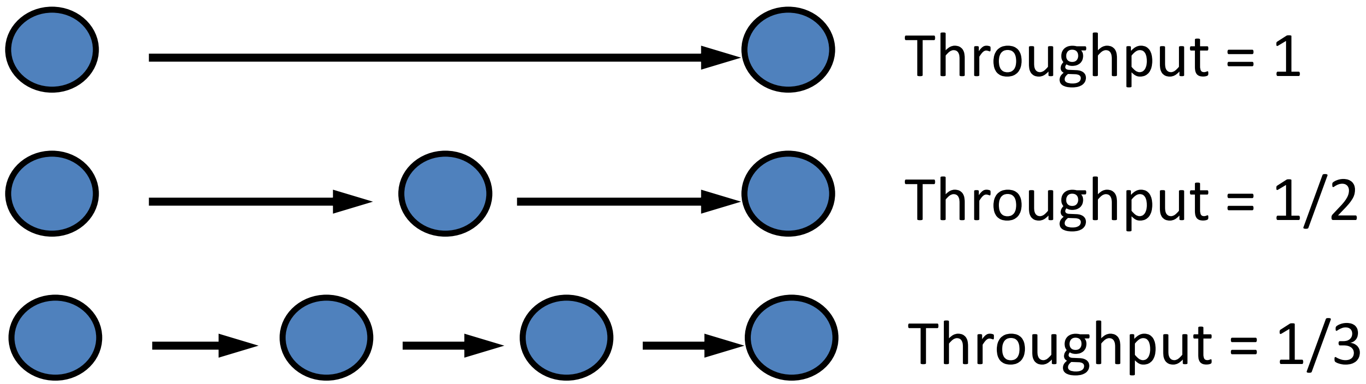

But is not good enough because different edges many have very different packet loss rates

Challenge: links are lossy and asymmetric

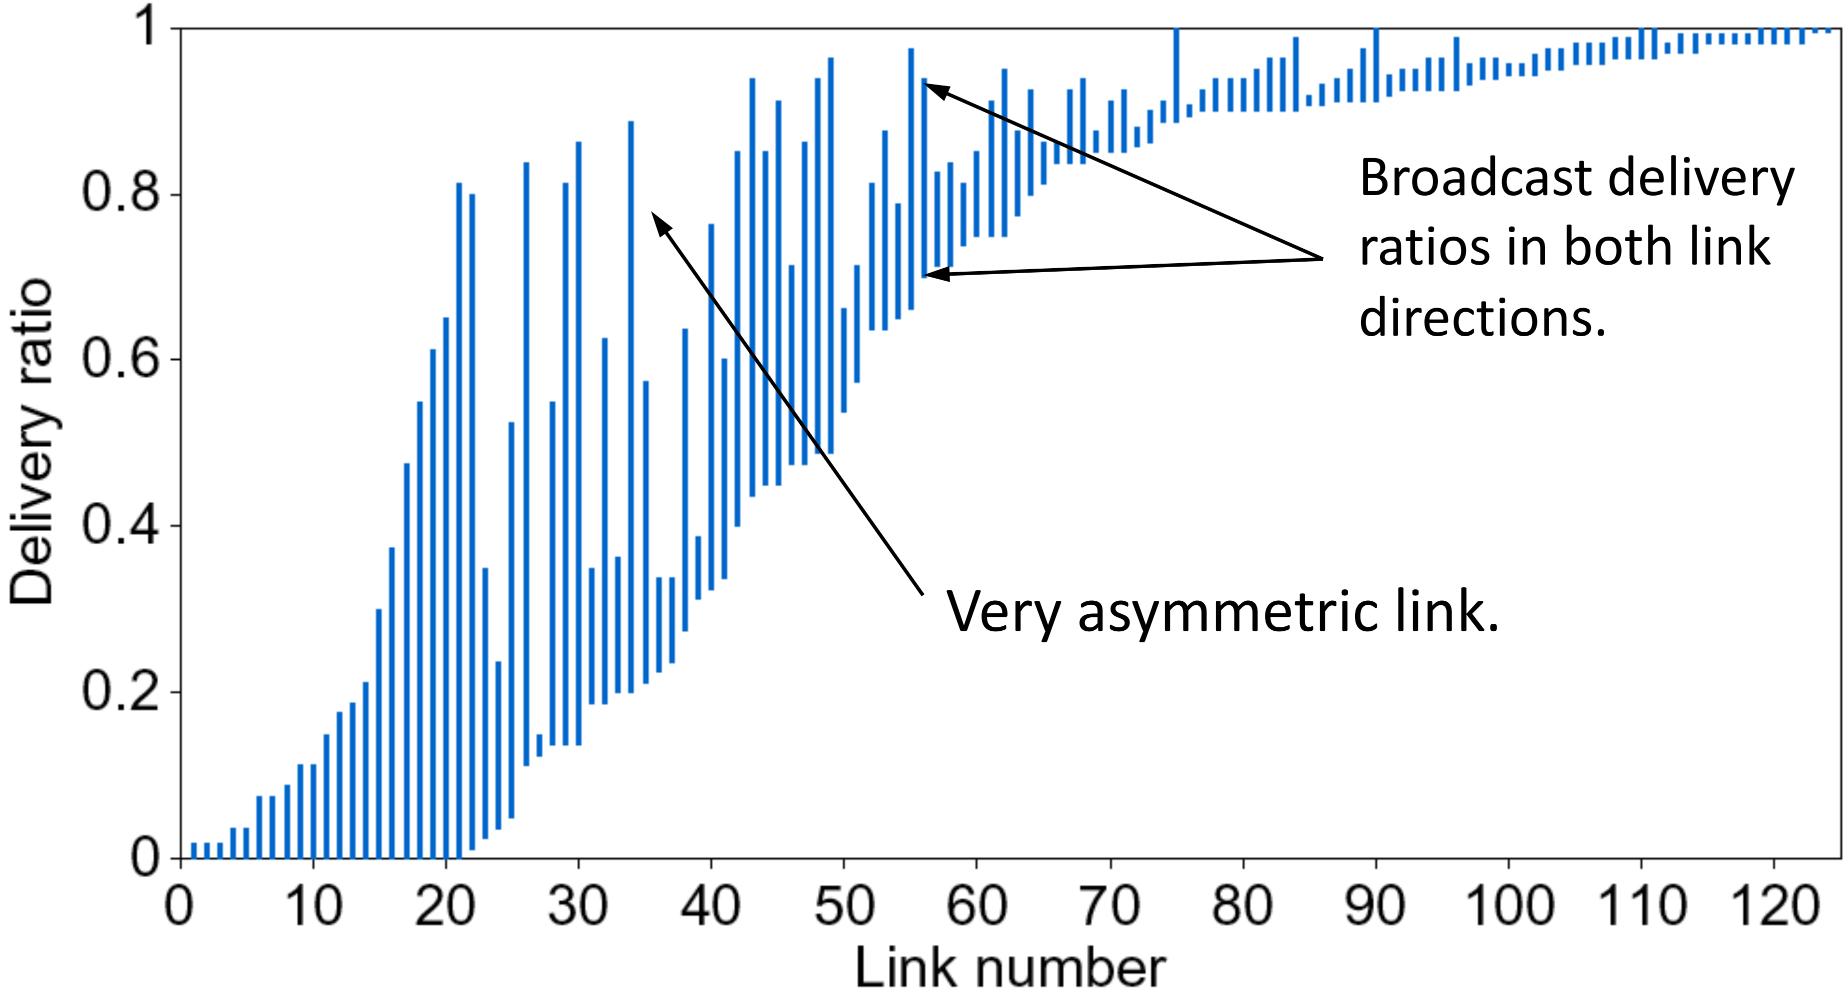

Different links have different loss rates

Further, the loss rate may be different in each direction

A straw-man route metric (2):

Maximize bottleneck throughput

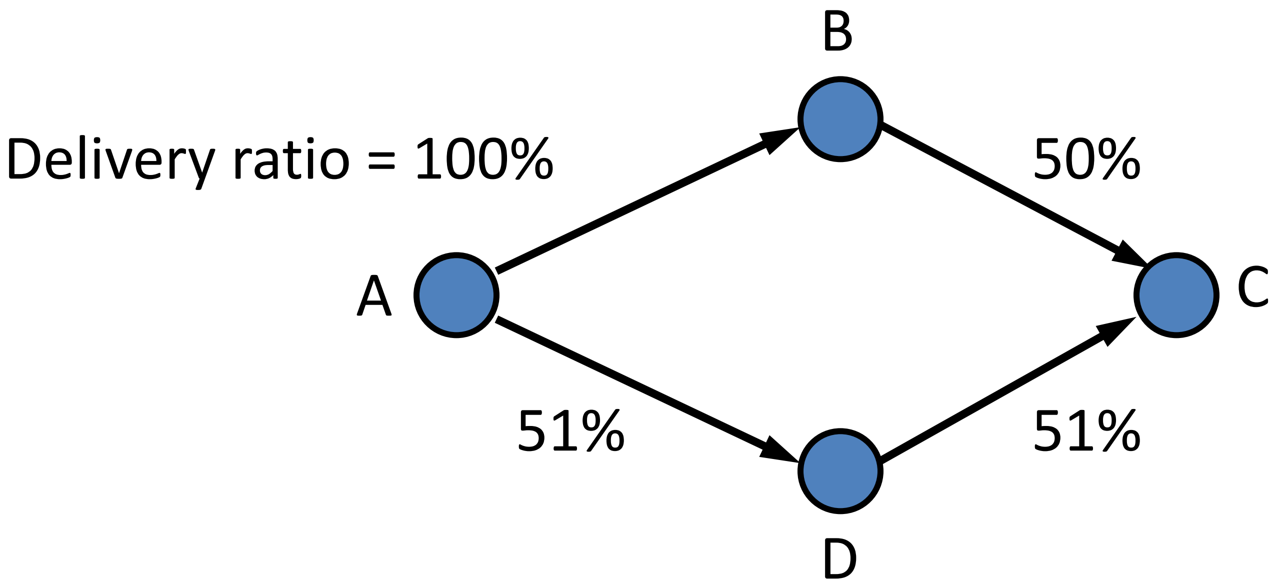

Bottleneck throughput:
- A-B-C = 50%
- A-D-C = <u>51%</u>

Actual throughput:
- A-B-C : ABBABBABB = <u>33%</u>
- A-D-C : AADDAADD = 25%

Key Idea: In a shared medium links are not independent

A straw-man metric (3):
Maximize end-to-end delivery ratio

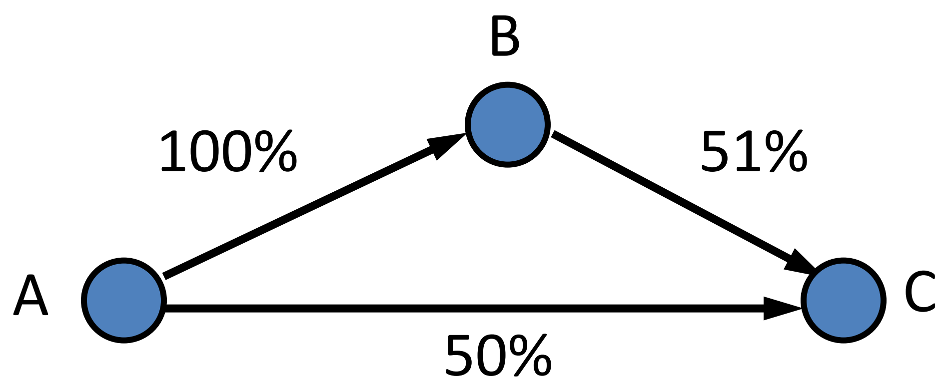

End-to-end delivery ratio:

- A-B-C = 51%
- A-C = 50%

Actual throughput:

- A-B-C : ABBABBABB = 33%
- A-C : AAAAAAAA = 50%

Key Idea: Again, links are not independent

Wireless routing metric: ETX

Minimize total transmissions per packet
(ETX, ‘Expected Transmission Count’)

Link throughput ≈ 1/ Link ETX

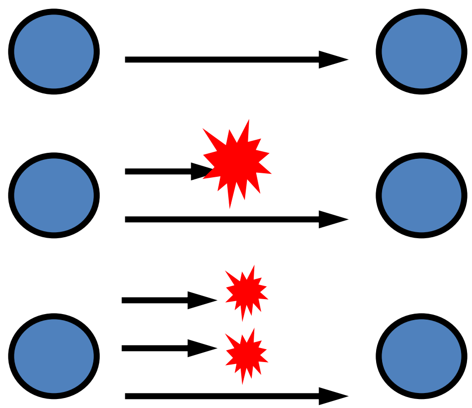

Delivery Ratio	Link ETX	Throughput
100%	1	100%
50%	2	50%
33%	3	33%

Route ETX

Route ETX = Sum of link ETXs

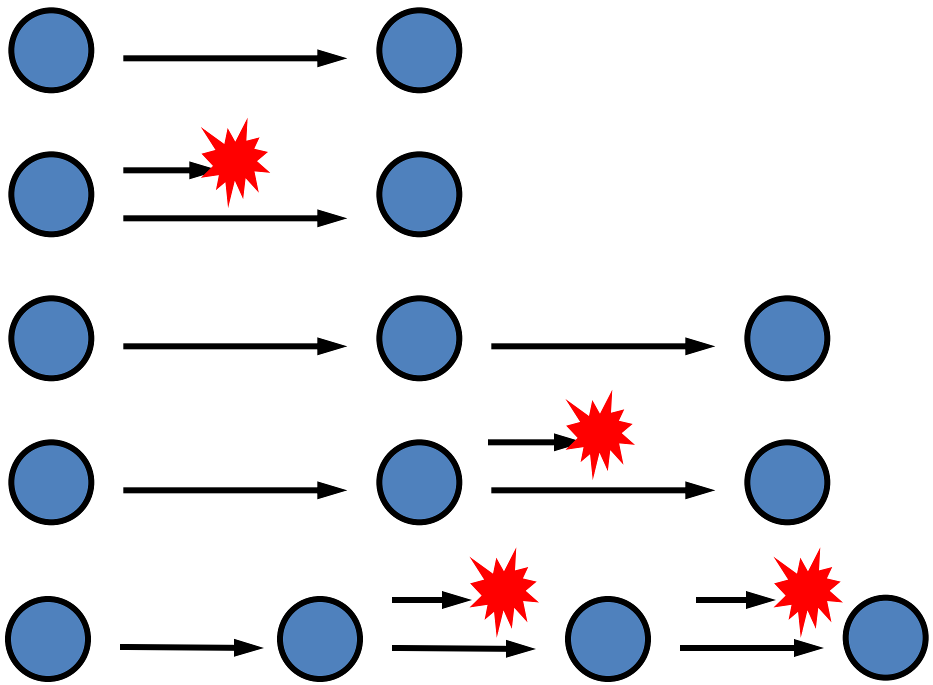

Route ETX	Throughput
1	100%
2	50%
2	50%
3	33%
5	20%

ETX Properties

- ETX predicts throughput for short routes (1, 2, and 3 hops)
- ETX quantifies loss
- ETX quantifies asymmetry

ETX Caveats

- It is really hard to measure link quality/loss
  - Changes as a function of load
  - Changes with time

- ETX ignores differences in bit-rate and packet size

  ETT = ETX *(pkt_size/link-bit-rate)

- ETX ignores spatial re-use (i.e., assumes all links interfere)

Lecture Outline

- Single Path Routing
- ⇨ Opportunistic Routing with ExOR
- Intra-Flow Network Coding with MORE

Context: Roofnet

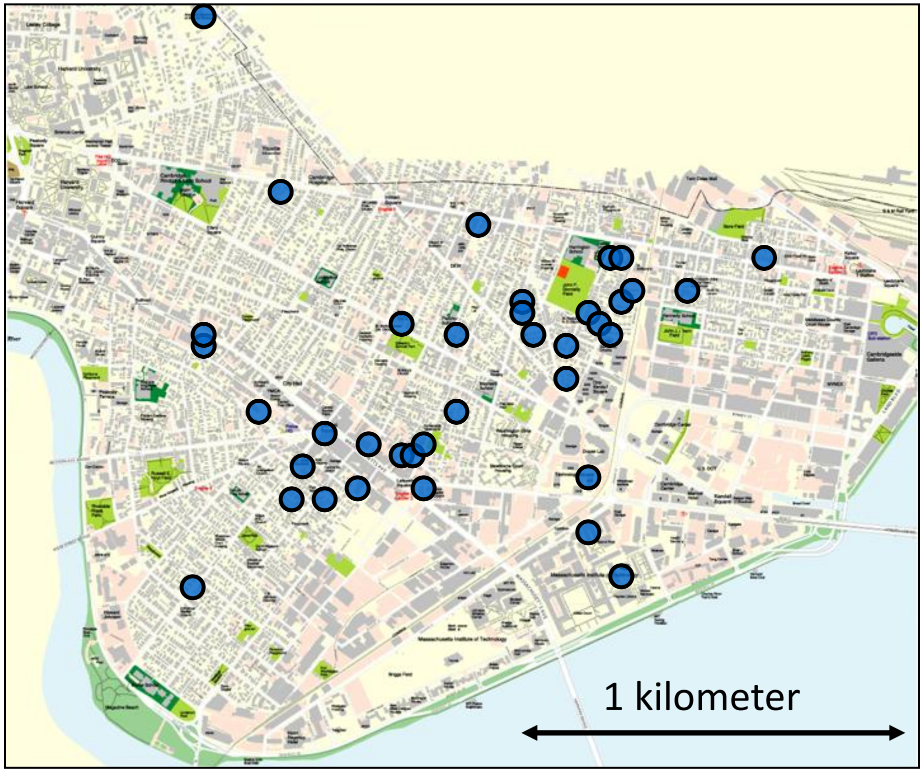

- Dense 802.11-based multi-hop network
- Goal is high-throughput in the presence of lossy links

Traditional routing

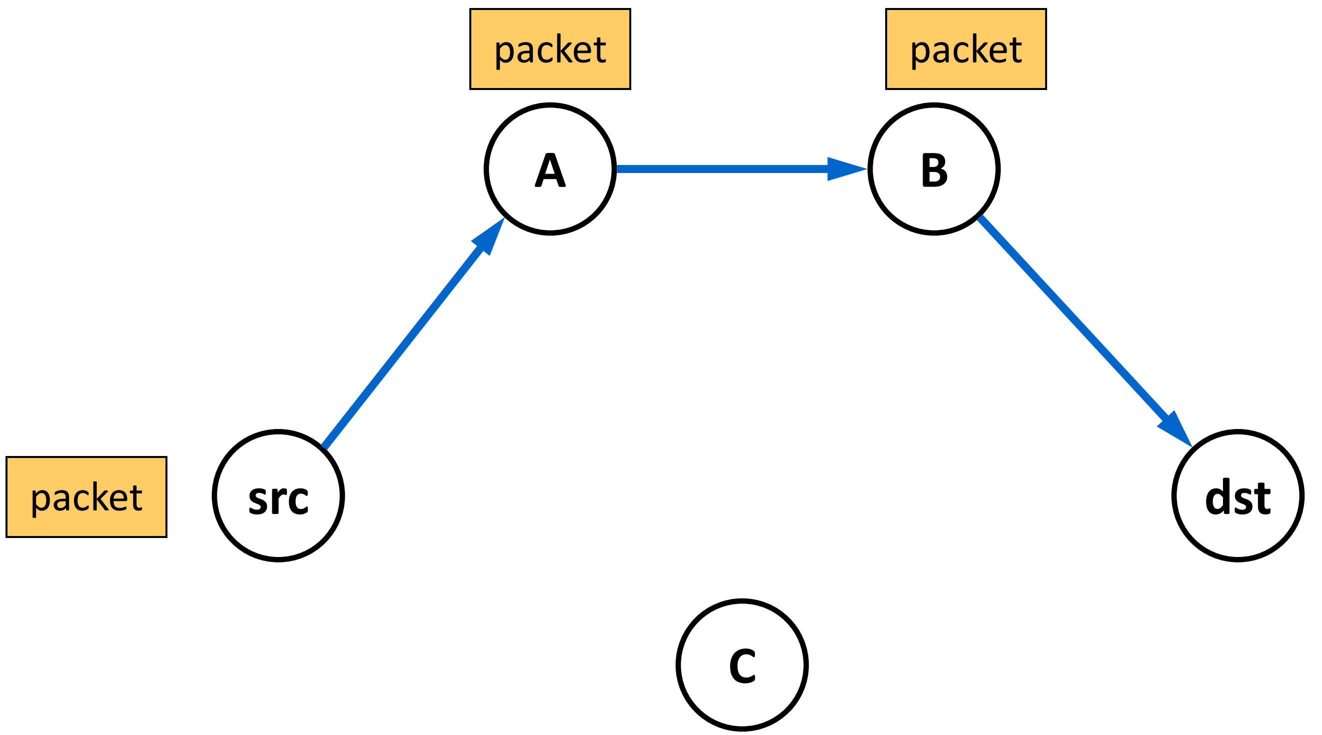

- Identify a route, forward over links
- Abstract radio to look like a wired link

Radios aren't wires

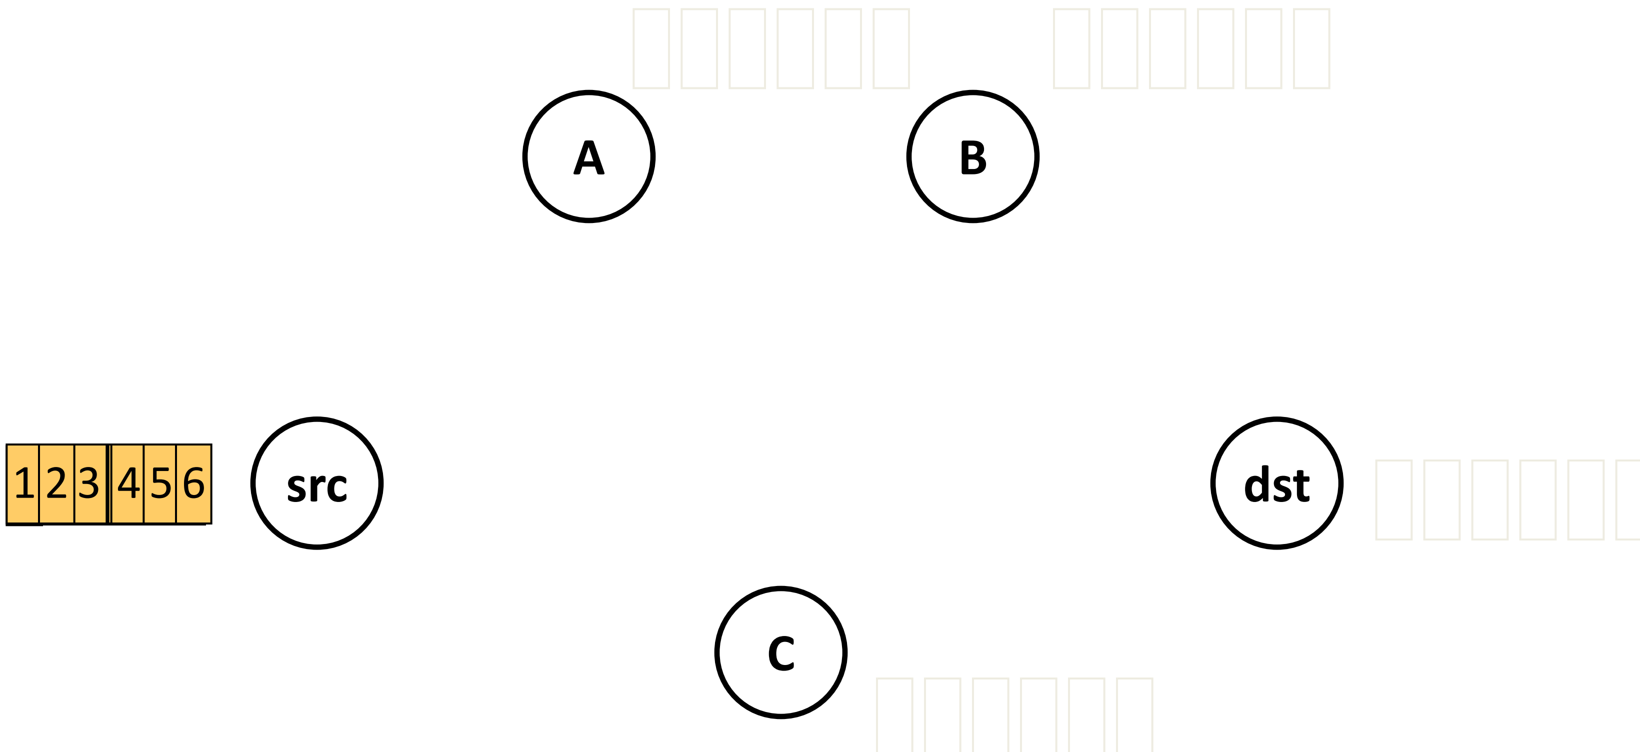

- Every packet is broadcast
- Receptions are probabilistic and independent (Spatial diversity)

ExOR Idea: exploit probabilistic broadcast

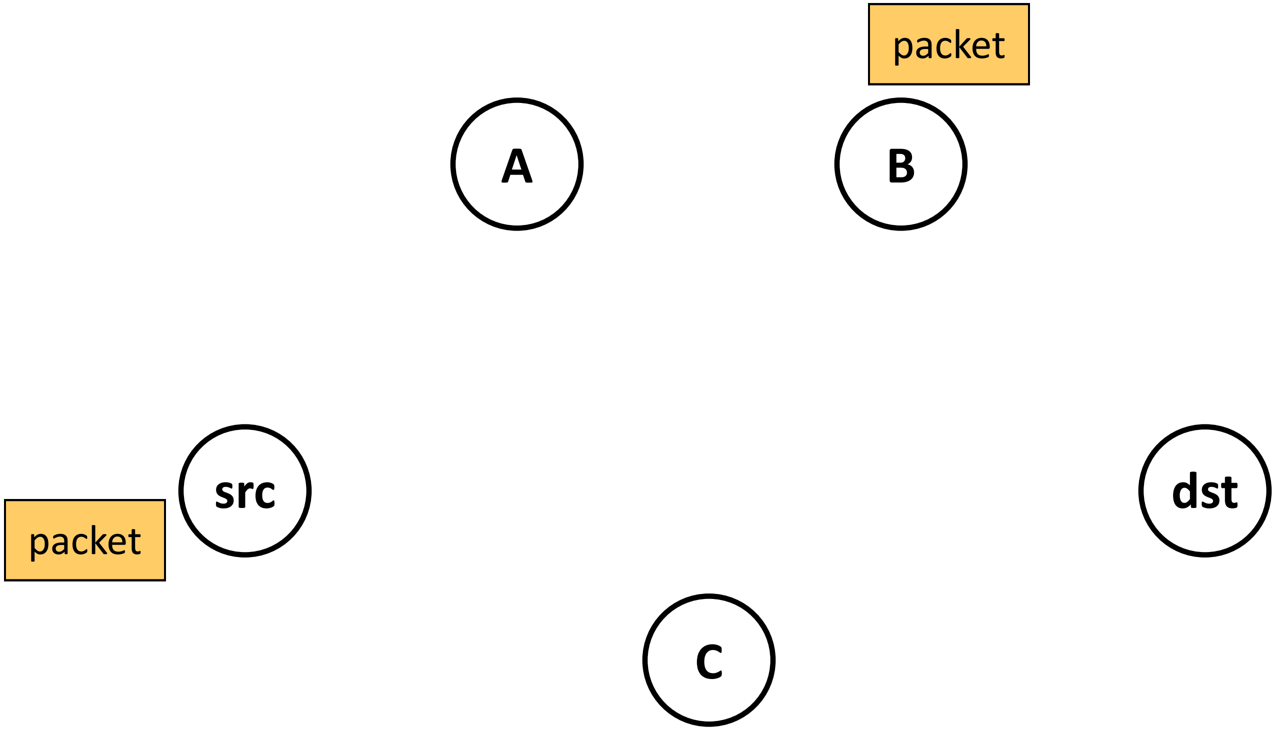

- Decide who forwards <u>after</u> reception
- Goal: for each packet, receiver closest to the destination should forward
- Challenge: agree efficiently on which node should forward, and avoid duplicate transmissions

Why ExOR might increase throughput (1)

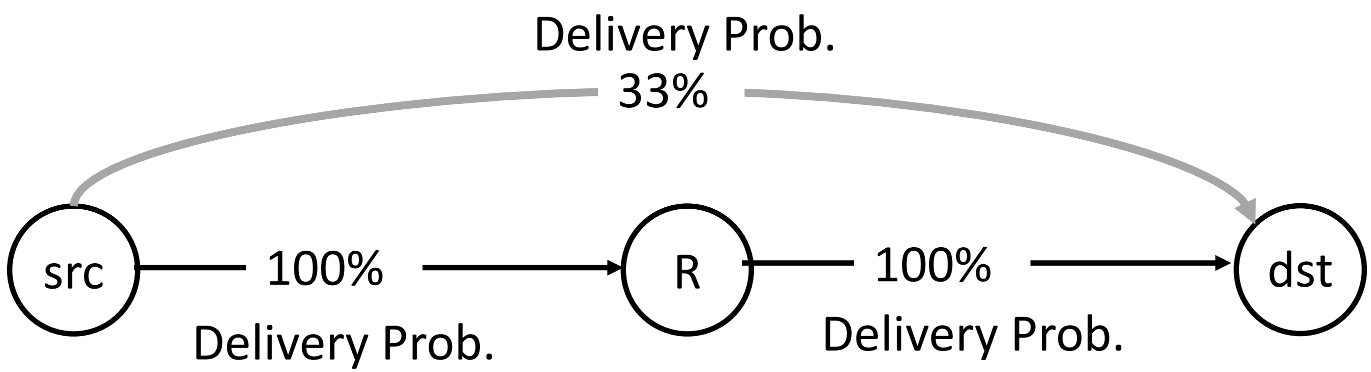

- Traditional routing picks the path via R → on average 2 tx per packet
- Throughput $\cong {}^{1}/_{\#\ \text{transmissions}}$
- Traditional routing ignores that 33% of the packets make it to the destination in one transmission
- ExOR exploits these opportunistic receptions → 1.67 tx per packet

Why ExOR might increase throughput (2)

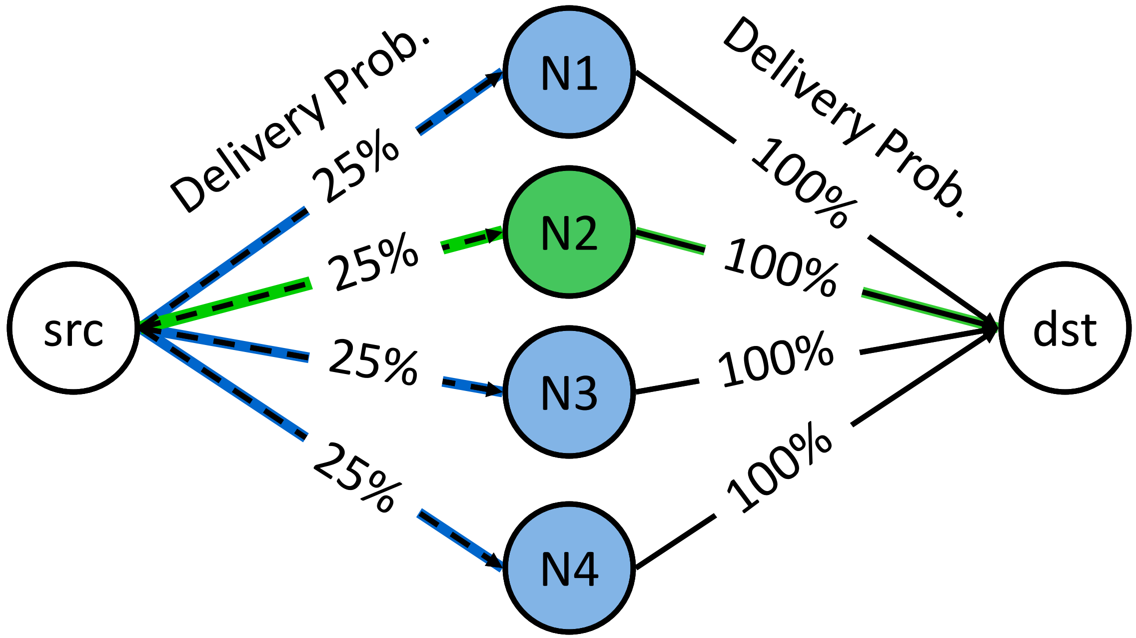

- Traditional routing: $^1/_{0.25} + 1 = 5$ tx
- ExOR: $^1/_{(1-(1-0.25)^4)} + 1 = 2.5$ transmissions

The ExOR Protocol

- Sends packets in batches
- Order the nodes according to their distance to destination
- Go in rounds until all packets in batch are delivered
- In each round,
  - Starting from the node closest to the destination, each node forwards the received packets that no node closer to the destination has received
  - Nodes learn what other nodes have received using ack summaries

ExOR batching

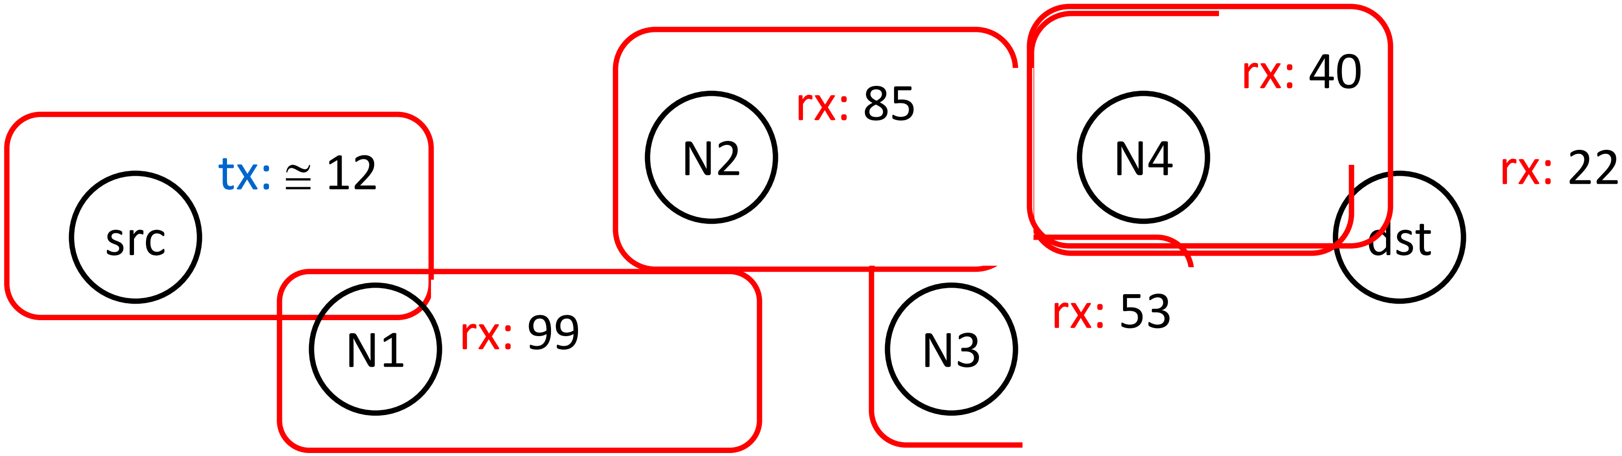

- Challenge: finding the closest node to have rx'd
- Send batches of packets for efficiency
- Node closest to the destination sends first
  - Other nodes listen, send remaining packets in turn
- Repeat schedule until destination has whole batch

Reliable summaries

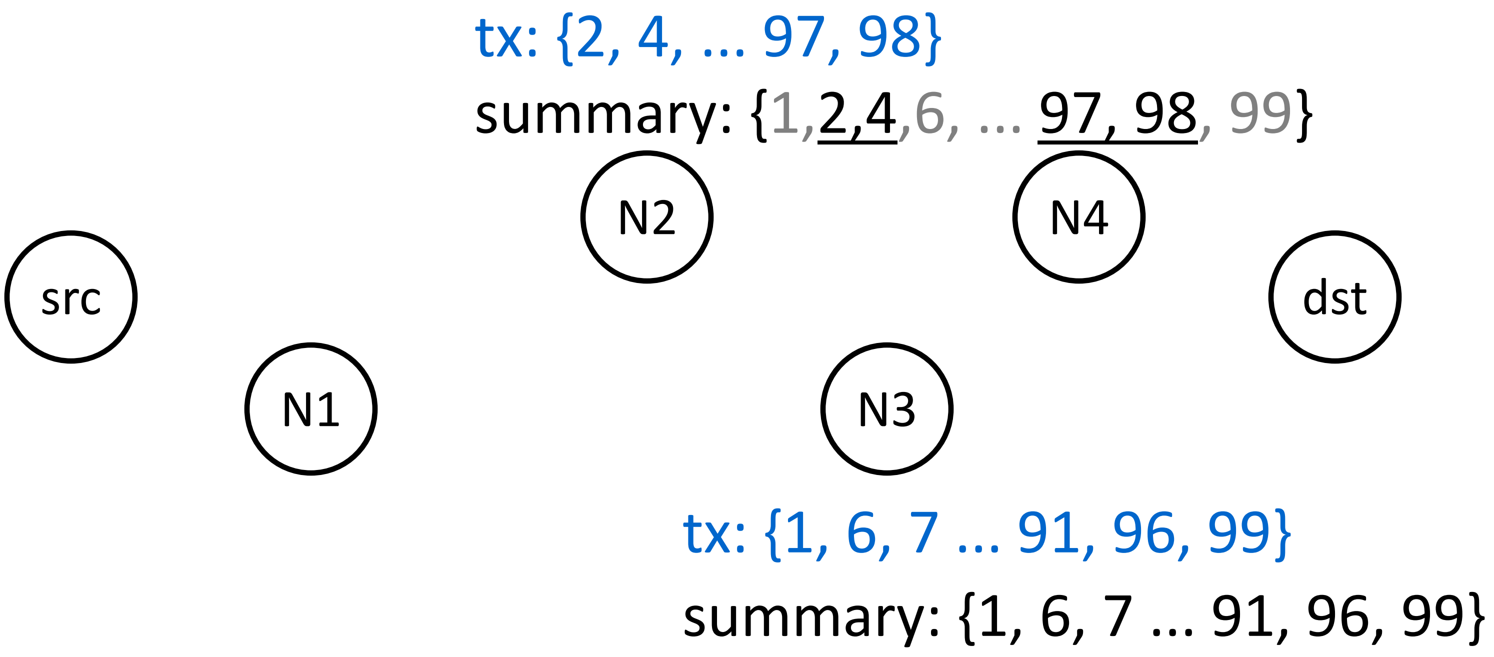

- Repeat summaries in every data packet (as meta data)
- Summaries are cumulative, i.e., a summary reports what all downstream nodes rx'd
- Repetition and accumulation ensure that upstream nodes learn what downstream nodes received and do not forward those packets

ExOR Evaluation

- Does ExOR increase throughput?
- When/why does it work well?

65 Roofnet node pairs

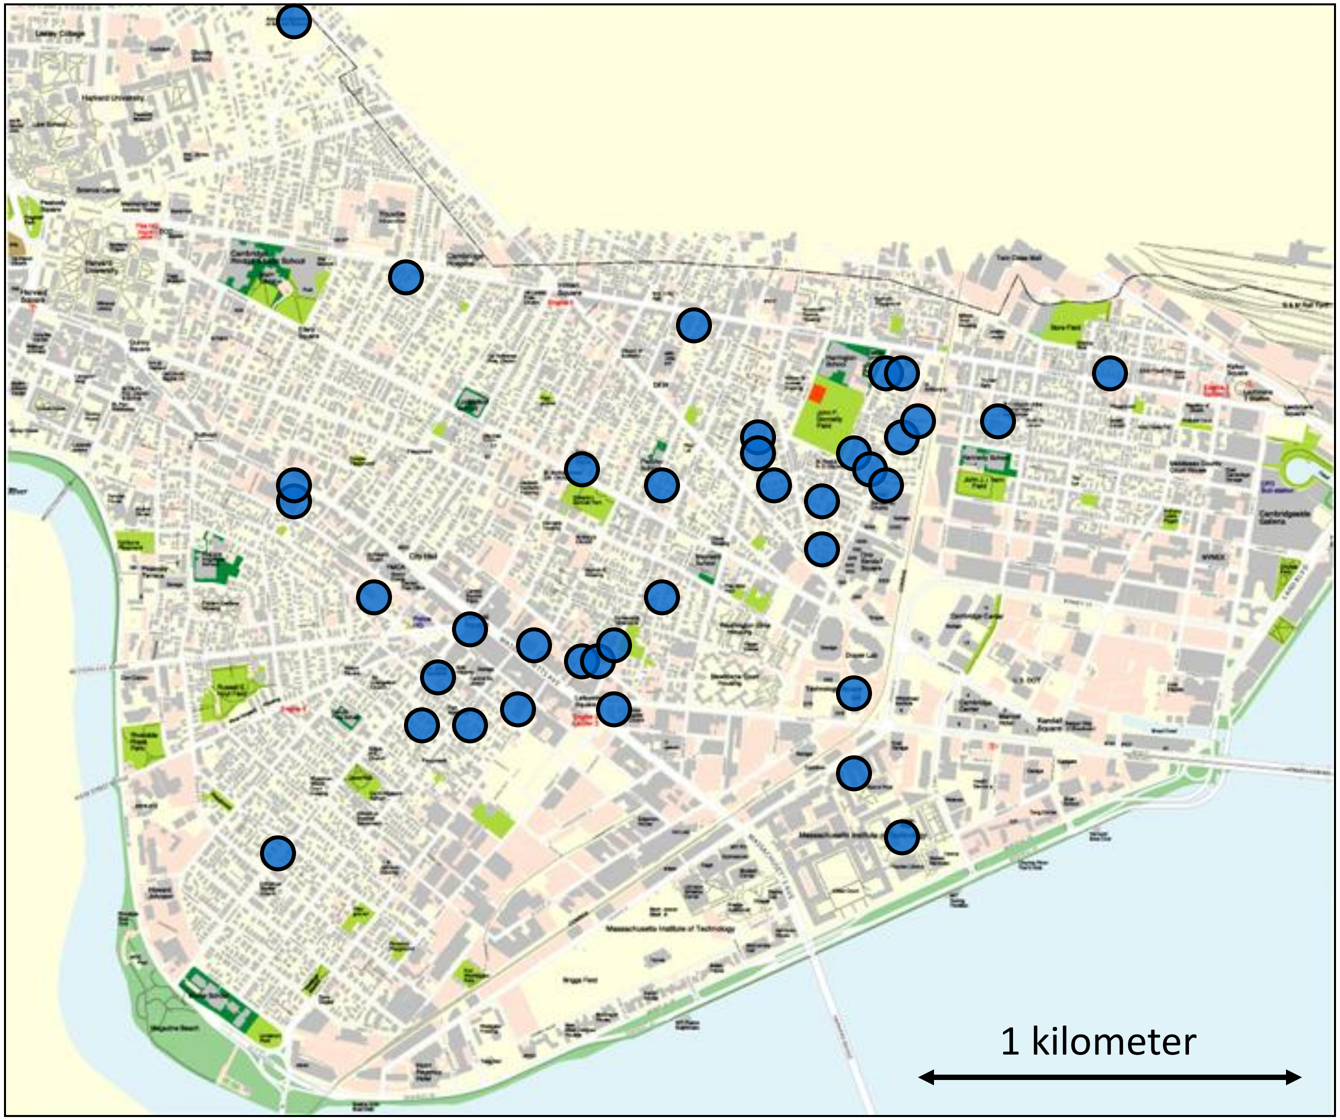

Evaluation Details

- 65 Node pairs
- 1.0MByte file transfer
- 1 Mbit/s 802.11 bit rate
- 1 KByte packets

Traditional Routing	ExOR
802.11 unicast with link-level retransmissions Hop-by-hop batching UDP, sending as MAC allows	802.11 broadcasts 100 packet batch size

ExOR: 2x overall improvement

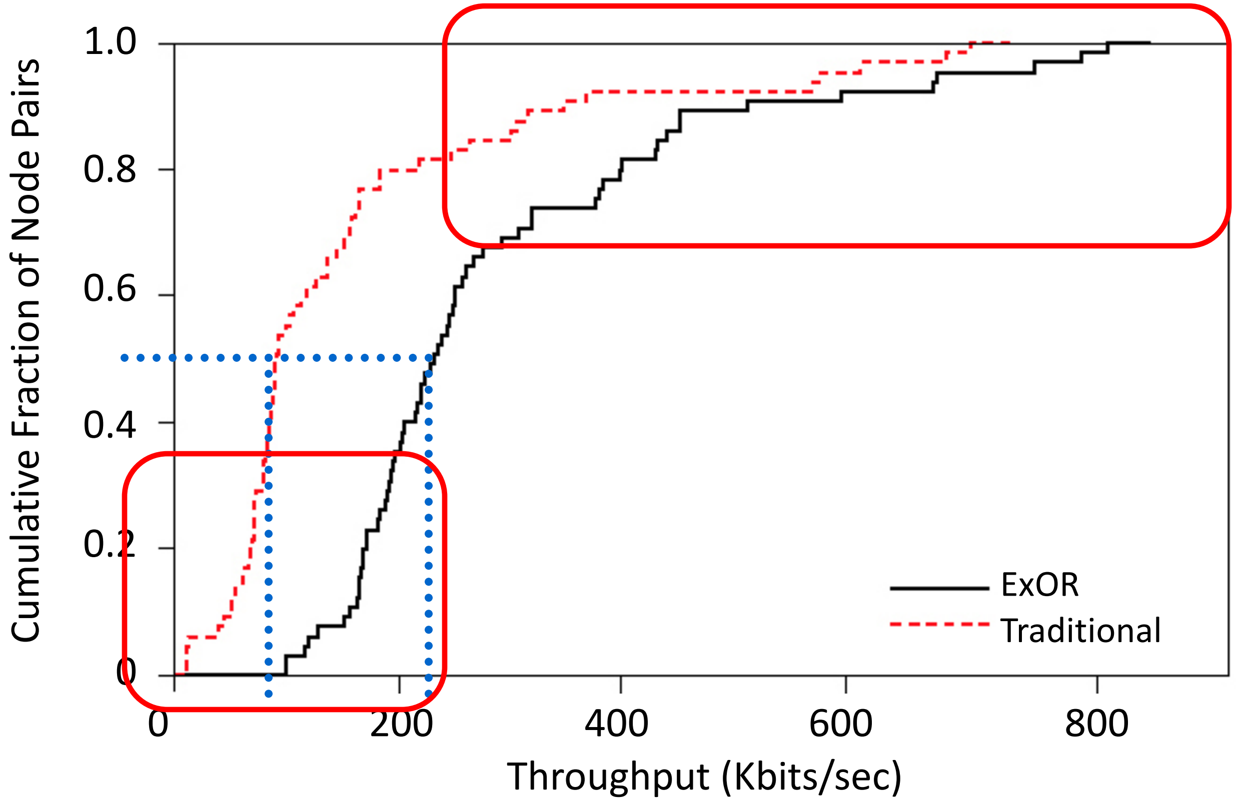

- Median throughputs: 240 Kbits/sec for ExOR,
  121 Kbits/sec for Traditional

25 Highest throughput pairs

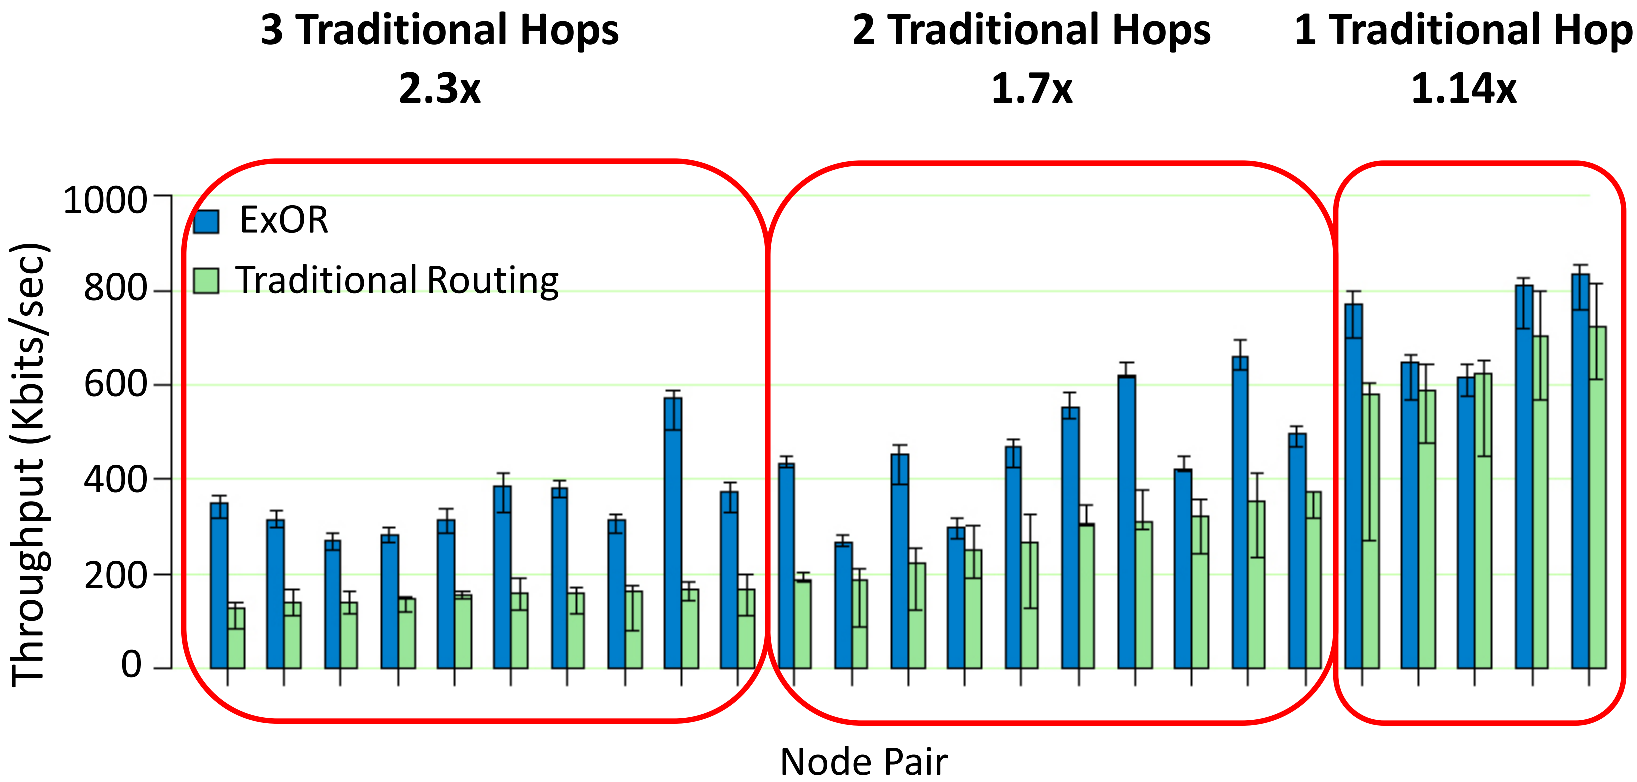

25 Lowest throughput pairs

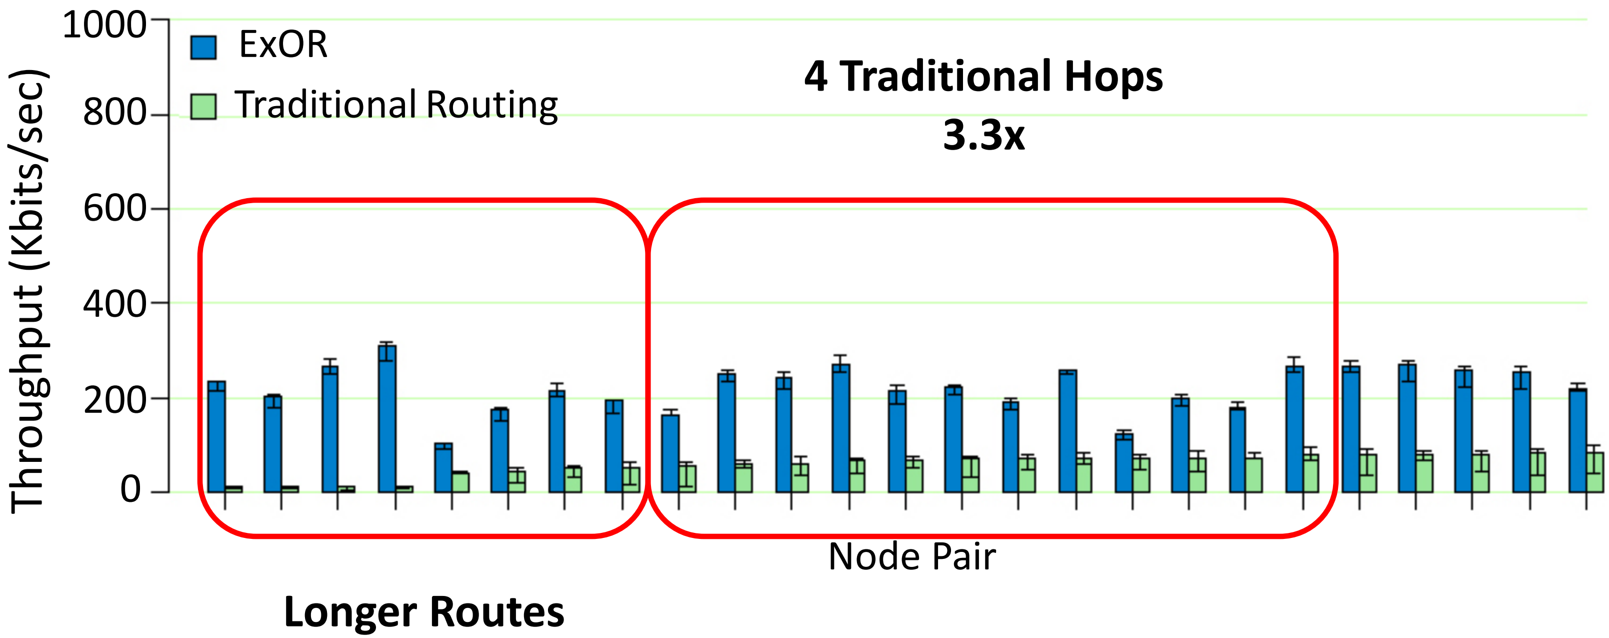

Gain is typically higher for the worst node pairs

ExOR uses links in parallel

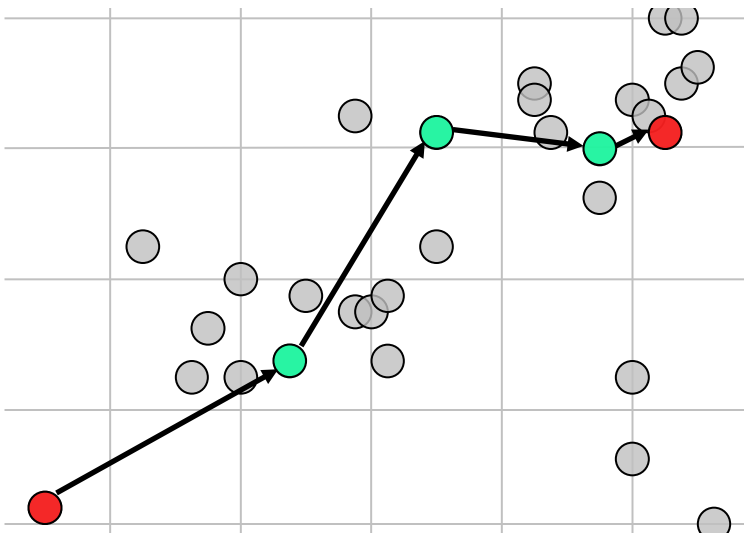

Traditional Routing
3 forwarders
4 links

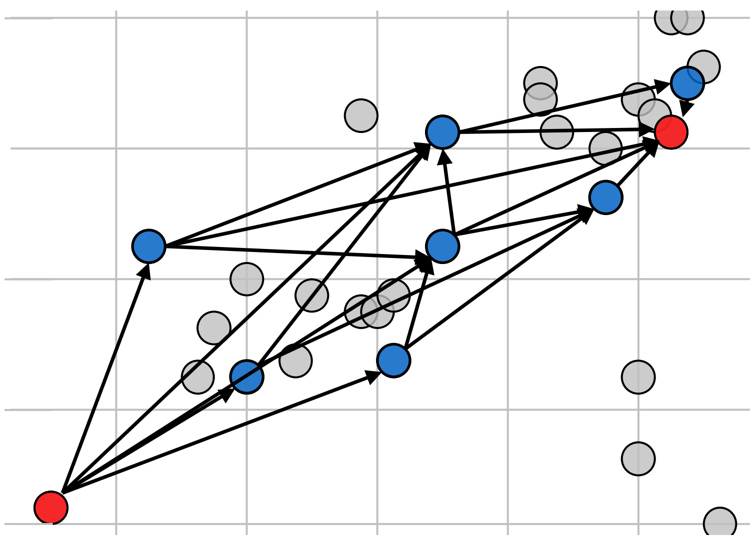

ExOR
7 forwarders
18 links

Limitations of ExOR

- Does not account for the fact that different links can support different bitrates
- Batching is ineffective with short flows or realtime traffic
- Only one node can transmit at any point in time, even if no contention occurs between nodes
- Requires global coordination of the nodes in the network

Lecture Outline

- Single Path Routing
- Opportunistic Routing with ExOR
- ⇨ Intra-Flow Network Coding with MORE

Sender Is in a Bad Spot

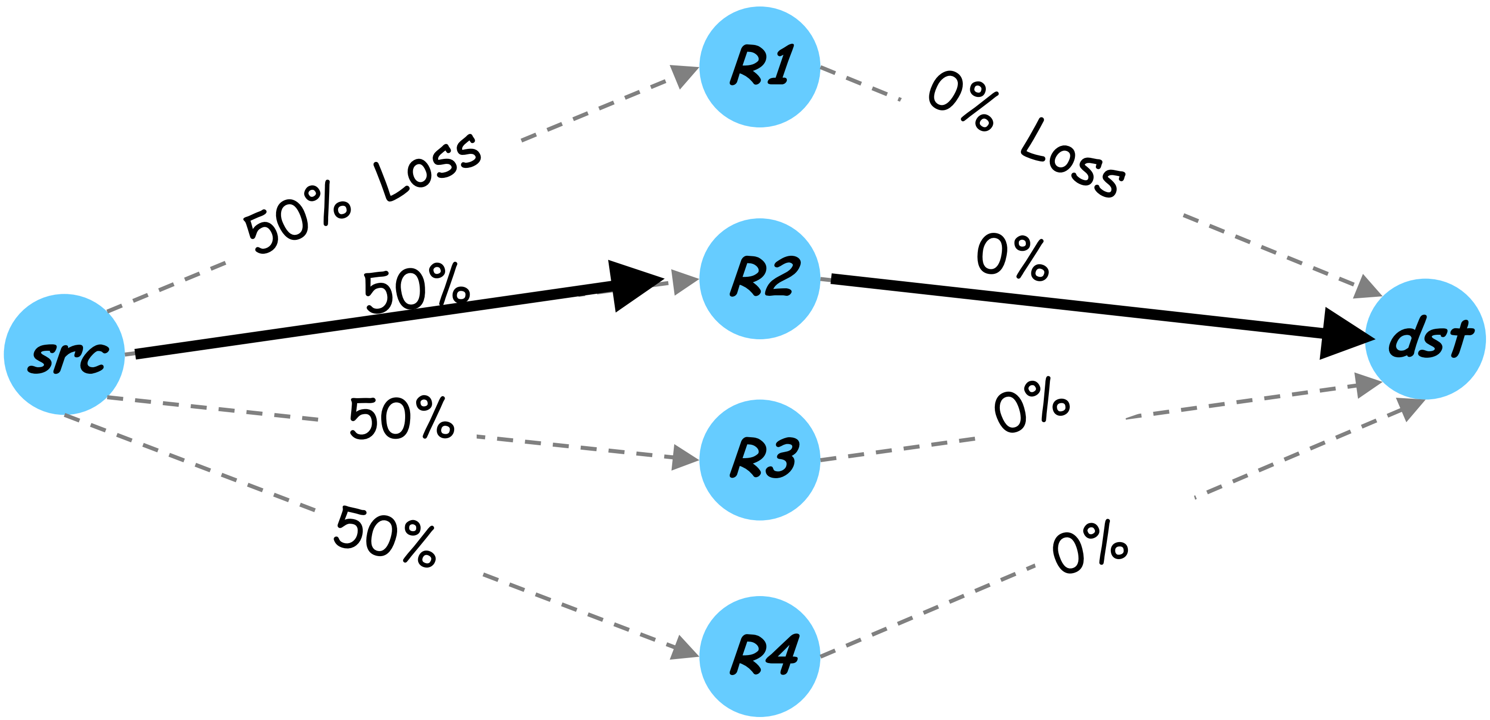

- Best single path → Prob. of loss 50%

Use Opportunistic Routing

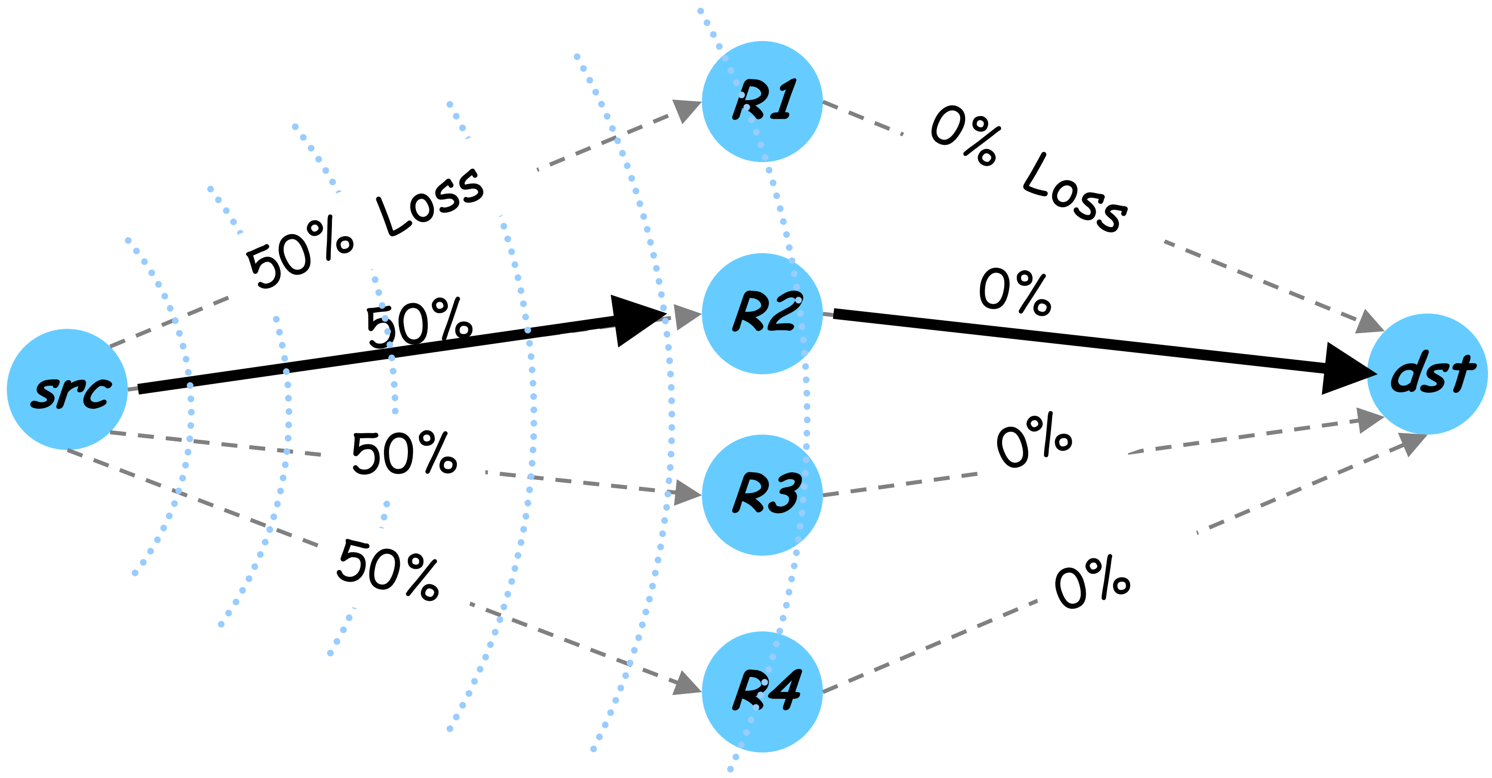

- Best single path → Prob. of loss 50%
- Opportunistic Routing can do better
  - Any router can forward packet → Prob. of loss $0.5^4 = 6\%$

Challenge

Overlap in received packets → Routers forward duplicates

Challenge

Overlap in received packets → Routers forward duplicates

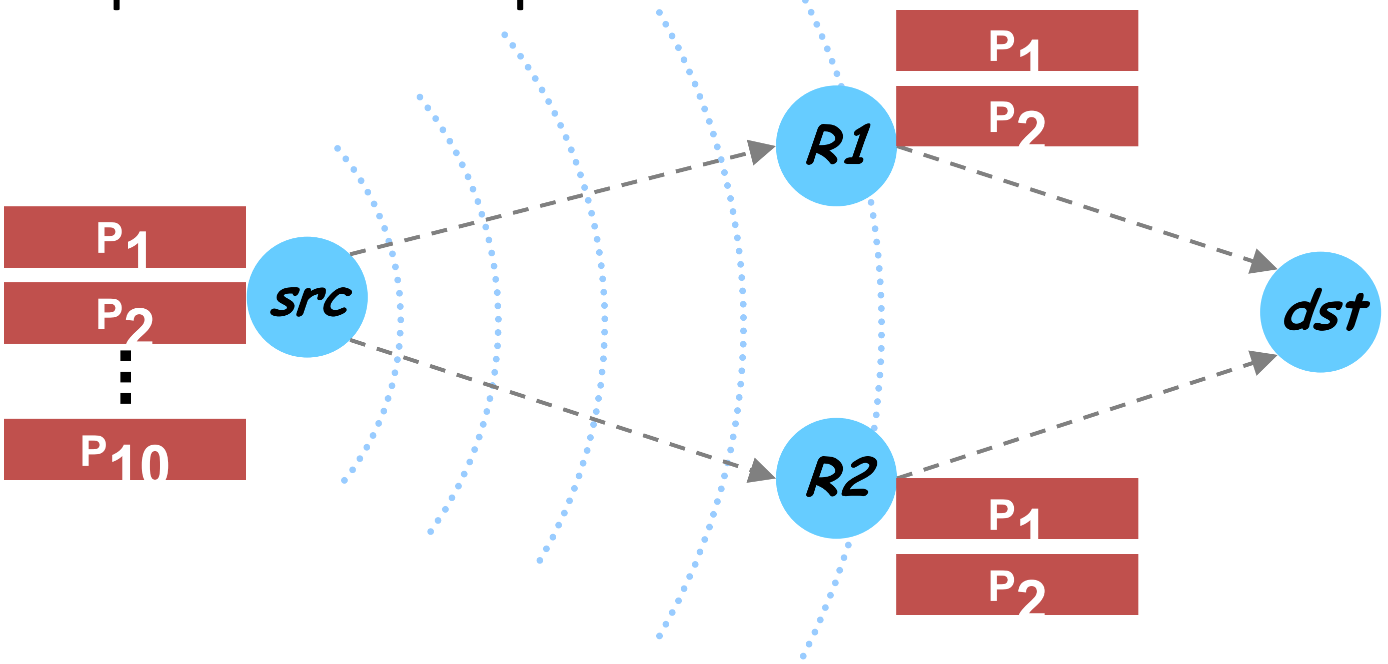

ExOR imposes a global scheduler:

- Nodes have to agree on who transmits what
- Only one node transmits at a time, others listen

Limitations of ExOR

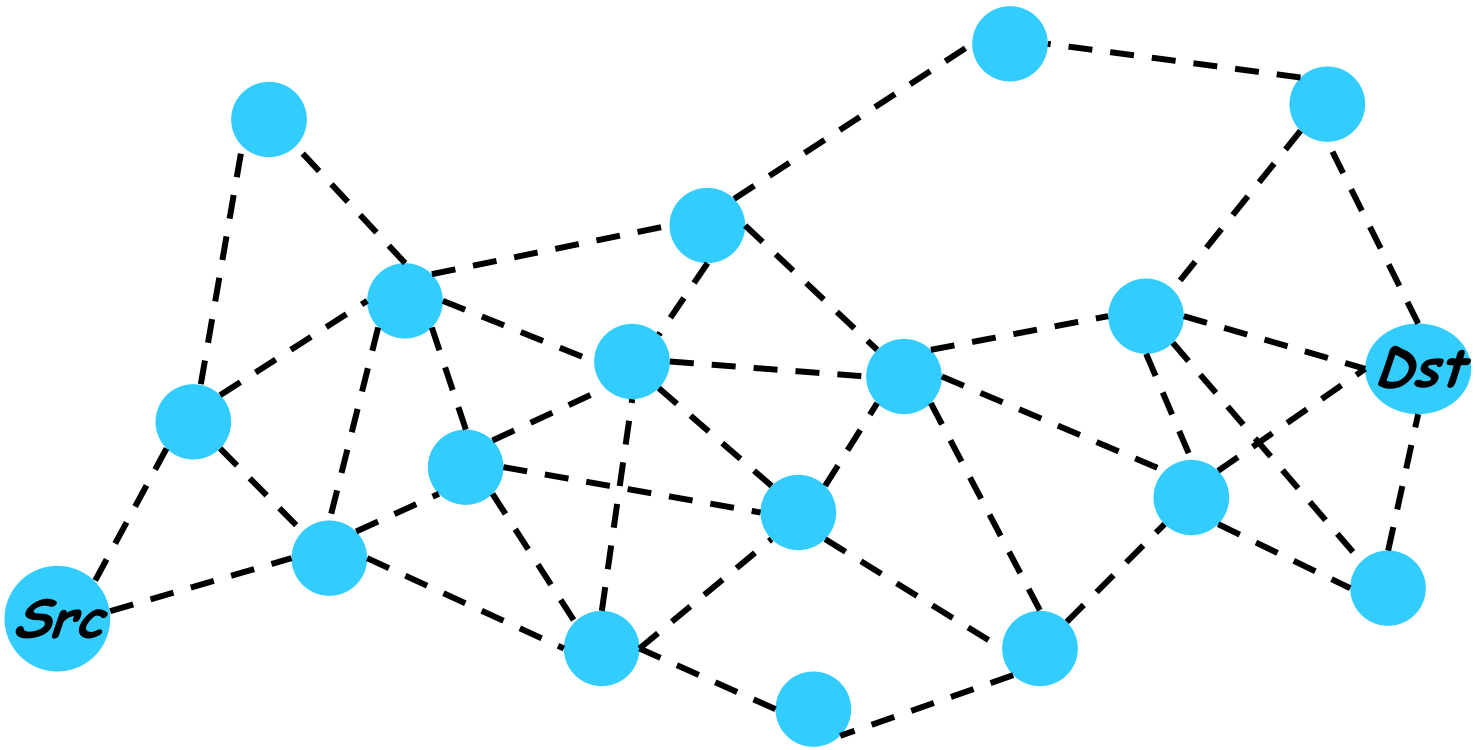

- Learning who received what → too much overhead
- Forcing only one transmitter at a time → prevents spatial reuse of the medium

Limitations of ExOR

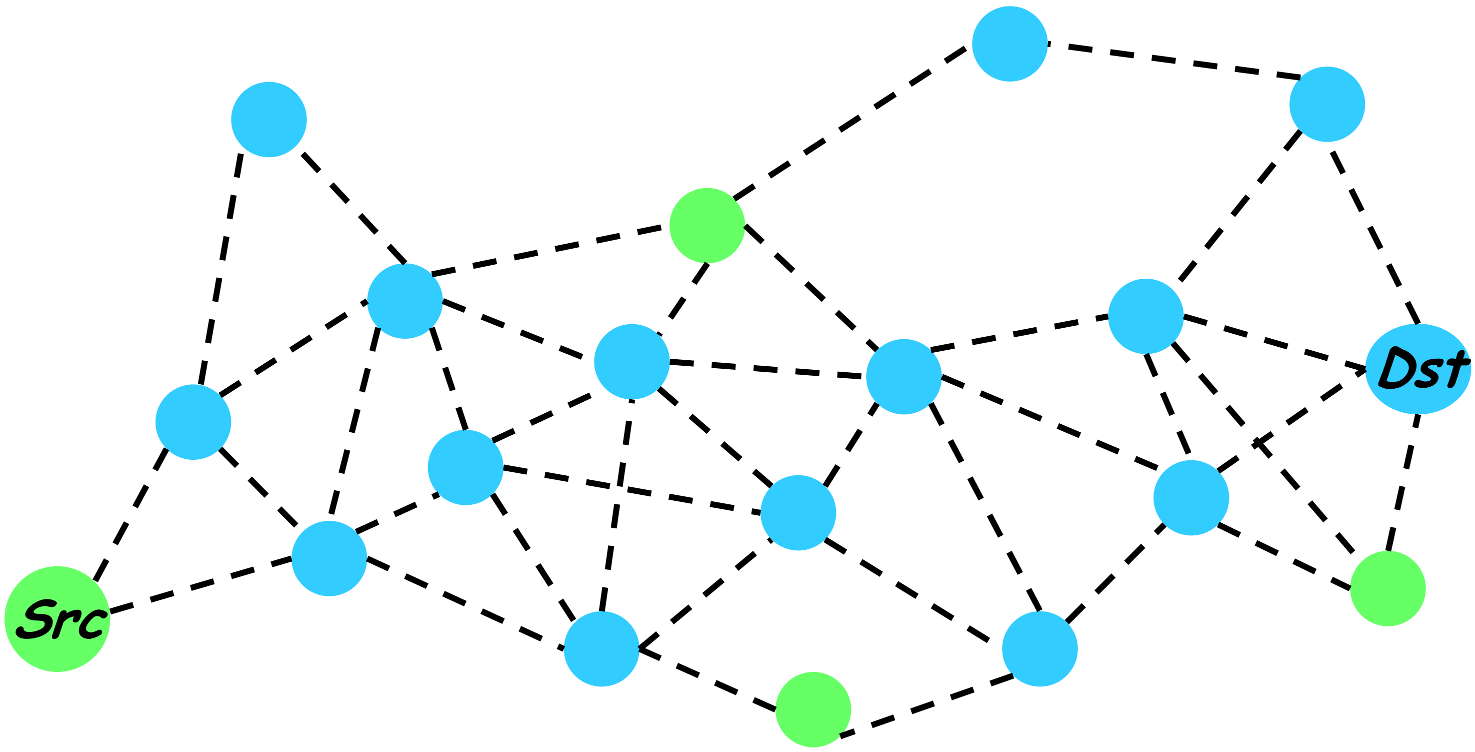

- Learning who received what → too much overhead
- Forcing only one transmitter at a time → prevents spatial reuse of the medium

Can we eliminate these problems?

Solution: Random Network Coding

Each router forwards random combinations of packets

Solution: Random Network Coding

Each router forwards random combinations of packets

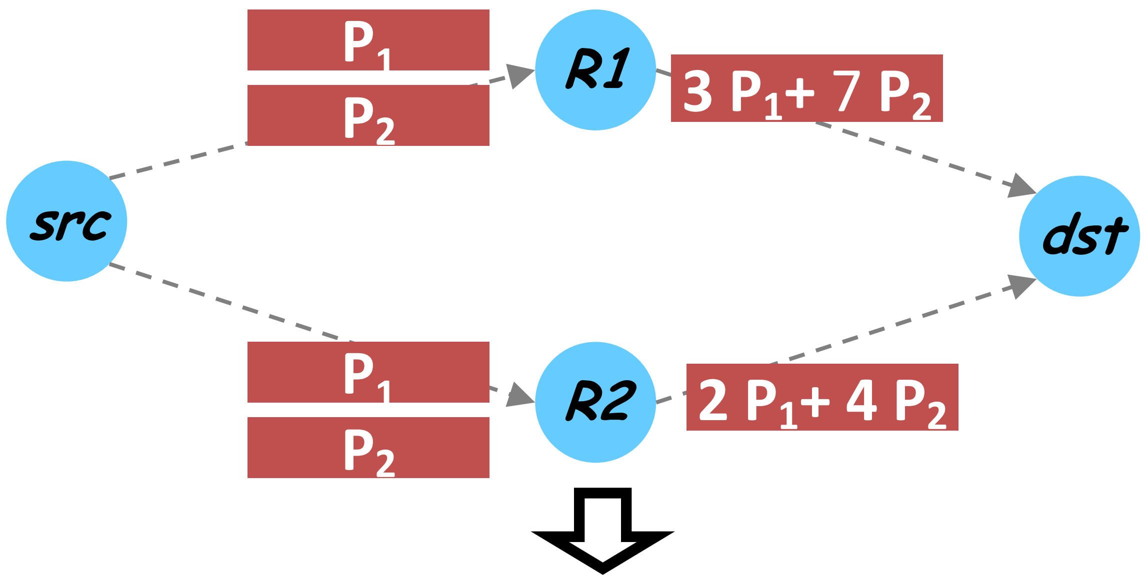

Randomness prevents duplicates

No need to know who received what

Can exploit spatial reuse

Network Coding Also Benefits Multicast

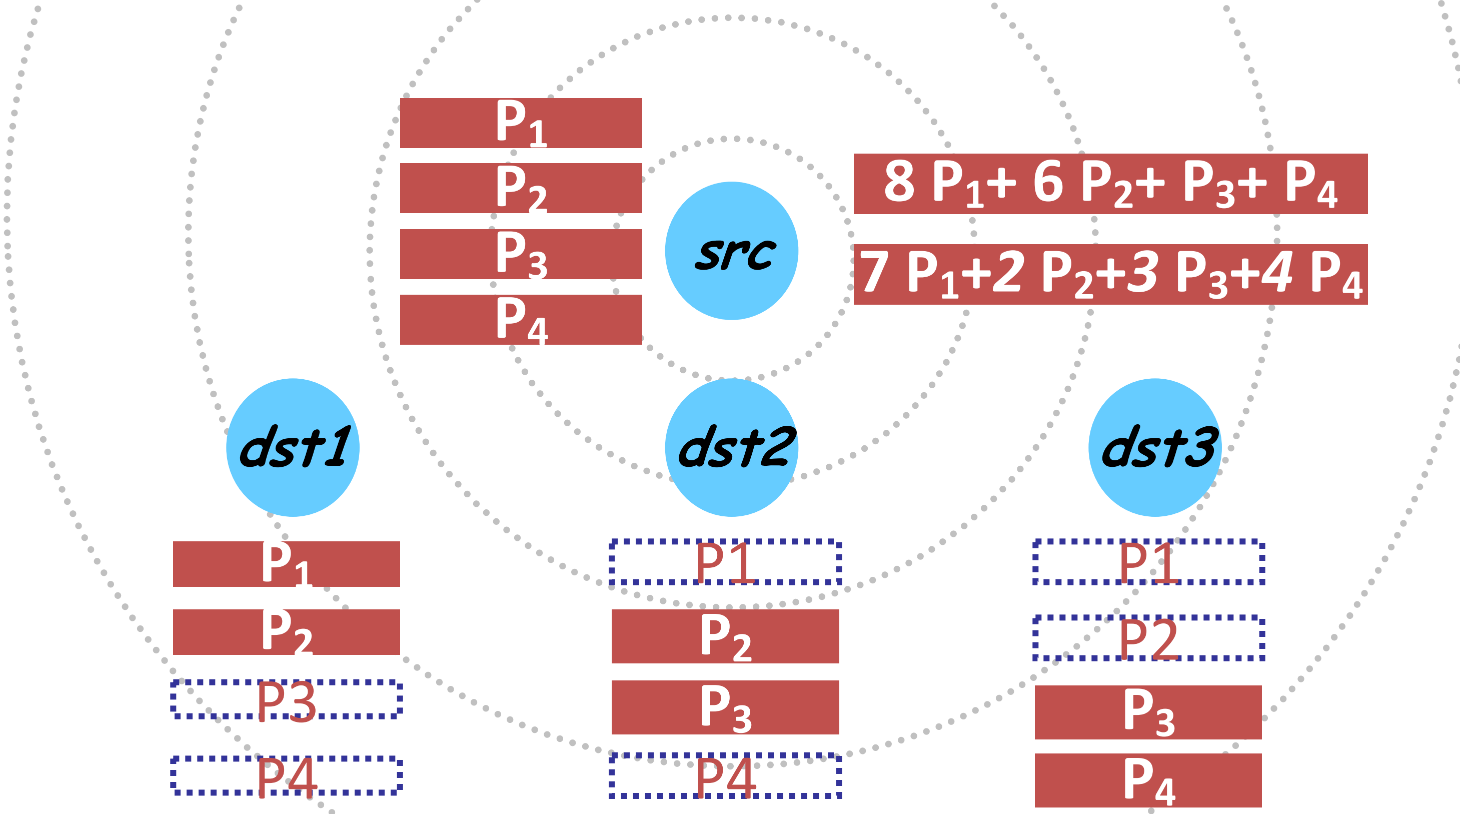

Without coding → source has to retransmit all 4 packets

With network coding → 2 packets are sufficient

MORE

- An opportunistic routing protocol that reduces overhead and enables spatial reuse

- It is based on network coding, where routers code packets together before forwarding them
  - »

How Does MORE Work?

- Source sends packets in batches
- Forwarders keep all heard (innovative) packets in a buffer
- Nodes transmit linear combinations of buffered packets

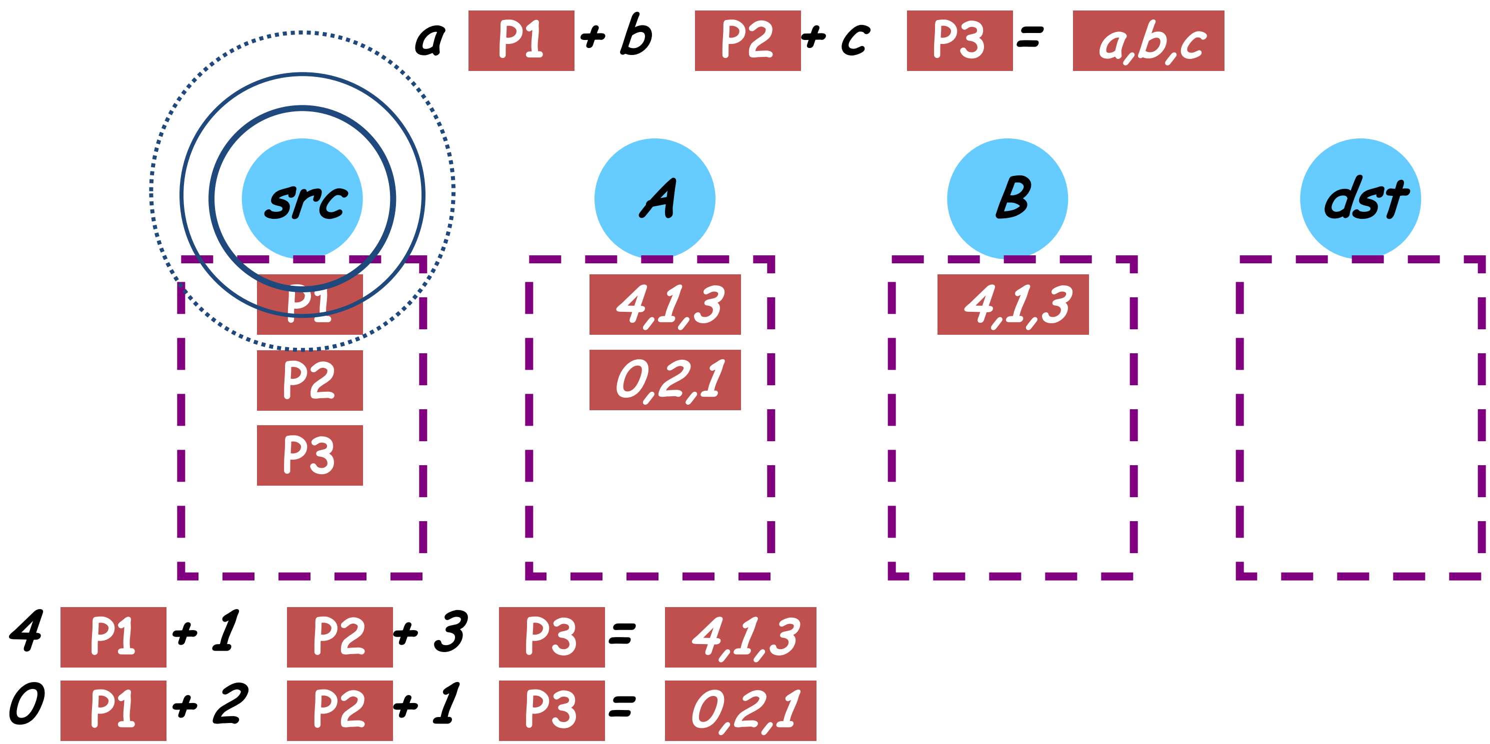

How Does MORE Work?

- Source sends packets in batches
- Forwarders keep all heard (innovative) packets in a buffer
- Nodes transmit linear combinations of buffered packets

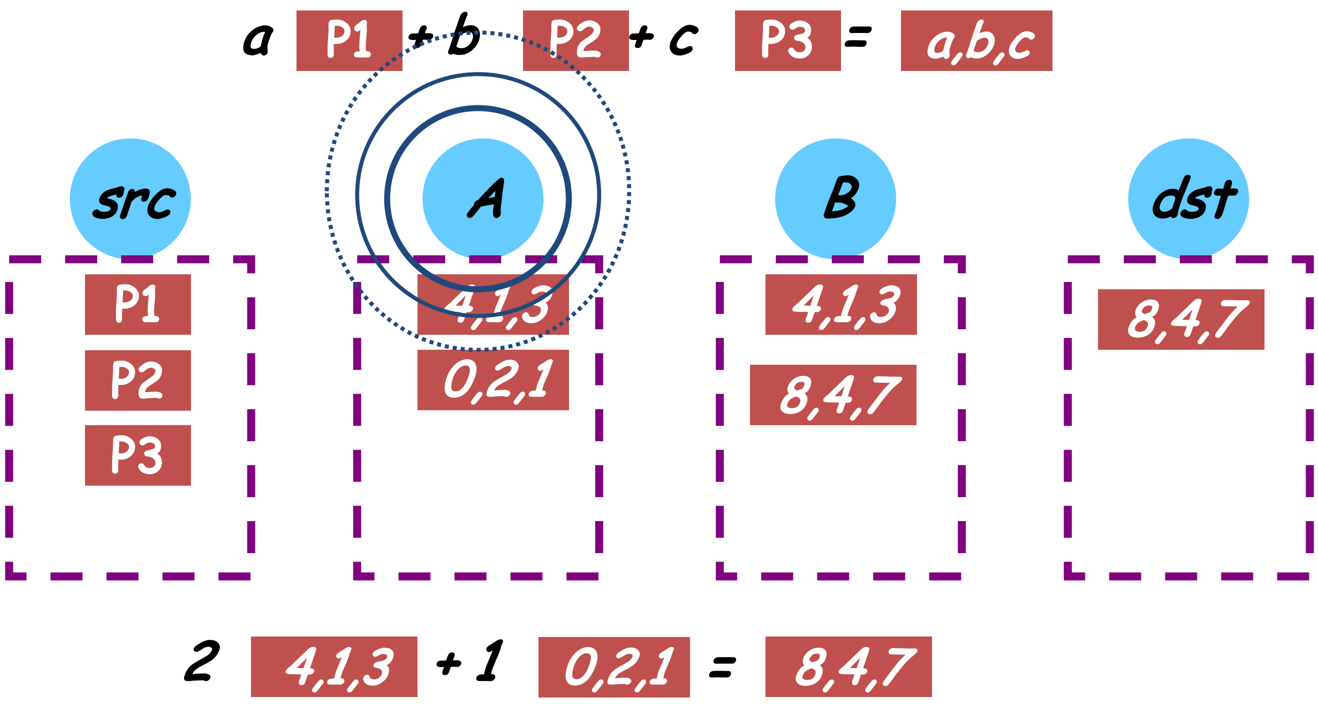

How Does MORE Work?

- Source sends packets in batches
- Forwarders keep all heard (innovative) packets in a buffer
- Nodes transmit linear combinations of buffered packets

- Destination decodes once it receives enough combinations
  - Say batch is 3 packets

1 P1 + 3 P2 + 2 P3 = 1,3,2

5 P1 + 4 P2 + 5 P3 = 5,4,5

4 P1 + 5 P2 + 5 P3 = 4,5,5

- Decoding is solving linear equations
- Once it decoded a batch, the destination acks the batch and the source moves to next batch

Dealing with Interference

MORE uses the 802.11 MAC as it is

- In contrast to ExOR where a node has to wait for downstream nodes to deliver their packets, any node that senses the medium available can transmit
- Thus, it allows spatial reuse

But How Do We Get the Most Throughput?

- Naïve approach transmits whenever it can
- Inefficient ☹

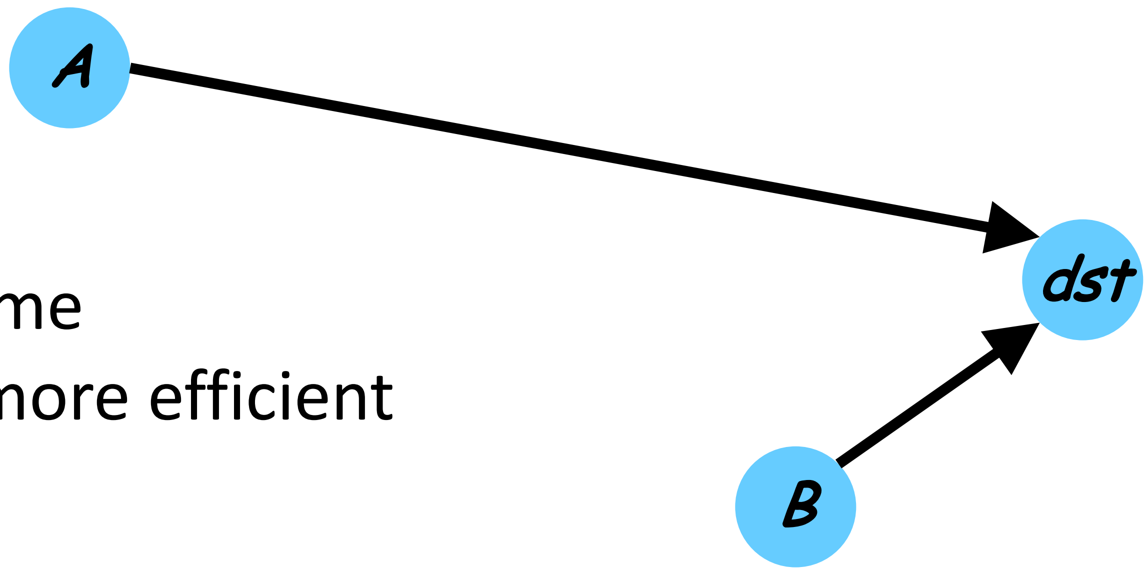

If *A* and *B* have same information, it is more efficient for B to send it

Need to control how much a node transmits

Probabilistic Forwarding

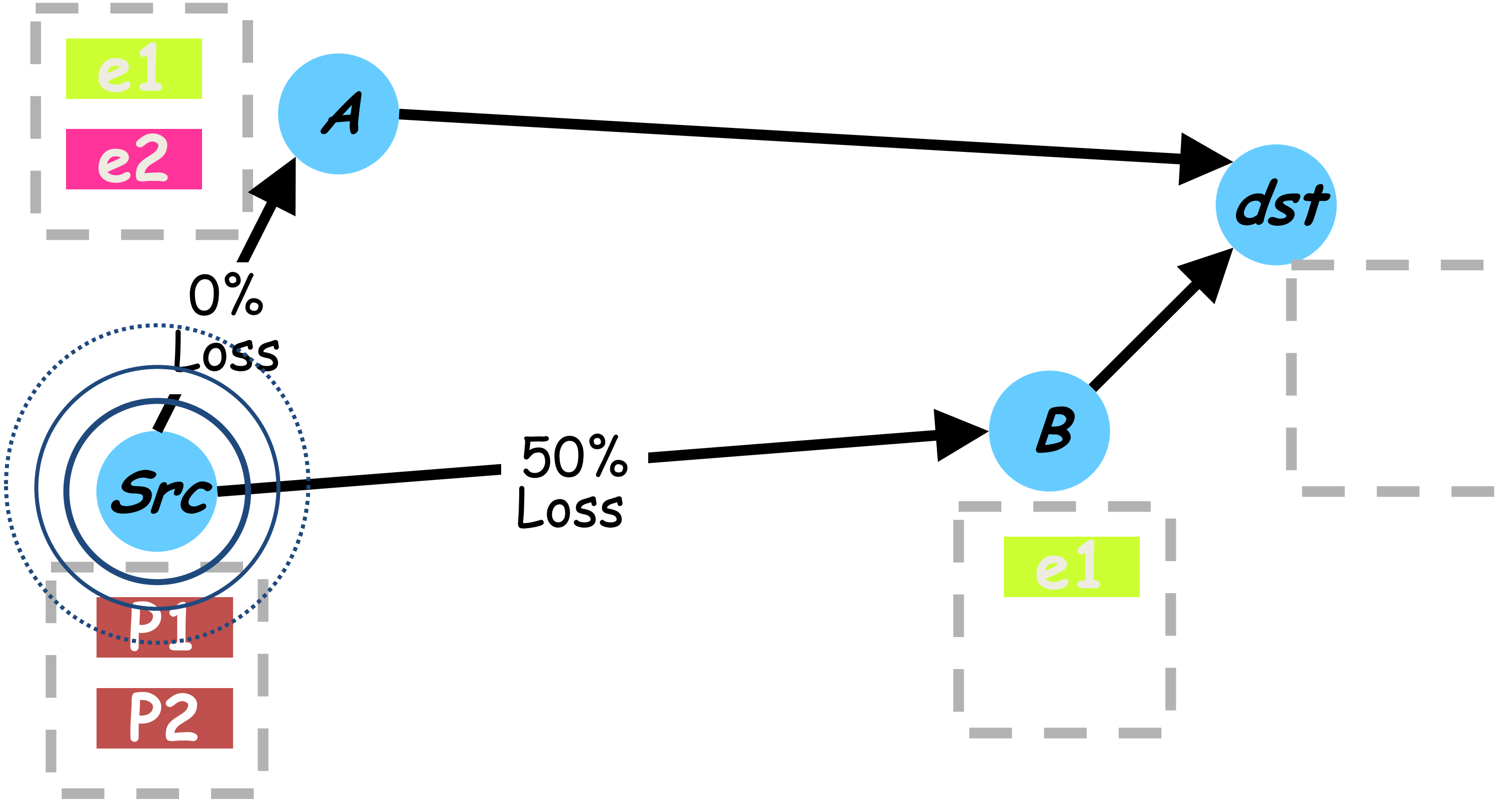

Probabilistic Forwarding

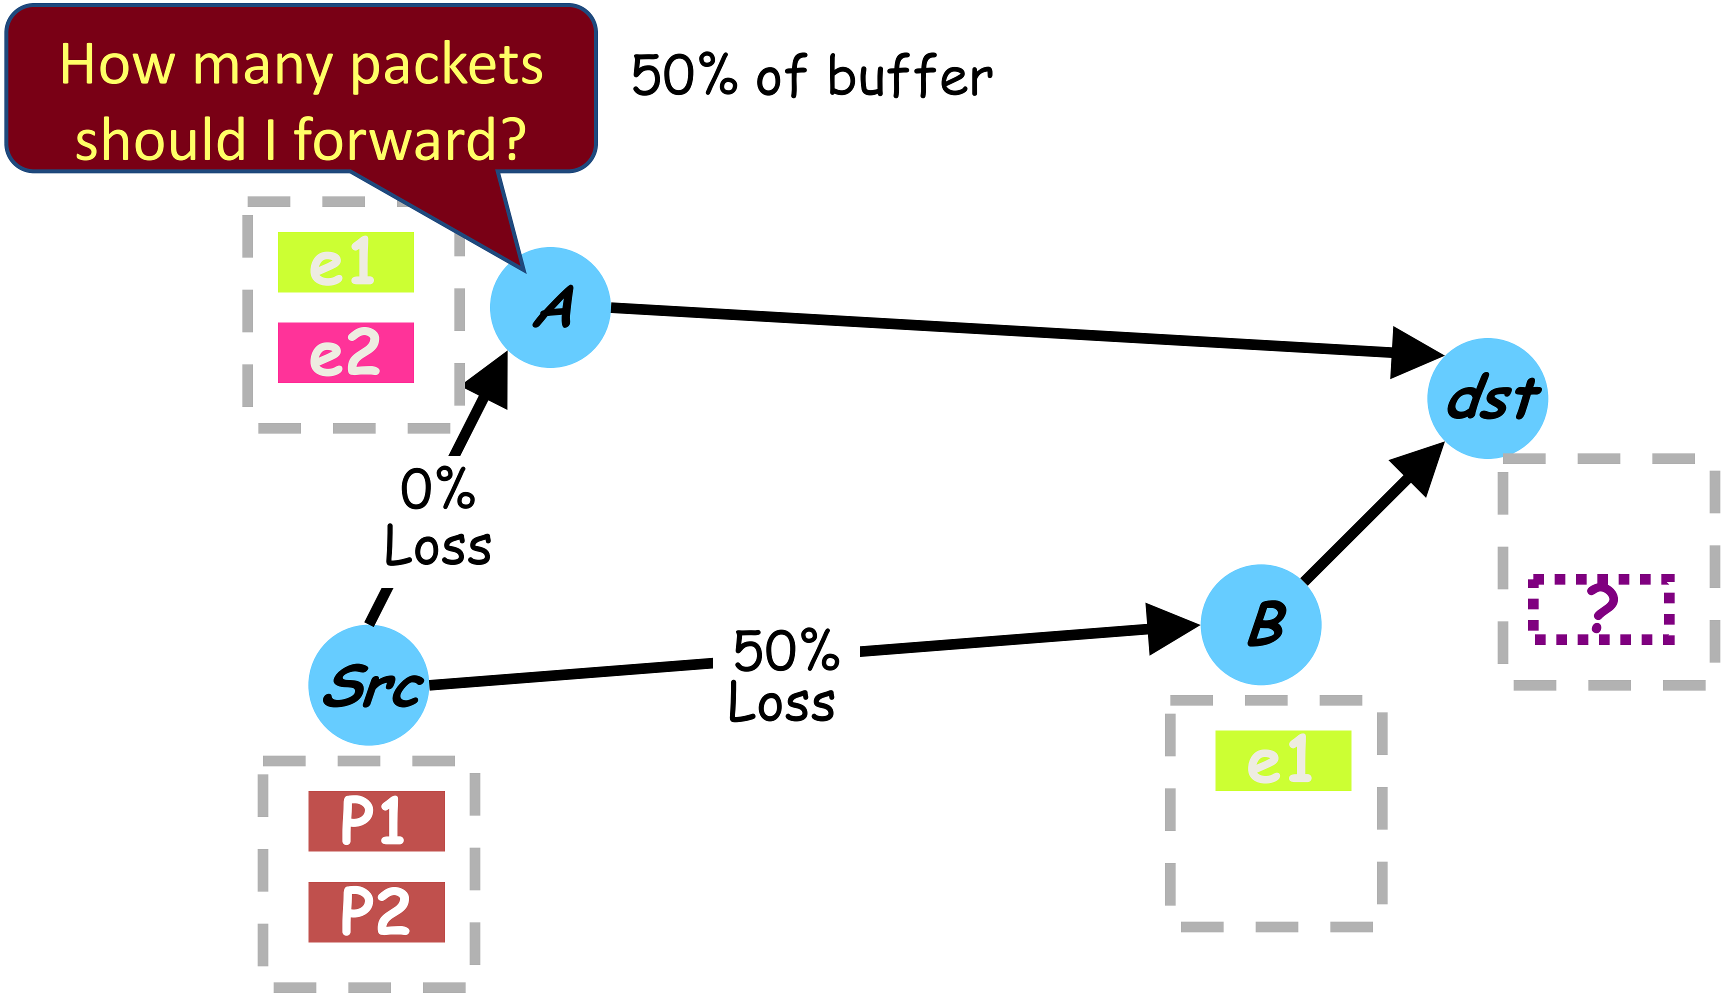

Need to know the fraction of overlap, not which packets overlap

Probabilistic Forwarding

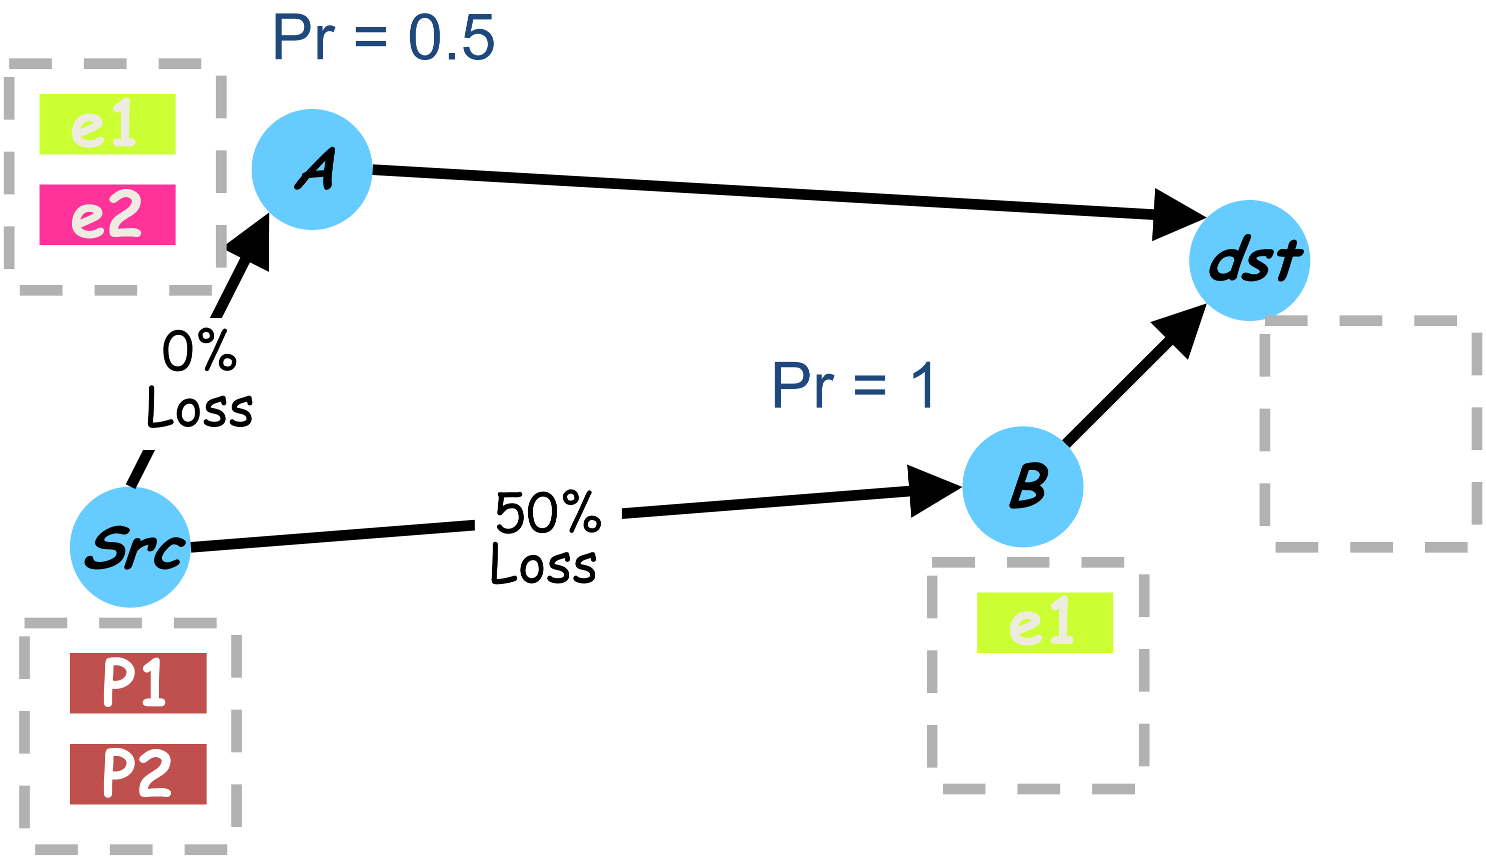

Compute forwarding probabilities without coordination using loss rates

ExOR Can't Do It

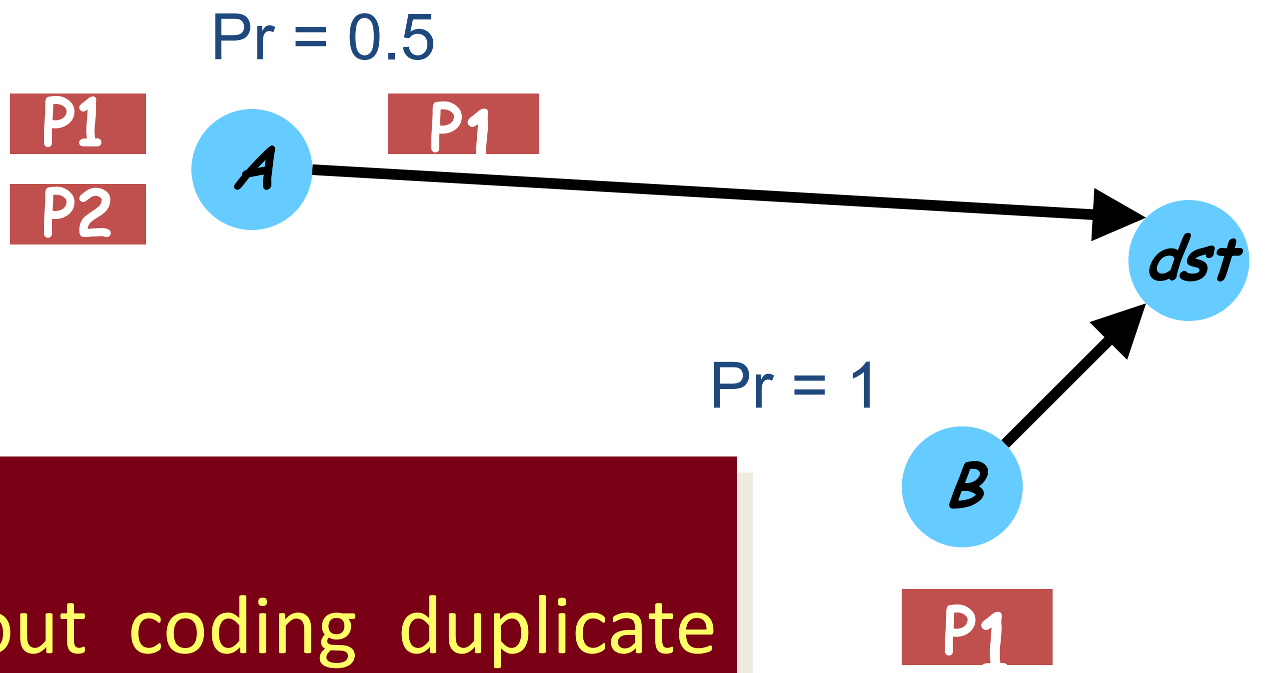

Without coding duplicate probability is 50%

Testbed

- 20-node testbed over three floors

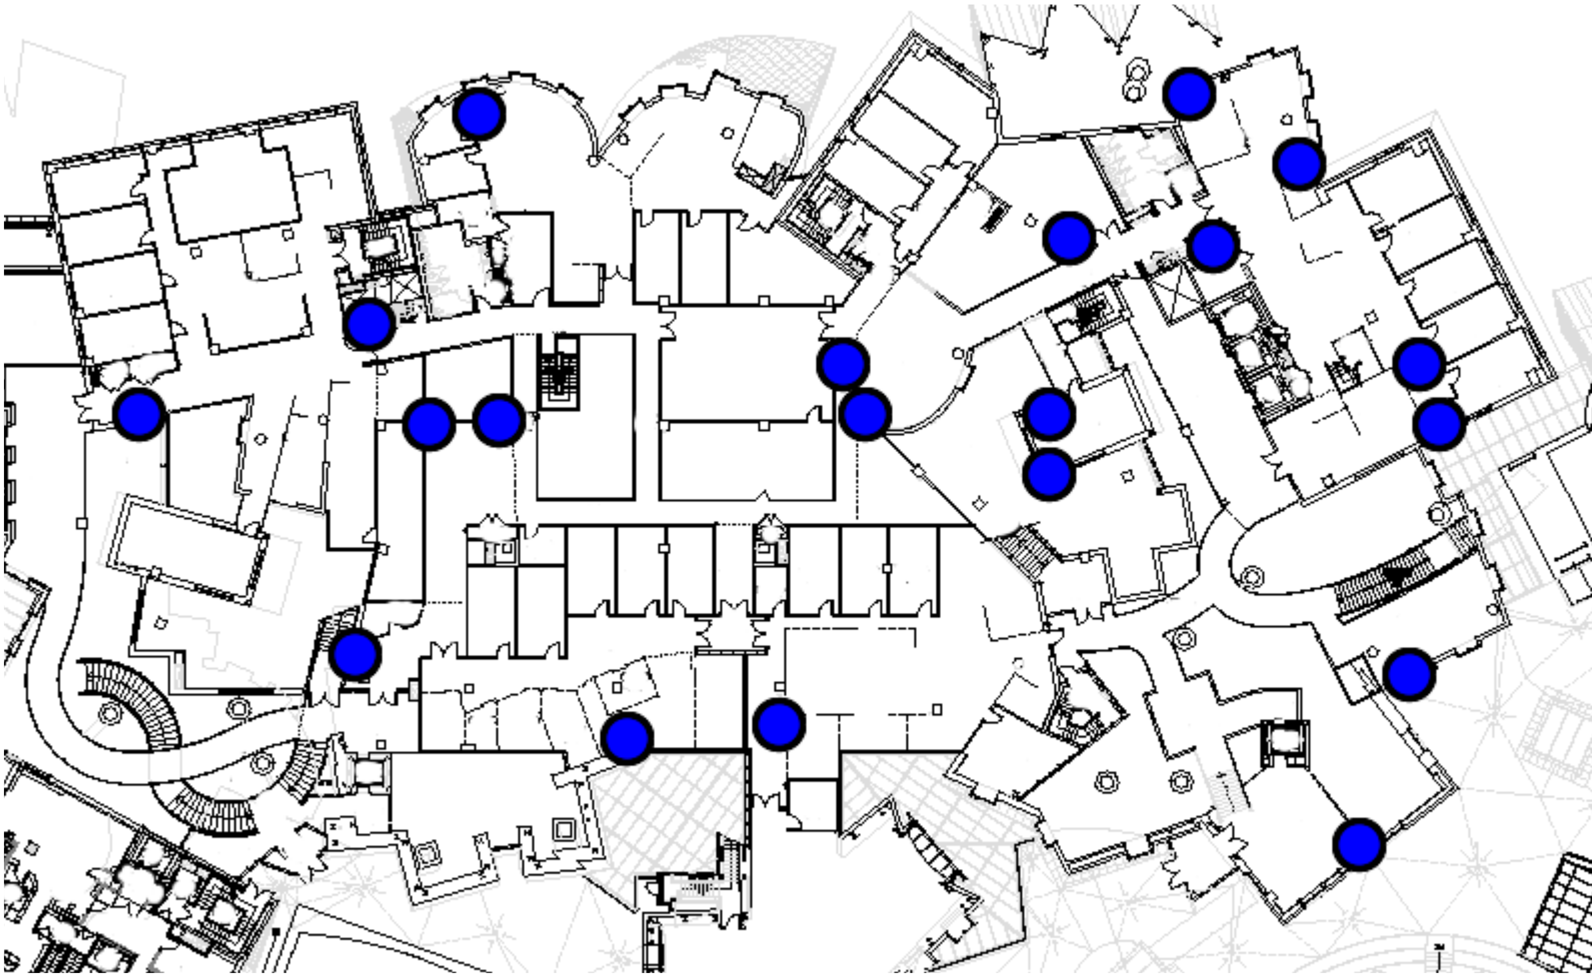

Testbed

- 20-node testbed over three floors

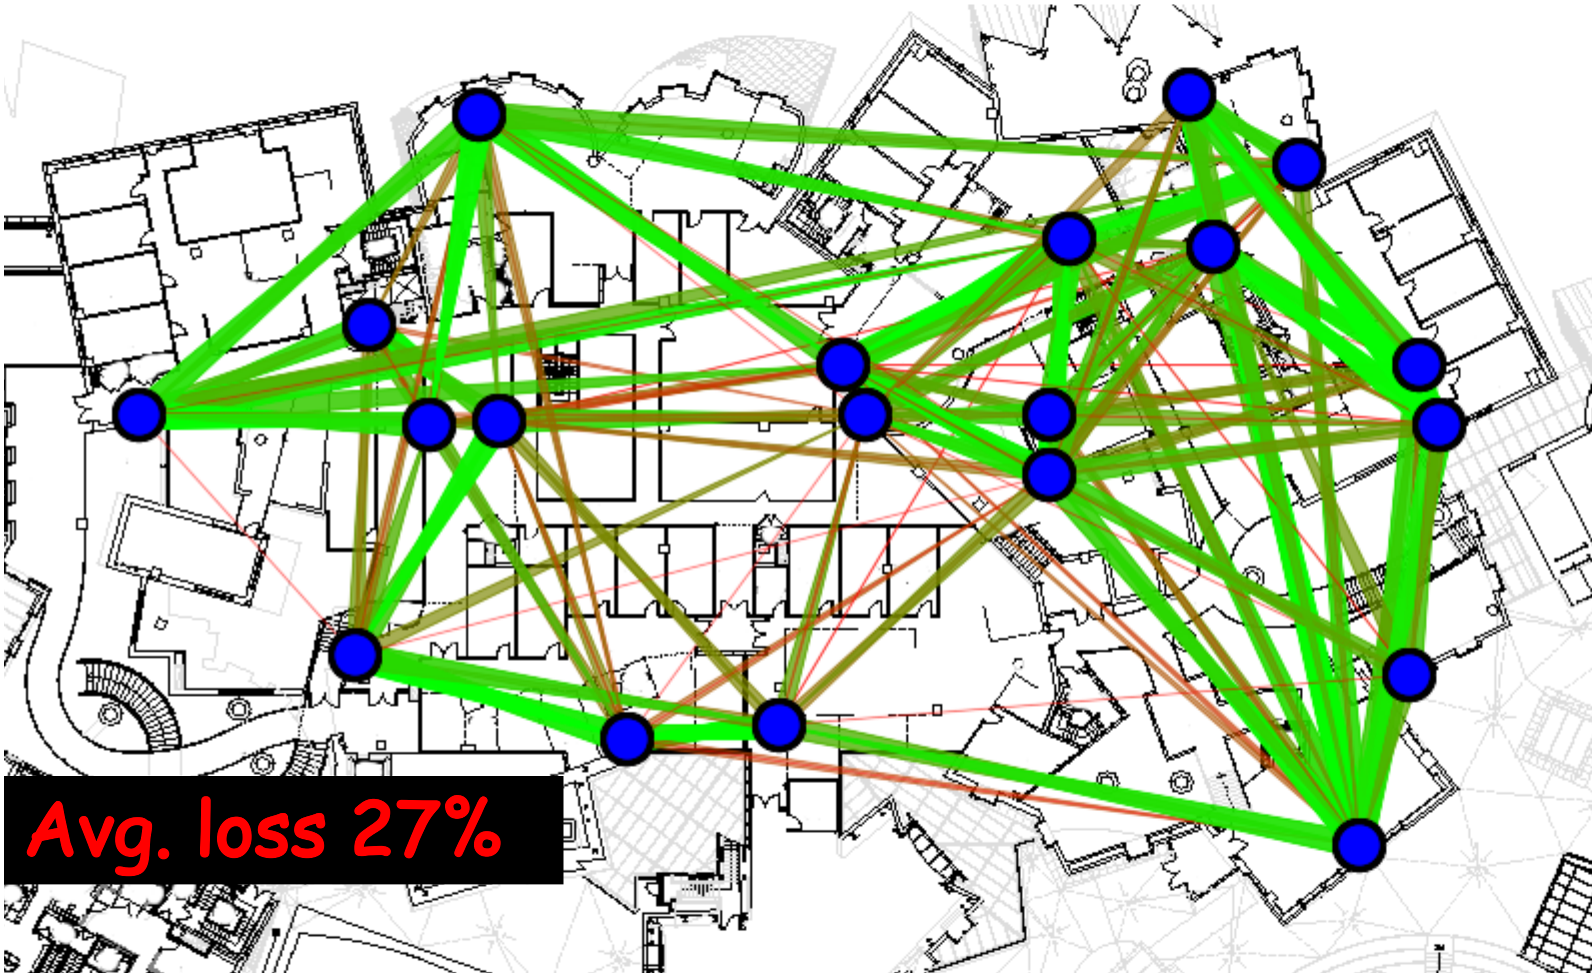

Does MORE Improve Wireless Throughput?

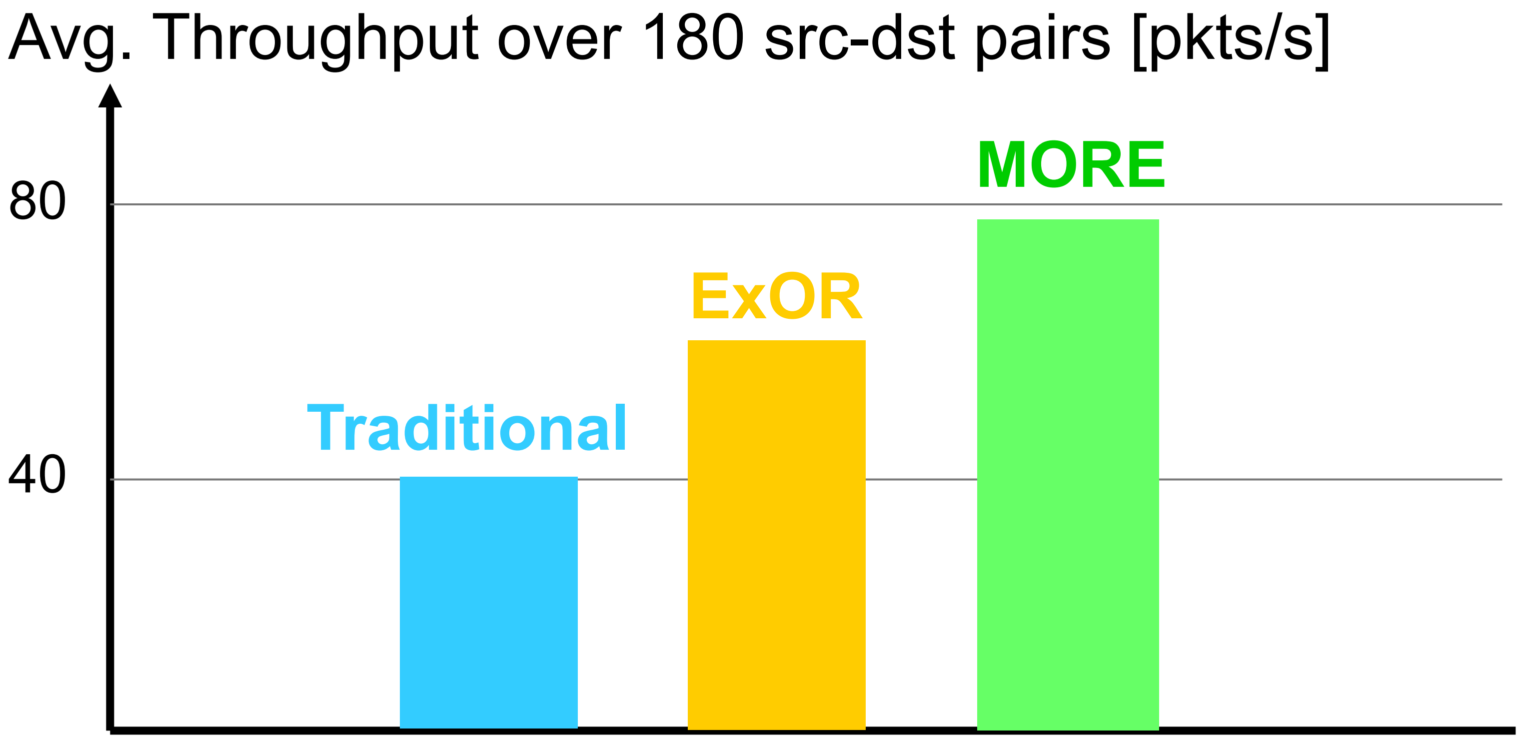

MORE's throughput is double current routing and is 22% better than ExOR

Throughput of All Source-Destination Pairs

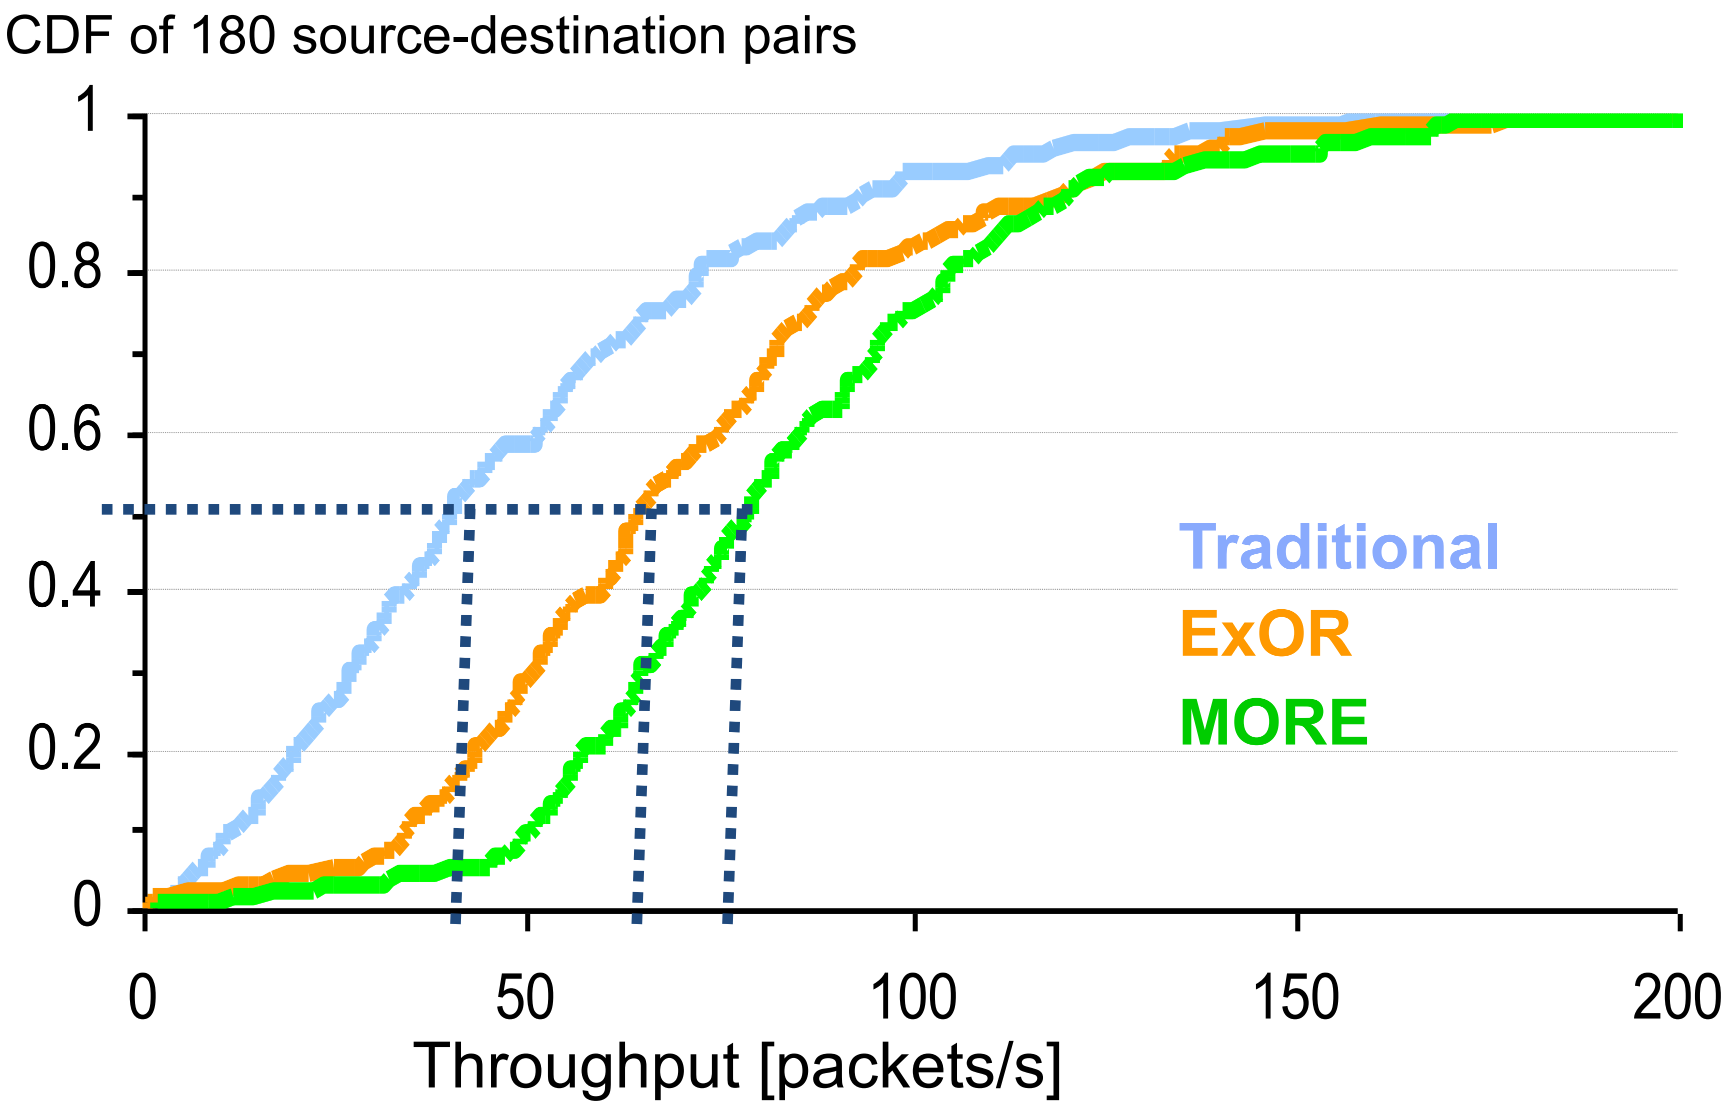

Zoom in on the worst 10%

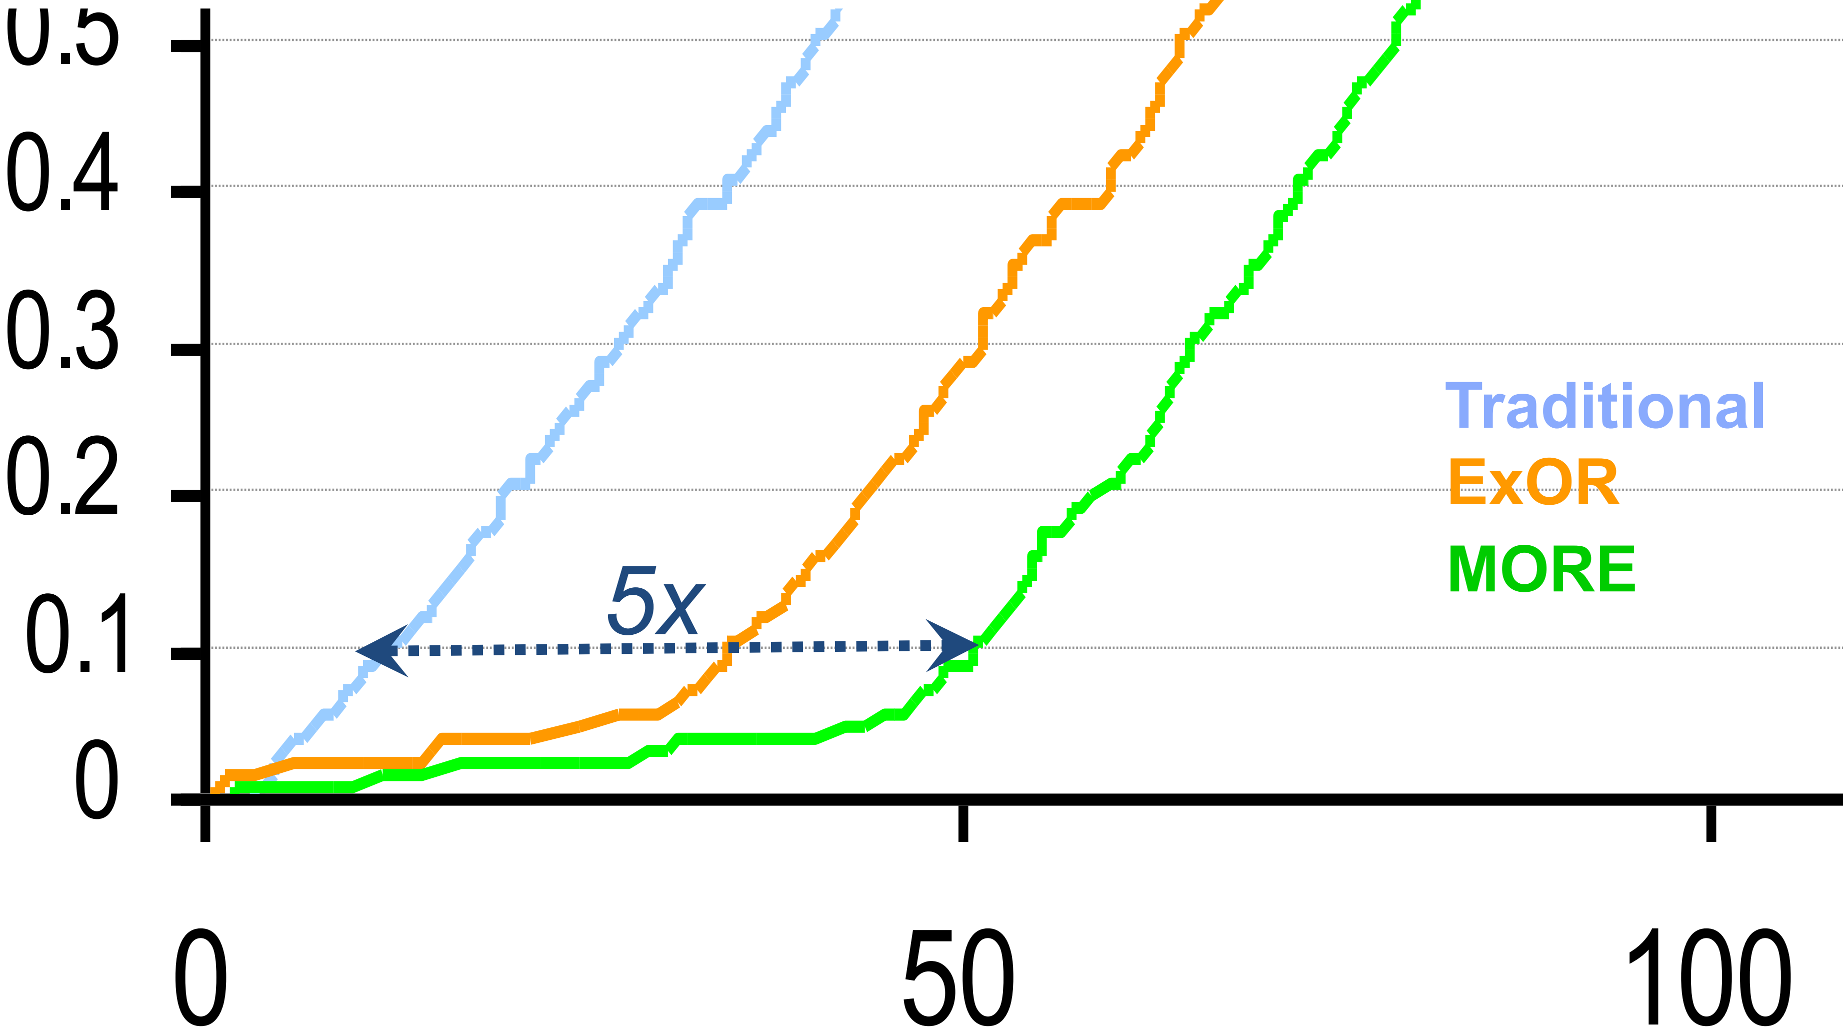

Similarly to ExOR, the gains are large for the worst connections MORE

Network Coding

- Requires less coordination
  - No scheduler
- More flexibility
  - One framework for unicast and multicast
- More throughput
  - 22% more than ExOR and 95% more than current shortest path routing

Limitations of MORE

- Similarly to ExOR, does not deal with the fact that different links can support different bit rates
- Batching is undesirable though the batch size is smaller in MORE