Routing in
Wireless Mesh Networks

Multi-hop Wireless Networks

	Stationary Nodes	Mobile Nodes
Motivating scenario	Community wireless networks (Mesh Networks)	Battlefield networks
Key challenge	Improving Network Capacity	Handling mobility, limited power.

Wireless Mesh Networks
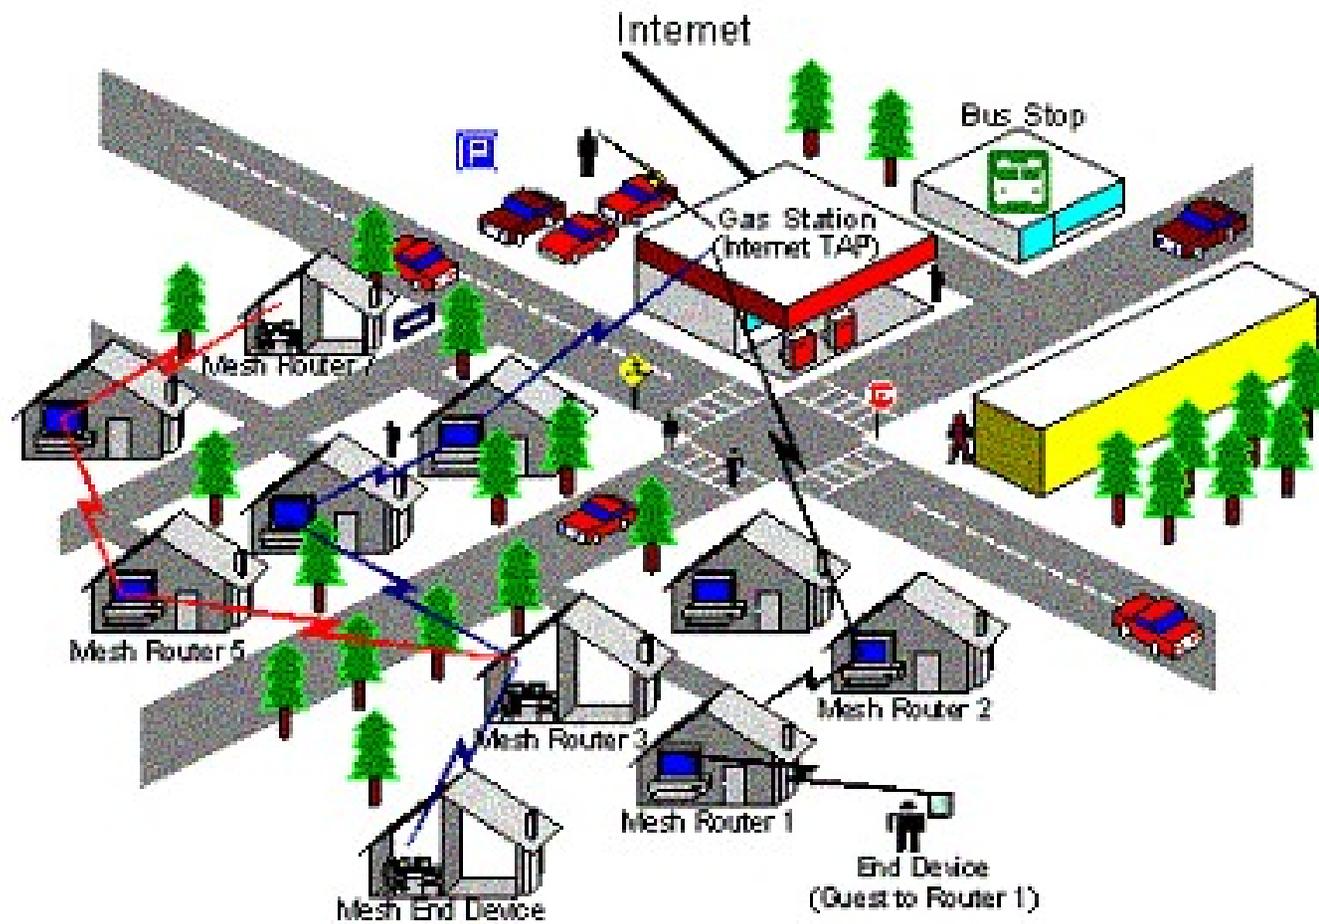
What do you think are good routing metrics?

Potential Ideas (and their cons)

❒ Product of per-link delivery ratios

❒ Throughput of a path's bottleneck link

❒ End-to-end delay

Potential Ideas (and their cons)

☐ Product of per-link delivery ratios

  ○ A perfect 2-hop route is viewed as better than a 1-hop route with 10% loss ratio

☐ Throughput of a path's bottleneck link

  ○ Same as above

☐ End-to-end delay

  ○ Changes with network load as queue lengths vary ... can cause oscillations

A High-Throughput Path Metric for Multi-Hop Wireless Routing

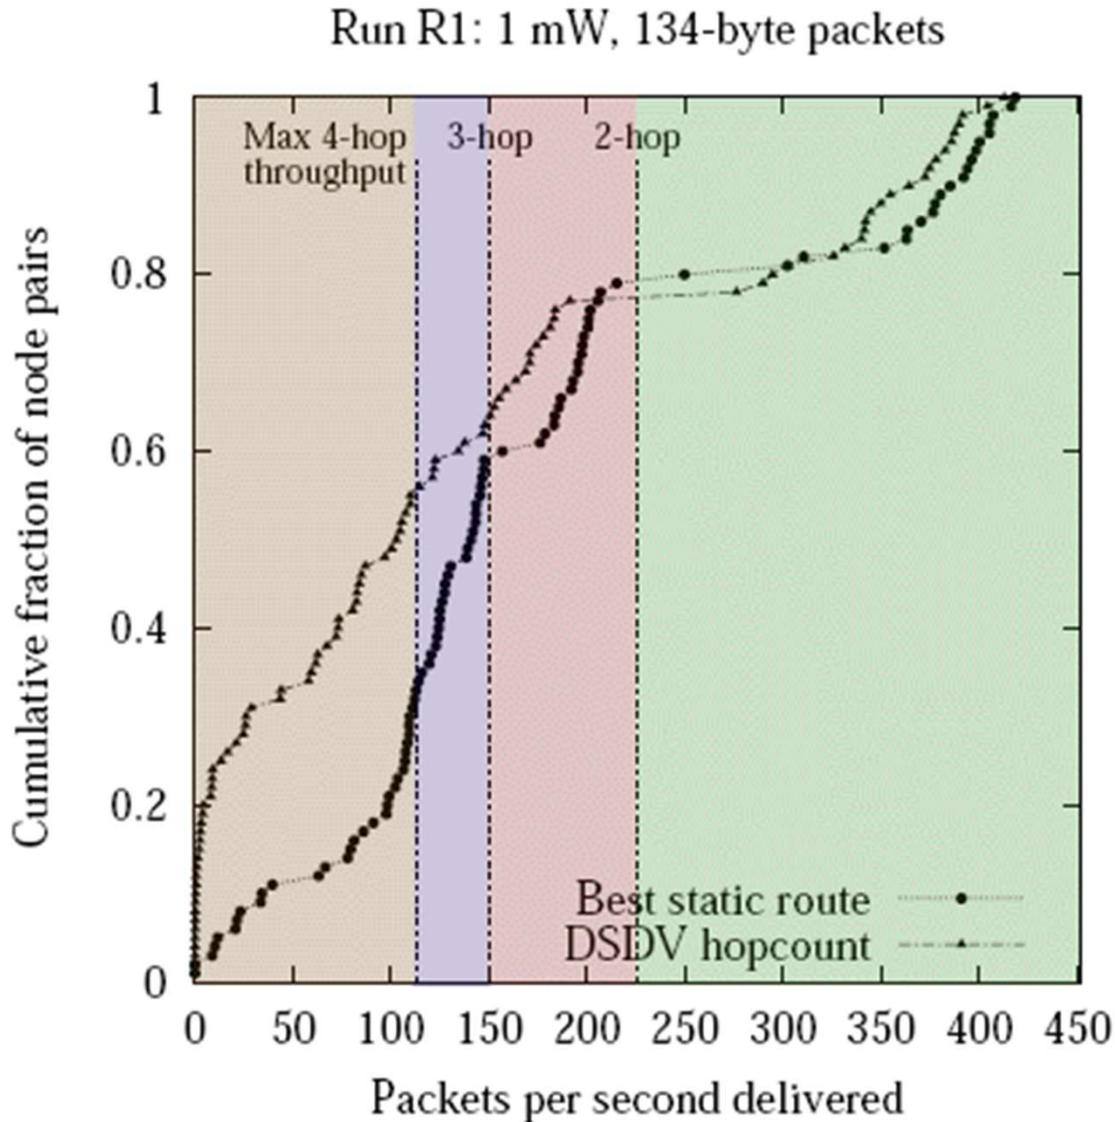D. S.J. Couto, D. Aguayo, J. Bicket, R. Morris

MIT

Hop Count Metric

Hop Count Metric

- Maximizes the distance traveled by each hop
  - Minimizes signal strength -> Maximizes the loss ratio
  - Uses a higher TxPower -> Interference
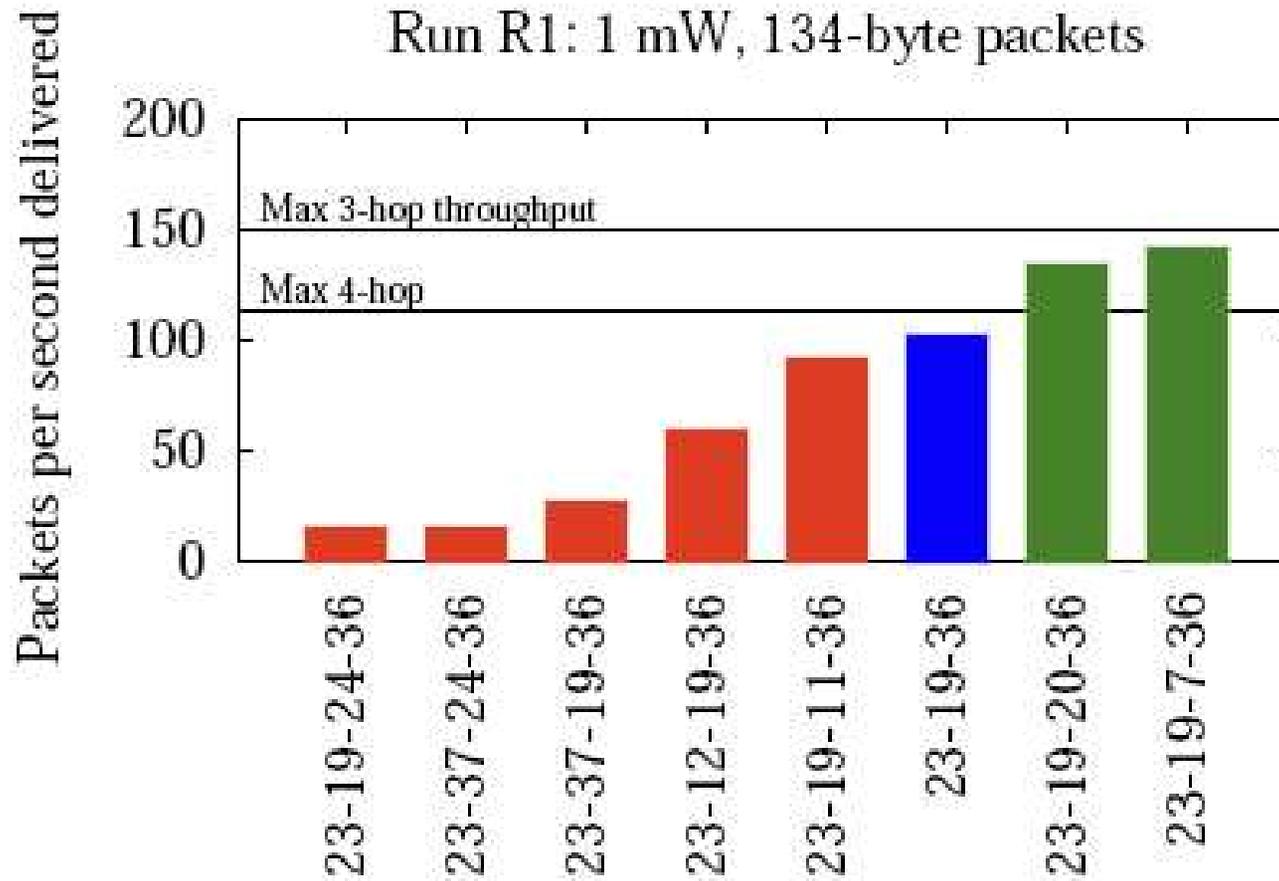- Possibly many shortest routes
- Avoid lossy links?
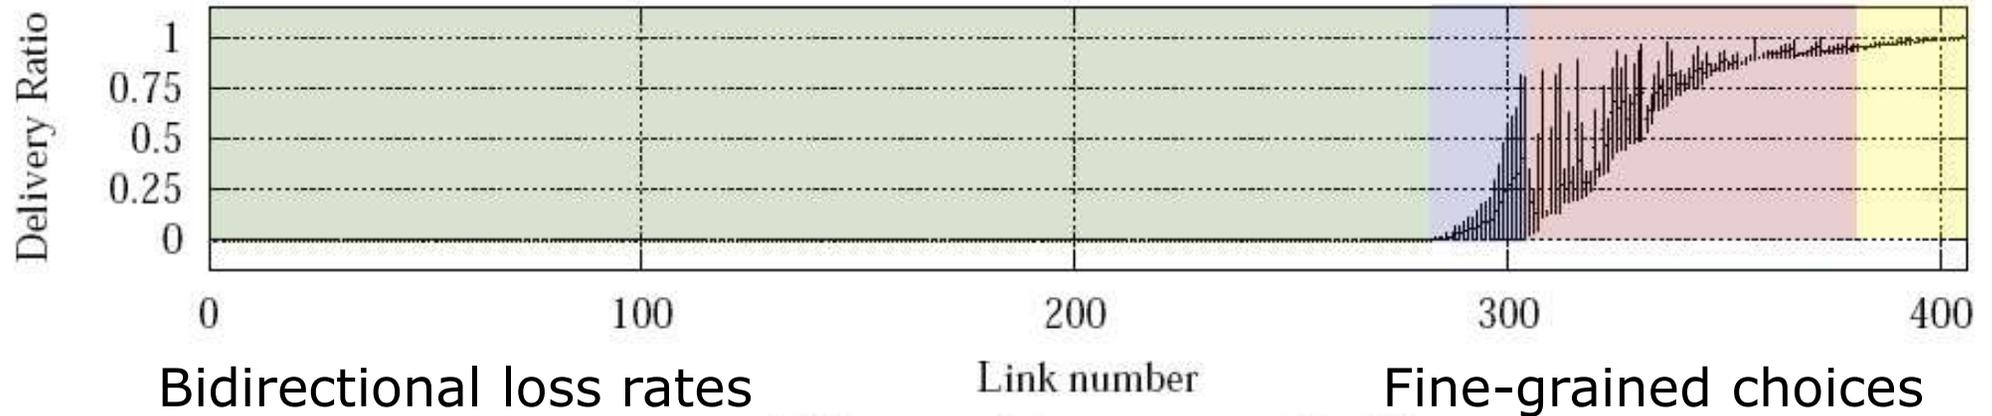
Hop Count vs. "Optimal"

x axis:

throughput

y axis:
fraction of
pairs with
less
throughput

Run R1: 1 mW, 134-byte packets

Max 4-hop throughput   3-hop   2-hop

Cumulative fraction of node pairs

Best static route
DSDV hopcount

Packets per second delivered

Hop Count Route Selection

Run R1: 1 mW, 134-byte packets

Motivation for a Better Routing Metric

(a) Pairwise delivery ratios at 1 mW

Bidirectional loss rates

Fine-grained choices

(b) Pairwise delivery ratios at 30 mW

Intermediate loss rates
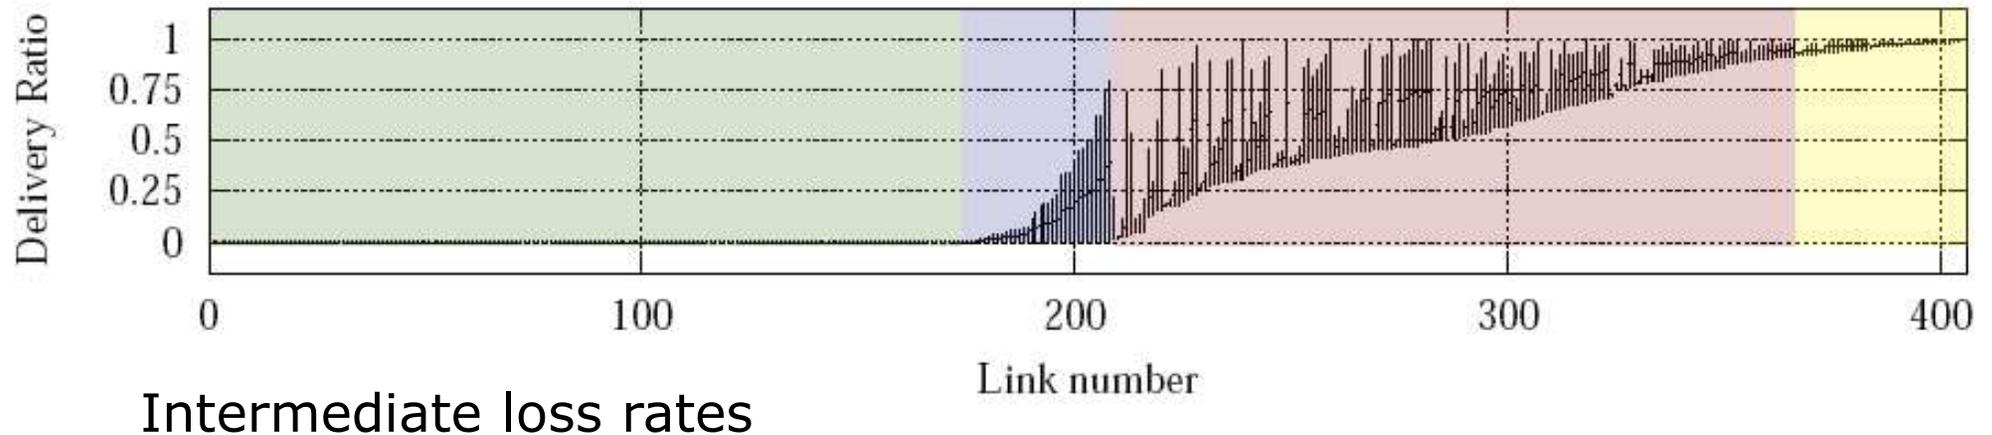
What metrics would you suggest to account for different link loss rates?

ETX

☐ The predicted number of data transmissions required to send a packet over a link

☐ The ETX of a path is the sum of the ETX values of the links over that path

☐ Examples:
  ○ ETX of a 3-hop route with perfect links is 3
  ○ ETX of a 1-hop route with 50% loss is 2

ETX continued...

□ Expected probability that a transmission is successfully received and acknowledged is $d_f \times d_r$

  ○ d_f is forward delivery ratio

  ○ d_r is reverse delivery ratio

□ Each attempt to transmit a packet is a Bernoulli trial, so...

$$\text{ETX} = \frac{1}{d_f \times d_r}$$

Hooray for ETX!

- Based on delivery ratios, which affect throughput
- Detects and handles asymmetry by incorporating loss ratios in each direction
- Uses precise link loss measurements to make fine-grained decisions between routes
  - Assumes you can measure these ratios precisely
- Penalizes routes with more hops, which have lower throughput due to inter-hop interference
  - Assumes loss rates are equal over links
- Tends to minimize spectrum use, which should maximize overall system capacity (reduce power too)
  - Each node spends less time retransmitting data

Acquiring ETX values

- Measured by broadcasting dedicated link probe packets with an average period τ (jittered by ±0.1τ)
- Delivery ratio: $r(t) = \dfrac{\text{count}(t - w, t)}{w/\tau}$

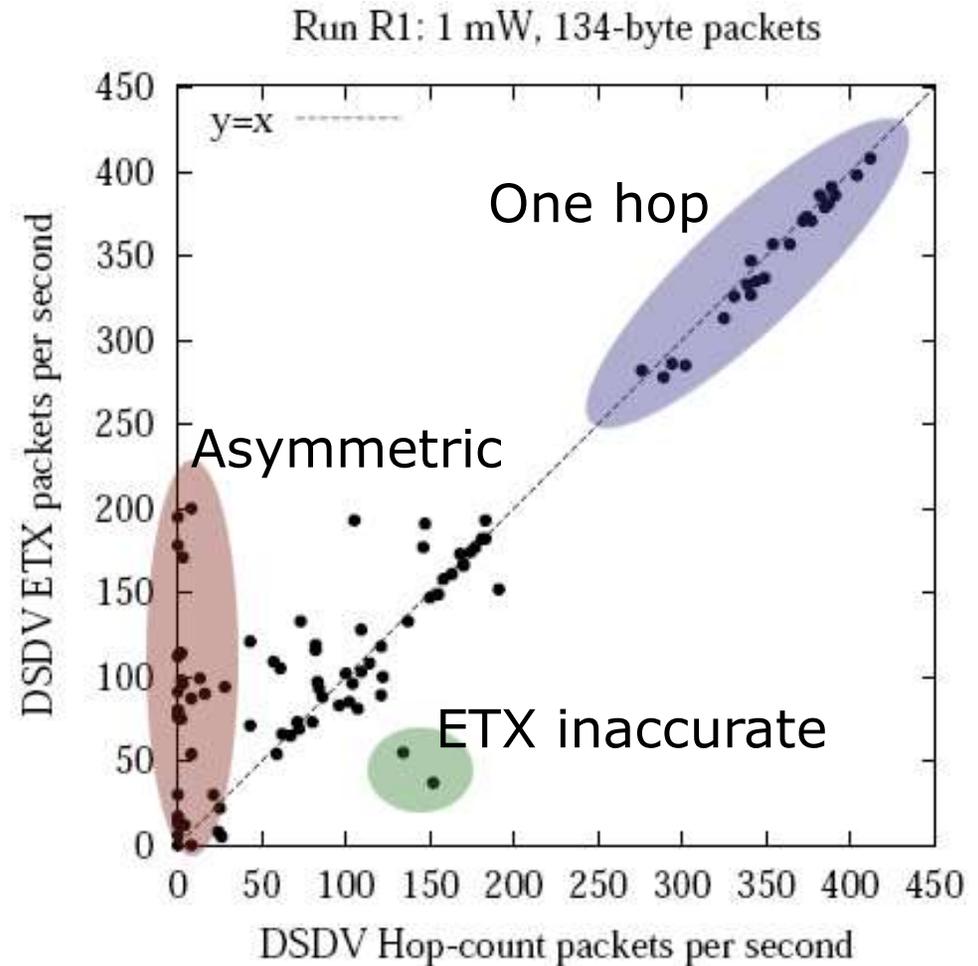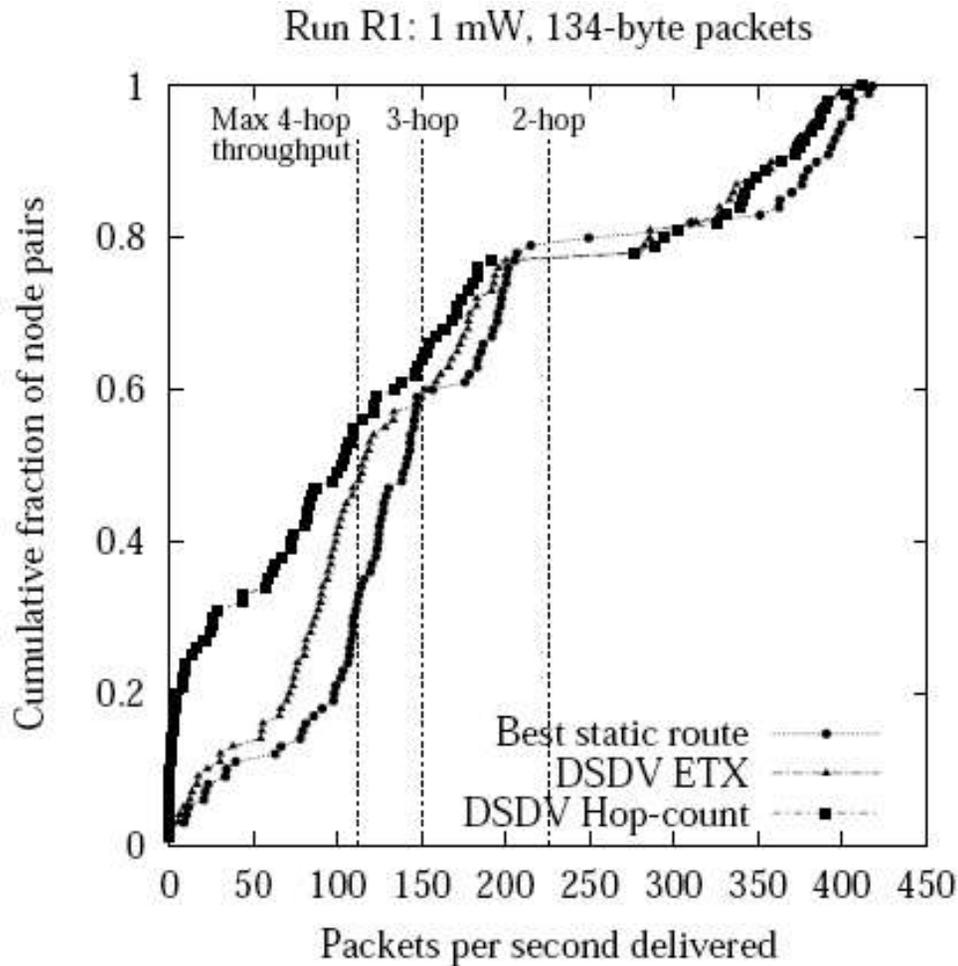 - $count(t\text{-}w,t)$ is the # of probes received during window w
  - w/τ is the # of probes that should have been received
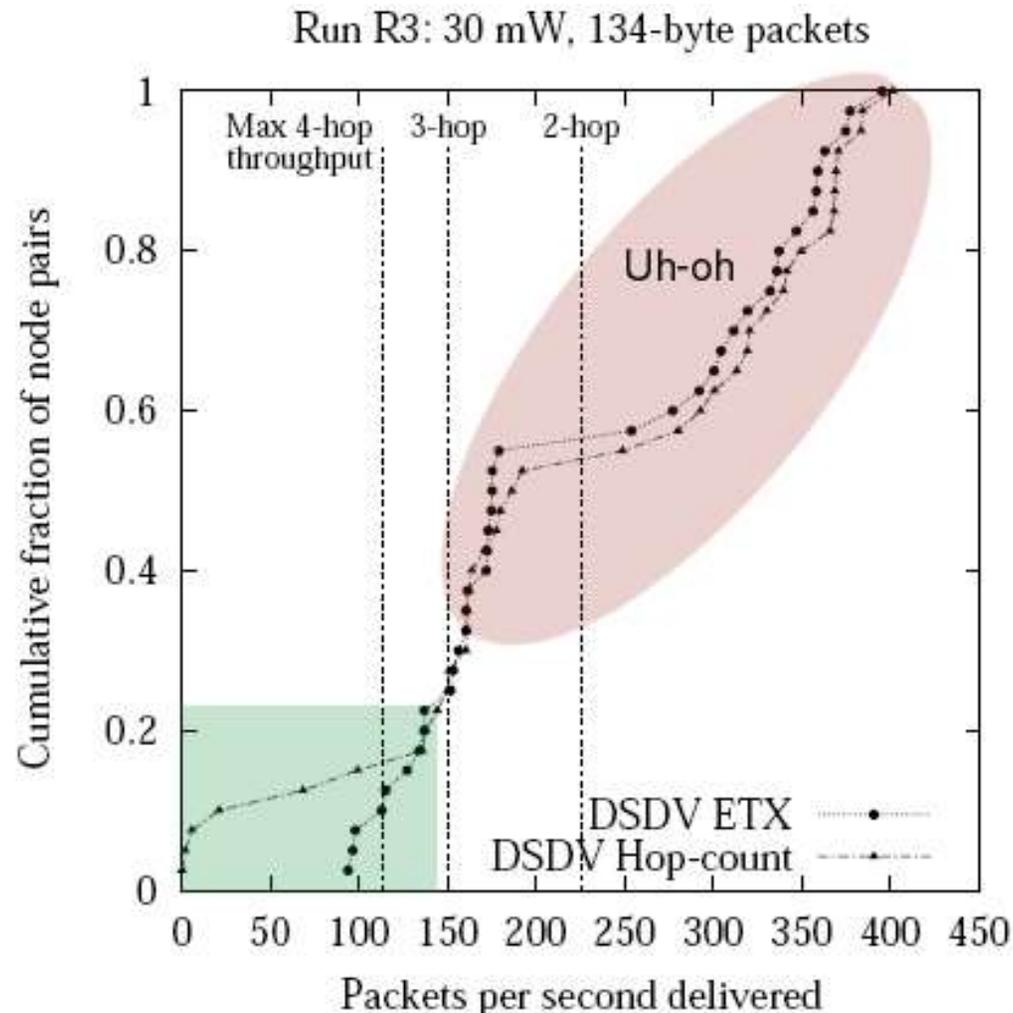- Each probe contains this information

Implementation and such…

- Modified DSDV and DSR
- τ = 1 packet per second, w = 10 sec
- Multiple queues (different priorities)
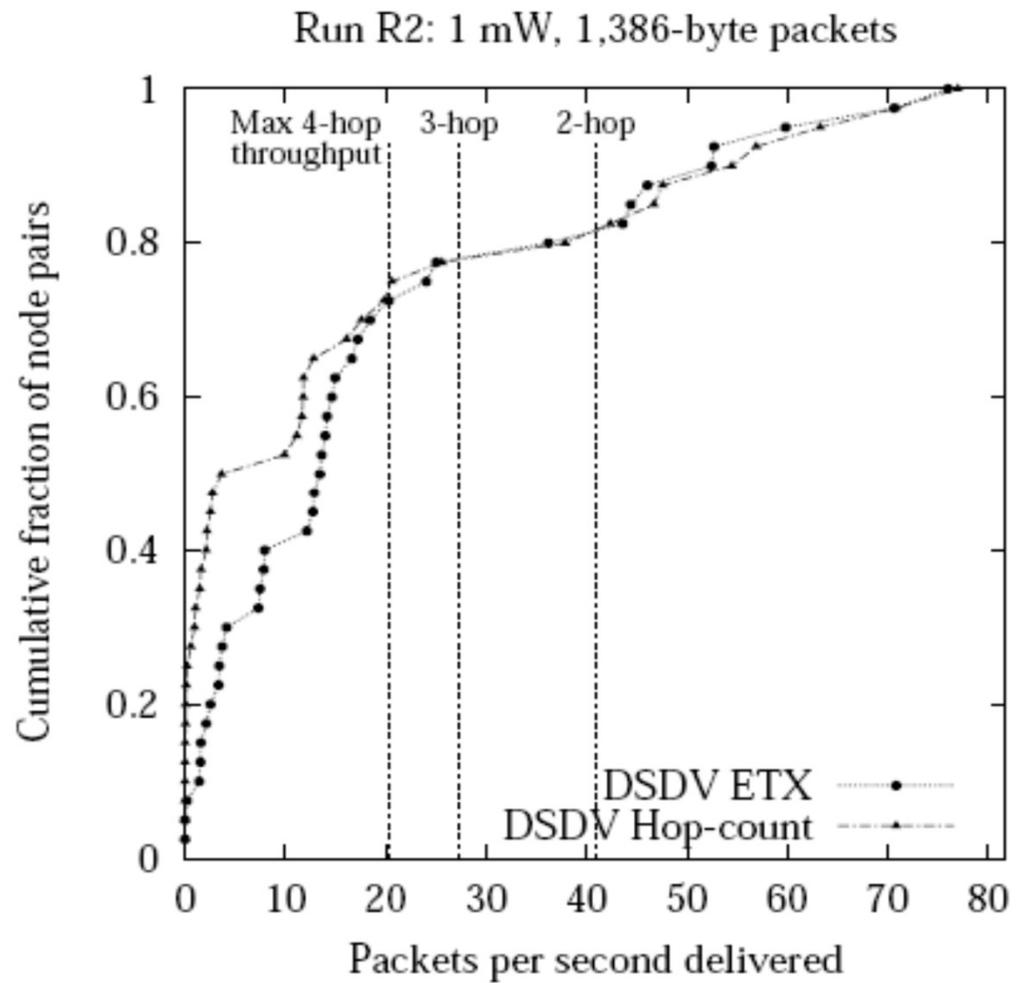 - Loss-ratio probes, protocol packets, data packets

DSDV Performance

Run R1: 1 mW, 134-byte packets

Max 4-hop throughput   3-hop   2-hop

Cumulative fraction of node pairs

Packets per second delivered

Best static route
DSDV ETX
DSDV Hop-count

Run R1: 1 mW, 134-byte packets

y=x

One hop

Asymmetric

ETX inaccurate

DSDV ETX packets per second

DSDV Hop-count packets per second

DSDV and High Transmit Power

Run R3: 30 mW, 134-byte packets

□ Less throughput advantage than when probe packets are smaller (134 bytes)

□ Why?

Run R2: 1 mW, 1,386-byte packets

Max 4-hop throughput   3-hop   2-hop

DSDV ETX
DSDV Hop-count

Cumulative fraction of node pairs

Packets per second delivered

Packet sizes continued

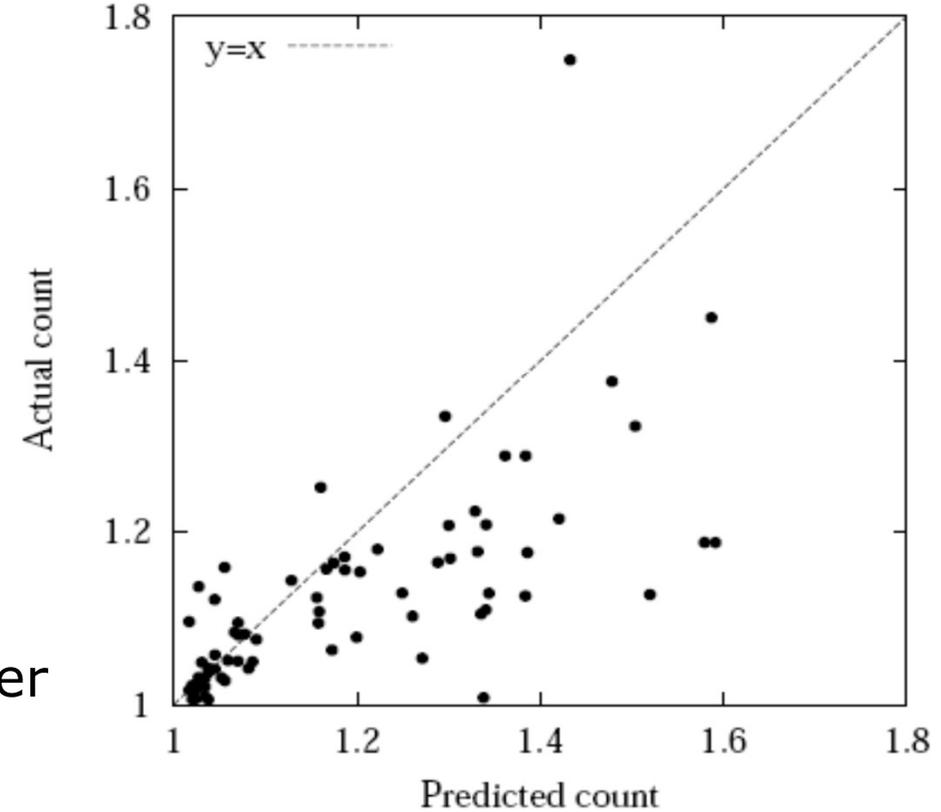Run R4: 30 mW, various packet sizes

ACK's smaller than probe packets

Run R5: 30 mW, 104-byte packets
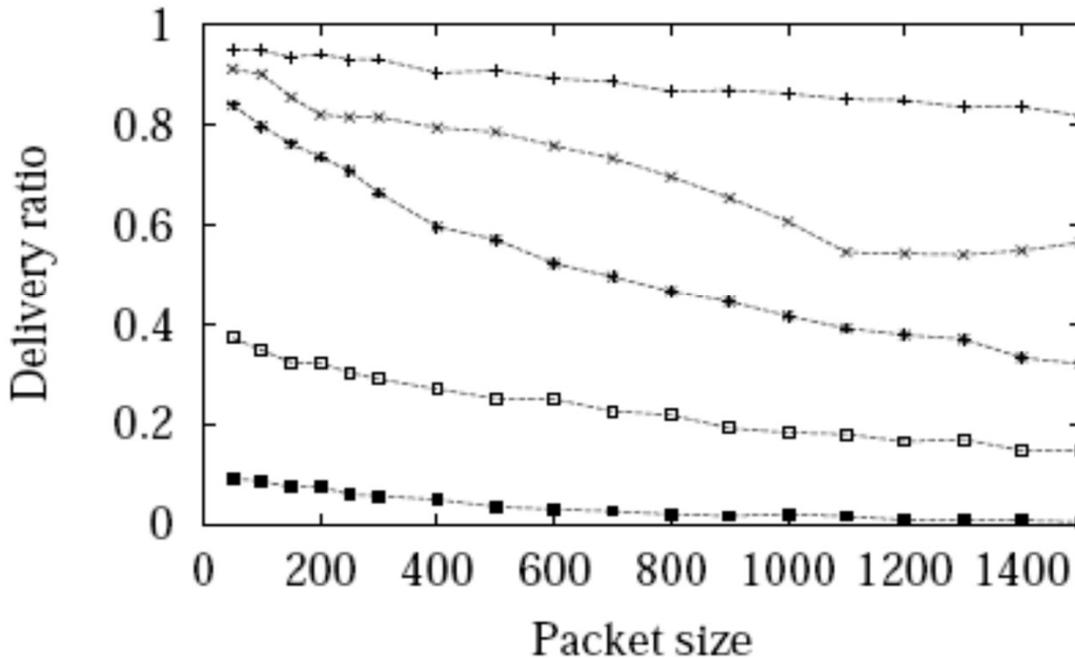
ETX underestimates ACK delivery ratios -> overestimates total number of transmissions per packet

Pros and cons?

Conclusions

□ Pros
  ○ ETX performs better or comparable to Hop Count Metric
    • Accounts for bi-directional loss rates
 ○ Can easily be incorporated into routing protocols

□ Cons
  ○ Only considers link loss rates
  ○ May not be best metric for all networks
    • Mobility
    • Power-limited
    • Adaptive rate (multi-rate)
    • Interference
  ○ Predictions of loss ratios not always accurate and incur overhead
  ○ Does not explicitly incorporate the interaction between routing changes and ETX change ➜ oscillation and select sub-optimal paths

Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks
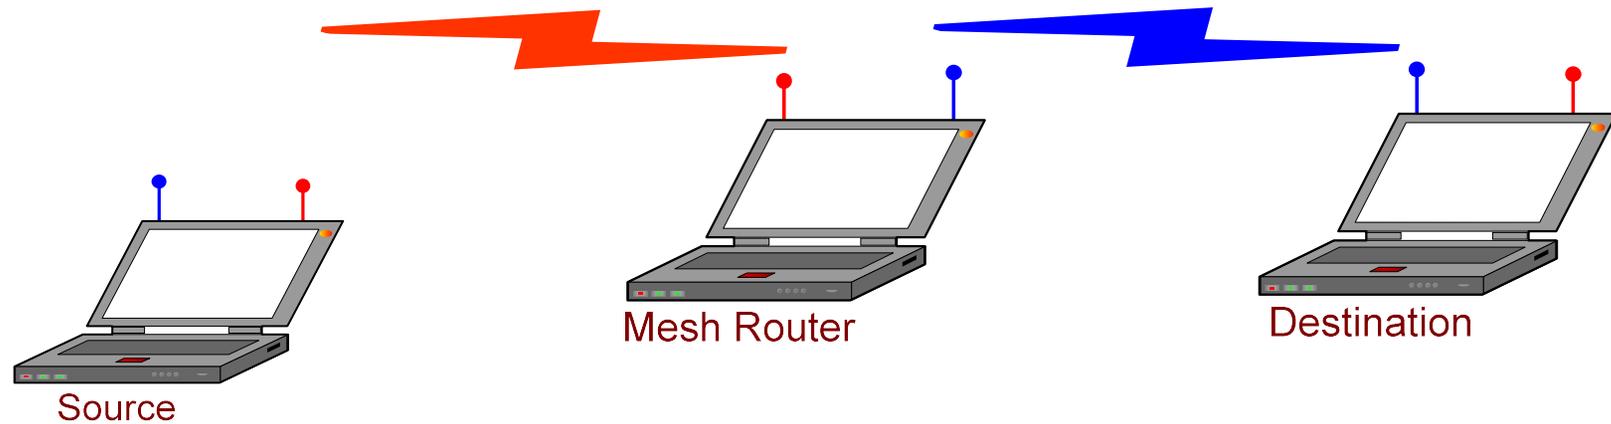
Richard Draves, Jitendra Padhye, and Brian Zill

Microsoft Research

Multi-Hop Networks with Single Radio

Source   Mesh Router   Destination
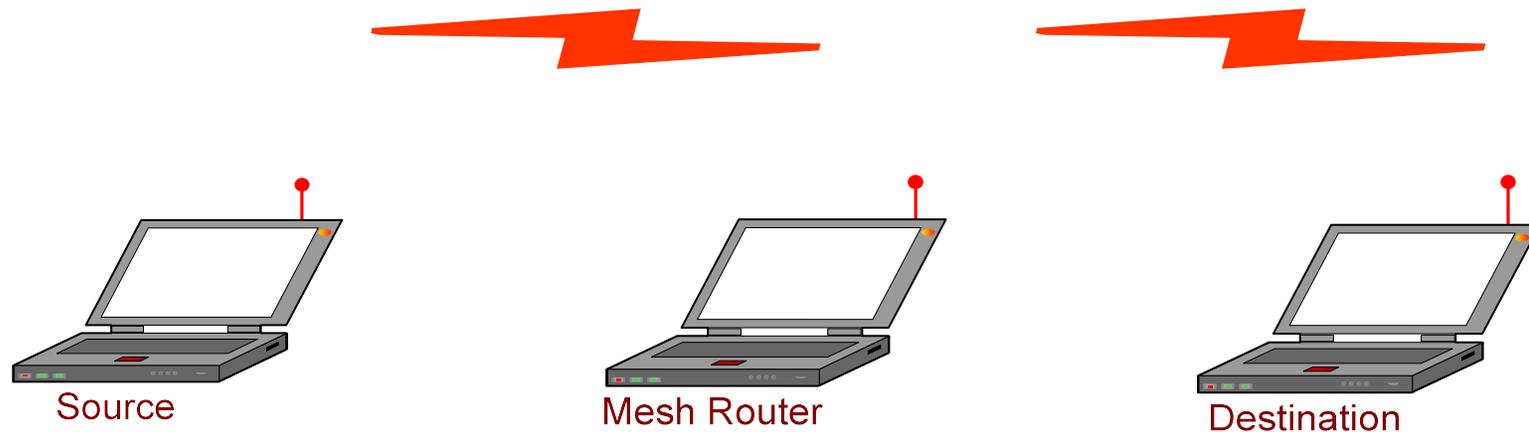
With a single radio, a node can not transmit and receive simultaneously.

Multi-Hop Networks with Multiple Radios

Source

**Mesh Router**

Destination

With two radios tuned to non-interfering channels, a node can transmit and receive simultaneously.

Other Advantages of Multiple Radios

☐ Increased robustness due to frequency diversity

  ○ e.g. 2.4GHz (802.11b) and 5GHz (802.11a) have different fading characteristics

☐ Possible tradeoff between range and data rate

  ○ Can be helpful during early deployment

Existing Routing Metrics are Inadequate

2 Mbps

12 Mbps

12 Mbps

Source

Mesh Router

Destination
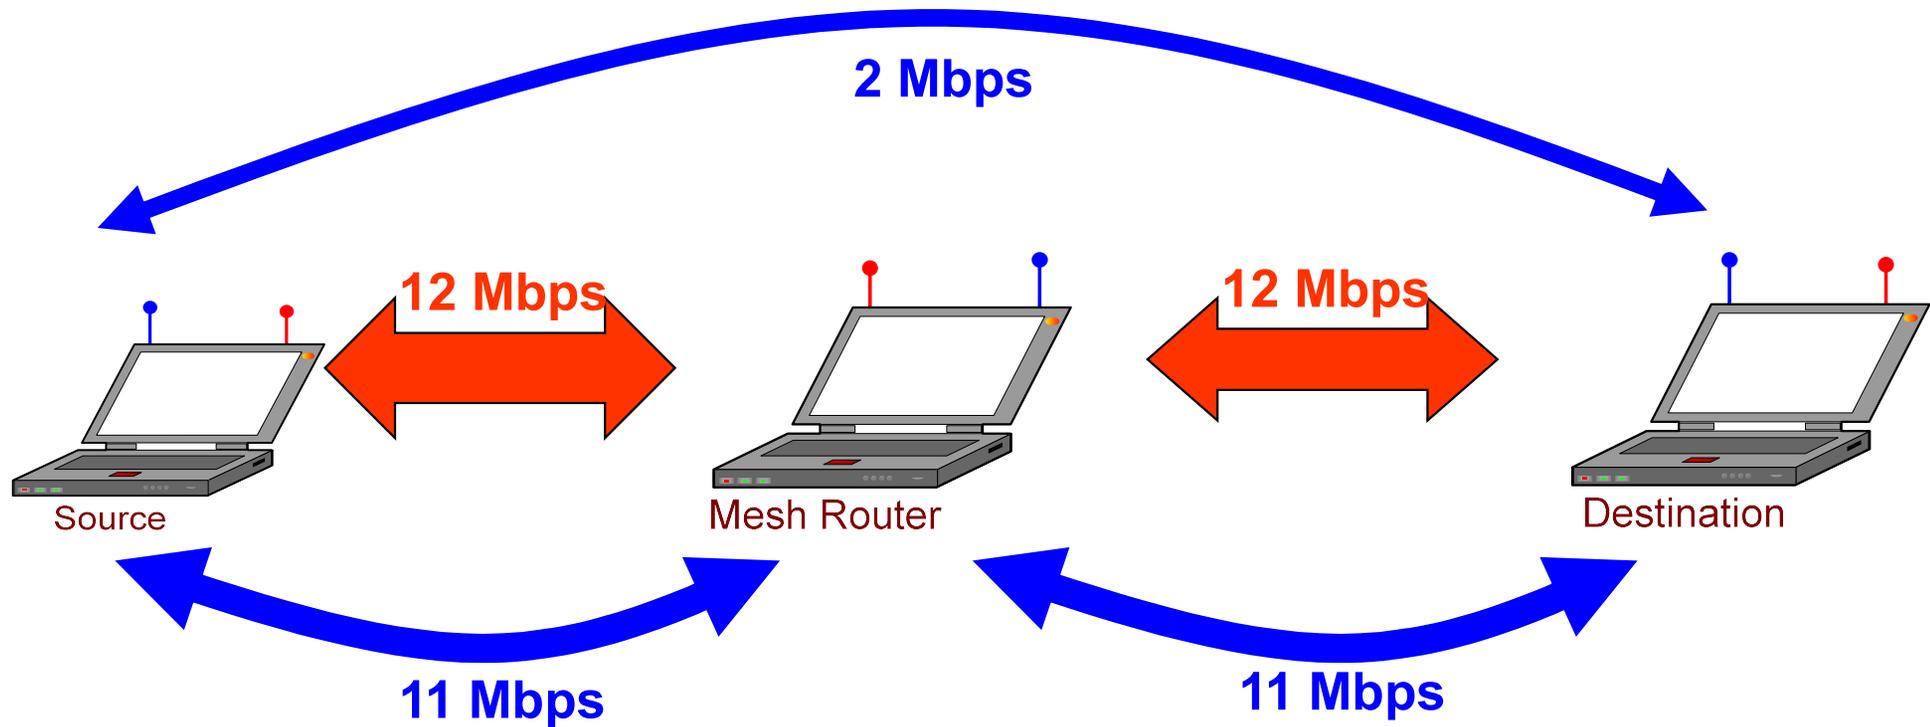
11 Mbps

11 Mbps

Shortest path: 2 Mbps

Path with fastest links: 6 Mbps

Best path: 11 Mbps

Contributions

☐ New routing metric for multi-radio mesh networks
  ○ Weighted Cumulative Expected Transmission Time (WCETT)

☐ Implementation of the metric in a link-state routing protocol
  ○ Multi-Radio Link-Quality source routing (MR-LQSR)

☐ Experimental evaluation of WCETT:
  ○ 24-node, multi-radio mesh testbed
  ○ 2 radios per node, 11a and 11g
  ○ Side-by-side comparison with:
    • Shortest path (HOP)
    • ETX (De Couto et. al. MOBICOM 2003)

Summary of Results

- WCETT makes judicious use of two radios
  - Over 250% better than HOP
  - Over 80% better than ETX

- Gains more prominent over shorter paths and in lightly-loaded scenarios

Outline of the talk

☐ Design of WCETT

☐ Experimental results

☐ Conclusion

Design of Routing Metric: Assumptions

- No power constraints

- Little or no node mobility
  - Relatively stable links

- Nodes have one or more 802.11 radios

- Multiple radios on a node are tuned to non-interfering channels
  - Channel assignment is fixed

Implementation Framework

- Implemented in a source-routed, link-state protocol
  - Multi-Radio Link Quality Source Routing (MR-LQSR)

- Nodes discovers links to its neighbors

- Measure quality of those links

- Link information floods through the network
  - Each node has "full knowledge" of the topology

- Sender selects "best path"
  - Packets are source routed using this path

Goal for Multi-Radio Routing Metric

Goal for Multi-Radio Routing Metric

Maximize throughput of a given flow:

- Prefer high-bandwidth, low-loss links

- When possible, select channel diverse paths

- Prefer shorter paths

Components of a Routing Metric

☐ Link Metric: Assign a weight to each link

WCETT: Prefer high-bandwidth, low-loss links

☐ Path Metric: Combine metrics of links on path

WCETT: Prefer short, channel-diverse paths

Is ETX a good enough link metric?

Link Metric: Expected Transmission Time (ETT)

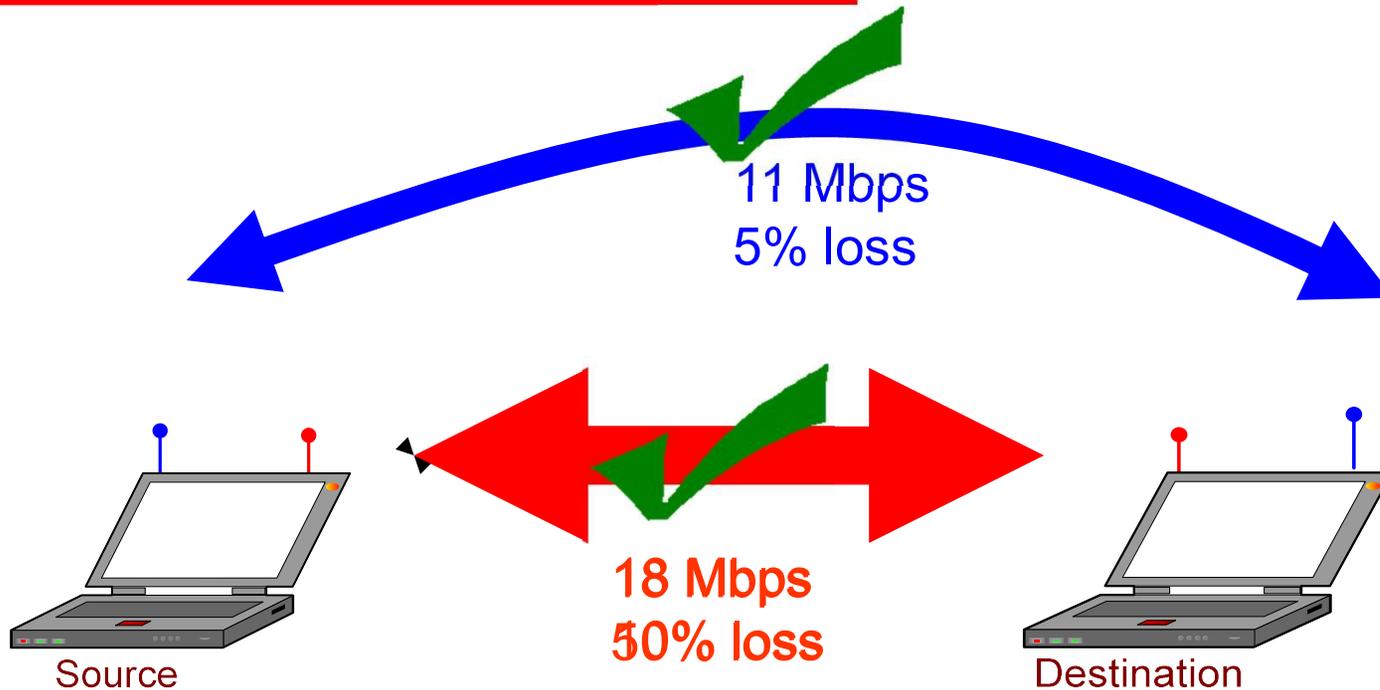- Link loss rate = p
  - Expected number of transmissions

$$ETX = \frac{1}{1-p}$$

- Packet size = S, Link bandwidth = B
  - Each transmission lasts for S/B

$$ETT = \left(\frac{S}{B}\right) * ETX$$

- Lower ETT implies better link

ETT: Illustration

11 Mbps
5% loss

18 Mbps
50% loss

Source

Destination

1000 Byte Packet

ETT : 0.77 ms

ETT : 0.89 ms

Any limitations with ETT as link metric?

How to design ETT-based path metric?

Combining Link Metric into Path Metric Proposal 1

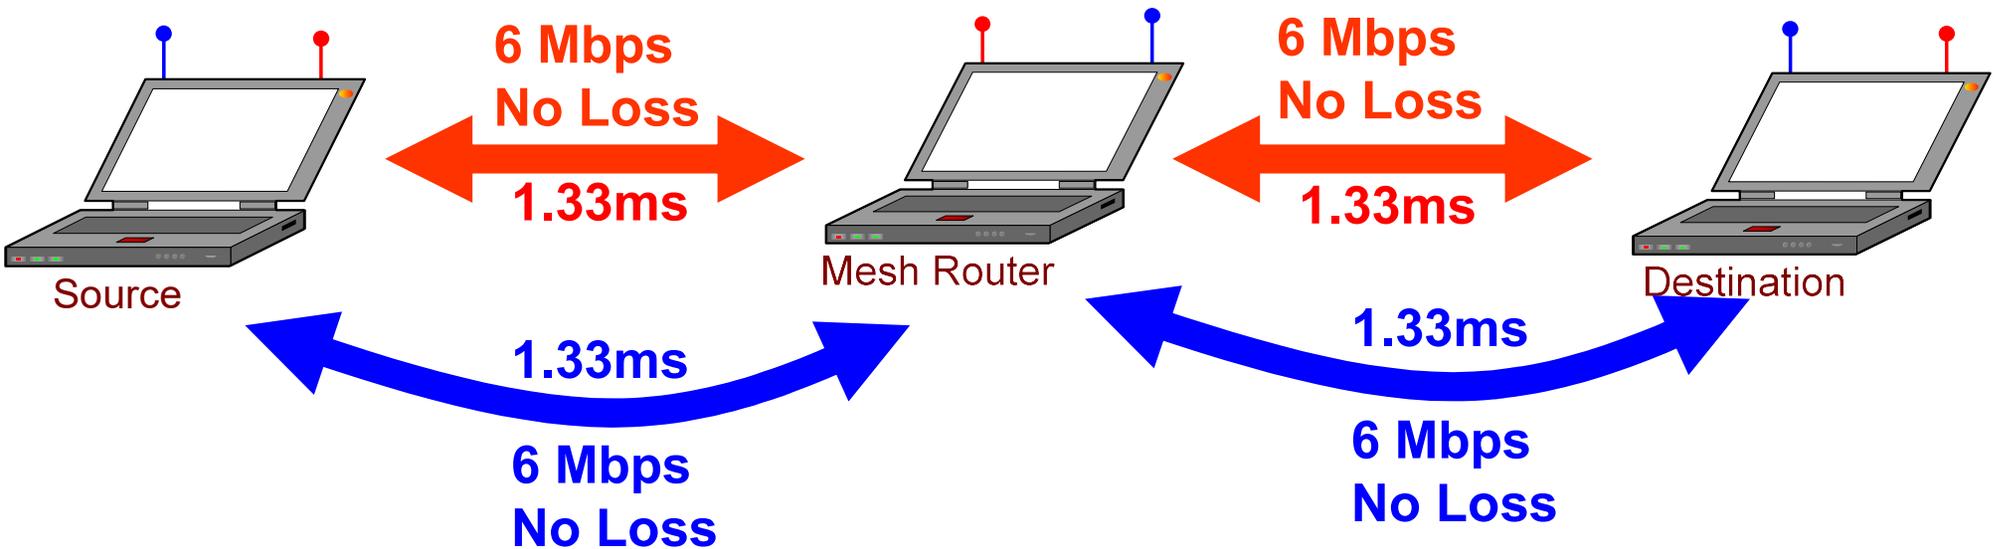❏ Add ETTs of all links on the path

❏ Use the sum as path metric

SETT = Sum of ETTs of links on path

(Lower SETT implies better path)

Pro: Favors short paths

Con: Does not favor channel diversity

SETT does not favor channel diversity

6 Mbps No Loss — 1.33ms

6 Mbps No Loss — 1.33ms

Source   Mesh Router   Destination

1.33ms

6 Mbps No Loss

1.33ms

6 Mbps No Loss

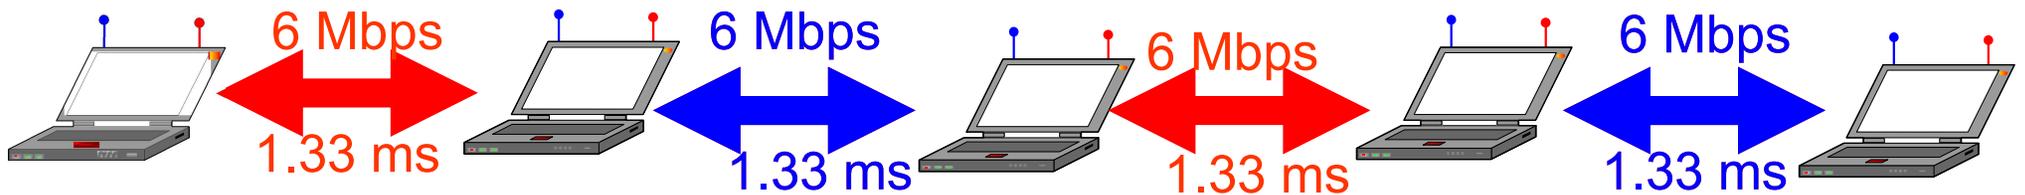| Path	Throughput	SETT
Red-Blue	6 Mbps	2.66 ms
Red-Red	3 Mbps	2.66 ms

Impact of Interference

- Interference reduces throughput

- Throughput of a path is lower if many links are on the same channel
  - Path metric should be worse for non-diverse paths

- Assumption: All links that are on the same channel interfere with one another
  - Pessimistic for long paths

Combining Link Metric into Path Metric Proposal 2

- ❑ Group links on a path according to channel
  - ○ Links on same channel interfere

- ❑ Add ETTs of links in each group

- ❑ Find the group with largest sum.
  - ○ This is the "bottleneck" group
  - ○ Too many links, or links with high ETT ("poor quality" links)

- ❑ Use this largest sum as the path metric
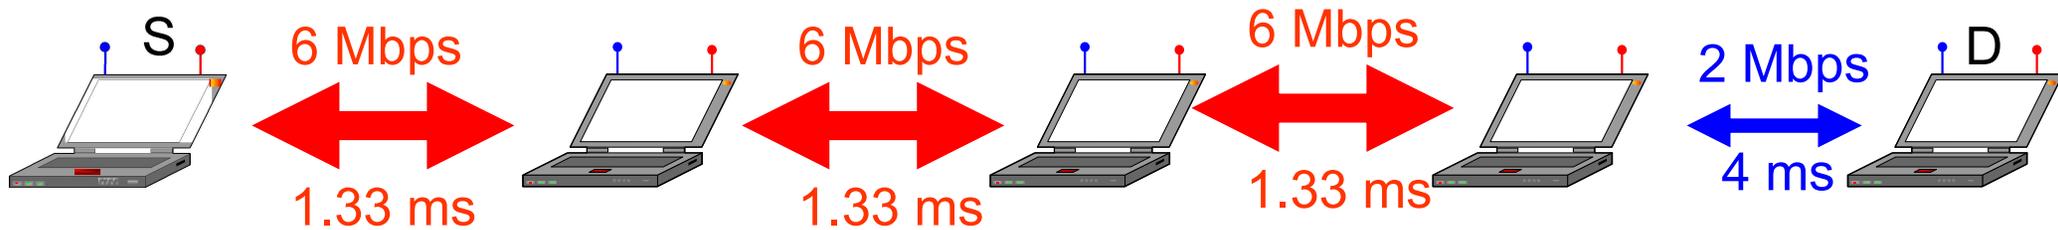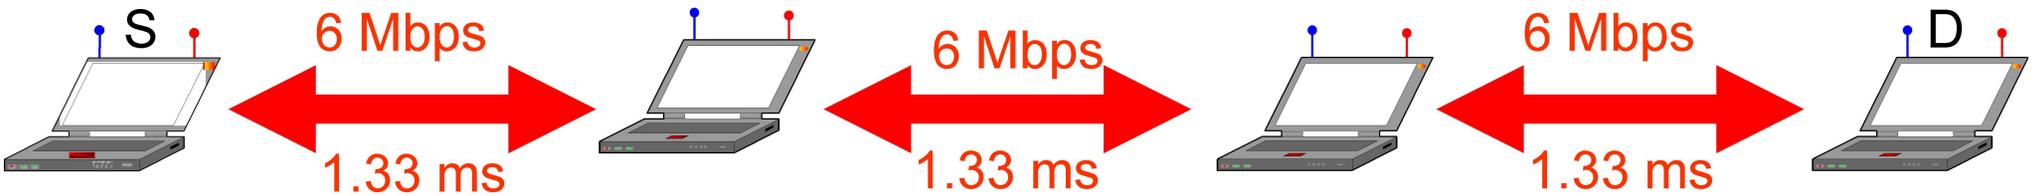 - ○ Lower value implies better path

"Bottleneck Group ETT" (BG-ETT)

BG-ETT Example

Path	Throughput	Blue Sum	Red Sum	BG-ETT
All Red	1.5 Mbps	0	5.33 ms	5.33 ms
1 Blue	2 Mbps	1.33 ms	4 ms	4 ms
Red-Blue	3 Mbps	2.66 ms	2.66 ms	2.66 ms

BG-ETT favors high-throughput, channel-diverse paths.

BG-ETT does not favor short paths

Path	Throughput	Blue Sum	Red Sum	BG-ETT
3-Hop	2 Mbps	0	4 ms	4 ms
4-Hop	2 Mbps	4 ms	4 ms	4 ms

Path Metric: Putting it all together

☐ SETT favors short paths

☐ BG-ETT favors channel diverse paths

Weighted Cumulative ETT (WCETT)

$$WCETT = (1-\beta) * SETT + \beta * BG\text{-}ETT$$

β is a tunable parameter

Higher value: More preference to channel diversity
Lower value: More preference to shorter paths

Any limitation on WCETT as path metric?

How to measure loss rate and bandwidth?
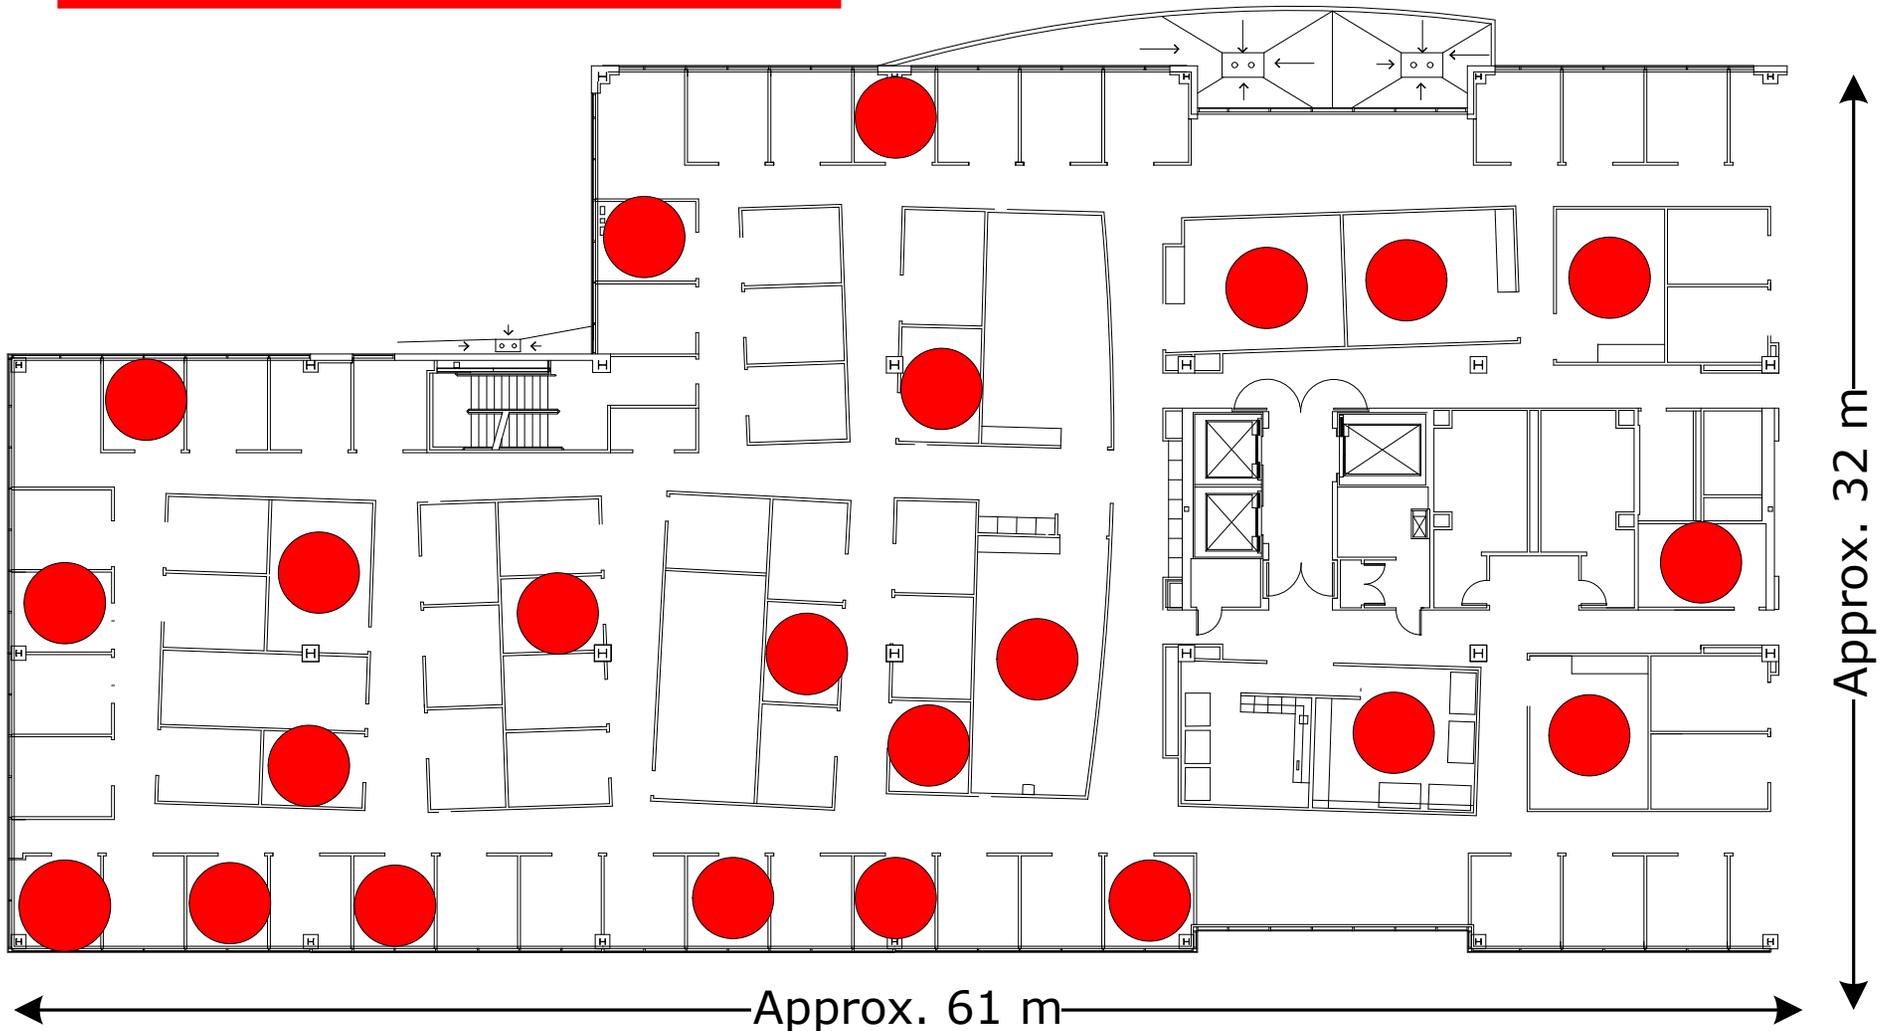
□ Loss rate measured using broadcast probes
  ○ Similar to ETX
  ○ Updated every second

□ Bandwidth estimated using periodic packet-pairs
  ○ Updated every 5 minutes

Outline of the talk

- ☐ Design of WCETT

- ☐ Experimental results

- ☐ Conclusion

Mesh Testbed

Approx. 32 m

Approx. 61 m

23 nodes running Windows XP.
Two 802.11a/b/g cards per node: Proxim and NetGear (Autorate)
Diameter: 6-7 hops.

Experiment Setting

❒ 2-Minute TCP transfer between 100 randomly selected node pairs (Out of 23x22 = 506)

❒ Only one transfer active at a time

❒ Performance metric:

   ○ **Median throughput of 100 transfers**
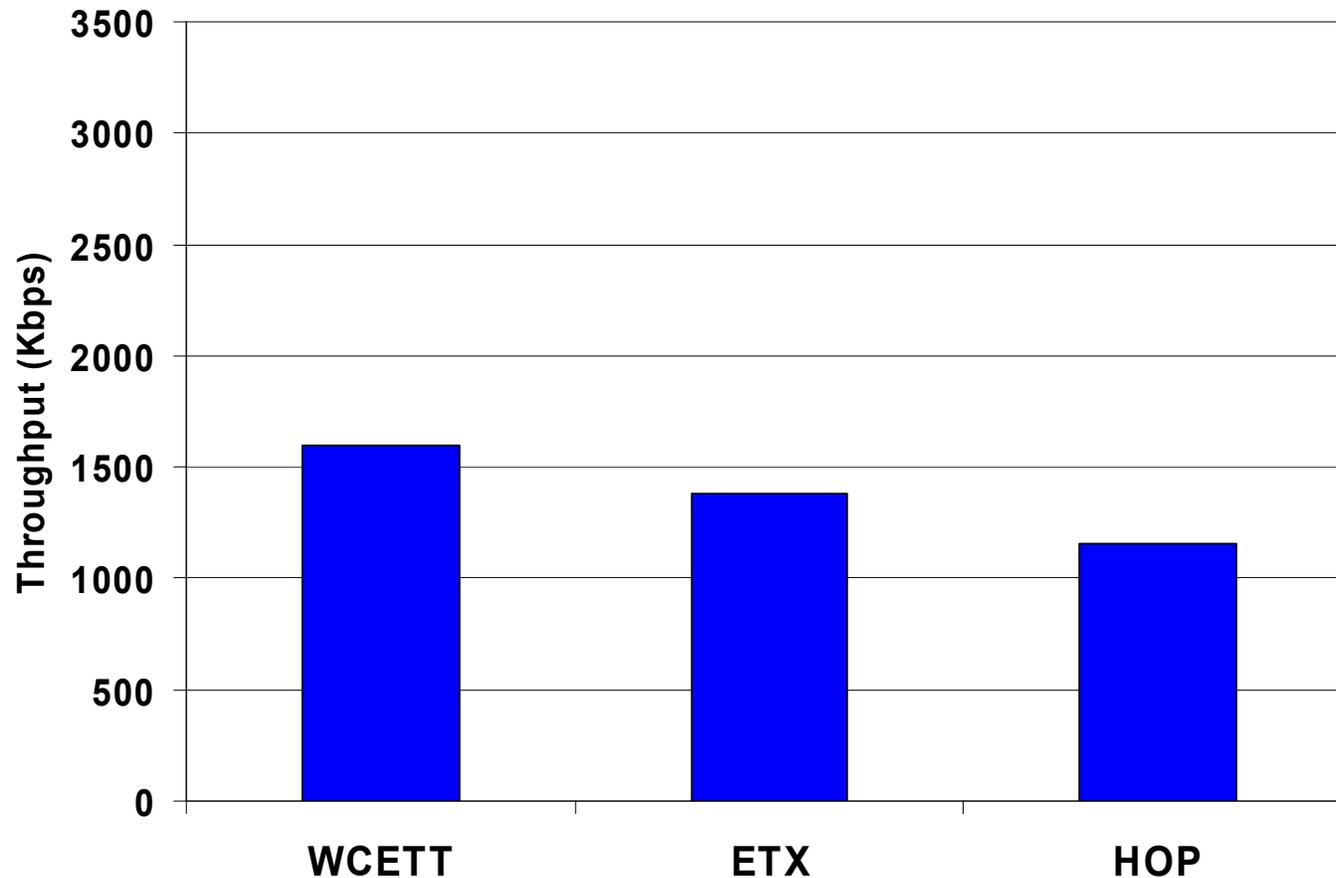
Baseline (Single Radio)
NetGear on 802.11a (Channel 36), Proxim OFF

Two Radio
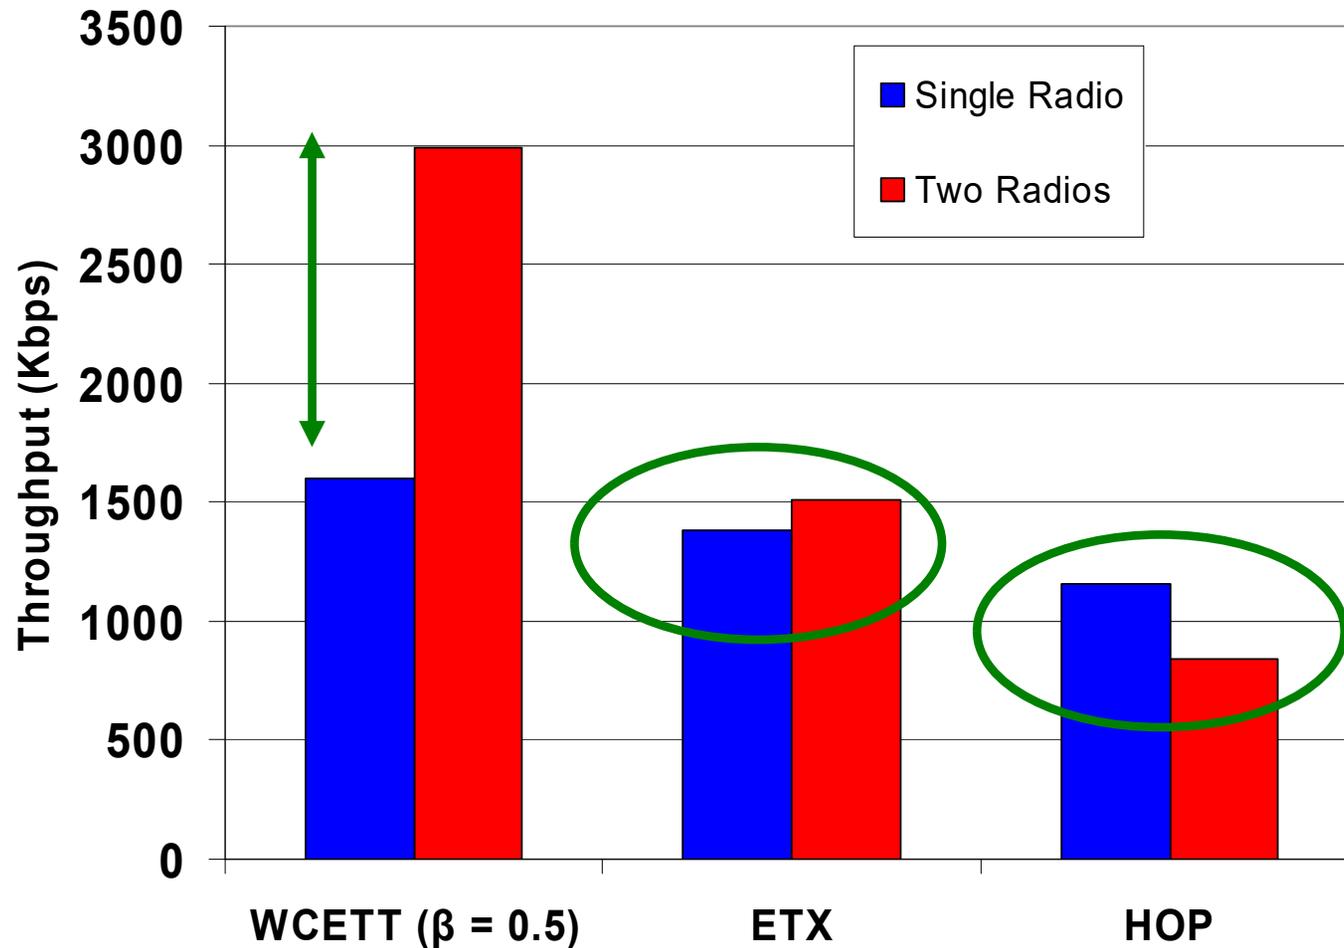**NetGear on 802.11a (Chan 36), Proxim on 802.11g (Chan 10)**

(802.11g radios have longer range, lower bandwidth)

Median Throughput (Baseline, single radio)

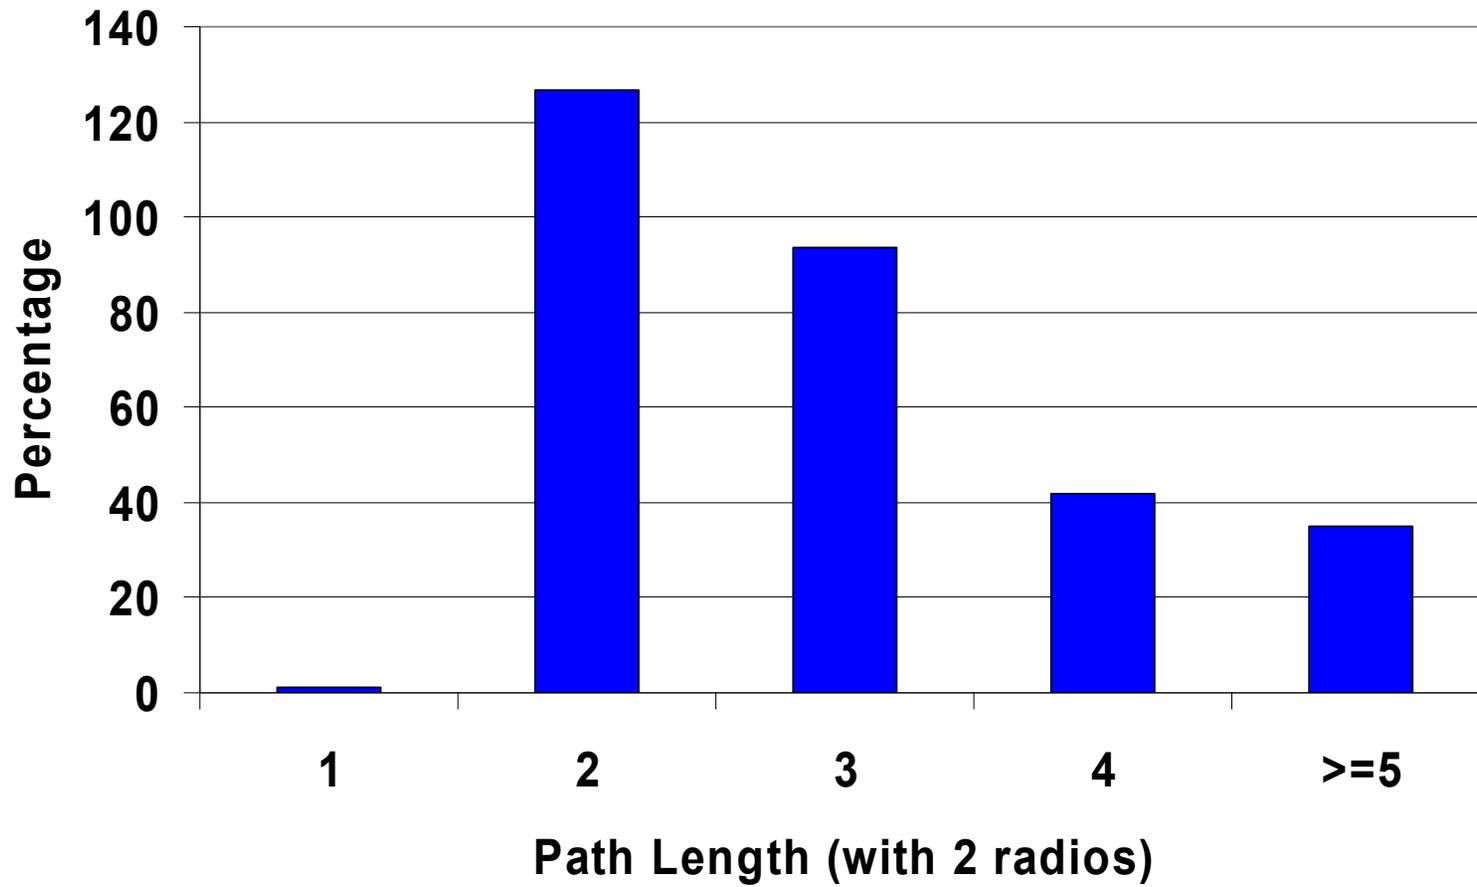**WCETT provides performance gain even with one radio.**

Median Throughput (Two radios)

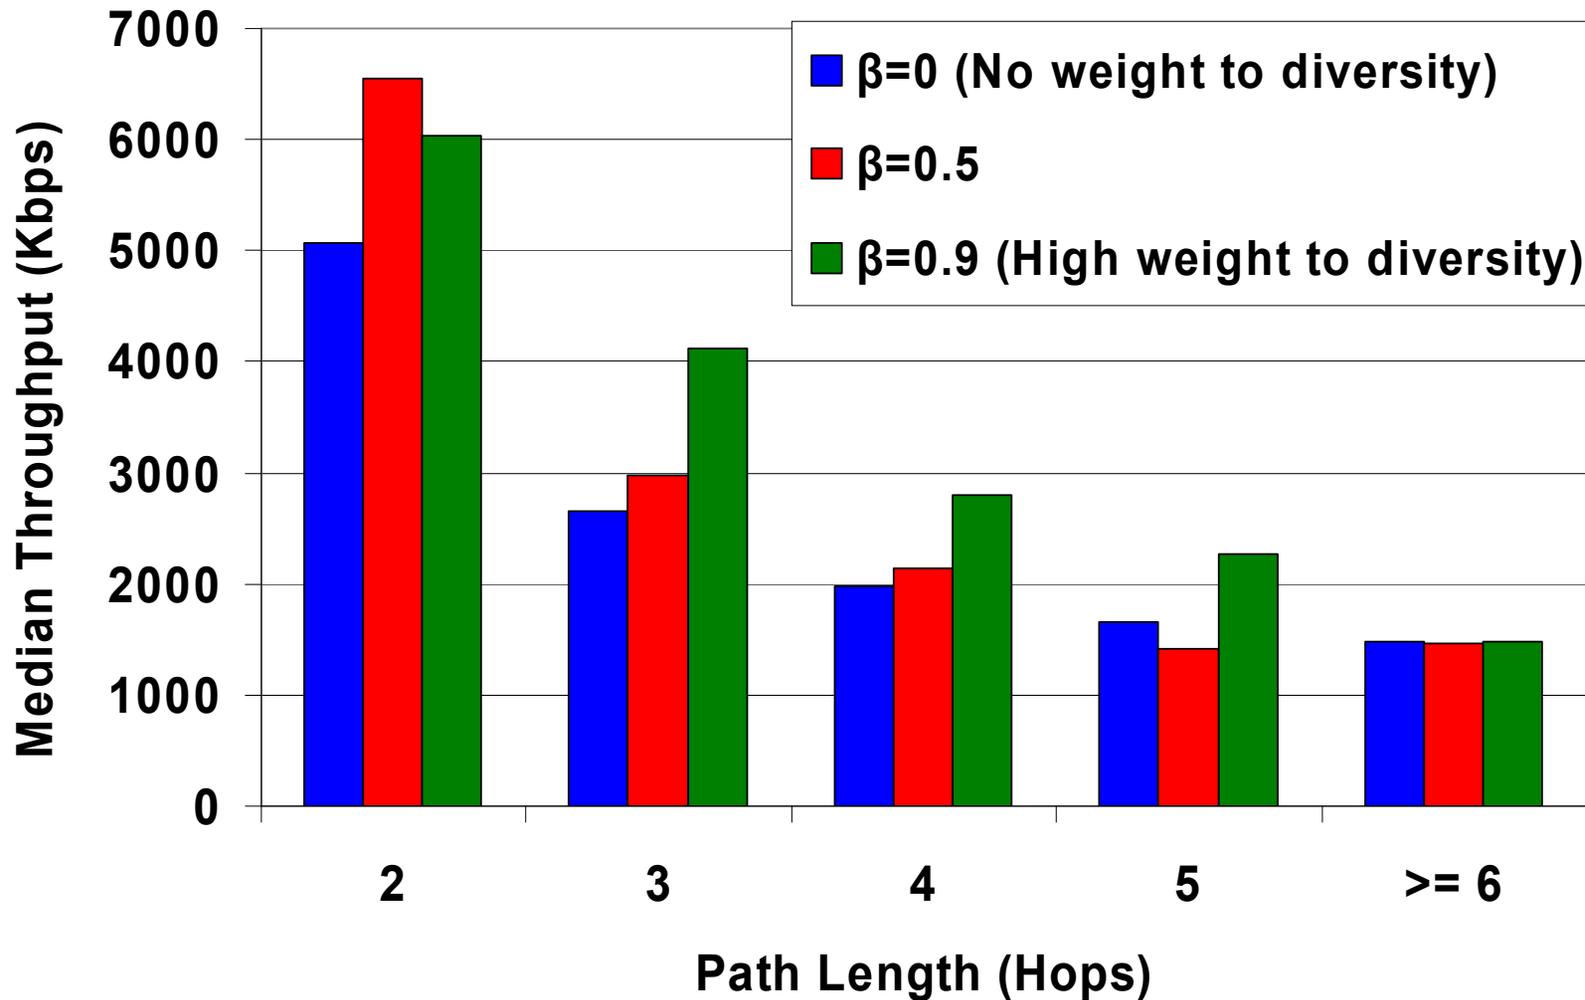WCETT n Performance of HOP worsens with 2nd radio! r baseline

Do all paths benefit equally with WCETT?

Improvement in Median Throughput over Baseline (1 radio)

Path Length (with 2 radios)

WCETT gains are more prominent for shorter paths
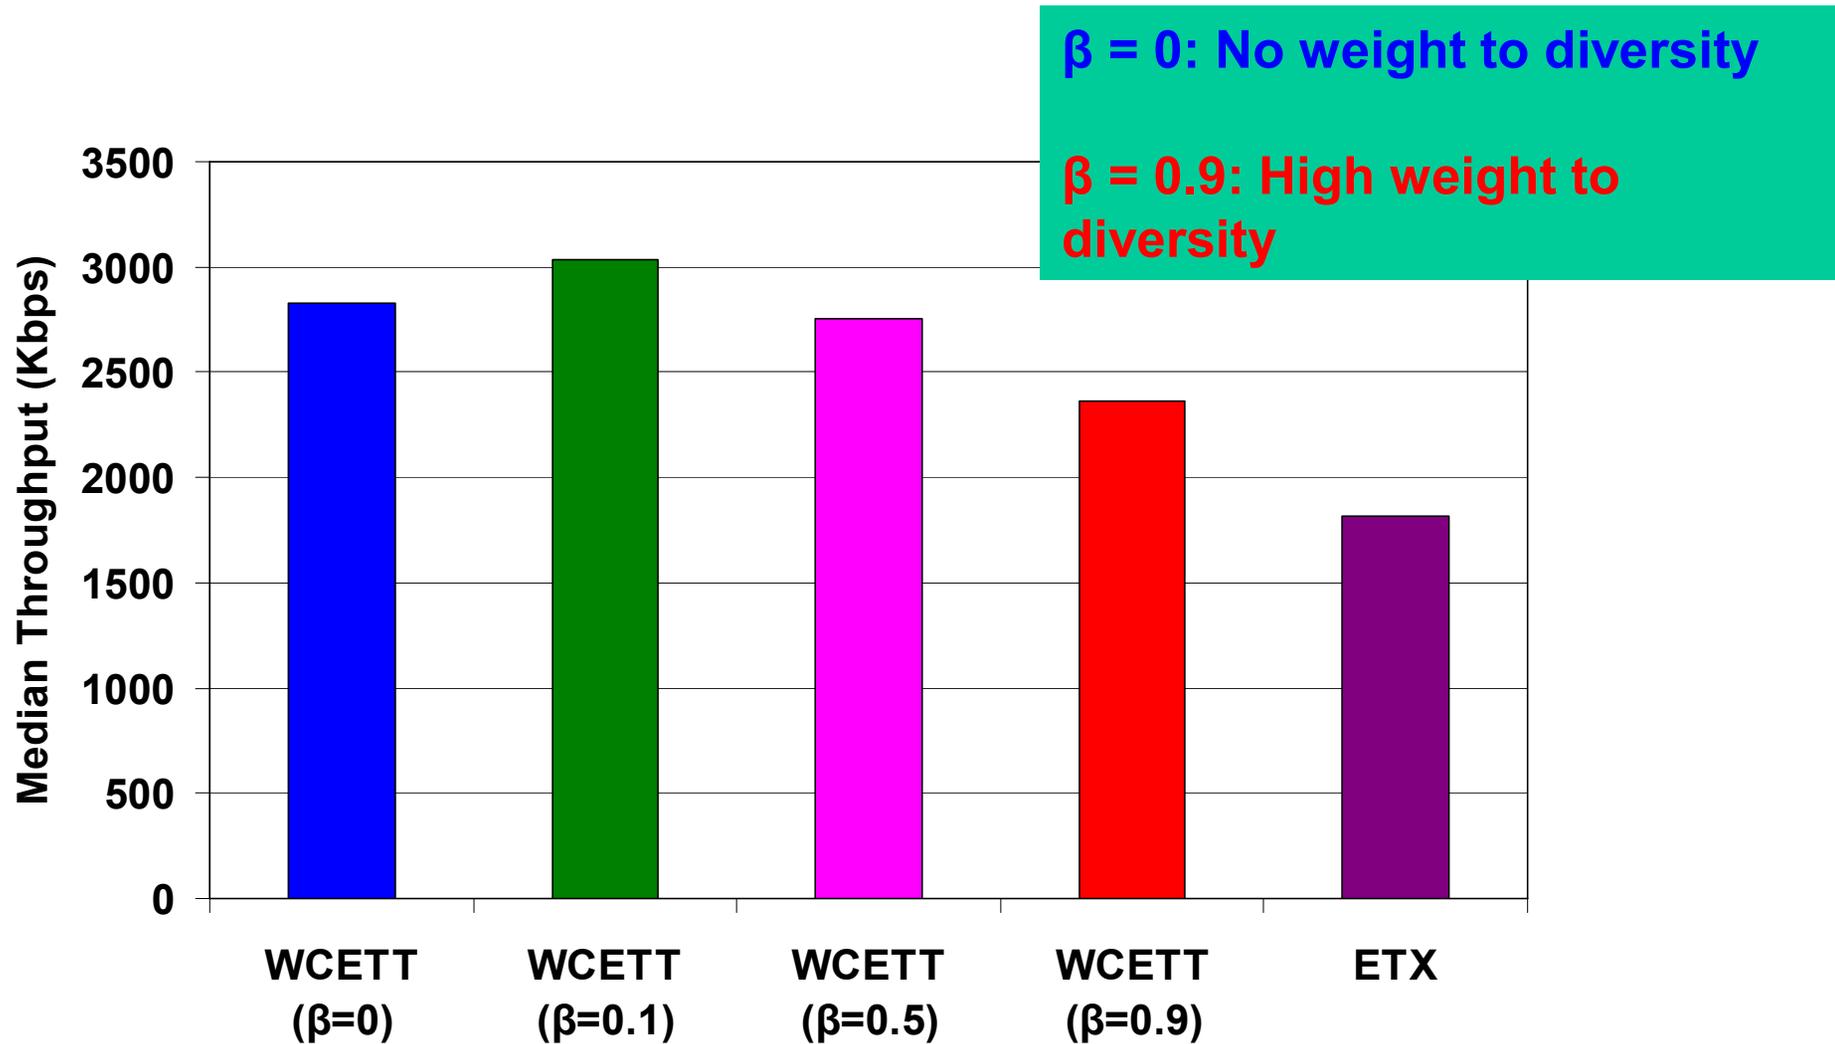
Impact of β value

Channel diversity is important; especially for shorter paths

Performance of Two Simultaneous Flows

- ❐ 2-Minute TCP transfer between 100 randomly selected node pairs

- ❐ Two transfers active at a time

- ❐ Two radios: Netgear: 36-a, Proxim: 10-g

- ❐ Performance metric: 2 x Median throughput

- ❐ Repeat for ETX and WCETT (β = 0, 0.5, 0.9)

Two simultaneous flows

β = 0: No weight to diversity

β = 0.9: High weight to diversity

Median Throughput (Kbps)

3500
3000
2500
2000
1500
1000
500
0

WCETT (β=0)   WCETT (β=0.1)   WCETT (β=0.5)   WCETT (β=0.9)   ETX

Throughput better for lower values of β

Conclusions

- Previously proposed routing metrics are inadequate in multi-radio scenario

- WCETT improves performance by judicious use of 2nd radio
  - Benefits are more prominent for shorter paths

- Optimal value of β depends on load

- Passive inference of loss rate and channel bandwidth