

Roofnet: An 802.11b Mesh Network

Brad Karp
UCL Computer Science



CS M038 / GZ06
14th January, 2011

Context: Mesh Networks

- Ad hoc networking
 - Mobile, highly dynamic topologies
 - Chief metrics: routing protocol overhead, packet delivery success rate, hop count
 - Largely evaluated in simulation
- Sensornets
 - Fixed, resource-impooverished nodes
 - Chief metric: energy consumption
- Mesh networks
 - Fixed, PC-class nodes
 - Motivation: shared Internet access in community
 - Chief metric: TCP throughput
 - Today: Roofnet, a real, deployed mesh network

Context: Mesh Networks

- Ad hoc networking
 - Mobile, highly dynamic topologies
 - Chief metrics: routing protocol overhead, packet delivery success rate, hop count

Why a multi-hop mesh vs. single-hop access points?

- Mesh networks
 - Fixed, PC-class nodes
 - Motivation: shared Internet access in community
 - Chief metric: TCP throughput
 - Today: Roofnet, a real, deployed mesh network

Roofnet: Design Choices

- Volunteer users host nodes at home
 - Open participation without central planning
 - No central control over topology
- Omni antennas
 - Ease of installation by naïve user: no choice of neighbors or aiming
 - Links interfere, likely low quality
- Multi-hop routing (not 1-hop APs)
 - Potentially better coverage, path diversity
 - Routing more complex, end-to-end loss higher
- Goal: high TCP throughput
 - Reachability alone less challenging on (nearly) static network

Roofnet: Design Choices

- Volunteer users host nodes at home
 - Open participation without central planning
 - No central control over topology

Stated non-goals for paper:

- Throughput of multiple flows
- Scalability in number of nodes
- Design of routing protocols
- Performance change over time
- Topology change as users join / leave network

Roofnet Deployment



- Each node: PC, 802.11b card, roof-mounted omni antenna

Node Addresses

- Autoconfiguration of wireless interface IP address
 - High byte: private (e.g., net 10) prefix
 - Roofnet nodes not reachable from Internet
 - Low 3 bytes are low 3 bytes of Ethernet MAC address
- NAT between wired Ethernet and Roofnet
 - Private addresses (net 192.168.1) for wired hosts
 - No address allocation coordination across Roofnet nodes required
 - Roofnet hosts can't connect to one another; only to Internet

Internet Gateways

- Roofnet node tries DHCP on wired Ethernet; then tries reaching Internet hosts; success indicates node is an Internet gateway
- NAT between wireless interface and wired Internet gateway interface
 - Why needed?
- Roofnet nodes track gateway used for each open TCP connection they originate
 - If best gateway changes, open connections continue to use gateway they already do
 - Why?

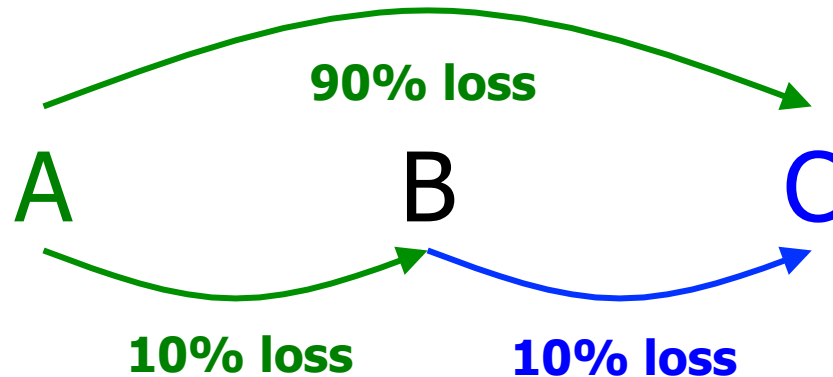
Routing Protocol

- Srcr: DSR-like protocol
- Each link has **metric** (not necessarily 1!)
- Data packets contain **full source routes** (**robust against loops; metric may be dynamic**)
- Nodes keep **database of link metrics**
 - Nodes write current link metric into source route of all packets they forward
 - Nodes flood route queries when cannot find route; queries accumulate link metrics
 - Nodes cache link metrics overheard in queries/responses
- Run Dijkstra's algorithm over database to compute source routes

Link Characteristics

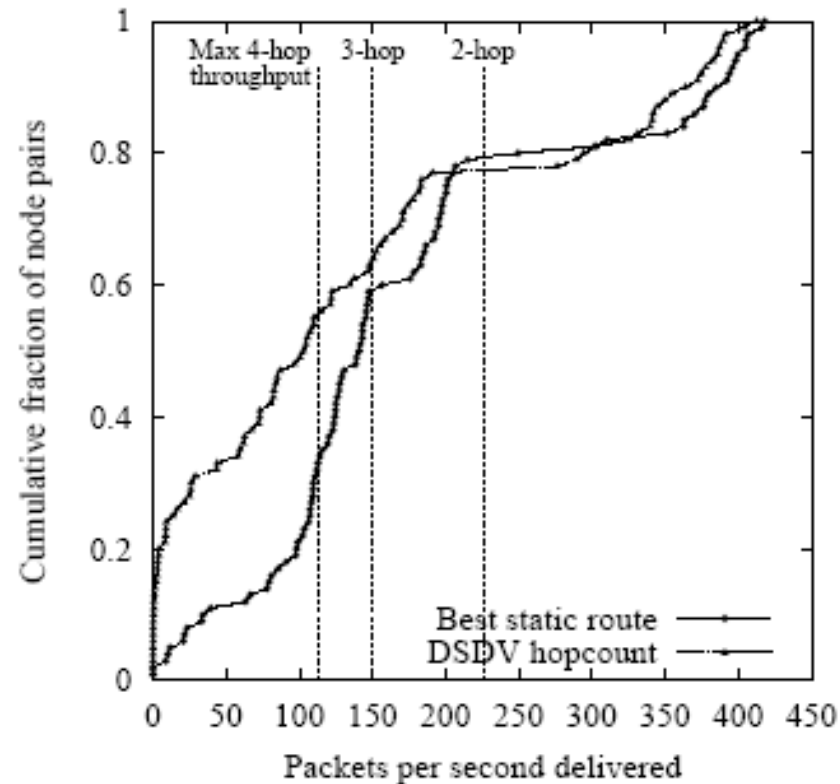
- Wired networks
 - Wired link offers bit error rate 10^{-12}
 - Links “all” (connected) or “nothing” (cut)
- Wireless networks
 - Bit error rate depends on SNR at receiver
 - Dependent on distance, attenuation, &c.
 - Ideal: radio mimics “all or nothing” links; beyond threshold distance, bit error rate approaches 1
 - Reality: links at every bit error (packet loss) rate
 - Are all hops created equal?

Varying Link Loss Rates: Example



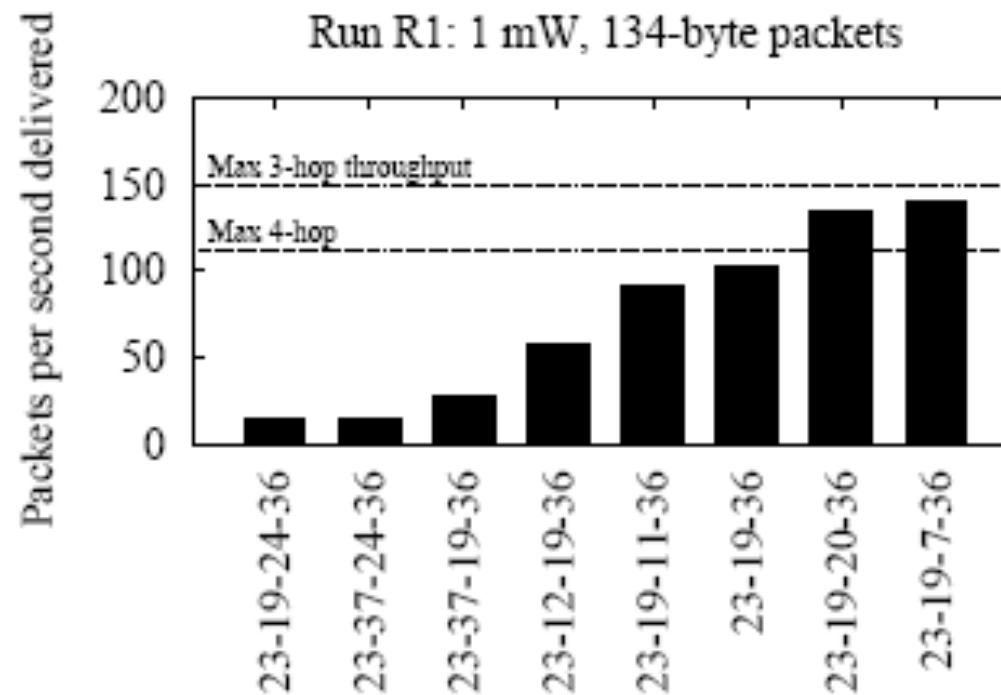
- A → C: 1 hop; **high loss**
- A → B → C: 2 hops; **lower loss**
- But does this happen in practice?

Hop Count and Throughput




- [DeCouto et al., 2003]; indoor predecessor to Roofnet
- 134-byte packets; theoretical 1-hop max = 451 pkts/s

Hop Count and Throughput (cont'd)



- [DeCouto et al., 2003]
- Shortest path not highest throughput
- 4-hop paths span wide range of throughputs

Wireless Link Loss Rates



- [DeCouto et al., 2003]
- Vertical bar ends: loss rates in each direction on one link
- Large fraction of links **very lossy in at least one direction**
- **Asymmetric** loss rates
- **Wide range** of loss rates

Link Metric: Straw Men

- Discard links with loss rate above a threshold?
 - Risks disconnecting nodes!
- Product of link delivery ratios as probability of end-to-end delivery?
 - Ignores inter-hop interference: prefers 2-hop route with 0% loss over 1-hop with 10% loss, when latter is nearly double the throughput
- Throughput of highest-loss link on path?
 - Also ignores inter-hop interference

ETX: Expected Transmissions

- **Link ETX:** predicted number of transmissions
- **Path ETX:** sum of link ETX values on path
- Calculate link ETX using **forward and reverse delivery ratios**
- To avoid retry, **data packet *and* ACK must succeed**
- $ETX = 1 / (d_f \times d_r)$
 - d_f = **forward delivery ratio (data packet)**
 - d_r = **reverse delivery ratio (ACK packet)**

ETX: Expected Transmissions

- **Link ETX:** predicted number of transmissions
- **Path ETX:** sum of link ETX values on path
- Calculate link ETX using **forward and reverse delivery ratios**

Does path ETX allow overlapping transmissions along a path?

Does path ETX offer equal accuracy for paths of all lengths?

– d_f = forward delivery ratio (data packet)

– d_r = reverse delivery ratio (ACK packet)

ETX: Measuring Loss Rates

- Periodically send broadcast probe packets of fixed size
- All nodes know sending rate of probes
- All nodes compute loss rate based on how many arrive per measurement interval
- Nodes enclose loss measurements in their probes (B tells A loss from $A \rightarrow B$)

Multi-Rate Radios

- ETX assumes **all radios run at same bit-rate**
- 802.11b rates: **{1, 2, 5.5, 11} Mbps**
- Cannot compare 2 transmissions at 1 Mbps with 2 at 2 Mbps
- Solution: use **time spent** rather than **transmission count**

ETT: Expected Transmission Time

- ACKs **always sent at 1 Mbps**
- Data packets **typically 1500 bytes**
- Nodes send **1500-byte broadcast probes** at **every bit rate b** (delivery ratio: $d_{f,b}$)
- Nodes send **60-byte (min size) broadcast probes** at **1 Mbps** (delivery ratio: d_r)
- At each bit-rate b , $ETX_b = 1 / (d_{f,b} \times d_r)$
- For packet of length S , $ETT_b = (S/b) \times ETX$
- Link ETT = $\min_b (ETT_b)$

ETT: Assumptions


- Path throughput t given by:

$$t = \frac{1}{\sum_i \frac{1}{t_i}}$$

- where t_i = throughput of hop i
- Underestimates throughput for long paths
 - distant nodes can send simultaneously
- Overestimates throughput for paths with heavy “self-collisions”


Auto Bit-Rate Selection

- Radio firmware automatically chooses bit-rate among {1, 2, 5.5, 11} Mbps
 - avoids bit-rates with high loss rates
- Undesirable policy!



Auto Bit-Rate Selection

- Radio firmware automatically chooses bit-rate among {1, 2, 5.5, 11} Mbps
 - avoids bit-rates with high loss rates
- Undesirable policy!



Auto Bit-Rate Selection (cont'd)

- Ideally, could choose exact bit-rate that at given SNR, gives **highest throughput and nearly zero loss**
- Instead, 802.11b bit-rates **quantized at roughly powers of two**
- Result: **over single hop, bit-rate $2R$ with up to 50% loss always higher-throughput than bit-rate R !**


Auto Bit-Rate Selection in RoofNet: SampleRate

- Samples delivery rates of **actual data packets** using **802.11** retransmit indication
- Occasionally sends packets at rates **other than current rate**
- Sends most packets at rate **predicted to offer best throughput** (as with ETT)
- Adjusts per-packet bit-rate **faster than ETT**
 - only 1 hop of information required
 - delivery ratio estimates not periodic, but per-packet

RoofNet Evaluation


- TCP always single flow at a time
- Multi-hop: 15-second, 1-way bulk TCP transfers between all pairs of nodes
- Single-hop: same, direct link between all pairs of nodes
- Loss matrix: loss rate between all pairs for 1500-byte broadcasts at each bit-rate
- No RTS/CTS (more later!)
- Background traffic: users always active

End-to-End Throughput



- Mean: 627 kbps; median: 400 kbps

End-to-End Throughput



- Mean: 627 kbps; median: 400 kbps
- Routing queries fail for 10% of pairs; link losses, retries fail

Hop Count, Throughput, Latency

Hops	Number of Pairs	Throughput (kbits/sec)	Latency (ms)
1	158	2451	14
2	303	771	26
3	301	362	45
4	223	266	50
5	120	210	60
6	43	272	100
7	33	181	83
8	14	159	119
9	4	175	182
10	1	182	218
no route	132	0	–
Avg: 2.9	Total: 1332	Avg: 627	Avg: 39

- Neighboring nodes interfere with one another

Theoretical Max Throughput (Lossless)

Rate	Max Throughput (kbits/sec)		
	1 Hop	2 Hops	3 Hops
1	890	445	297
2	1634	817	545
5.5	3435	1718	1145
11	5013	2506	1671

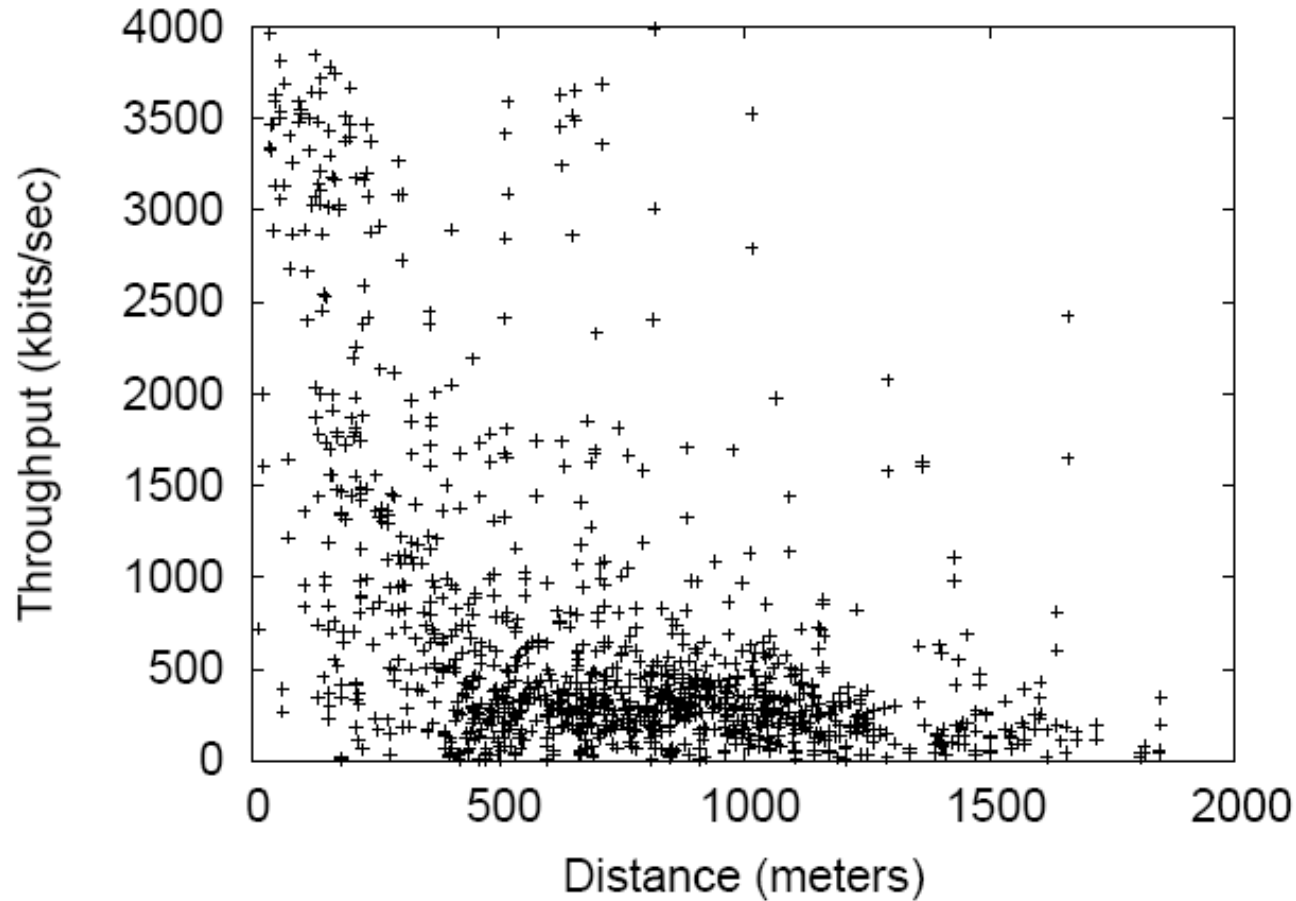
- Computed analytically, assuming hops don't forward in parallel
- One-hop routes seem to use 5.5 Mbps
- Longer routes far **slower than predicted**

User Experience: Mean Throughput from Gateway

Hops	Number of nodes	Throughput (kbits/sec)	Latency (ms)
1	12	2752	9
2	8	940	19
3	5	552	27
4	7	379	43
5	1	89	37
Avg: 2.3	Total: 33	Avg: 1395	Avg: 22


- Latency: 84-byte ping; interactive use OK
- Acceptable throughput, even 4 hops out

Link Quality vs. Distance: All Links



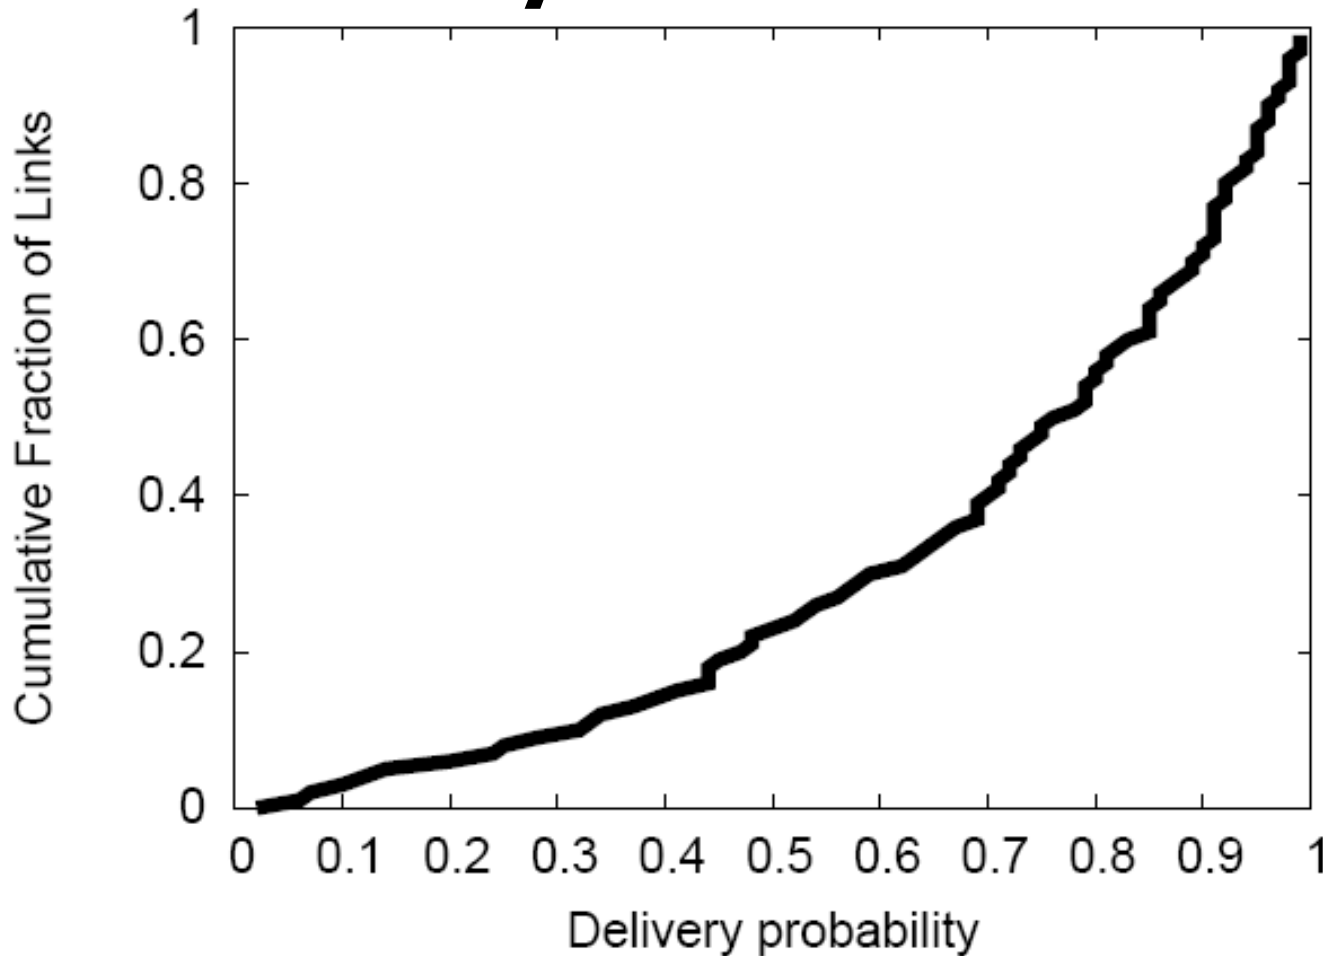
- Single-hop TCP workload
- Many links ca. 500 kbps of varying lengths
- A few short, high-throughput links; a very few long, high-throughput links

Link Quality vs. Distance: Srcr Links



- Multi-hop TCP workload
- Srcr favors short, fast links

SampleRate Bit-Rate Choice: Lossy Links Useful




- Median: 0.8; 20%+ loss links used half the time

Density Evaluation


- Want to evaluate Roofnet with varying numbers of nodes (== varying density)
- One-hop TCP throughput known by measurement
- Using path ETT formula, can estimate multi-hop TCP throughput for any path
- Choose random node subsets, compute estimated throughput using only subset member nodes in paths

Node Density and Connectivity



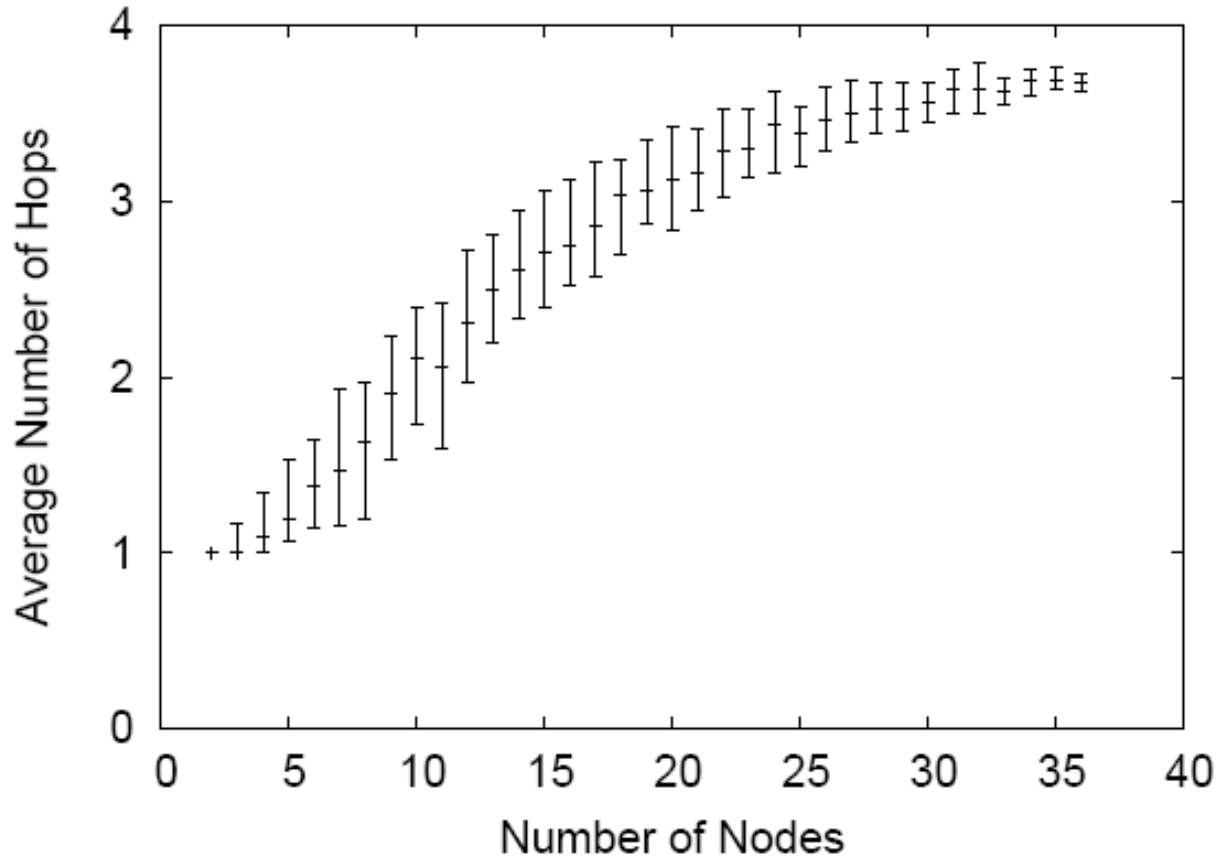
- 25th, 50th, 75th %iles over 100 random subsets
- Connected = ≥ 1 kbyte / s throughput

Node Density and Throughput




- Why does throughput increase?

Node Density and Path Length



- Increasing density **increases diversity: adds short, low-loss links!**

Diversity in Node Use: “Meshness”

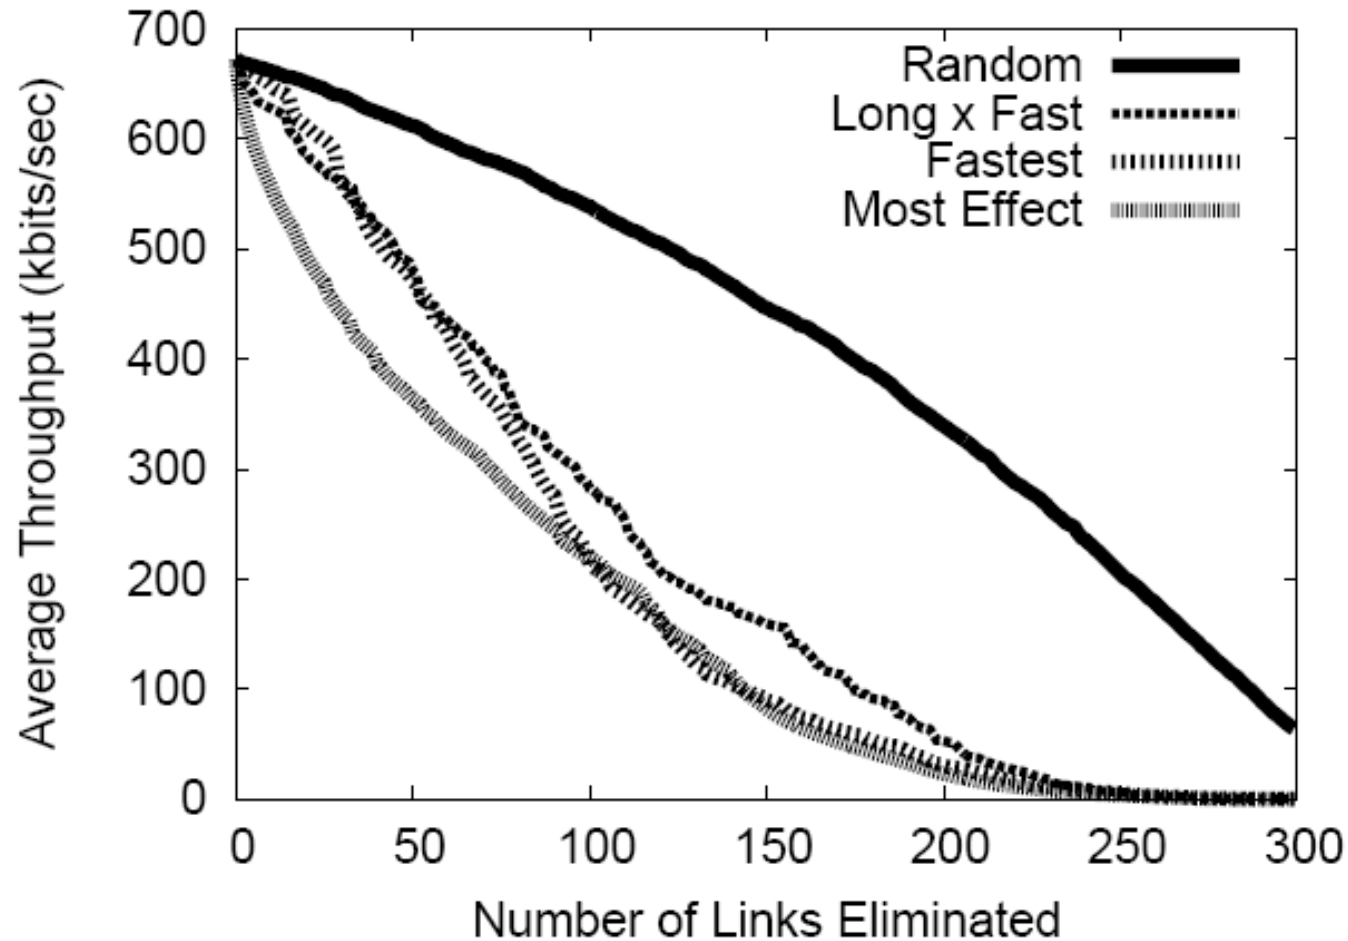


- Most nodes route via a **diverse set of neighbors**

Mesh Robustness Evaluation: Sensitivity to Eliminated Links

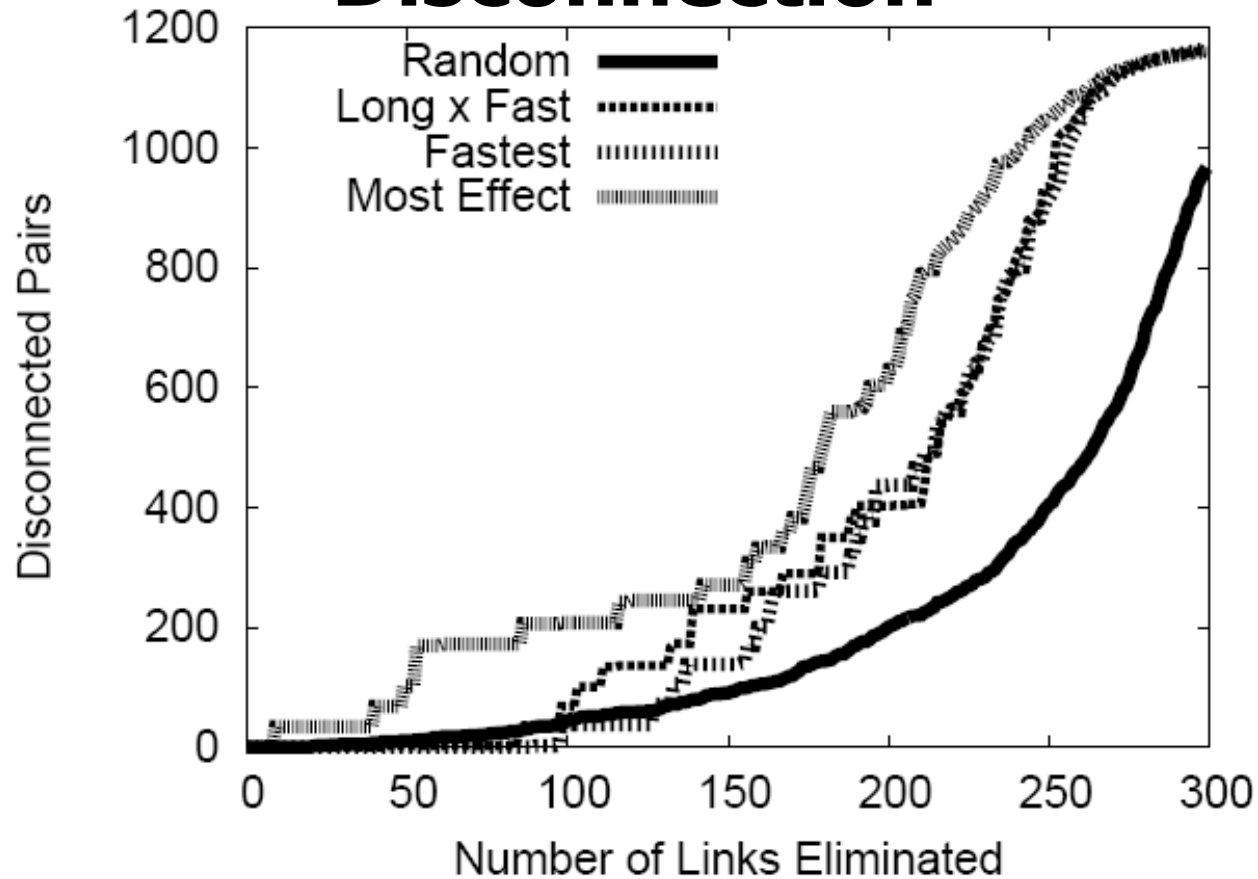
- Know single-hop TCP throughputs for all node pairs
- Try eliminating links, compute multi-hop throughputs analytically (ETT path equation)
- Orders of link removal:
 - **Most Effect:** link that decreases average throughput most
 - **Long x Fast:** link with greatest product distance x tput
 - **Fastest:** link with greatest throughput
 - **Random:** mean of 40 simulations, deleted in random order

Link Elimination Sensitivity: Average Throughput




- Best few links matter a lot
- Over 50 links lost before throughput halved

Link Elimination Sensitivity: Disconnection



- Long & fast links more essential to connectivity than fastest links

Node Elimination Sensitivity: Average Throughput



- Eliminate nodes that appear in the most all-pairs routes
- First two eliminations **reduce throughput by 43%**;
thereafter **more gradual**

Why not Access Points?

- Mesh networking is far from perfect
 - **Complexity** of multi-hop routing and path selection, vs. single-hop access point choice
 - **Interference** between neighboring forwarding hops
 - **Loss** substantially increases with path length
- **Could we do better with same hardware?**
 - Place nodes as before
 - Same goal: Internet access for all nodes
 - Constrain topology to access point case: **all nodes one hop from an Internet gateway**

Evaluation Strategy: Multi-Hop vs. AP

- Add gateways to the network one by one
- “Optimal”: at each step, add gateway that maximizes number of nodes that becomes newly connected with non-zero throughput
- “Random”: use randomly selected set of gateways of designated size; repeat for 250 trials; take median set (by # of connected nodes)
- Break ties by mean throughput

Optimal Gateway Placement

GWs	Multi-Hop		Single-Hop	
	Conn	Throughput (kbits/sec)	Conn	Throughput (kbits/sec)
1	37	781	23	174
2	37	1450	32	824
3	37	1871	34	1102
4	37	2131	36	1140
5	37	2355	37	1364
6	37	2450	37	2123
7	37	2529	37	2312
8	37	2614	37	2475
9	37	2702	37	2564
10	37	2795	37	2659
⋮	⋮	⋮	⋮	⋮
15	37	3197	37	3180
20	37	3508	37	3476
25	37	3721	37	3658

- Complete coverage: **5 GWs in single-hop; 1 GW in multi-hop**
- Multi-hop offers **greater throughput at any number of gateways (why?)**

Random Gateway Placement

GWs	Multi-Hop		Single-Hop	
	Conn	Throughput (kbits/sec)	Conn	Throughput (kbits/sec)
1	34	760	10	535
2	35	1051	17	585
3	35	1485	22	900
4	35	2021	25	1260
5	36	1565	28	1221
6	36	1954	30	1192
7	36	1931	31	1662
8	37	1447	32	1579
9	37	1700	33	1627
10	37	1945	34	1689
⋮	⋮	⋮	⋮	⋮
15	37	2305	36	1714
20	37	2509	36	2695
25	37	2703	37	2317

- Complete coverage: **8 GWs for multi-hop; 25 for single-hop**

Random Gateway Placement

GWs	Multi-Hop		Single-Hop	
	Conn	Throughput (kbits/sec)	Conn	Throughput (kbits/sec)
1	24	760	10	525


For few gateways, random placement with multi-hop outperforms optimal placement with single-hop

For many gateways, optimal placement with single-hop outperforms random placement with multi-hop

⋮	⋮	⋮	⋮	⋮
15	37	2305	36	1714
20	37	2509	36	2695
25	37	2703	37	2317

- Complete coverage: **8 GWs for multi-hop; 25 for single-hop**

Forwarding Creates Interference



- Multi-hop throughput **less than predicted**
- Reason: **interference between successive forwarding hops**

RTS/CTS Don't Prevent Interference

Hops	Pairs	Average throughput without	Throughput with
1	3	2094	1735
2	5	836	725
3	6	314	312

- Mean throughputs for node pairs separated by paths of various lengths
- Single-hop: RTS/CTS **just overhead**
- Multi-hop: RTS/CTS **don't improve throughput**