Distributed Mesh Wireless Networks

COS 418: *Distributed Systems*
Lecture 20

Kyle Jamieson

[Selected content adapted from B. Karp]

Today

1. **Roofnet: An unplanned Wi-Fi Mesh network**
   - Wireless mesh link measurements
   - Routing and bit rate selection
   - End-to-end performance evaluation

2. Advertisement for COS-598A

2

Context, *ca.* 2000-2005

- **Mobile ad hoc networking** research
  - Mobile, hence highly dynamic topologies
  - Chief metrics: routing protocol overhead, packet delivery success rate, hop count
  - Largely evaluated in simulation
- **Today: Roofnet,** a real mesh network **deployment**
  - Fixed, PC-class nodes
  - Motivation: shared Internet access in community
  - Chief metric: TCP throughput
  - "Test of time" system, led to **Cisco Meraki**

3

Roofnet: Design Choices

1. Volunteer users host nodes at home
   - Open participation **without central planning**
   - No central **control** over topology
2. Omnidirectional rather than directional antennas
   - **Ease of installation**: no choice of neighbors/aiming
   - Links **interfere**, likely **low quality**
3. Multi-hop routing (not single-hop hot spots)
   - Improved **coverage** (path diversity)
   - Must build a routing protocol
4. Goal: high TCP throughput

4

Roofnet: Goals and non-goals

- Each part of the mesh architecture had been **previously examined** in isolation
- **Paper contribution:** A systematic evaluation of whether their architecture can achieve the goal of providing Internet access
- **Stated non-goals for paper:**
  - Throughput of multiple concurrent flows
  - Scalability in number of nodes
  - Design of routing protocols

5

Roofnet deployment

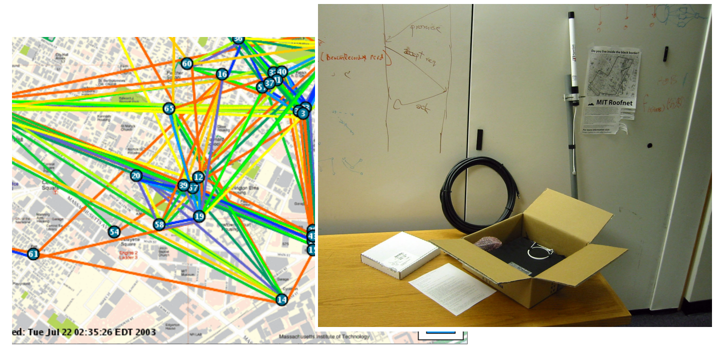

- Each node: PC, 802.11b card, roof-mounted omni antenna

6

Hardware design

- PC Ethernet interface provides wired Internet for user
- Omnidirectional antenna in **azimuthal** direction
  - 3 dB **vertical** beam width of 20 degrees
    - Wide beam sacrifices gain but **removes the need** for perfect vertical **antenna orientation**
- **802.11b** radios (*Intersil Prism 2.5* chipset)
  - 200 mW transmit power
  - All share same 802.11 channel (frequency)

7

Node addresses

- Auto-configuration of wireless interface IP address
  - High byte: private (*e.g.,* net 10) prefix
    - Roofnet nodes **not reachable from Internet**
  - Low three bytes: low 24 bits of Ethernet address
- NAT between wired Ethernet and Roofnet
  - Private addresses (192.168.1/24) for wired hosts
    - Can't connect to one another; only to Internet
  - **Result: No address allocation coordination** across Roofnet boxes required

8

Internet gateways

- Node sends DHCP request on Ethernet then tests reachability to Internet hosts
  - Success indicates node is an **Internet gateway**
    - Gateways **translate** between Roofnet and Internet IP address spaces
- 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
- If a Roofnet gateway fails, **existing** TCP connections through that gateway **will fail**

9

Links: Wired *v.* wireless

- **Wired** links
  - Most wired links offer bit error rate *ca.* **10^{-12}**
  - Links are **"all"** (connected) or **"nothing"** (cut)
- **Wireless** links
  - Bit error rate depends on **signal to interference plus noise ratio** (SNR) at receiver
  - Dependent on distance, attenuation, interference
- **Would like:** Wireless links like wired links

10

Example: Varying link loss rates

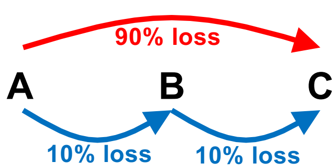

- A → C: 1 hop; **high loss**
- A → B → C: 2 hops; **lower loss**
- But **does this happen** in practice?

11

Hop count and throughput (1)

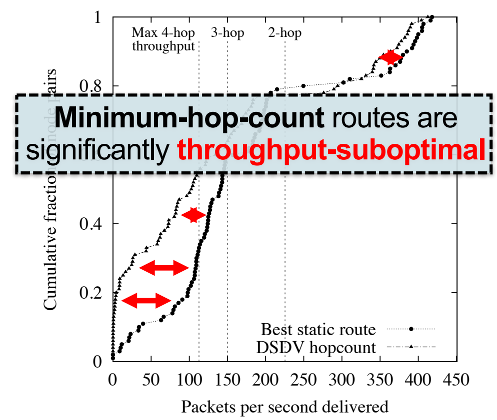

12

Hop count and throughput

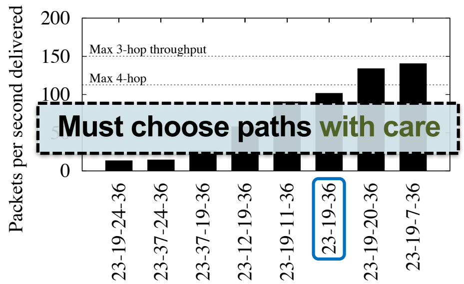

Must choose paths with care

- Two-hop path is **suboptimal**
- Some **3-hop paths better, some worse** than 2-hop

13

Link loss is high and asymmetric

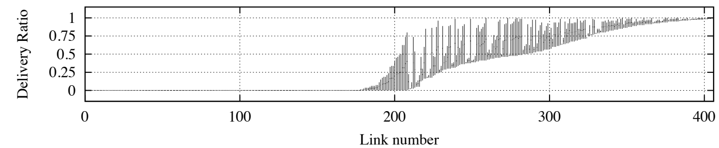

- Vertical bar ends = loss rate on 1 link in each direction
- Many links **asymmetric** and **very lossy in ≥ 1 way**
- **Wide range** of loss rates

14

Today

1. **Roofnet: An unplanned Wi-Fi Mesh network**
   - Wireless mesh link measurements
   - **Routing and bit rate selection**
   - End-to-end performance evaluation

2. Advertisement for COS-598A

15

Routing protocol: Srcr (1)

- Each link has an associated ***metric*** (not necessarily 1!)
- Data packets contain **source routes**
- Nodes keep **database of link metrics**
  - Nodes write current metric into source route of all forwarded packets
  - Nodes flood ***route queries*** when they can't find a route; queries accumulate link metrics
    - Route queries contain route from requesting node
  - Nodes cache overheard link metrics
- **Dijkstra's algorithm** computes source routes

16

Routing Protocol: Srcr (2)

- Gateways periodically flood queries for a non-existent destination address
  - Everyone learns route to the gateway
  - When a node sends data to gateway, gateway learns route back to the node
- Flooded queries **might not follow the best route**; solution:
  1. Add link metric info in query's source route to database
  2. Compute **best route** from query's source
  3. Replace **query's path from source** with **best route**
  4. Rebroadcast the modified query

17

Link metric: Strawmen

- *Discard links with loss rate above a threshold?*
  - Risks **unnecessarily disconnecting** nodes
- *Product of link delivery rates → prob. of e2e delivery?*
  - **Ignores inter-hop interference**
    - Prefers 2-hop, 0% loss route over 1-hop, 10% loss route (but latter is **double throughput**)
- *Throughput of highest-loss link on path?*
  - Also **ignores inter-hop interference**

18

ETX: Expected Transmission Count

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

19

Measuring link delivery ratios

- Nodes periodically send broadcast ***probe*** packets
  - All nodes know the **sending period** of probes
  - All nodes **compute loss rate** based on how many probes arrive, per measurement interval
- Nodes **enclose these loss measurements** in their transmitted probes
  - *e.g.* **B** tells node **A** the link delivery rate from **A** to **B**

20

Multi-bitrate radios

- ETX assumes all radios run at same bit-rate
  - But 802.11b rates: {1, 2, 5.5, 11} Mbit/s
- **Can't compare** two transmissions at 1 Mbit/s with two at 2 Mbit/s
- **Solution:** Use expected **time** spent on a packet, rather than transmission count

21

ETT: Expected Transmission Time

- ACKs always sent at 1 Mbps, data packets 1500 bytes
- Nodes send 1500-byte broadcast probes at every bit rate b to compute ***forward link delivery rates*** $\boldsymbol{d_f(b)}$
  - Send 60-byte (min size) probes at 1 Mbps → $\boldsymbol{d_r}$
- At each bit-rate b, $ETX_b = 1 / (d_f(b) \times d_r)$
- For packet of length S, $ETT_b = (S / b) \times ETX_b$
- **Link ETT** = $\min_b (ETT_b)$

22

ETT: Assumptions

- Path throughput estimate t is given by $t = \dfrac{1}{\sum_{\text{hop } i \in \text{path}} \frac{1}{t_i}}$
  - t_i = throughput of hop i
- *Does ETT maximize throughput?* No!

1. Underestimates throughput for long (≥ 4-hop) paths
   - Distant nodes can send simultaneously
2. Overestimates throughput when transmissions on different hops collide and are lost

23

Auto bit-rate selection

- Prism radio firmware (*ca.* 2005) automatically chose bit-rate among {1, 2, 5.5, 11} Mbps
  - Avoids bit-rates with high loss rates
- Undesirable policy!

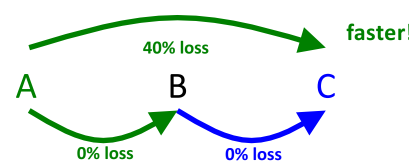

24

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 are **quantized at roughly powers of two**
- **Result:** Over a single hop, bit-rate **2R** with **up to 50% loss** always **higher throughput** than bit-rate **R**!

25

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 route selection
  - Only one hop of information required
  - Delivery ratio estimates not periodic, but per-packet

26

Today

1. **Roofnet: An unplanned Wi-Fi Mesh network**
   - Wireless mesh link measurements
   - Routing and bit rate selection
   - **End-to-end performance evaluation**
2. Advertisement for COS-598A

27

Roofnet evaluation

Datasets:

1. Multi-hop TCP: 15-second, 1-way bulk TCP transfers between **all node pairs**
2. Single-hop TCP: same, direct link between **all node pairs**
3. Loss matrix: loss rate between **all node pairs** for 1500-byte broadcasts at each bit-rate

- TCP flows, always a **single flow at a time**
- But background traffic present: **users always active**

28

Wide spread of end-to-end throughput

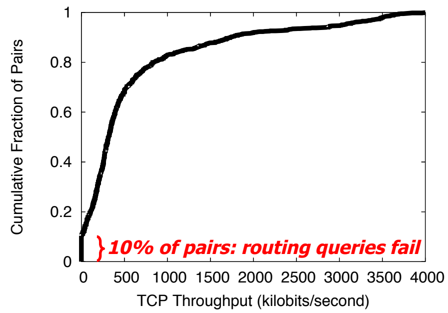

- Multi-hop TCP dataset
- Mean: 627 kbps; median: 400 kbps

29

End-to-end throughput by hop count

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

- **Higher** hop count correlates with **lower** throughput
  - Neighboring nodes **interfere** with one another

30

Comparing with computed throughput

All-pairs empirical

Hops	Number of Pairs	Throughput (kbits/sec)	Latency (ms)
1	158	2451	14
2	303	771	26
3	301	362	45

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

Theoretical, loss-free maximum throughput

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

31

Forwarding indeed creates interference

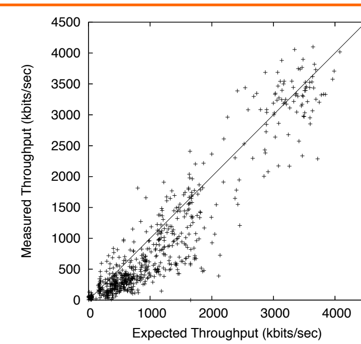

- Multi-hop **measured** throughput often **less** than predicted
- **Reason:** Interference between successive forwarding hops

32

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; okay for **interactive use**
- **Acceptable throughput** (379 Kbit/sec), even four hops out

33

What link ranges/speeds to expect?

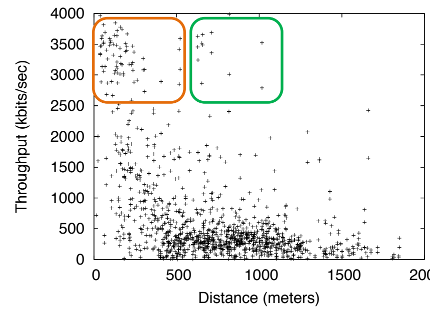

- Single-hop TCP workload
- Many links of varying lengths support ≈ 500 Kbit/s
- A **few short and fast** links; **very few long and fast** links

34

Which network links does Srcr use?

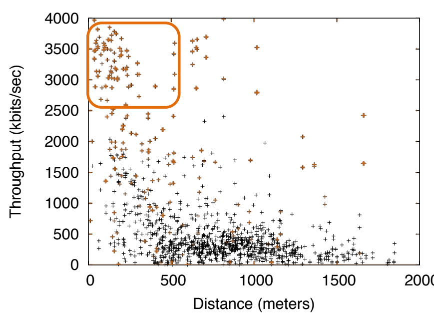

- **Multi-hop TCP workload: links Srcr uses in red,** all others (single-hop TCP) in black
- Srcr somewhat favors **short, fast** links

35

Lossy Links are Useful

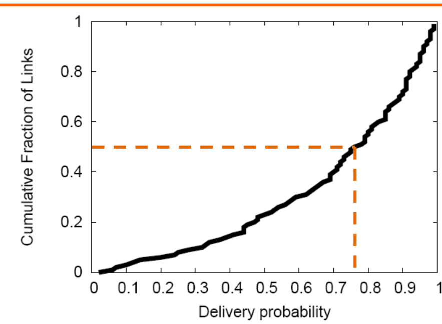

- Delivery probability for links Srcr uses, **at the bit rate SampleRate chooses**
- **>25%-loss links** used **half** the time

36

Diversity in node use: "Meshness"

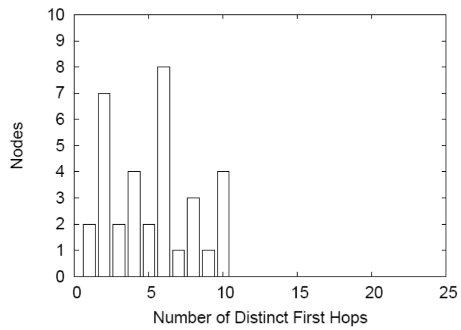

- Most nodes route via a **diverse set of neighbors**

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 the same hardware?*
  - Place nodes as before
  - **Same goal:** Internet access for all nodes
  - Constrain topology to **access point (AP)** case
    - All nodes are one hop from an Internet gateway AP

Evaluation strategy: Multi-hop *v.* AP

- Add gateways (GWs) to the network one by one
- **"Optimal":** at each step, add the GW that **maximizes** number of **newly connected** nodes
- **"Random":** use randomly selected set of GWs of designated size; repeat for 250 trials; take median set (by number of connected nodes)

Optimal AP (GW) placement

GWs	Multi-Hop Conn	Multi-Hop Throughput (kbits/sec)	Single-Hop Conn	Single-Hop 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 for single-hop versus 1 for multi-hop
- Multi-hop is **faster** for any number of gateways
  - Can **use short, high-quality** links

Random AP (GW) placement

GWs	Multi-Hop Conn	Multi-Hop Throughput (kbits/sec)	Single-Hop Conn	Single-Hop 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

- More realistic scenario
- **Complete coverage:** **eight** GWs for multi-hop, **25** for single-hop
  - Route query failure (no retransmissions)
- For ≤ 5 GWs, randomly chosen multi-hop GWs outperform optimally chosen single-hop APs

Roofnet: Concluding thoughts

- Network's architecture designed for **ease of deployment**
  - Omni-directional antennas, self-configuring software, multi-hop routing
- Performance evaluation showed that an **unplanned mesh works well**

Today

1. Roofnet: An unplanned Wi-Fi Mesh network
2. **Advertisement for COS-598A**

COS-598A *Wireless Networking and Sensing Systems*

- Graduate-level seminar open to advanced undergraduates (see me to discuss)
- Explores **recent developments** in:
  - **Wireless** data communication **networks**
  - **Wireless sensing** systems
- Introduces you to the wireless physical layer
  - In a way that is **accessible** for students with solely a **computer systems and networking** background

www.cs.princeton.edu/courses/archive/spring17/cos598A

COS-598A: Topics and goals

- TCP over wireless
- **Rateless** error control codes
- Wi-Fi based **localization**
- Indoor **radar**
- **Full-duplex** wireless

- **Goal:** Understand the state of the art in the above areas
  - Develop taste in research

- **Goal:** Investigate novel ideas in the above areas thru project

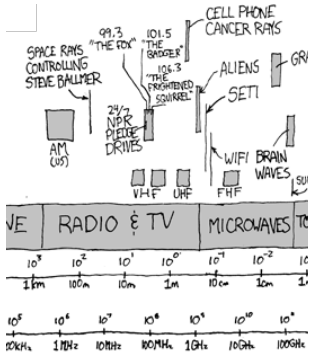

45

Wednesday topic:
Big Data Processing
Graph processing (GraphLab)

46