COS-598A, Spring 2017, Kyle Jamieson

Lecture 3: Roofnet

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 using fixed, PC-class nodes
- Motivation: shared Internet access in community
- Chief metric: TCP throughput
- "Test of time" system, led to Cisco Meraki

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
    - Goal: high TCP throughput

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

Roofnet deployment

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

Node addressing

- 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

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

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

Example: Varying link loss rates

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

Hop count and throughput (1)

Minimum-hop-count routes are significantly throughput-suboptimal

Hop count and throughput

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

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

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

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

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

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 / (df × dr)
  - df = forward link delivery ratio (data packet)
  - dr = reverse link delivery ratio (ack packet)

- Path ETX: sum of the link ETX values on a path

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

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

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 df(b)
  - Send 60-byte (min size) probes at 1 Mbps dr

- At each bit-rate b, $ETX_b = 1 / (df(b) \times dr)$
- For packet of length S, $ETT_b = (S / b) \times ETX_b$
- Link ETT = min*b (ETT*b)

ETT: Assumptions

- Path throughput estimate t is given by
  - t_i = throughput of hop i

- Does ETT maximize throughput? No!
- Underestimates throughput for long (≥ 4-hop) paths
  - Distant nodes can send simultaneously

- Overestimates throughput when transmissions on different hops collide and are lost

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!
- 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!

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

Roofnet evaluation

- Datasets:
- Multi-hop TCP: 15-second, 1-way bulk TCP transfers between all node pairs
- Single-hop TCP: same, direct link between all node pairs
- 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

Wide spread of end-to-end throughput

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

End-to-end throughput by hop count

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

Comparing with computed 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

Forwarding indeed creates interference

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

User experience: Mean throughput from gateway

- Latency: 84-byte ping; okay for interactive use
- Acceptable throughput (379 Kbit/sec), even four hops out

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

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

Lossy Links are Useful

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

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

- 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

- 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