Ad Hoc Routing

CMU 15-744

David Andersen

Ad Hoc Routing

- Goal: Communication between wireless nodes
  - No external setup (self-configuring)
  - Often need multiple hops to reach dst

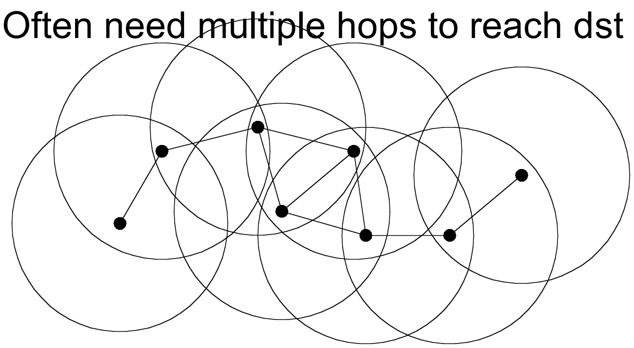

Challenges and Variants

- Poorly-defined “links”
  - Probabilistic delivery, etc. Kind of n^2 links
- Time-varying link characteristics
- No oracle for configuration (no ground truth configuration file of connectivity)
- Low bandwidth (relative to wired)
- Possibly mobile
- Possibly power-constrained

Goals

- #0: Provide connectivity!
- Low consumption of memory, bandwidth, {possibly power}
- Scalable with numbers of nodes
- Localized effects of link failure
  - (For scalability and stability)

Standard Routing: DV and LS

- DV protocols may form loops
  - Very wasteful in wireless: bandwidth, power
  - Loop avoidance sometimes complex
- LS protocols: high storage and communication overhead – particularly when potentially n^2 links!
- More links in wireless (e.g., clusters) - may be redundant → higher protocol overhead

Problems Using DV or LS

- Periodic updates waste power
  - Tx sends portion of battery power into air
  - Reception requires less power, but periodic updates prevent mobile from "sleeping"
- Convergence may be slower in conventional networks but must be fast in ad-hoc networks and be done without frequent updates

Design Space

- 1) How to disseminate information about links and to send packets along a path
- 2) How to decide which path to use from many possibilities
  - (How good is a particular path?)
  - Really early models: binary
  - Early models: If deliver > x%, good
  - New models: ETX/ETT (wait for it)

- Base knowledge: Every node knows about neighbors because they can hear them directly. (Periodic beacons, transmissions, etc.)

Evaluating Ad Hoc Protocols

- Parameter question: How much mobility?
  - And what mobility model?
  - Consider reality: Random waypoint? Clustered?
    - Businesspeople wandering to/from work vs. soldiers, etc.
- Link model
  - Early research all used spherical propagation, etc.
    - Tended to binary “working” or “not working”
  - More modern uses traces from real deployments or more realistic models

Proposed Protocols

- Destination-Sequenced Distance Vector (DSDV)
  - DV protocol, destinations advertise sequence number to avoid loops, not on demand
- Temporally-Ordered Routing Algorithm (TORA)
  - On demand creation of hbh routes based on link-reversal
- Dynamic Source Routing (DSR)
  - On demand source route discovery
- Ad Hoc On-Demand Distance Vector (AODV)
  - Combination of DSR and DSDV: on demand route discovery with hbh routing

DSR Concepts

- Source routing
  - No need to maintain up-to-date info at intermediate nodes
- On-demand route discovery
  - No need for periodic route advertisements

DSR Components

- Route discovery
  - The mechanism by which a sending node obtains a route to destination
- Route maintenance
  - The mechanism by which a sending node detects that the network topology has changed and its route to destination is no longer valid

DSR Route Discovery

- Route discovery - basic idea
  - **Source** broadcasts route-request to **Destination**
  - Each node forwards request by adding own address and re-broadcasting
  - Requests propagate outward until:
    - Target is found, or
    - A node that has a route to Destination is found

C Broadcasts Route Request to F

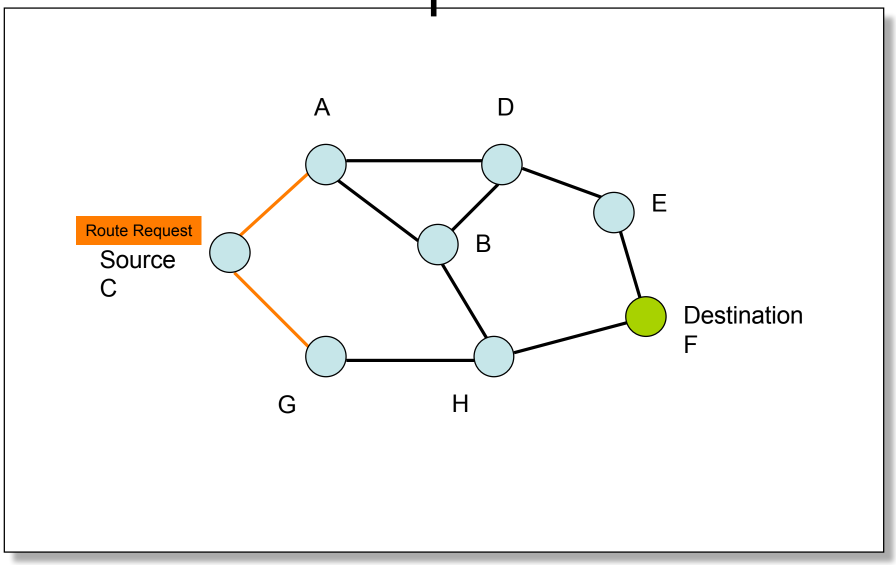

C Broadcasts Route Request to F

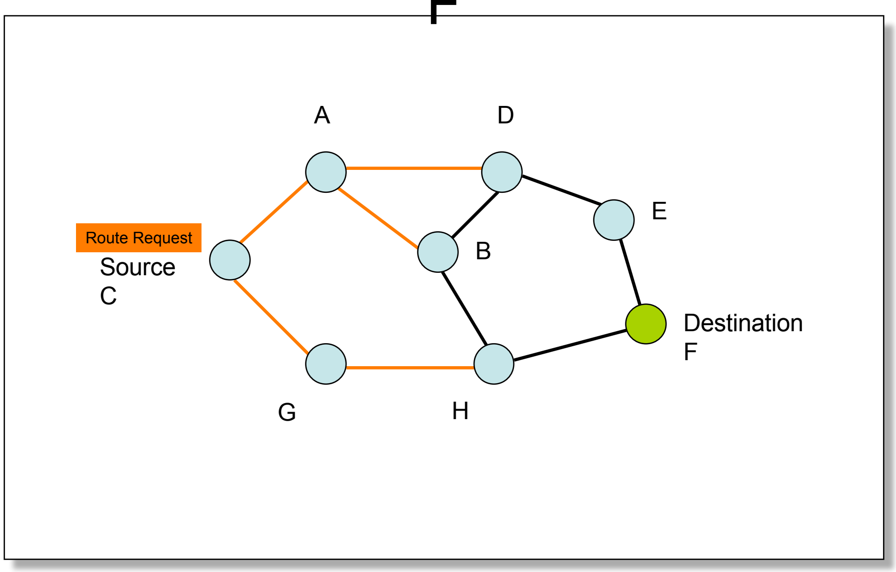

H Responds to Route Request

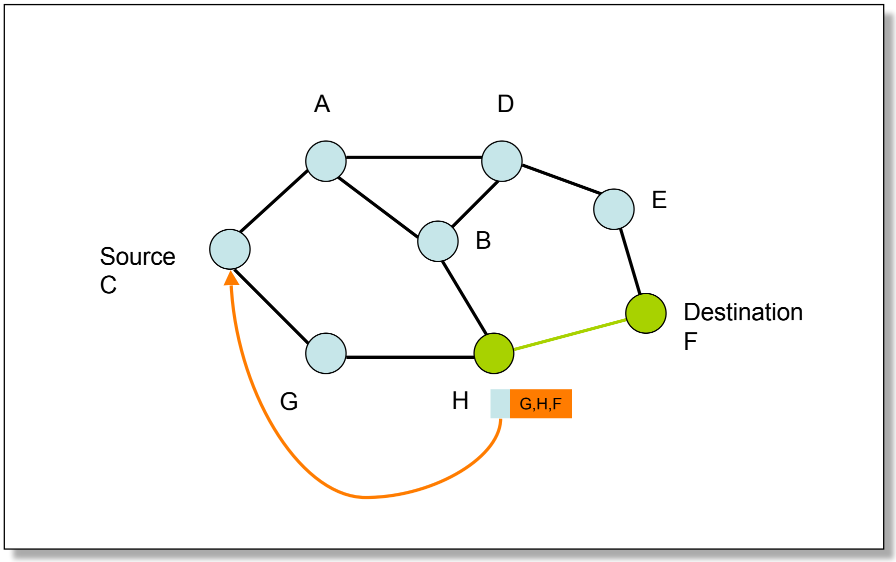

C Transmits a Packet to F

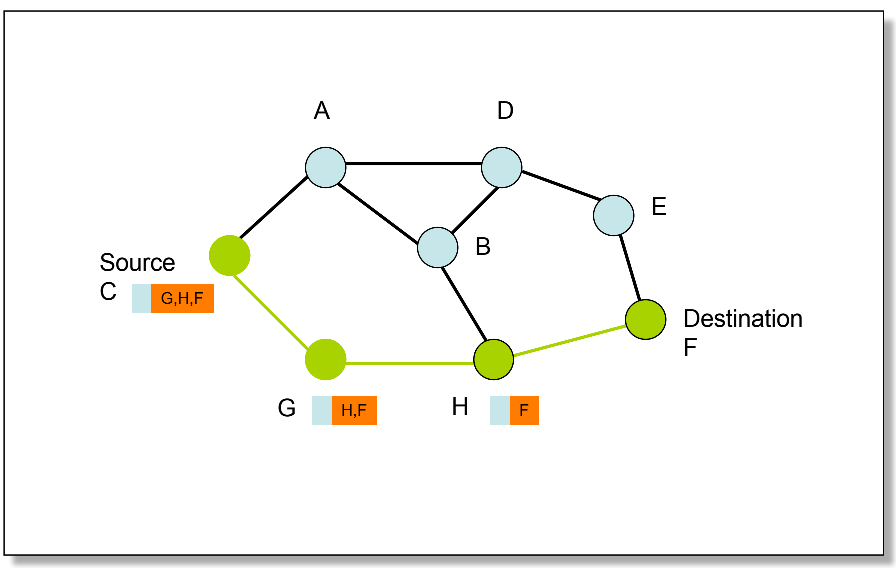

Forwarding Route Requests

- A request is forwarded if:
  - Node is not the destination
  - Node not already listed in recorded source route
  - Node has not seen request with same sequence number
  - IP TTL field may be used to limit scope
- Destination copies route into a Route-reply packet and sends it back to **Source**

Route Cache

- All source routes learned by a node are kept in Route Cache
  - Reduces cost of route discovery
- If intermediate node receives RR for destination and has entry for destination in route cache, it responds to RR and does not propagate RR further
- Nodes overhearing RR/RP may insert routes in cache

Sending Data

- Check cache for route to destination
- If route exists then
  - If reachable in one hop
    - Send packet
  - Else insert routing header to destination and send
- If route does not exist, buffer packet and initiate route discovery

Discussion

- Source routing is good for on demand routes instead of a priori distribution
- Route discovery protocol used to obtain routes on demand
  - Caching used to minimize use of discovery
- Periodic messages avoided
- But need to buffer packets
- How do you decide between links?

Deciding Between Links

- Most early protocols: Hop Count
  - Link-layer retransmission can mask some loss
  - But: a 50% loss rate means your link is only 50% as fast!
- Threshold?
  - Can sacrifice connectivity. ☹
  - Isn't a 90% path better than an 80% path?
- Real life goal: Find highest throughput paths

Forwarding Packets is expensive

- Throughput of 802.11b =~ 11Mbits/s
  - In reality, you can get about 5.
- What is throughput of a chain?
  - A -> B -> C ?
  - A -> B -> C -> D ?
  - Assume minimum power for radios.
- Routing metric should take this into account! Affects throughput

ETX

- Measure each link's delivery probability with broadcast probes (& measure reverse)
- P(delivery) = (df * dr)   (ACK must be delivered too…)
- Link ETX = 1 / P(delivery)
- Route ETX = Σ link ETX
- (Assumes all hops interfere - not true, but seems to work okay so far)

ETX: Sanity Checks

- ETX of perfect 1-hop path: 1
- ETX of 50% delivery 1-hop path: 2
- ETX of perfect 3-hop path: 3

- (So, e.g., a 50% loss path is better than a perfect 3-hop path! A threshold would probably fail here…)

ETT

- What if links @ different rates?
- ETT – expected *transmission time*
  – ETX / Link rate
  = 1 / (P(delivery) * Rate)

SampleRate

- What is best rate for link?
  – The one that maximizes ETT for the link!
  – SampleRate is a technique to adaptively figure this out. (See new Roofnet paper)

ETX measurement results

- Delivery *is* probabilistic
  - A 1/r^2 model wouldn't really predict this!
  - Sharp cutoff (by spec) of "good" vs "no" reception. Intermediate loss range band is just a few dB wide!
- Why?
  - Biggest factor: Multi-path interference
    - 802.11 receivers can suppress reflections < 250ns
    - Outdoor reflections delay often > 1 \mu sec
    - Delay offsets == symbol time look like valid symbols (large interferece)
    - Offsets != symbol time look like random noise
    - Small changes in delay == big changes in loss rate

Take home points

- Value of implementation & measurement
  - Simulators did not "do" multipath
    - Routing protocols dealt with the simulation environment just fine
    - Real world behaved differently and *really* broke a lot of the proposed protocols that worked so well in simulation!
- Rehash: Wireless differs from wired…
- Metrics: Optimize what matters; hop count often a very bad proxy in wireless
- What we didn't look at: routing protocol overhead
  - One cool area: Geographic routing
  - See extra reading listed on Web page. ☺