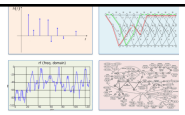

INTRODUCTION TO EECS II

DIGITAL COMMUNICATION SYSTEMS

Network Routing: II

Lecture 19

6.02 Fall 2010

November 15, 2010

- Link-state routing
- Routing "around" failures: HELLO protocol & periodic advertisements/integration
- Things that can go wrong: routing loops

Link-State Routing

- Advertisement step:
  - Information about its links to its neighbors
  - Neighbors re-send on *their* links → **flooding**
  - Result: Each node discovers map of the network
- Integration: Each node runs the same shortest path algorithm over its map
  - If each node implements computation correctly and each node has the same map, then routing tables will be correct

Integration Step: Dijkstra's Algorithm (Example)

Suppose we want to find paths from A to other nodes

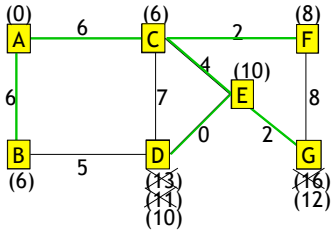

Link-State Advertisements and Flooding

- Periodically send LSA (Link-State Advertisement) [Node, seq#, [(nbr1, linkcost1), (nbr2, linkcost2), …] to all neighbors
- If seq > last_heard:
  - save seq, LSA; rebroadcast LSA to neighbors
- LSAs aren't sent reliably (i.e., no ACKs/retries)
- Periodic messages handle dynamism: state in each node is "soft" and *times out if not refreshed*

[F, seq, (G, 8), (C, 2)]

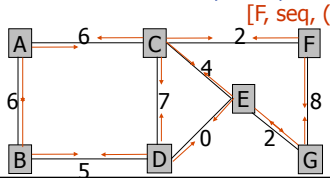

Distance Vector: Pros and Cons

- + Simple protocol
- + Works well for small networks
- - Works only on small networks

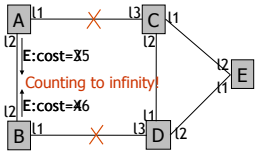

But what if A had sent advert. to B before B sends advert to A?

Suppose link AC fails.
When A discovers failure, it sends E: cost = INFINITY to B.
B sends 'E: cost=2' to A
A installs E: cost=3 in table.
Now suppose link BD fails.
B discovers it, then installs E: cost = INFINITY in its table.
Sends info to A, A installs E: cost = INFINITY in its table

Distance Vector Cons (Cont.)

- Solving counting to infinity problem needs "infinity" to be small
  - Otherwise, convergence time too long
- One approach: *split horizon* advertisements
  - Don't advertise route to X if you got the installed route from X
  - Better still, advertise a cost of INFINITY to whoever advertised the route to you (so there's no chance they will use the route you advertise)

Path Vector Routing

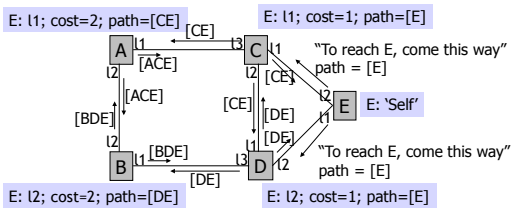

- For each advertisement, run "integration step"
  - E.g., pick shortest, cheapest, quickest, etc.
- **Ignore advertisements with own address in path vector**
  - Avoids routing loops that "count to infinity"

Routing Loop in Link State Protocol

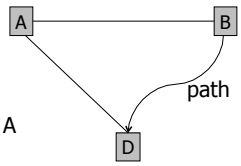

B to D is via A
Link AD fails
A's LSA to B is lost
A now uses B to get to D
But B continues to use A
Routing loop!

Summary

- The network layer implements the "glue" that achieves connectivity
- Forwarding entails a routing table lookup; the table is built using routing protocol
- Distance-vector protocol: distributes route computation; each node advertises path cost if its best route to neighbors
- Link-state protocol: floods neighbor information; centralizes route computation using shortest-path algorithm at each node