

Link is Down at Time 4?



Paritioned Network at Time 10



used you for their route.

6.02 Spring 2011

Lecture 20, Slide #3

Count to Infinity



6.02 Spring 2011

Path-vector (PV) routing



Fixing "Count to Infinity"

- Problem
 - Node C's route to A breaks. C sets cost to ∞
 - But at next round of advertisements, hears of lower-cost routes from neighbors, not know the neighbor's routes used C itself to get to A.
- Solution
 - In addition to reporting costs in advertisements, also report routing path as discovered incrementally by Bellman-Ford
 - Called "path-vector"
 - Modify Bellman-Ford update with new rule: nodes should ignore advertised routes that contain itself in the routing path
 - Pros: count-to-infinity "problem" is solved (routing tables eliminate routes to unreachable nodes more quickly)
 - Cons: advertisement overhead is larger

6.02 Spring 2011

Lecture 20. Slide #6

Paritioned Network at Time 10: PV



6.02 Spring 2011

Link-State Routing

- Advertisement step
 - Send information about its <u>links</u> to its neighbors (aka *link state advertisement* or LSA):

[seq#, [(nbhr1, linkcost1), (nbhr2, linkcost2), ...]

- Do it periodically (liveness, recover from lost LSAs)
- Integration
 - If seq# in incoming LSA > seq# in saved LSA for source node: update LSA for node with new seq#, neighbor list rebroadcast LSA to neighbors (→ flooding)
 - Remove saved LSAs if seq# is too far out-of-date
 - Result: Each node discovers current map of the network
- Build routing table
 - Periodically 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

6.02 Spring 2011

Lecture 20, Slide #9



LSA Flooding

- LSA travels each link in each direction
 - Don't bother with figuring out which link LSA came from
- Termination: each node rebroadcasts LSA exactly once
- · All reachable nodes eventually hear every LSA

Finding shortest paths from A:

B: [(A,19), (C,11), (D, 4)] C: [(A, 7), (B,11), (D,15), (E, 5)]

D: [(B, 4), (C,15), (E,13)] E: [(C, 5), (D,13)]

A: [(B,19), (C, 7)]

- Time required: number of links to cross network

6.02 Spring 2011

LSAs:

Lecture 20, Slide #10

Dijkstra's Shortest Path Algorithm

- Initially
 - nodeset = [all nodes] = set of nodes we haven't processed
 - spcost = {me:0, all other nodes: ∞ } # shortest path cost
 - routes = {me:--, all other nodes: ?} # routing table
- while nodeset isn't empty:
 - find u, the node in nodeset with smallest spcost
 - remove u from nodeset
 - for v in [u's neighbors]:
 - d = spcost(u) + cost(u,v) # distance to v via u
 - if d < spcost(v): # we found a shorter path!
 - spcost[v] = d
 - routes[v] = routes[u] (or if u == me, enter link from me to v)
- Complexity: N = number of nodes, L = number of links
 - Finding u (N times): linear search=O(N), using heapq=O(log N)
 - Updating spcost: O(L) since each link appears twice in neighbors

6.02 Spring 2011

Dijkstra Example



	_											
Step	и	Nodeset	spcost					route				
			Α	B	С	D	Ε	Α	B	С	D	Ε
0		[A,B,C,D,E]	0	œ	8	x	x		?	?	?	?
1	А	[B,C,D,E]	0	19	7	x	8		LO	L1	?	?
2	С	[B,D,E]	0	18	7	22	12		L1	L1	L1	L1
3	Е	[B,D]	0	18	7	22	12		L1	L1	L1	L1
4	В	[D]	0	18	7	22	12		L1	L1	L1	L1
5	D	[]	0	18	7	22	12		L1	L1	L1	L1

Why is Network Routing Hard?

- Inherently distributed problem
 - Information about links and neighbors is local to each node, but we want global reach
- Efficiency: want reasonably good paths, and must find them without huge overhead
- Handling failures and "churn"
 - Must tolerate link, switch, and network faults
 - Failures and recovery could be arbitrarily timed, messages could be lost, etc.
- Scaling to large size very hard (later courses)
 - And on the Internet, many independent, competing organizations must cooperate
 - Mobility makes the problem harder

```
6.02 Spring 2011
```

Lecture 20, Slide #13

Hierarchical Routing



- Internet: collection of domains/networks
- Inside a domain: Route over a graph of routers
- · Between domains: Route over a graph of domains
- Address: concatenation of "Domain Id", "Node Id"

6.02 Spring 2011

٠

Lecture 20, Slide #14

Pros and Cons

Advantages

- Scalable
 - Smaller tables
 - Smaller messages
- Delegation
 - Each domain can run its own routing protocol

Disadvantages

- Mobility is difficult
 - Address depends on geographic location
- Sup-optimal paths
 - E.g., in the figure, the shortest path between the two machines should traverse the yellow domain. But hierarchical routing goes directly between the green and blue domains, then finds the local destination → path traverses more routers.

Summary

- The network layer implements the "glue" that achieves connectivity
 - Does addressing, forwarding, and routing
- Forwarding entails a routing table lookup; the table is built using routing protocol
- DV protocol: distributes route computation; each node advertises its best routes to neighbors
- LS protocol: distributes (floods) neighbor information; centralizes route computation using shortest-path algorithm