

INTRODUCTION TO EECS II
**DIGITAL
COMMUNICATION
SYSTEMS**

6.02 Spring 2012 Lecture #21

- Link-state routing
- Routing around failures

6.02 Spring 2012
Lecture 21, Slide #1

Link-State Routing

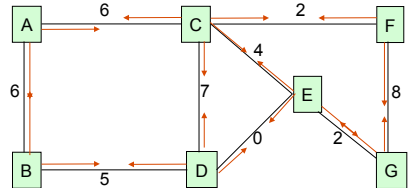
- Advertisement step
 - Send information about its links to its neighbors (aka **link state advertisement** or LSA):

[node, seq#, [(nbr1, linkcost1), (nbr2, 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 - Remove saved LSAs if seq# is too far out-of-date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 (e.g., Dijkstra's alg)
 - If each node implements computation correctly and each node has the same map, then routing tables will be correct

6.02 Spring 2012
Lecture 21, Slide #2

LSA Flooding

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

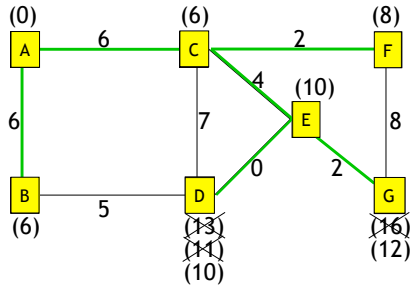


- Periodically originate LSA
- LSA travels each link in each direction
 - Don't bother with figuring out which link LSA came from
- Termination: each node rebroadcasts LSA exactly once
 - Use sequence number to determine if new, save latest seq
- Multiple opportunities for each node to hear any given LSA
 - Time required: number of links to cross network

6.02 Spring 2012
Lecture 21, Slide #3

Integration Step: Dijkstra's Algorithm (Example)

Suppose we want to find paths from A to other nodes



6.02 Spring 2012
Lecture 21, Slide #4

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)

6.02 Spring 2012


Lecture 21, Slide #5

Another Example

Finding shortest paths from A:

LSAs:

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



Step	u	Nodeset	spcost					route				
			A	B	C	D	E	A	B	C	D	E
0		[A,B,C,D,E]	0	∞	∞	∞	∞	--	?	?	?	?
1	A	[B,C,D,E]	0	19	7	∞	∞	--	L0	L1	?	?
2	C	[B,D,E]	0	18	7	22	12	--	L1	L1	L1	L1
3	E	[B,D]	0	18	7	22	12	--	L1	L1	L1	L1
4	B	[D]	0	18	7	22	12	--	L1	L1	L1	L1
5	D	[]	0	18	7	22	12	--	L1	L1	L1	L1

6.02 Spring 2012

Lecture 21, Slide #6


Failures

- Problems: Links and switches could fail
 - Advertisements could get lost
 - Routing loop
 - A sequence of nodes on forwarding path that has a cycle (so packets will never reach destination)
 - Dead-end: route does not actually reach destination
 - Loops and dead-ends lead to routes not being valid
- Solution
 - HELLO protocol to detect neighbor liveness
 - Periodic advertisements from nodes
 - Periodic integration at nodes
 - Leads to eventual convergence to correct state (see Chapter 18)

6.02 Spring 2012

Lecture 21, Slide #7

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!
 Must wait for eventual arrival of correct LSAs to fix loop

6.02 Spring 2012

Lecture 21, Slide #8

Distance-Vector: Pros, Cons, and Loops

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



Suppose link AC fails.
 When A discovers failure, it sends E: cost = INFINITY to B.
 B advertises E: cost=2 to A
 A sets E: cost=3 in its table

Now suppose link BD fails.
 B discovers it, then sets E: cost = INFINITY.
 Sends info to A, A sets E: cost = INFINITY.

But what if A had advertised to B before B advertised to A?

6.02 Spring 2012

Lecture 21, Slide #9

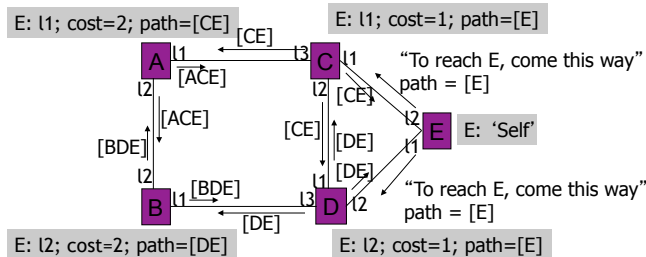
Fixing “Count to Infinity” with Path Vector Routing

- In addition to (or instead of) reporting costs, advertise the *path* discovered incrementally by the Bellman-Ford update rule
- Called “path-vector”
- Modify Bellman-Ford update with new rule: a node should ignore any advertised route that contains itself in the advertisement

6.02 Spring 2012

Lecture 21, Slide #10

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”

6.02 Spring 2012

Lecture 21, Slide #11

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
 - Path-vector: include path, not just cost, in advertisement to avoid “count-to-infinity”
- LS protocol: distributes (floods) neighbor information; centralizes route computation using shortest-path algorithm

6.02 Spring 2012

Lecture 21, Slide #12