


1 Distance-Vector Routing



The nodes in the above network share routes with each other using distance-vector routing. Below are the initial routing tables for each node, and a table showing the costs for each of their neighboring links. Links without a distance provided have an implicit distance of 1.

S1	
Dest.	NextHop, Dist.
A	Direct, 1

S2	
Dest.	NextHop, Dist.
B	Direct, 1

S3	
Dest.	NextHop, Dist.
C	Direct, 1

S4	
Dest.	NextHop, Dist.
D	Direct, 1

The following questions indicate events that happen consecutively. You can assume that no other events occur other than the ones specified. Note that all blanks may not be necessary.

EVENT: *S3 advertises its routes to S1 and S4.*

(1) What do the routing tables for S1 and S4 look like after receiving S3's routes? (You may not need to fill in all the rows)

S1	Dest.	NextHop, Dist.
	A	Direct, 1
	C	S3, 2

S4	Dest.	NextHop, Dist.
	D	Direct, 1
	C	S3, 3

(2) Which nodes among S1 and S4 are expected to advertise their routes after receiving S3's routes? (Assuming the advertising routes on a routing table change optimization is being used)

Solution: Since both S1 and S4 updated their shortest paths, they will both advertise their routing tables.

EVENT: S1 advertises its routes to S2, S3, and S4.

- (3) What do the routing tables for S2, S3, and S4 look like after receiving S1's routes? (You may not need to fill in all columns)

	Dest.	NextHop, Dist.
S2	B	Direct, 1
	A	S1, 3
	C	S1, 4

	Dest.	NextHop, Dist.
S3	C	Direct, 1
	A	S1, 2

	Dest.	NextHop, Dist.
S4	D	Direct, 1
	C	S3, 3
	A	S1, 8

EVENT: S4 advertises its routes to S1 and S3. **EVENT:** S1 advertises its routes to S2 and S4.

- (4) At this point, what path does S2 use to reach D, and what is the cost?

Solution: S2 → S1 → S4 with a cost of 10.

S1 has only heard about a route to D from S4. S1 in turn advertises this route to S2.

EVENT: S3 advertises its routes to S1 and S4.


- (5) What do the routing tables for S1 and S2 look like now?

	Dest.	NextHop, Dist.
S1	A	Direct, 1
	C	S3, 2
	D	S3, 4

	Dest.	NextHop, Dist.
S2	B	Direct, 1
	A	S1, 3
	C	S1, 4
	D	S1, 10

Solution: S2 has not heard about S3's routing update to S1 about a lower cost route to D yet

2 Split Horizon and Poisoned Reverse



All **unlabeled** links have a cost of 1. The parts of the question do *not* build on each other.

- (1) Assume that the routers use **split horizon**. Say that S4 advertises (A: 2, C: 1) to S3. Assuming that S3 has received no other advertisements, what does S3 now tell S4 about S3's path to A?

Solution: Nothing. Split Horizon means that we never tell a neighbor about paths that go through that neighbor. So in this case, S3 doesn't tell S4 about its path to A.

- (2) Assume that the routers use **poisoned reverse**. Routing tables have not converged and S3 believes its shortest path to A is through S1 (this path is S3-S1-S2 of length 4). S3 advertises its routes to S4. Now,

S4 advertises to S3. S4 bases this advertisement off of its routing table which has: (B: 2, A: 2, C : 1). After recomputing its routes, S3 advertises its routes to S4. What is the advertised distance to A?

Solution: S3 will tell S4 that its distance to A is infinitely long, because S3's new shortest route goes through S4.

- (3) Consider the simple topology (A-S1-S2-S3). After the routing tables have converged, link S1-S2 goes down. When S2 advertises to S3 (A: ∞), is this an act of **poisoning a route** or **poisoned reverse**?

Solution: A is **poisoning a route**. Namely, it tells C that its distance is ∞, not because A's new path goes through C, but because A actually has no route now.

- (4) **Poisoning a route** and **poisoned reverse** might sound similar, but actually we can think of one of them as being “honest” while the other one is “lying.” Which one tells the truth, and which one tells a white lie to keep the network functioning?

Solution: **Poisoned reverse** encourages routers to tell a white lie. With poisoned reverse, we tell a neighbor that we have no path to a certain destination if our path goes through that neighbor. Since we actually do have a path, our message is not strictly true. On the other hand, **poisoning a route** happens when a link goes down, and we actually lose our path to some destination. Thus, we're telling the truth when we advertise a distance of ∞ to this destination (given that a infinitely long path is equivalent to no path).

3 Count to Infinity (Guided Q)



- (1) For part 1 of this question there is **no** split-horizon or poisoned reverse, and advertisements are only sent periodically (aka when it is explicitly stated).

- (a) What do the routing tables look like once S1, S2 and S3 converge?

S1		S2		S3	
Dest.	NextHop, Dist.	Dest.	NextHop, Dist.	Dest.	NextHop, Dist.
A	Direct, 1	A	S1, 2	B	Direct, 1
B	S2, 3	B	S3, 2	A	S2, 3

- (b) What periodic advertisement will S1 and S2 send to each other? (One such message is given as an example)

From	To	(Destination, Distance)
S1	S2	(A, 1)
S1	S2	(B, 3)
S2	S1	(A, 2)
S2	S1	(B, 2)

EVENT: The link between S2 and S3 goes down.

- (c) What will S1 and S2 send to each other?

From	To	(Destination, Distance)
S1	S2	(A, 1)
S1	S2	(B, 3)
S2	S1	(A, 2)
S2	S1	(B, 2)

Solution: (The tables did not change because no routes have expired yet.)

EVENT: S2's route to B *finally* expires

- (d) After S1 and S2 exchange advertisements again, what will their routing tables look like?

S1		S2	
Dest.	NextHop, Dist.	Dest.	NextHop, Dist.
A	Direct, 1	A	S1, 2
B	S2, 3	B	S1, 4

EVENT: S1's route to B expires

- (e) After S1 and S2 exchange advertisements again, what will their routing tables look like?

S1		S2	
Dest.	NextHop, Dist.	Dest.	NextHop, Dist.
A	Direct, 1	A	S1, 2
B	S2, 5	B	S1, 4

- (f) Is this good? **Solution: NO! This is called count to infinity. Both switches think they have a path to B for a long time after the path ceases to exist.**

- (2) For part 2 of this question there is **split-horizon**, but **no** poisoned reverse and advertisements are only sent periodically (aka when it is explicitly stated). Also, all dropped links are back up and the routing state starts out converged!

- (a) What will S1 and S2 send to each other after everything has converged?

From	To	(Destination, Distance)
S1	S2	(A, 1)
S2	S1	(B, 2)

EVENT: The link between S2 and S3 goes down.

- (b) What will S1 and S2 send to each other?

From	To	(Destination, Distance)
S1	S2	(A, 1)
S2	S1	(B, 2)

EVENT: S2's route to B *finally* expires

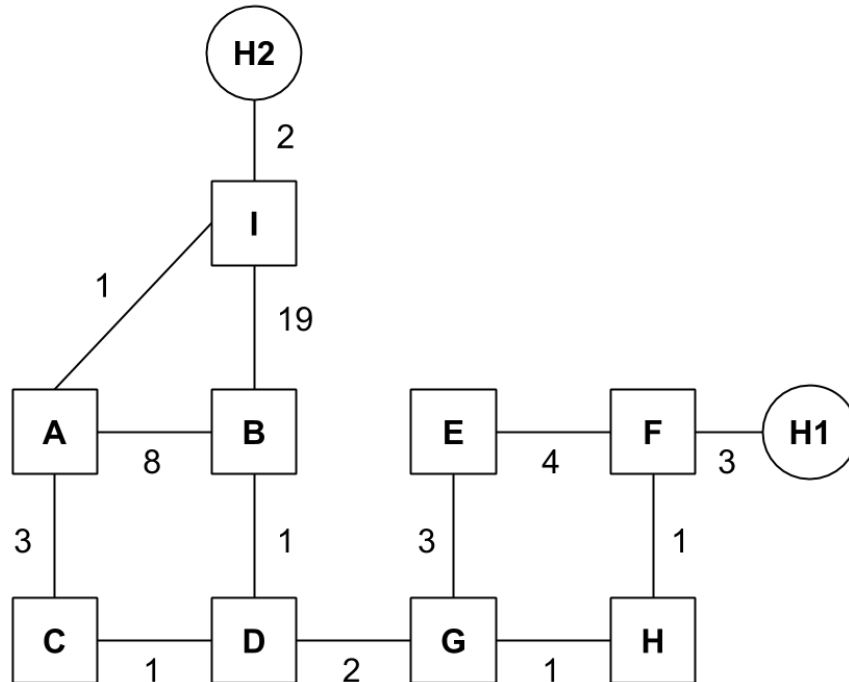
- (c) After S1 and S2 exchange advertisements again, what will their routing tables look like?

S1		S2	
Dest.	NextHop, Dist.	Dest.	NextHop, Dist.
A	Direct, 1	A	S1, 2
B	S2, 3		

(d) Will this end well?

Solution: Yep! S1's route to B will expire because it has not been updated in a while.

4 Link-State Routing



For this problem, assume the network is running a link state routing protocol, minimizing total route latency. The following questions indicate events that happen consecutively.

(1) After convergence, what route does **Router D** think its packet will take to **Host H1**?

Solution: D → G → H → F → H1

EVENT: Link GH goes down.

(2) **Router G** and **Router H** have recomputed their routes, but have not yet sent updates to other routers. What route does **Router D** think its packet will take to **Host H1**?

Solution: Same, D → G → H → F → H1

(3) What route does it actually take?

Solution: D → G → E → F → H1

(4) Assume all nodes are now aware of the new network state and have recomputed their routes. What route does a packet take from **Router B** to **Host H2**?

Solution: B → D → C → A → I → H2

EVENT: *The cost of link AI increases to 100.*

- (5) **Router A** and **Router I** recompute their routes, but have not yet sent updates to other routers. What route does **Router A** think its packet will take to **Host H2**?

Solution: A → C → D → B → I → H2

- (6) What route does it actually take?

Solution: A → C → A (*loop*)


- (7) Which additional routers must receive the routing updates and recompute their routes for all routers to be able to successfully send packets to **Host H2**?

Solution: **Router B, Router C, Router D**

- (8) All routers except **Router B** have received the routing updates and recomputed their routes. Which routers can successfully send packets to **Host H2**?

Solution: None of the routers (except **Router I**)

5 L3 Link State Routing Extra Problem



Suppose we have an L3 network with the topology shown above, and the routing algorithm used is link state.

1. After convergence, what is the path cost from H1 to H2, and what are all the possible paths with this cost?

Solution: The cost is 15. All paths are $H1 \rightarrow S \rightarrow [A, B, C] \rightarrow D \rightarrow [E, F] \rightarrow T \rightarrow H2$.

2. Suppose that a control message (a message used by the routing algorithm) takes 1 second to propagate along a link, regardless of link cost. What individual link failure inside the network would cause the longest delay to reconvergence, and what is that delay?

Solution: A link SB failure would cause a 3 second delay. The last node that learns about the failure is node T, and the update must traverse through at least 3 links to get to T, which will take 3 seconds. Any other link failure within the network can propagate to every other node in less than 3 seconds. Host links could also take 3 seconds to propagate.

3. Suppose you have the ability to take down individual nodes. Which nodes would you take down in order to partition the network? If you can't partition the network, which nodes would you take down to increase path costs from H1 to H2 maximally? In each part, suppose you can take down:
 - (a) a single node (excluding S and T).

Solution: D. Can't partition, and removing D increases path costs the most (by 1).

Solution: E and F. Removing them will partition T from the network.

4. Which single link's cost (if any) should you double in order to increase the path cost from H1 to H2?

Solution: Doubling the cost of H1-S or T-H2 will increase the total path cost by 1. It's not possible with any other link, since there is always an alternate, cheaper path (see part 1).