Lecture 5.3 - Link-State Routing Exercises ANS

1 Link-State Routing

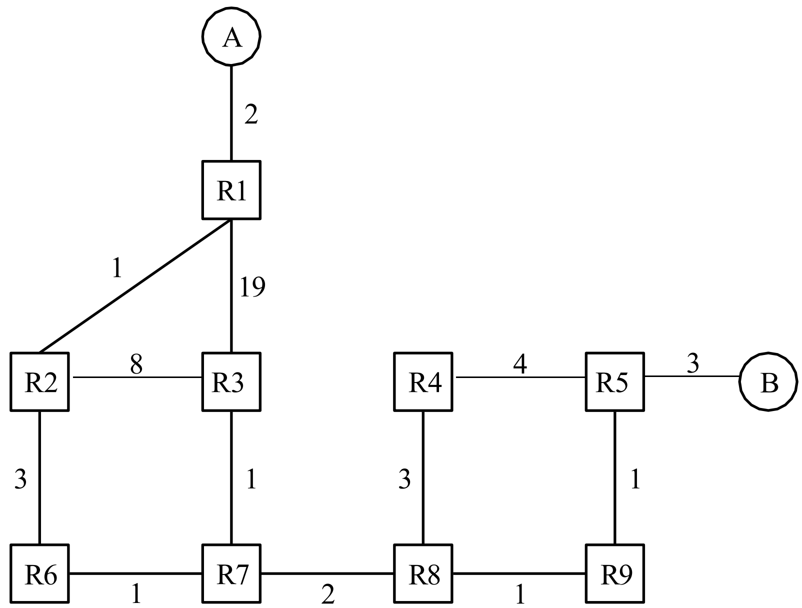

For this problem, assume the network is running a link-state routing protocol, minimizing total route latency. The following questions indicate events that happen consecutively.

1.1 After convergence, what route does R7 think its packet will take to Host B?

R7 → R8 → R9 → R5 → B

EVENT: *Ge R8-to-R9 link goes down.*

1.2 R8 and R9 have recomputed their routes but have not yet sent updates to other routers. What route does R7 think its packet will take to Host B?

Same, R7 → R8 → R9 → R5 → B

1.3 What route does it actually take?

R7 → R8 → R4 → R5 → B

1.4 Assume all nodes are now aware of the new network state and have recomputed their routes. What route does a packet take from R3 to Host A?

R3 → R7 → R6 → R2 → R1 → A

EVENT: *Ge cost of the R1-to-R2 link increases to 100.*

1.5 R2 and R1 recompute their routes but have not yet sent updates to other routers. What route does R2 think its packet will take to Host A?

R2 → R6 → R7 → R3 → R1 → A

1.6 What route does it actually take?

R2 → R6 → R2 *(loop) ()*

1.6 Justification: Although R2 has updated its own routing table, R6 still believes that R2 is the best next hop to reach Host A. This inconsistent state causes packets forwarded from R2 to R6 to be sent back to R2, forming a transient routing loop.

1.7 Which additional routers must receive the routing updates and recompute their routes for all routers to be able to successfully send packets to Host A?

R3, R6, R7

1.7 Justification: R3, R6, R7 must receive routing updates. If R7 has not received routing update of R1-to-R2 link cost 100, then it still thinks its shortest path to A is R7-R6-R2-R1. But R6 thinks its shortest path to A is R6-R7-R3-R1, so there is a routing loop R6-R7. Even if routers R4, R5, R8, R9 did not receive routing updates, they can still send to A, since all packets go though R7.

1.8 All routers except R3 have received the routing updates and recomputed their routes. Which routers can successfully send packets to Host A?

None of the routers (except R1)

1.8 Justification: Because R3 has not yet recomputed its routes, it continues to forward packets toward Host A through R2 using stale link-state information based on the old R1–R2 link cost. Other routers forward traffic toward A through R3 (based on updated R1–R2 link cost). When those packets reach R3 they are forwarded to R2, forming a routing loop R2-R6-R7-R3. Inconsistent local views during partial convergence mean that reachability to A from most routers is broken; only R1, which is directly connected to Host A, is guaranteed to successfully deliver packets to A.

2 More L3 Link State

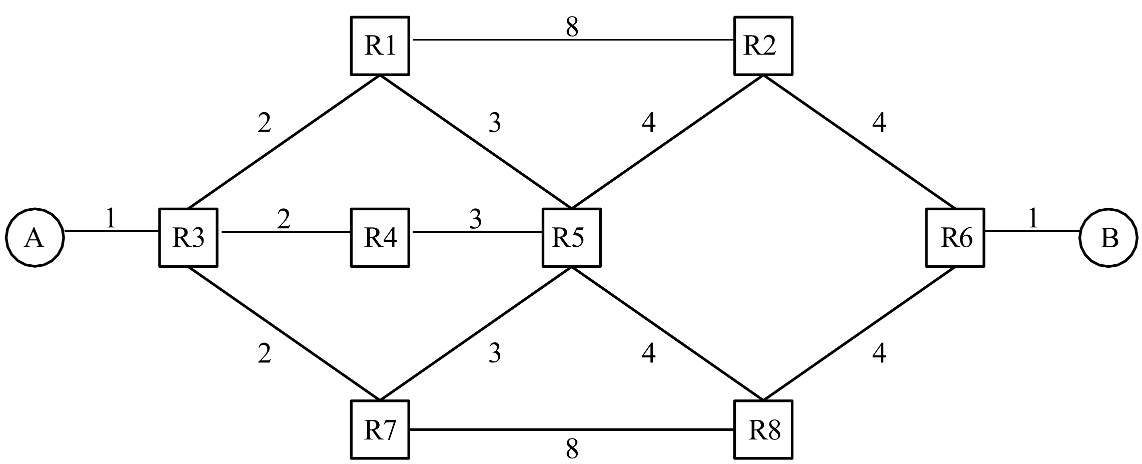

2.1 After convergence, what is the path cost from A to B, and what are all the possible paths with this cost?

The cost is 15. All paths are A → R3 → [R1, R4, R7] → R5 → [R2, R8] → R6 → B.

2.2 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?

ANS: A failure of the R3–R4 link results in the longest delay because the resulting link-state update must traverse the greatest number of hops to reach the furthest router, R6. Since control messages take one second per hop regardless of link cost, this produces a maximum delay of three seconds.

2.3 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 A to B maximally?

In each part, suppose you can take down:

(a) A single node (excluding R3 and R6).

ANS: R5. R5 lies on all shortest paths between A and B, making it a critical articulation point for low-cost routing. Removing R5 does not fully partition the network but forces all remaining paths to detour, increasing the total path cost (by 1).

(b) Two nodes (excluding R3 and R6).

ANS: R2 and R8. R2 and R8 together form the only connections between R6 and the rest of the network. Removing both nodes disconnects R6 entirely, thereby partitioning the network.

2.4 Which single link's cost (if any) should you double in order to increase the path cost from A to B?

ANS: Doubling the cost of A-R3 or R6-B will increase the total path cost by 1. It's not possible with any other link, since there is always an alternate, cheaper path.

Q3. Link-State

Consider the following network topology. All routers are running the link-state protocol. Each subpart is independent unless otherwise stated.

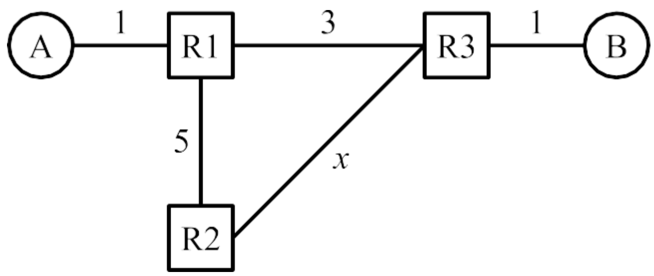

Q3.1 In total, how many "hello" advertisements get sent between the routers? Do not count periodically re-sent advertisements.

ANS: 6. R1 sends a hello to R2, R3.

R2 sends a hello to R1, R3.

R3 sends a hello to R1, R2.

Note that hosts are not counted here, since the routing protocol is performed by the routers, not the hosts.

Q3.2 Assume R1 doesn't know about the R1-to-R3 link. All other routers know the full topology. If __________, then a packet sent from A to B would get stuck in a routing loop. Fill in the blank with a strict inequality

ANS: $x > 8$. R1 is unaware of the R1-to-R3 link, so it forwards A-to-B packets to R2. R2 knows the full topology. If R2 forwards via R1, the cost to B is $5 + 3 + 1 = 9$. If R2 forwards via R3, the cost to B is $x + 1$. A loop forms if R2 decides to forward the packet back to R1. This occurs if the cost via R1 is less than the cost via R3. Writing this as an inequality, we get $9 < x + 1$, hence $x > 8$.

Q3.3 What feature of the link-state protocol will help update the forwarding table to remove these routing loops?

A) Routers send poison advertisements.
B) Packets have a TTL field, and the packet is dropped when the TTL reaches 0.
C) Advertisements are periodically re-sent.
D) None of the above. The routing loops in the above scenarios will stay forever.

ANS: C). Option A is false, because poison advertisements are not used at all in the link-state protocol. Option B is false, as the TTL field does not explicitly update the forwarding table. Option C is correct. Eventually advertisements will be re-sent and all routers will learn the correct network topology.

Q3.5 What is a possible cause of the scenario where R1 doesn't know about the R1-to-R3 link?

A) Advertisement(s) get dropped.
B) User packets get dropped.
C) R1 is not running a correct shortest-path algorithm.
D) The scenario where R1 doesn't know about a link will never happen in the link-state protocol

ANS: A). The advertisement from R3 to R1 could've been dropped, which leads to R1 not knowing about its link to R3. Option B is false, as user packets don't affect routing. Option C is false, as the shortest-path algorithm is a separate step from learning about the network topology using advertisements.

Q3.6 Suppose a routing loop has formed between R2 and R3. R2 forwards a user packet to R3, and then the same packet is sent back to R2. What will R2 do as a result of receiving this packet?

A) Forward the packet to R3, and remove an entry from its forwarding table.
B) Forward the packet to R3, and add an entry to its forwarding table.
C) Forward the packet to a different router (not R3), and leave the forwarding table unchanged.
D) Forward the packet to R3, and leave the forwarding table unchanged.

ANS: D). Routers use destination-based forwarding. If R2 forwards a packet to R3 and then receives that packet again, R2 will still make the same forwarding decision, since the destination of that packet is still the same.