Lectures 5.2 5.3 Distance Vector & Link State Routing Quiz ANS

1. What is the Bellman-Ford equation used for?
A) Calculating the bandwidth of a link
B) Computing the shortest path in Distance-Vector routing
C) Encrypting packets for secure transmission
D) Converting IPv4 addresses to IPv6

ANS: B
The Bellman-Ford equation ($d_x(y) = \min \{ c(x,v) + d_v(y) \}$) is the foundation of Distance-Vector routing, allowing a node to calculate the least-cost path to a destination by using neighbor costs.

2. In Distance-Vector routing, what information does a node send to its neighbors?
A) The entire network topology map
B) Only the cost to its direct neighbors
C) Its entire distance vector (estimates to all destinations)
D) The list of all packets it has received

ANS: C
Each node shares its own distance vector (its estimated costs to all known destinations) with its immediate neighbors. It does NOT share the full topology or path details.

3. What triggers a node to send a routing update in a Distance-Vector protocol?
A) Only when a timer expires
B) Only when the administrator manually reboots the router
C) When a local link cost changes or it receives an update from a neighbor
D) Every time it forwards a packet

ANS: C
Distance-Vector is iterative and asynchronous. Updates are triggered by local link cost changes or by receiving a changed distance vector from a neighbor.

4. What is the 'Count-to-Infinity' problem?
A) A loop where routing updates bounce back and forth, incrementing costs indefinitely
B) When a router runs out of memory
C) When a packet takes too many hops to reach the destination
D) When the network bandwidth is infinite

ANS: A
Count-to-Infinity happens when bad news (like a link failure) propagates slowly. Routers believe they can reach the destination through each other, creating a loop where the cost keeps increasing.

5. Which technique prevents a routing loop between two immediate neighbors?

A) Poisoned Reverse
B) Infinite Bandwidth
C) Manual Configuration
D) Using IPv6

ANS: A
Poisoned Reverse prevents 2-node loops. If Z routes through Y to reach X, Z tells Y that its distance to X is infinity, so Y won't try to route back through Z to get to X.

6. Does Poisoned Reverse solve all counting-to-infinity loops?
A) Yes, it solves all loops permanently
B) No, it only works for loops involving 2 nodes
C) No, it only works for loops involving 3 or more nodes
D) It creates more loops

ANS: B
Poisoned Reverse works for loops between two immediate neighbors. It cannot detect or prevent larger loops involving three or more nodes (e.g., A -> B -> C -> A).

7. What defines the 'state' of a node in Distance-Vector routing?
A) Its physical location
B) Its distance vector (costs to all destinations) and its forwarding table
C) The amount of RAM it has
D) The number of cables plugged in

ANS: B
The node's state consists of its own distance vector estimates (D_x), the vectors it received from neighbors (D_v), and its local link costs ($c(x,v)$).

8. How does 'Good News' propagate in Distance-Vector routing?
A) Very slowly (count-to-infinity)
B) Fast
C) It never propagates
D) It causes loops

ANS: B
Good news (e.g., a link cost decreasing) travels fast because the new lower cost is immediately advertised and accepted by neighbors.

9. How does 'Bad News' propagate in Distance-Vector routing?
A) Very fast
B) Slowly
C) Instantly
D) It is ignored

ANS: B

Bad news (e.g., link failure or cost increase) travels slowly. Routers may cling to outdated information from neighbors (the 'Count-to-Infinity' problem), delaying convergence.

10. What is a key disadvantage of Distance-Vector routing?

A) It requires a central server

B) It suffers from slow convergence and routing loops

C) It uses too much bandwidth for good news

D) It cannot handle more than 5 routers

ANS: B

The main drawbacks are slow convergence (due to bad news propagation) and the risk of transient routing loops during that convergence time.

11. In Distance-Vector, does a node know the full network topology?

A) Yes, it has a complete map

B) No, it only knows costs to neighbors and estimates to destinations

C) Yes, but only for Class A networks

D) No, it knows nothing at all

ANS: B

DV is decentralized. A node knows its own local links and the vectors shared by neighbors, but it does NOT have a global map of the topology.

12. What does RIP stand for?

A) Routing Information Protocol

B) Real-time Internet Protocol

C) Rapid Inter-domain Protocol

D) Rest In Peace

ANS: A

RIP (Routing Information Protocol) is a classic example of a Distance-Vector protocol used in older or smaller networks.

13. In Distance-Vector, if neighbor V advertises a cost of 5 to destination Y, and the link cost c(x,v) is 2, what is x's cost to Y via V?

A) 2

B) 5

C) 7

D) 3

ANS: C

The cost is the sum of the link cost to the neighbor plus the neighbor's cost to the destination: 2 + 5 = 7.

14. When does the Distance-Vector algorithm terminate?

A) Never, it runs forever

B) When there are no more updates to send (quiescence)

C) After exactly 10 iterations
D) When the administrator stops it

ANS: B
The algorithm is self-terminating. If no nodes exchange updates (because no costs changed), the algorithm enters a quiet state until the next topology change.

15. Why is Poisoned Reverse not a complete solution?
A) It is too expensive to implement
B) It only prevents loops of size 2, not larger loops
C) It stops forwarding all packets
D) It requires manual configuration

ANS: B
Poisoned reverse only affects the advertisement back to the immediate neighbor. It cannot prevent a loop like X->Y->Z->X.

16. What is the fundamental difference between Link-State (LS) and Distance-Vector (DV)?
A) LS is slower
B) LS provides every router with a complete map of the network topology
C) LS only works for small networks
D) LS does not use costs

ANS: B
In Link-State routing, every router floods information so that all routers build an identical, complete map (graph) of the entire network.

17. Which algorithm is typically used in Link-State routing to compute paths?
A) Bellman-Ford
B) Dijkstra's Algorithm
C) Diffie-Hellman
D) Round Robin

ANS: B
Link-State protocols (like OSPF) use Dijkstra's algorithm to compute the shortest path tree from the source node to all other nodes.

18. What information is flooded in a Link-State protocol?
A) The entire forwarding table
B) Distance vectors to all destinations
C) Link State Advertisements (LSAs) containing neighbors and link costs
D) User data packets

ANS: C
Routers flood Link State Advertisements (LSAs) which describe only their local connectivity (who their neighbors are and the cost to reach them).

19. What is the time complexity of Dijkstra's algorithm with a naive implementation?
A) O(1)
B) O(N^2)
C) O(N log N)
D) O(N)

ANS: B
The standard implementation of Dijkstra's algorithm has a complexity of O(N^2), where N is the number of nodes.

20. How does a router in LS ensure it has the latest topology information?
A) By polling the central server
B) Through Reliable Flooding of LSAs
C) By guessing
D) By rebooting every 5 minutes

ANS: B
Reliable flooding ensures that LSAs from every router reach every other router in the network, allowing them to construct the same graph.

21. What happens if different routers in LS have different views of the topology (inconsistent state)?
A) Nothing bad happens
B) Routing loops can form
C) The network becomes faster
D) Security is improved

ANS: B
If routers disagree on the topology (transient inconsistency), they may compute conflicting paths, leading to temporary routing loops or black holes.

22. In Link-State, does a router send its forwarding table to neighbors?
A) Yes, always
B) No, it only sends link state info (LSAs)
C) Only to the root node
D) Only during startup

ANS: B
Forwarding tables are computed locally. Routers exchange raw topology data (LSAs), not the processed tables.

23. What prevents an LSA from looping forever during flooding?
A) Sequence numbers and checking if the LSA is new
B) It loops forever until power off
C) LSAs have no TTL
D) The administrator deletes them

ANS: A
Routers track LSA sequence numbers. If a router receives an LSA it has already seen (or an older one), it discards it, preventing infinite flooding loops.

24. Why might route oscillation occur in Link-State routing?
A) Because cables are loose
B) If link costs depend on traffic load
C) Because Dijkstra's algorithm is random
D) It never happens

ANS: B
If cost is based on traffic (congestion), routers will switch to a low-cost idle path. That path immediately becomes congested, raising its cost, causing routers to switch back. This creates oscillation.

25. Comparing LS and DV: Which one generates more control traffic initially?
A) DV
B) LS (due to flooding the whole topology)
C) They are exactly the same
D) Neither generates traffic

ANS: B
Link-State requires flooding LSA information from every node to every other node, creating a burst of traffic to build the initial map. DV only talks to neighbors.

26. Comparing LS and DV: Which one typically converges faster?
A) DV
B) LS
C) They are the same
D) DV never converges

ANS: B
Link-State usually converges faster (O(N^2) calculation) once LSAs are flooded. DV suffers from the count-to-infinity problem which can delay convergence significantly.

27. What is the robustness issue with Distance Vector?
A) A single router broadcasting incorrect costs can screw up the whole network
B) It is too secure
C) It requires too much CPU
D) It fails if one cable breaks

ANS: A
In DV, a router tells neighbors its computed path costs. If it lies (e.g., claiming a 0 cost path to everywhere), that error propagates through the network.

28. What is the robustness issue with Link State?
A) A router can advertise incorrect link costs, but only for its own links

B) It crashes the internet
C) It cannot handle large tables
D) It deletes files

ANS: A
In LS, a router can only lie about its own direct links. It cannot fabricate a path cost for a remote part of the network. This makes LS slightly more robust to bad data than DV.

29. Why don't we use Link-State for the entire Internet (Inter-domain)?
A) Dijkstra is too hard to spell
B) Privacy and Scalability
C) It is too fast
D) IS-IS is patented

ANS: B
Link-State requires revealing your internal topology to everyone (Privacy issue) and flooding LSAs to millions of routers (Scalability issue). This is why BGP (Path Vector) is used instead.

30. In Link-State routing, what is Dijkstra's Algorithm used to calculate?
A) The physical distance between routers in kilometers
B) The shortest path tree from the source node to all other nodes in the network
C) The total number of packets dropped in the last hour
D) The encryption key for secure communication

ANS: B
Once a router has the complete network graph (topology), it runs Dijkstra's algorithm locally to compute the shortest (least-cost) path to every other destination, which is then used to populate its forwarding table.