CHAPTER 20

PRACTICE SET

Questions

Q20-1. According to the principle we mentioned in the text, the shortest path is the inverse of the original one. The shortest path is G → E → B → A.

Q20-3. Link-state routing uses Dijkstra's algorithm to first create the shortest-path tree before creating the forwarding table. The algorithm needs to have the complete LSDB to start.

Q20-5. The three ASs described in the text are *stub*, *multihomed*, and *transient*. The first two do not allow transient traffic; the third does. The stub and multihomed ASs are similar in that they are either the sink or source of traffic; the first is connected to only one other AS, but the second is connected to more than one ASs.

Q20-7. The source and destination IP addresses in datagrams carrying payloads between the hosts are the IP addresses of the hosts; the IP addresses carrying routing update packets between routers are IP addresses of the routing interfaces from which the packets are sent or received. This shows that a router needs as many IP addresses as it has interfaces.

Q20-9. Although RIP is running as a process using the service of the UDP, the process is called a *daemon* because it is running all the time in the background. Each router acts both as a client and a server; it acts as a client when there is a message to send; it acts as a server when a message arrives.

Q20-11. OSPF divides an AS into areas, in which routing in each area is independent from the others; the areas only exchange a summary of routing information between them. RIP, on the other hand, considers the whole AS as one single entity.

Q20-13. In RIP, each router just needs to share its distance vector with its neighbor. Since each router has one type of distance vector, we need only one update message. In OSPF, each router needs to share the state of its links with every other router. Since a router can have several types of links (a router link, a network link, …), we need several update messages.

Q20-15. The type of payload can be determined from the value of the protocol field. The protocol field value for ICMP is 01; for OSPF, it is 89.

Q20-17. It cannot. A link needs to be advertised in a router link LSP; a network needs to be advertised in a network link LSP.

Q20-19. BGP is designed to create semi-permanent communication between two BGP speakers; this requires the service of TCP. A connection is made between the two speakers and remains open, while the messages are exchanged between them. UDP cannot provide such a service.

Q20-21. The following shows the use of each attribute:

a. The LOCAL-PREF is used to implement the organization policy.

b. The AS-PATH defines the list of autonomous systems through which the destination can be reached.

c. The NEXT-HOP defines the next router to which the data packet should be forwarded.

Problems

P20-1. We have

$$D_{xy} = \min \{(c_{xa} + D_{ay}), (c_{xb} + D_{by}), (c_{xc} + D_{cy}), (c_{xd} + D_{dy})\}$$

$$D_{xy} = \min \{(2 + 5), (1 + 6), (3 + 4), (1 + 3)\} = \min \{7, 7, 7, 4\} = 4$$

P20-3.

a. The *hello message* (type 1) is used by a router to introduce itself to neighboring routers and to introduce already-known neighboring routers to other neighbors.

b. The *data description message* (type 2) is sent in response to a hello message. A router sends its full LSDB to the newly joined router.

c. The *link-state request message* (type 3) is sent by a router that needs information about a specific LS.

d. The *link-state update message* (type 4) is sent by a router to other routers for building the LSDB. There are five different versions of this message to announce different link states.

e. The *link-state acknowledge message* (type 5) is sent by a router to announce the receiving of a link-state update message. This message is used to provide reliability for the main message used in OSFP.

P20-5. Two nodes, A and D, see the changes. These two nodes update their vectors immediately. We assume that changes in each round are fired in the order A, B, C, D. The following shows that the internet is actually stable after two rounds of updates, but more updates are fired to assure the system is stable. We have shown only three columns for each forwarding table, but RIP usually uses more than columns. Also note that we have used the yellow color to show the changed in cost field, which triggers updates. The cost and the next hop fields participate in updating.

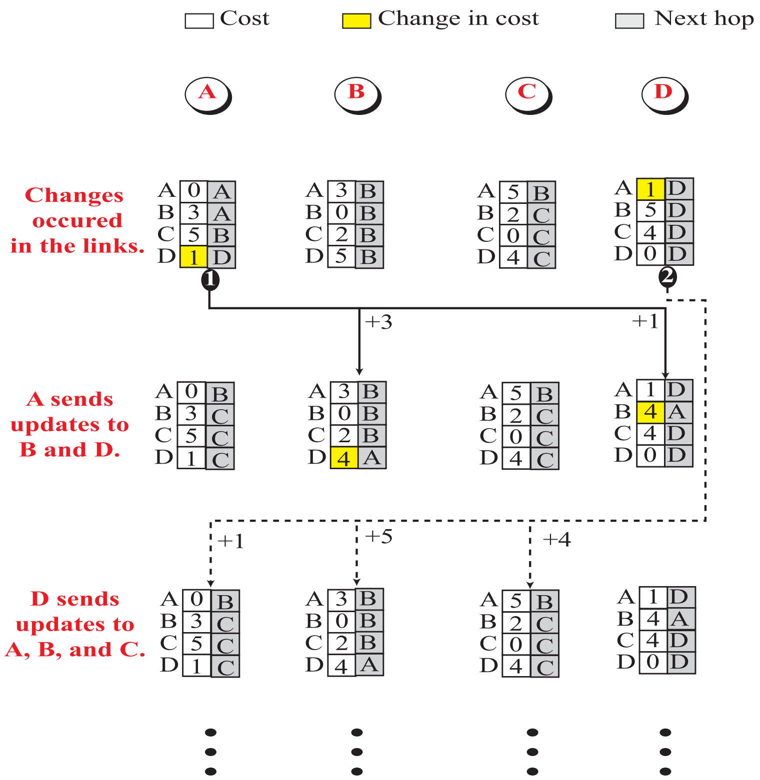

P20-7. The number of operations in each iteration of the algorithm is n, in which n is the number of nodes in the network. In computer science, this complexity is written as $O(n)$ and is referred to as Big-O notation.

P20-9. The following shows the advertisement in each case (a triplet defines the destination, cost, and the next hop):

a. From A to B: (A, 0, A), (C, 4, A).

b. From C to D: (A, 4, C), (C, 0, C).

c. From D to B: (C, 6, D), (D, 0, D).

d. From C to A: (B, 8, D), (C, 0, C), (D, 6, C).

P20-11. The forwarding table for node A can be made using the least-cost tree, as shown below:

Forwarding table for node A

Destination	Cost	Next hop
A	0	—
B	2	—
C	7	B
D	3	—
E	6	B
F	8	B
G	9	B

P20-13. The following shows the steps to create the shortest path tree for node B and the forwarding table for this node.

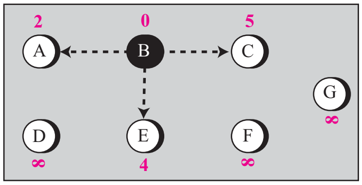

Initialization

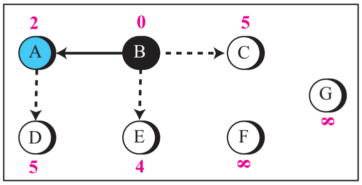

Iteration 1

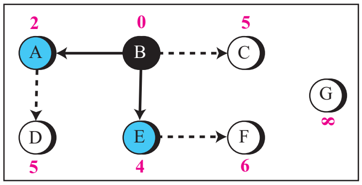

Iteration 2

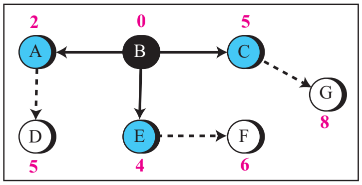

Iteration 3

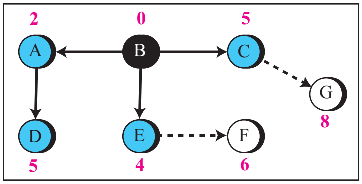

Iteration 4

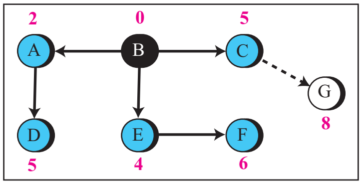

Iteration 5

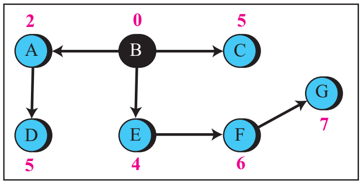

Iteration 6

Destination	Cost	Next hop
A	2	—
B	0	—
C	5	—
D	5	A
E	4	—
F	6	E
G	8	E

Forwarding table for node B

P20-15. The number of searches in each iteration of Dijkstra's algorithm is different. In the first iteration, we need n number of searches, in the second iteration, we need $(n - 1)$, and finally in the last iteration, we need only one. In other words, the total number of searches for each node to find its own shortest-path tree is

$$\text{Number of searches} = n + n - 1 + n - 2 + n - 3 + \ldots + 3 + 2 + 1 = n\,(n + 1) / 2$$

The series can be calculated if it is written twice: once in order and once in the reverse order. We then have n items, each of value $(n + 1)$, which results in $n(n + 1)$. However, we need to divide the result by 2. In computer science, this complexity is written as $O(n^2)$ and is referred to as Big-O notation.

P20-17. Router R1, using its OSPF forwarding table, knows how to forward a packet destined for N4. R1 announces this reachability to R5 using an eBGP session. R5 adds an entry to its RIP forwarding table that shows R1 as the next router for any packet destined for N4.