Computer Networks

Day - 9

Network Layer: Routing

UNICAST ROUTING PROTOCOLS

- *A routing table can be either static or dynamic.*
- *A static table is one with manual entries.*
- *A dynamic table is one that is updated automatically when there is a change somewhere in the Internet.*
- *A routing protocol is a combination of rules and procedures that lets routers in the Internet inform each other of changes.*

Autonomous systems

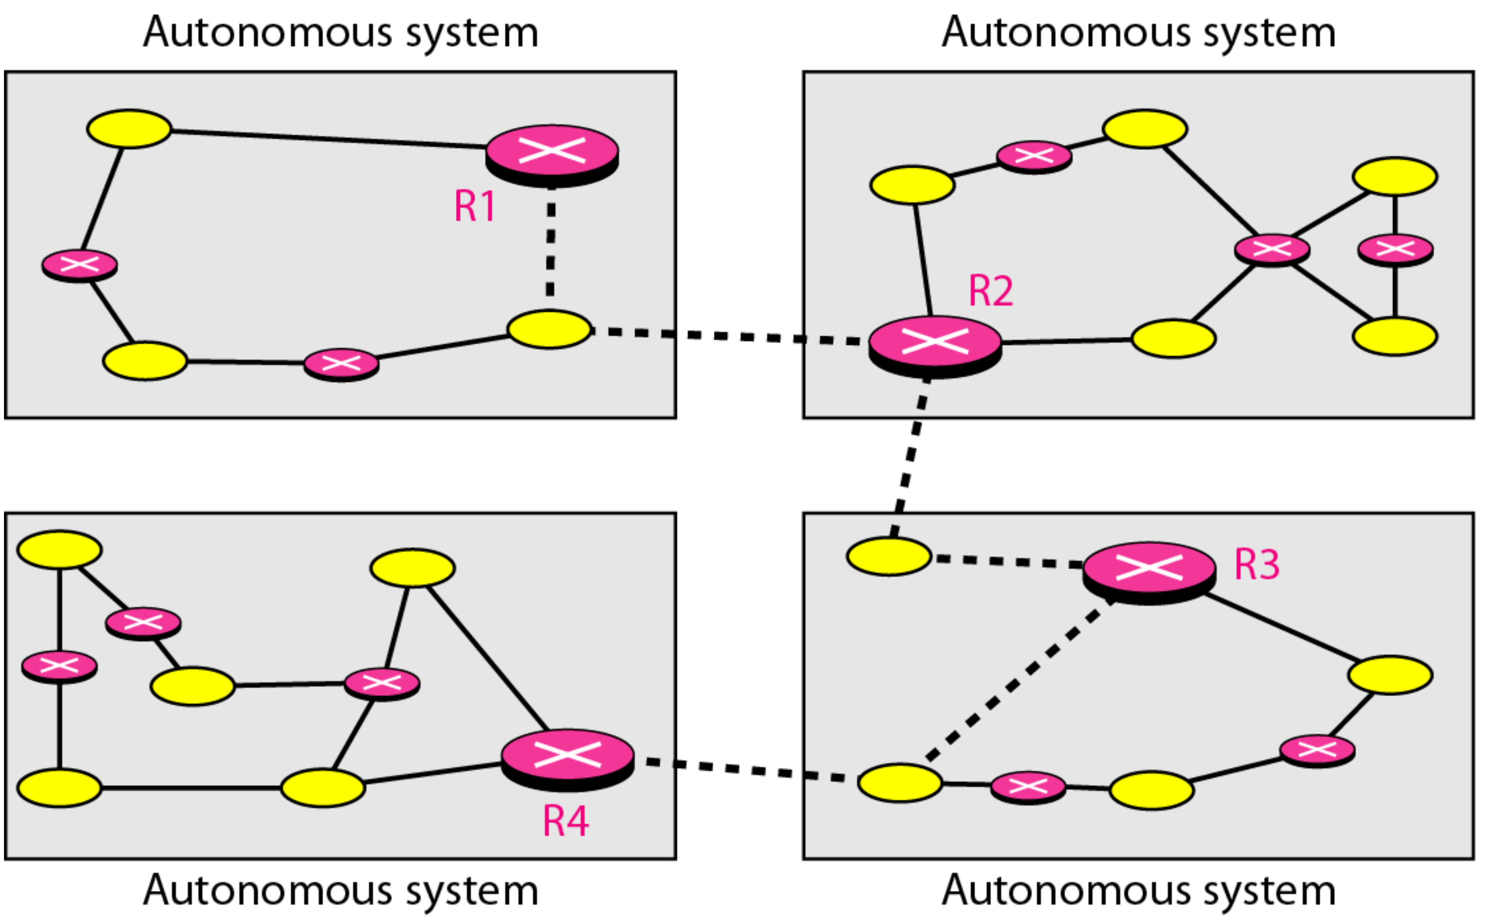

Popular routing protocols

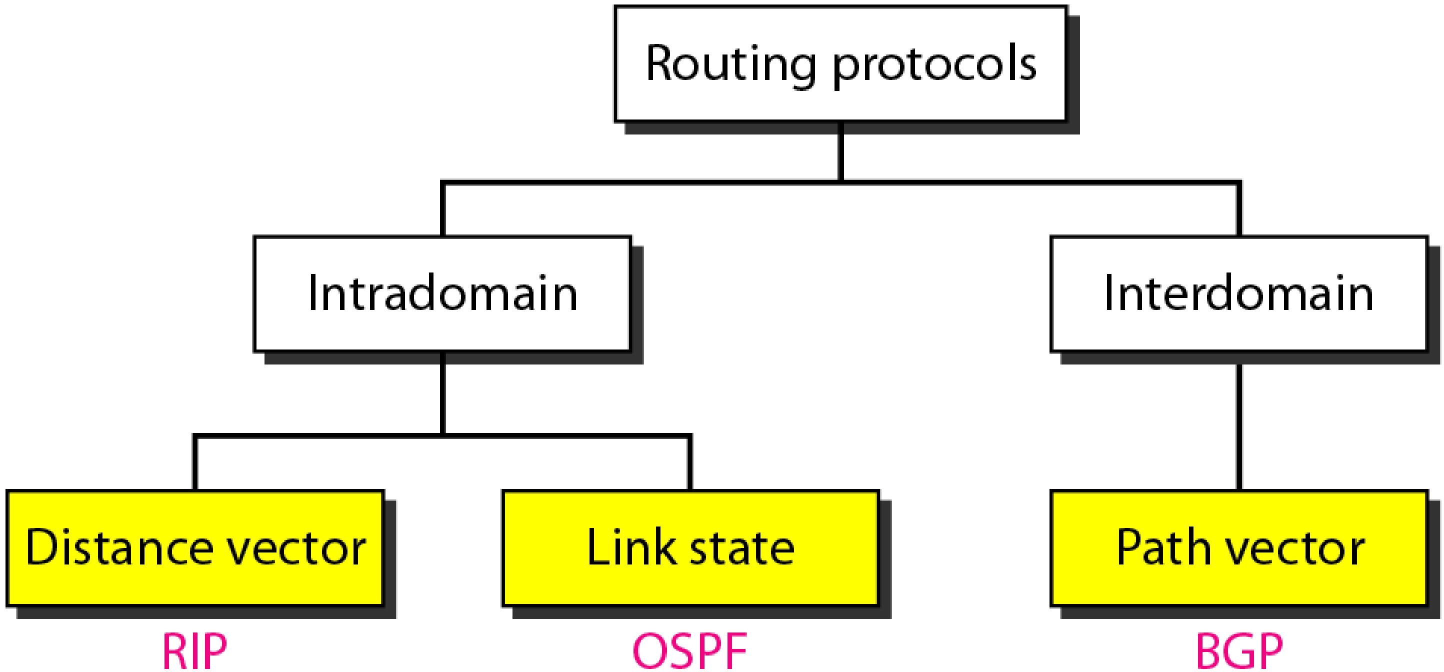

Distance vector routing

Bellman Ford Algorithm or Ford Fulkerson Algorithm

$$D_{ij} = \text{minimum } \{(c_{i1} + D_{1j}), (c_{i2} + D_{2j}), \ldots (c_{iN} + D_{Nj})\}$$

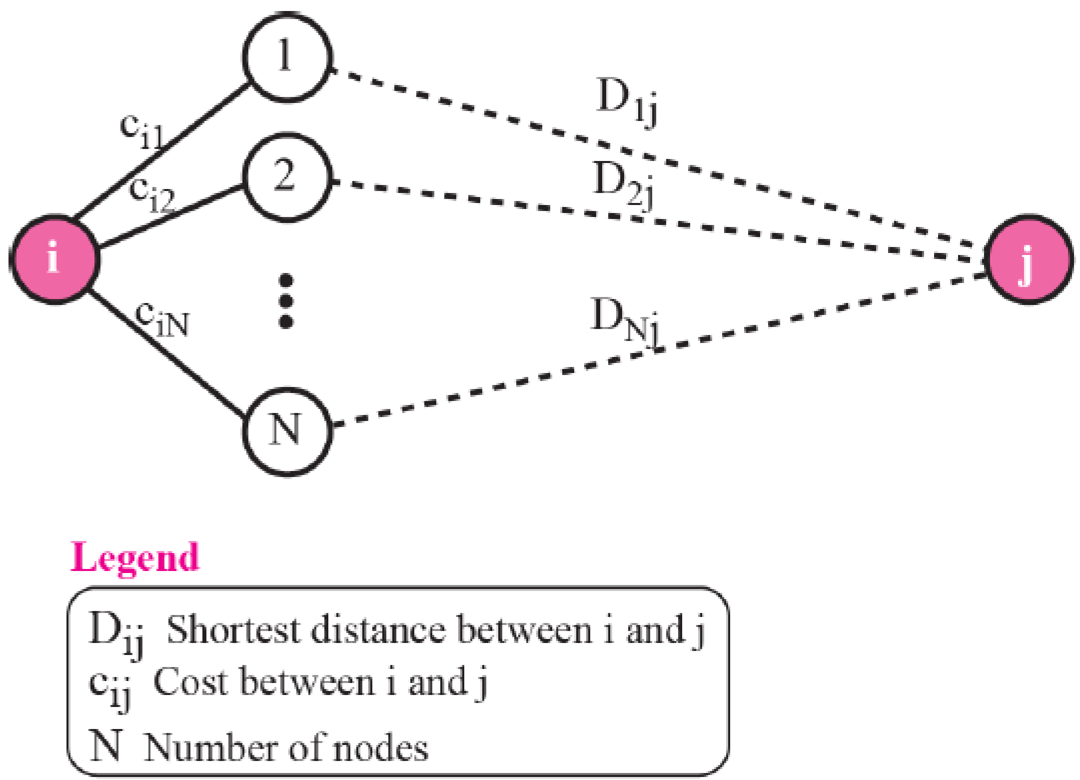

Distance vector routing tables

To	Cost	Next
A	0	—
B	5	—
C	2	—
D	3	—
E	6	C

A's table

To	Cost	Next
A	5	—
B	0	—
C	4	—
D	8	A
E	3	—

B's table

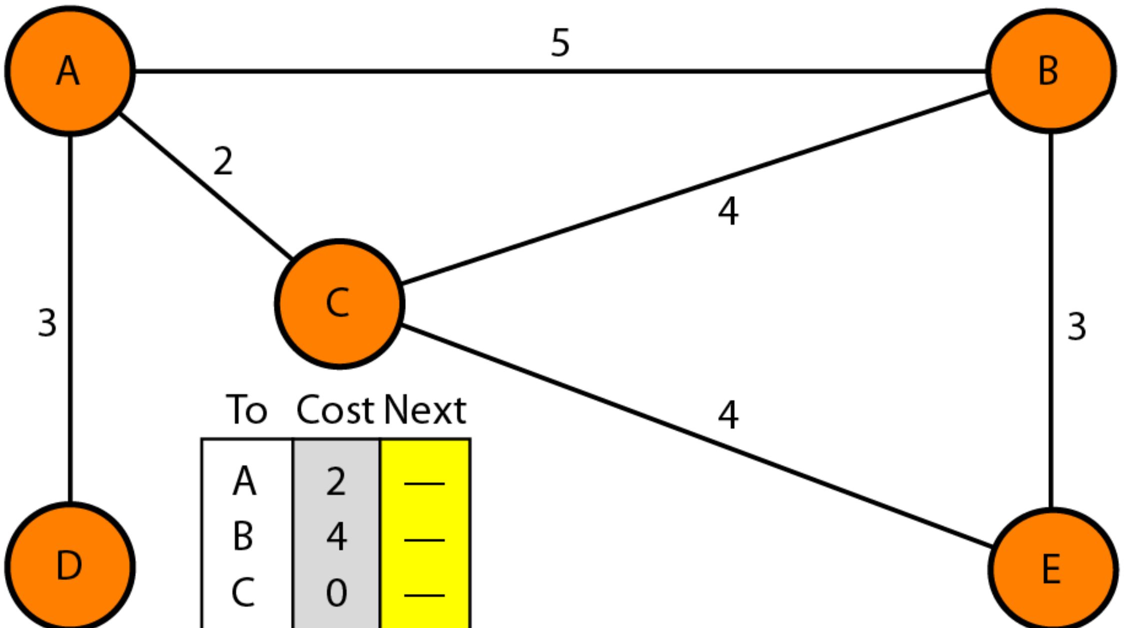

To	Cost	Next
A	3	—
B	8	A
C	5	A
D	0	—
E	9	A

D's table

To	Cost	Next
A	2	—
B	4	—
C	0	—
D	5	A
E	4	—

C's table

To	Cost	Next
A	6	C
B	3	—
C	4	—
D	9	C
E	0	—

E's table

Initialization of tables in distance vector routing

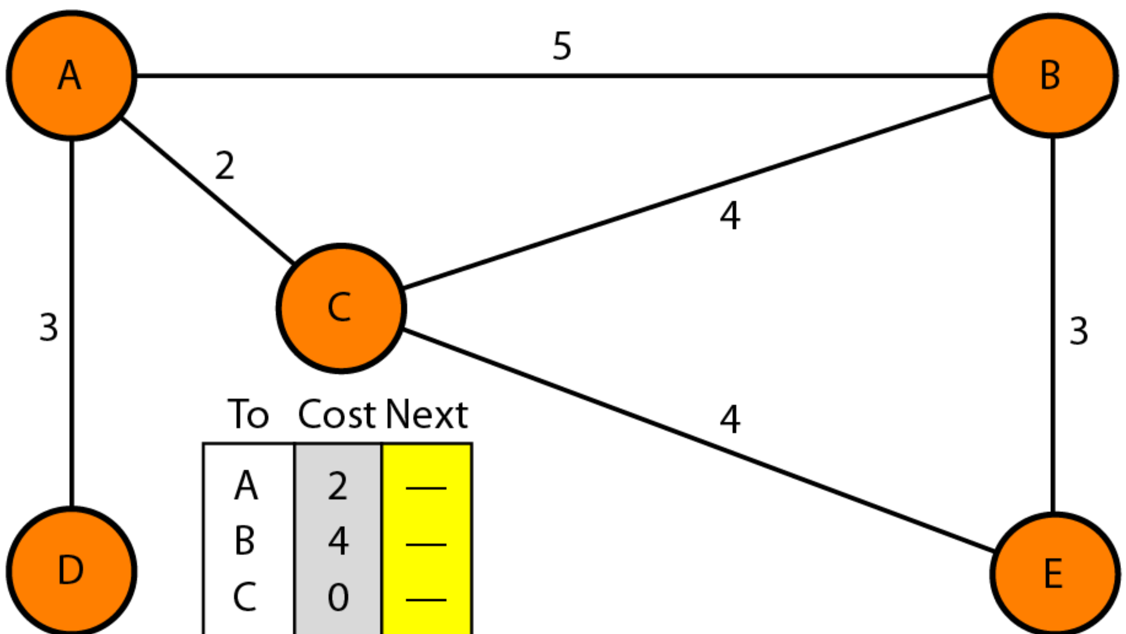

A's table

To	Cost	Next
A	0	—
B	5	—
C	2	—
D	3	—
E	∞	

B's table

To	Cost	Next
A	5	—
B	0	—
C	4	—
D	∞	
E	3	—

D's table

To	Cost	Next
A	3	—
B	∞	
C	∞	
D	0	—
E	∞	

C's table

To	Cost	Next
A	2	—
B	4	—
C	0	—
D	∞	
E	4	—

E's table

To	Cost	Next
A	∞	
B	3	B
C	4	C
D	∞	
E	0	D

In distance vector routing, each node shares its routing table with its immediate neighbors periodically (30s) and when there is a change.

Updating in distance vector routing

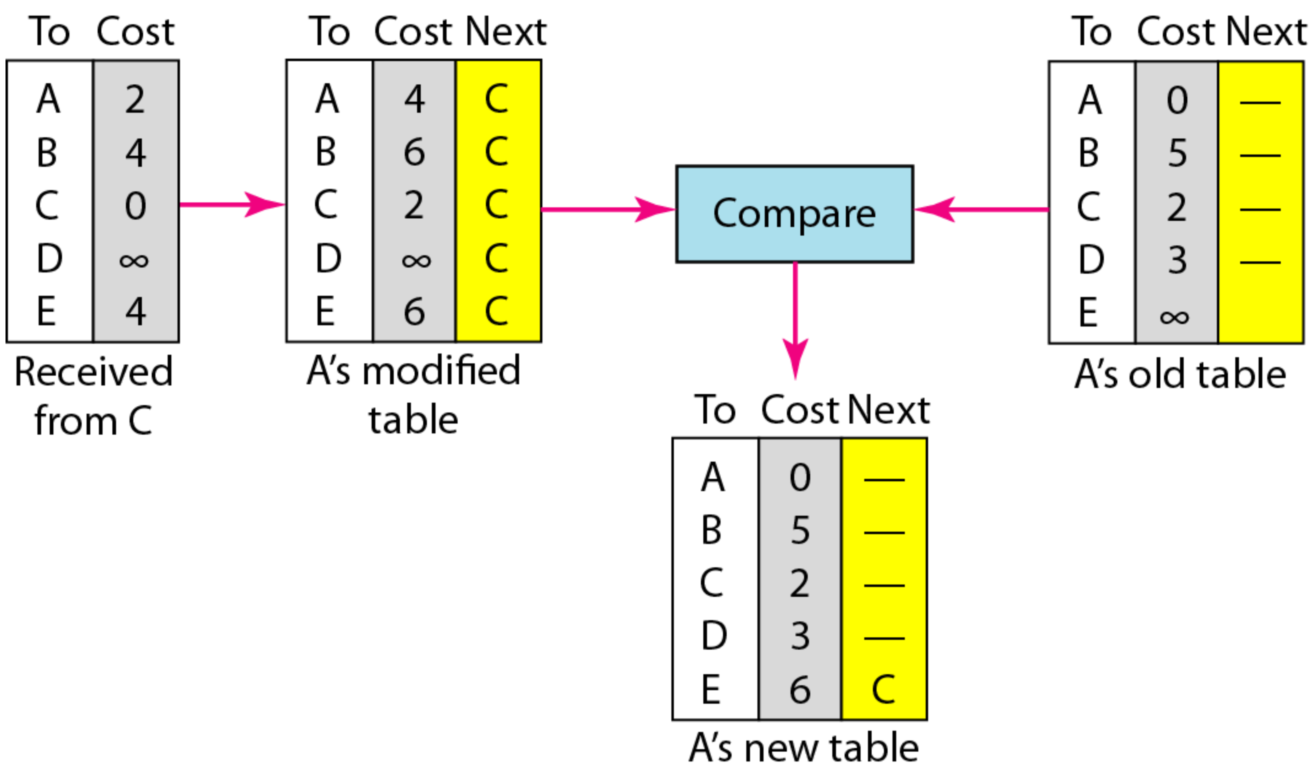

Two-node instability

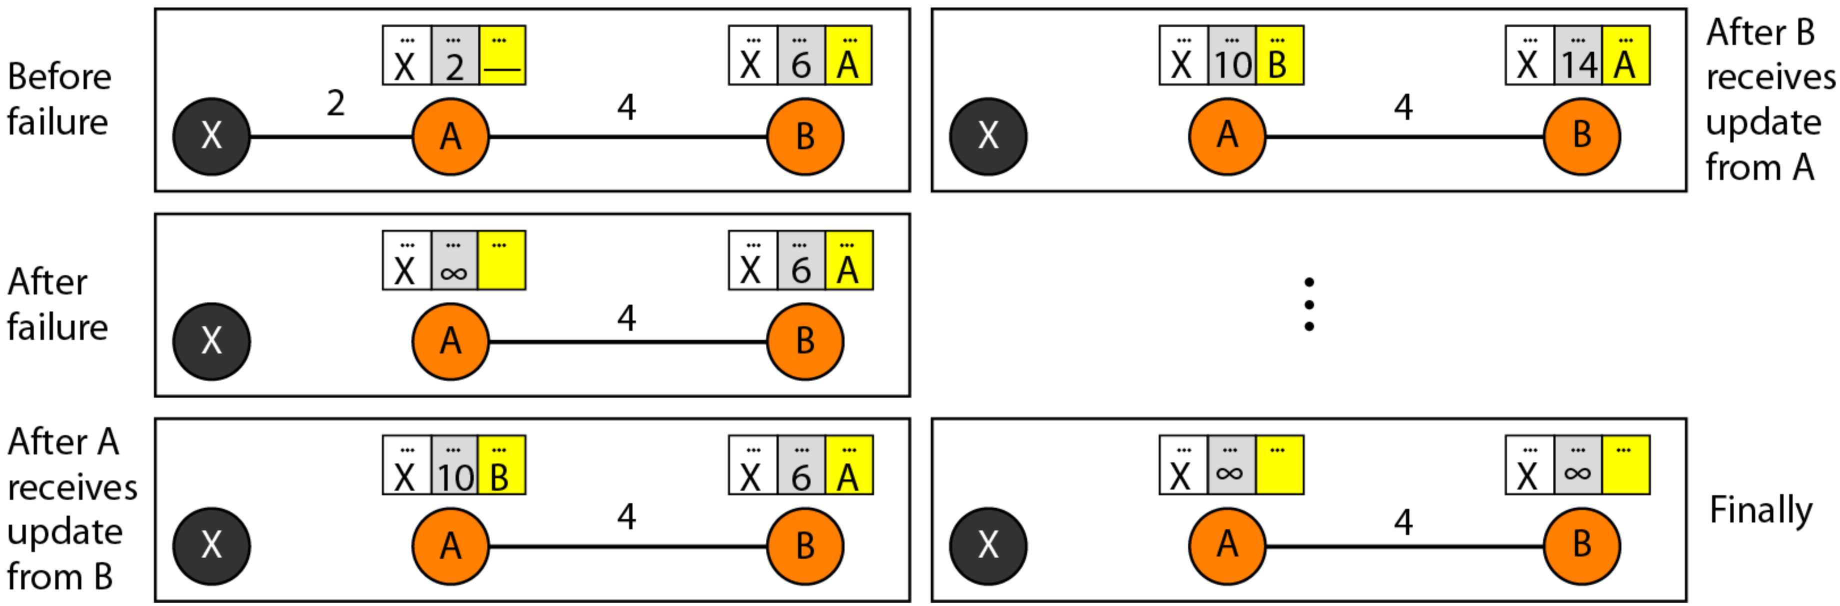

Two-Node Instability (1)

- Defining Infinity
  - Most implementations define 16 as infinity
- Split Horizon
  - Instead of flooding the table through each interface, each node sends only part of its table through each interface
  - E.g. node B thinks that the optimum route to reach X is via A, it does not need to advertise this piece of information to A

Two-Node Instability (2)

- Split Horizon and Poison Reverse
  - One drawback of Split Horizon
    - Normally, the DV protocol uses a timer and if there is no news about a route, the node deletes the route from its table
    - In the previous e.g., node A cannot guess that this is due to split horizon or because B has not received any news about X recently
  - Poison Reverse
    - Node B can still advertise the value for X, but if the source of information is A, it can replace the distance with infinity as a warning

Three-node instability

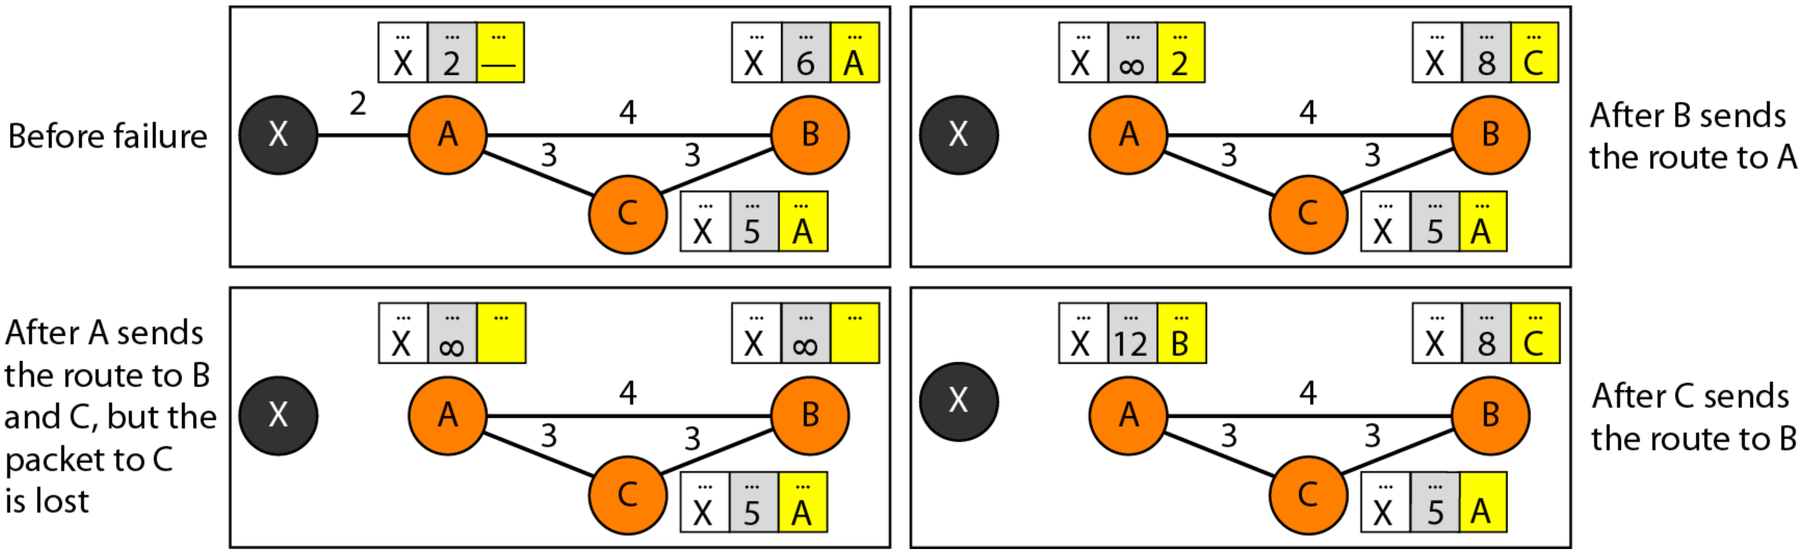

Example of a domain using RIP

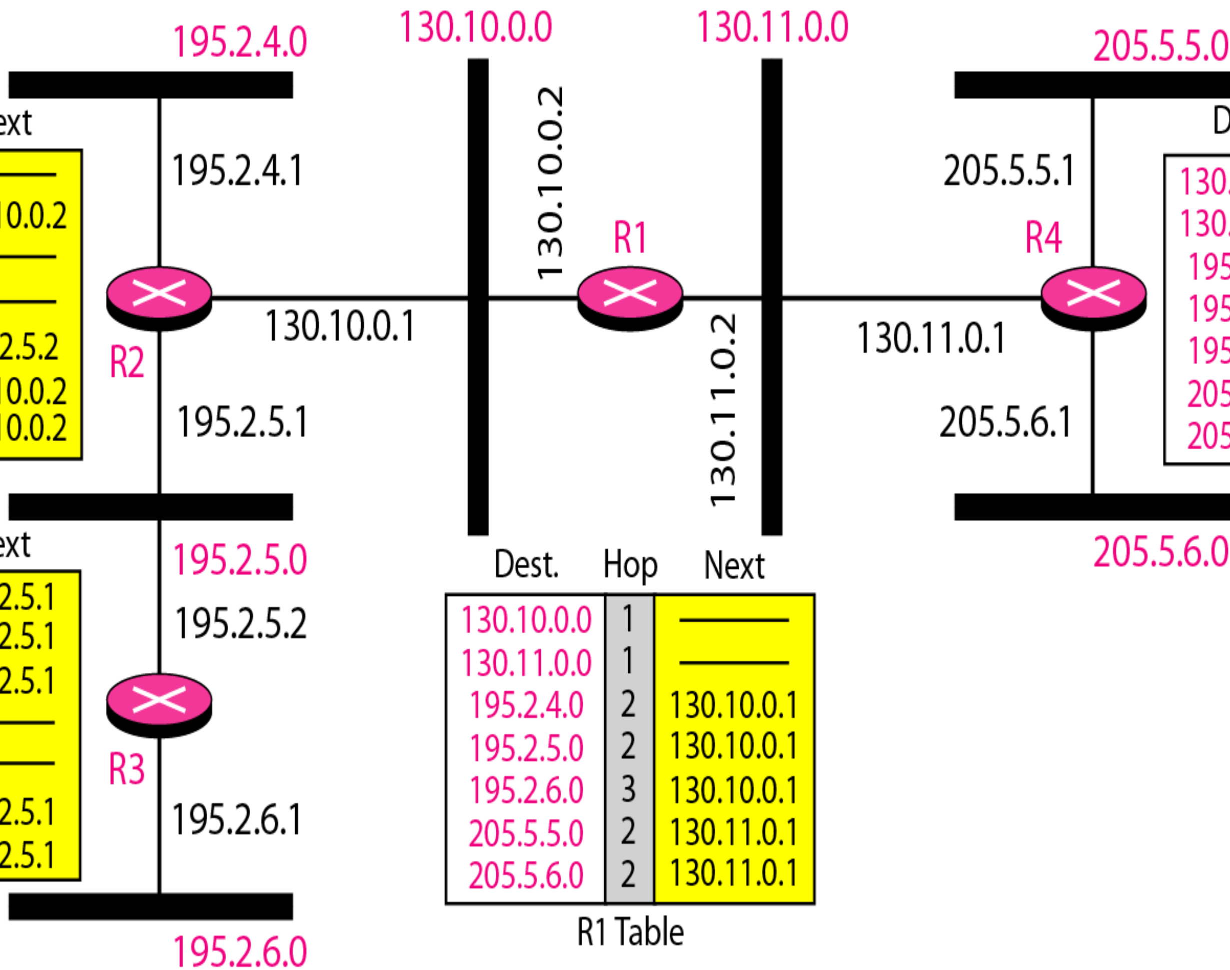

Dest.	Hop	Next
130.10.0.0	1	——
130.11.0.0	2	130.10.0.2
195.2.4.0	1	——
195.2.5.0	1	——
195.2.6.0	2	195.2.5.2
205.5.5.0	3	130.10.0.2
205.5.6.0	3	130.10.0.2

R2 Table

Dest.	Hop	Next
130.10.0.0	2	195.2.5.1
130.11.0.0	3	195.2.5.1
195.2.4.0	2	195.2.5.1
195.2.5.0	1	——
195.2.6.0	1	——
205.5.5.0	4	195.2.5.1
205.5.6.0	4	195.2.5.1

R3 Table

Dest.	Hop	Next
130.10.0.0	1	——
130.11.0.0	1	——
195.2.4.0	2	130.10.0.1
195.2.5.0	2	130.10.0.1
195.2.6.0	3	130.10.0.1
205.5.5.0	2	130.11.0.1
205.5.6.0	2	130.11.0.1

R1 Table

Dest.	Hop	Next
130.10.0.0	2	130.11.0.2
130.11.0.0	1	——
195.2.4.0	3	130.11.0.2
195.2.5.0	3	130.11.0.2
195.2.6.0	4	130.11.0.2
205.5.5.0	1	——
205.5.6.0	1	——

R4 Table

LINK STATE ROUTING

- Link state routing has a different philosophy from that of distance vector routing.

- In link state routing, if each node in the domain has the entire topology of the domain—the list of nodes and links, how they are connected including the type, cost (metric), and the condition of the links (up or down)—the node can use the Dijkstra algorithm to build a routing table.

Concept of link state routing

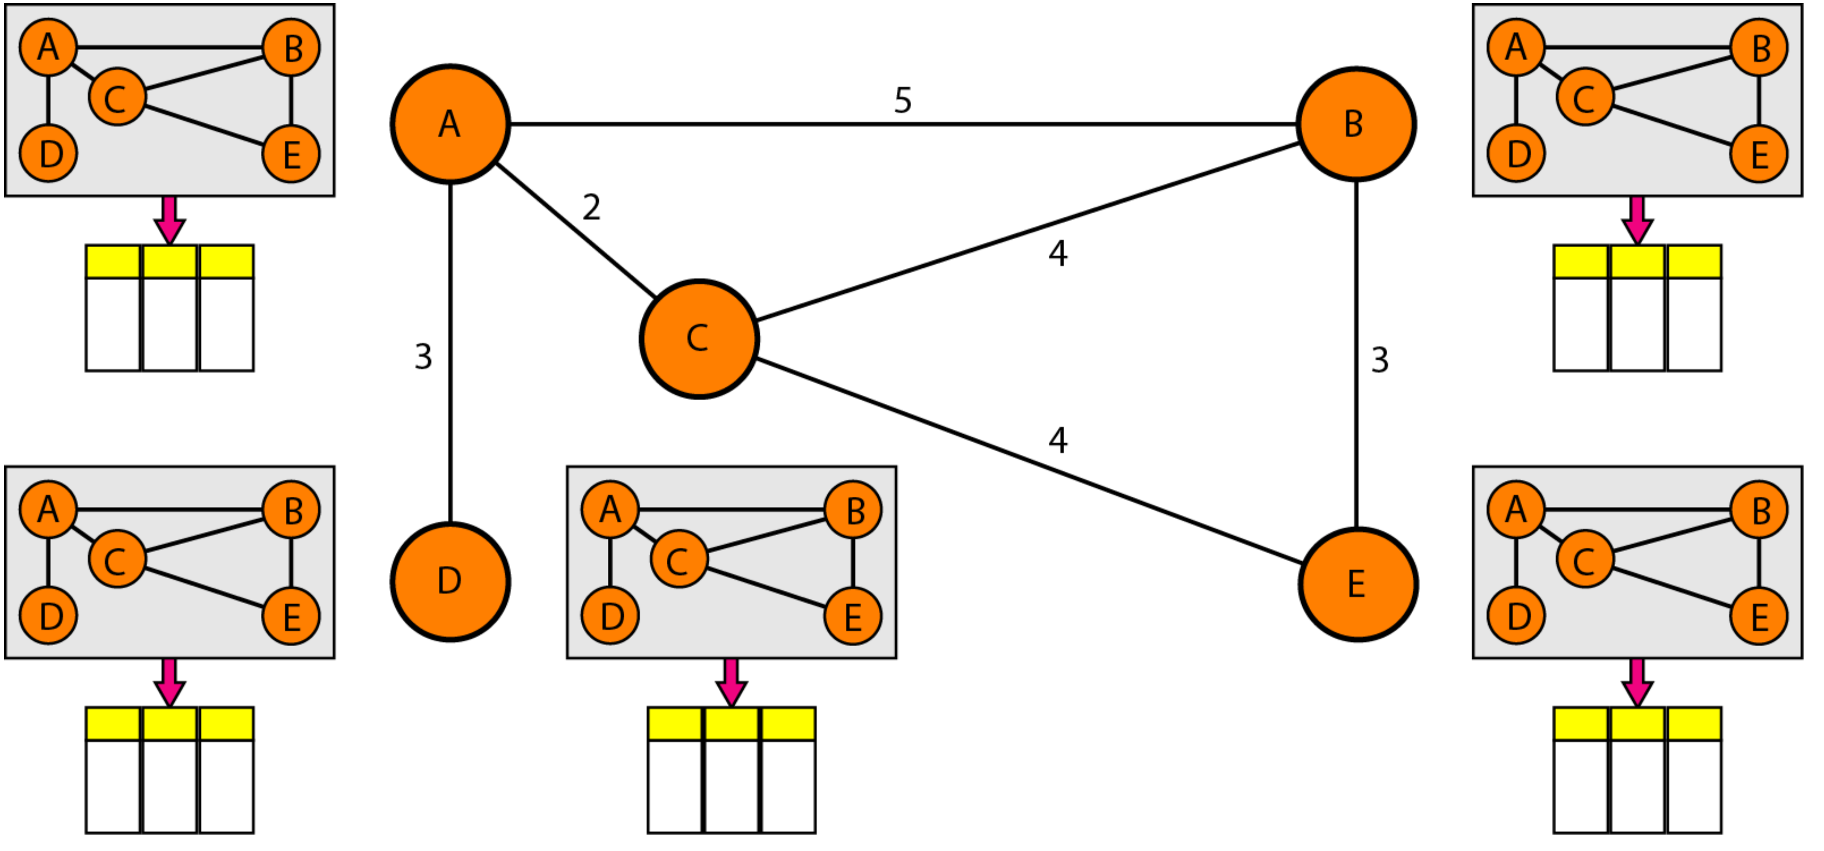

Link state knowledge

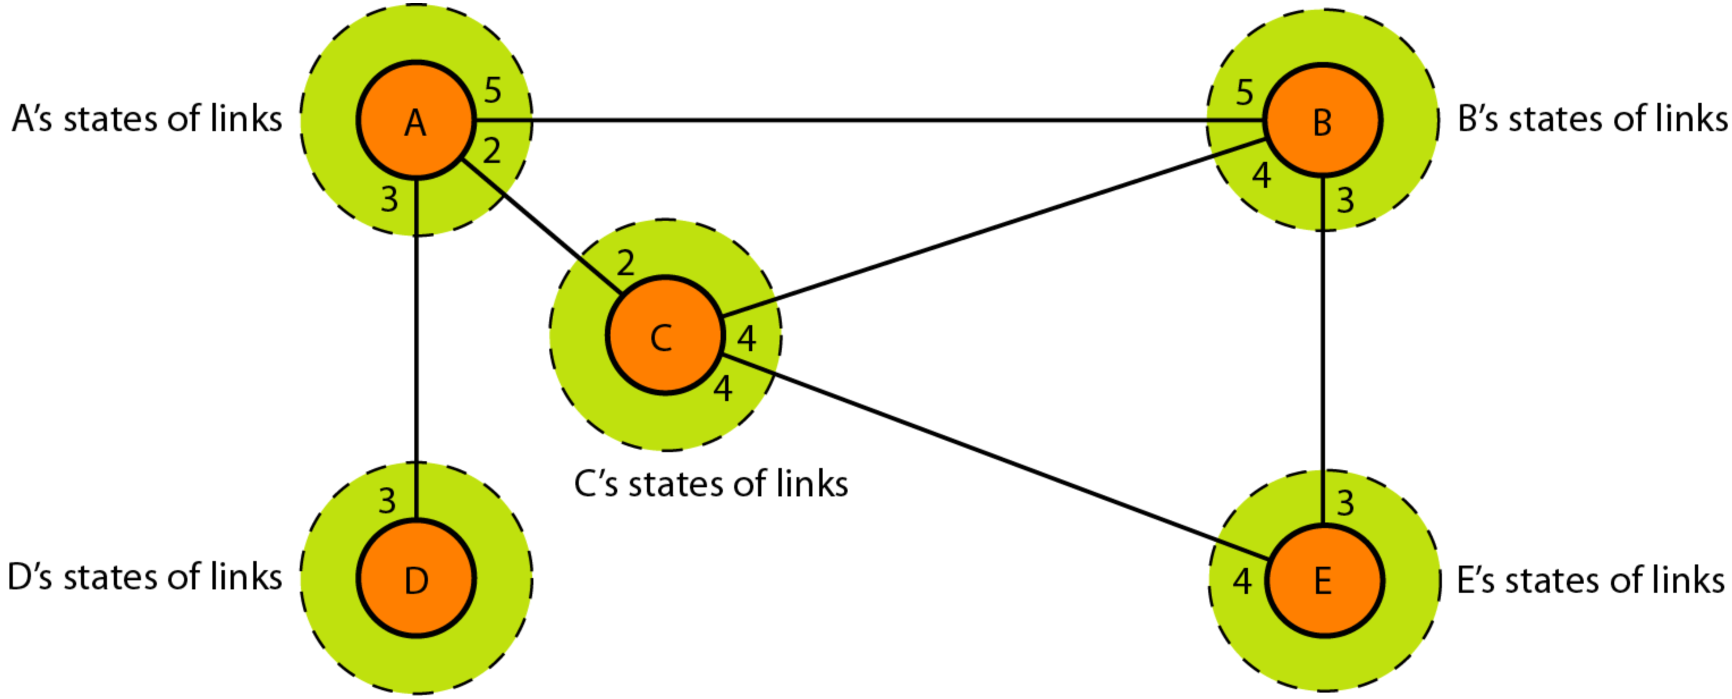

Building Routing Tables

- Creation of the states of the links by each node, called the link state packets (LSP)
- Dissemination of LSPs to every other routers, called flooding (efficiently)
- Formation of a shortest path tree for each node
- Calculation of a routing table based on the shortest path tree

Creation of LSP

- LSP data: E.g. the node ID, the list of links, a sequence number, and age.
- LSP Generation
  - When there is a change in the topology of the domain
  - On a periodic basis
    - There is no actual need for this type of LSP, normally 60 minutes or 2 hours

Dijkstra algorithm

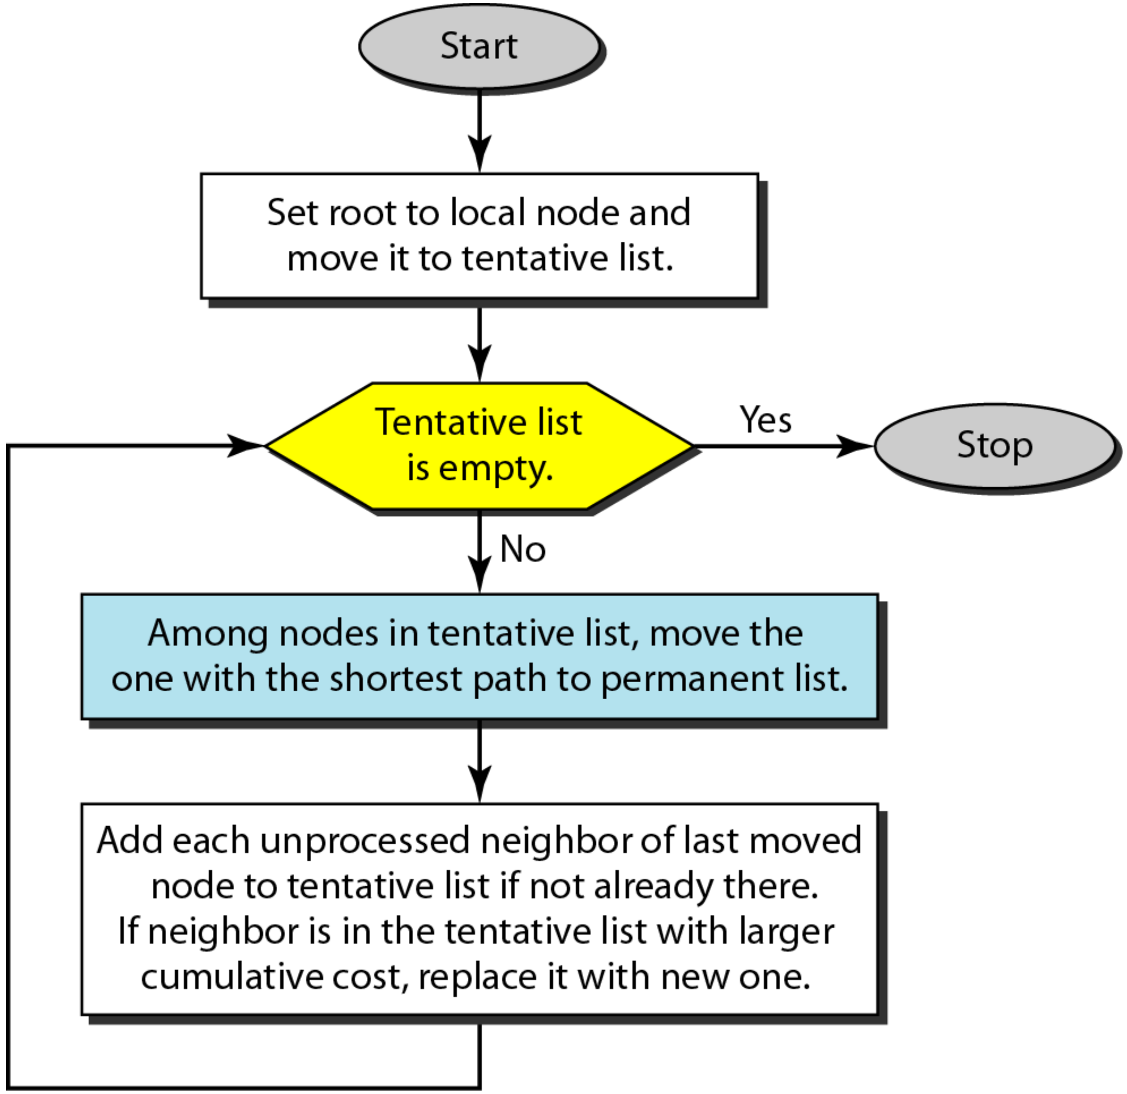

Example of formation of shortest path tree

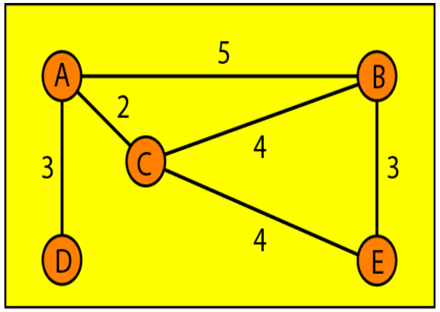

Topology

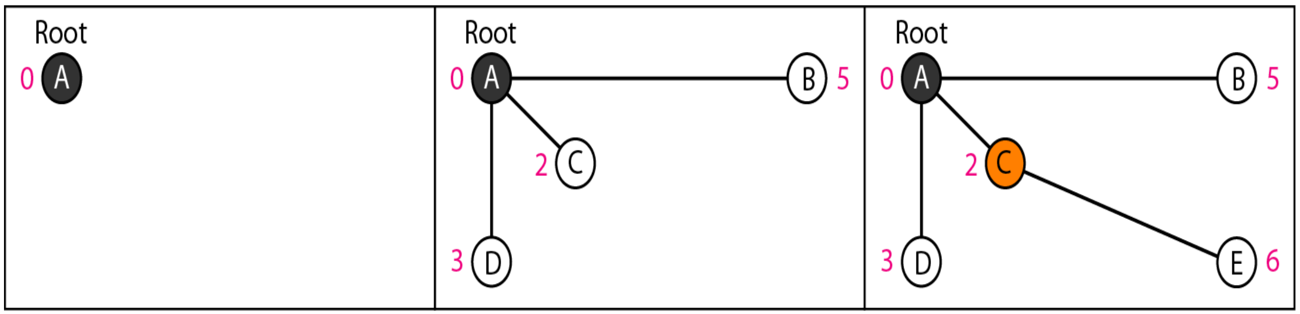

1. Set root to A and move A to tentative list.

2. Move A to permanent list and add B, C, and D to tentative list.

3. Move C to permanent and add E to tentative list.

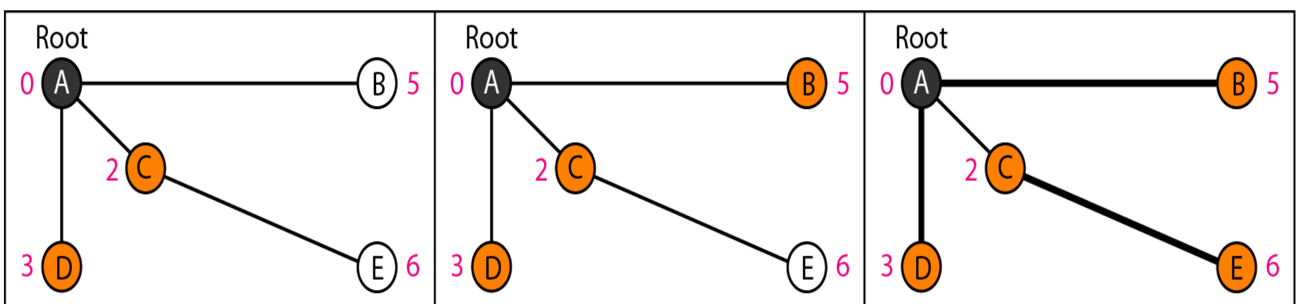

4. Move D to permanent list.

5. Move B to permanent list.

6. Move E to permanent list (tentative list is empty).

Node	*Cost*	*Next Router*
A	0	—
B	5	—
C	2	—
D	3	—
E	6	C

Forming shortest path three for router A in a graph

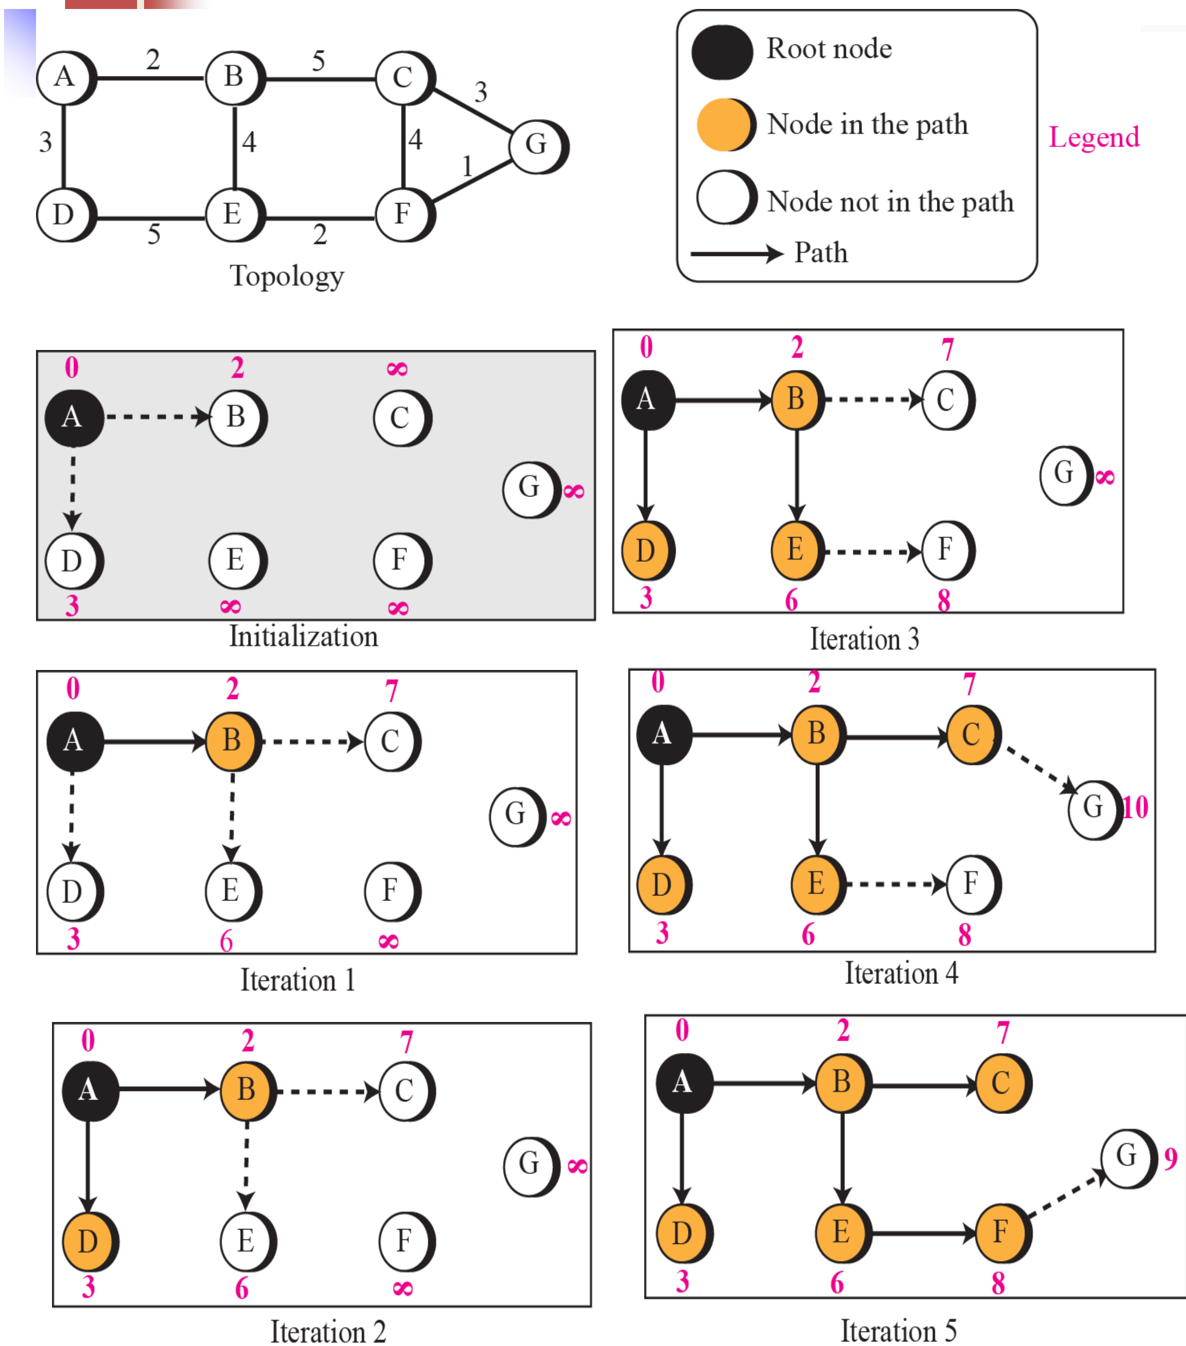

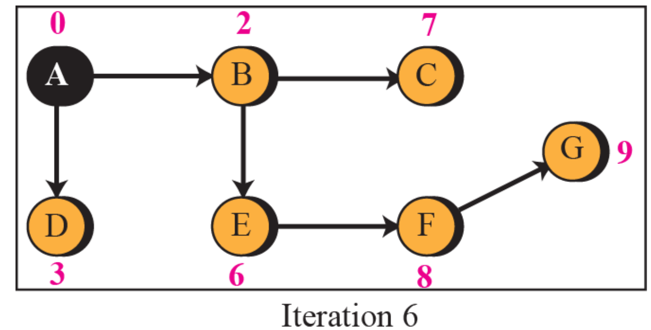

Table 11.4 *Routing Table for Node A*

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

Areas in an autonomous system

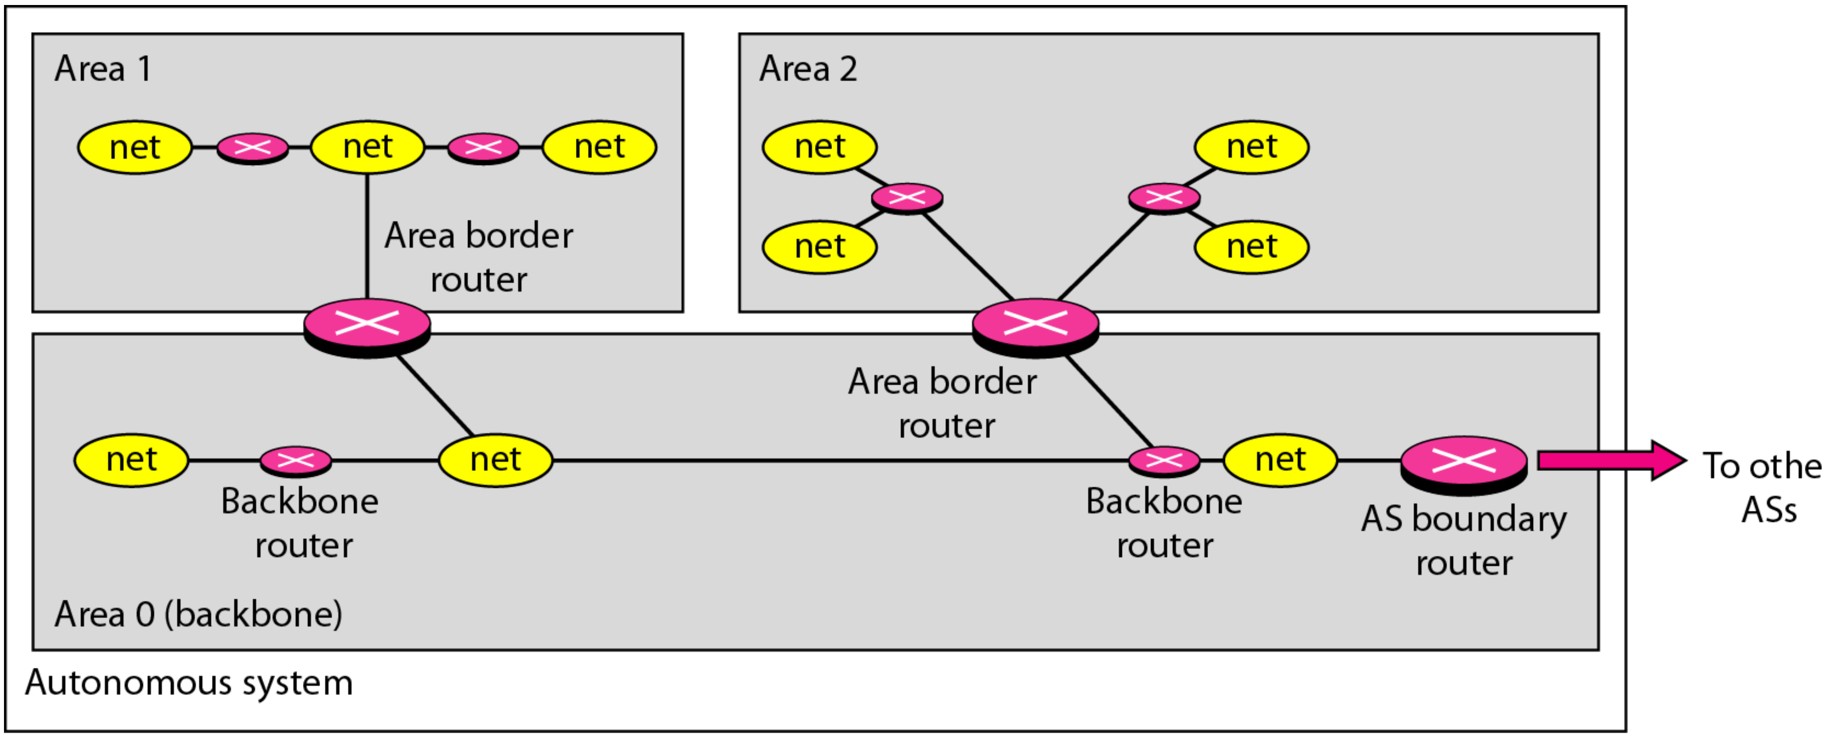

Types of links

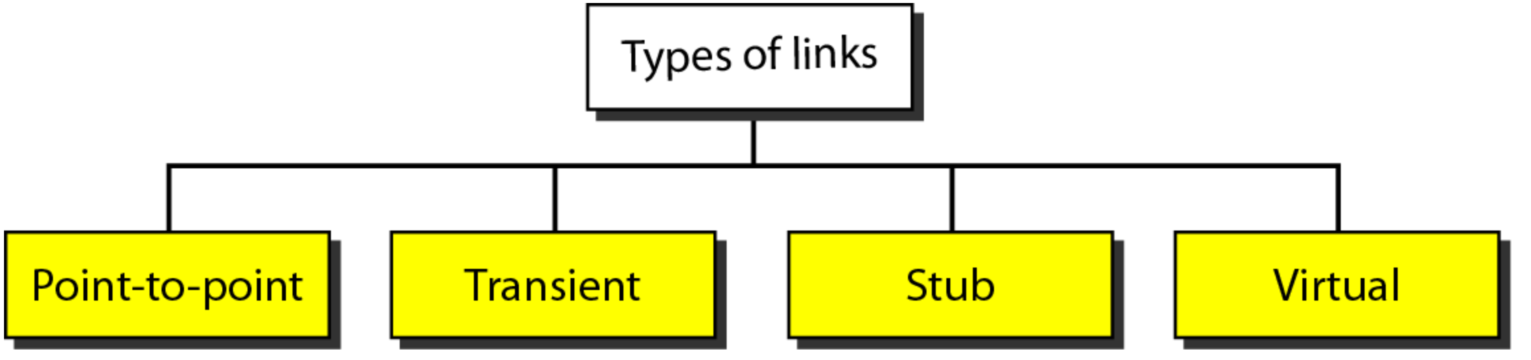

Point-to-point link

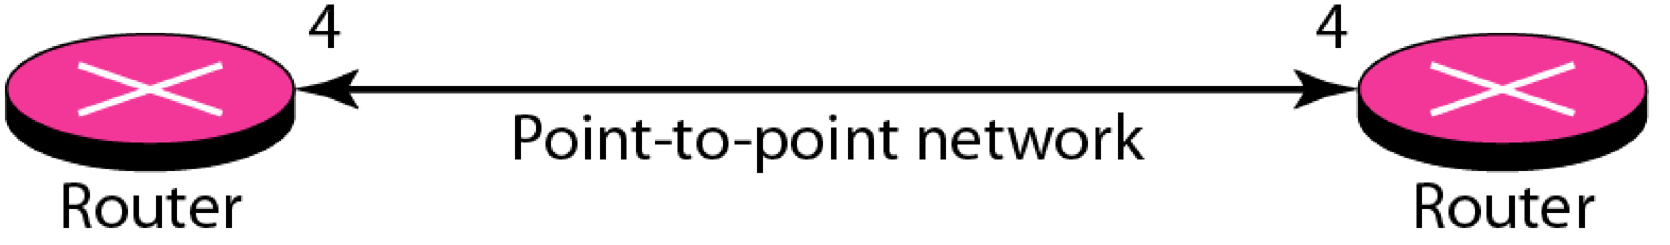

Transient link

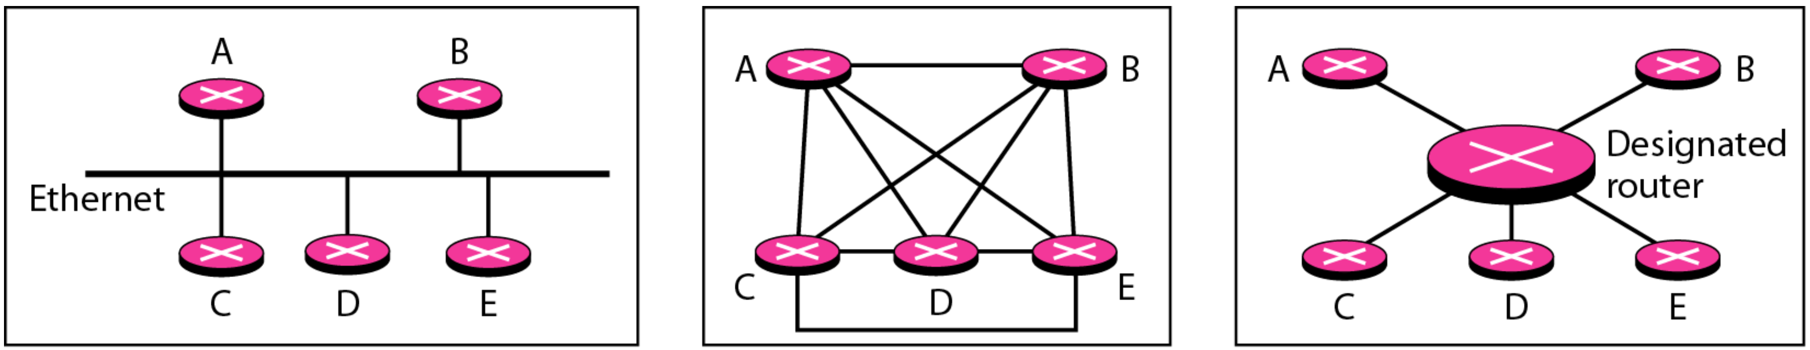

a. Transient network b. Unrealistic representation c. Realistic representation

Stub link

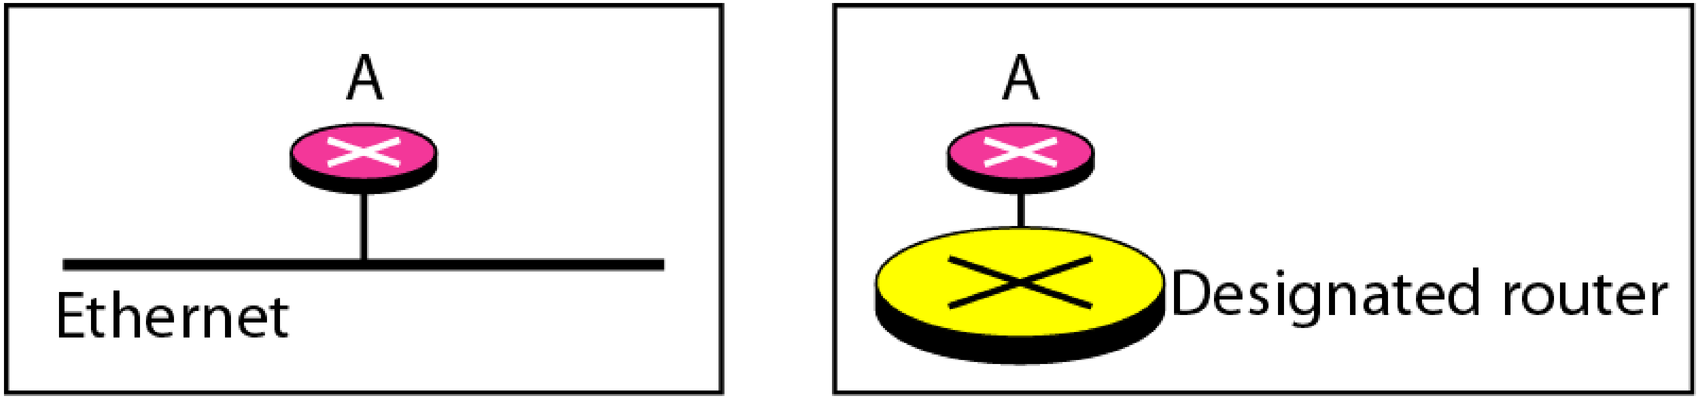

a. Stub network

b. Representation

Example of an AS and its graphical representation in OSPF

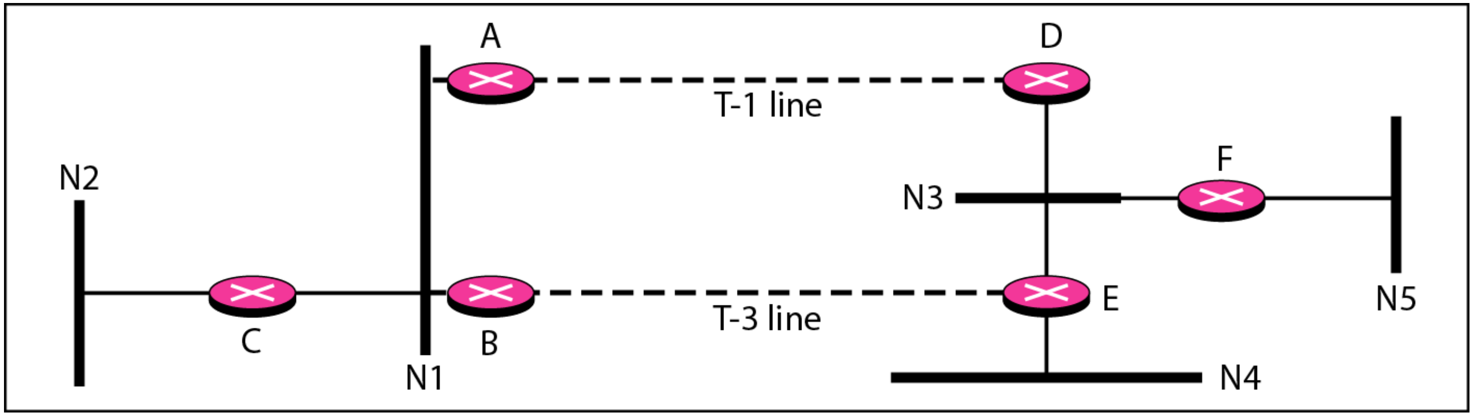

a. Autonomous system

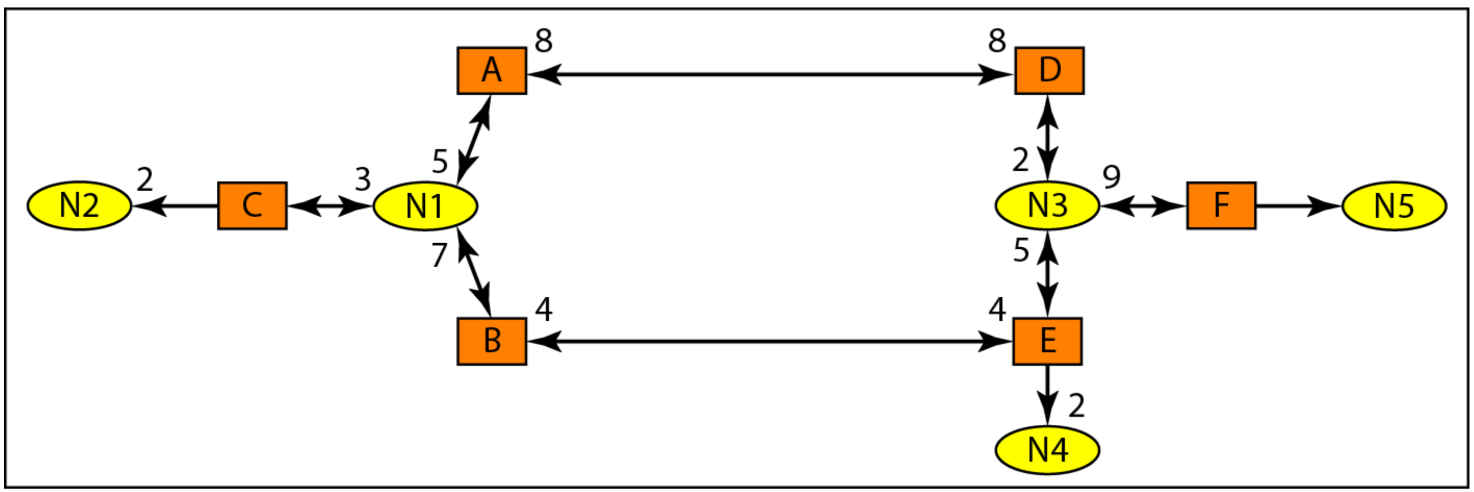

b. Graphical representation

PATH VECTOR ROUTING

- Distance vector and link state routing are both interior routing protocols.
- They can be used inside an autonomous system. Both of these routing protocols become intractable when the domain of operation becomes large.
- Distance vector routing is subject to instability if there is more than a few hops in the domain of operation.
- Link state routing needs a huge amount of resources to calculate routing tables. It also creates heavy traffic because of flooding.
- There is a need for a third routing protocol which we call path vector routing.

Initial routing tables in path vector routing

Dest.	Path
A1	AS1
A2	AS1
A3	AS1
A4	AS1
A5	AS1

A1 Table

Dest.	Path
C1	AS3
C2	AS3
C3	AS3

C1 Table

Dest.	Path
D1	AS4
D2	AS4
D3	AS4
D4	AS4

D1 Table

Dest.	Path
B1	AS2
B2	AS2
B3	AS2
B4	AS2

B1 Table

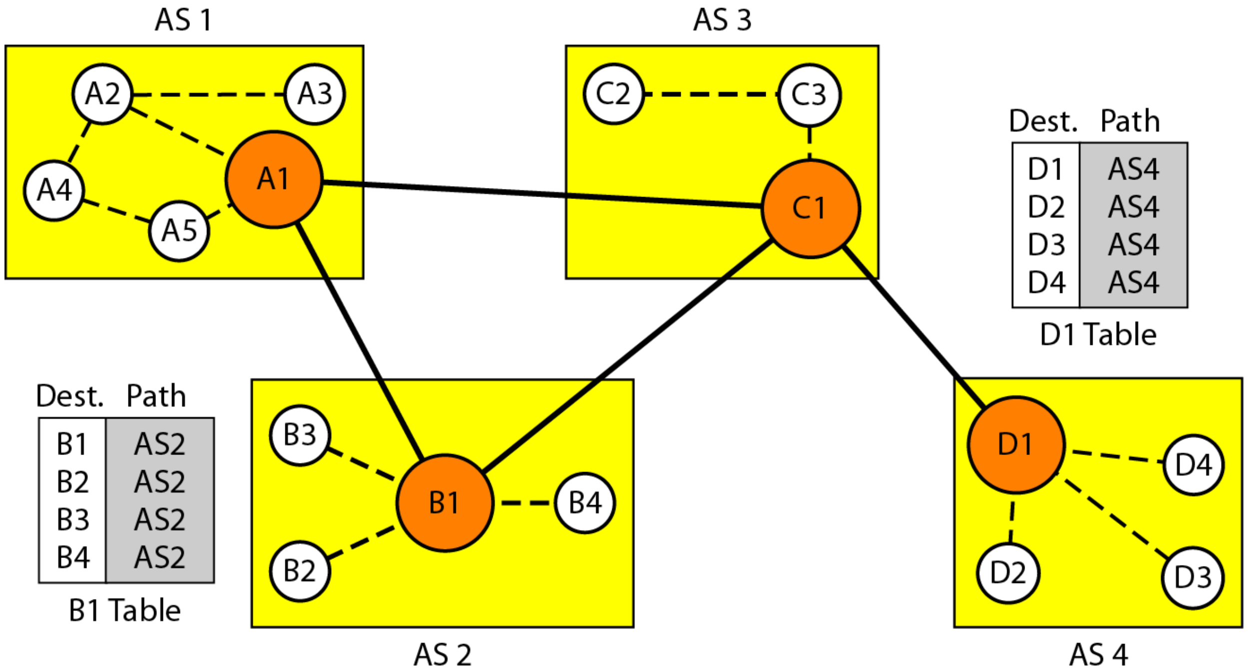

Stabilized tables for three autonomous systems

Dest.	Path
A1	AS1
…	
A5	AS1
B1	AS1-AS2
…	…
B4	AS1-AS2
C1	AS1-AS3
…	…
C3	AS1-AS3
D1	AS1-AS2-AS4
…	…
D4	AS1-AS2-AS4

A1 Table

Dest.	Path
A1	AS2-AS1
…	
A5	AS2-AS1
B1	AS2
…	…
B4	AS2
C1	AS2-AS3
…	…
C3	AS2-AS3
D1	AS2-AS3-AS4
…	…
D4	AS2-AS3-AS4

B1 Table

Dest.	Path
A1	AS3-AS1
…	
A5	AS3-AS1
B1	AS3-AS2
…	…
B4	AS3-AS2
C1	AS3
…	…
C3	AS3
D1	AS3-AS4
…	…
D4	AS3-AS4

C1 Table

Dest.	Path
A1	AS4-AS3-AS1
…	
A5	AS4-AS3-AS1
B1	AS4-AS3-AS2
…	…
B4	AS4-AS3-AS2
C1	AS4-AS3
…	…
C3	AS4-AS3
D1	AS4
…	…
D4	AS4

D1 Table

Internal and external BGP sessions

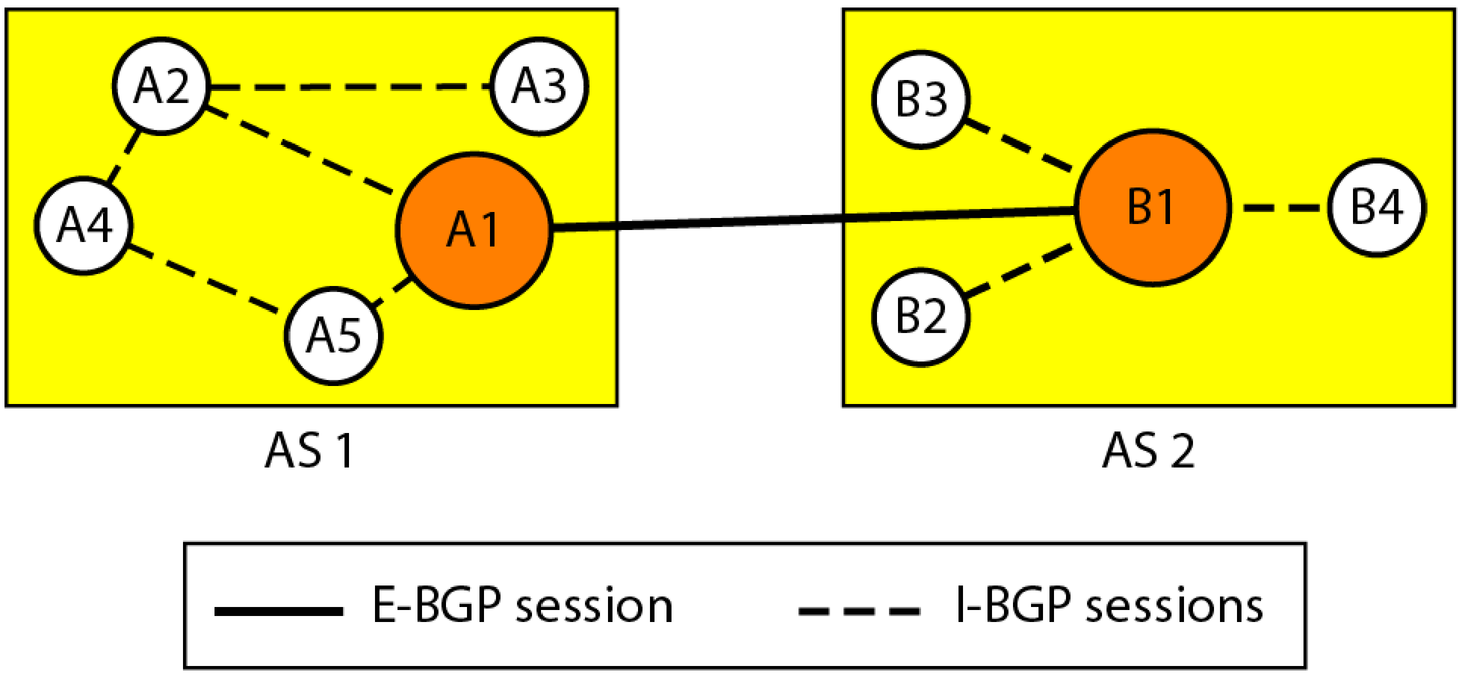

BGP supports classless addressing and CIDR.

Unicasting

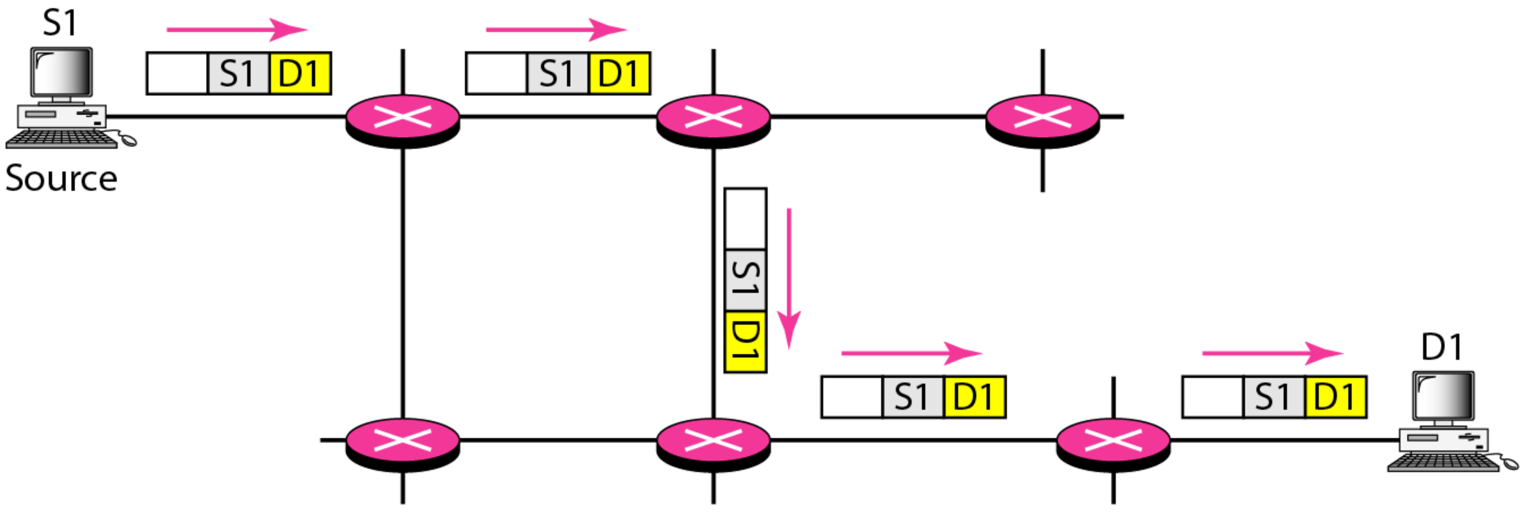

In unicasting, the router forwards the received packet through only one of its interfaces.

Multicasting

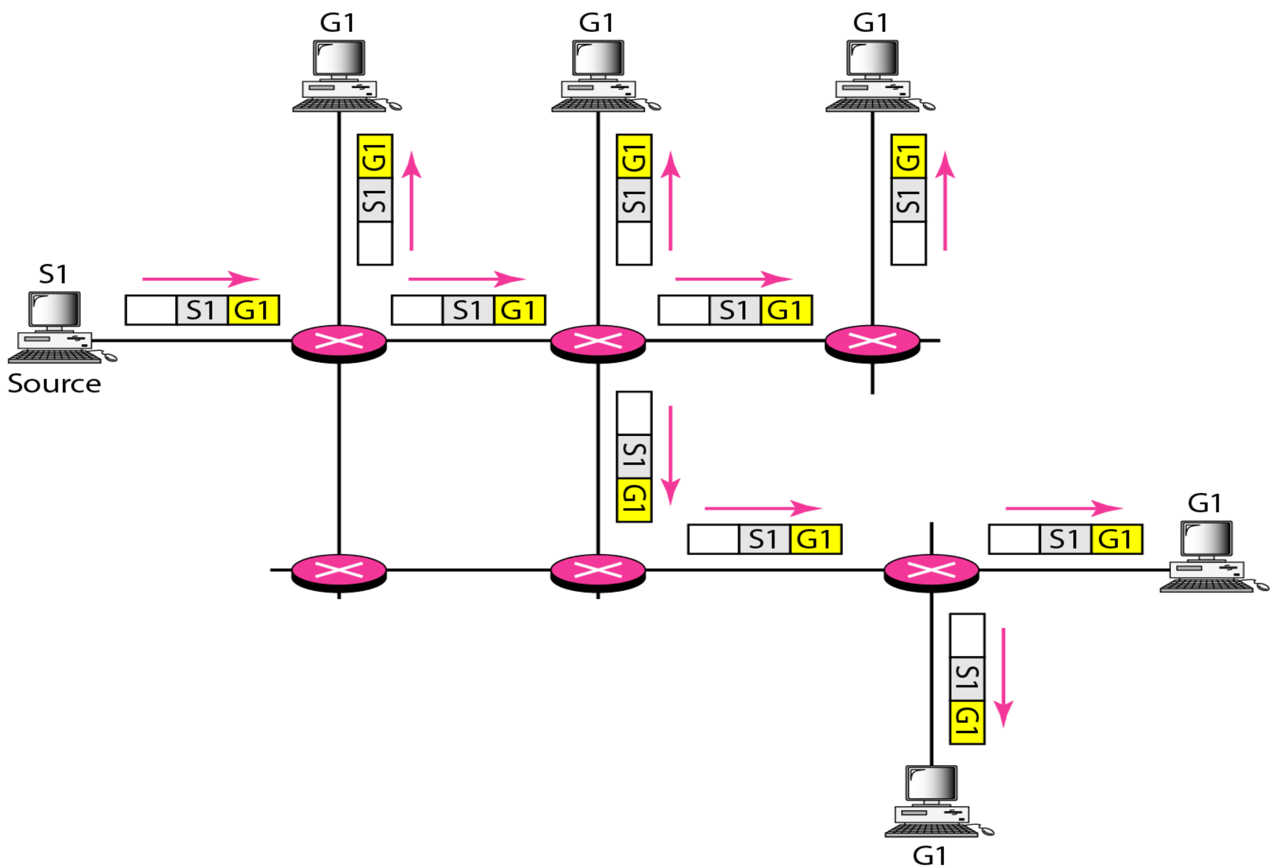

In multicasting, the router may forward the received packet through several of its interfaces.

Multicasting versus multiple unicasting

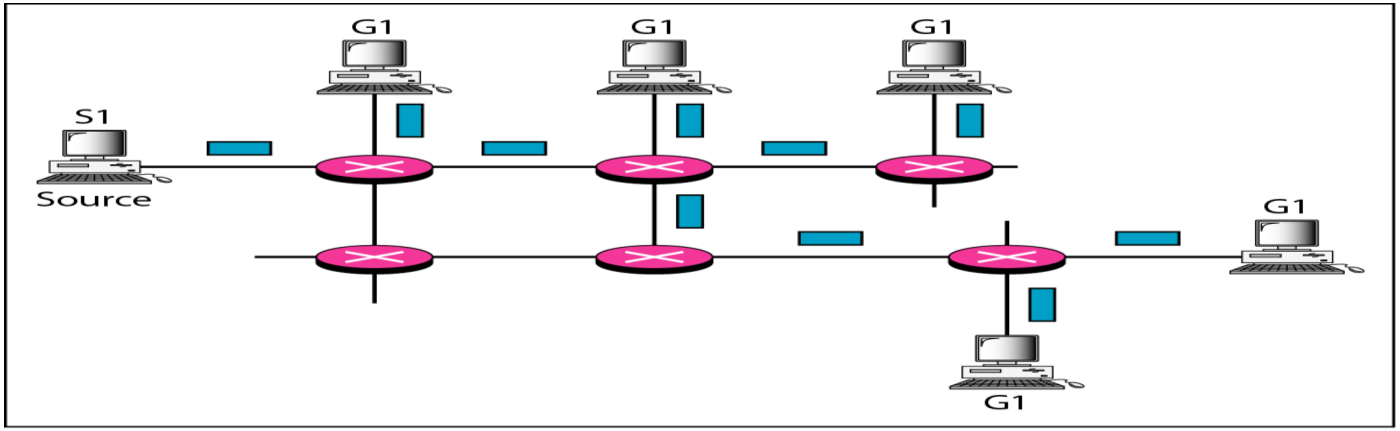

a. Multicasting

Consider a network with five nodes, N1 to N5, as shown below.

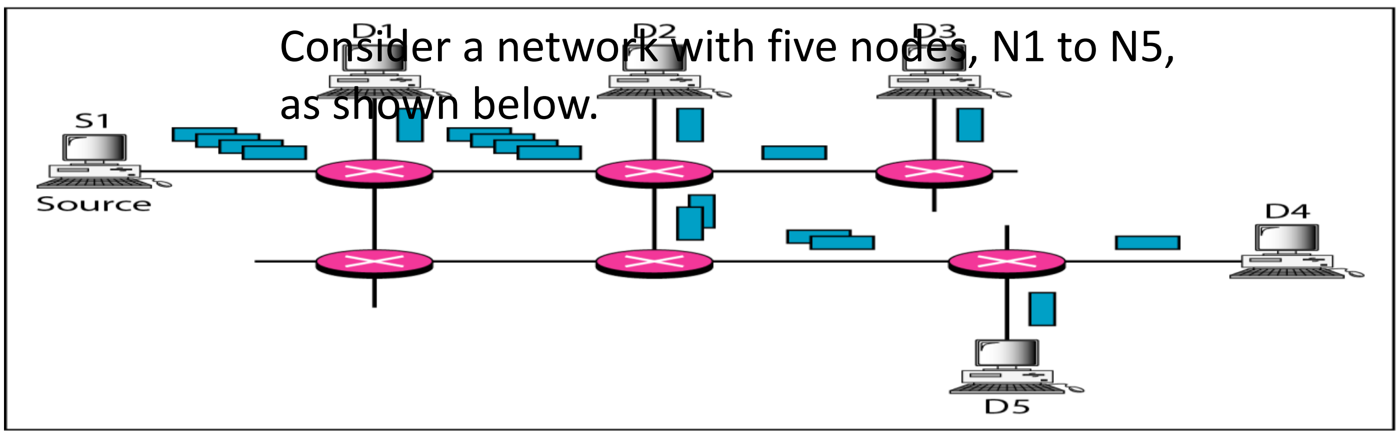

b. Multiple unicasting

Emulation of multicasting through multiple unicasting is not efficient and may create long delays, particularly with a large group.

Exercise

1. Consider a network with five nodes, N1 to N5, as shown below.

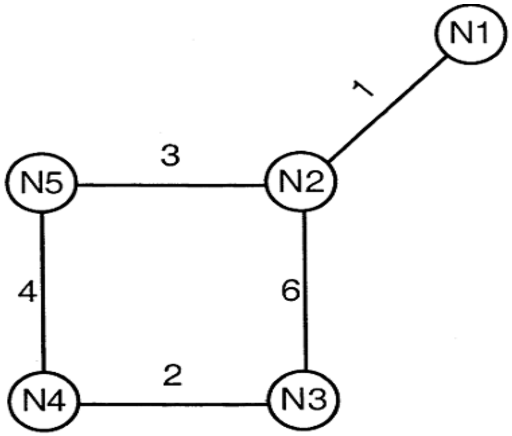

The network uses a Distance Vector Routing protocol. Once the routes have stabilized, the distance vectors at different nodes are as following. N1: (0, 1, 7, 8, 4)
N2: (1, 0, 6, 7, 3) N3: (7, 6, 0, 2, 6) N4: (8, 7, 2, 0, 4) N5: (4, 3, 6, 4, 0)
Each distance vector is the distance of the best known path at the instance to nodes, N1 to N5, where the distance to itself is 0. Also, all links are symmetric and the cost is identical in both directions. In each round, all nodes exchange their distance vectors with their respective neighbors. Then all nodes update their distance vectors. In between two rounds, any change in cost of a link will cause the two incident nodes to change only that entry in their distance vectors. The cost of link N2-N3 reduces to 2(in both directions). After the next round of updates, what will be the new distance vector at node, N3.

(A) (3, 2, 0, 2, 5)
(B) (3, 2, 0, 2, 6)
(C) (7, 2, 0, 2, 5)
(D) (7, 2, 0, 2, 6)

Exercise

1. **Consider a network with five nodes, N1 to N5, as shown below.**

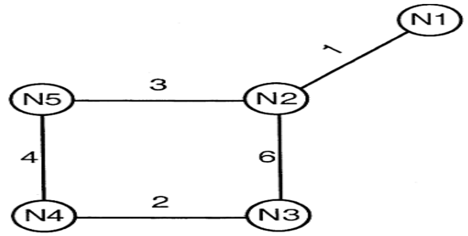

The network uses a Distance Vector Routing protocol. Once the routes have stabilized, the distance vectors at different nodes are as following. N1: (0, 1, 7, 8, 4)
N2: (1, 0, 6, 7, 3) N3: (7, 6, 0, 2, 6) N4: (8, 7, 2, 0, 4) N5: (4, 3, 6, 4, 0)
Each distance vector is the distance of the best known path at the instance to nodes, N1 to N5, where the distance to itself is 0. Also, all links are symmetric and the cost is identical in both directions. In each round, all nodes exchange their distance vectors with their respective neighbors. Then all nodes update their distance vectors. In between two rounds, any change in cost of a link will cause the two incident nodes to change only that entry in their distance vectors. 52. The cost of link N2-N3 reduces to 2(in both directions). After the next round of updates, what will be the new distance vector at node, N3.

(A) (3, 2, 0, 2, 5)
(B) (3, 2, 0, 2, 6)
(C) (7, 2, 0, 2, 5)
(D) (7, 2, 0, 2, 6)

N3: (7, 2, 0, 2, 6)

N3 knows (1,0,2,7,3) from N2 and (8,7,2,0,4) from N4.

N3: (3,2,0,2,5)

Exercise

2. Consider the same data as given in previous question.

After the update in the previous question, the link N1-N2 goes down. N2 will reflect this change immediately in its distance vector as cost, infinite. After the NEXT ROUND of update, what will be cost to N1 in the distance vector of N3?

(A) 3

(B) 9

(C) 10

(D) Infinite

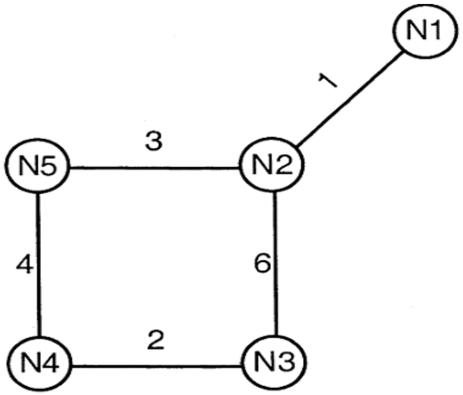

Exercise

2. Consider the same data as given in previous question.

After the update in the previous question, the link N1-N2 goes down. N2 will reflect this change immediately in its distance vector as cost, infinite. After the NEXT ROUND of update, what will be cost to N1 in the distance vector of N3?

(A) 3

(B) 9

(C) 10

(D) Infinite

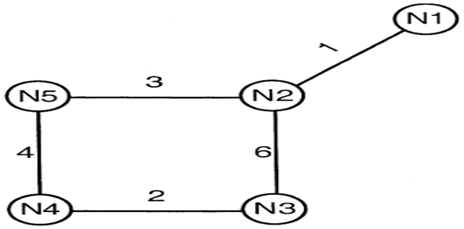

N1: (0,infinity,3,8,4)
N2: (infinity,0,2,4,3)
N3: (3,2,0,2,5)
N4: (8,4,2,0,4)
N5: (4,3,5,4,0)

After the next round of exchange, N3 knows the distance vector (infinity,0,2,4,3) from N2 and (8,4,2,0,4) from N4.

So N3 will update its distance vector to (10,2,0,2,5).

Exercise

3. Consider a network with 6 routers R1 to R6 connected with links having weights as shown in the following diagram

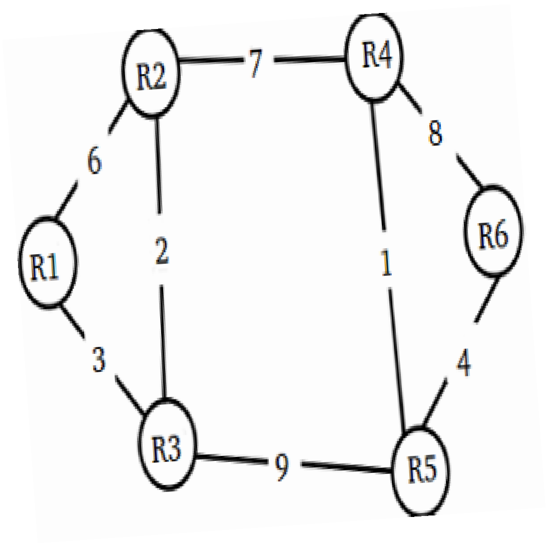

All the routers use the distance vector based routing algorithm to update their routing tables. Each router starts with its routing table initialized to contain an entry for each neighbor with the weight of the respective connecting link. After all the routing tables stabilize, how many links in the network will never be used for carrying any data?

(A) 4

(B) 3

(C) 2

(D) 1

Sol: We can check one by one all shortest distances. When we check for all shortest distances for Ri we don't need to check its distances to R0 to Ri-1 because the network graph is undirected.

Following will be distance vectors of all nodes.

Shortest Distances from R1 to R2, R3, R4, R5 and R6

R1 (5, 3, 12, 12, 16)

Links used: R1-R3, R3-R2, R2-R4, R3-R5, R5-R6

Shortest Distances from R2 to R3, R4, R5 and R6

R2 (2, 7, 8, 12)

Links used: R2-R3, R2-R4, R4-R5, R5-R6

Shortest Distances from R3 to R4, R5 and R6

R3 (9, 9, 13)

Links used: R3-R2, R2-R4, R3-R5, R5-R6

Shortest Distances from R4 to R5 and R6

R4 (1, 5)

Links used: R4-R5, R5-R6

Shortest Distance from R5 to R6

R5 (4)

Links Used: R5-R6

If we mark, all the used links one by one, we can see that following links are never used.

R1-R2

R4-R6

Exercise

3. Consider a network with 6 routers R1 to R6 connected with links having weights as shown in the following diagram

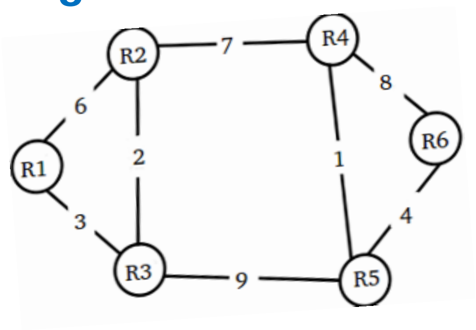

All the routers use the distance vector based routing algorithm to update their routing tables. Each router starts with its routing table initialized to contain an entry for each neighbour with the weight of the respective connecting link. After all the routing tables stabilize, how many links in the network will never be used for carrying any data?

(A) 4

(B) 3

(C) 2

(D) 1

Exercise

4. Suppose the weights of all unused links in the previous question are changed to 2 and the distance vector algorithm is used again until all routing tables stabilize. How many links will now remain unused?

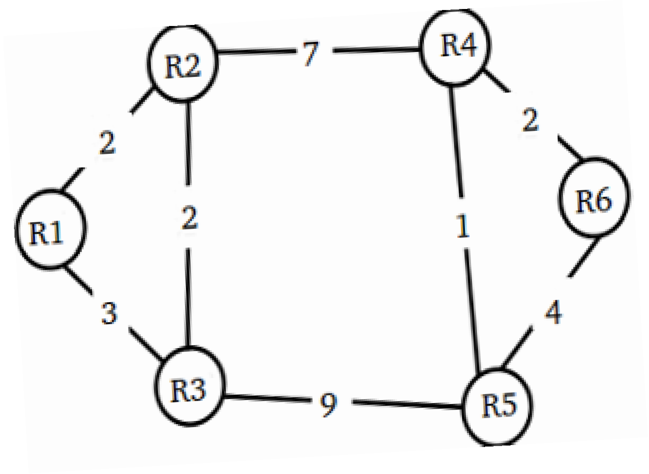

(A) 0

(B) 1

(C) 2

(D) 3

Sol: The distance vectors of all nodes
R1 (2, 3, 9, 10, 11)
Links used: R1-R2, R1-R3, R2-R4, R4-R5, R4-R6
R2 (2, 7, 8, 9)
Links used: R2-R3, R2-R4, R4-R5, R4-R6
R3 (9, 9, 11)
Links used: R3-R2, R2-R4, R3-R5, R4-R6
R4 (1, 2)
Links used: R4-R5, R4-R6
R5 (3)
Links Used: R5-R4, R4-R6
If we mark, all the used links one by one, we can see that all links are used except the following link.
R5-R6

Exercise

4. Suppose the weights of all unused links in the previous question are changed to 2 and the distance vector algorithm is used again until all routing tables stabilize. How many links will now remain unused?

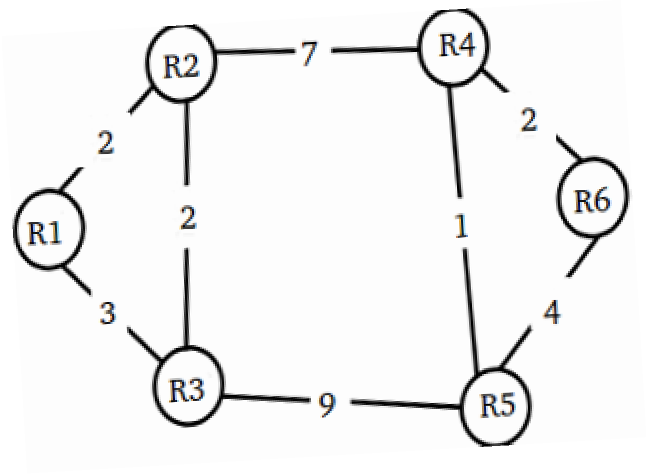

(A) 0

(B) 1

(C) 2

(D) 3

Exercise

5. Which one of the following is TRUE about interior Gateway routing protocols – Routing Information Protocol (RIP) and Open Shortest Path First (OSPF)

(A) RIP uses distance vector routing and OSPF uses link state routing

(B) OSPF uses distance vector routing and RIP uses link state routing

(C) Both RIP and OSPF use link state routing

(D) Both RIP and OSPF use distance vector routing

Exercise

5. Which one of the following is TRUE about interior Gateway routing protocols – Routing Information Protocol (RIP) and Open Shortest Path First (OSPF)

(A) RIP uses distance vector routing and OSPF uses link state routing

(B) OSPF uses distance vector routing and RIP uses link state routing

(C) Both RIP and OSPF use link state routing

(D) Both RIP and OSPF use distance vector routing

Exercise

6. Consider the following three statements about link state and distance vector routing protocols, for a large network with 500 network nodes and 4000 links.
[S1] The computational overhead in link state protocols is higher than in distance vector protocols.
[S2] A distance vector protocol (with split horizon) avoids persistent routing loops, but not a link state protocol.
[S3] After a topology change, a link state protocol will converge faster than a distance vector protocol.

Which one of the following is correct about S1, S2, and S3 ?

(A) S1, S2, and S3 are all true.

(B) S1, S2, and S3 are all false.

(C) S1 and S2 are true, but S3 is false

(D) S1 and S3 are true, but S2 is false

Exercise

6. Consider the following three statements about link state and distance vector routing protocols, for a large network with 500 network nodes and 4000 links.
[S1] The computational overhead in link state protocols is higher than in distance vector protocols.
[S2] A distance vector protocol (with split horizon) avoids persistent routing loops, but not a link state protocol.
[S3] After a topology change, a link state protocol will converge faster than a distance vector protocol.

Which one of the following is correct about S1, S2, and S3 ?

(A) S1, S2, and S3 are all true.

(B) S1, S2, and S3 are all false.

(C) S1 and S2 are true, but S3 is false

(D) S1 and S3 are true, but S2 is false

Exercise

7. What is routing algorithm used by OSPF routing protocol?

(A) Distance vector

(B) Flooding

(C) Path vector

(D) Link state

Exercise

7. What is routing algorithm used by OSPF routing protocol?

(A) Distance vector

(B) Flooding

(C) Path vector

(D) Link state

Exercise

8. In OSPF, a ………. link is a network is connected to only one router.

A. point-to-point

B. transient

C. stub

D. multipoint

Exercise

8. In OSPF, a ………. link is a network is connected to only one router.

A. point-to-point

B. transient

C. stub

D. multipoint

Exercise

9. In unicast routing, each router in the domain has a table that defines a ……… path tree to possible destinations.

A. average

B. longest

C. shortest

D. very longest

Exercise

9. In unicast routing, each router in the domain has a table that defines a .... path tree to possible destinations.

A. average

B. longest

C. shortest

D. very longest

Exercise

10. In OSPF, a ……… link is a network with several routers attached to it.

A. point-to-point

B. transient

C. stub

D. multipoint

Exercise

10. In OSPF, a ……… link is a network with several routers attached to it.

A. point-to-point

B. transient

C. stub

D. multipoint

Transport Layer

PROCESS-TO-PROCESS DELIVERY

- *The transport layer is responsible for process-to-process delivery—the delivery of a message, from one process to another.*
- *Two processes communicate in a client/server relationship.*

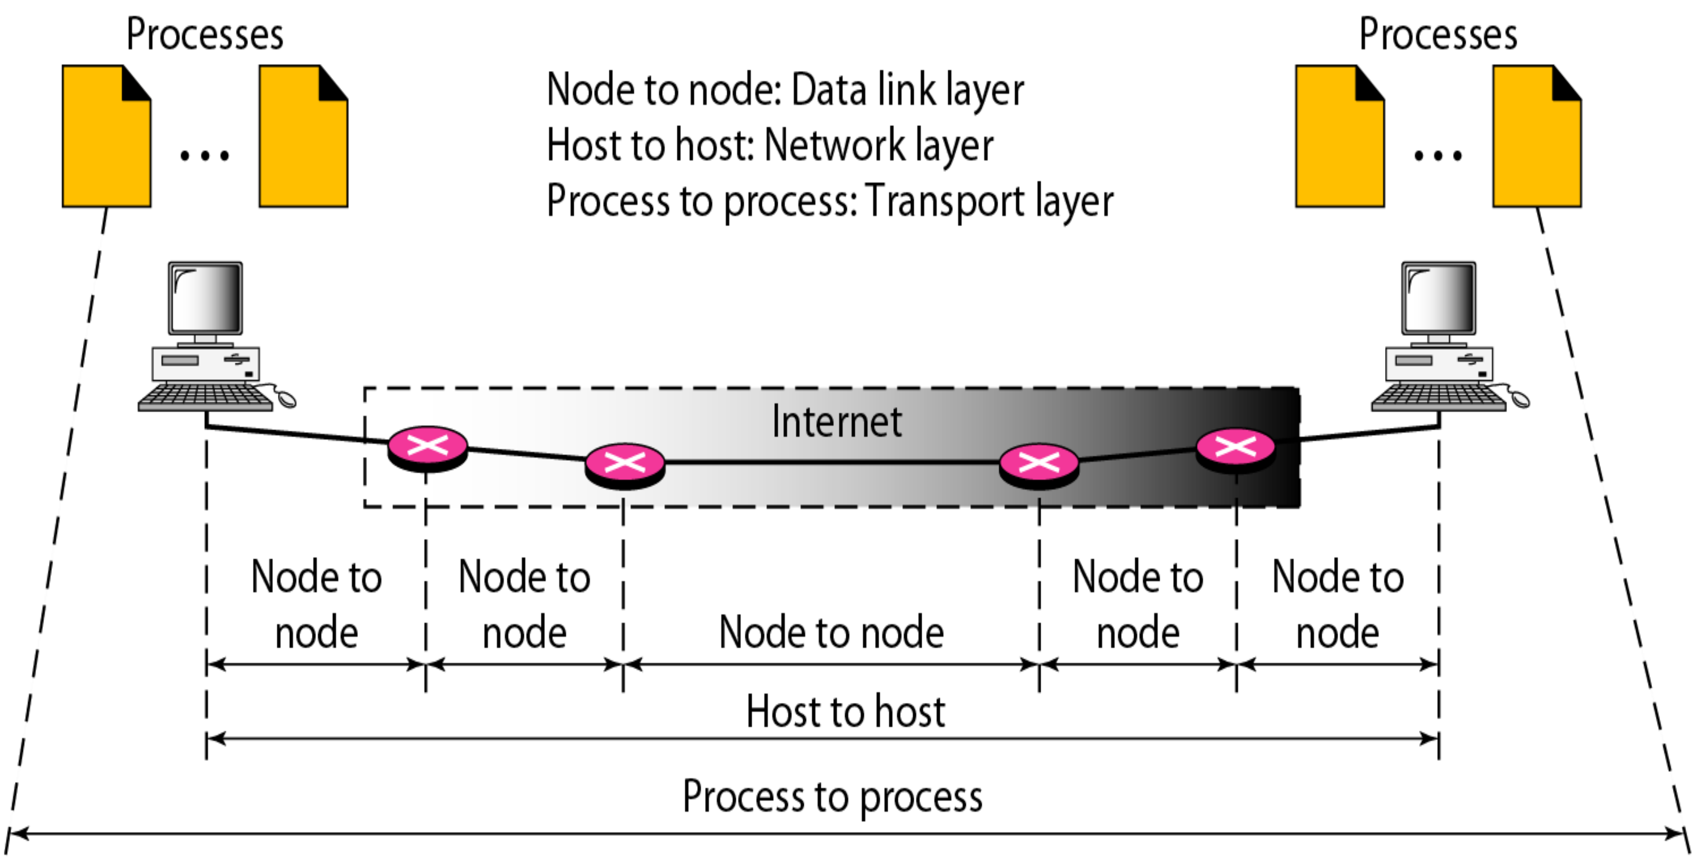

Client/Server Paradigm

Port numbers

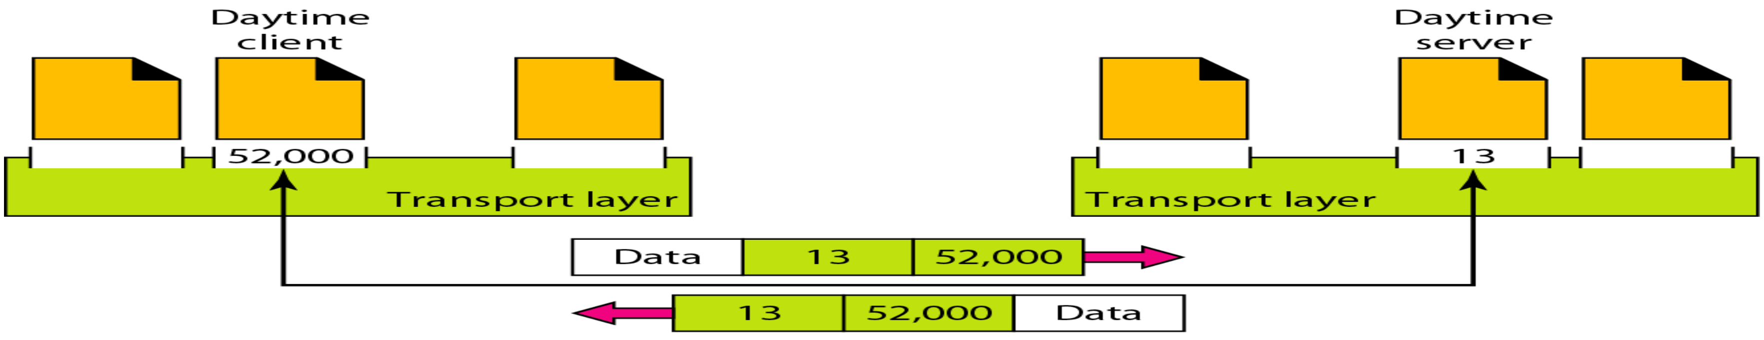

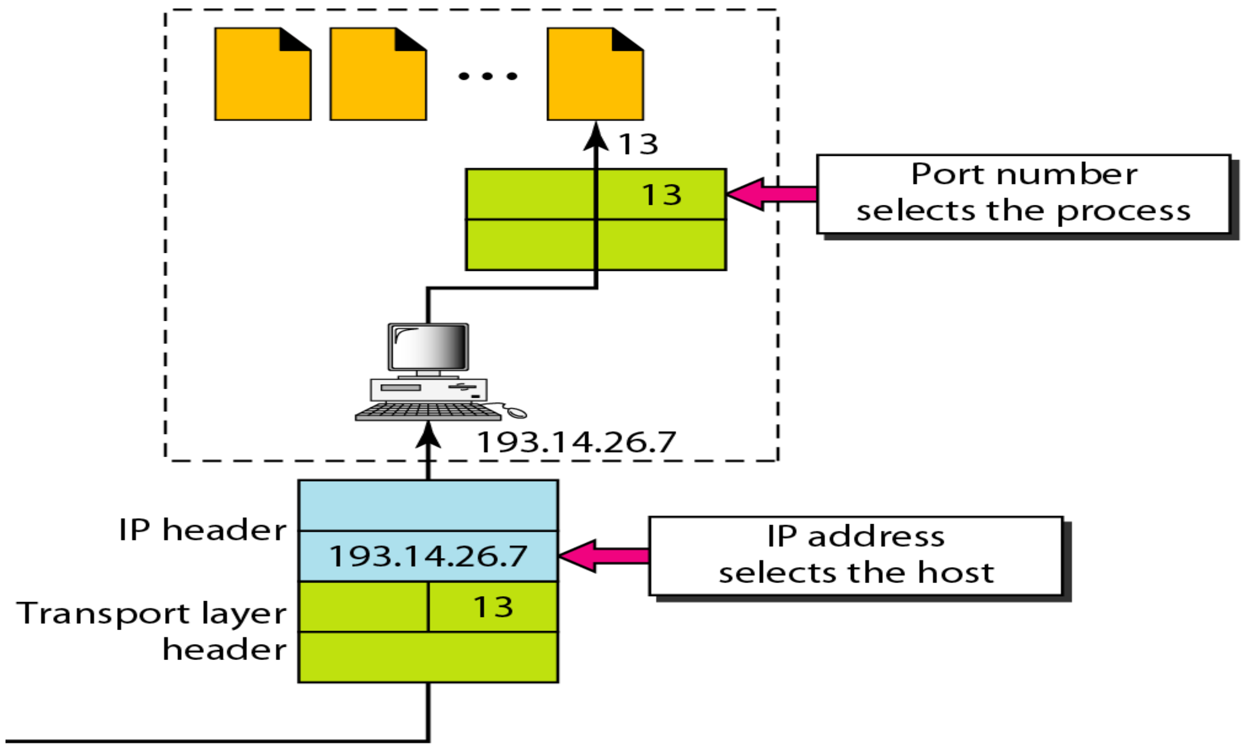

IP addresses versus port numbers

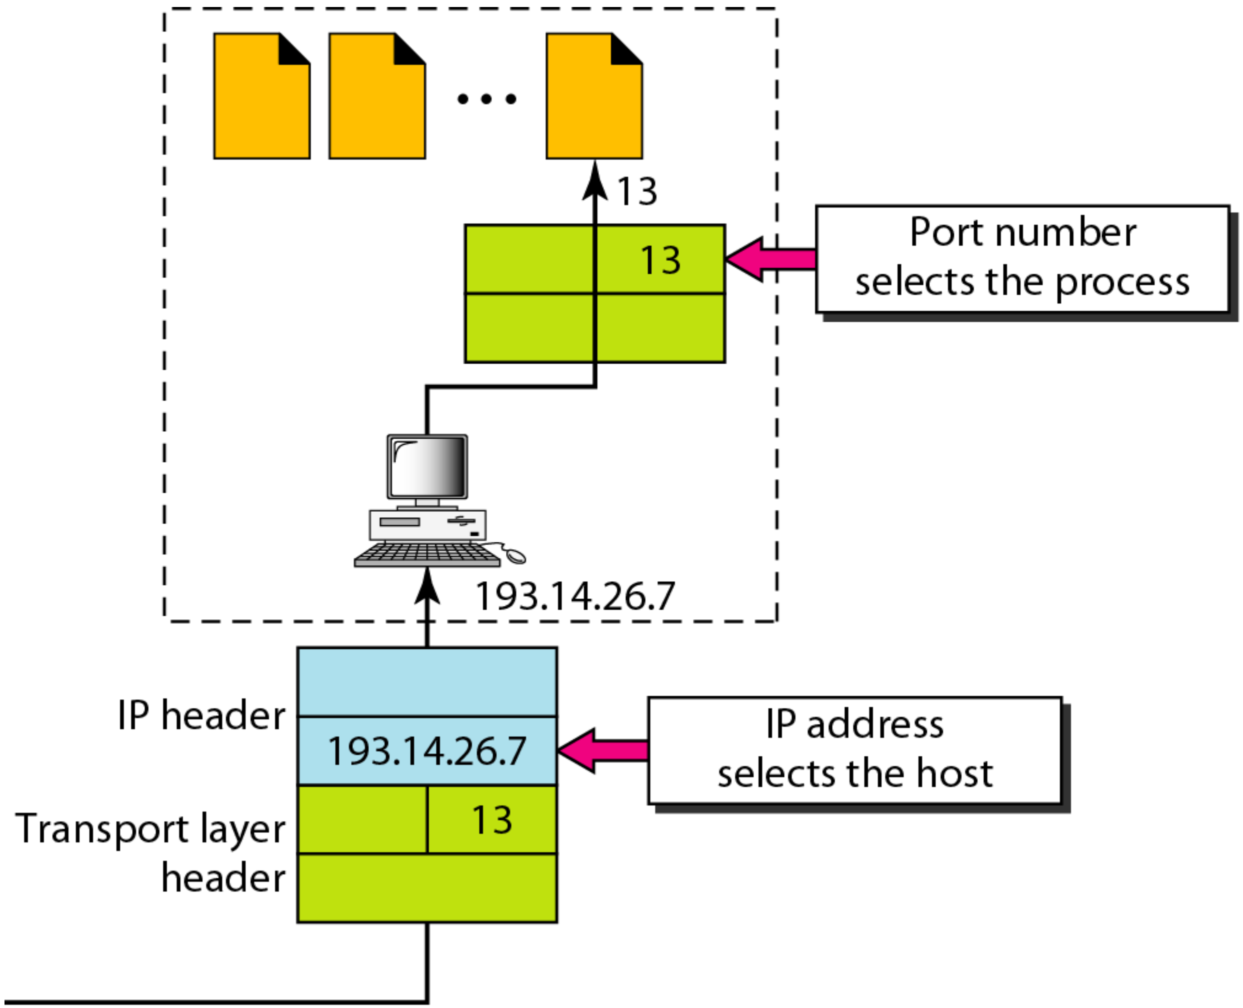

Port Numbers: IANA ranges

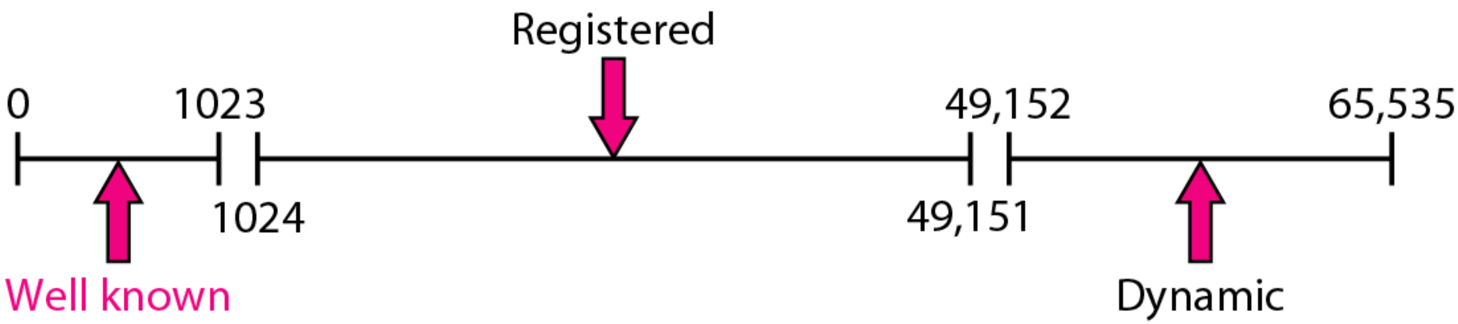

Socket address

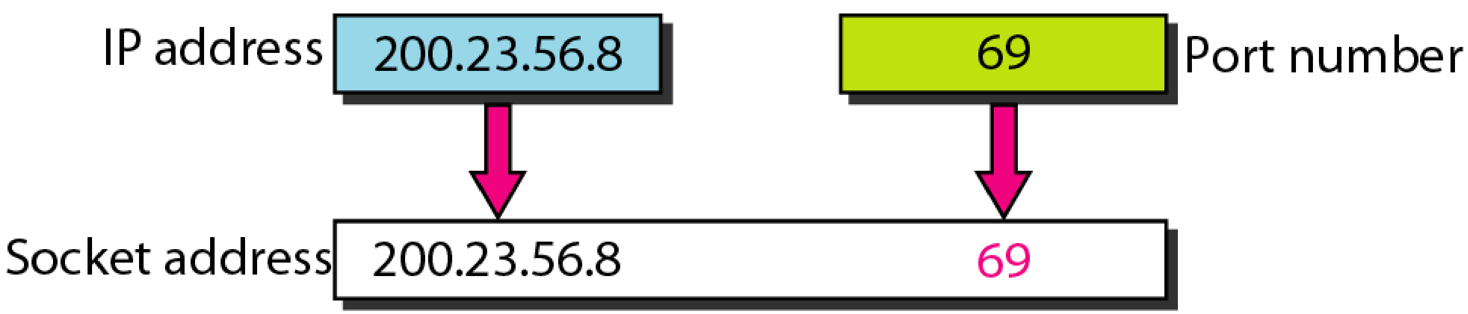

Multiplexing and demultiplexing

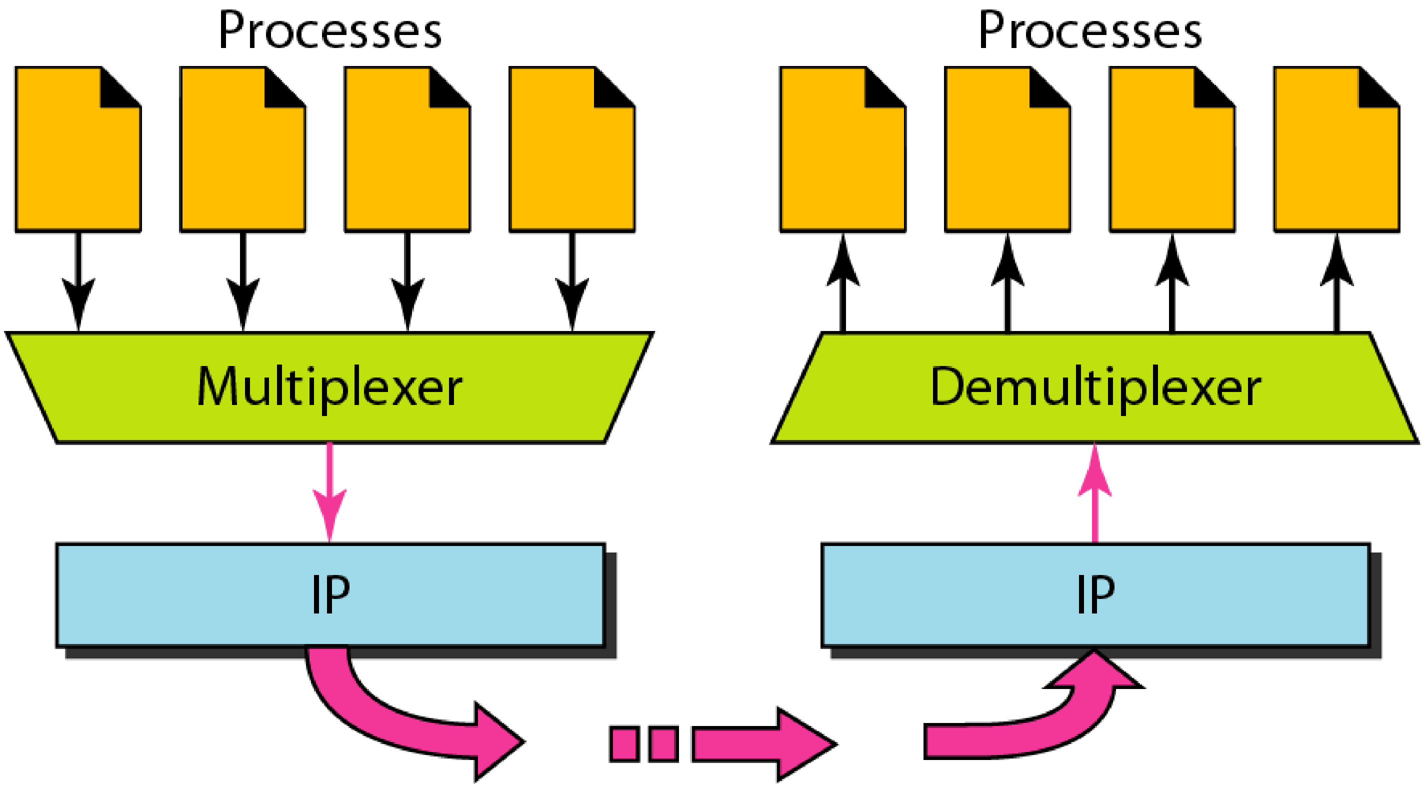

Error control

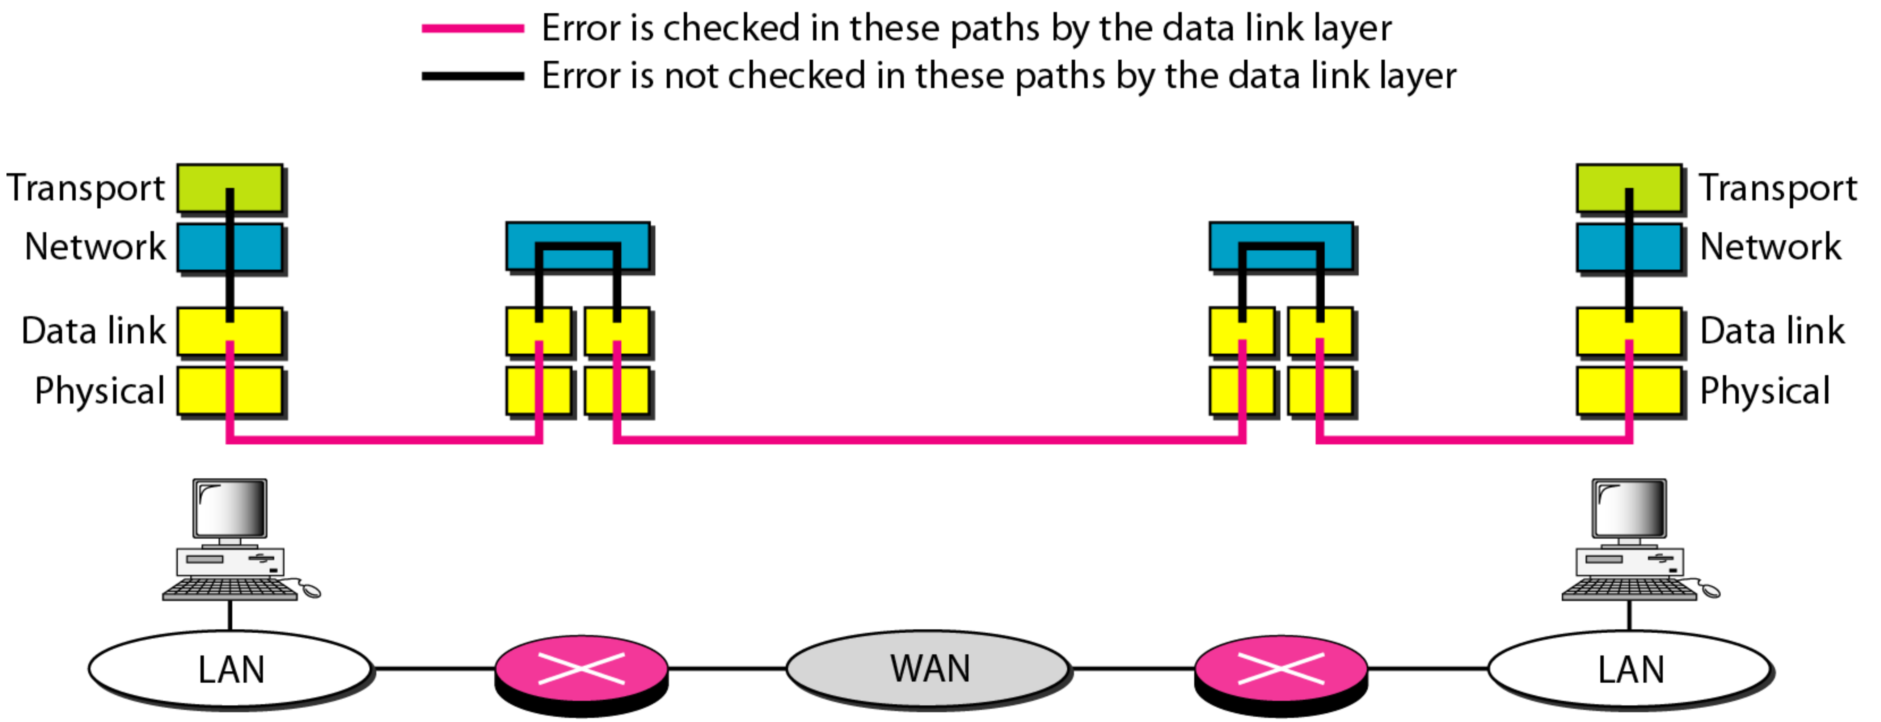

Position of UDP, TCP, and SCTP in TCP/IP suite

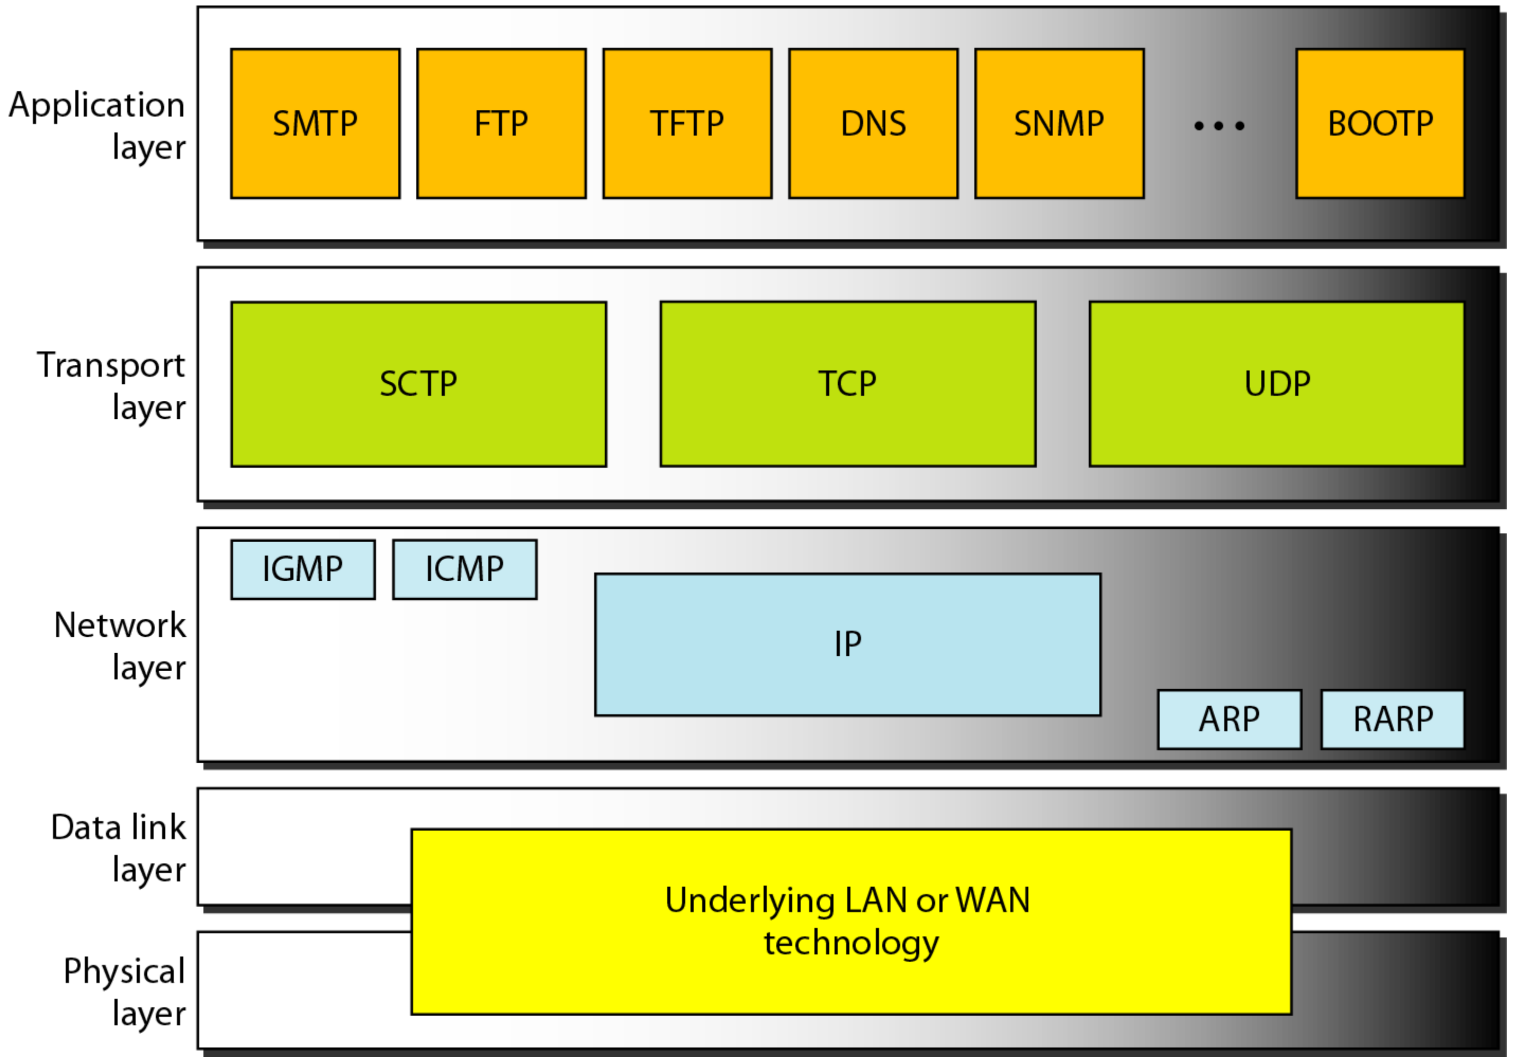

USER DATAGRAM PROTOCOL (UDP)

- *The User Datagram Protocol (UDP) is called a connectionless, unreliable transport protocol.*
- *It does not add anything to the services of IP except to provide process-to-process communication instead of host-to-host communication.*

Well-known ports used with UDP

Port	*Protocol*	*Description*
7	Echo	Echoes a received datagram back to the sender
9	Discard	Discards any datagram that is received
11	Users	Active users
13	Daytime	Returns the date and the time
17	Quote	Returns a quote of the day
19	Chargen	Returns a string of characters
53	Nameserver	Domain Name Service
67	BOOTPs	Server port to download bootstrap information
68	BOOTPc	Client port to download bootstrap information
69	TFTP	Trivial File Transfer Protocol
111	RPC	Remote Procedure Call
123	NTP	Network Time Protocol
161	SNMP	Simple Network Management Protocol
162	SNMP	Simple Network Management Protocol (trap)

User datagram format

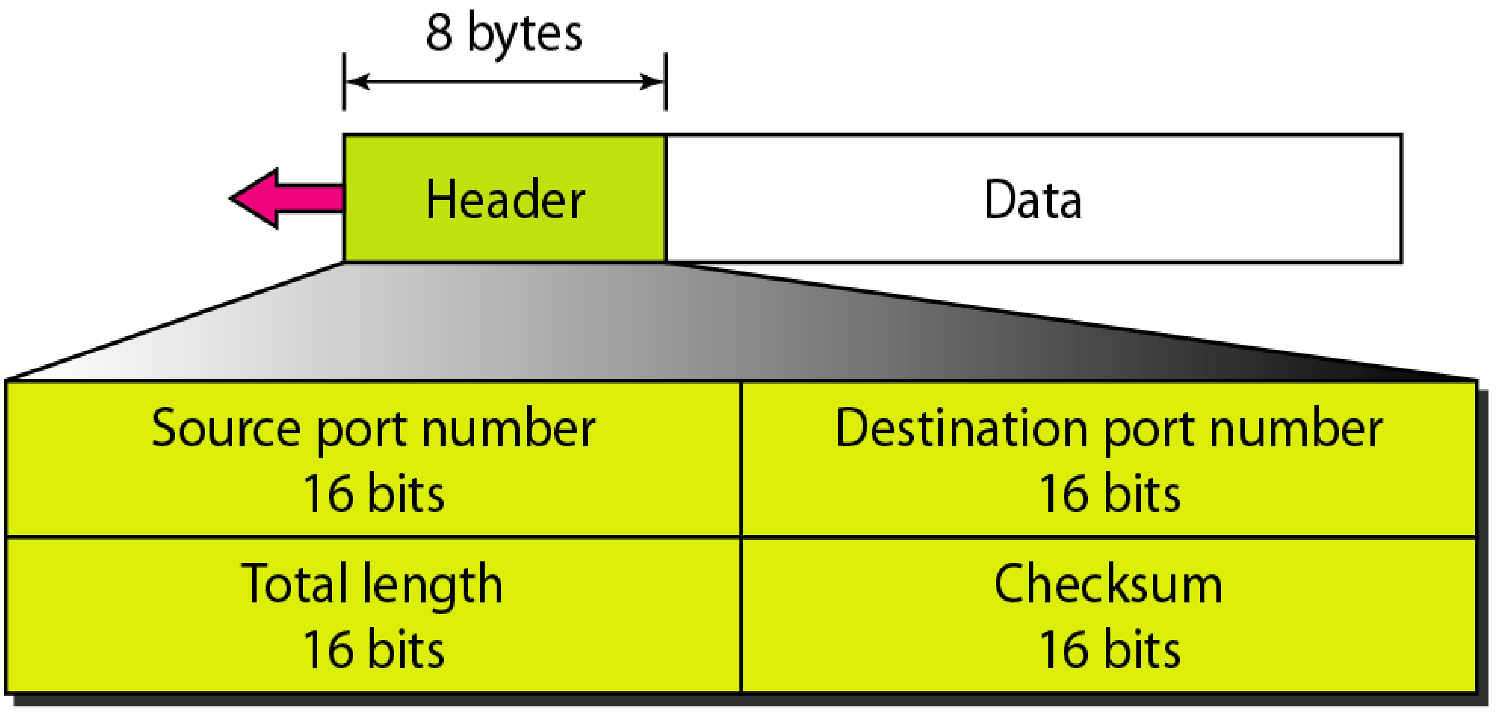

UDP length = IP length – IP header's length

Pseudoheader for checksum calculation

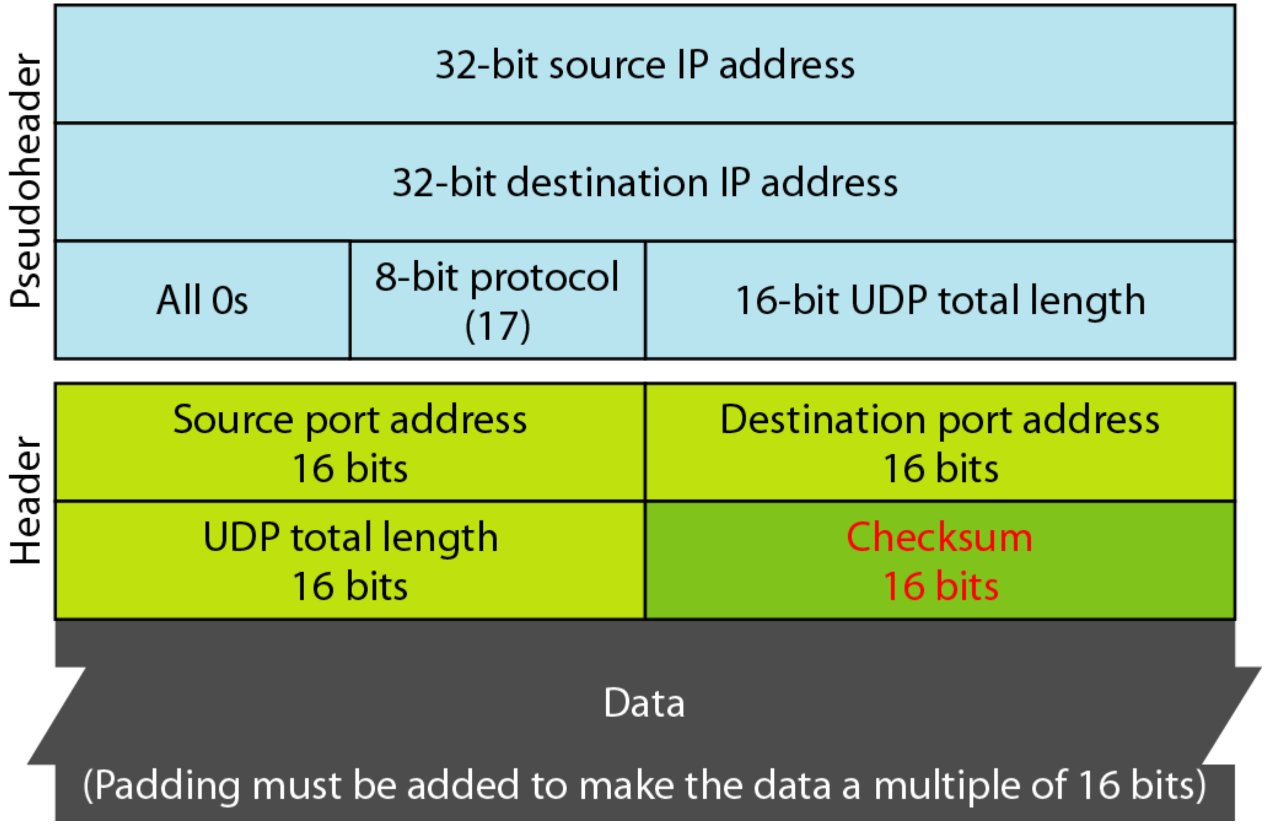

Checksum calculation of a simple UDP user datagram

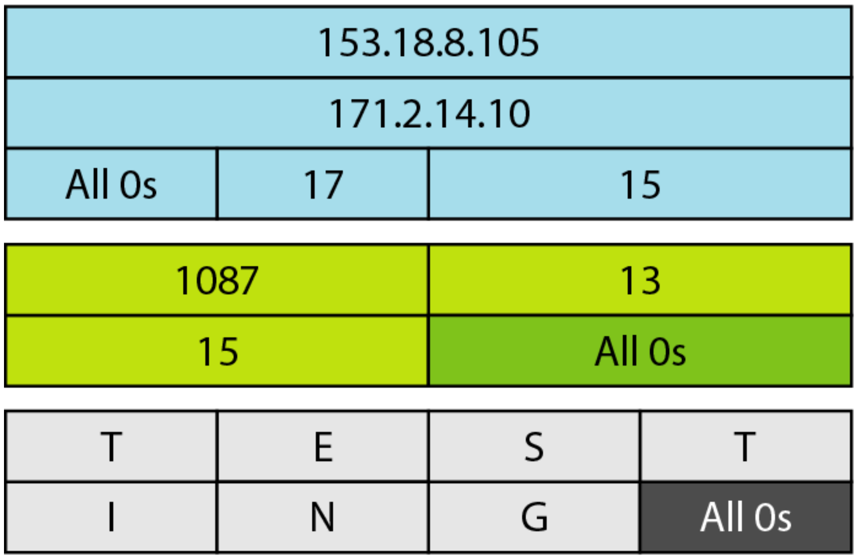

10011001 00010010	→	153.18
00001000 01101001	→	8.105
10101011 00000010	→	171.2
00001110 00001010	→	14.10
00000000 00010001	→	0 and 17
00000000 00001111	→	15
00000100 00111111	→	1087
00000000 00001101	→	13
00000000 00001111	→	15
00000000 00000000	→	0 (checksum)
01010100 01000101	→	T and E
01010011 01010100	→	S and T
01001001 01001110	→	I and N
01000111 00000000	→	G and 0 (padding)
10010110 11101011	→	Sum
01101001 00010100	→	Checksum

Queues in UDP

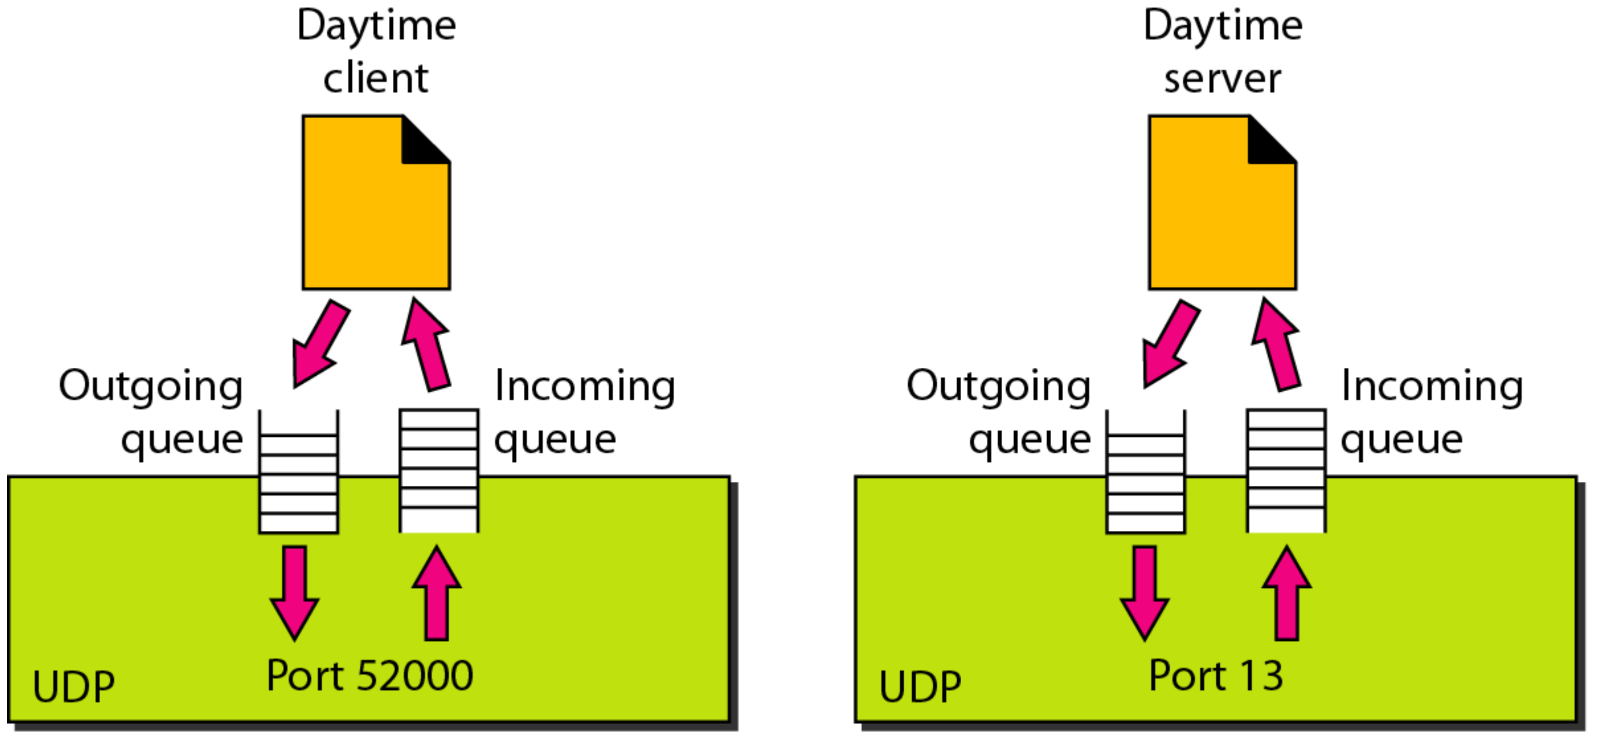

Uses of UDP

• Suitable for a process that requires simple Request-Response communication with little concern on flow & error Control

• TFTP includes internal flow and error control

•Suitable for Multicasting

•Management process such as SNMP

•Used for Rote update protocols like RIP

Socket application program interfaces (APIs) are the network standard for TCP/IP.

UDP Sockets: System Call

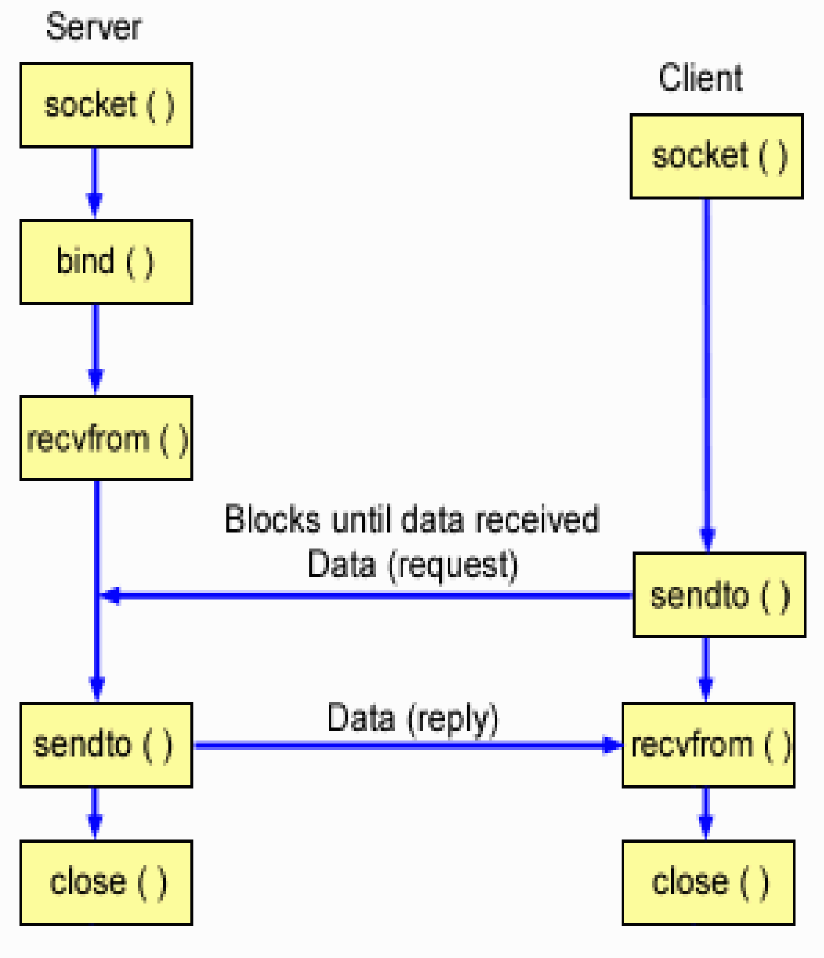

TCP Sockets: System Call

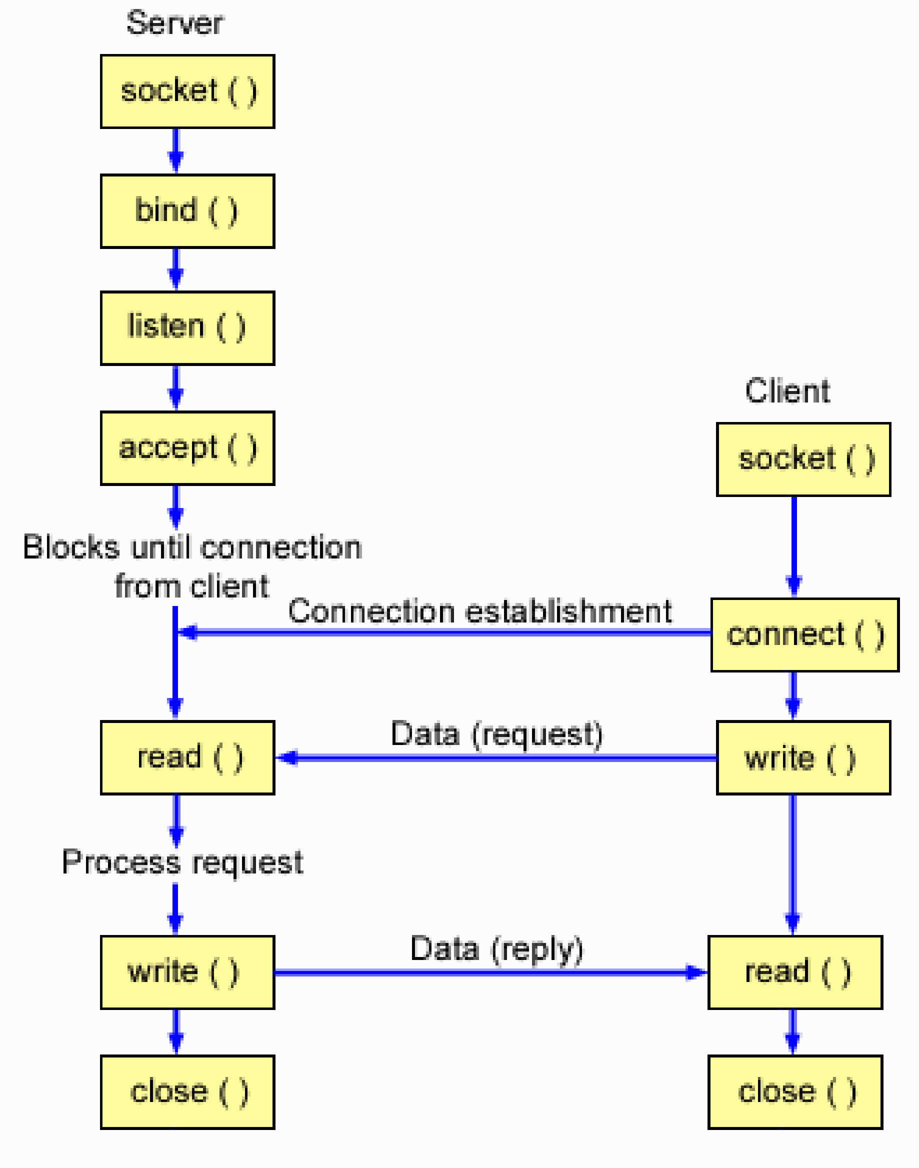

Exercise

1. Which of the following functionalities must be implemented by a transport protocol over and above the network protocol?

   (A) Recovery from packet losses

   (B) Detection of duplicate packets

   (C) Packet delivery in the correct order

   (D) End to end connectivity

Exercise

1. **Which of the following functionalities must be implemented by a transport protocol over and above the network protocol?**

   (A) Recovery from packet losses

   (B) Detection of duplicate packets

   (C) Packet delivery in the correct order

   (D) End to end connectivity

Exercise

2. The transport layer protocols used for real time multimedia, file transfer, DNS and email, respectively are:

(A) TCP, UDP, UDP and TCP

(B) UDP, TCP, TCP and UDP

(C) UDP, TCP, UDP and TCP

(D) TCP, UDP, TCP and UDP

Exercise

2. The transport layer protocols used for real time multimedia, file transfer, DNS and email, respectively are:

(A) TCP, UDP, UDP and TCP

(B) UDP, TCP, TCP and UDP

(C) UDP, TCP, UDP and TCP

(D) TCP, UDP, TCP and UDP

Exercise

3. Which one of the following uses UDP as the transport protocol?

(A) HTTP

(B) Telnet

(C) DNS

(D) SMTP

Exercise

3. Which one of the following uses UDP as the transport protocol?

(A) HTTP

(B) Telnet

(C) DNS

(D) SMTP

Exercise

4. Packets of the same session may be routed through different paths in:

(a) TCP, but not UDP

(b) TCP and UDP

(c) UDP, but not TCP

(d) Neither TCP nor UDP

Exercise

4. Packets of the same session may be routed through different paths in:

(a) TCP, but not UDP

(b) TCP and UDP

(c) UDP, but not TCP

(d) Neither TCP nor UDP

Exercise

5. The receiver of the data controls the amount of data that are to be sent

by the sender is referred to as ___________

a) Flow control

b) Error control

c) Congestion control

d) Error detection

Exercise

5. The receiver of the data controls the amount of data that are to be sent by the sender is referred to as ____________

a) Flow control

b) Error control

c) Congestion control

d) Error detection

Exercise

6. Which of the following is NOT true about User Datagram Protocol in transport layer?

a) **Works well in unidirectional communication, suitable for broadcast information.**

b) **It does three way handshake before sending datagrams**

c) **It provides datagrams, suitable for modeling other protocols such as in IP tunneling or Remote Procedure Call and the Network File System**

d) **The lack of retransmission delays makes it suitable for real-time applications**

Exercise

6. Which of the following is NOT true about User Datagram Protocol in transport layer?

a) Works well in unidirectional communication, suitable for broadcast information.

b) It does three way handshake before sending datagrams

c) It provides datagrams, suitable for modeling other protocols such as in IP tunneling or Remote Procedure Call and the Network File System

d) The lack of retransmission delays makes it suitable for real-time applications

Exercise

7. Which one of the following socket API functions converts an unconnected active TCP socket into a passive socket?

a) Connect

b) Bind

c) Listen

d) Accept

Exercise

7. Which one of the following socket API functions converts an unconnected active TCP socket into a passive socket?

a) **Connect**

b) **Bind**

c) **Listen**

d) **Accept**

Exercise

8. Identify the correct order in which a server process must invoke the function calls accept, bind, listen, and recv according to UNIX socket API

a) listen, accept, bind, recv

b) bind, listen, accept, recv

c) bind, accept, listen, recv

d) accept, listen, bind, recv

Exercise

8. Identify the correct order in which a server process must invoke the function calls accept, bind, listen, and recv according to UNIX socket API

a) listen, accept, bind, recv

b) bind, listen, accept, recv

c) bind, accept, listen, recv

d) accept, listen, bind, recv