Routing Protocols

Introduction

- Placing a datagram in its route to its destination.
- Routing requires a host or a router to have a routing table.
- When a host has a datagram to send or when a router has received a datagram to be forwarded, it looks at this table to find the route to the final destination.

Forwarding versus Routing

- Forwarding
  - To select an output port based on destination address and routing table
- Routing
  - Process by which routing table is built

Forwarding table VS Routing table

- Forwarding table
  - Used when a packet is being forwarded and so must contain enough information to accomplish the forwarding function
  - A row in the forwarding table contains the mapping from a network number to an outgoing interface and some MAC information, such as Ethernet Address of the next hop
- Routing table
  - Built by the routing algorithm as a precursor to build the forwarding table
  - Generally contains mapping from network numbers to next hops

Forwarding table VS Routing table

(a)

Prefix/Length	Next Hop
18/8	171.69.245.10

(b)

Prefix/Length	Interface	MAC Address
18/8	if0	8:0:2b:e4:b:1:2

(a) Routing table

(b) Forwarding table

Interdomain vs Intradomain Routing

- Autonomous System
  - A group of networks and routers under the authority of a single administration.
- Intradomain Routing
  - Routing Inside autonomous systems
- Interdomain Routing
  - Routing between autonomous systems.

Interdomain vs Intradomain Routing

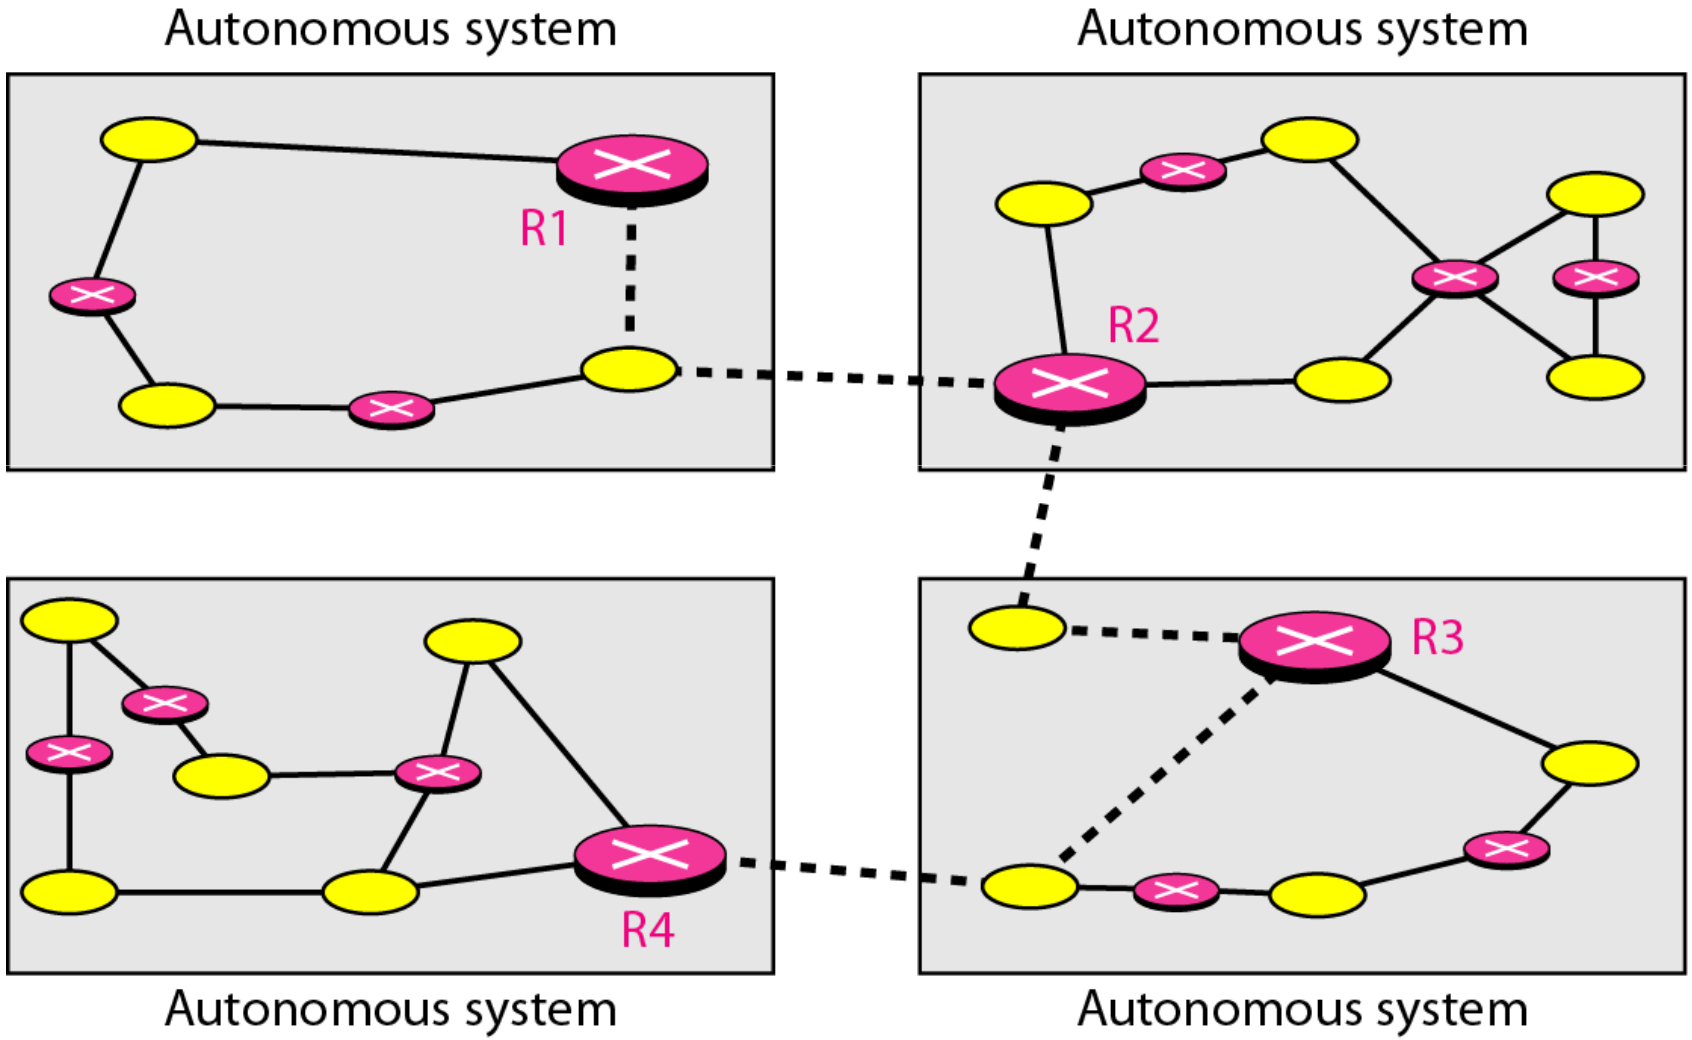

Network as a Graph

- The network is considered as a graph
- The basic problem of routing is to find the lowest-cost path between any two nodes
  - Where the cost of a path equals the sum of the costs of all the edges that make up the path

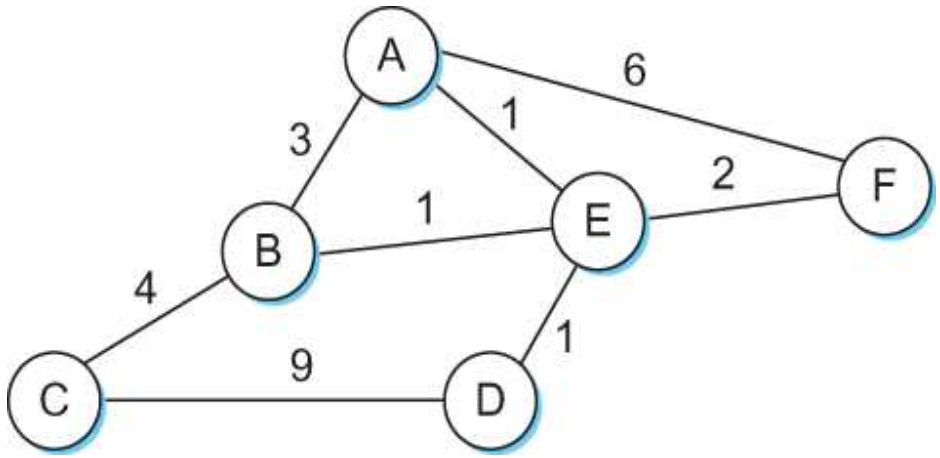

Network as a Graph

- For a simple network, calculate all shortest paths and load them into some nonvolatile storage on each node.
- Such a static approach has several shortcomings
  - It does not deal with node or link failures
  - It does not consider the addition of new nodes or links
  - It implies that edge costs cannot change

Network as a Graph

- What is the solution?
  - Need a distributed and dynamic protocol

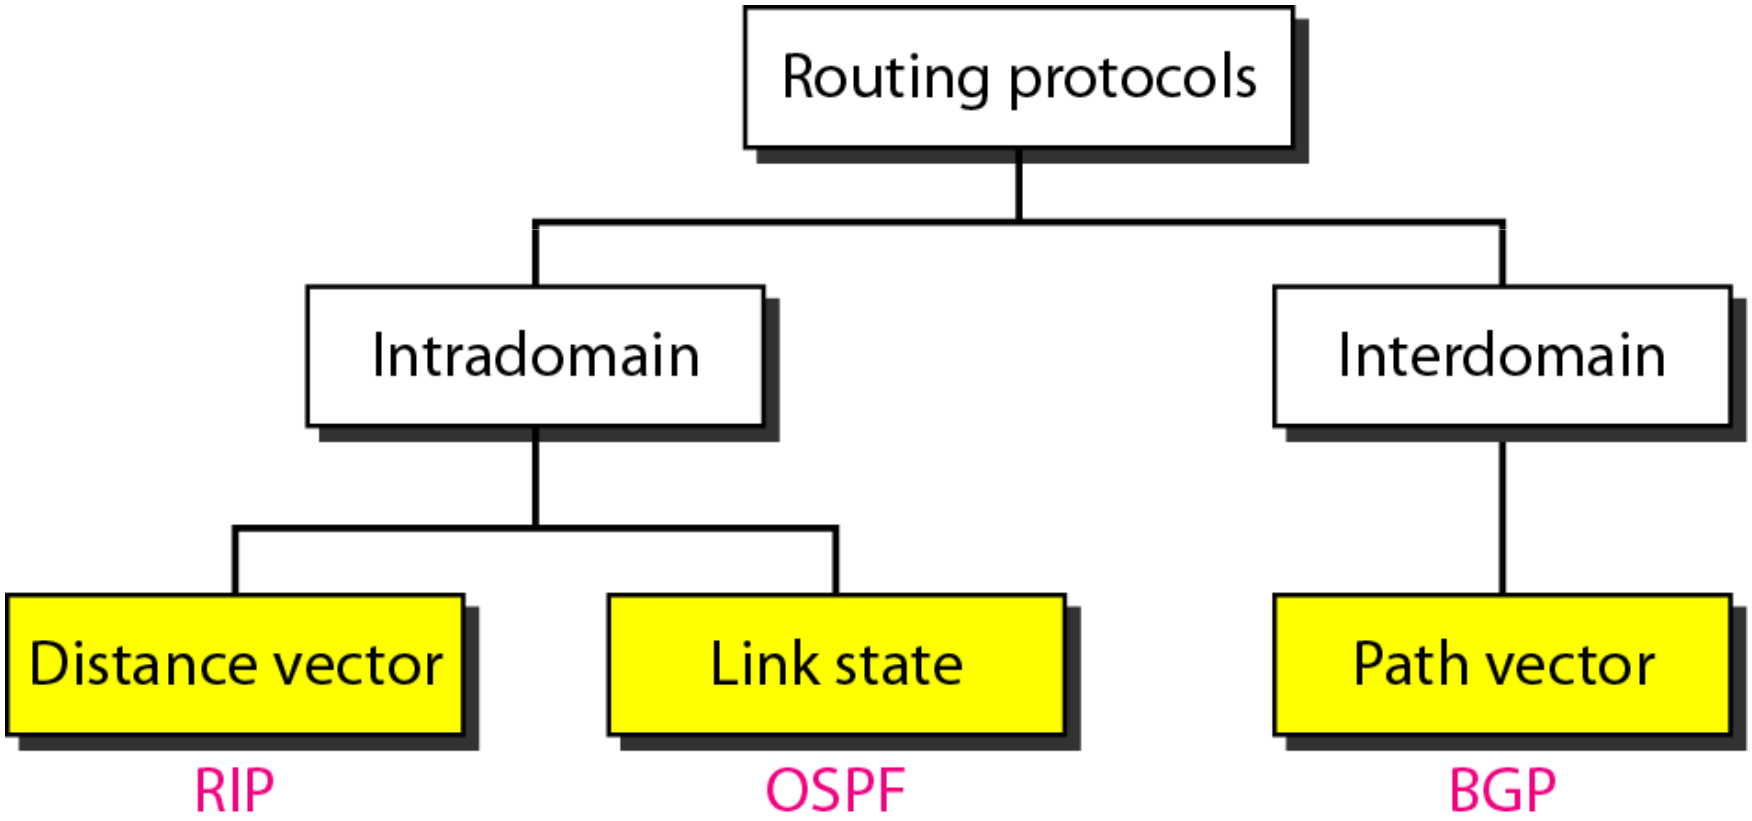

Distance Vector Routing

- Commonly called as Distributed Bellman Ford Routing Algorithm and the Ford-Fulkerson Algorithm
- Completely decentralized algorithm
  - No node has complete information about the costs of all network links
- Gradual calculation of path by exchanging information with neighbors

Distance Vector Routing

- Each node constructs a one-dimensional array (vector) containing the distances (costs) to all other nodes (as it relates to its knowledge) and distributes it to its immediate neighbors.
- Each node knows the cost of links to its immediate neighbor.
- If no link exists between two nodes, the cost between the nodes is marked as infinity.

Representation

- The cost from X to Y via Z is the cost from X to Z plus the "minimum" cost from Z to Y.

- $D^x(Y,Z) = C(X, Z) + \min_W \{ D^z(Y,W) \}$
  - $D^x(Y,Z)$ → Cost from X to Y via Z
  - $C(X, Z)$ → Cost from X to Z
  - $\min_W \{ D^z(Y,W) \}$ → Minimum cost from Z to Y via W.

- Minimum cost from Z to Y computed by taking all possible paths into consideration.

Routing Table

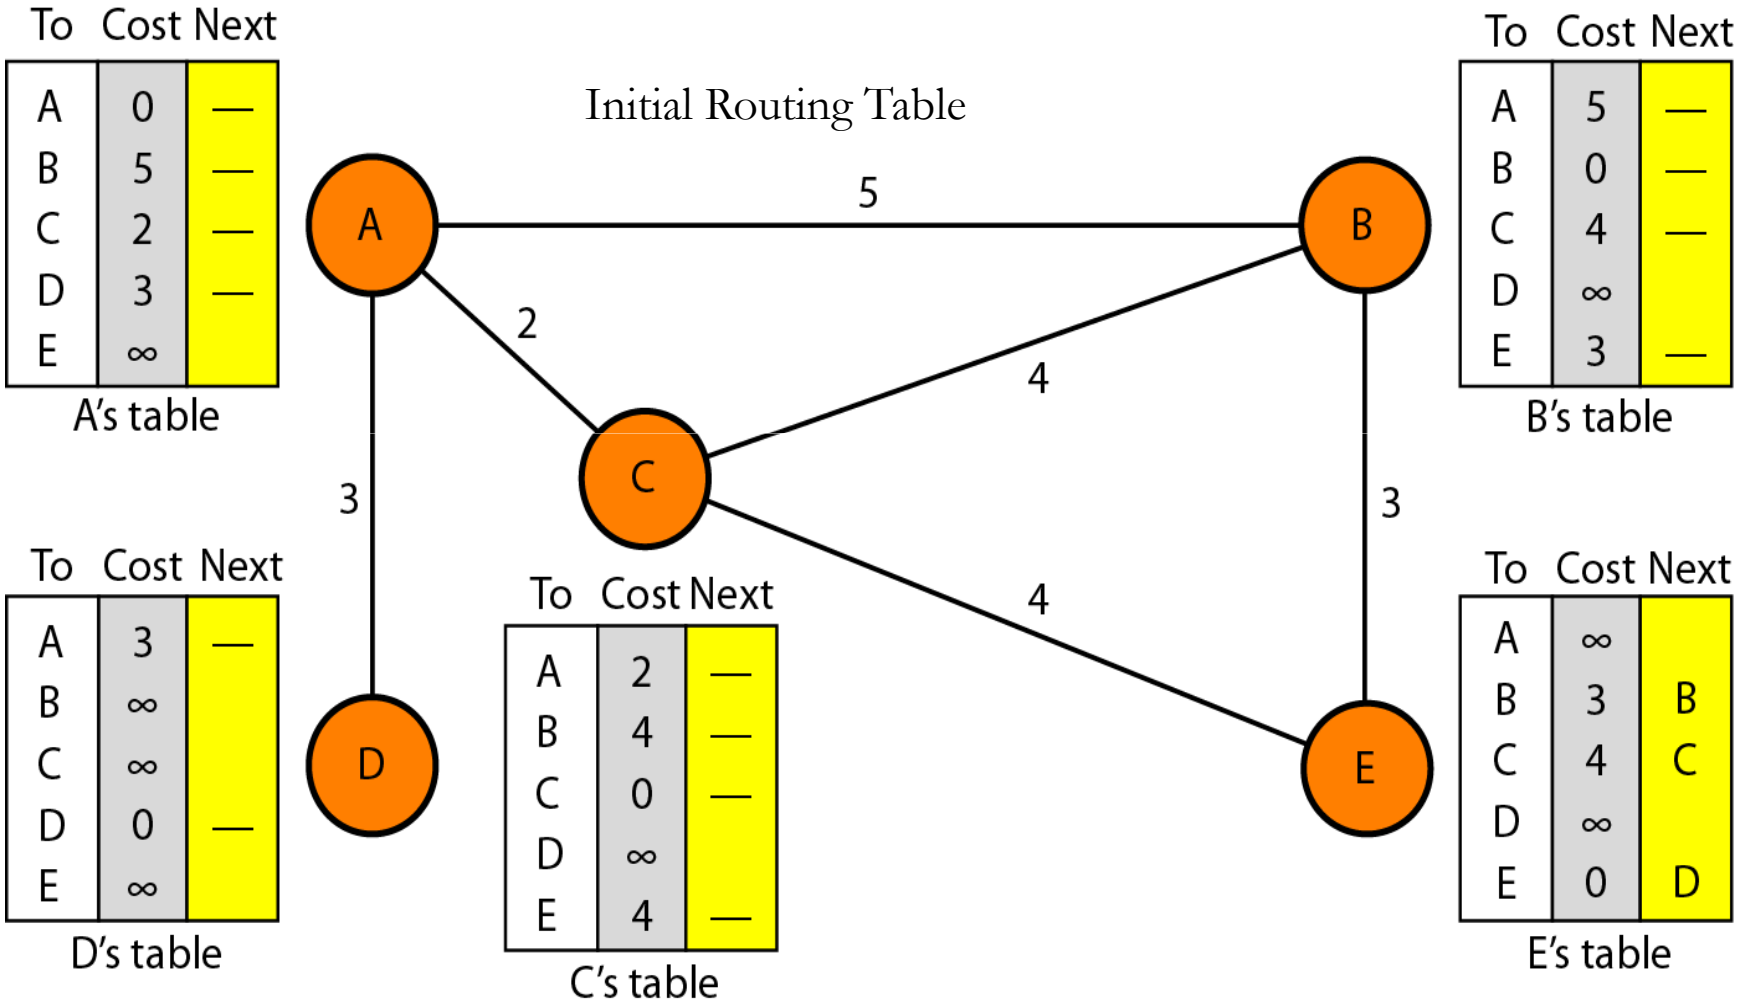

Routing Table

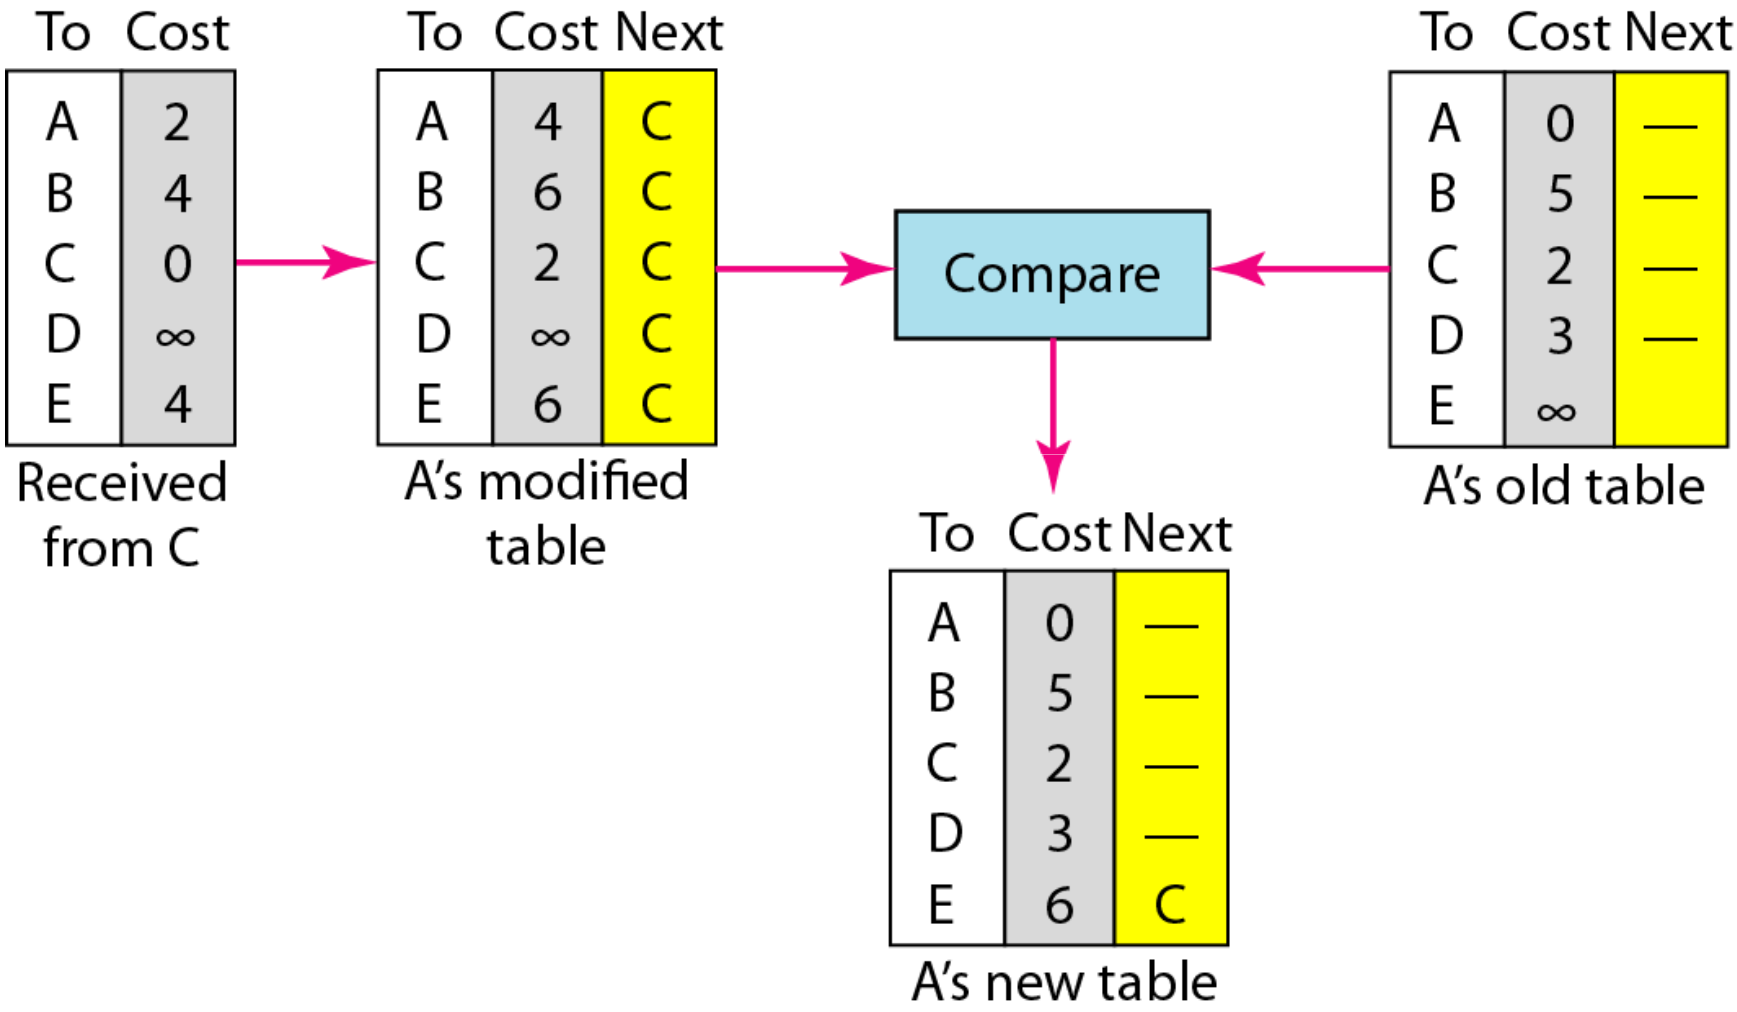

Routing Table

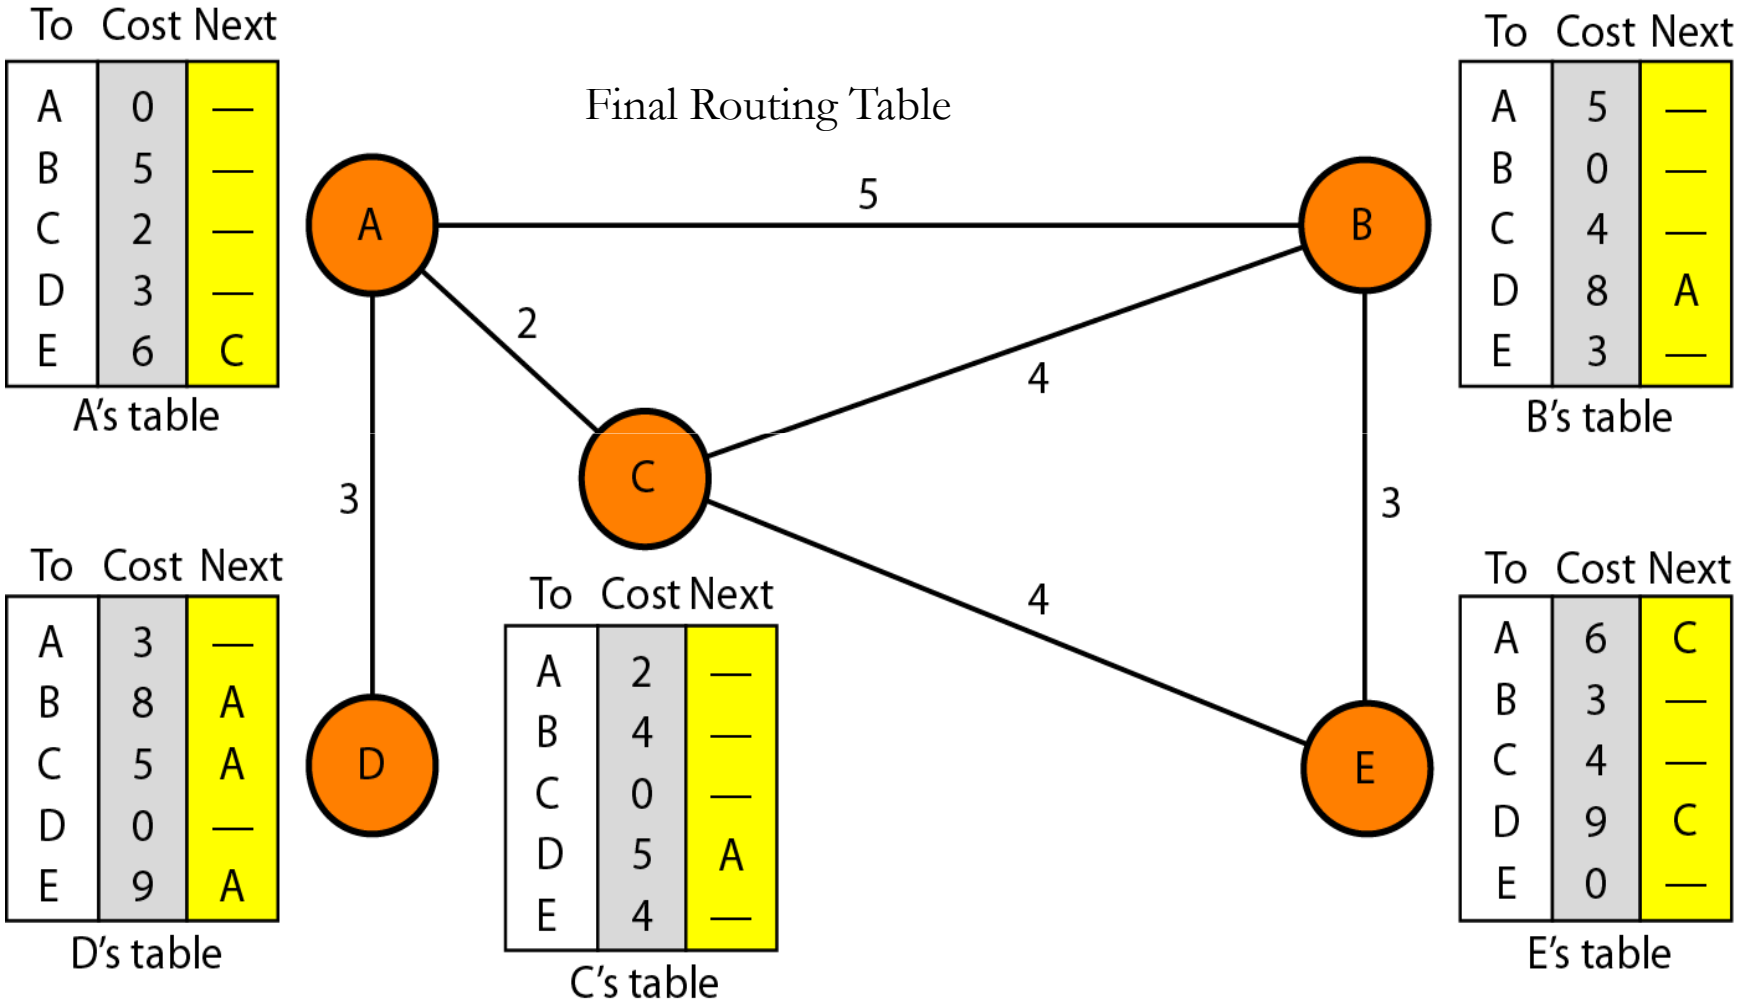

Convergence

- In the absence of topological changes -- few exchanges between neighbors before complete routing table is formed.
- As this table is consistent, Convergence is achieved.
- Notice -- no centralized authority
- There is no one node in the network that has all the information in the table—each node only knows about the contents of its own routing table

Routing Updates

- When are routing updates sent ?
- Periodic Updates
  - Even if nothing has changed, send periodically. Every T seconds.
  - Main reason is to let other nodes know that the sender is alive.
  - Refresh information that might be needed if some of the current routes become unavailable.
- Triggered Updates
  - When a node receives an update from one of its neighbors which may lead to a change in its routing tables (could be due to change in link cost).
- Note: typically order of periodicity is seconds to several minutes.

Link/Node Failures

- Nodes that notice a failure, first send new lists of distances to neighbors. And the system settles down quickly.
- How do nodes detect failure ?
  - A node continually tests the link to another node by sending a control packet and seeing if it receives an acknowledgment
  - A node determines that the link (or the node at the other end of the link) is down if it does not receive the expected periodic routing update for the last few update cycles

What happens when link fails?

- Let link from F to G fail.
- F sets new distance to G to ∞ ; sends update to A.
- A was initially routing to G via F. So it now sets link cost to G to ∞ .
- Next update from C; A learns that C has 2 hop path to G.
- A now can reach G in 3 hops via C.
- A sends an update to F.
- Thus, F now, can reach G via A in 4 hops.

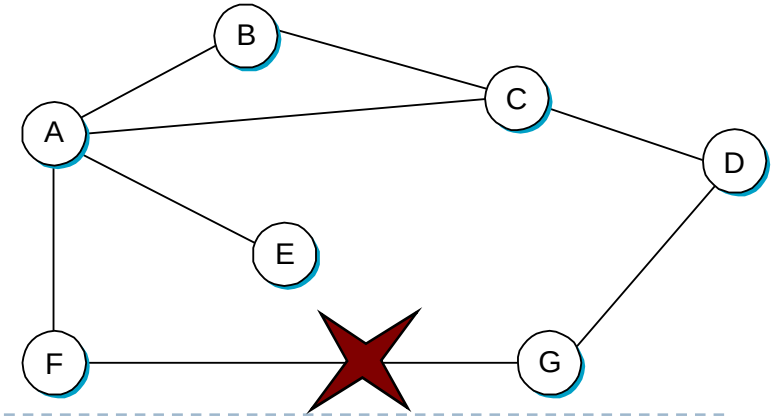

Count to Infinity

- A goes down.
- None of the nodes actually knows that A is unreachable
- This situation is known as the count-to-infinity problem.

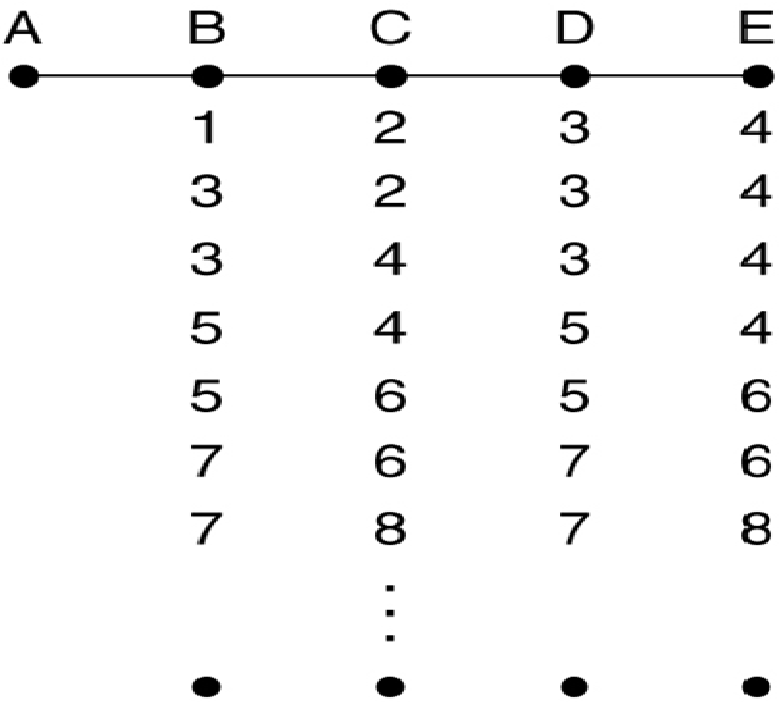

B	C	D	E	
1	2	3	4	Initially
3	2	3	4	After 1 exchange
3	4	3	4	After 2 exchanges
5	4	5	4	After 3 exchanges
5	6	5	6	After 4 exchanges
7	6	7	6	After 5 exchanges
7	8	7	8	After 6 exchanges
	⋮			

Split Horizon

- One solution would be to approximate ∞ to say 16 hops.
- When a node sends a routing table update to its neighbors, it "does not" send those routes it learned from "a particular" neighbor, back to that neighbor.
  - For example, C had A 2 B. When it sends a route update to B, it does not include this.
- In split horizon with poison reverse, this update is reported but the link weight is set to ∞
  - For example C sends (A, ∞) to B.

Does this work ?

- Works only if there is a loop with only 2 nodes.
- Typically, in static networks where link failures/node failures are rare, this may be enough.

Purging Routing Entries

- Each routing entry has a time-to-live (or TTL) field.
- A counter, initially set to MaxTTL.
- This is then decremented and if TTL = 0, then, time to purge the entry.

Routing Information Protocol

- An Example of Distance Vector Routing.
- In the Internet, goal of routers is to learn how to forward packets to various networks.

An Example of RIP

- Routers advertise the cost of reaching networks.
- In this example, C's update to A would indicate that C can reach Networks 2 and 3 with cost 0, Networks 5 and 6 with cost 1 and Network 4 with cost 2.

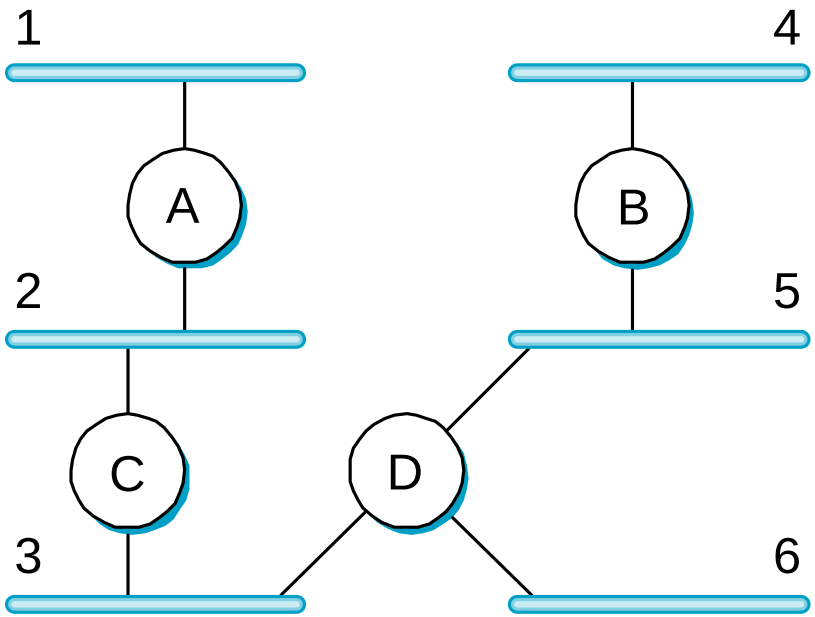

Another Example

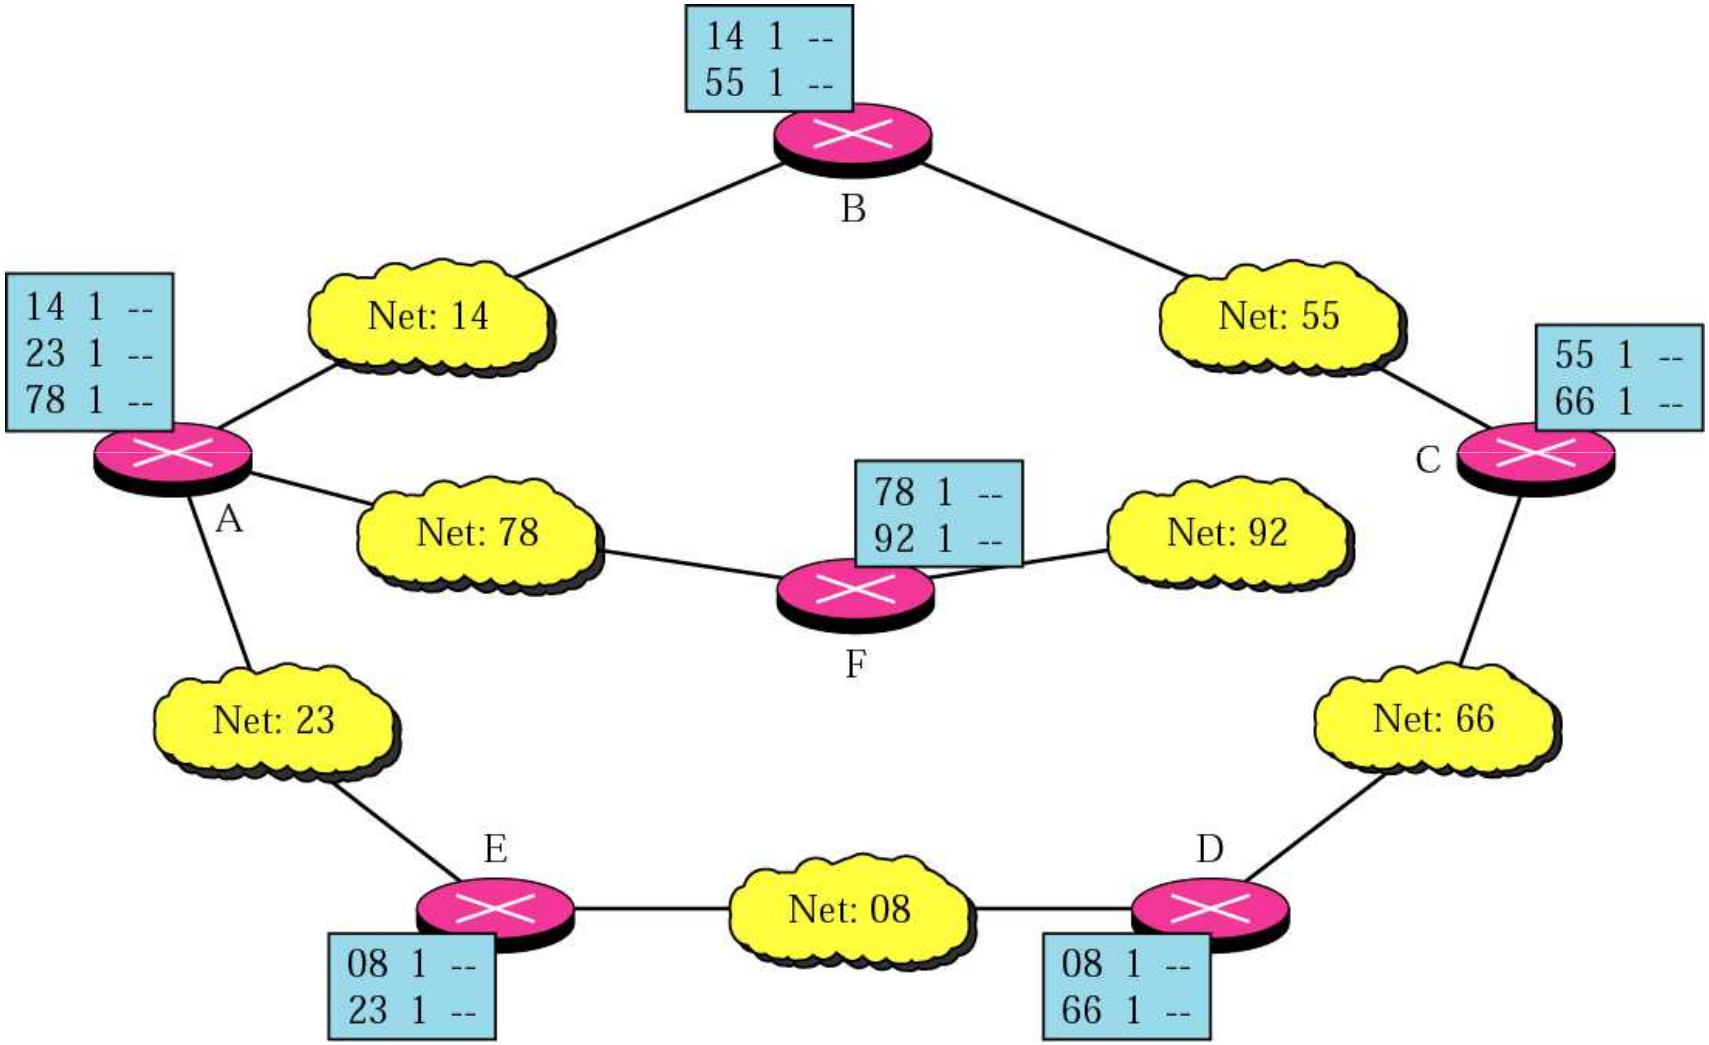

Another Example

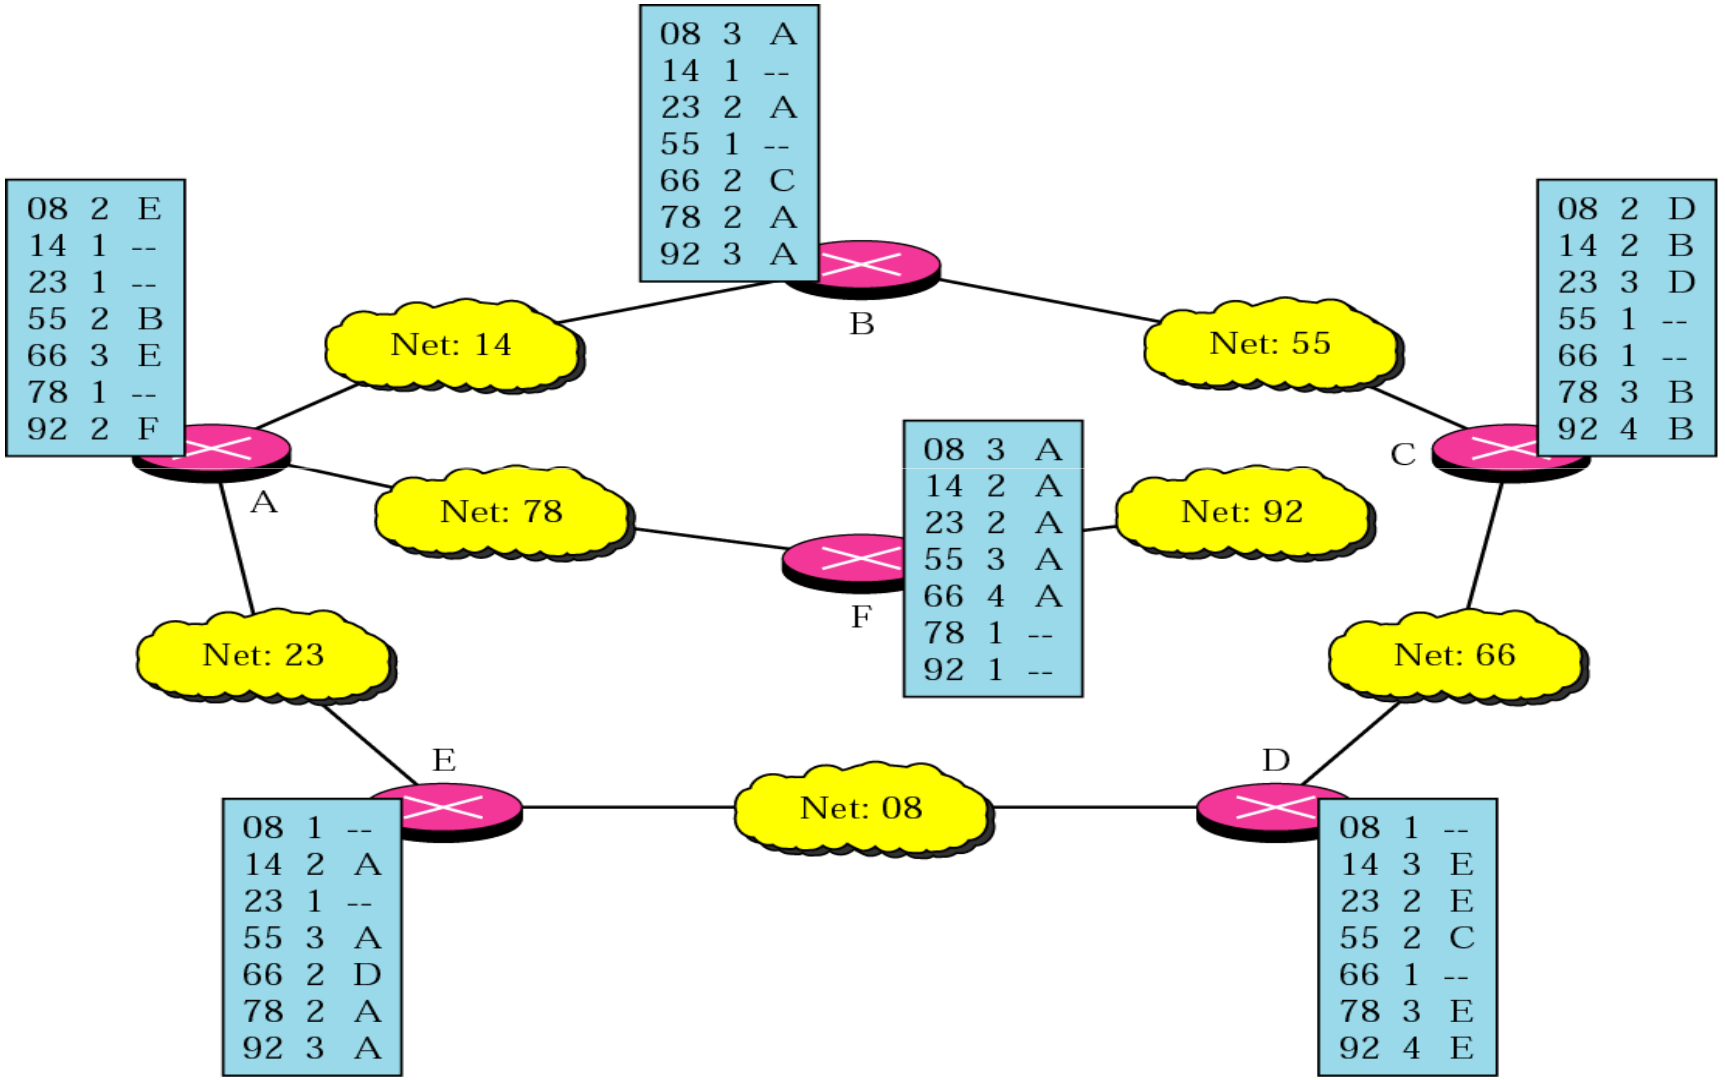

RIP Packet

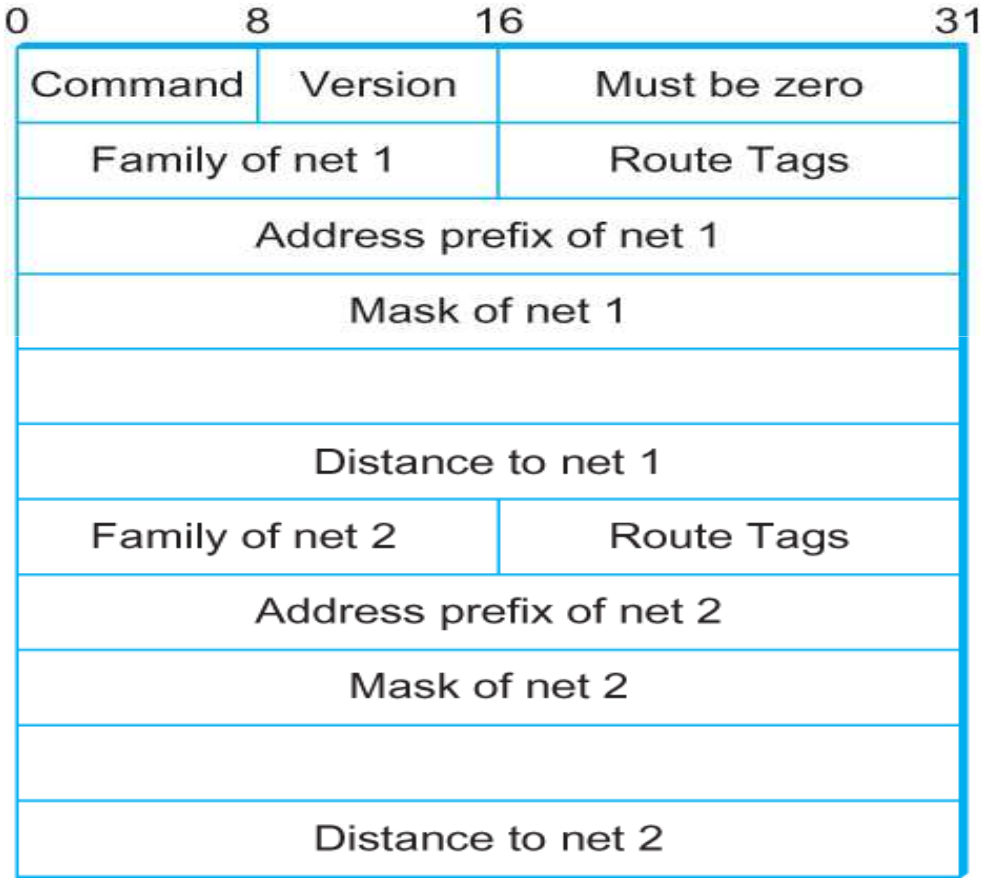

- Essentially contains the networks that can be reached and the distances to each.

Other RIP Details

- Routing tables are exchanged every 30 seconds using the RIP advertisement.
- If a router does not hear from its neighbor once every 180 seconds, the neighbor is deemed unreachable.
- The router that detects this will modify its routing table and propagate the information.
- RIP packets are sent using UDP