ETSF10
Routing part 1

Jens A Andersson

Electrical and Information Technology

Routing

- Introduction
- Inside the Router
- Unicast Routing
  - Intra Domain Routing
  - Inter Domain Routing
- MANET and AdHoc routing
- Multicast Routing

Figure 8.1 ***Circuit-switched network***

Circuit-switched routing

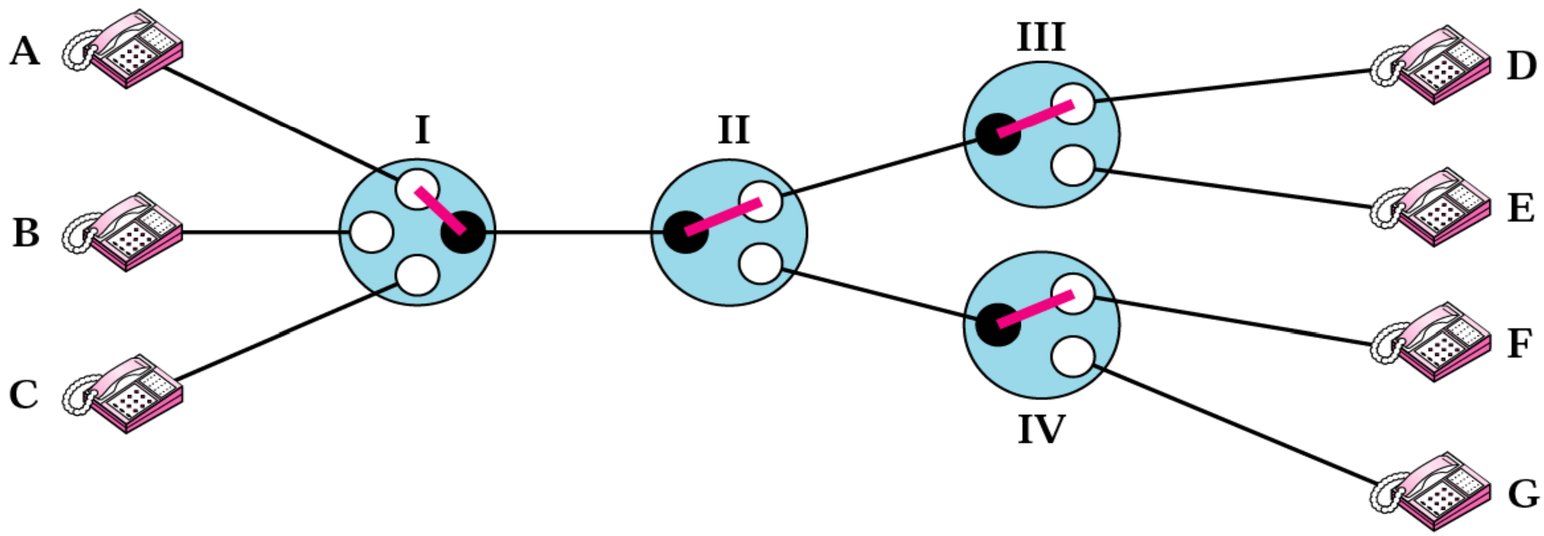

Figure 21.1 **Unicasting**

Packet Switched Routing

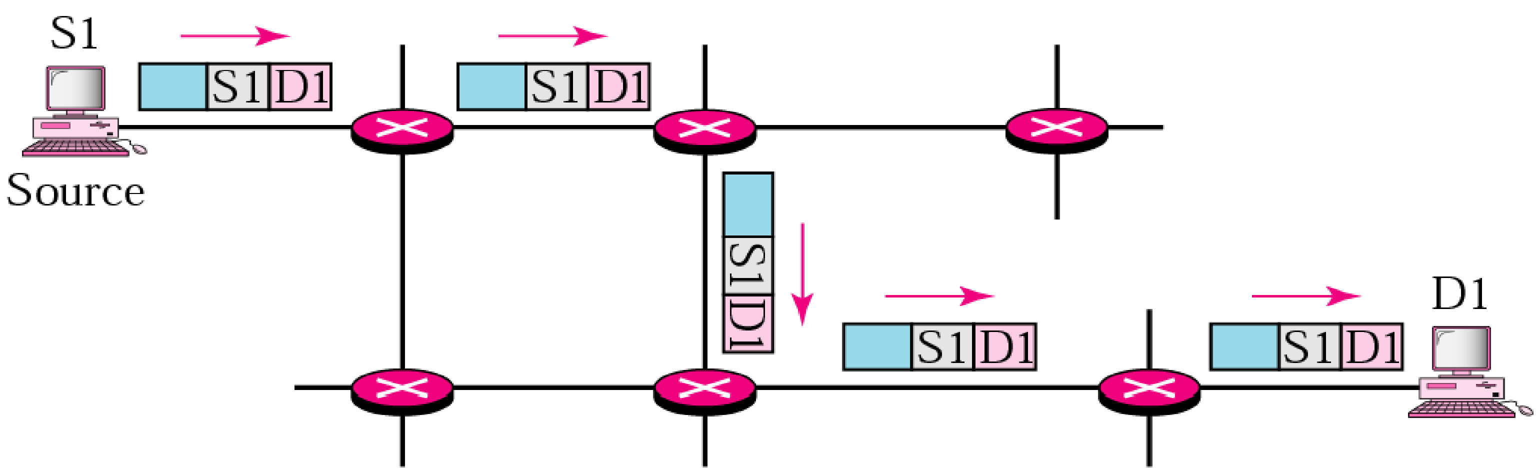

VPN, VPC

- Virtual Private Network
- Virtual Path Connection

- Build a virtual and private network overlaying a public network
- Often by tunnelling (encapsulation of packets inside packets)
- Often encrypted

Choosing an Optimal Path

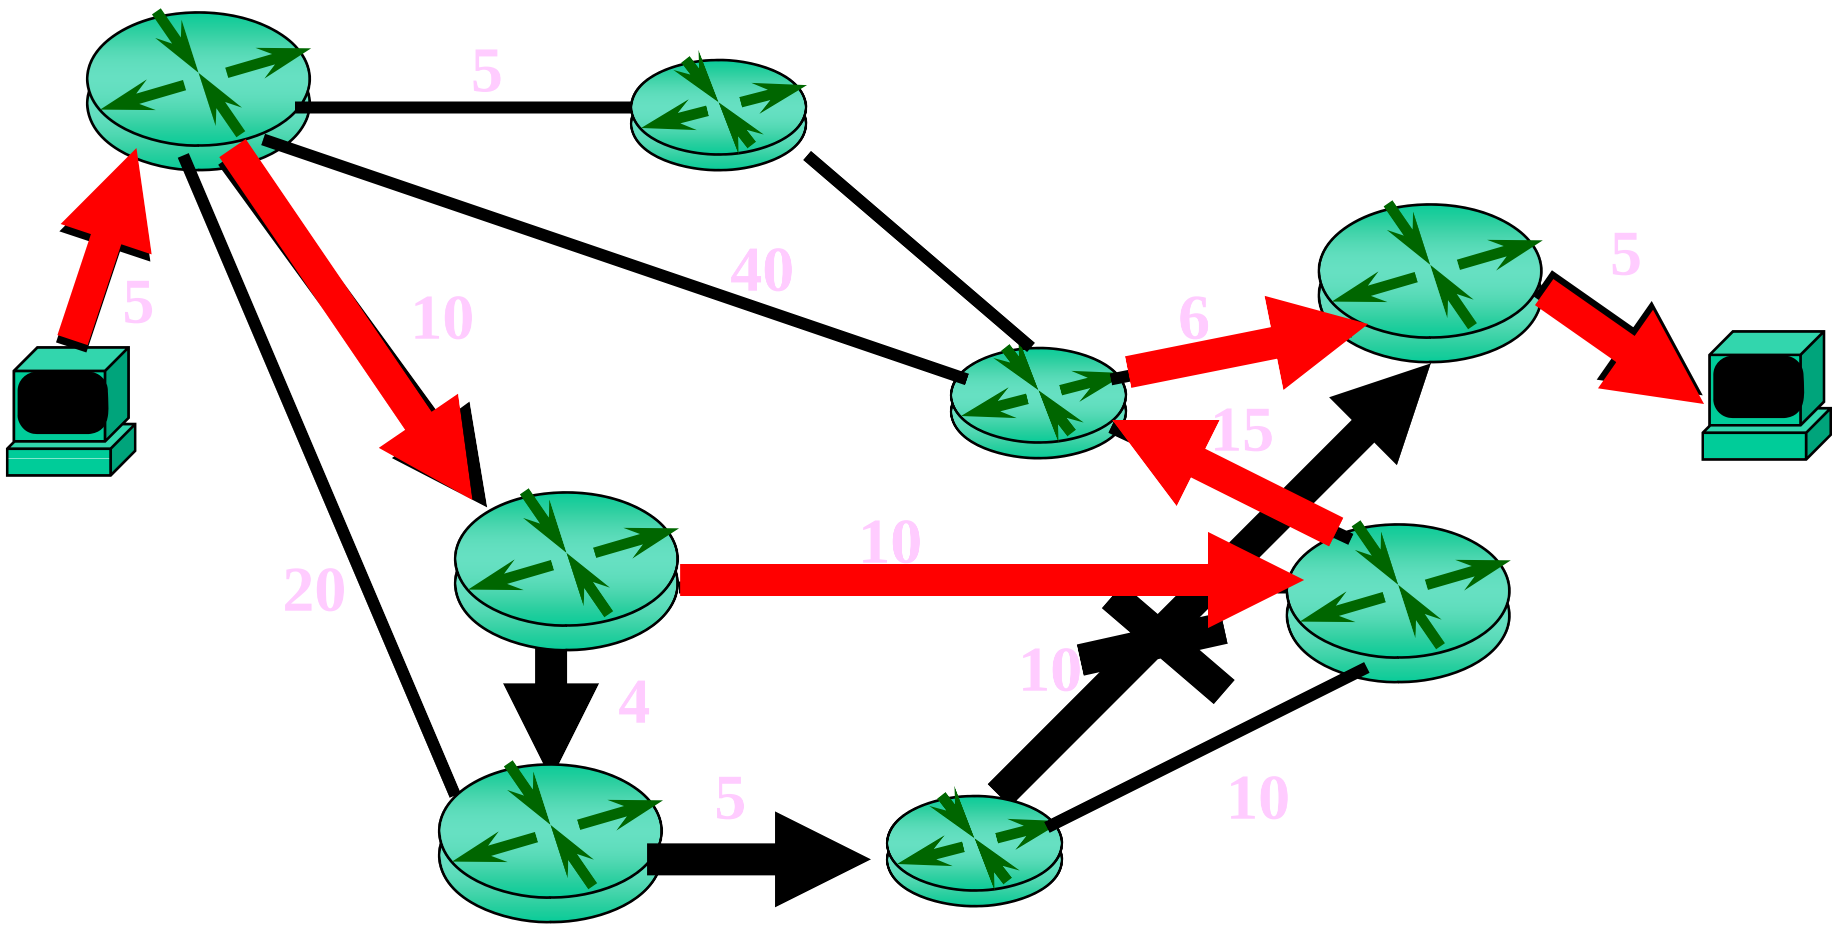

Router

- A router is a type of internetworking device that passes data packets between networks, based on Layer 3 or Network Layer addresses.
- A router has the ability to make intelligent decisions regarding the best path for delivery of data on the network.

Local Routing

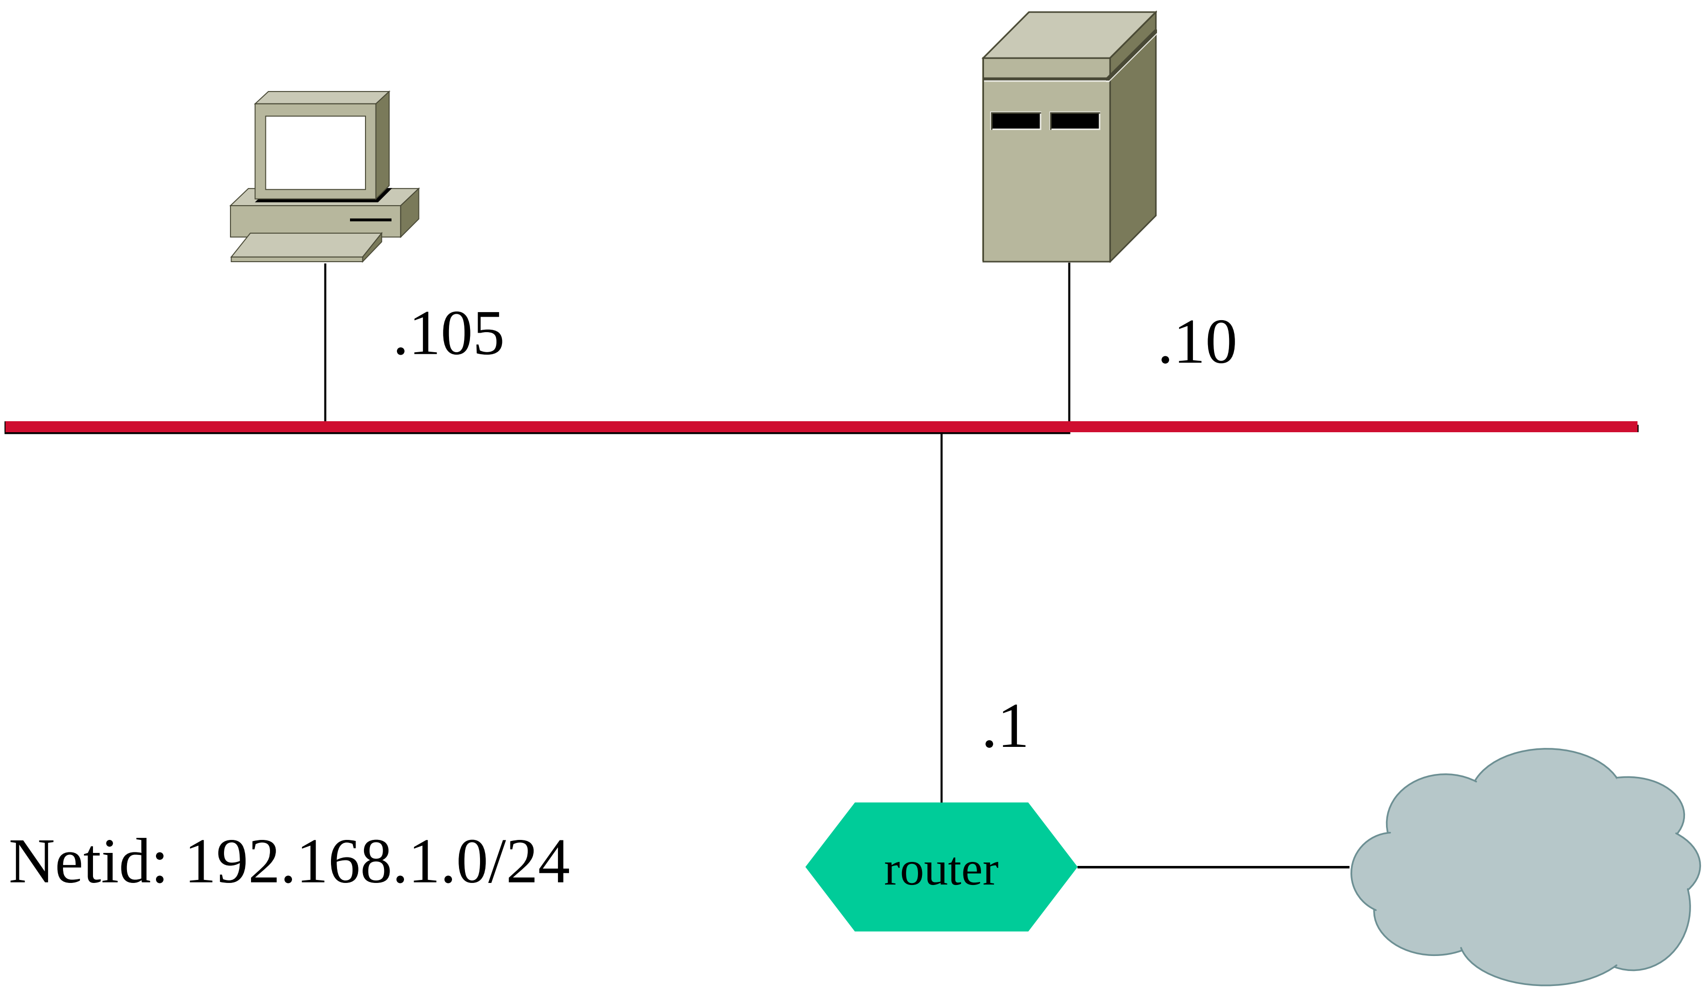

Routing

- Introduction
- Inside the Router
- Unicast Routing
  - Intra Domain Routing
  - Inter Domain Routing
- MANET and AdHoc routing
- Multicast Routing

Router Architecture Overview

Two key router functions:

- run *routing* algorithms/protocol (RIP, OSPF, BGP)
- *switching* datagrams from incoming to outgoing link

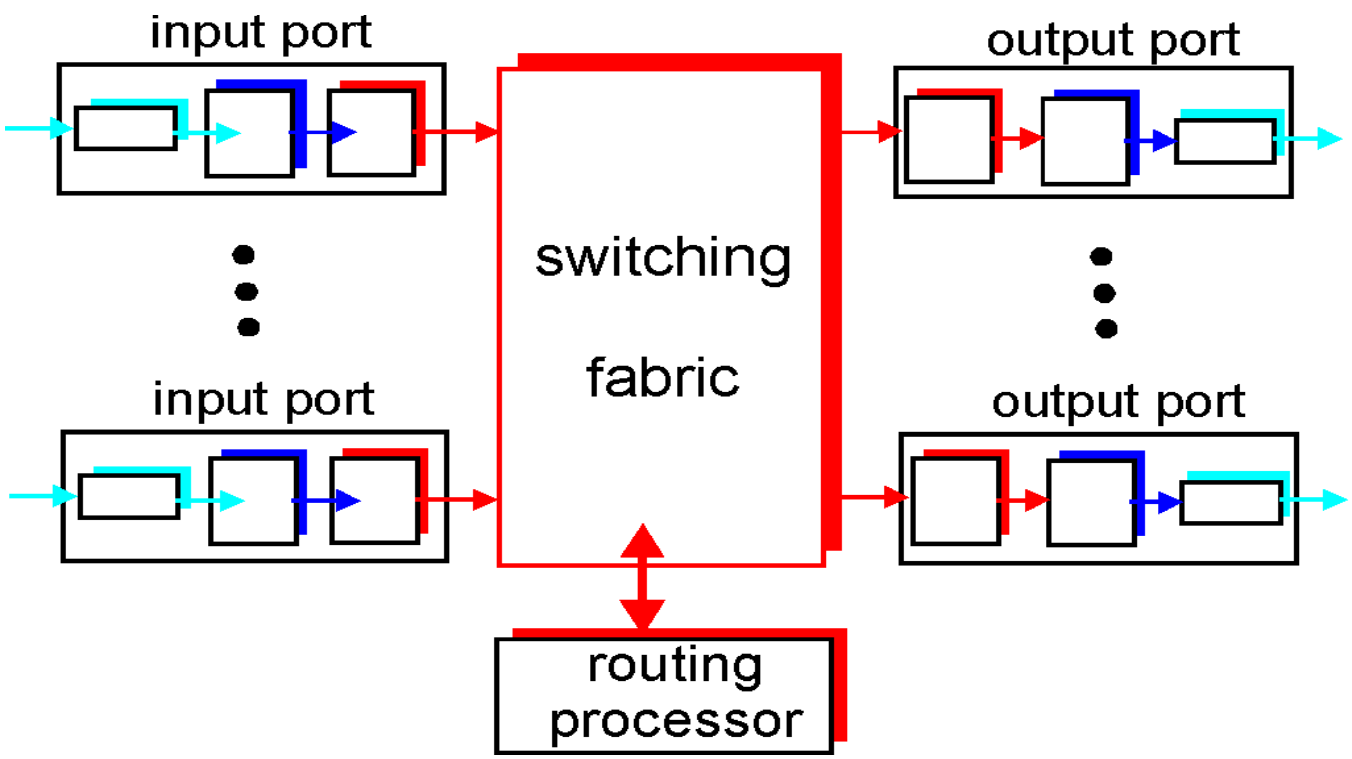

Input Port Functions

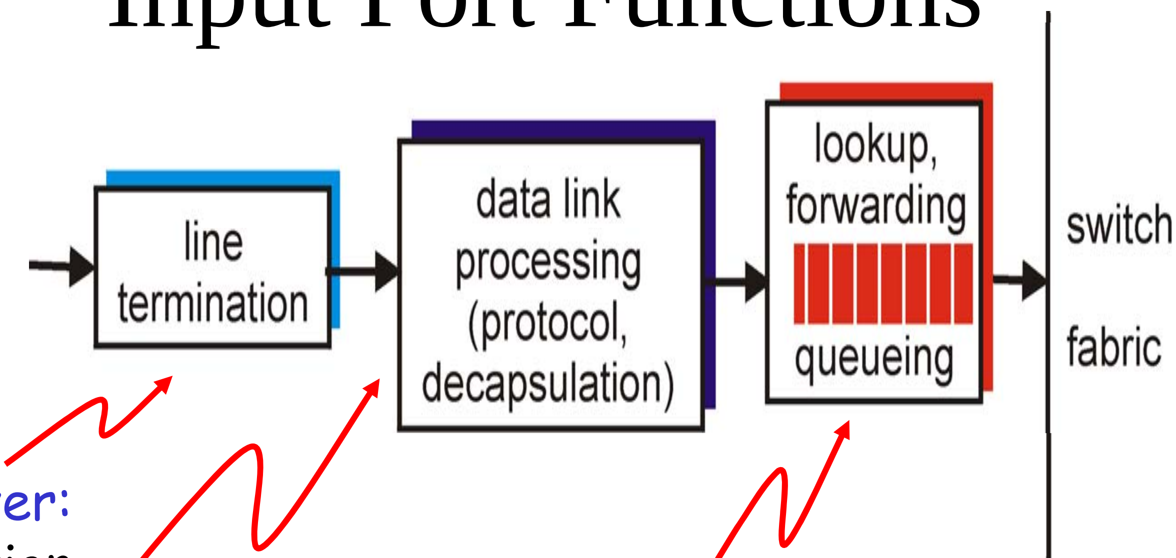

Physical layer: bit-level reception

Data link layer: e.g., Ethernet see chapter 5

Decentralized switching:

- given datagram dest., lookup output port using routing table in input port memory
- goal: complete input port processing at ‘line speed’
- queuing: if datagrams arrive faster than forwarding rate into switch fabric

Output Ports

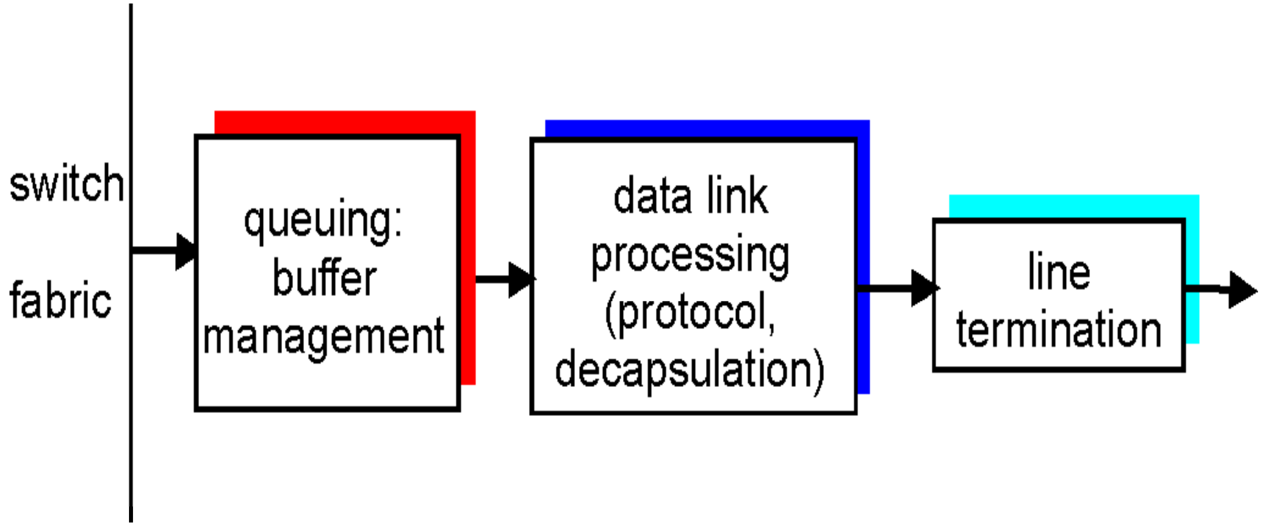

- *Buffering* required when datagrams arrive from fabric faster than the transmission rate
- *Scheduling discipline* chooses among queued datagrams for transmission (priority)

Three types of switching fabrics

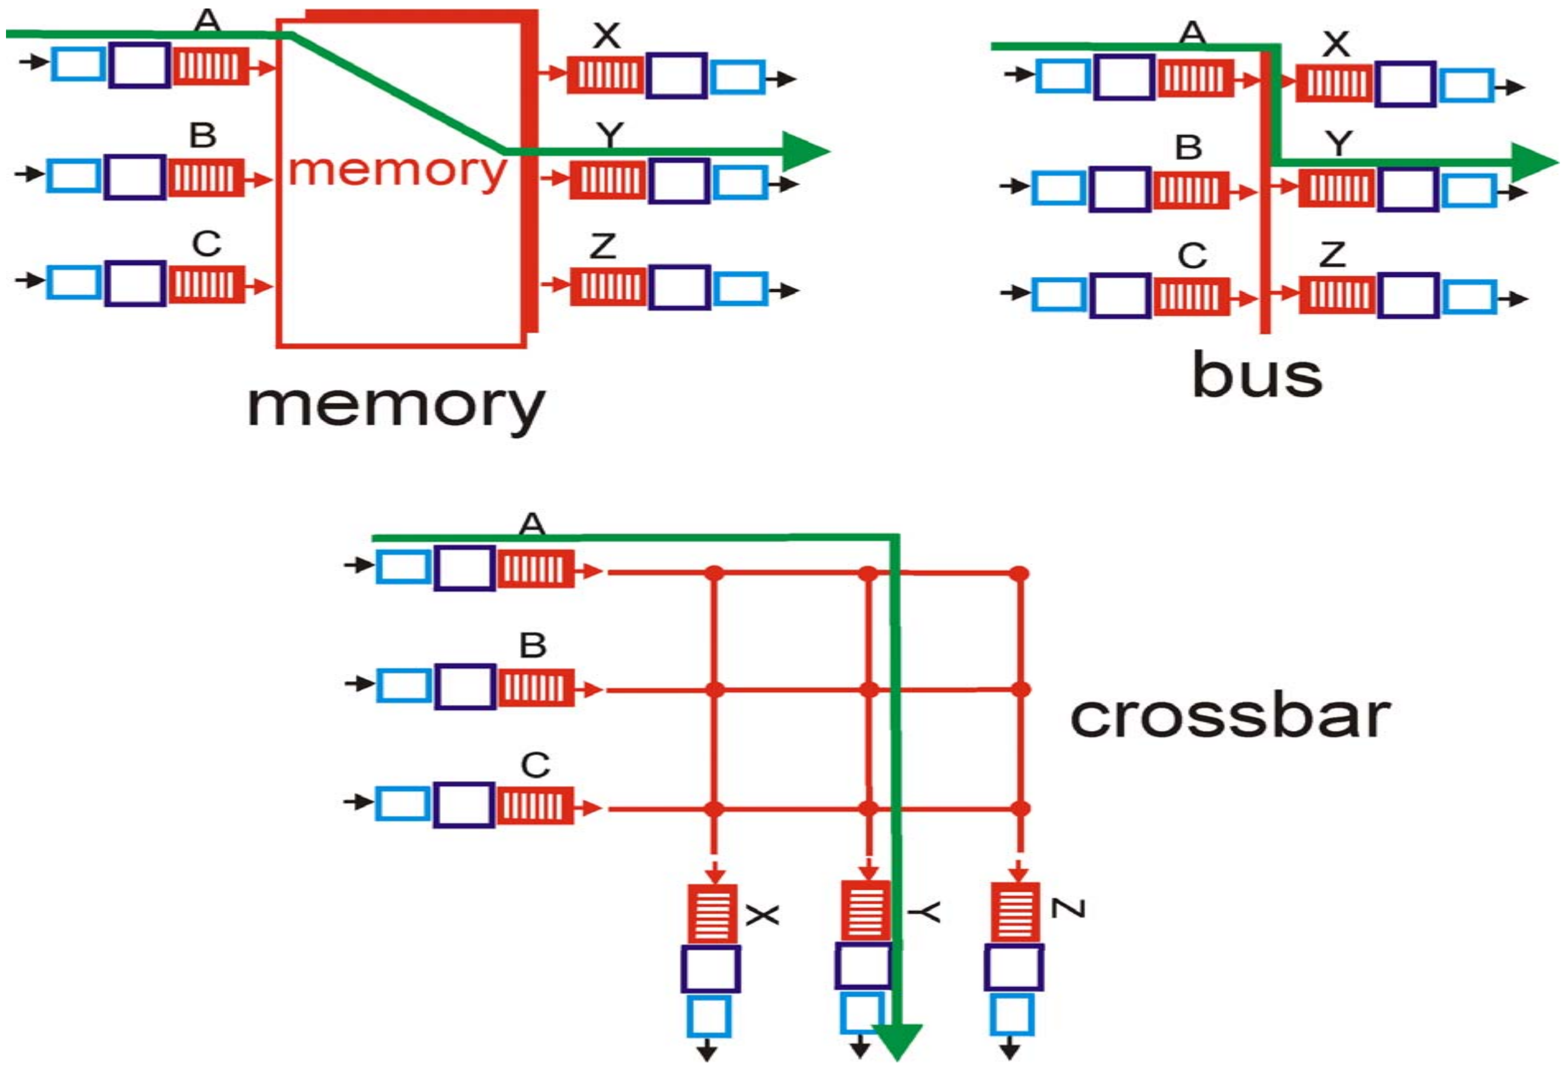

Switching Via Memory

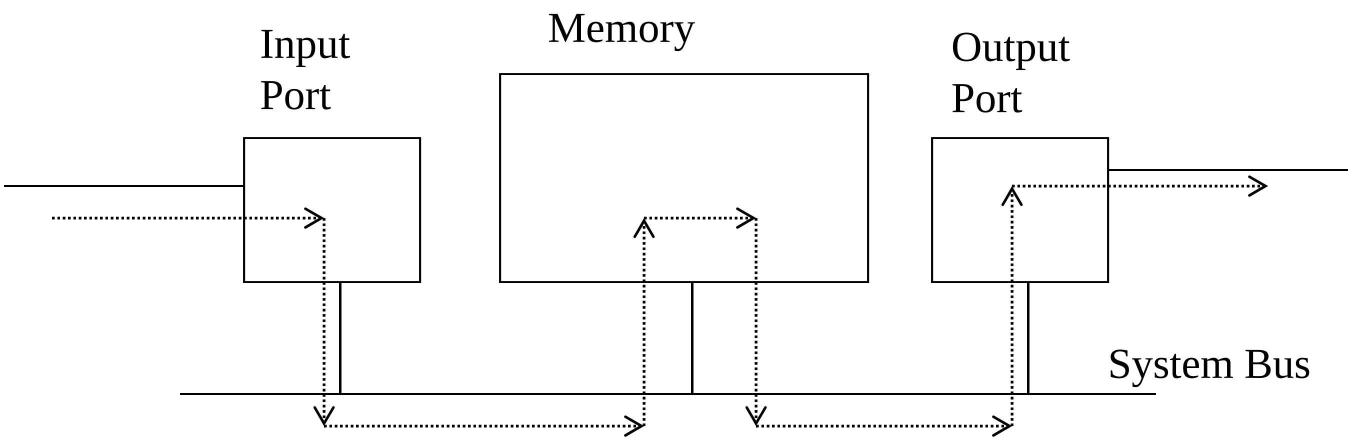

First generation routers:

- packet copied by system's (single) CPU
- speed limited by memory bandwidth (2 bus crossings per datagram

Modern routers:

- input port processor performs lookup, copy into memory

Switching Via a Bus

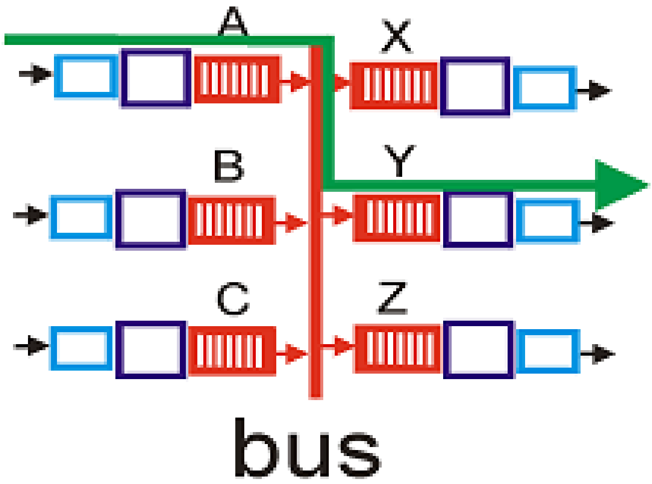

- datagram from input port memory
  to output port memory via a shared bus
- bus contention: switching speed limited by bus bandwidth

Switching Via An Interconnection Network

- Crossbar, fabric
- overcome bus bandwidth limitations
- Banyan networks, other interconnection nets initially developed to connect processors in multiprocessor
- Advanced design: fragmenting datagram into fixed length cells, switch cells through the fabric.
- Modern routers are switches running at wire speed

Input Port Queuing

- Fabric slower that sum of input ports -> queueing may occur at input queues
- Head-of-the-Line (HOL) blocking: queued datagram at front of queue prevents other datagrams in queue from moving forward
- *queueing delay and loss due to input buffer overflow!*
- Random Early Discard (RED) tampers with TCP to reduce load

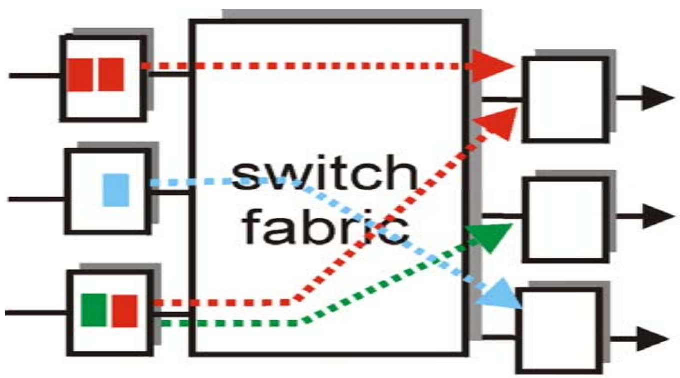

output port contention at time t - only one red packet can be transferred

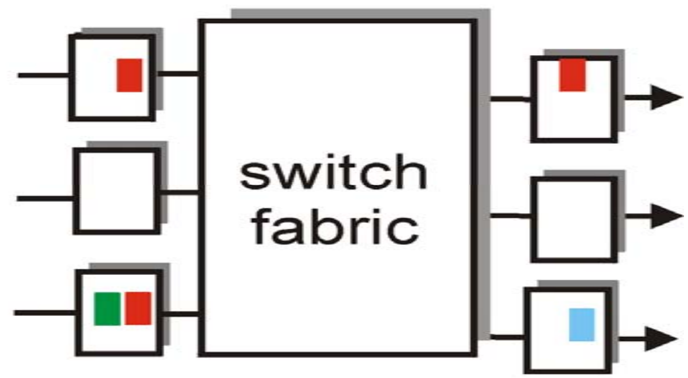

green packet experiences HOL blocking

Output port queueing

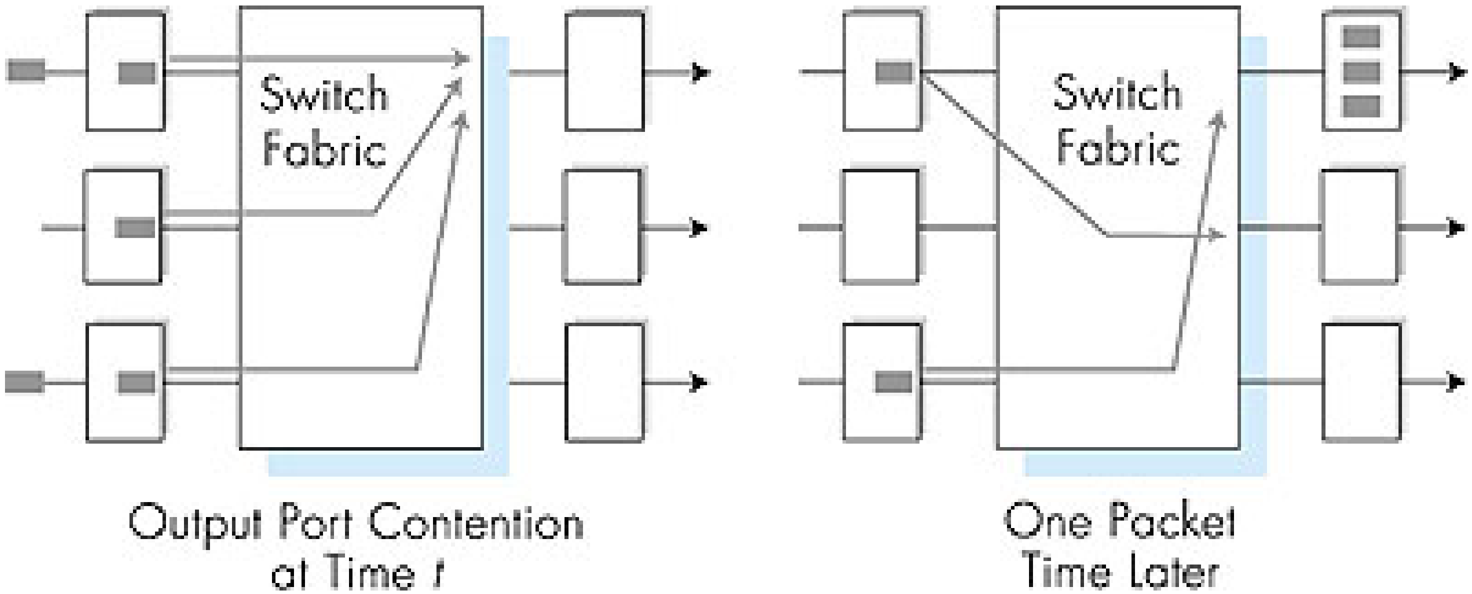

- buffering when arrival rate via switch exceeds output line speed
- *queueing (delay) and loss due to output port buffer overflow!*

Routing

- Introduction
- Inside the Router
- Unicast Routing
  - Intra Domain Routing
  - Inter Domain Routing
- MANET and AdHoc routing
- Multicast Routing

Unicast Routing

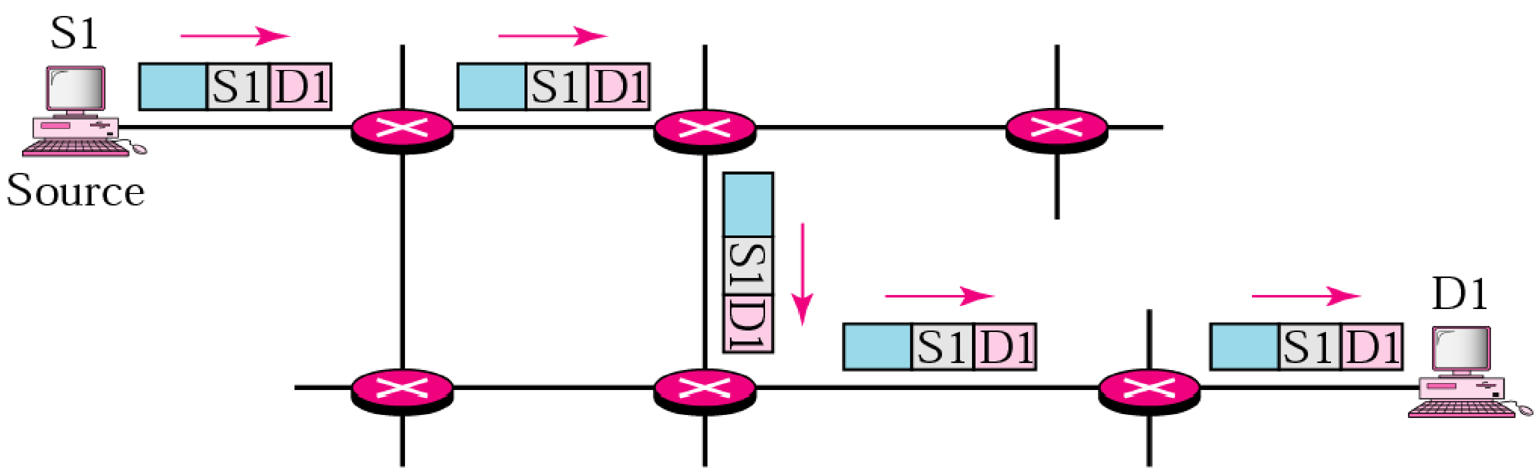

Reminder

- Routing performed on net ids
- One net id per link! (at least)
- (What if links without net ids?)

Flooding

- Routing without routing information
- Forward each packet to all ports except incoming port
- Detect and remove looping packets
- Very robust
- This is not broadcast

Technology of Distributed Routing

Link State

- Topology information is flooded within the routing domain
- Best end-to-end paths are computed locally at each router.
- Best end-to-end paths determine next-hops.
- Based on minimizing some notion of distance
- Works only if policy is shared and uniform
- Examples: OSPF, IS-IS

Distance Vector

- Each router knows little about network topology
- Only best next-hops are chosen by each router for each destination network.
- Best end-to-end paths result from composition of all next-hop choices
- Does not require any notion of distance
- Does not require uniform policies at all routers
- Examples: RIP, BGP

Routing Tables and Forwarding Table

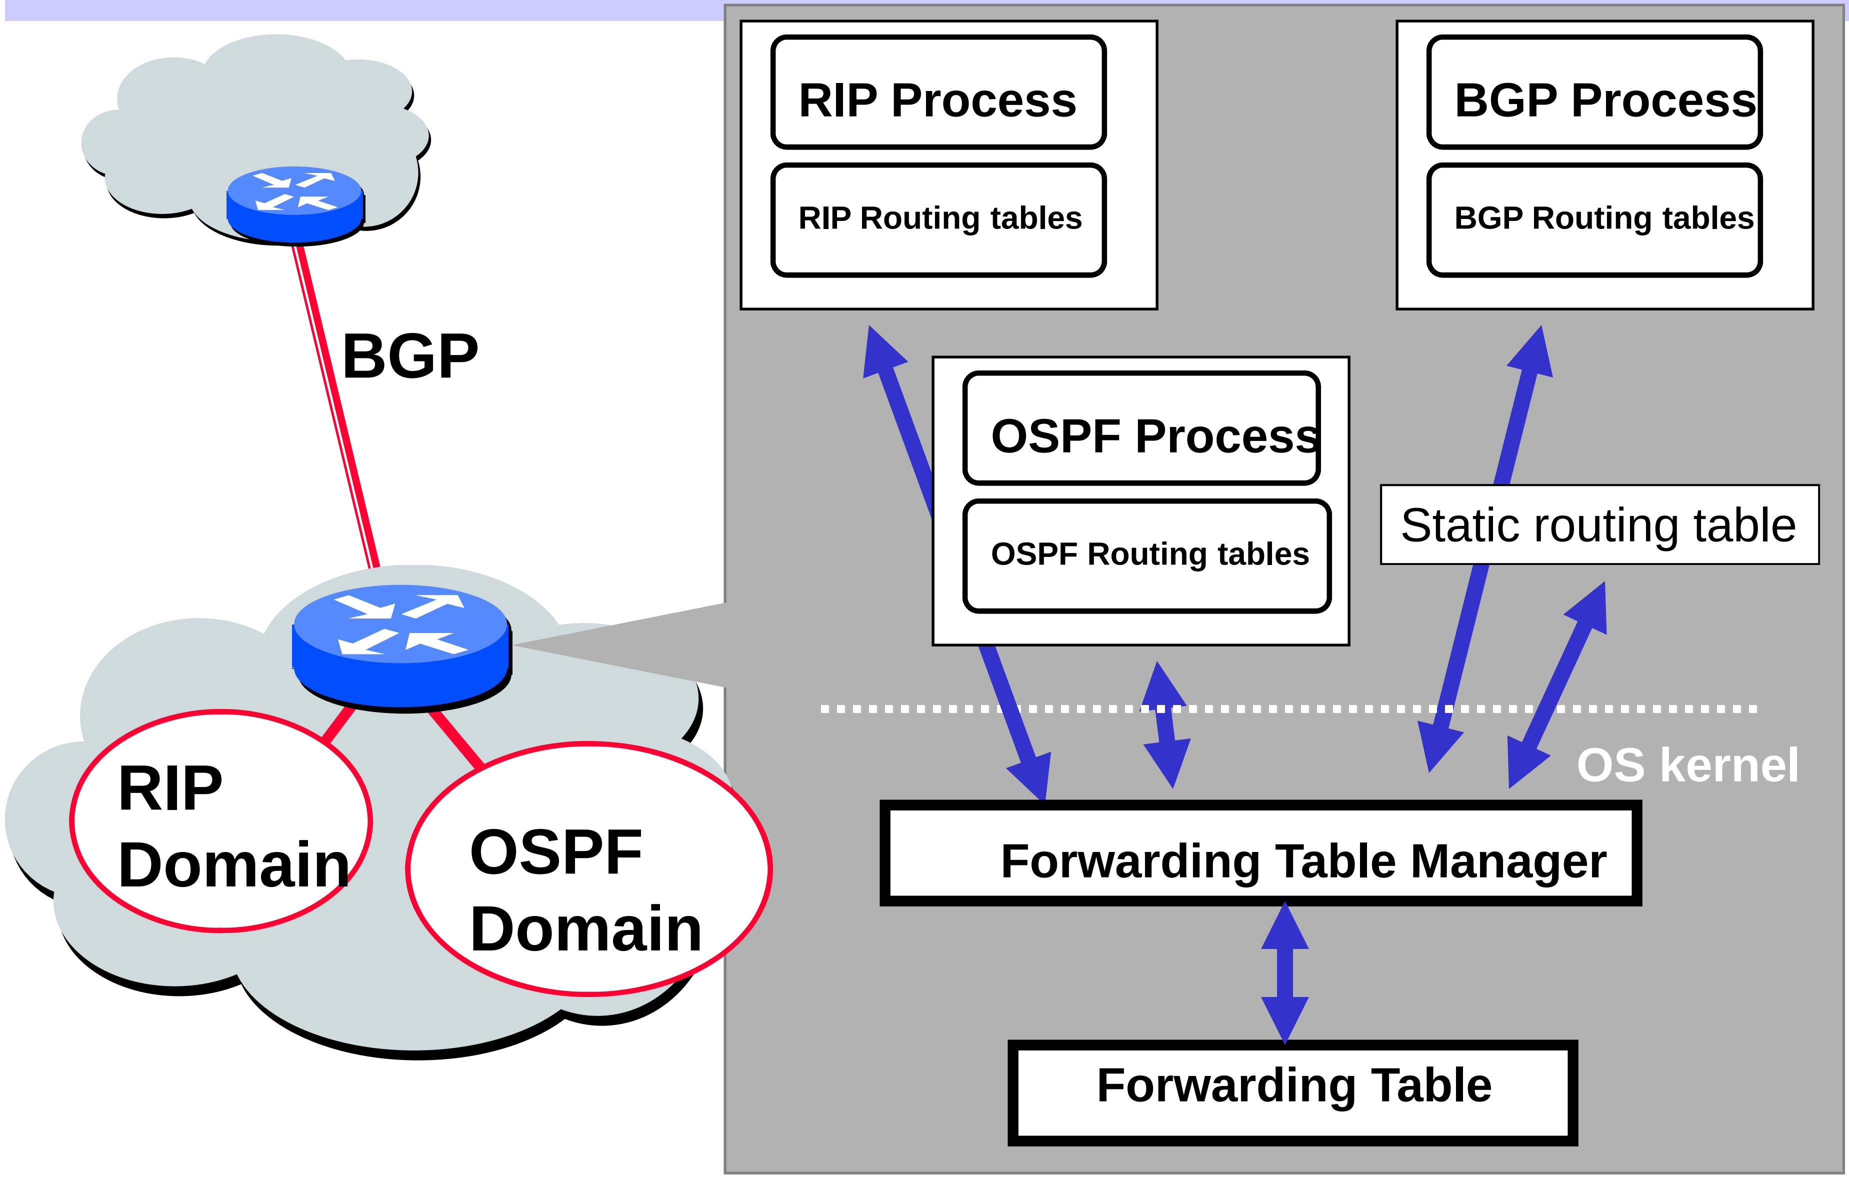

The Link Metric

- Possible metrics
  - hop count
  - inverse of the link bandwidth
  - delay
  - dynamically calculated
  - administratively assigned
  - combination
- Traffic should be monitored and metrics adjusted

Routing Protocol

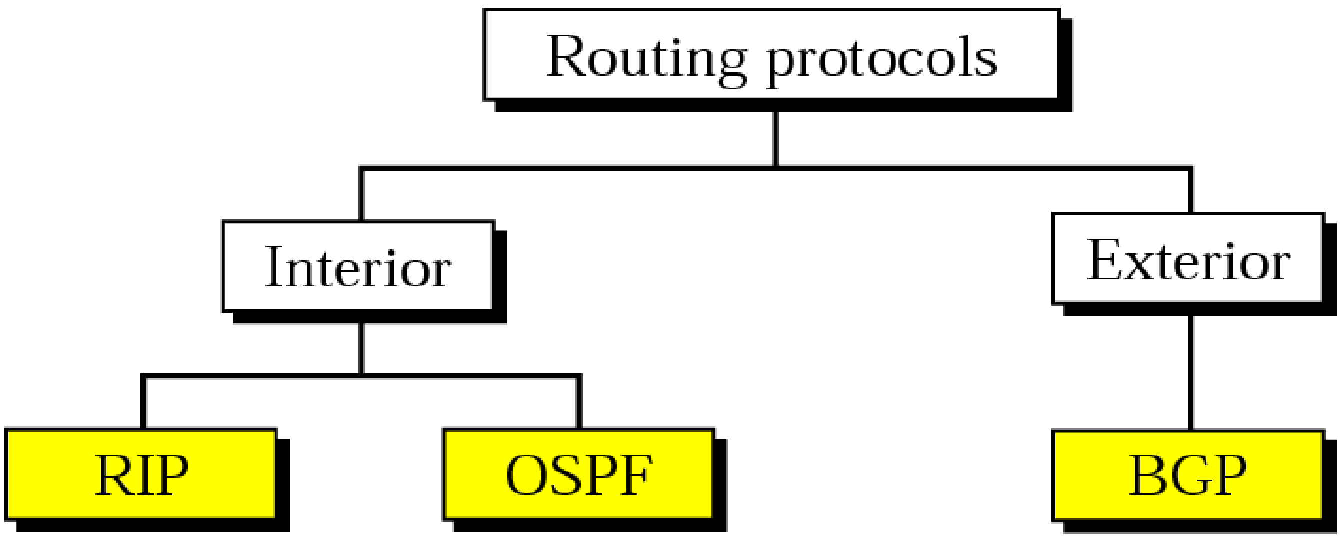

Figure 21-18

The Concept of Distance Vector Routing

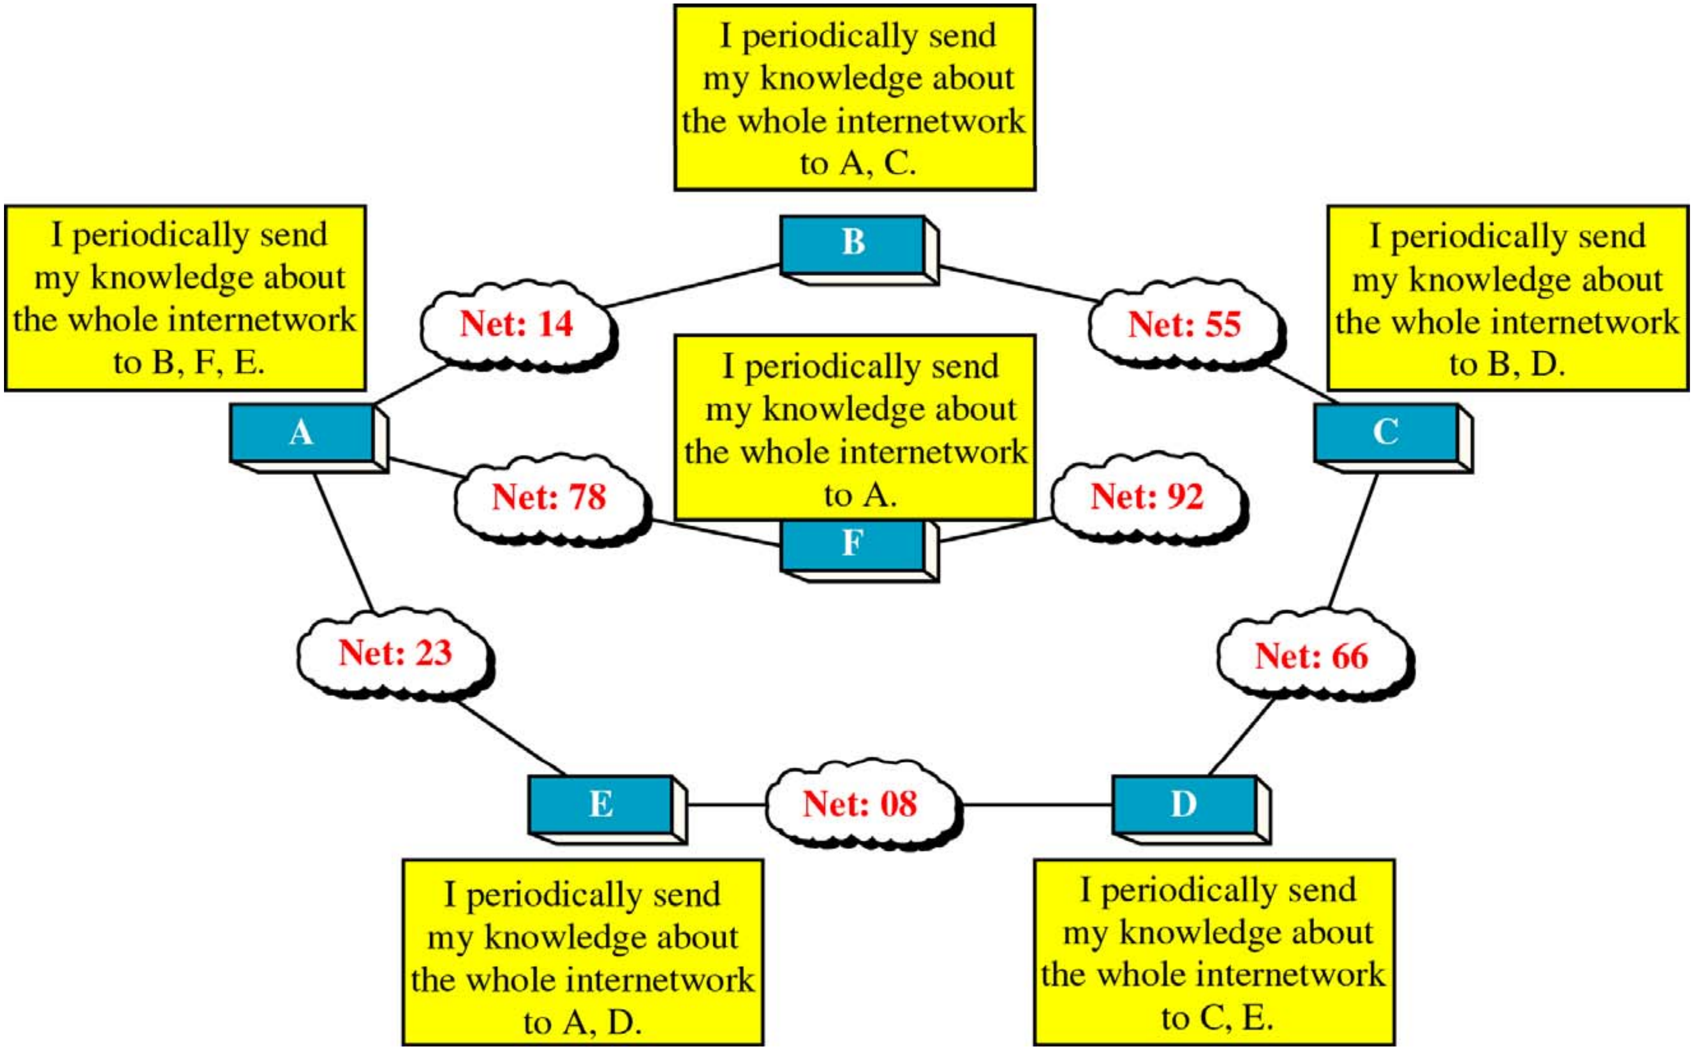

Table 21.1** **A distance vector routing table

Distance Vector Routing Table

Destination	Hop Count	Next Router	Other information
163.5.0.0	7	172.6.23.4	
197.5.13.0	5	176.3.6.17	
189.45.0.0	4	200.5.1.6	
115[illegible].0.0	6	131.4.7.19	

Network id

Metric

Host id of interface

RIP (Routing Information Protocol)

- Distance vector algorithm
- Included in BSD-UNIX Distribution in 1982
- Distance metric: # of hops (max = 15 hops)
  - *Can you guess why?*

- Distance vectors: exchanged among neighbours every 30 sec via Response Message (also called **advertisement**)
- Each advertisement: list of up to 25 destination nets within AS

RIP Updating Algorithm

(1) if (advertised destination not in table) then

    update table

(2) else

(2.a) if (advertised next-hop = next-hop in table) then

    replace entry

(2.b) else

(2.b.i) if (advertised hop count < hop count in table) then

    replace entry

(2.b.ii) else

    do nothing

Figure 21-23

Example 21.1

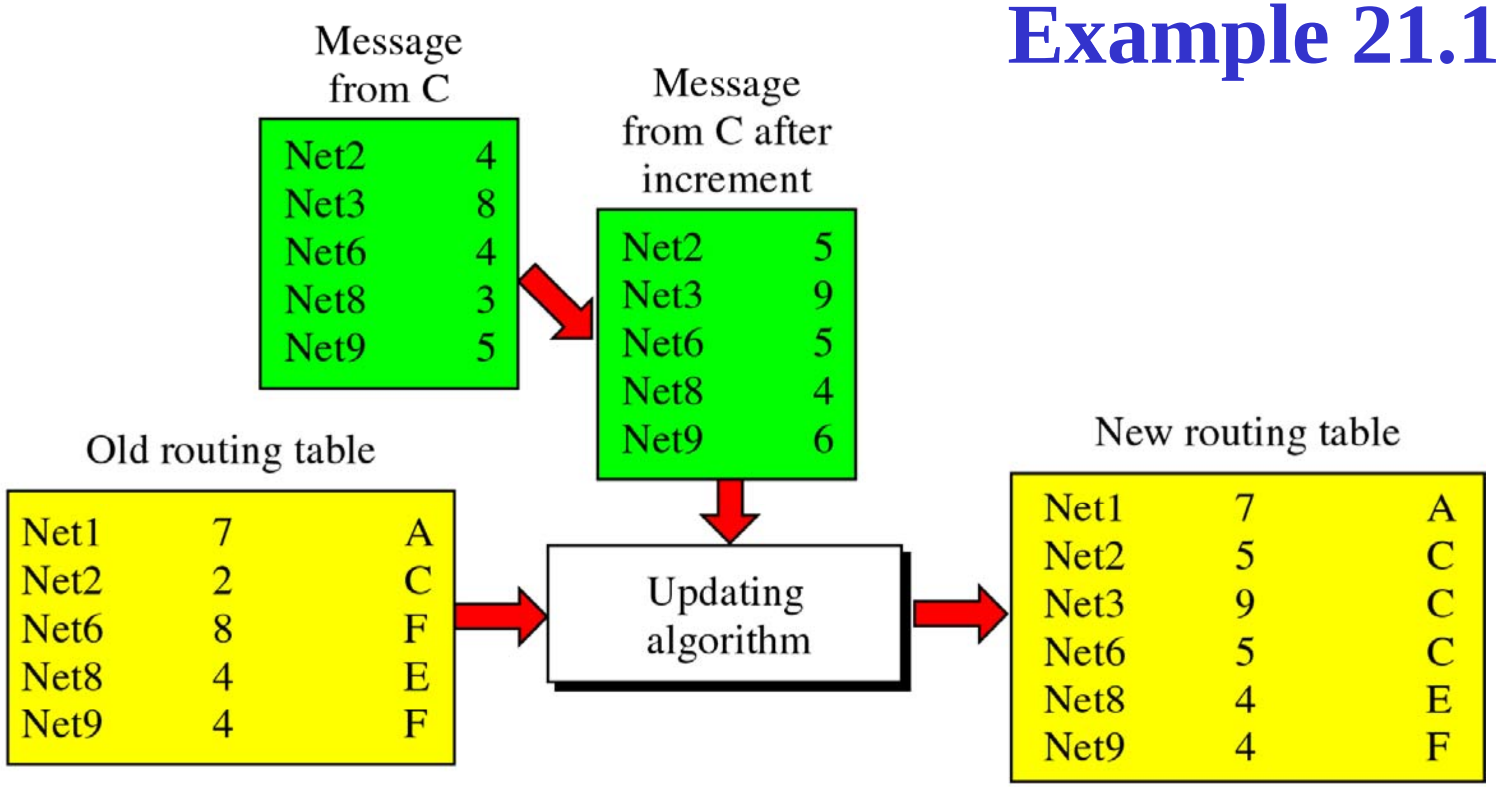

Rules

Net2: Replace **(Rule 2.a)**
Net3: Add **(Rule 1)**
Net6: Replace **(Rule 2.b.i)**
Net8: No change **(Rule 2.b.ii)**
Net9: No change **(Rule 2.b.ii)**

Note that there is no news about Net1 in the advertised message, so none of the rules apply to this entry.

RIP: Link Failure and Recovery

If no advertisement heard after 180 sec --> neighbor/link declared dead

- routes via neighbour invalidated
- new advertisements sent to neighbours
- neighbours in turn send out new advertisements (if tables changed)
- link failure info quickly propagates to entire net
- poison reverse used to prevent ping-pong loops (infinite distance = 16 hops)

The Count to Infinity Problem

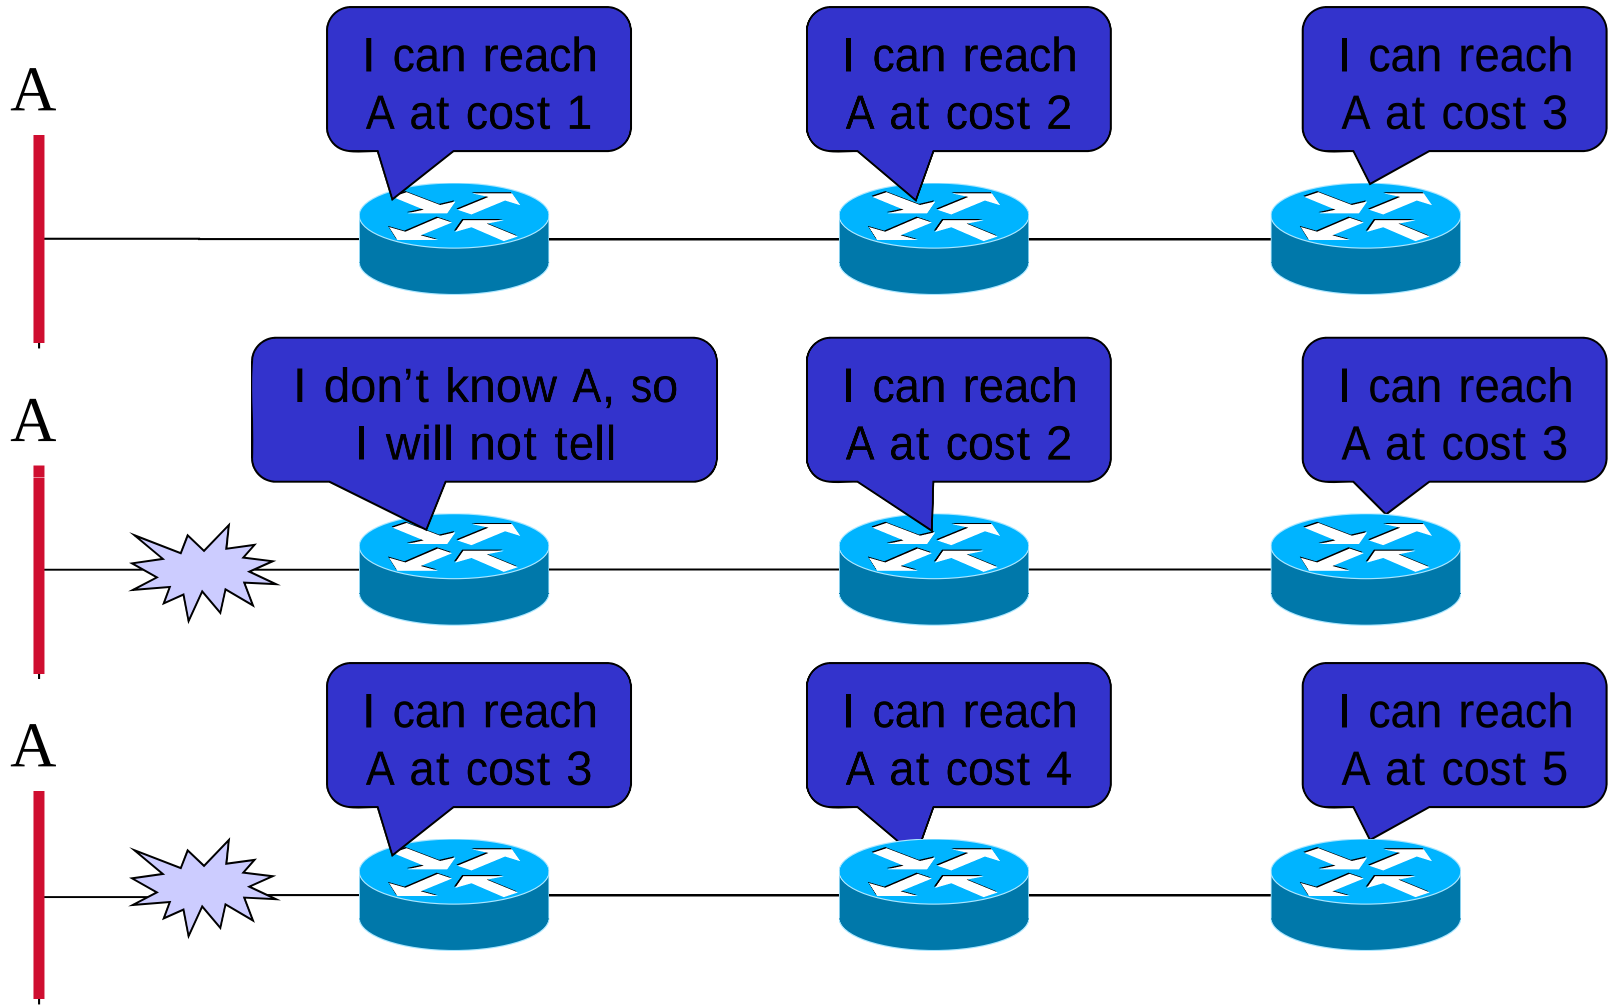

The Count to Infinity Problem (cont)

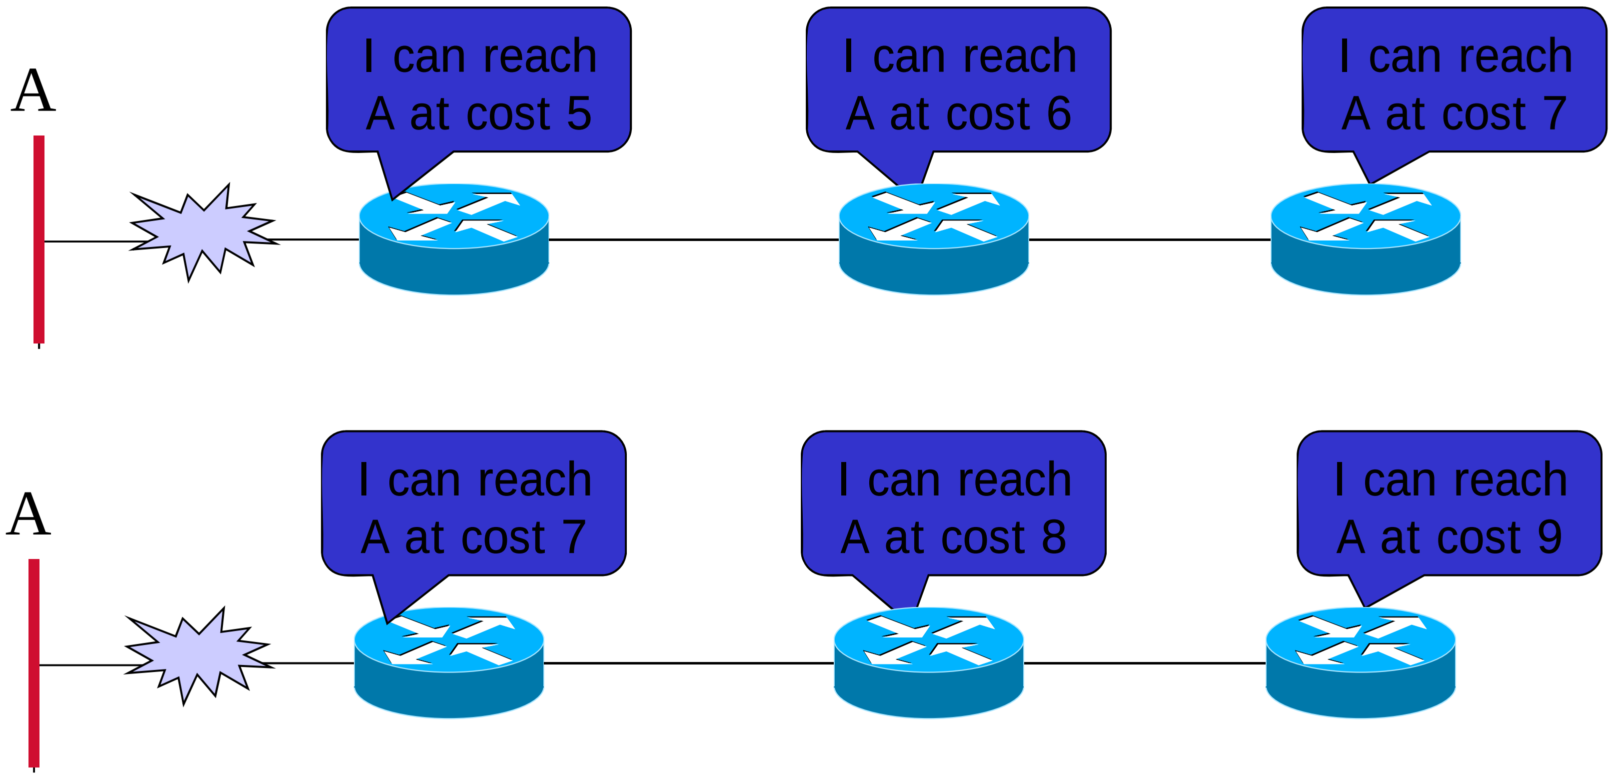

Split Horizon

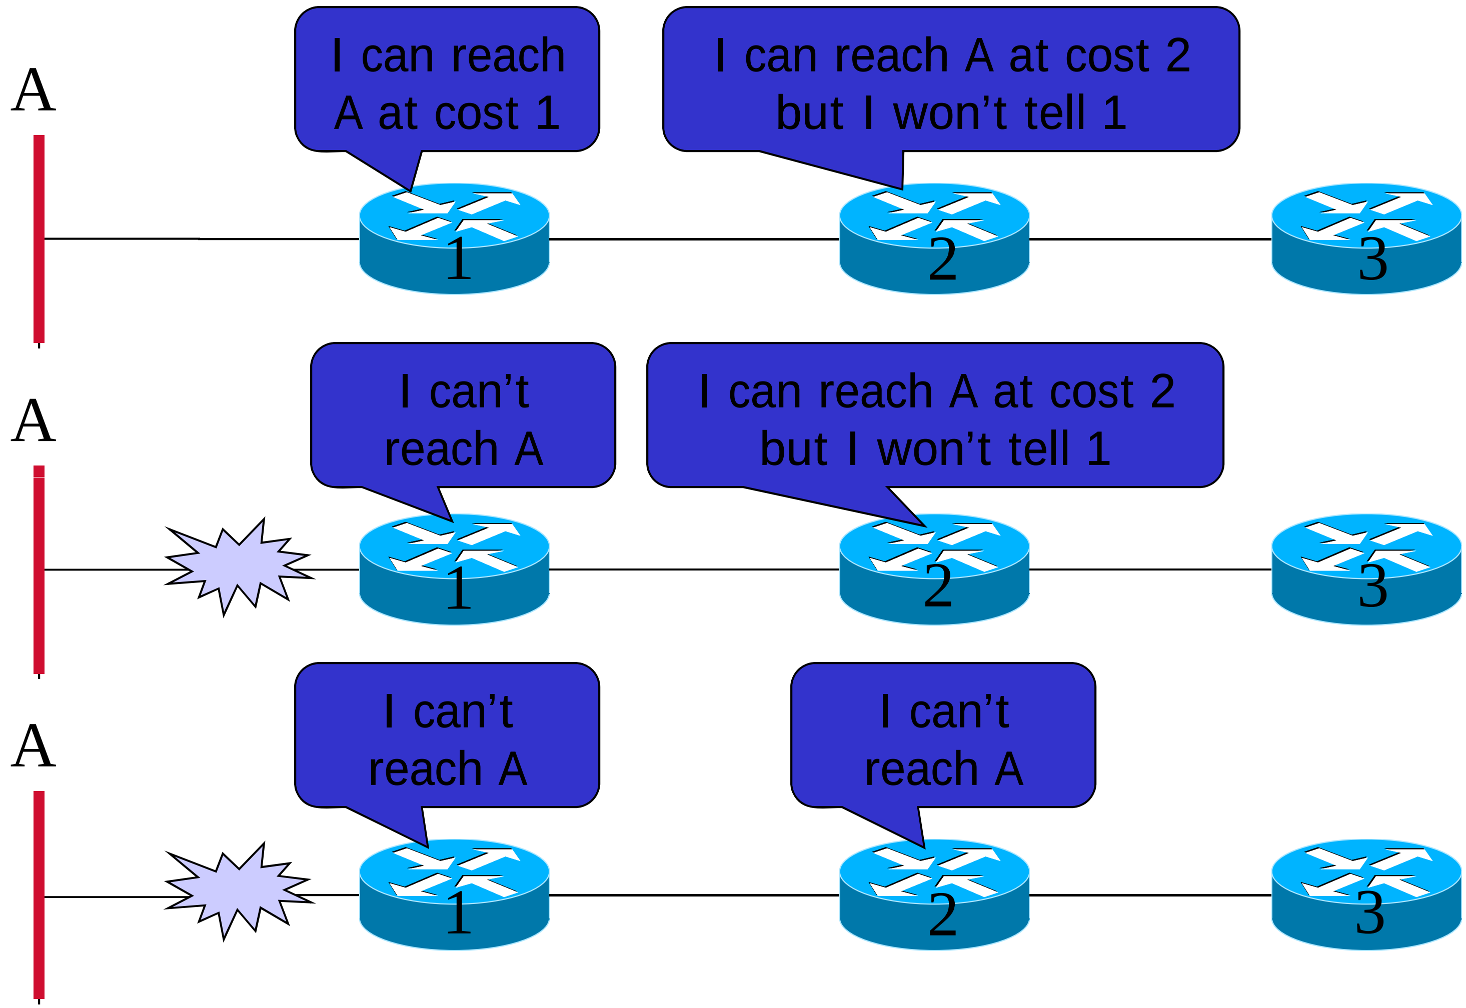

Split Horizon with Poison Reverse

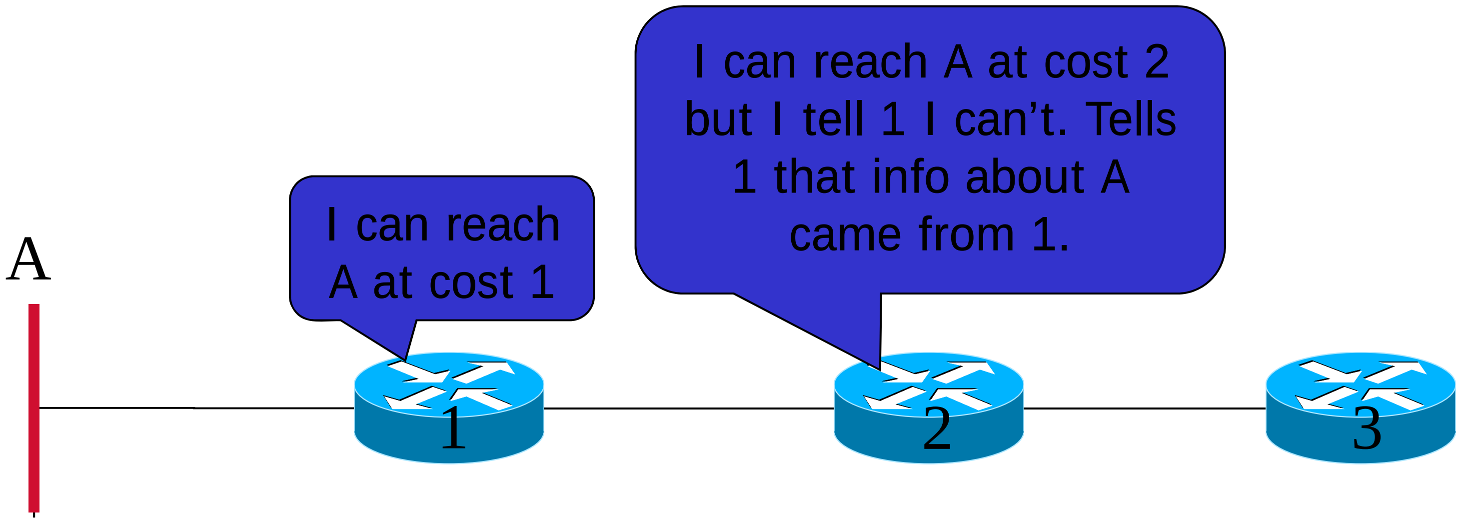

Figure 21-24

Concept of Link State Routing

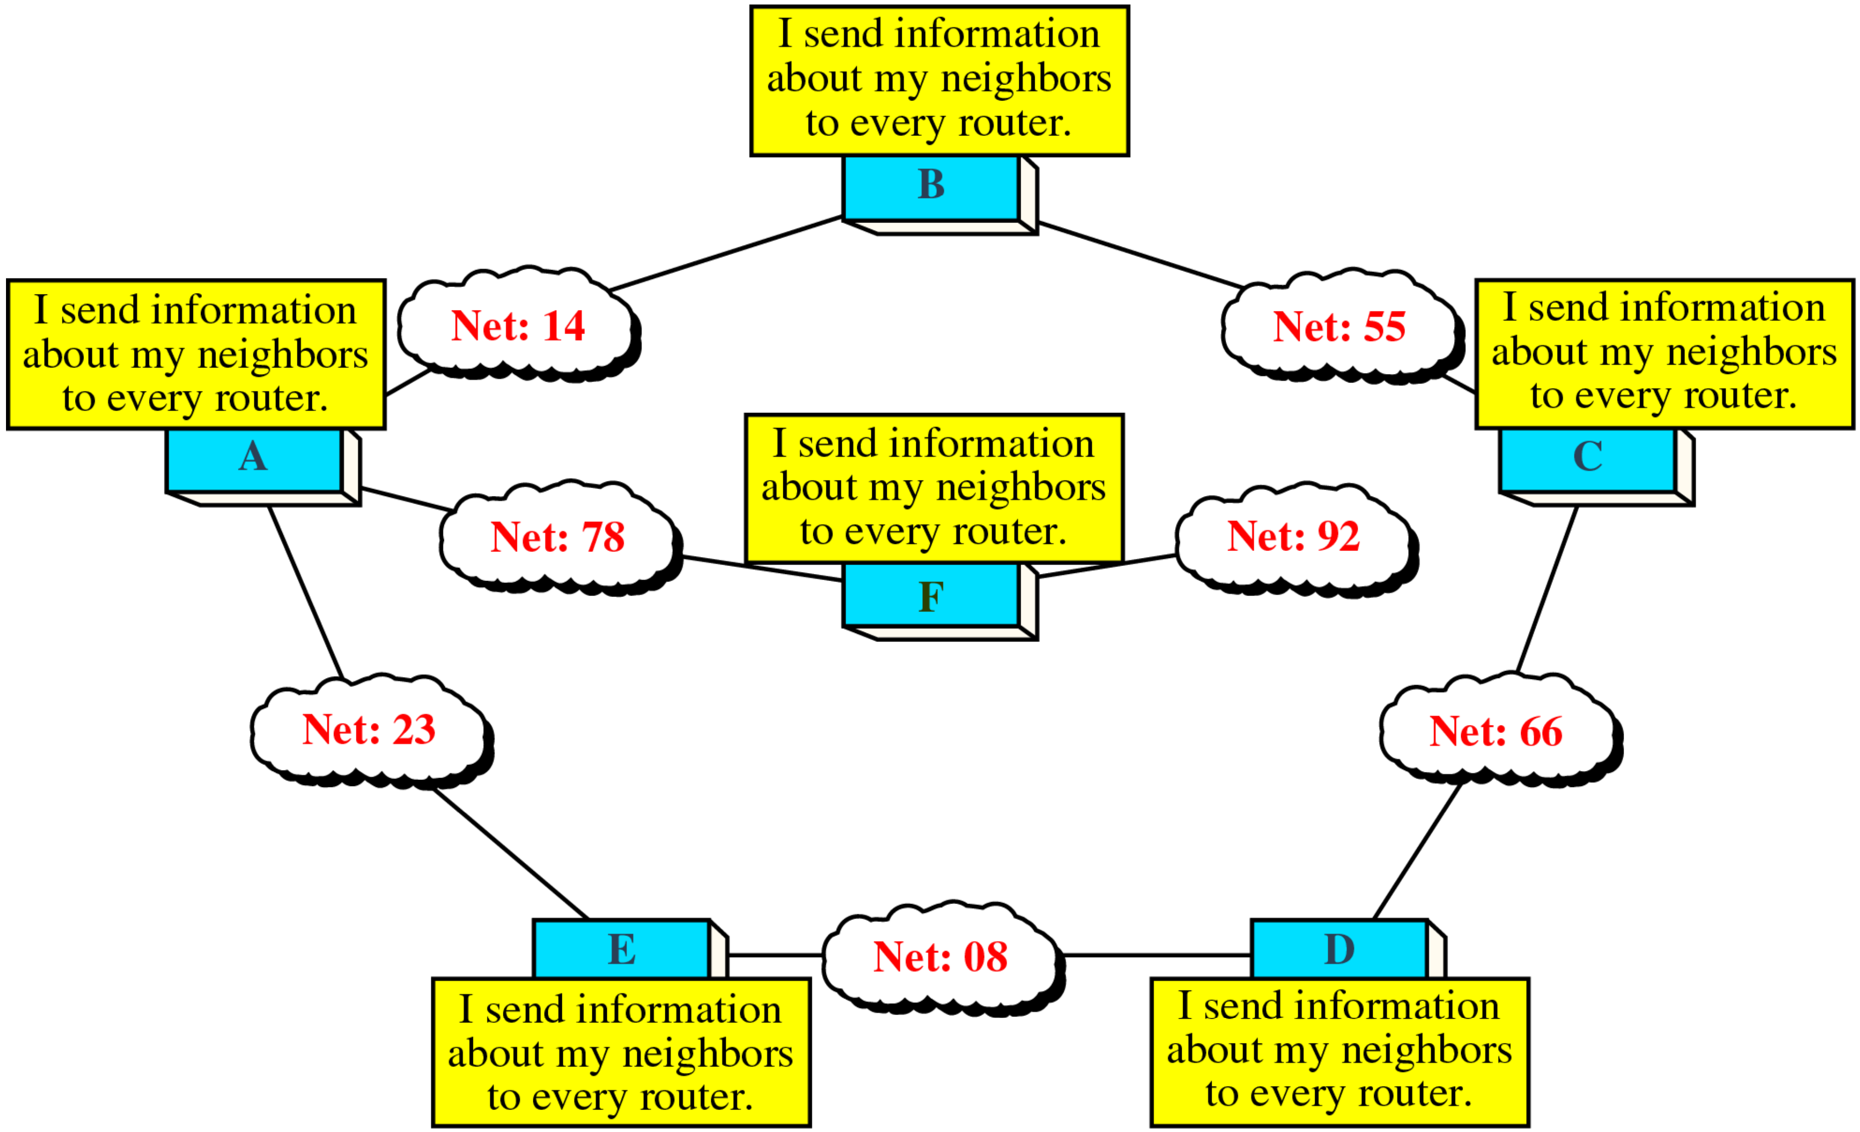

Figure 21-28

Link State Database

Advertiser	Network	Cost	Neighbor
A	14	1	B
A	78	3	F
A	23	2	E
B	14	4	A
B	55	2	C
C	55	5	B
C	66	2	D
D	66	5	C
D	08	3	E
E	23	3	A
E	08	2	D
F	78	2	A
F	92	3	—

The Dijkstra Algorithm

1. Identify the root (the node itself)
2. Attach all neighbour nodes temporarily
3. Make arc and node with least cumulative cost permanent
4. Choose this node
5. Repeat 2 and 3 until all nodes are permanent

Table 21.2 Link state routing table for router A

Link State Routing Table

Network	Cost	Next Router	Other Information
N1	5	C	
N2	7	D	
N3	10	B	
N4	11	D	
N5	15	C	

Network id

Metric

Host id of interface

Figure 21.7 Areas in an autonomous system

OSPF hierarchy

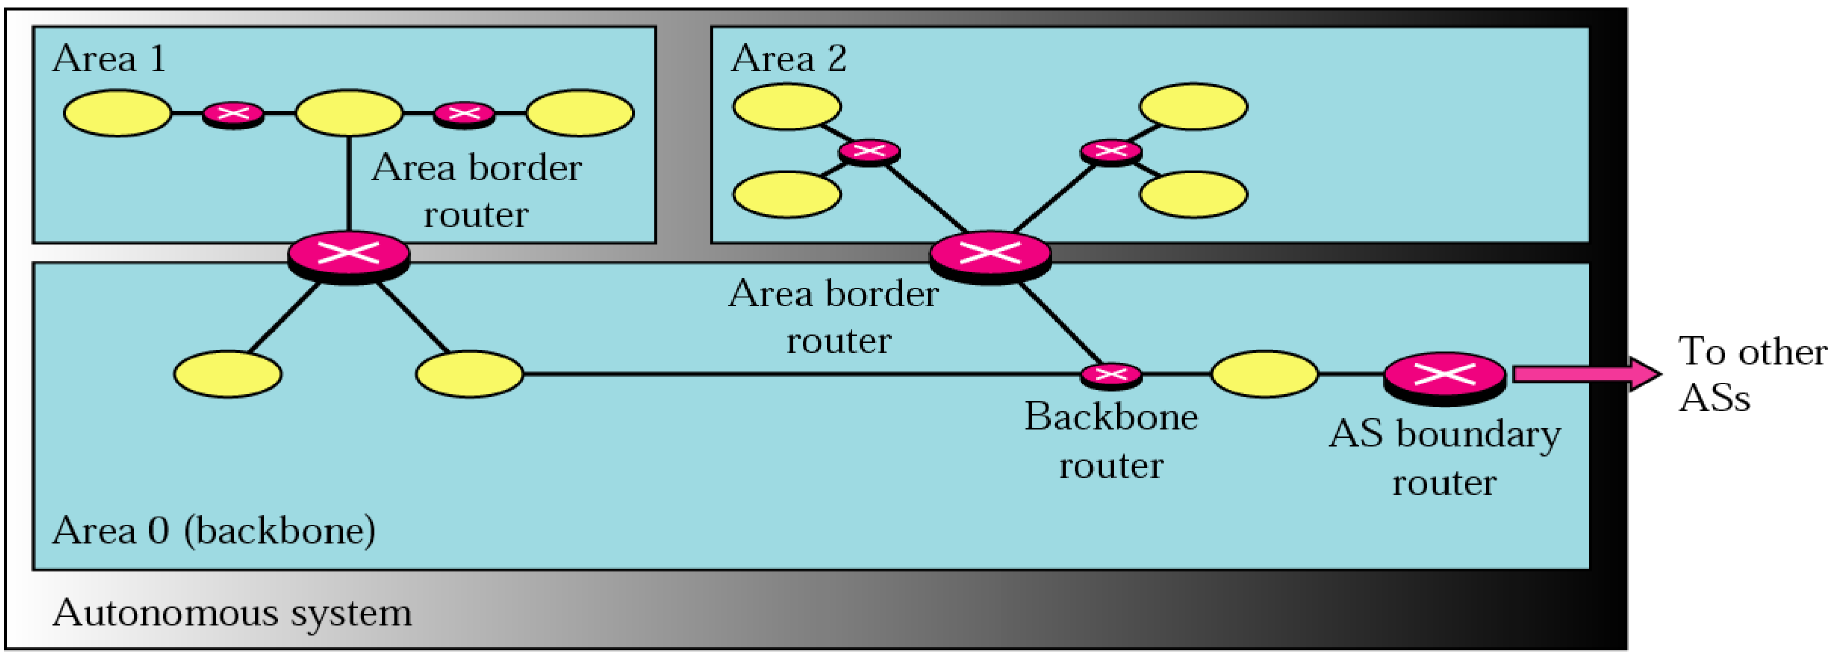

Types of Links

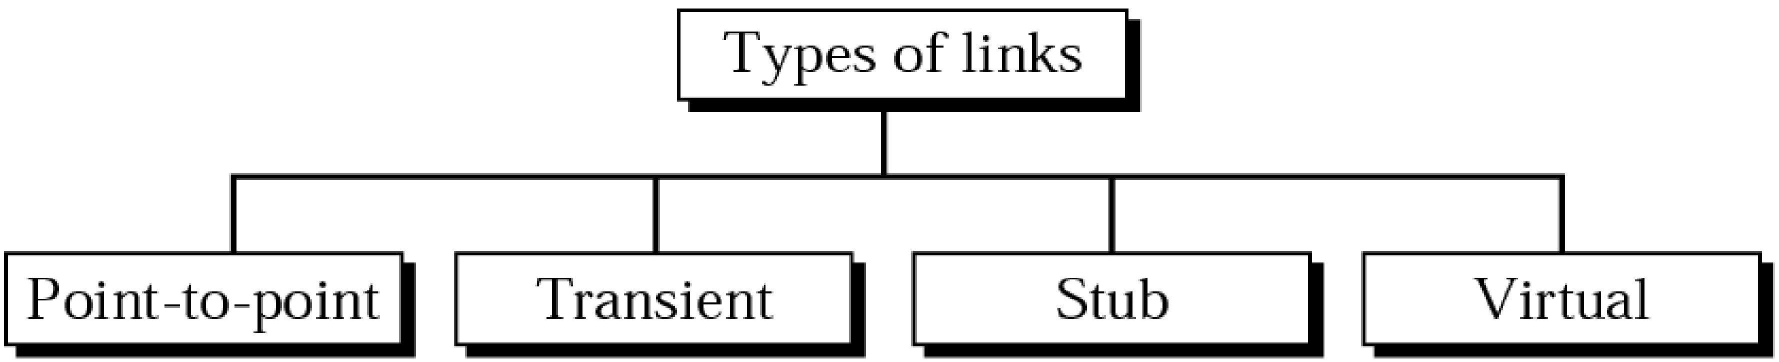

Point-to-Point Link

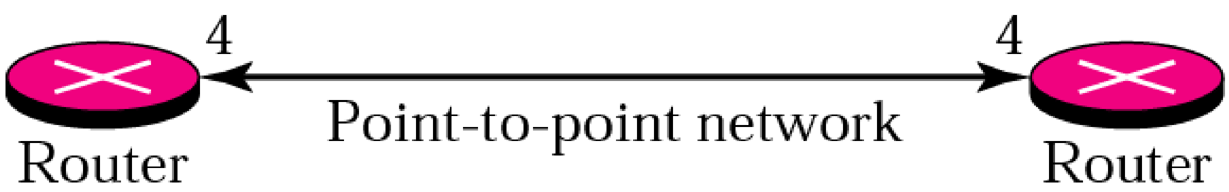

Figure 21.10 Transient link

Transient Link

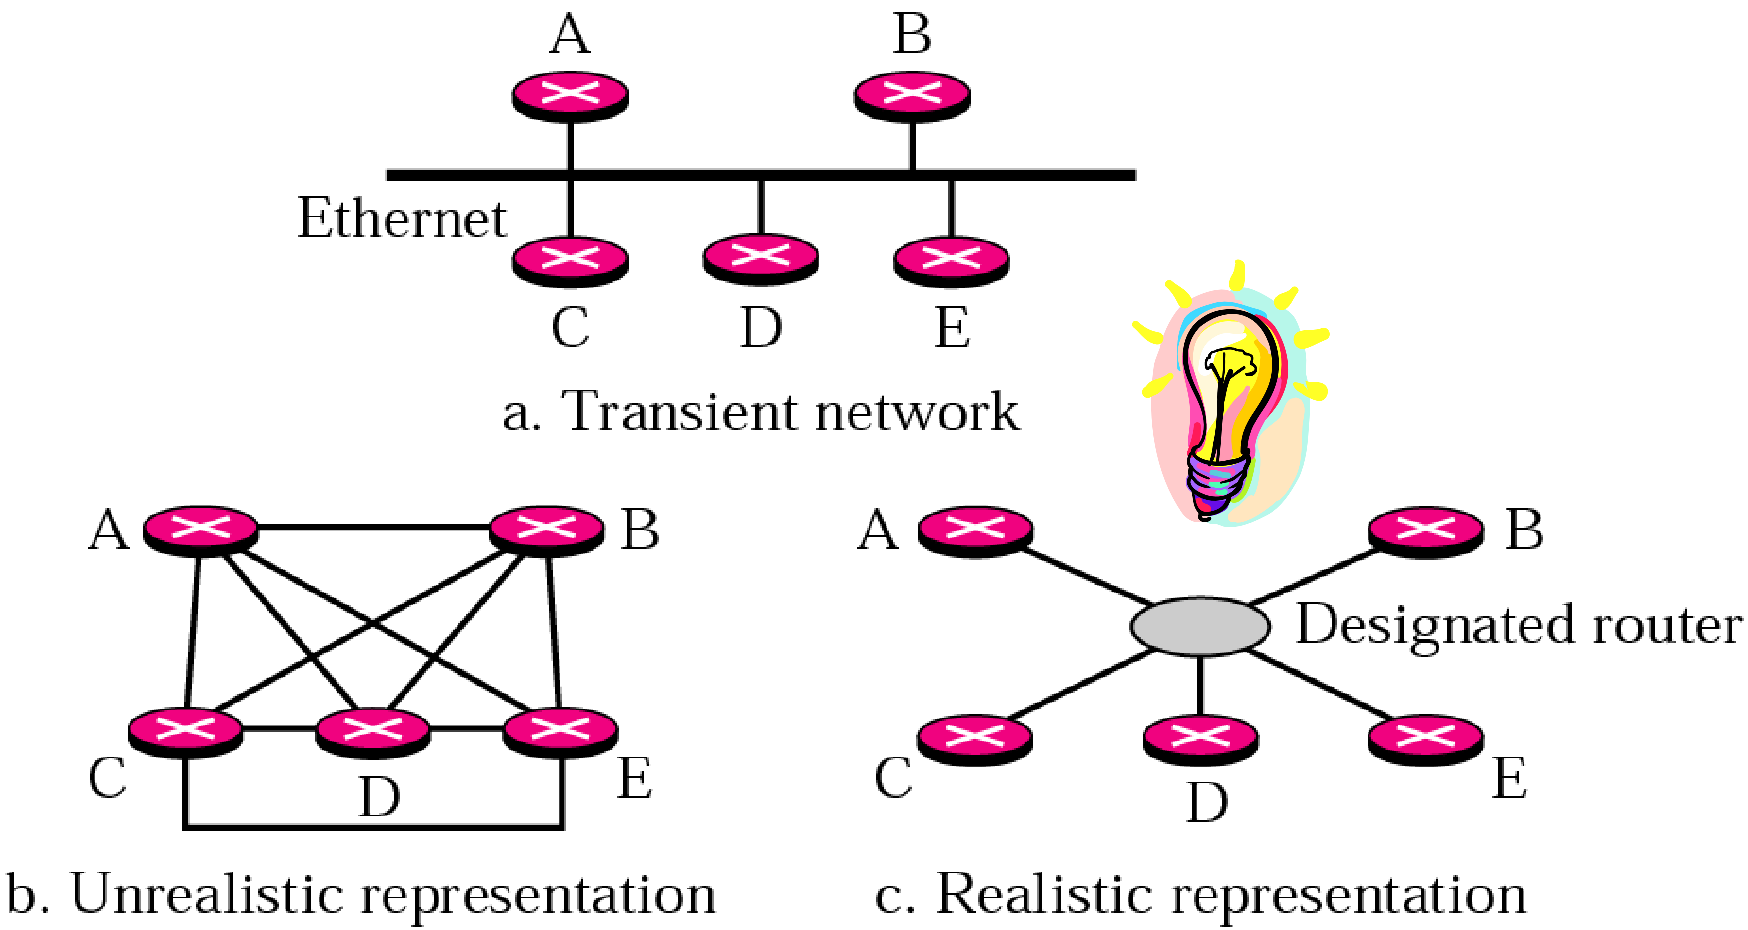

Stub Link

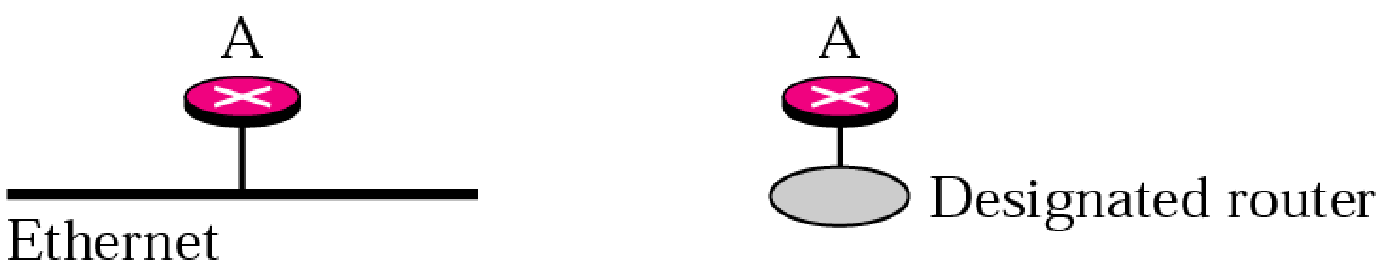

a. Stub network b. Representation

Types of LSAs

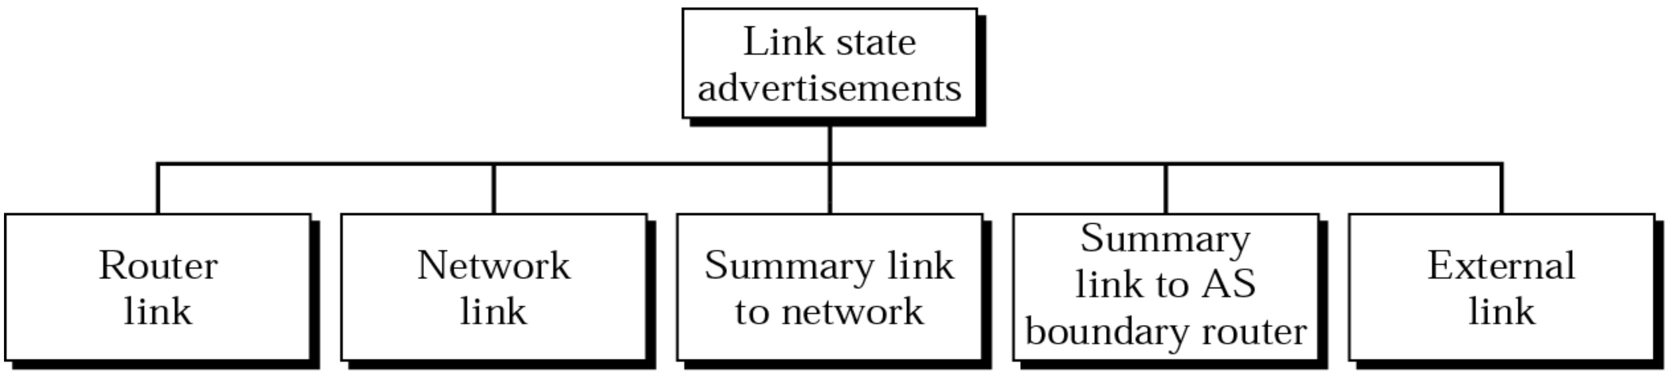

Router Link

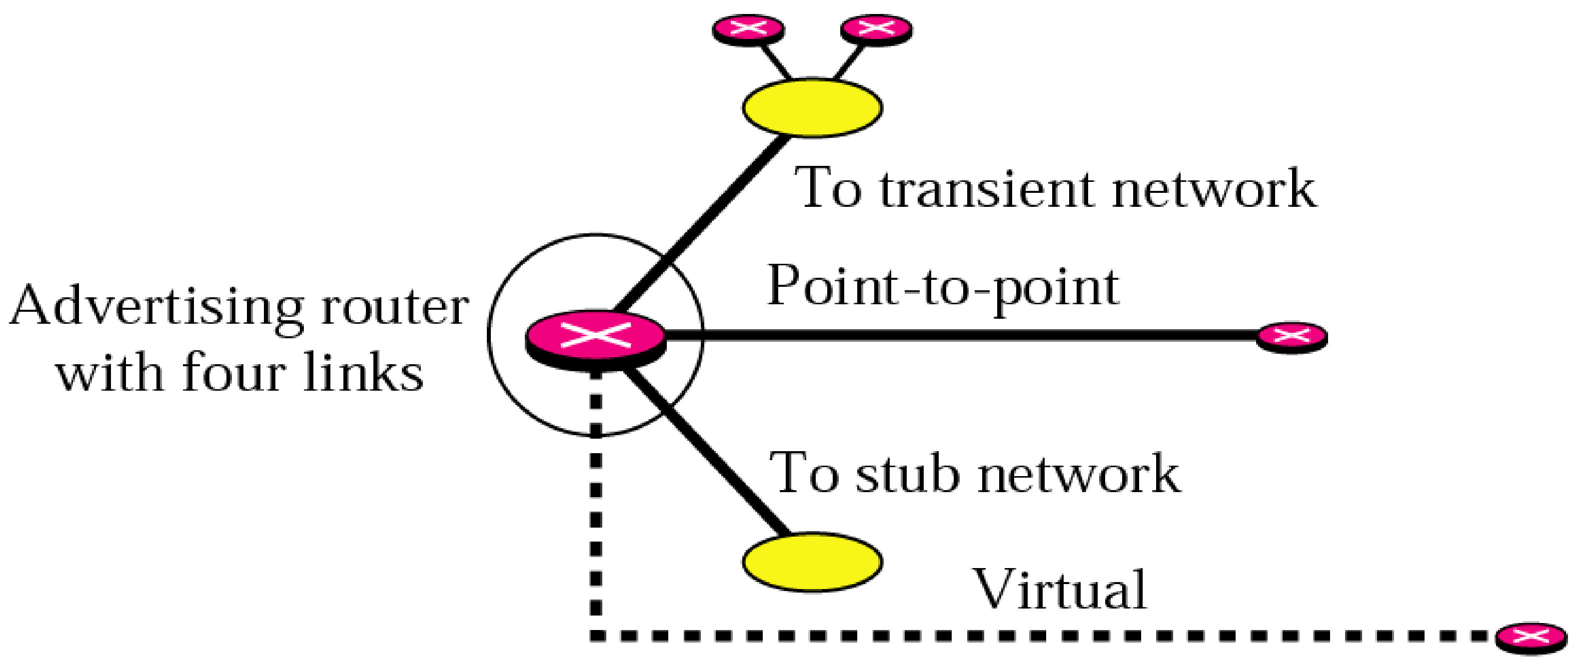
To transient network
Point-to-point
Advertising router
with four links
To stub network
Virtual

Figure 21.16 **Network link**

Network Links

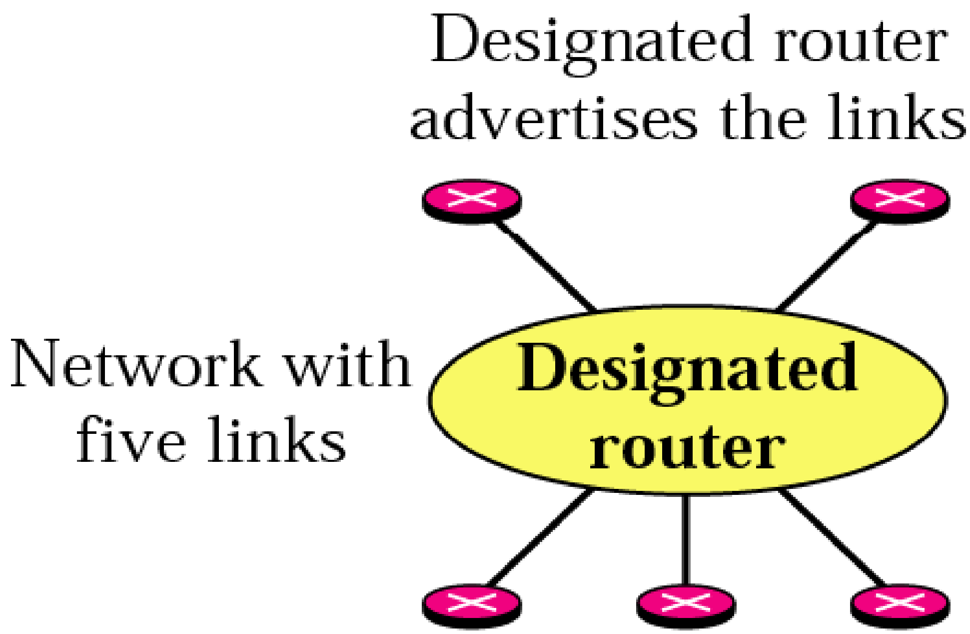

Summary Link to Network

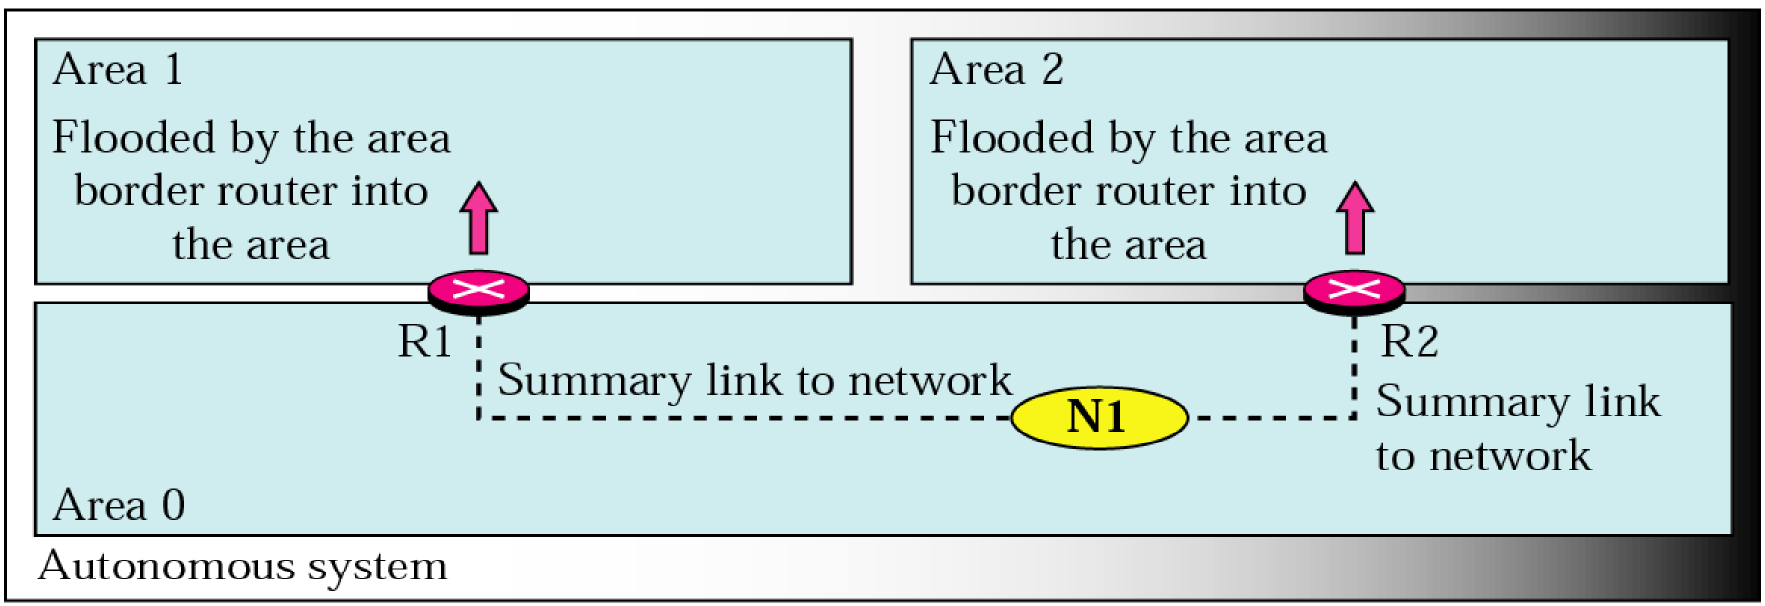

Summary link to AS boundary router

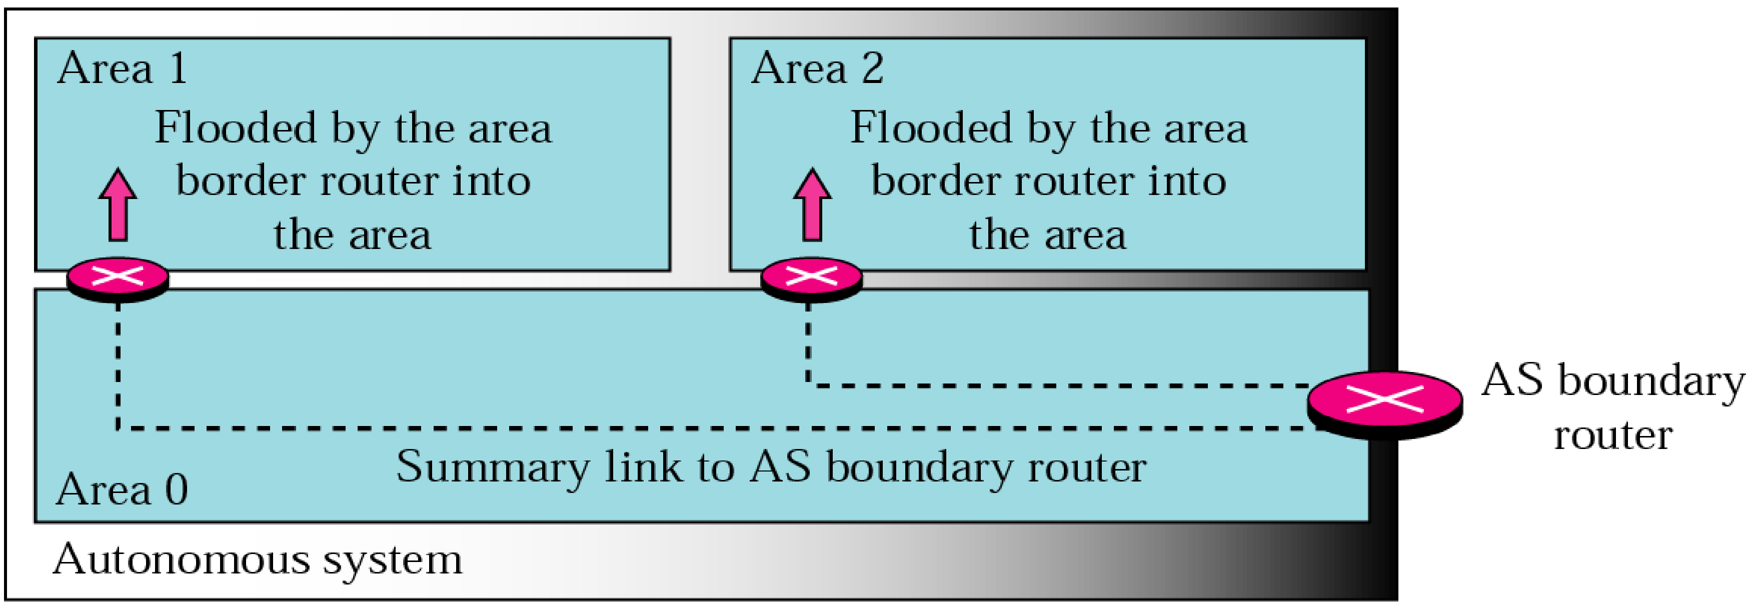

External Link

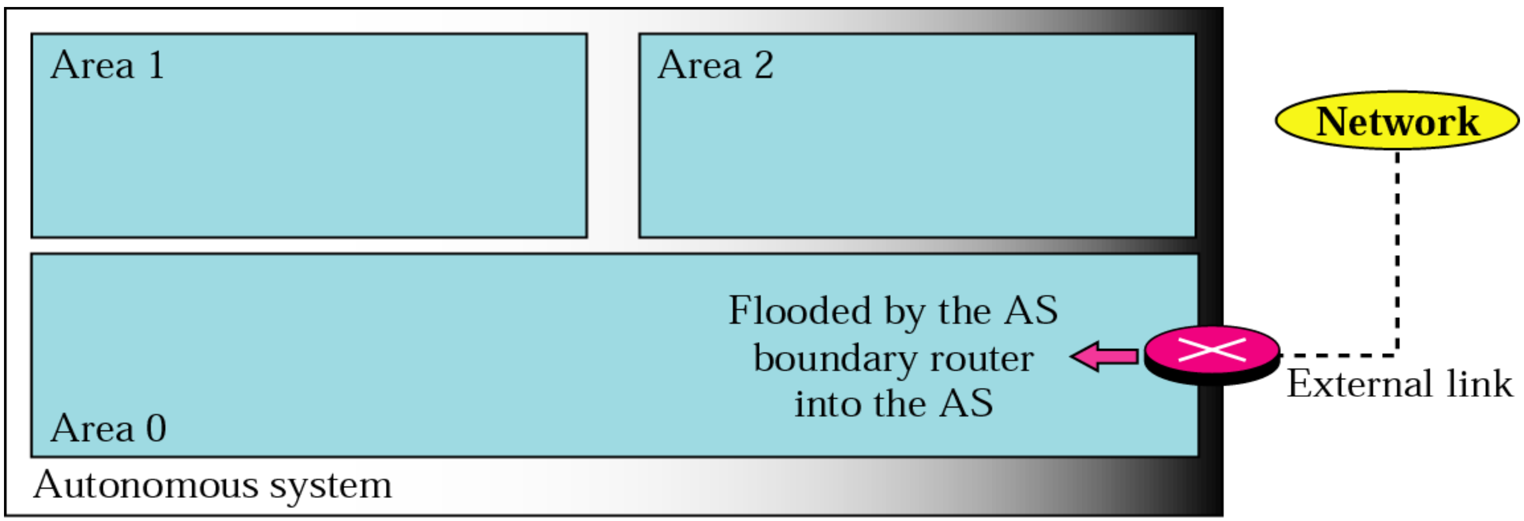

Comparison of LS and DV algorithms

Message complexity

- LS: with n nodes, E links, O(nE) msgs for full knowledge, changes sent to all nodes
- DV: exchange between neighbours only

Speed of Convergence

- LS: $O(n^2)$ algorithm requires O(nE) msgs
  - may have oscillations
- DV: convergence time varies
  - may be routing loops
  - count-to-infinity problem

Robustness: what happens if router malfunctions?

LS:

- node can advertise incorrect *link* cost for its "own" links
- node can break broadcast path or change broadcasted info
- each node computes only its *own* table

DV:

- DV node can advertise incorrect *path* cost to any path
- each node's table used by others
  - error propagate thru network

Figure 22.7 *Address aggregation*

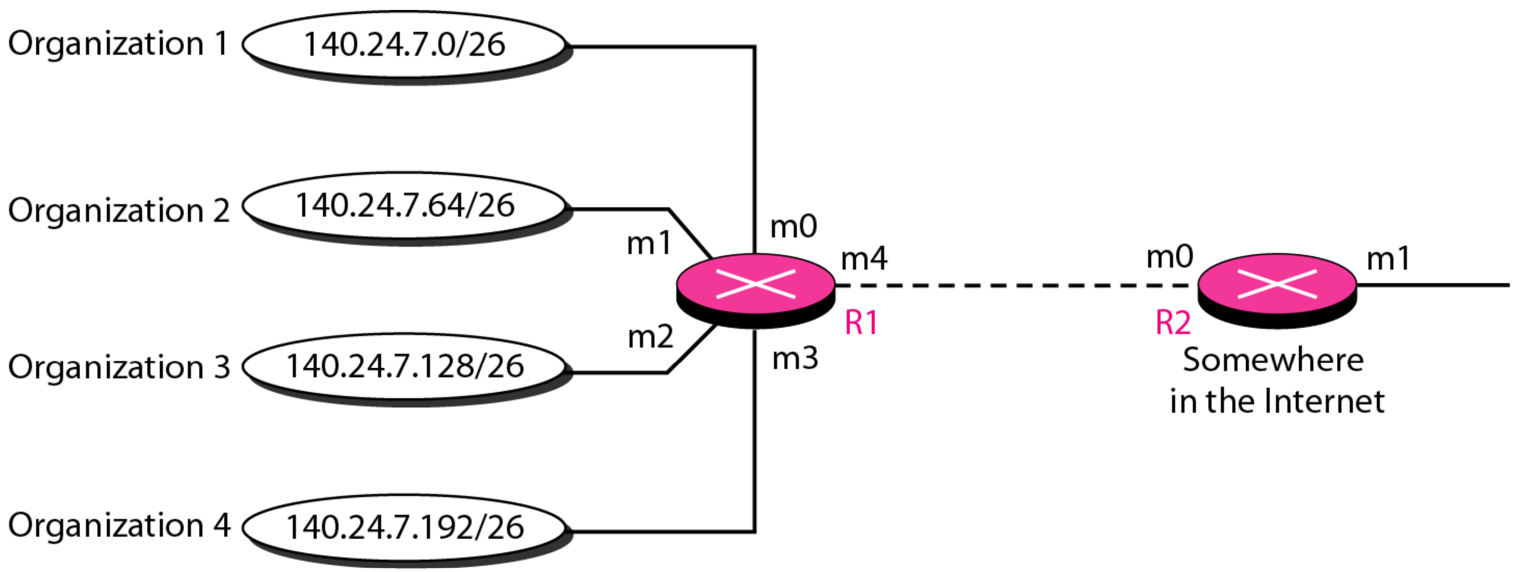

Mask	Network address	Next-hop address	Interface
/26	140.24.7.0	----------	m0
/26	140.24.7.64	----------	m1
/26	140.24.7.128	----------	m2
/26	140.24.7.192	----------	m3
/0	0.0.0.0	Default	m4

Routing table for R1

Mask	Network address	Next-hop address	Interface
/24	140.24.7.0	----------	m0
/0	0.0.0.0	Default	m1

Routing table for R2

Figure 22.8 *Longest mask matching*

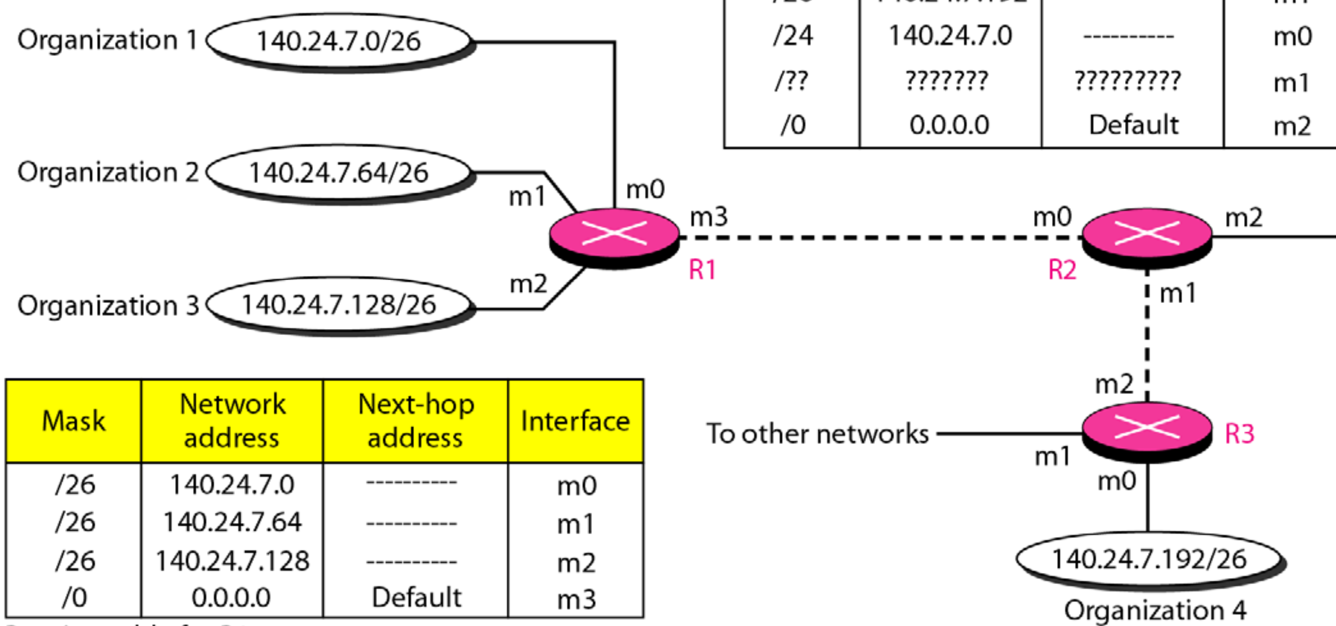

Routing table for R2

Mask	Network address	Next-hop address	Interface
/26	140.24.7.192	----------	m1
/24	140.24.7.0	----------	m0
/??	???????	?????????	m1
/0	0.0.0.0	Default	m2

Mask	Network address	Next-hop address	Interface
/26	140.24.7.0	----------	m0
/26	140.24.7.64	----------	m1
/26	140.24.7.128	----------	m2
/0	0.0.0.0	Default	m3

Routing table for R1

Mask	Network address	Next-hop address	Interface
/26	140.24.7.192	----------	m0
/??	???????	?????????	m1
/0	0.0.0.0	Default	m2

Routing table for R3

Figure 22.9 *Hierarchical routing with ISPs*

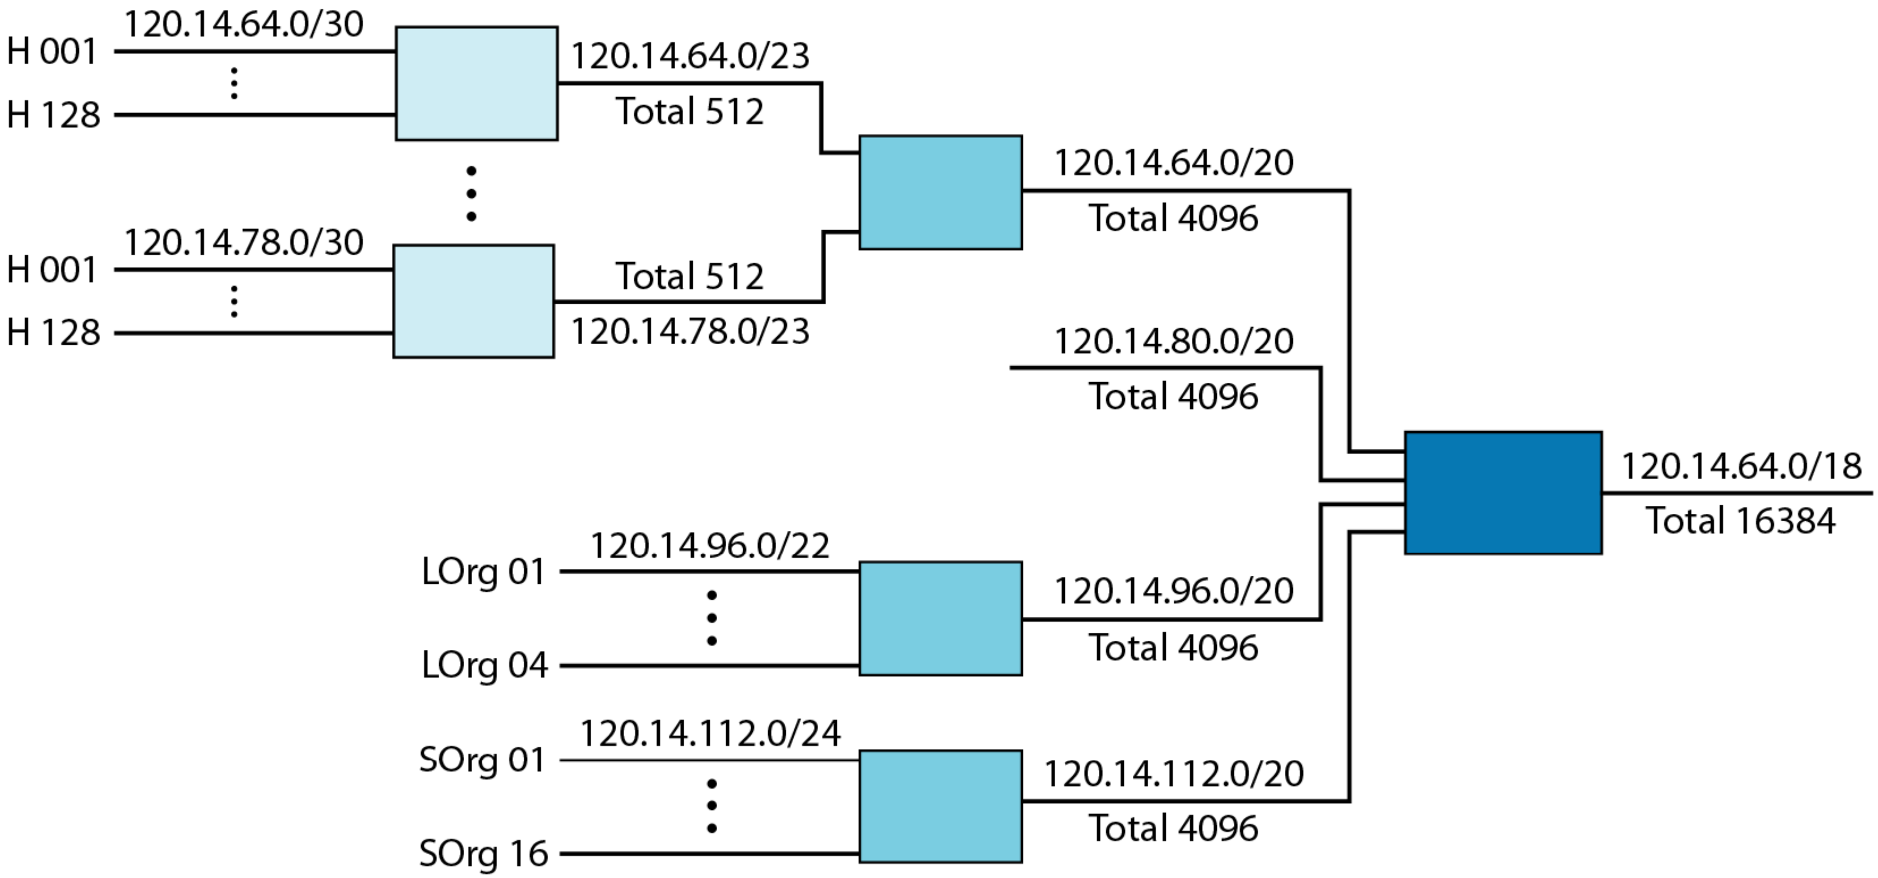

Source Address routing

- Route according to source address?!
- Example:
  - Customers share subnet (different ISPs; unlikely)
  - Customers have different traffic policies

Routing

- Introduction
- Inside the Router
- Unicast Routing
  - Intra Domain Routing
  - Inter Domain Routing
- MANET and AdHoc routing
- Multicast Routing (Lecture 3)

MANET routing

- Mobile Ad hoc NETwork
- Ad hoc = dynamic network structure
  - members come and go
- Network nodes move!
- Each host also a routing node

An Ad Hoc Network

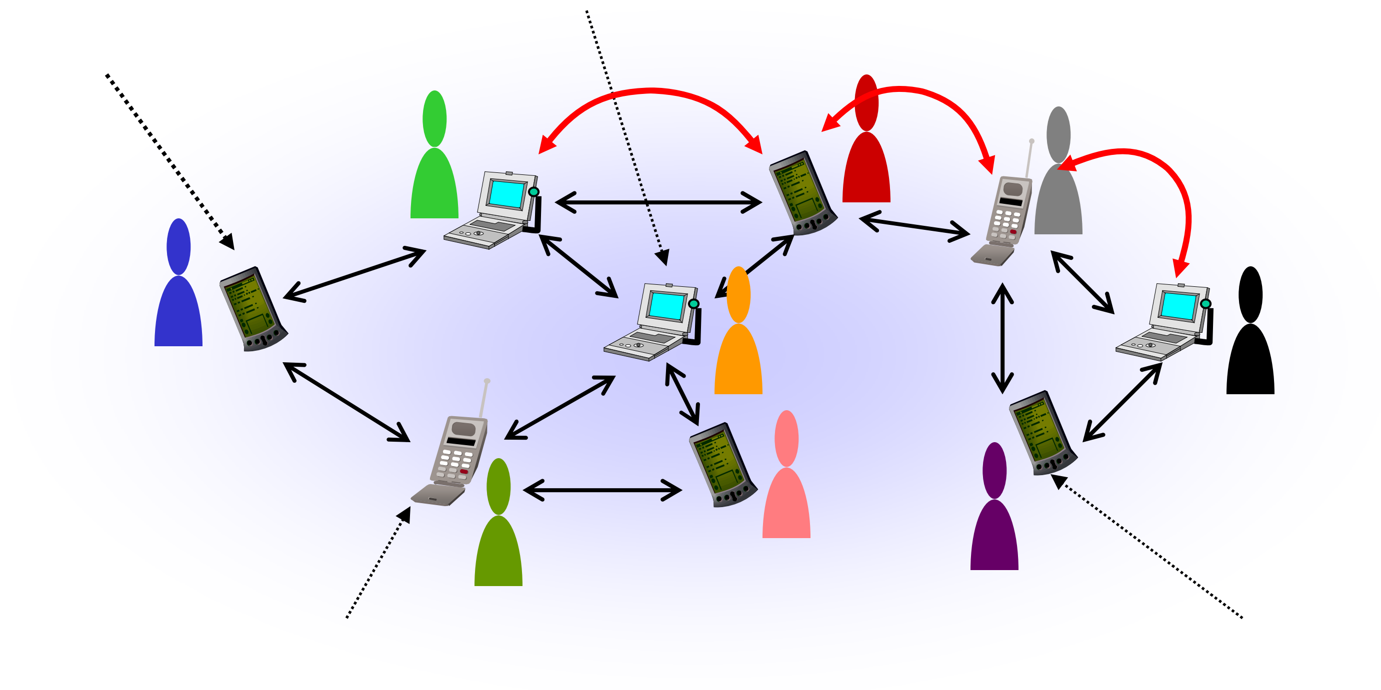

AdHoc: Special considerations

- Are all stations willing to forward other’s packets?
- Do I trust all members in this Ad Hoc net?
- Power consumption is one metric
- Forwarding capacity is one metric
- MANET: Nodes are moving!

AdHoc routing

- Proactive
  - complete routing info at hand all the time
  - no special action before sending
  - lots of energy lost in keeping track of paths never used
- Reactive/On demand
  - find best path when connection needed
  - only used paths are exploited
  - delay before connection can be used