計算機ネットワーク

開講クォーター: 1-2Q

曜日・時限: 火7-8限

講義室: 1Q @ W834, 2Q @ W931

横田理央

rioyokota@gsic.titech.ac.jp

参考書

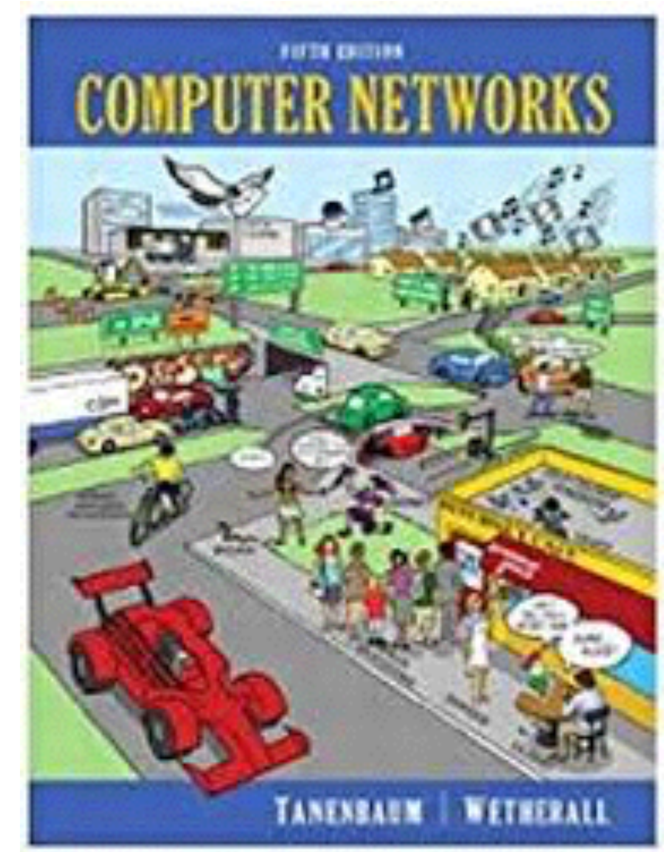

教科書

講義日程（2Q）

		授業計画		課題
06/14	第９回	ネットワーク層１ ルーティング・輻輳制御	5章	ルーティングの種類を理解し 輻輳制御手法を説明できる
06/21	第１０回	ネットワーク層２ インターネットとサービス品質	5章	インターネットの制御プロトコルを理解し ネットワーク間の接続について説明できる
06/28	第１１回	トランスポート層１ トランスポート・プロトコルの要素	6章	誤り制御とフロー制御を理解し 輻輳制御について説明できる
07/05	第１２回	トランスポート層２ UDP と TCP	6章	TCP の信頼性を理解し TCP のコネクション管理を説明できる
07/12	第１３回	アプリケーション層 DNS, 電子メール, www	7章	DNS, 電子メール, www のしくみを理解し ストリーミング，P2P について説明できる
07/26	第１４回	ネットワークセキュリティ１ 対称鍵暗号，公開鍵暗号	8章	暗号アルゴリズムを理解し SHA-1,2 と RSA について説明できる
08/02	第１５回	ネットワークセキュリティ２ デジタル署名，認証プロトコル	8章	電子メール,Web のセキュリティ の脅威について把握できる

Store-and-Forward Packet Switching

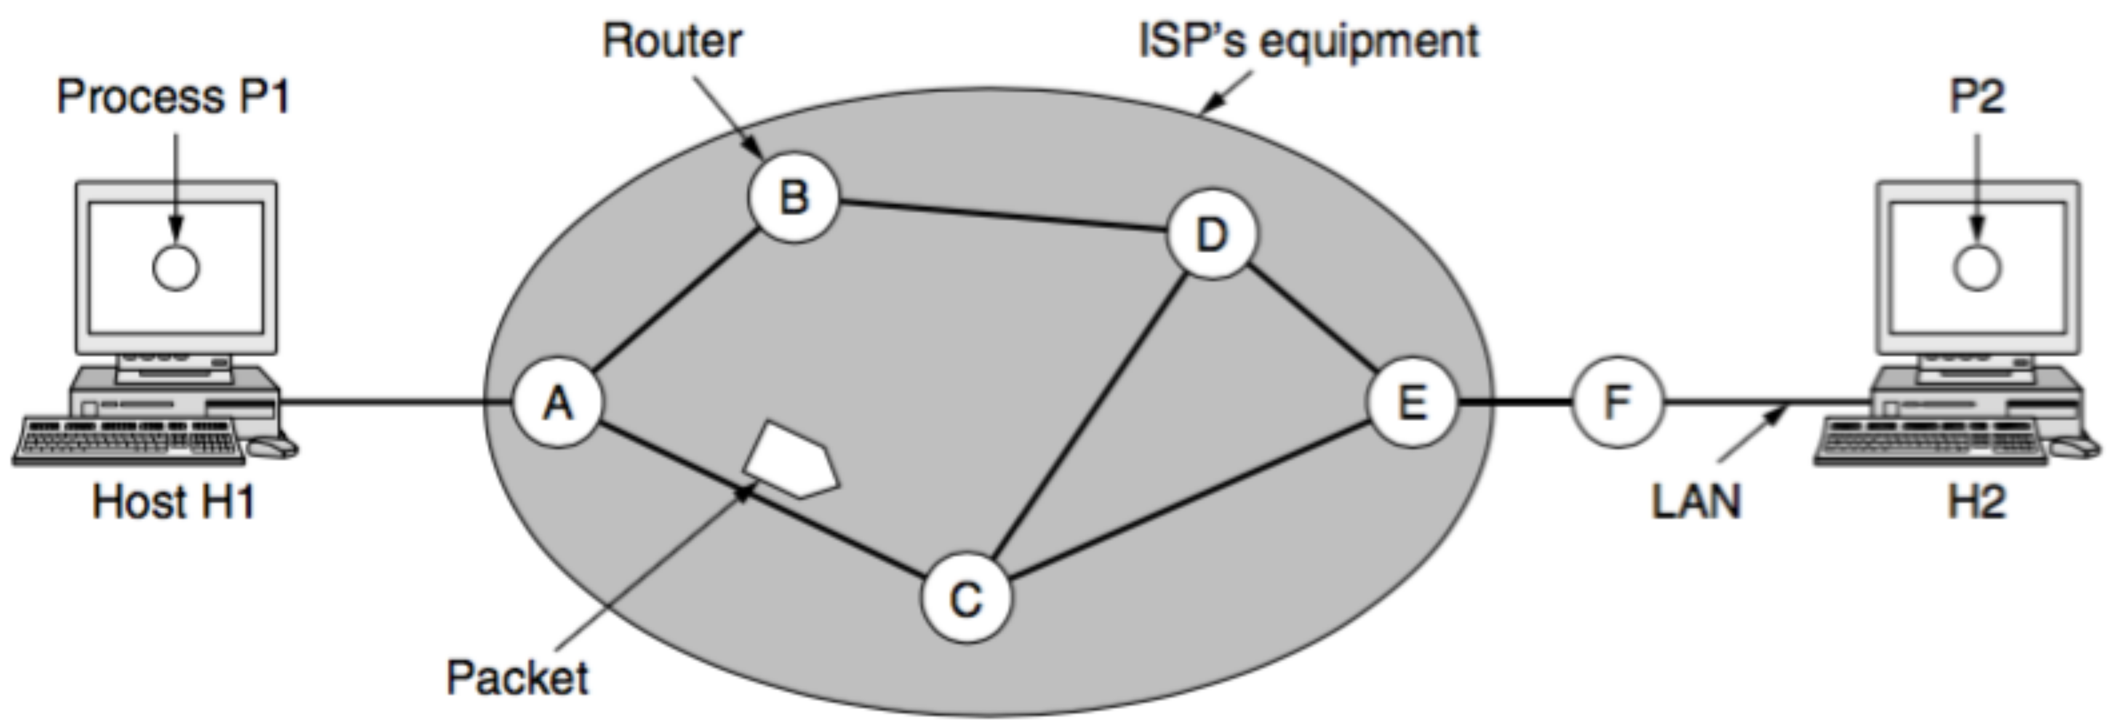

1. A host with a packet to send transmits it to the nearest router

2. The packet is stored there until it has fully arrived

3. Then it is forwarded to the next router along the path

4. This is repeated until it reaches its destination host

Services Provided to the Transport Layer

1. The services should be independent of the router technology
2. The transport layer should be shielded from the number, type, and topology of the routers
3. The network addresses made available to the transport layer should use a uniform numbering plan, even across LANs and WANs

<u>Connectionless</u>

Example: Internet

Reasoning: No packet ordering or flow control should be done

<u>Connection-oriented</u>

Example: Telephone

Reasoning: Quality of service for realtime traffic is important

Implementation of Connectionless Service

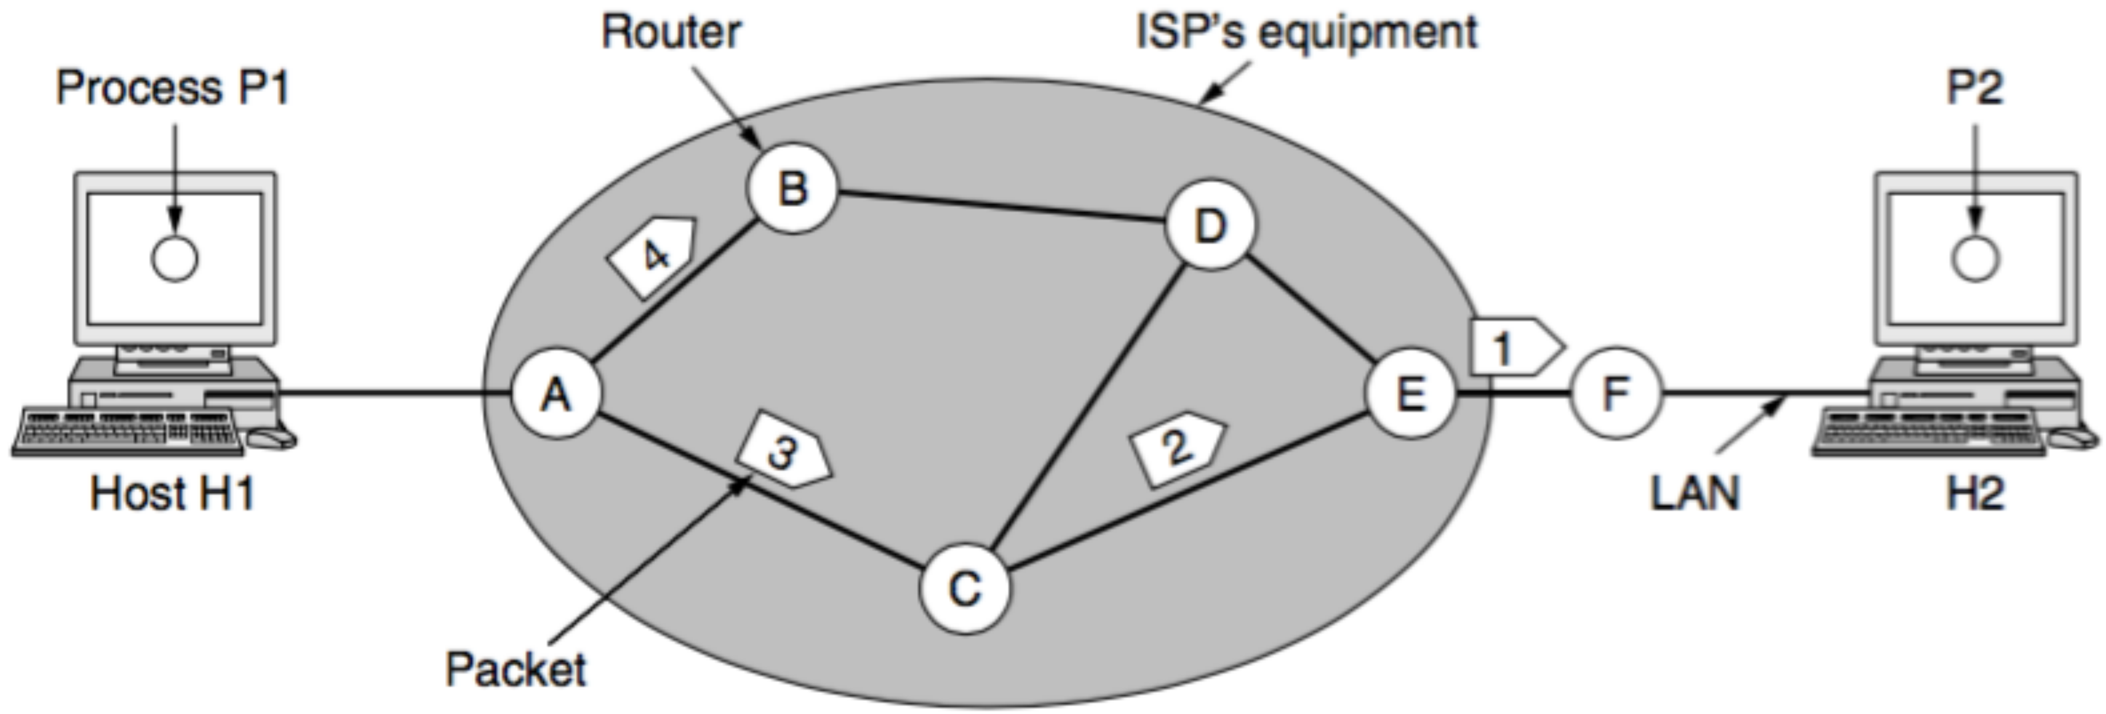

A's table (initially)

Dest.	Line
A	–
B	B
C	C
D	B
E	C
F	C

A's table (later)

A	–
B	B
C	C
D	B
E	B
F	B

C's table

A	A
B	A
C	–
D	E
E	E
F	E

E's table

A	C
B	D
C	C
D	D
E	–
F	F

Datagram network

1. Every router has a table of destinations and outgoing lines
2. The routing table is updated before packet 4 was sent

Implementation of Connection-Oriented Service

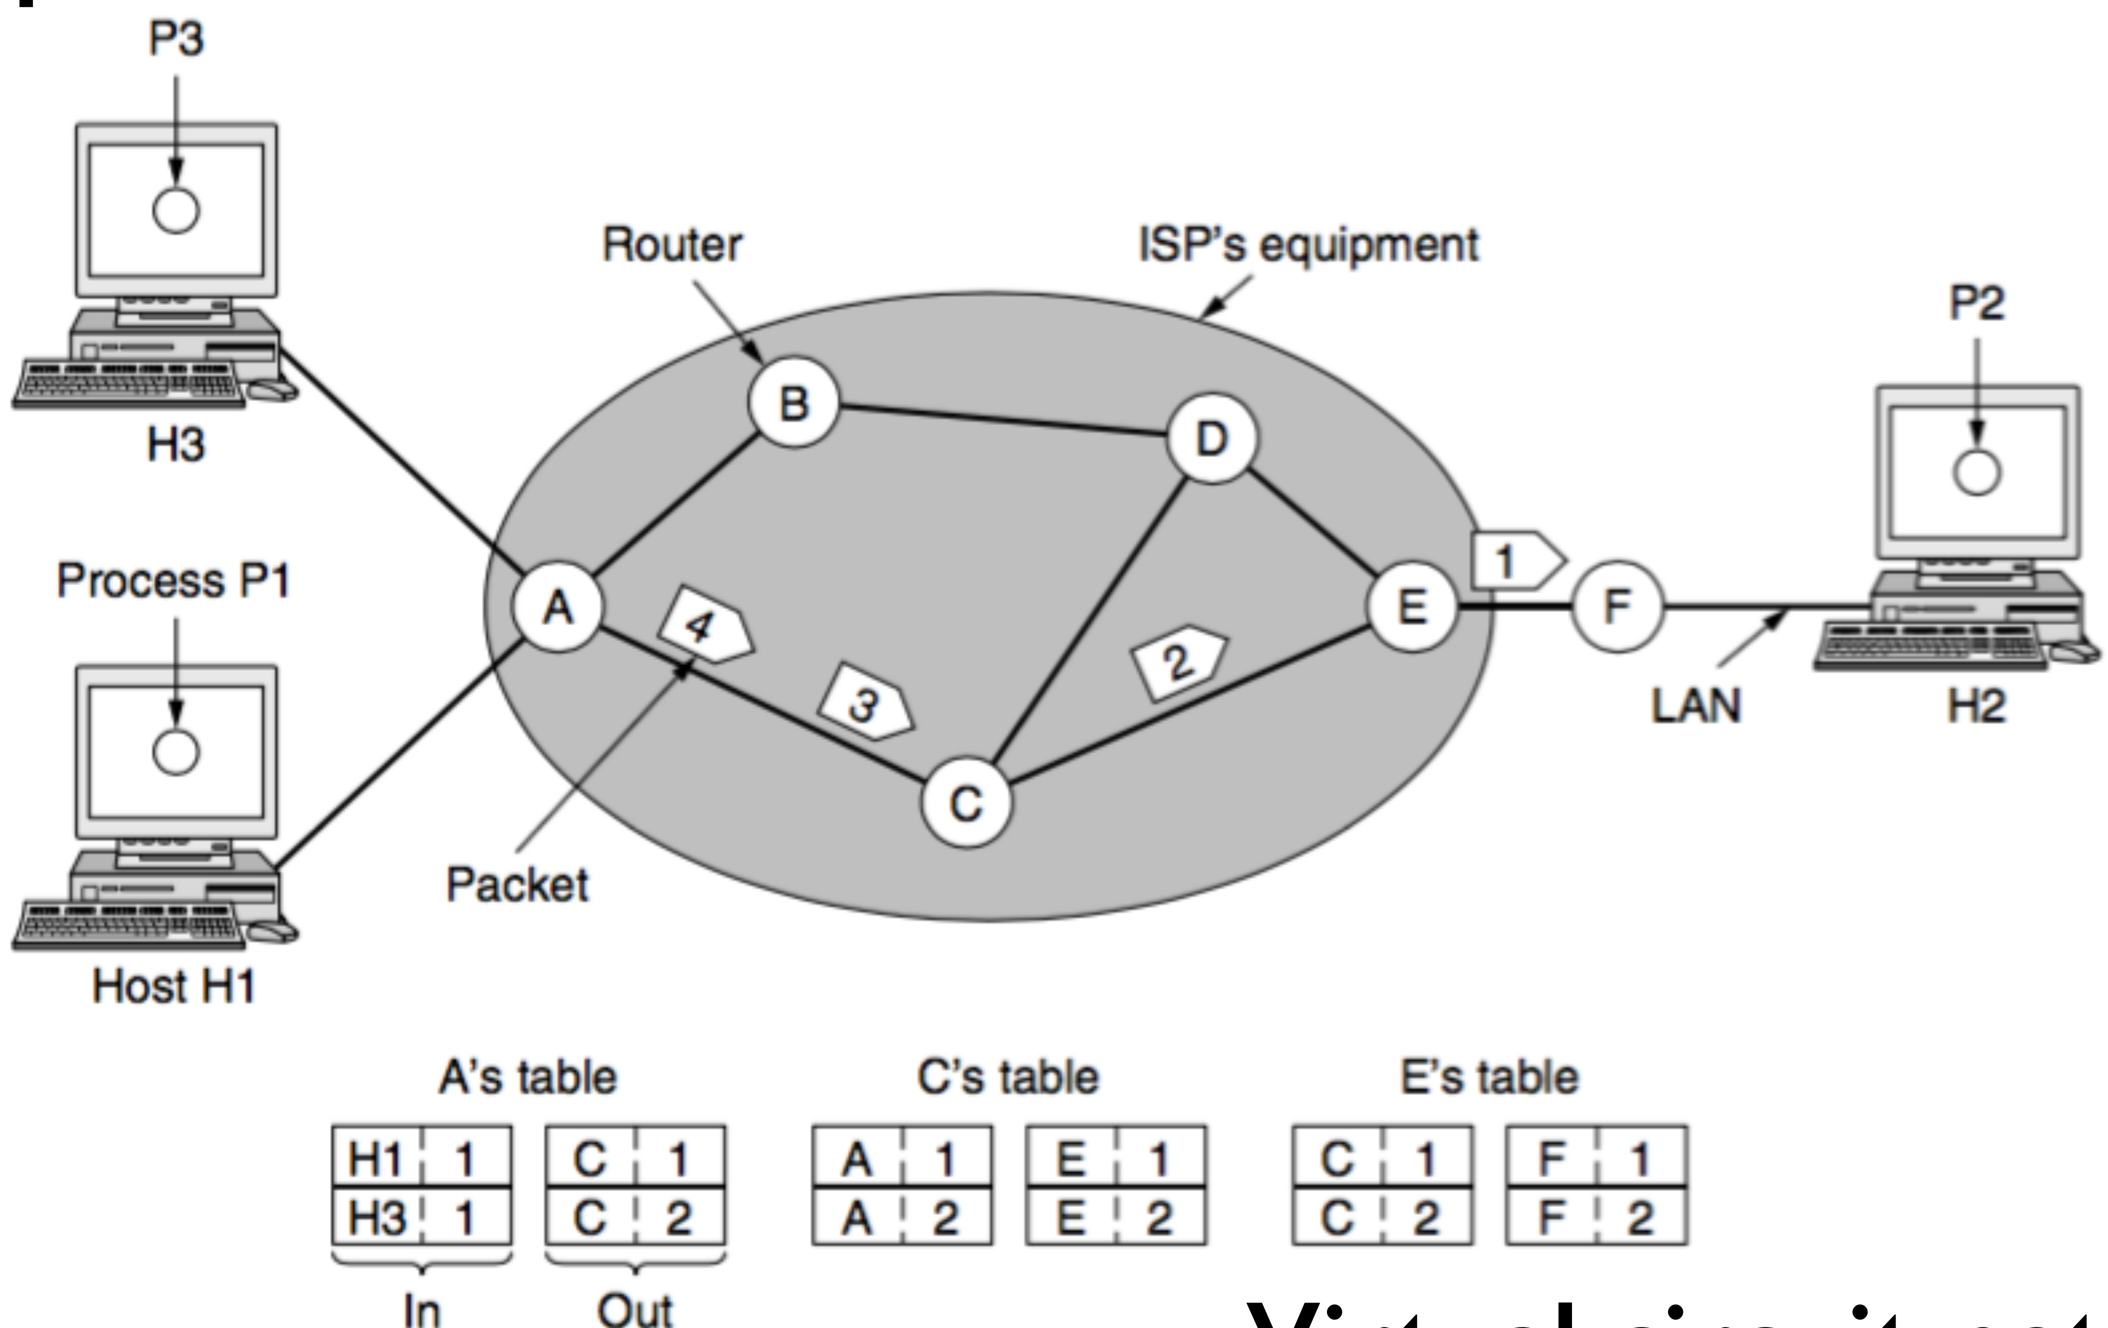

Virtual circuit network

1. A route from the source to the destination is chosen as part of the setup of the connection
2. This virtual circuit is stored as tables inside the routers
3. No updates to the tables until connection is terminated

Virtual-Circuit vs. Datagram Networks

Issue	Datagram network	Virtual-circuit network
Circuit setup	Not needed	Required
Addressing	Each packet contains the full source and destination address	Each packet contains a short VC number
State information	Routers do not hold state information about connections	Each VC requires router table space per connection
Routing	Each packet is routed independently	Route chosen when VC is set up; all packets follow it
Effect of router failures	None, except for packets lost during the crash	All VCs that passed through the failed router are terminated
Quality of service	Difficult	Easy if enough resources can be allocated in advance for each VC
Congestion control	Difficult	Easy if enough resources can be allocated in advance for each VC

Credit card transactions VPN

Routing Algorithms

<u>Desirable properties of routing algorithms</u>

1. Correctness
2. Simplicity
3. Robustness
4. Stability
5. Fairness
6. Efficiency

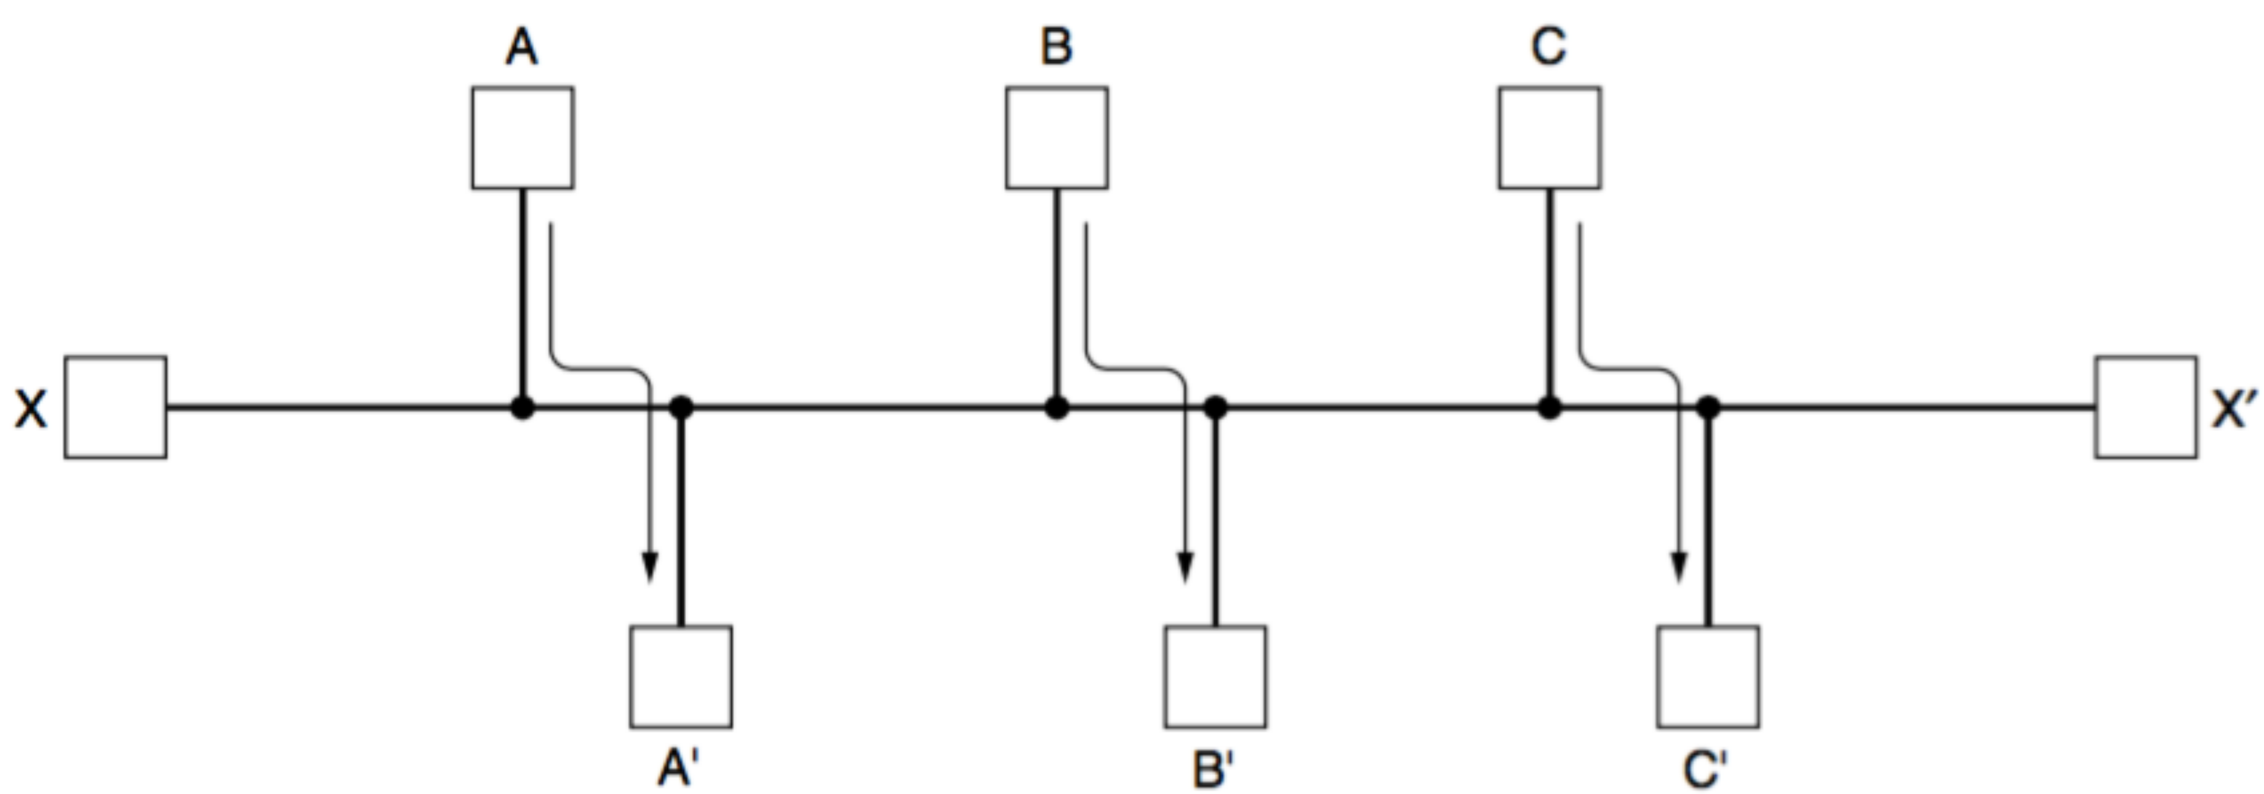

Routing Algorithms

<u>What do we seek to optimize?</u>

1. Minimizing mean packet delay
2. Maximizing total network throughput
3. Minimize the distance a packet has to travel
4. Minimize the number of hops the packet must take

<u>The optimality principle</u>

If router J is on the optimal path from router I to router K, then the optimal path from J to K also falls along the same route.

<u>Sink Tree</u>

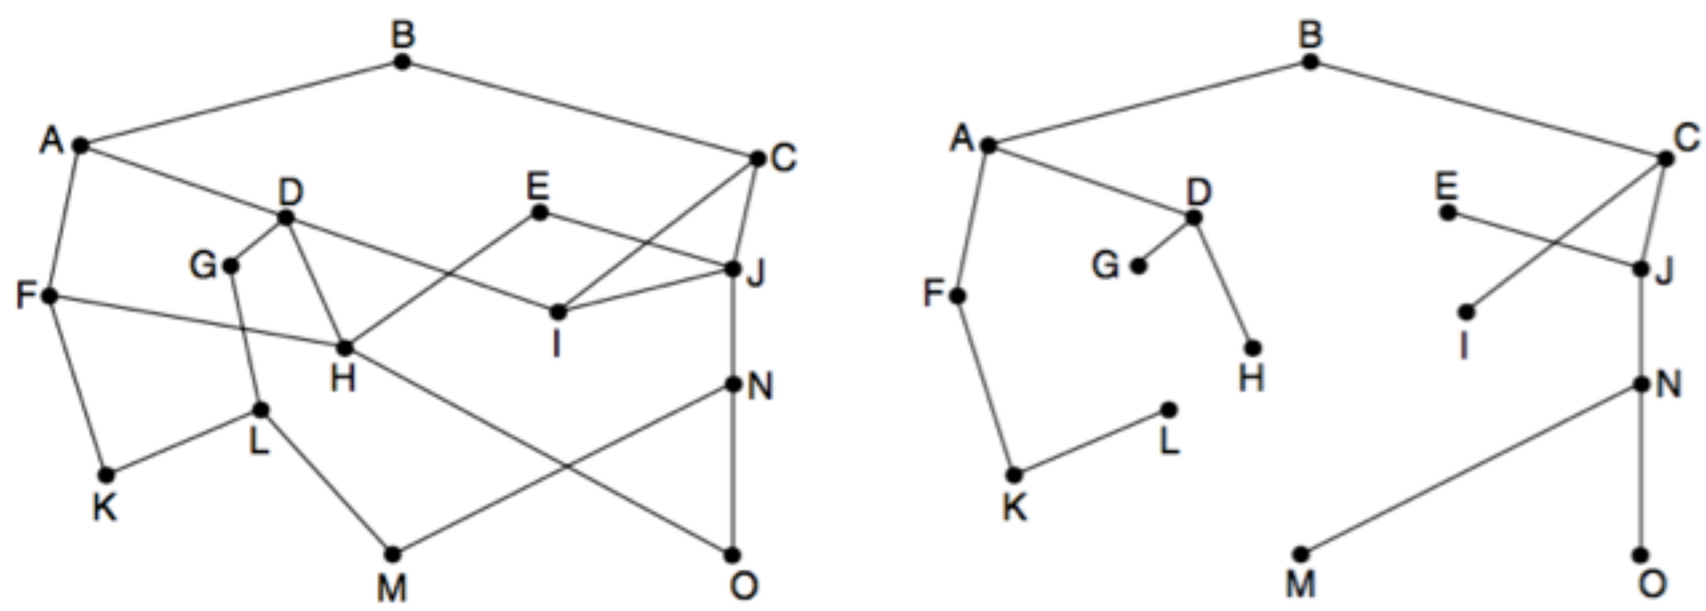

(a) (b)

Shortest Path Algorithm

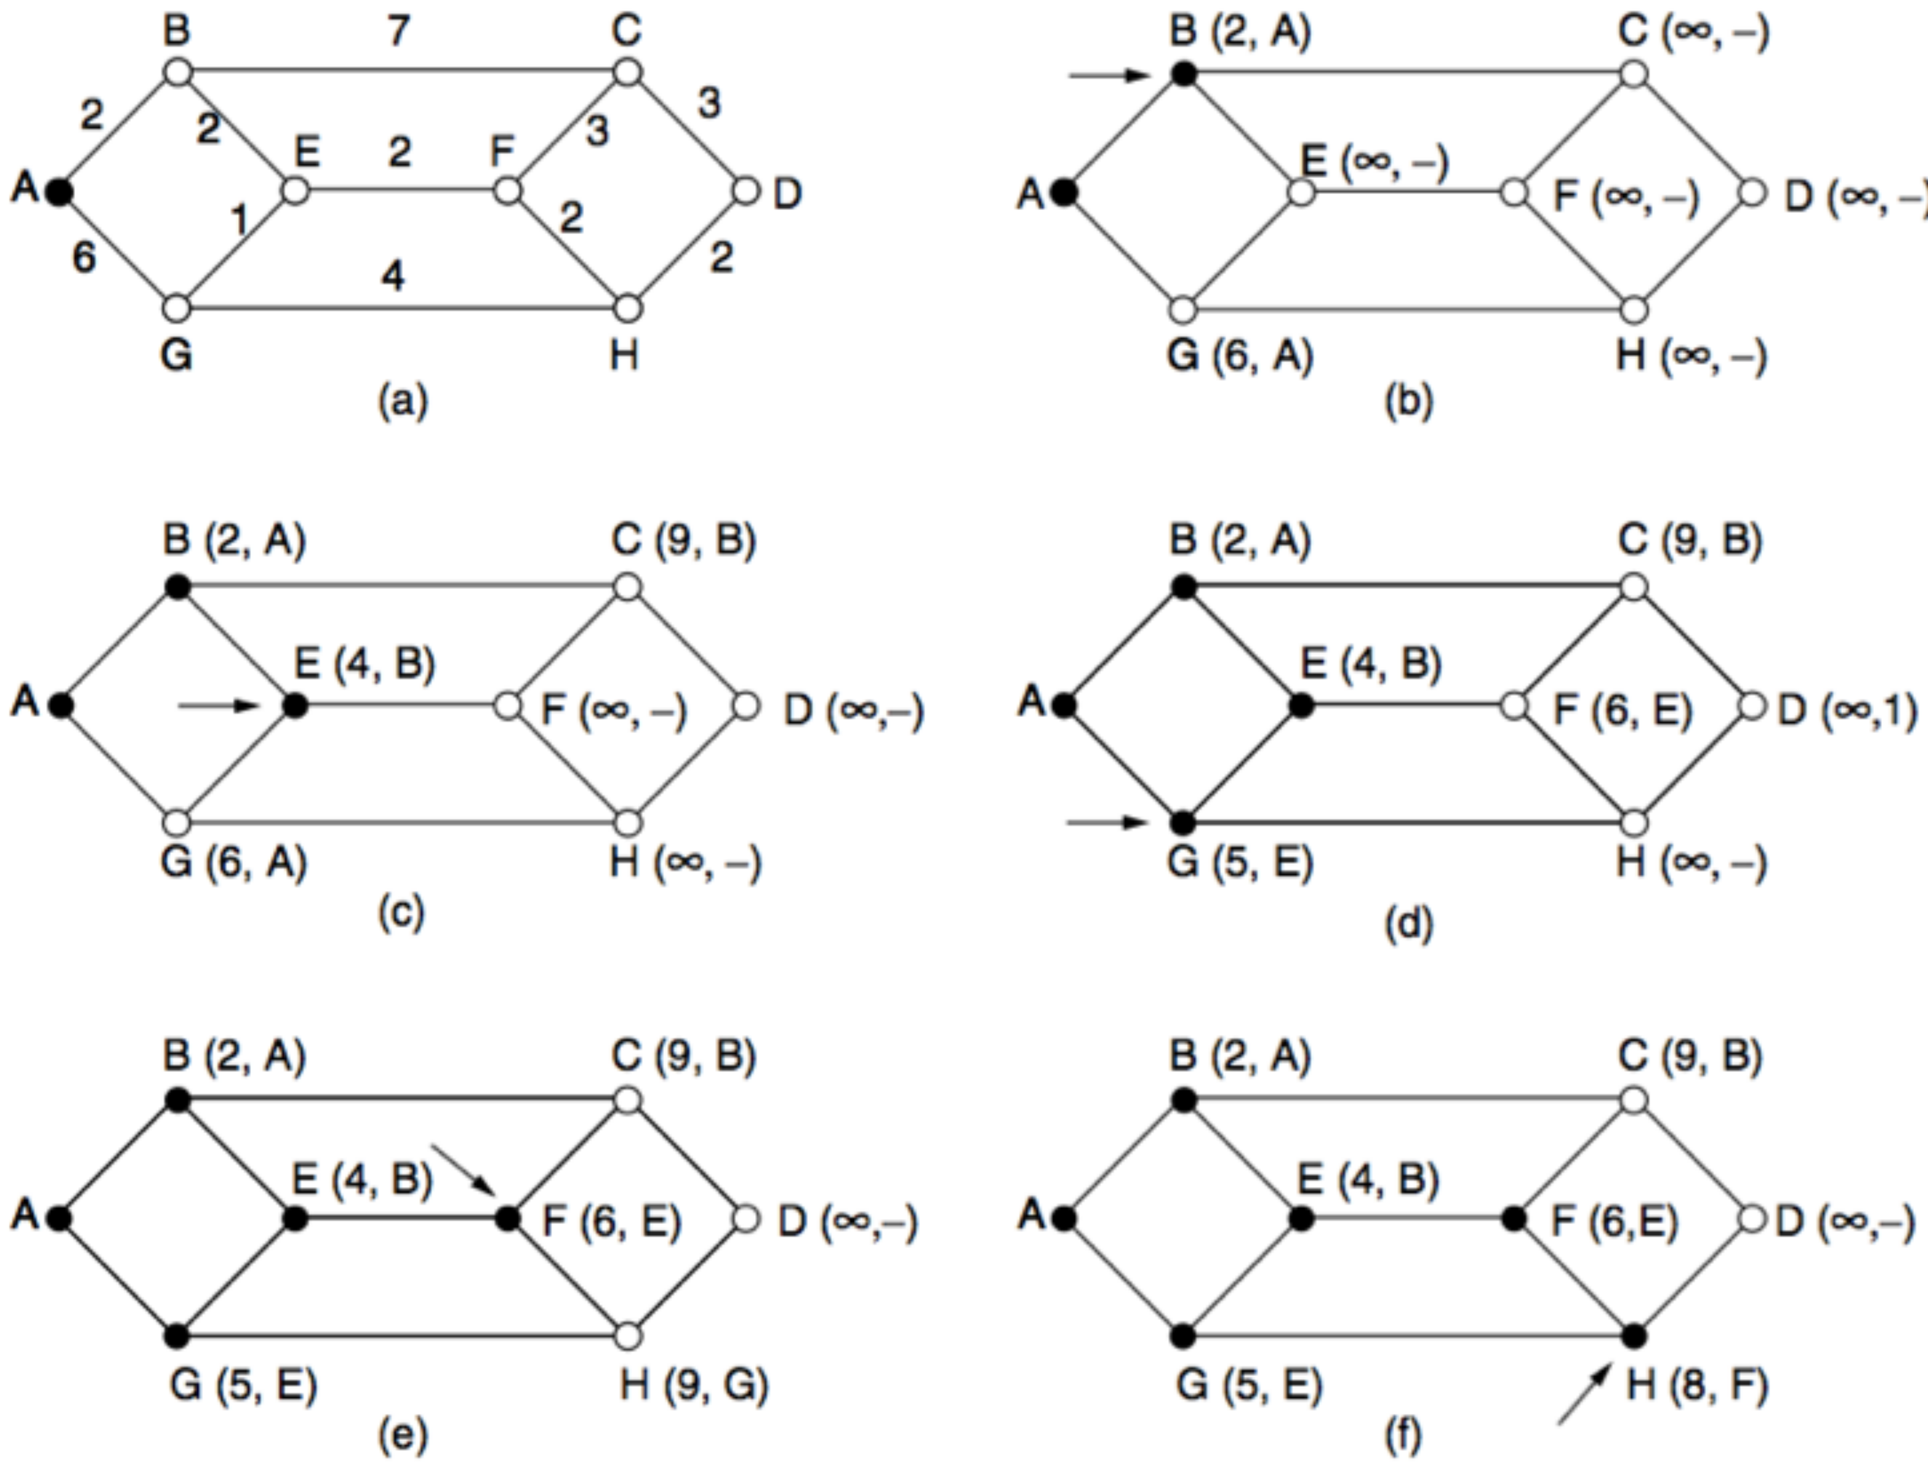

Examine all the tentatively labeled nodes in the whole graph and make the one with the smallest label permanent

Flooding

Every incoming packet is sent out on every outgoing line except the one it arrived on.

- It ensures that a packet is delivered to every node in the net- work.
- Flooding is tremendously robust
- flooding can be used as a building block for other routing algorithms that are more efficient but need more in the way of setup

Keep track of which packets have been flooded, to avoid sending them out a second time.

One way to achieve this goal is to have the source router put a sequence number in each packet it receives from its hosts. Each router then needs a list per source router telling which sequence numbers originating at that source have already been seen.

Distance Vector Routing (Bellman-Ford Routing)

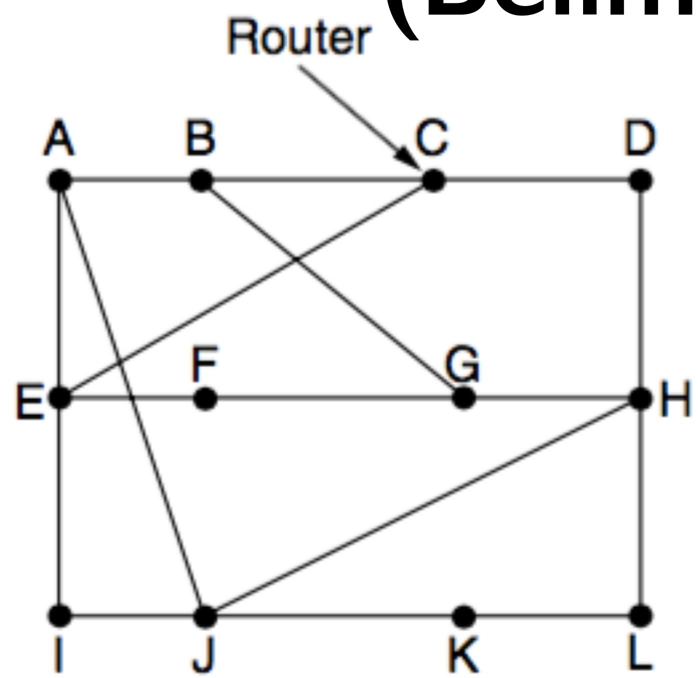

(a)

To	A	I	H	K	New estimated delay from J	Line
A	0	24	20	21	8	A
B	12	36	31	28	20	A
C	25	18	19	36	28	I
D	40	27	8	24	20	H
E	14	7	30	22	17	I
F	23	20	19	40	30	I
G	18	31	6	31	18	H
H	17	20	0	19	12	H
I	21	0	14	22	10	I
J	9	11	7	10	0	–
K	24	22	22	0	6	K
L	29	33	9	9	15	K
	JA delay is 8	JI delay is 10	JH delay is 12	JK delay is 6	New routing table for J	

Vectors received from J's four neighbors

(b)

Each router maintain a table giving the best known distance to each destination and which link to use to get there, which are updated by exchanging information with the neighbors.

Count-to-Infinity Problem

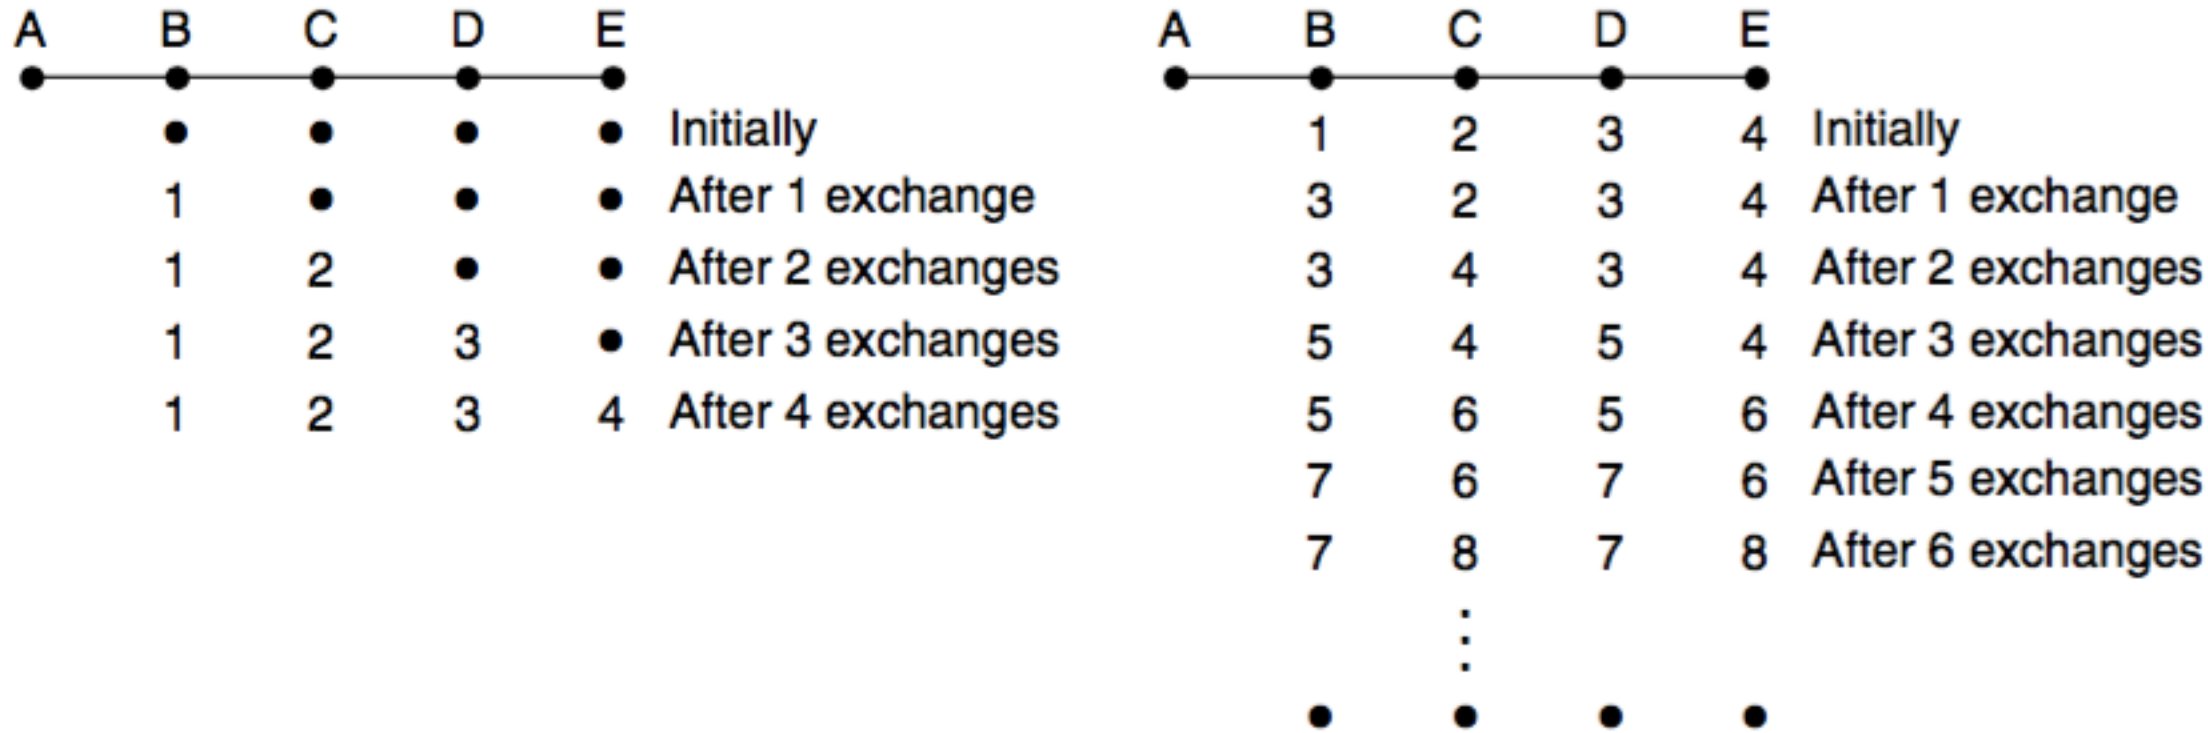

(a) (b)

Suppose A is down initially and all the other routers know this

Good news spreads at one hop per exchange

Bad news travels slowly while all routers work their way up to infinity

It is wise to set infinity to the longest path plus 1

Link State Routing

1. Discover its neighbors and learn their network addresses
2. Set the distance or cost metric to each of its neighbors
3. Construct a packet telling all it has just learned
4. Send this packet to and receive packets from all other routers
5. Compute the shortest path to every other router.

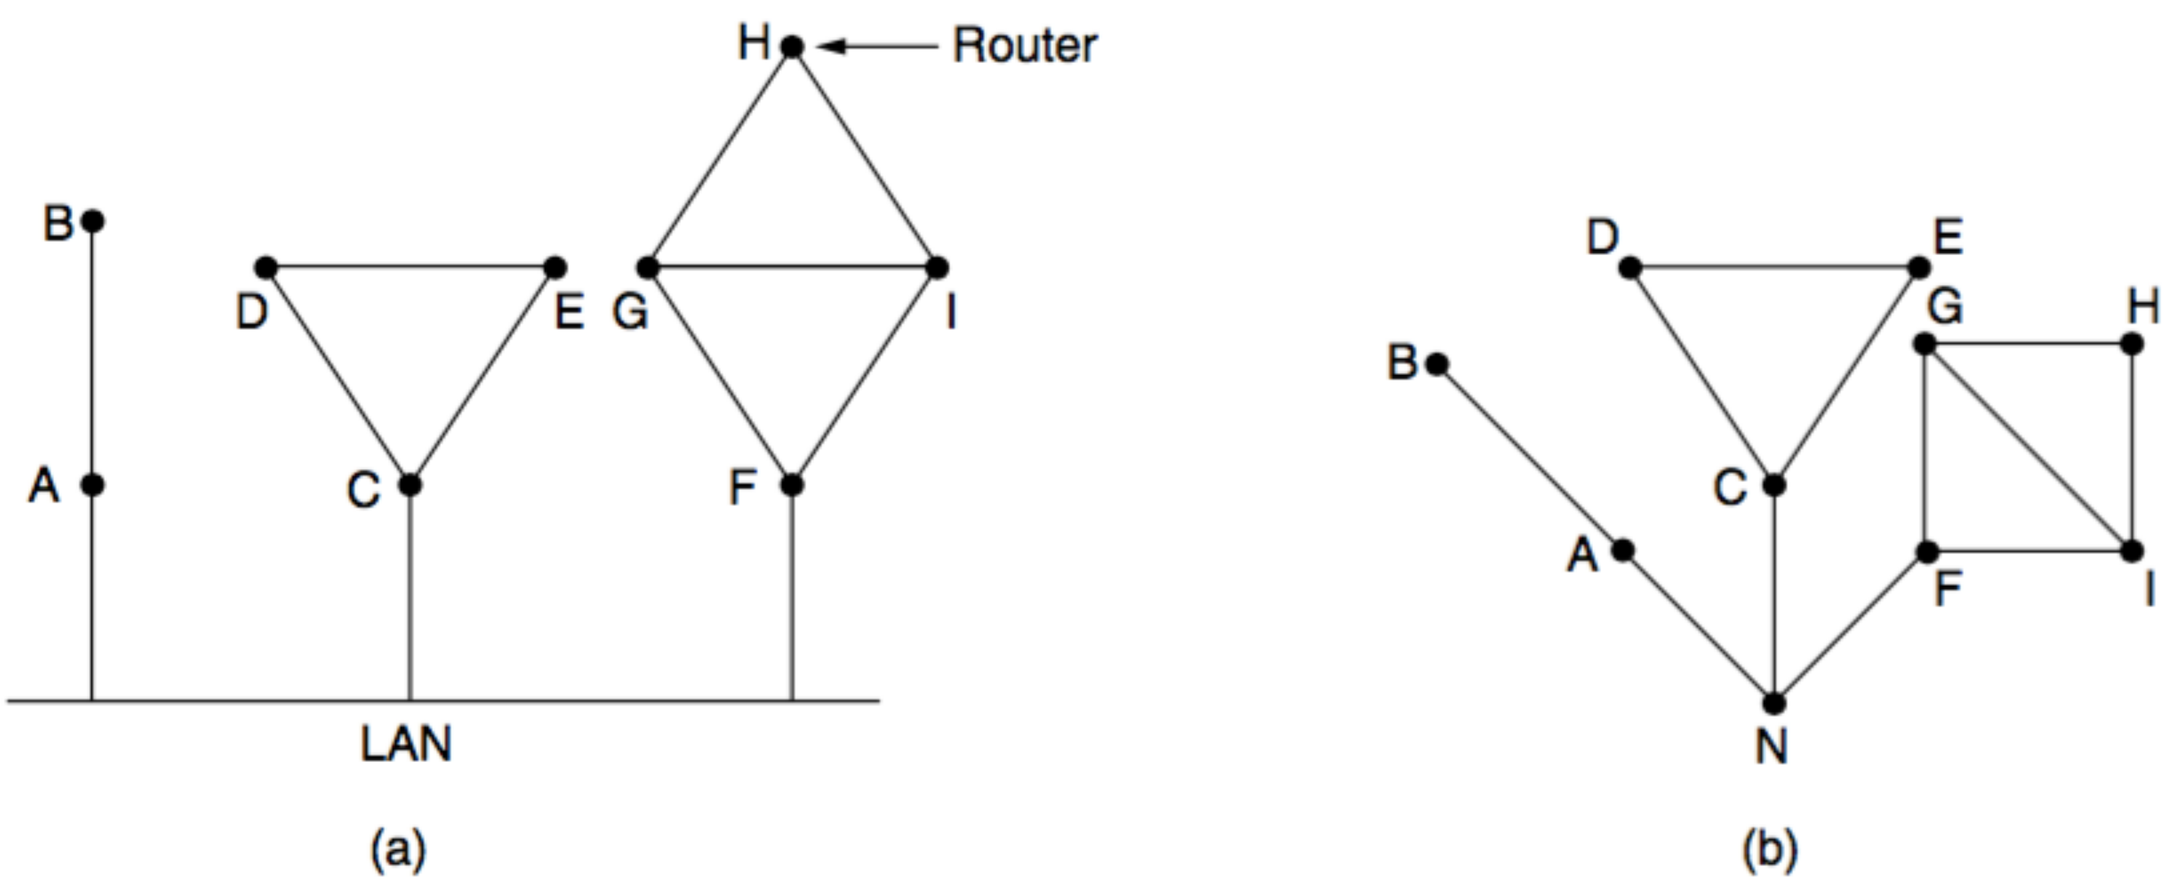

Link State Routing

<u>Setting Link Costs</u>

A common choice is to make the cost inversely proportional to the bandwidth of the link.

Delay of the links may be factored into the cost

<u>Building Link State Packets</u>

1. Identity of the sender
2. Sequence number
3. Age
4. List of neighbors

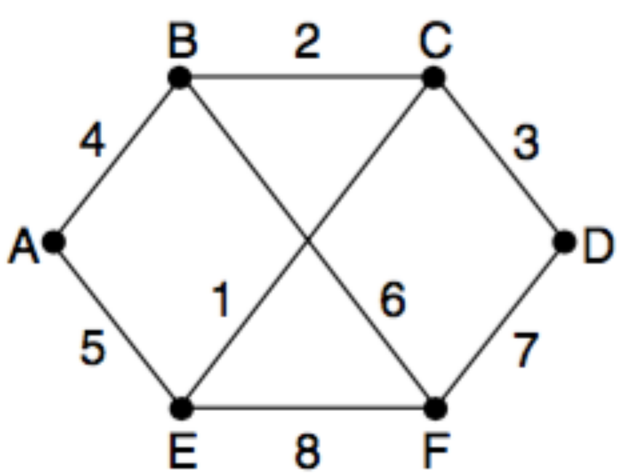

(a)

Link State Packets

A	B	C	D	E	F
Seq.	Seq.	Seq.	Seq.	Seq.	Seq.
Age	Age	Age	Age	Age	Age
B 4	A 4	B 2	C 3	A 5	B 6
E 5	C 2	D 3	F 7	C 1	D 7
	F 6	E 1		F 8	E 8

(b)

Link State Routing

<u>Distributing the Link State Packets</u>

1. Use flooding to distribute the link state packets to all routers
2. Each packet contains a 32-bit sequence number
3. Routers keep track of all the (source router, sequence) pairs they see
4. New link state packets are checked against the list of previous packets
5. If it is duplicate it is discarded
6. If the sequence number is old it is discarded
7. Include the age of each packet after the sequence number
8. Decrement the age once per second
9. When the age hits zero, the information from that router is discarded

Link State Routing

Distributing the Link State Packets

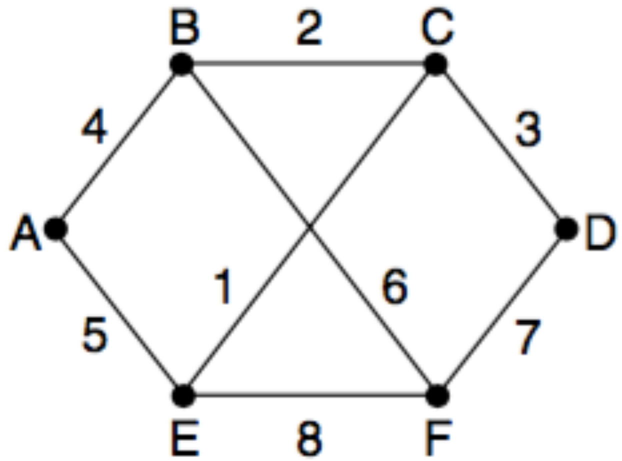

(a)

Link State Packets

A	B	C	D	E	F
Seq.	Seq.	Seq.	Seq.	Seq.	Seq.
Age	Age	Age	Age	Age	Age
B 4	A 4	B 2	C 3	A 5	B 6
E 5	C 2	D 3	F 7	C 1	D 7
	F 6	E 1		F 8	E 8

(b)

			Send flags			ACK flags			
Source	Seq.	Age	A	C	F	A	C	F	Data
A	21	60	0	1	1	1	0	0	
F	21	60	1	1	0	0	0	1	
E	21	59	0	1	0	1	0	1	
C	20	60	1	0	1	0	1	0	
D	21	59	1	0	0	0	1	1	

Data structure used by router B

Link State Routing

Computing the New Routes

1. Once a router has accumulated a full set of link state packets, it can construct the entire network graph
2. Every link is represented twice, once for each direction
3. Dijkstra's algorithm can be run locally to construct the shortest paths
4. For n routers, with k neighbors, the memory required is proportional to kn
5. In many practical situations, link state routing works well because it does not suffer from slow convergence problems

Link State Protocol

IS-IS (Intermediate State-Intermediate State)

OSPF (Open Shortest Path First)

Hierarchical Routing

1. Routers are divided into regions
2. Each router knows all the details about how to route packets within its own region
3. For huge networks it may be necessary to group the regions into clusters, the clusters into zones, the zones into groups
4. The optimal number of levels for an N router network is ln N

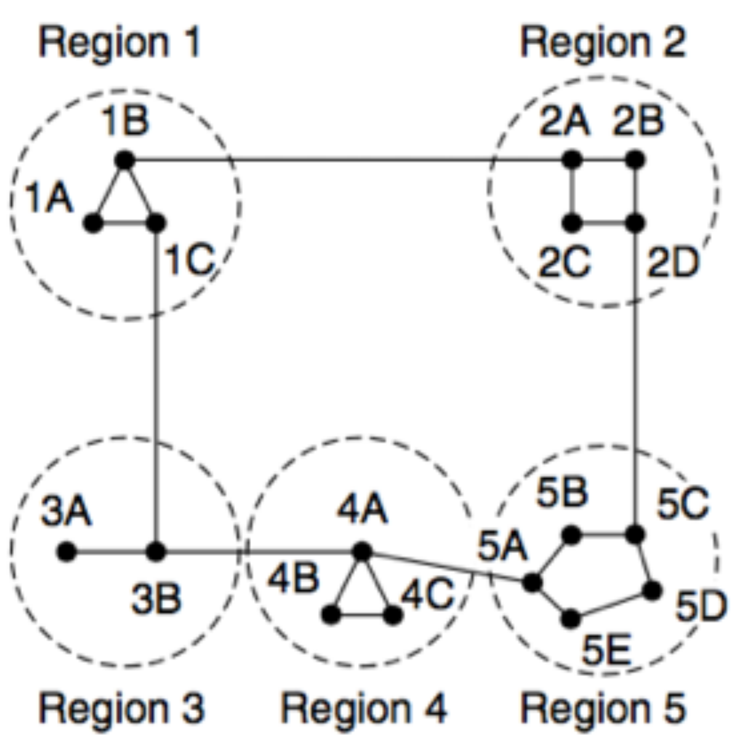

(a)

Full table for 1A

Dest.	Line	Hops
1A	–	–
1B	1B	1
1C	1C	1
2A	1B	2
2B	1B	3
2C	1B	3
2D	1B	4
3A	1C	3
3B	1C	2
4A	1C	3
4B	1C	4
4C	1C	4
5A	1C	4
5B	1C	5
5C	1B	5
5D	1C	6
5E	1C	5

(b)

Hierarchical table for 1A

Dest.	Line	Hops
1A	–	–
1B	1B	1
1C	1C	1
2	1B	2
3	1C	2
4	1C	3
5	1C	4

(c)

Broadcast Routing

Broadcasting
Send packets to all destinations in the network

Multidestination Routing
Network bandwidth is used more efficiently

Flooding
Uses links efficiently with a decision rule that is relatively simple

Reverse Path Forwarding
Only forwards if packet comes from the shortest path

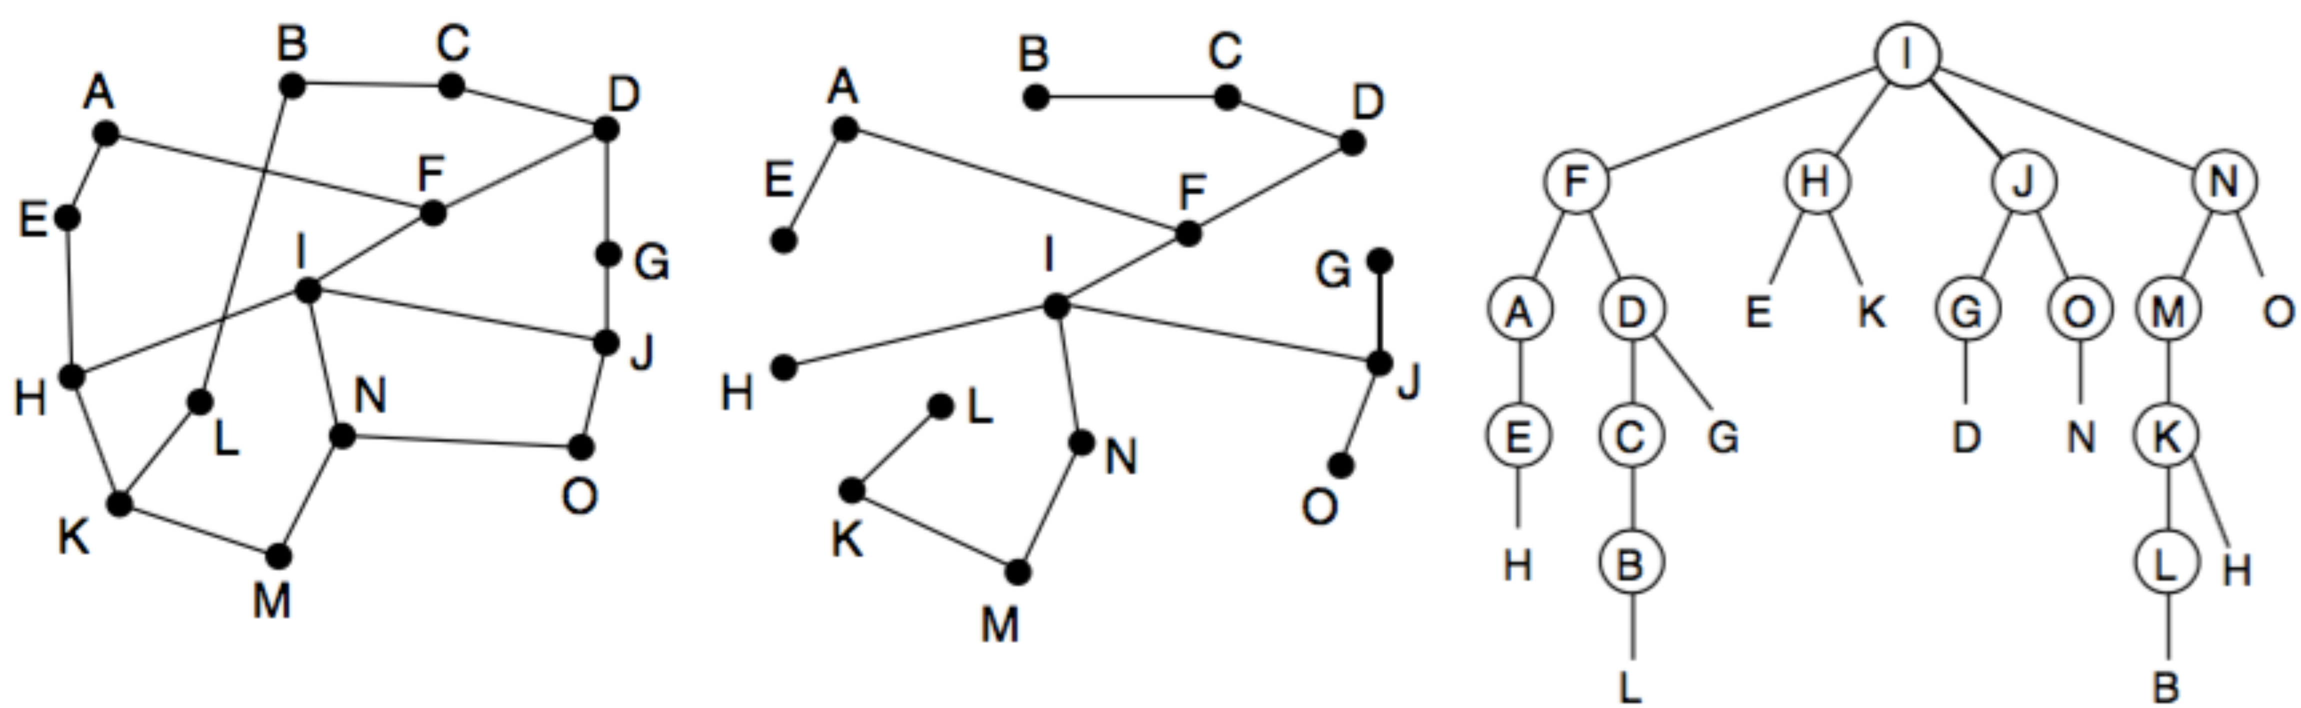

(a) (b) (c)

Multicast Routing

Prune the spanning tree into multicast groups

Link state routing: MOSPF (Multicast OSPF)

Distance vector routing: DVMRP (Distance Vector Multicast Routing Protocol)

Network

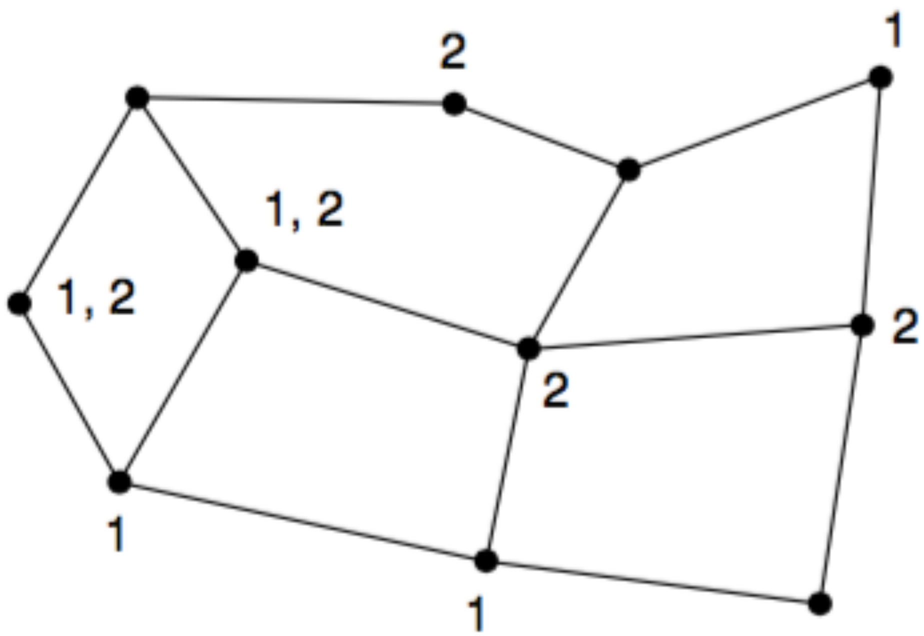

(a)

Spanning Tree

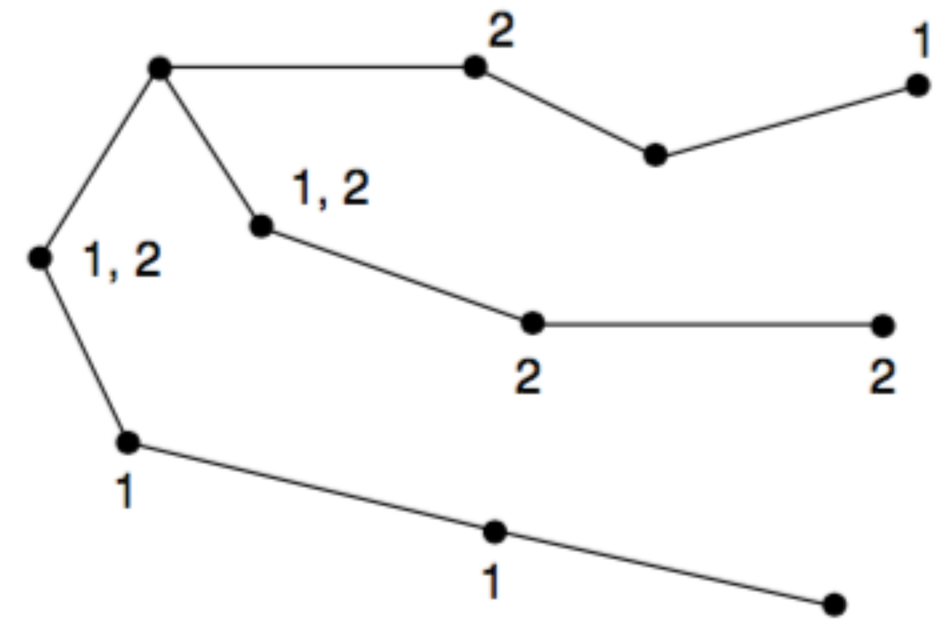

(b)

Broadcast Tree for Group 1

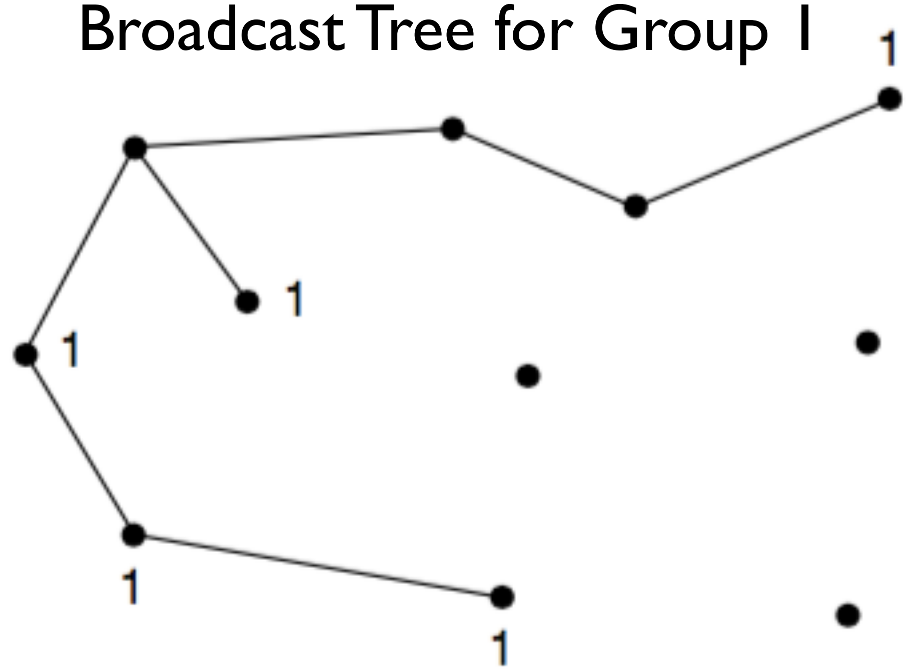

(c)

Broadcast Tree for Group 2

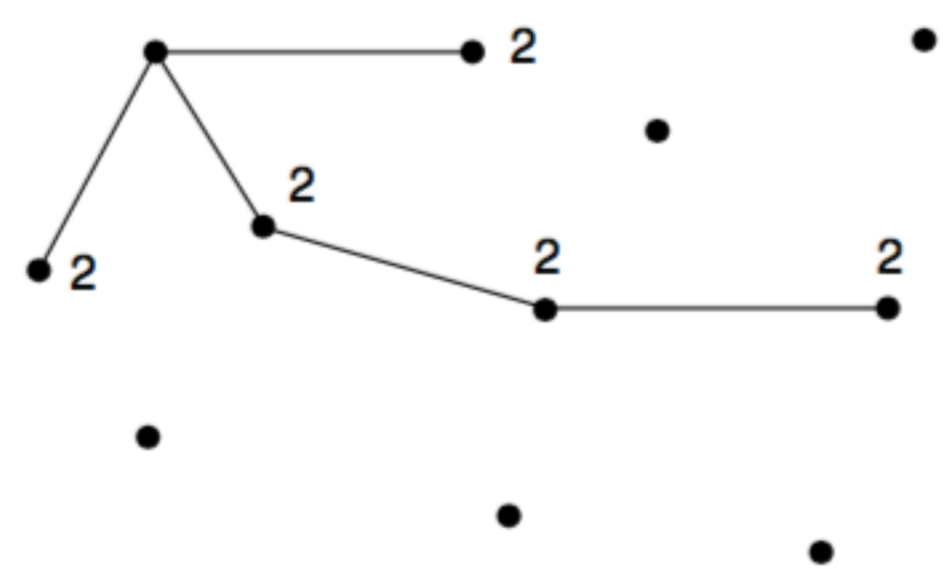

(d)

Core-based Trees

If the network has n groups, each with an average of m nodes
Broadcast Tree: Stores mn trees
Core-based Tree: Stores n trees

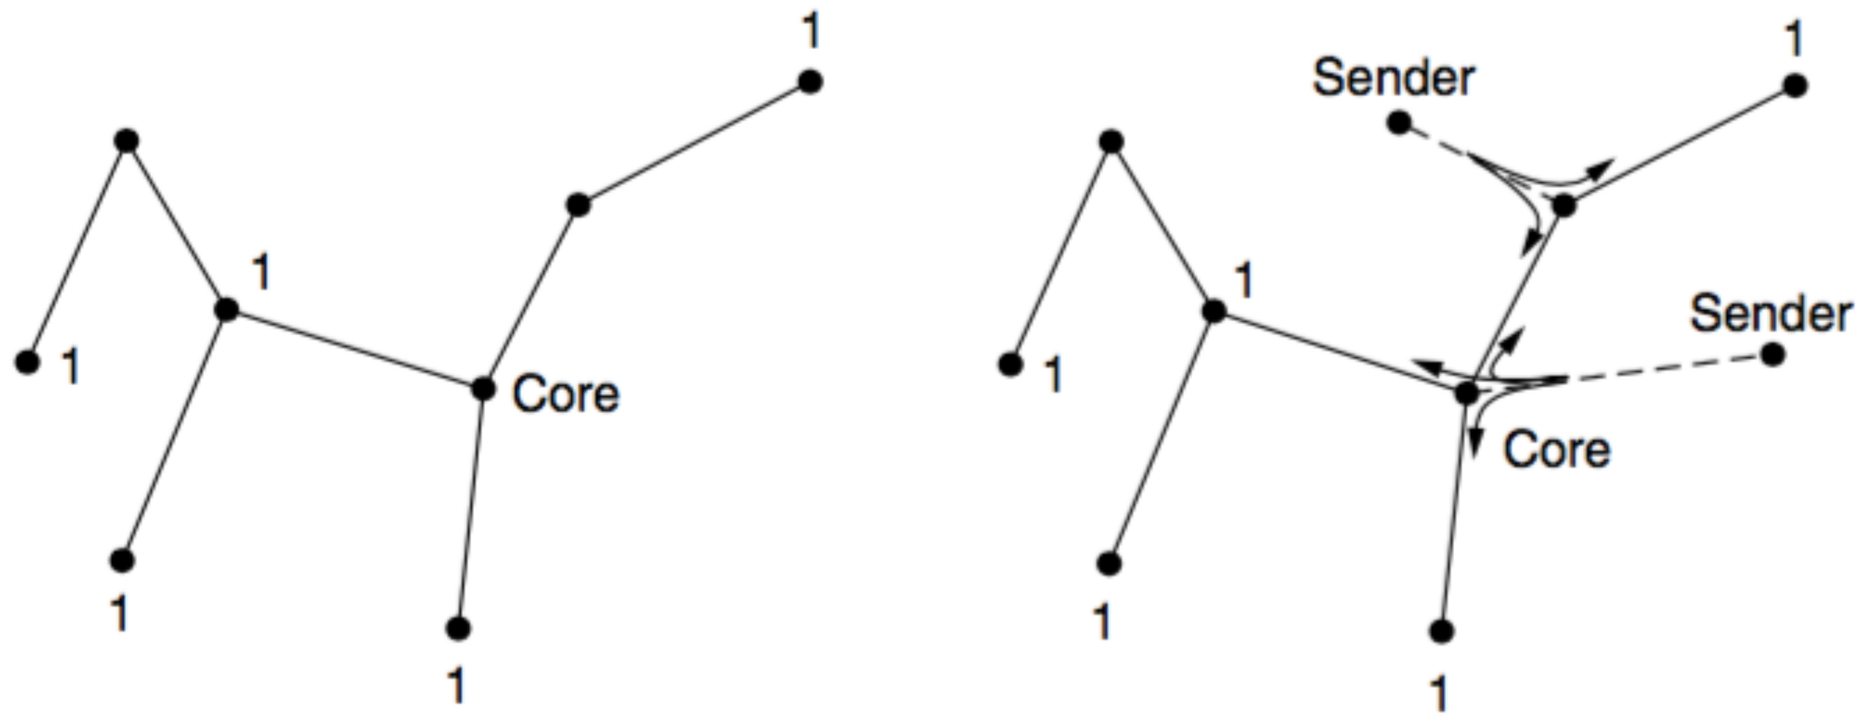

(a) (b)

Core-based tree protocol: PIM (Protocol Independent Multicast)

Anycast Routing

Broadcast: All destinations
Unicast: Single destination
Multicast: Multiple destinations
Anycast: Nearest member of the group

Anycast routing for Group 1 | Topology seen by the routing protocol

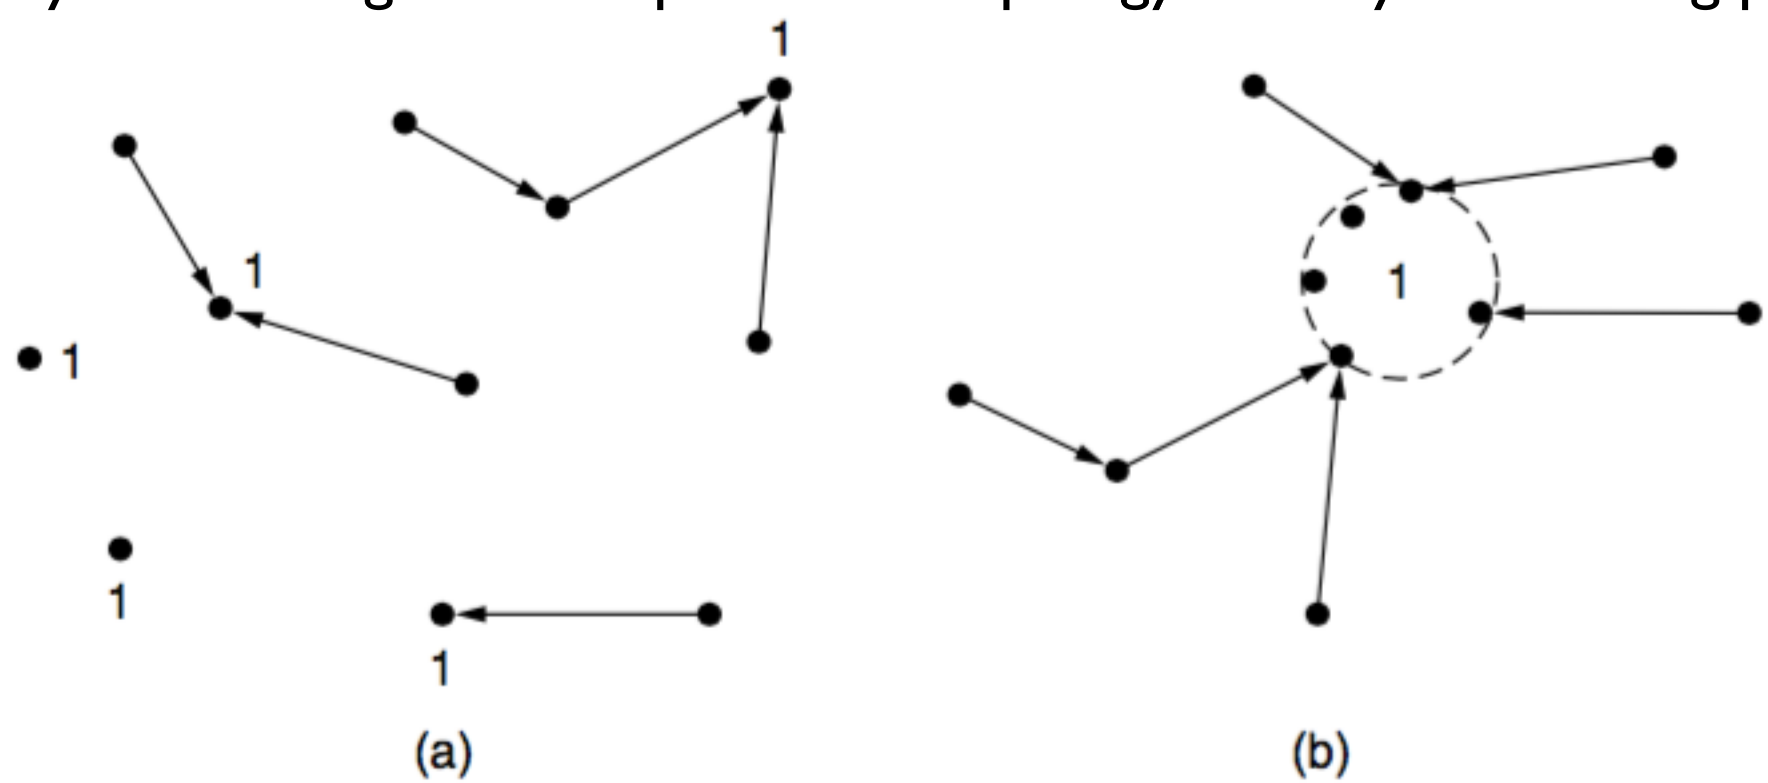

Routing for Mobile Hosts

We should not compute new routes every time a mobile device moves

Mobile phone number: 1-212-5551212
United States (country code 1) and Manhattan (212)

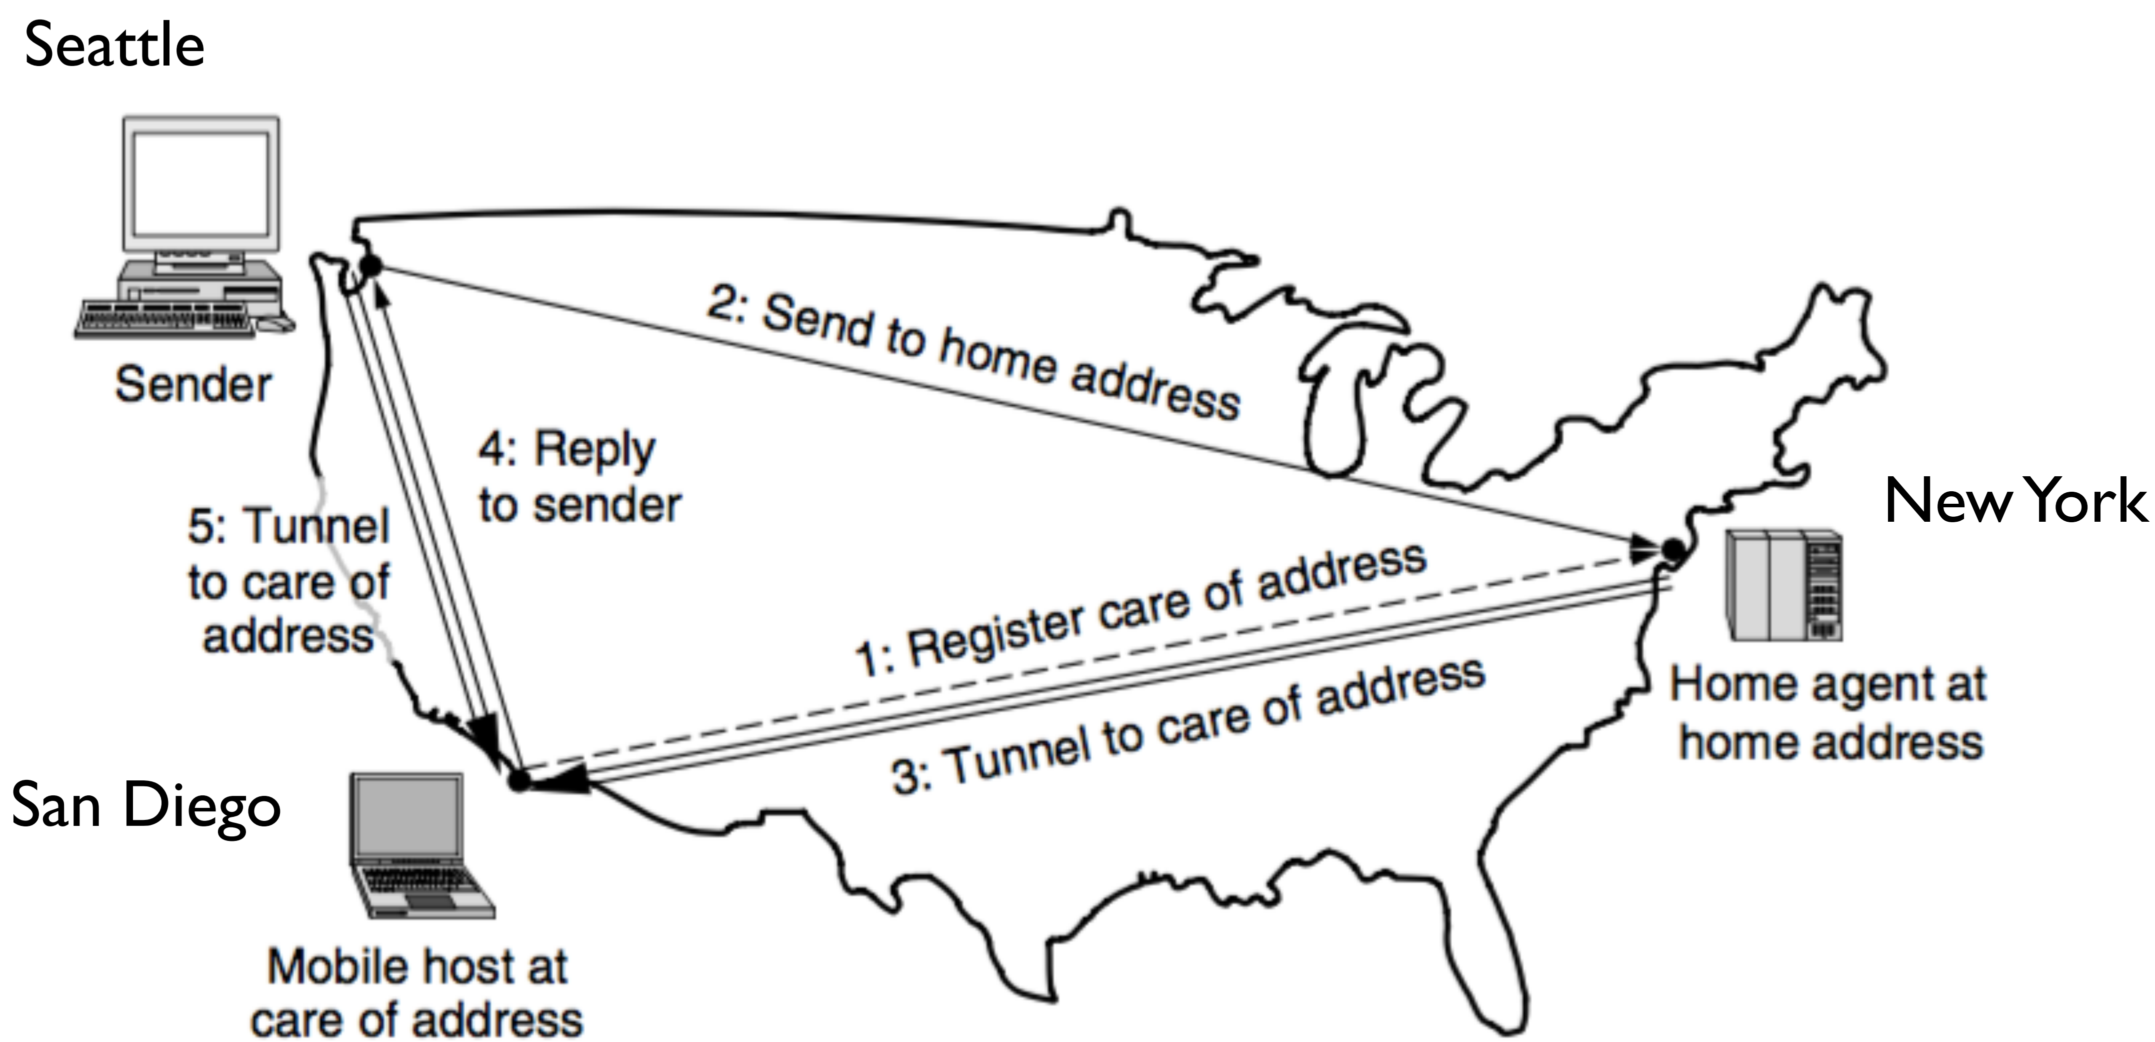

Routing in Ad Hoc Networks

Mobile Network: Routers are stationary

MANET (Mobile Ad hoc NETworks): Routers also move

<u>Routing Protocols</u>

AODV (Ad hoc On-demand Distance Vector)

- Routes to a destination are discovered on demand
- ROUTE REQUEST packet is broadcast using flooding
- ROUTE REPLY packet is unicast to the sender along the reverse of the path

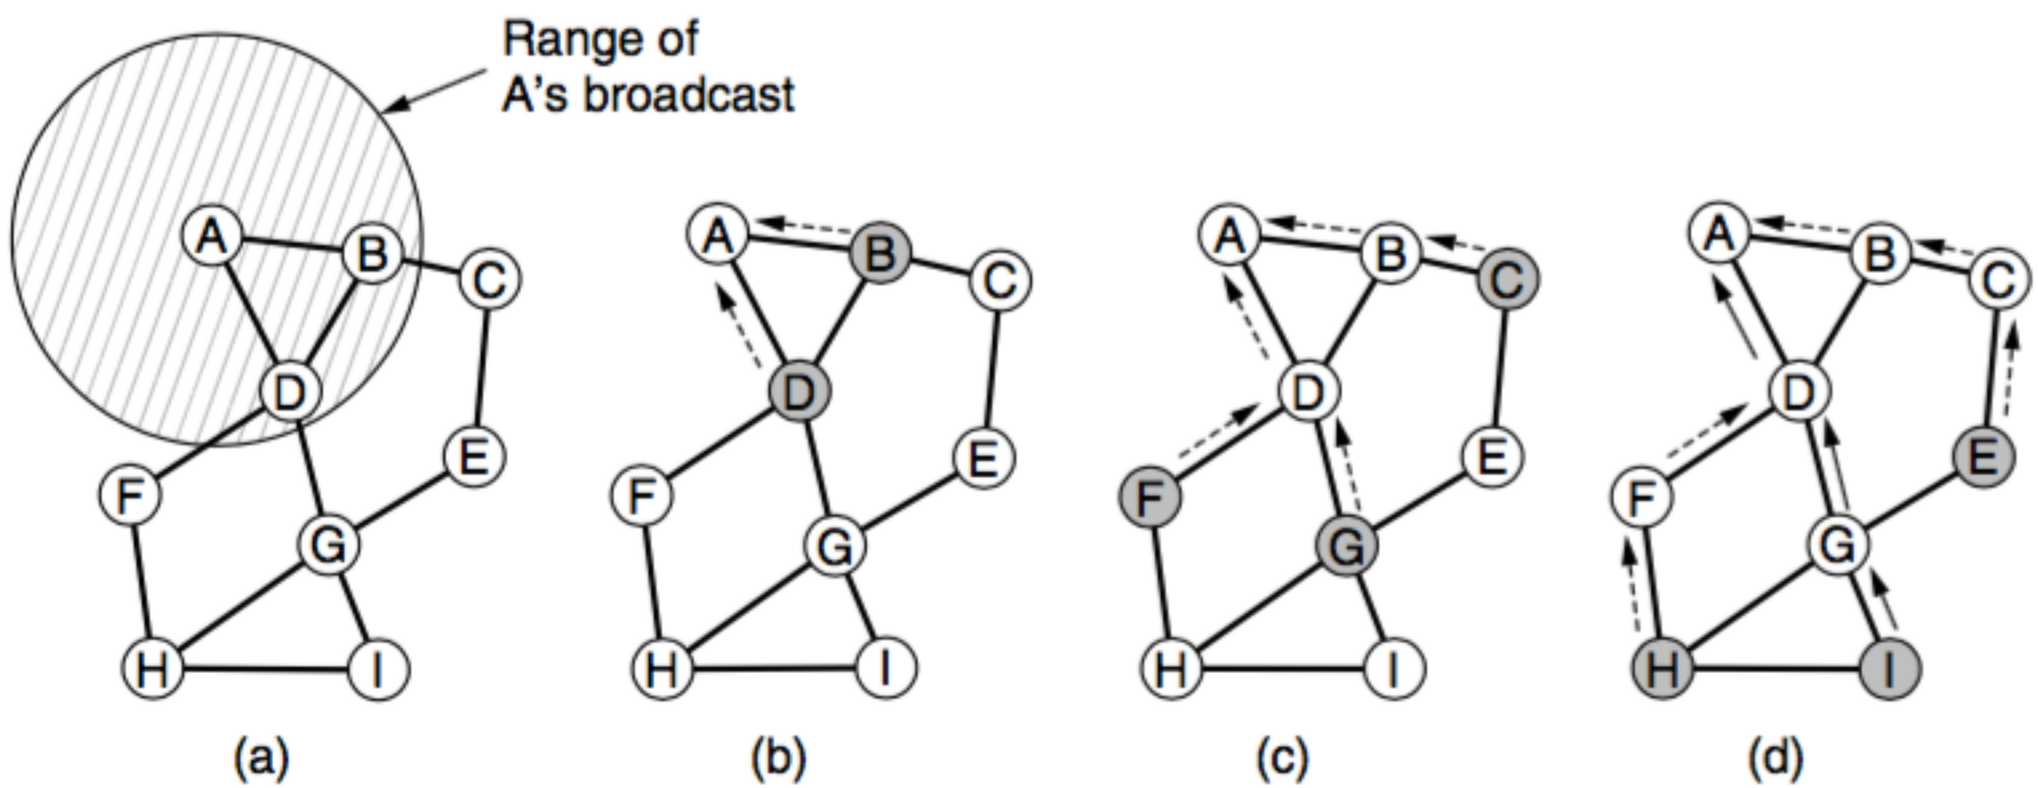

(a) Range of *A*'s broadcast. (b) After *B* and *D* receive it. (c) After *C*, *F*, and *G* receive it. (d) After *E*, *H*, and *I* receive it. The shaded nodes are new recipients. The dashed lines show possible reverse routes. The solid lines show the discovered route.

Congestion Control Algorithms

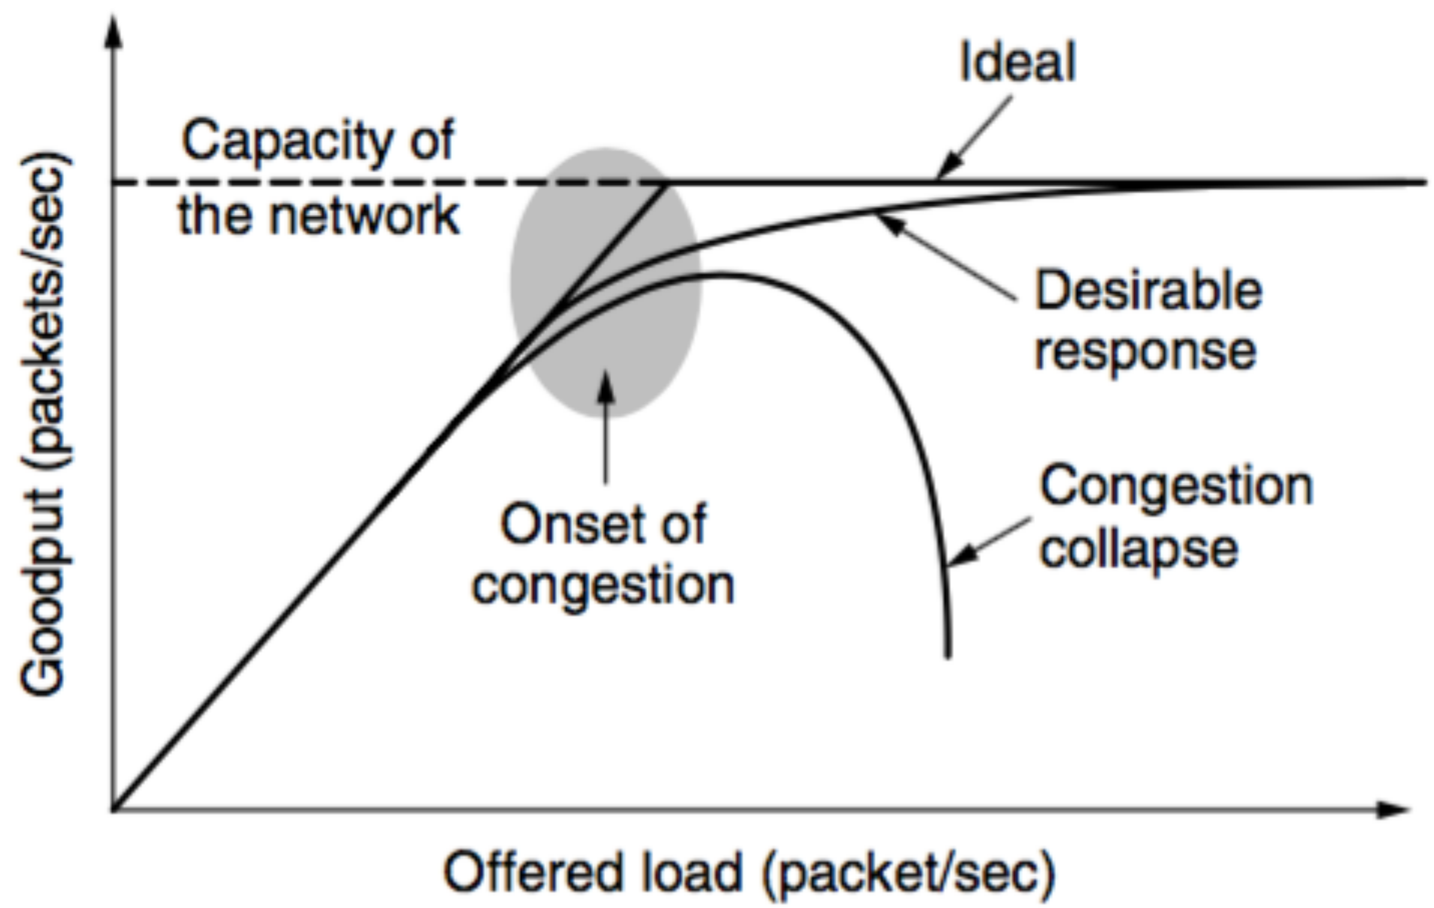

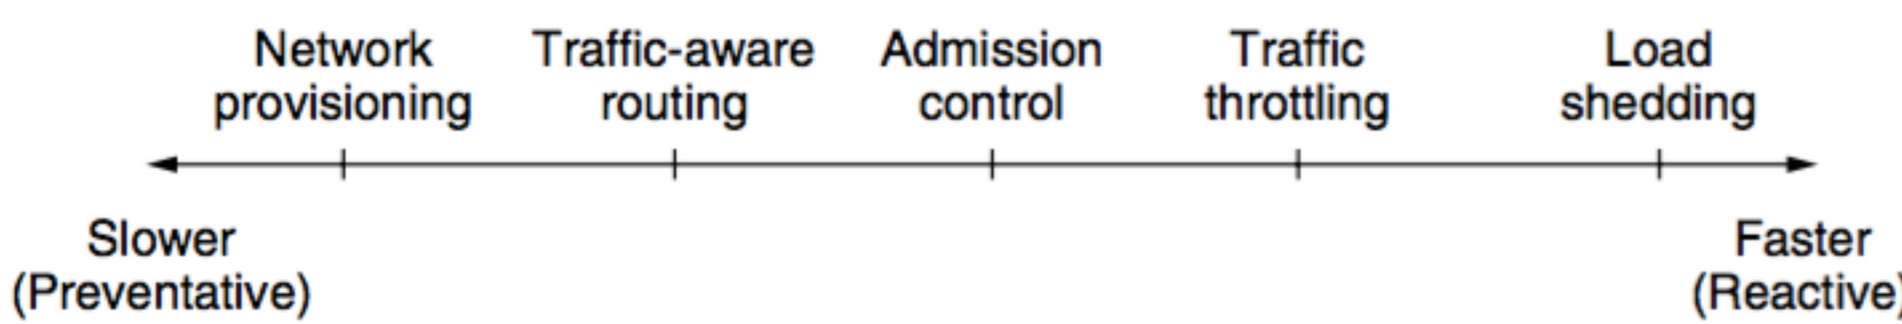

Timescale of approaches to congestion control

Traffic-aware Routing

Set the link weight to be a function of the (fixed) link bandwidth and propagation delay plus the (variable) measured load or average queuing delay

Internet routing protocols do not generally adjust their routes depending on the load.

Example of oscillating routing tables

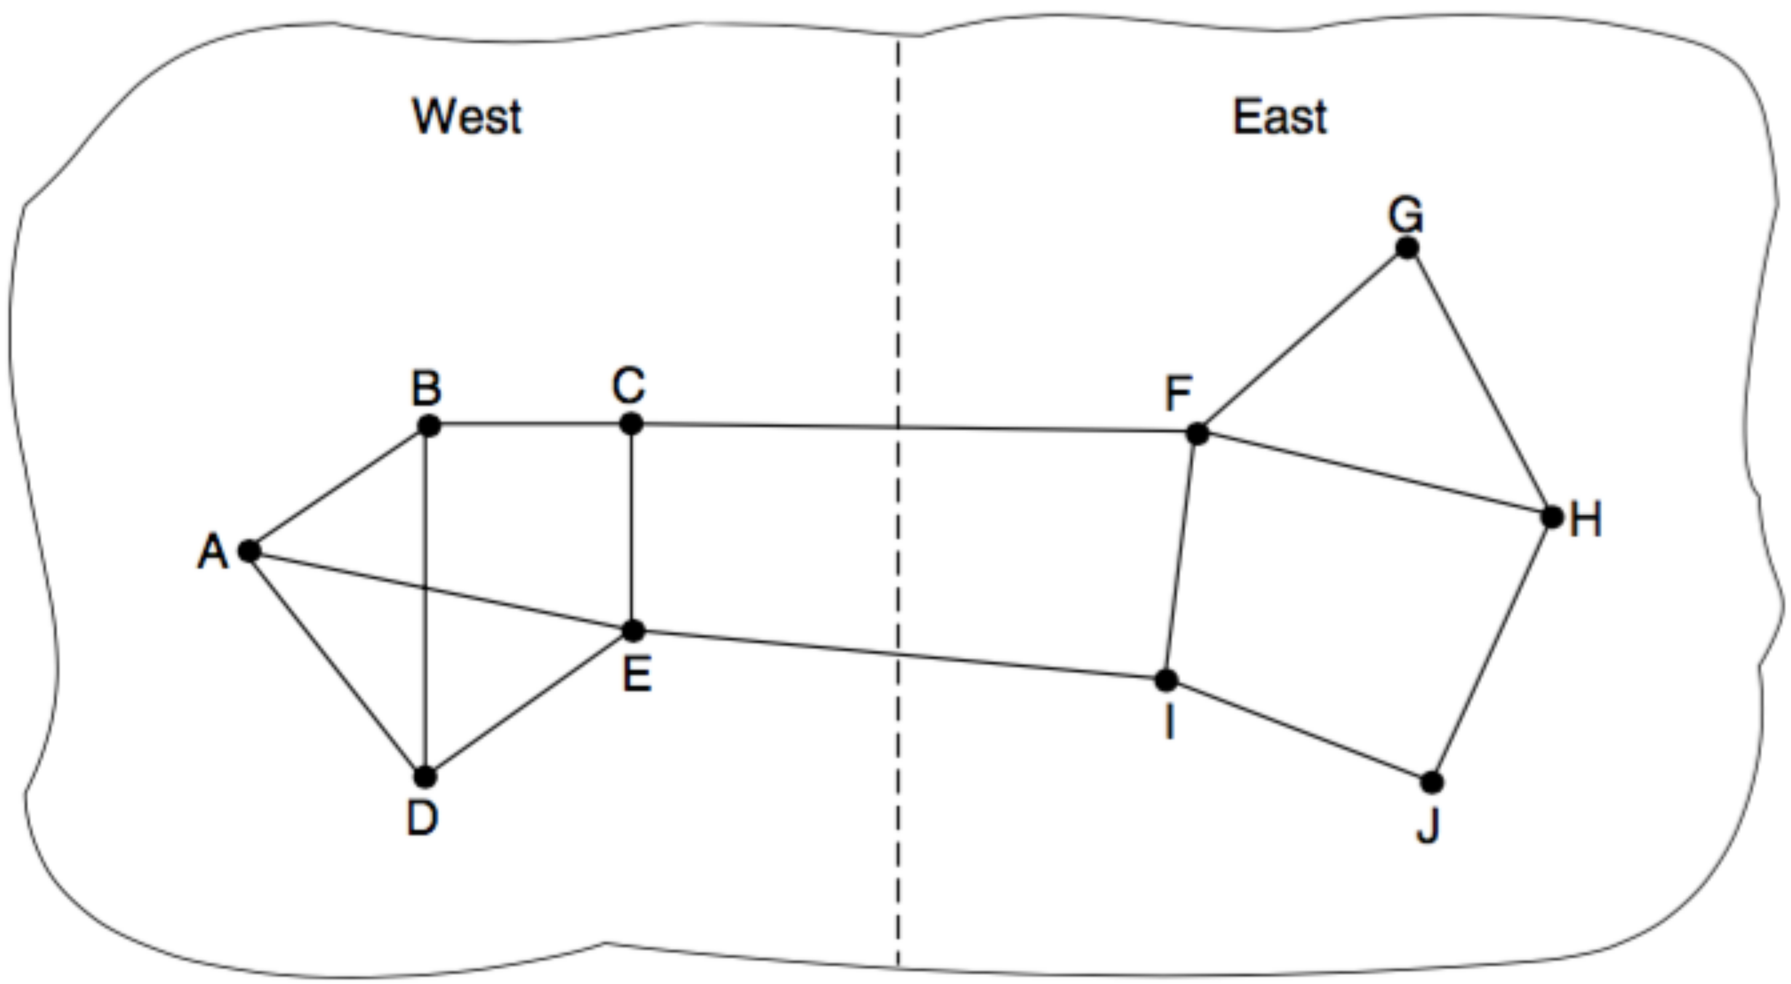

Admission Control

Do not set up a new virtual circuit unless the network can carry the added traffic without becoming congested.

Measurements of past behavior that capture the statistics of transmissions can be used to estimate the number of virtual circuits to admit.

Redraw the network without the congested routers.

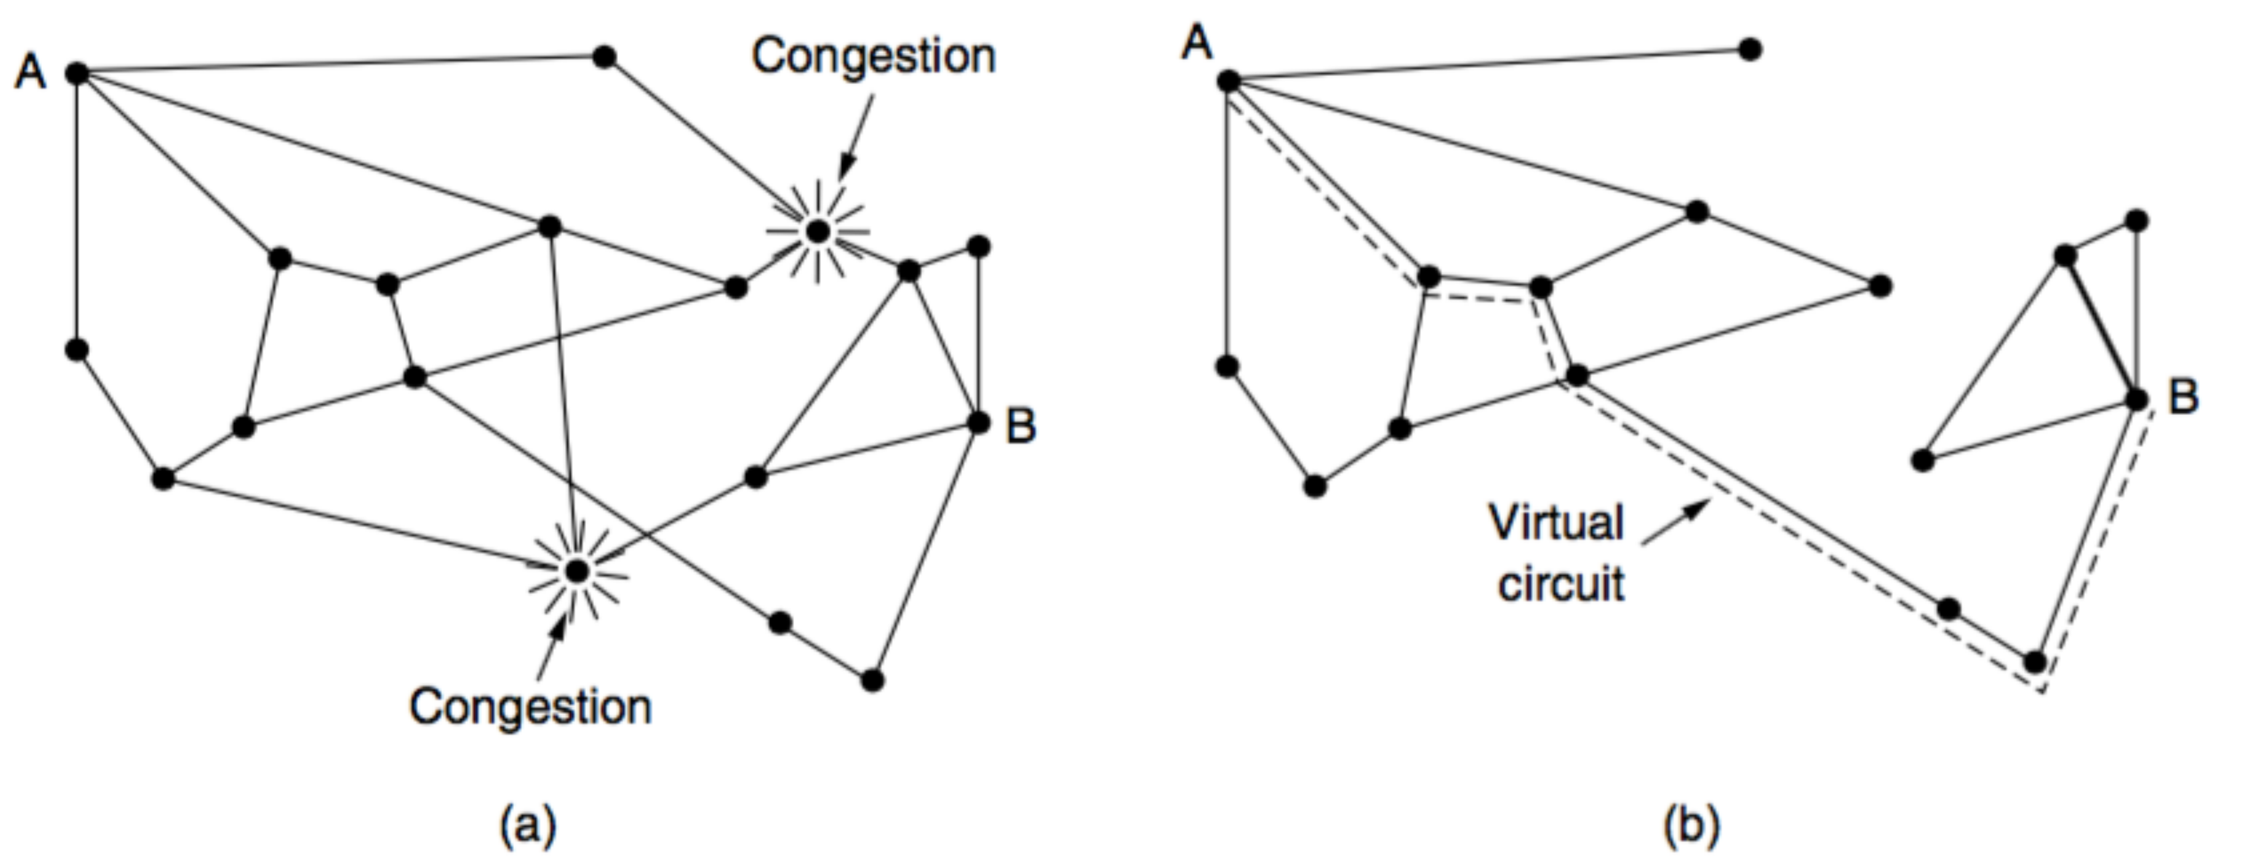

Traffic Throttling

Predict congestion by monitoring resources

- Utilization of the output links
- Buffering of queued packets in the router
- Number of packets that are lost due to insufficient buffering

EWMA (Exponentially Weighted Moving Average)

$$d_{\text{new}} = \alpha d_{\text{old}} + (1 - \alpha)s$$

Instantaneous queue length: s

How fast the router forgets history: α

Queueing delay: d

Traffic Throttling

Once the congestion is detected, routers must notify the senders

Choke packets

1. The router selects a congested packet and sends a choke packet back to the source host
2. When the source host gets the choke packet, it is required to reduce the traffic sent to the specified destination

Explicit Congestion Notification

1. Tag packets to indicate congestion
2. The destination can note that there is congestion and inform the sender when it sends a reply packet

Two bits in the IP packet header are used to record whether the packet has experienced congestion

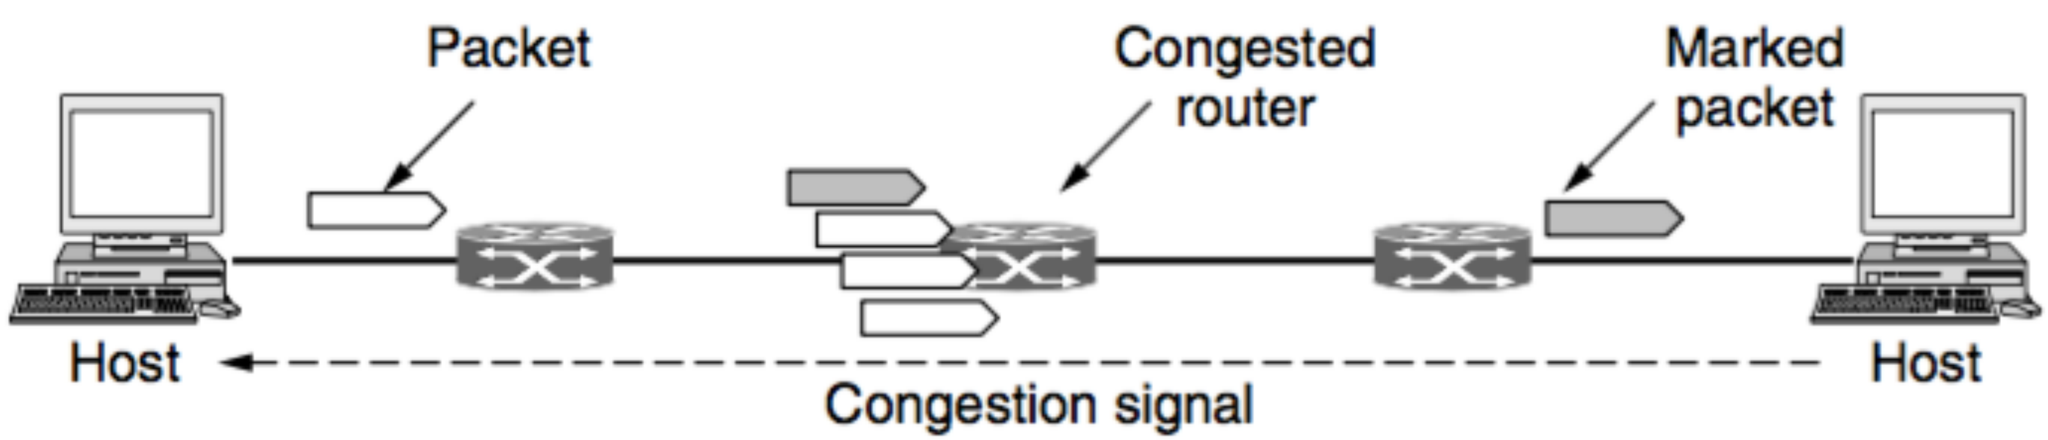

Traffic Throttling

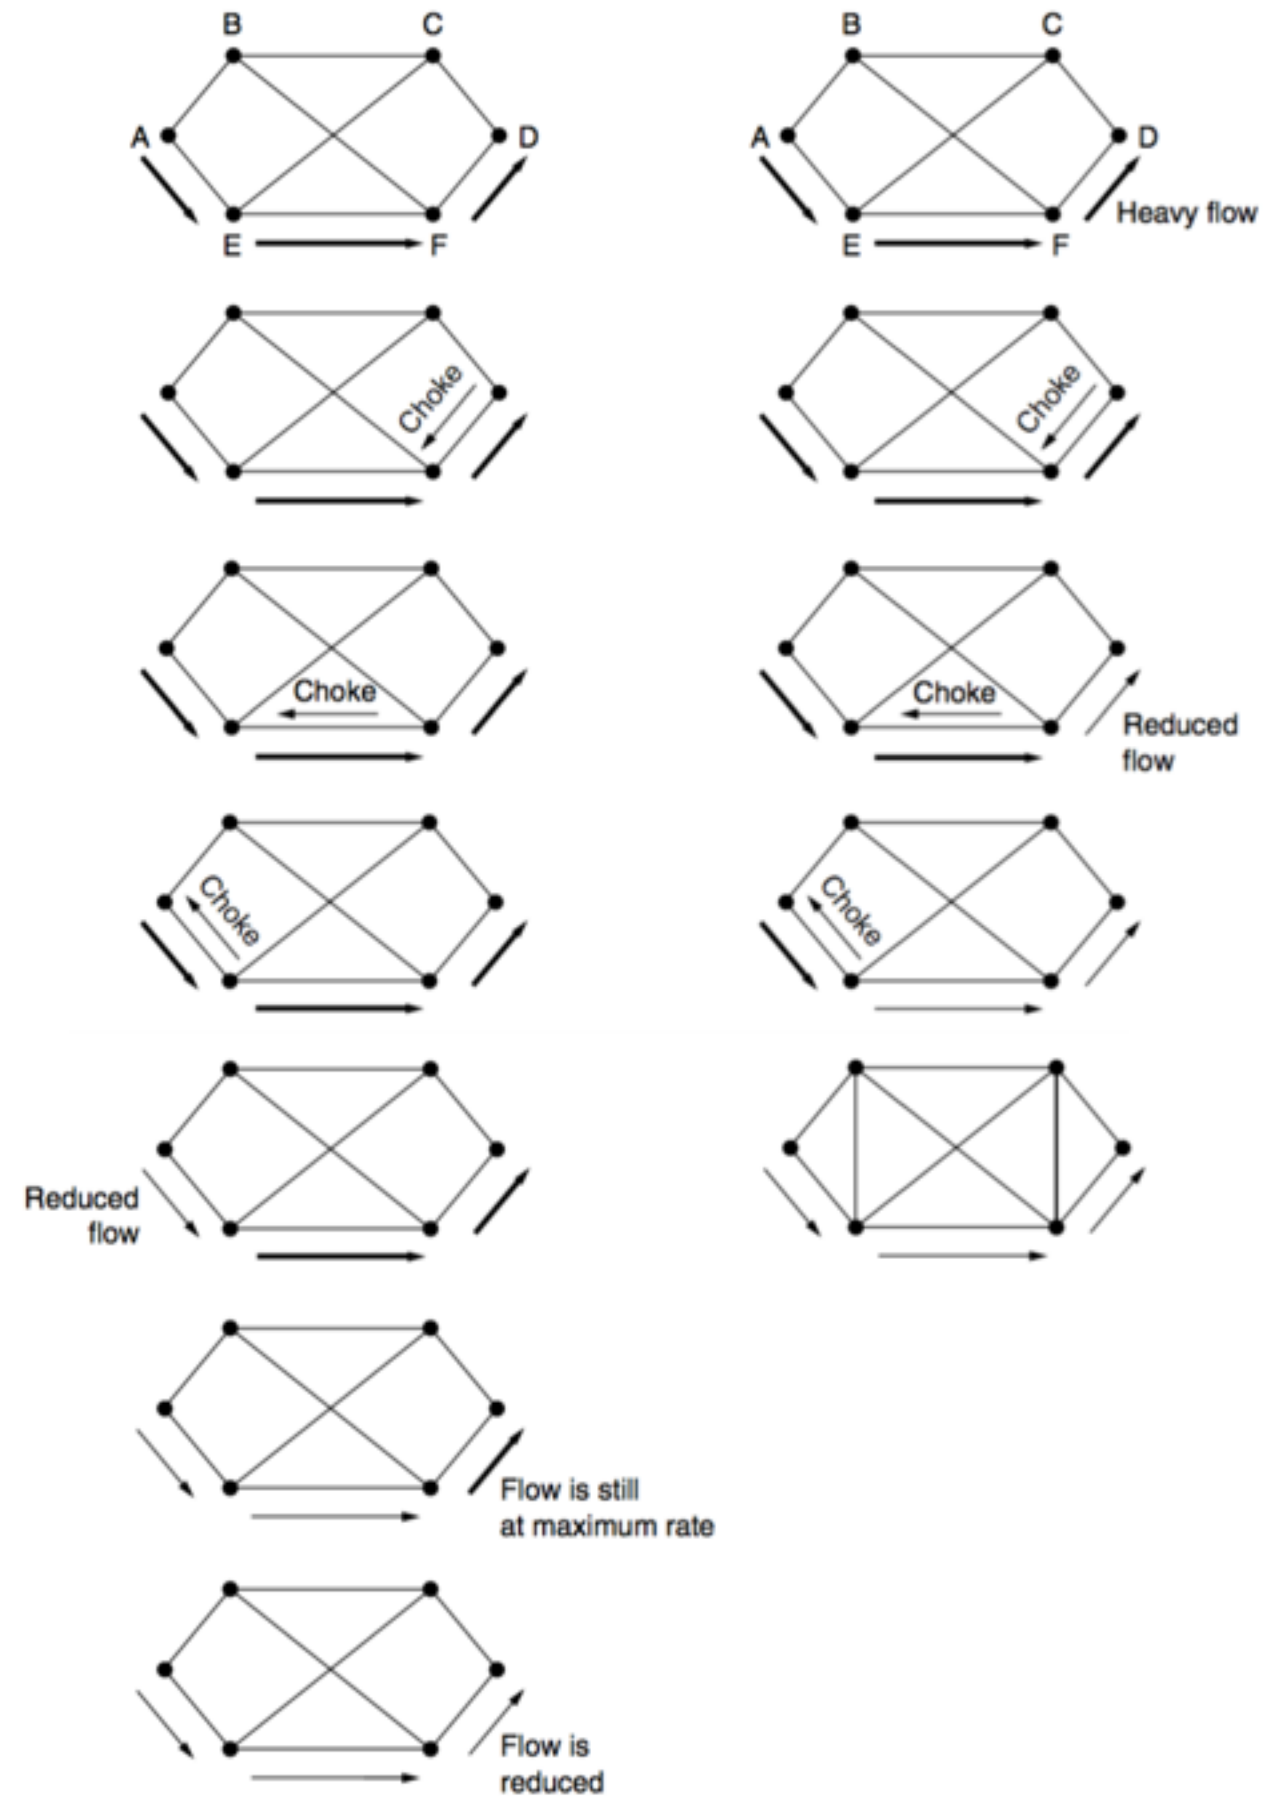

Hop-by-Hop Backpressure

Have the choke packet take effect at every hop it passes through

Load Shedding

Routers that are congested throw packets away

File transfer: old data is more valuable

Real time: new data is more valuable

Random Early Detection

1. Routers maintain a running average of their queue lengths
2. When the average queue length on exceeds a threshold, a small fraction of the packets are dropped at random
3. The lost packet is delivering the same message as a choke packet, but implicitly

Quality of Service

Four Issues to Ensure Quality of Service

1. What applications need from the network.
2. How to regulate the traffic that enters the network.
3. How to reserve resources at routers to guarantee performance.
4. Whether the network can safely accept more traffic.

Application Requirements

Flow of packets can be characterized by bandwidth, delay, jitter, and loss.

Application	Bandwidth	Delay	Jitter	Loss
Email	Low	Low	Low	Medium
File sharing	High	Low	Low	Medium
Web access	Medium	Medium	Low	Medium
Remote login	Low	Medium	Medium	Medium
Audio on demand	Low	Low	High	Low
Video on demand	High	Low	High	Low
Telephony	Low	High	High	Low
Videoconferencing	High	High	High	Low

Quality of Service

Variety of Quality of Service in Networks

1. Constant bit rate (e.g., telephony).
2. Real-time variable bit rate (e.g., compressed videoconferencing).
3. Non-real-time variable bit rate (e.g., watching a movie on demand).
4. Available bit rate (e.g., file transfer).

Traffic Shaping

A technique for regulating the average rate and burstiness of a flow of data that enters the network.
Customer and provider have a SLA (Service Level Agreement).
Packets in excess of the agreed pattern might be dropped by the network.

Leaky and Token Buckets

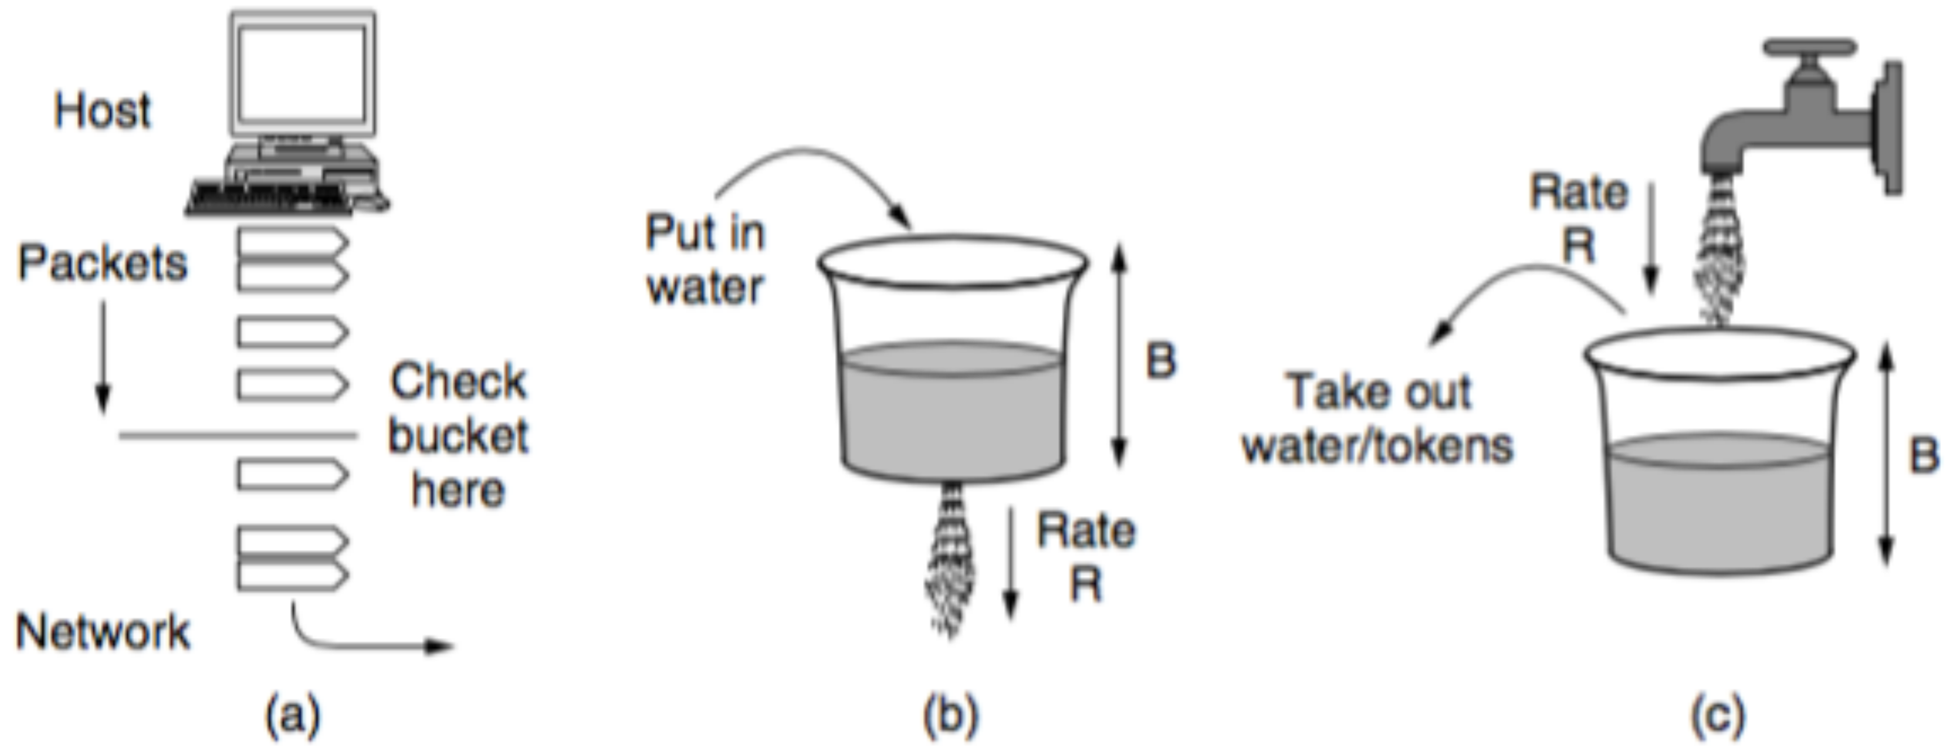

Leaky Bucket   Token Bucket

Leaky Bucket

If a packet arrives when the bucket is full:

1. It is queued until enough water leaks out (OS is shaping the traffic)
2. Hold it to be discarded (Provider network interface that is policing traffic)

Token Bucket

No more than a fixed number of tokens, B, can accumulate in the bucket, and if the bucket is empty, we must wait until more tokens arrive before we can send another packet.

Leaky and Token Buckets

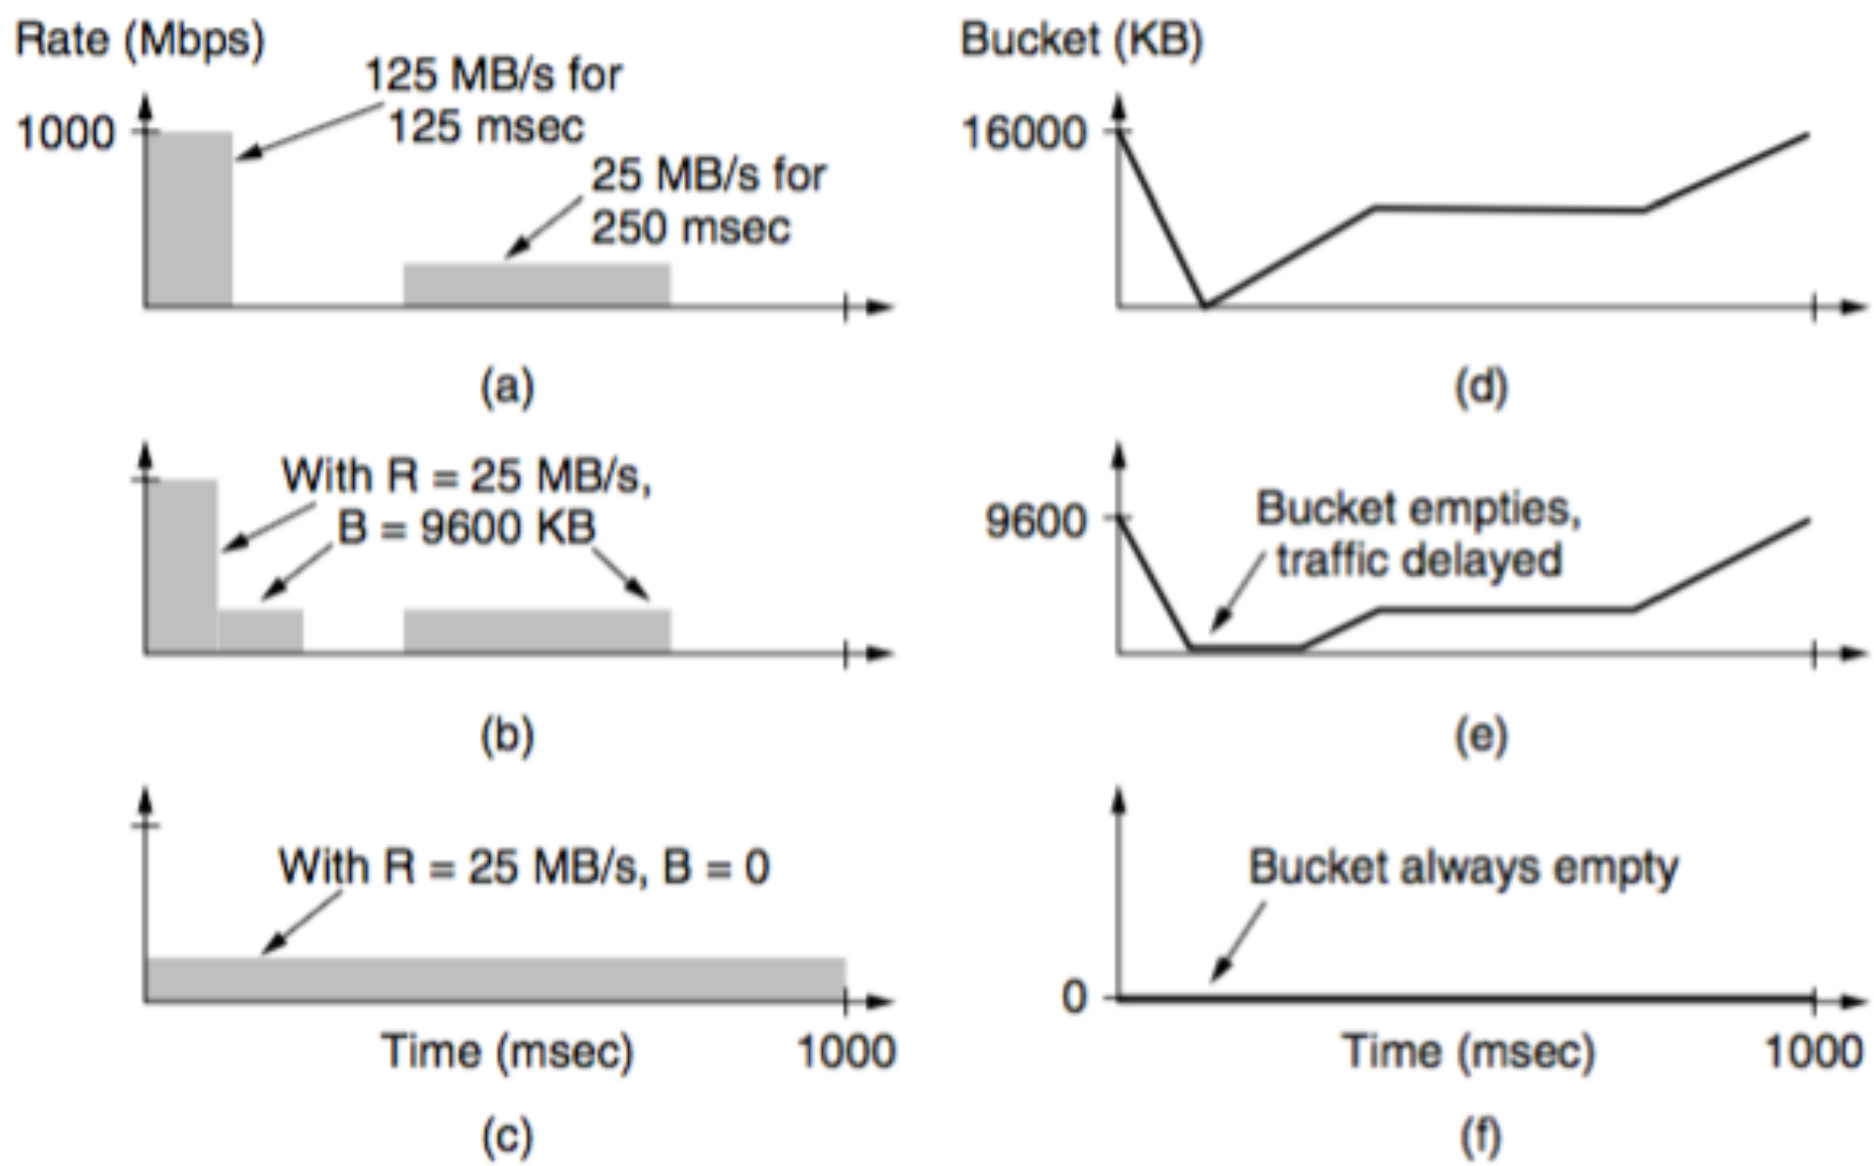

Burst length: S [sec]
Maximum output rate: M [bytes/sec]
Token bucket capacity B [bytes]
Token arrival rate: R [bytes/sec]

$$S = B/(M - R).$$

B = 9600 KB
M = 125 MB/sec
R = 25 MB/sec

S = 94 msec

講義日程（2Q）

		授業計画		課題
06/14	第９回	ネットワーク層１ ルーティング・輻輳制御	5章	ルーティングの種類を理解し 輻輳制御手法を説明できる
06/21	第１０回	ネットワーク層２ インターネットとサービス品質	5章	インターネットの制御プロトコルを理解し ネットワーク間の接続について説明できる
06/28	第１１回	トランスポート層１ トランスポート・プロトコルの要素	6章	誤り制御とフロー制御を理解し 輻輳制御について説明できる
07/05	第１２回	トランスポート層２ UDP と TCP	6章	TCP の信頼性を理解し TCP のコネクション管理を説明できる
07/12	第１３回	アプリケーション層 DNS, 電子メール, www	7章	DNS, 電子メール, www のしくみを理解し ストリーミング，P2P について説明できる
07/26	第１４回	ネットワークセキュリティ１ 対称鍵暗号，公開鍵暗号	8章	暗号アルゴリズムを理解し SHA-1,2 と RSA について説明できる
08/02	第１５回	ネットワークセキュリティ２ デジタル署名，認証プロトコル	8章	電子メール,Web のセキュリティ の脅威について把握できる