

# The Network Layer

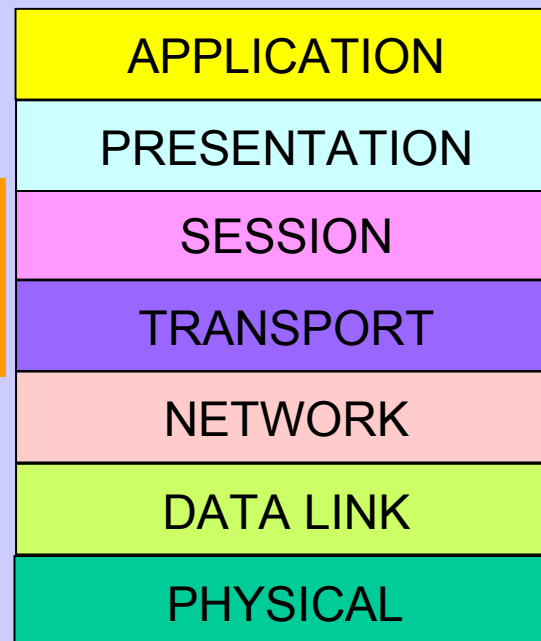
# review

## ISO/OSI's network model

How many layers have the OSI's model divided the network architecture into?

Seven layers

What are they from the bottom to the top?



# Description of the network layer

- a) **The network layer is concerned with getting p-packets from the source all the way to the destination.**
- **To achieve its goals, the network layer must know about the topology of the communication subnet and choose appropriate paths through it. It must also take care to choose routes to avoid overloading some of the communication lines and routers while leaving others idle.**

# Chapter 5 The Network Layer

## 5.1 Network Layer Design Issues


## 5.2 Routing Algorithms

## 5.6 The Network Layer in the Internet

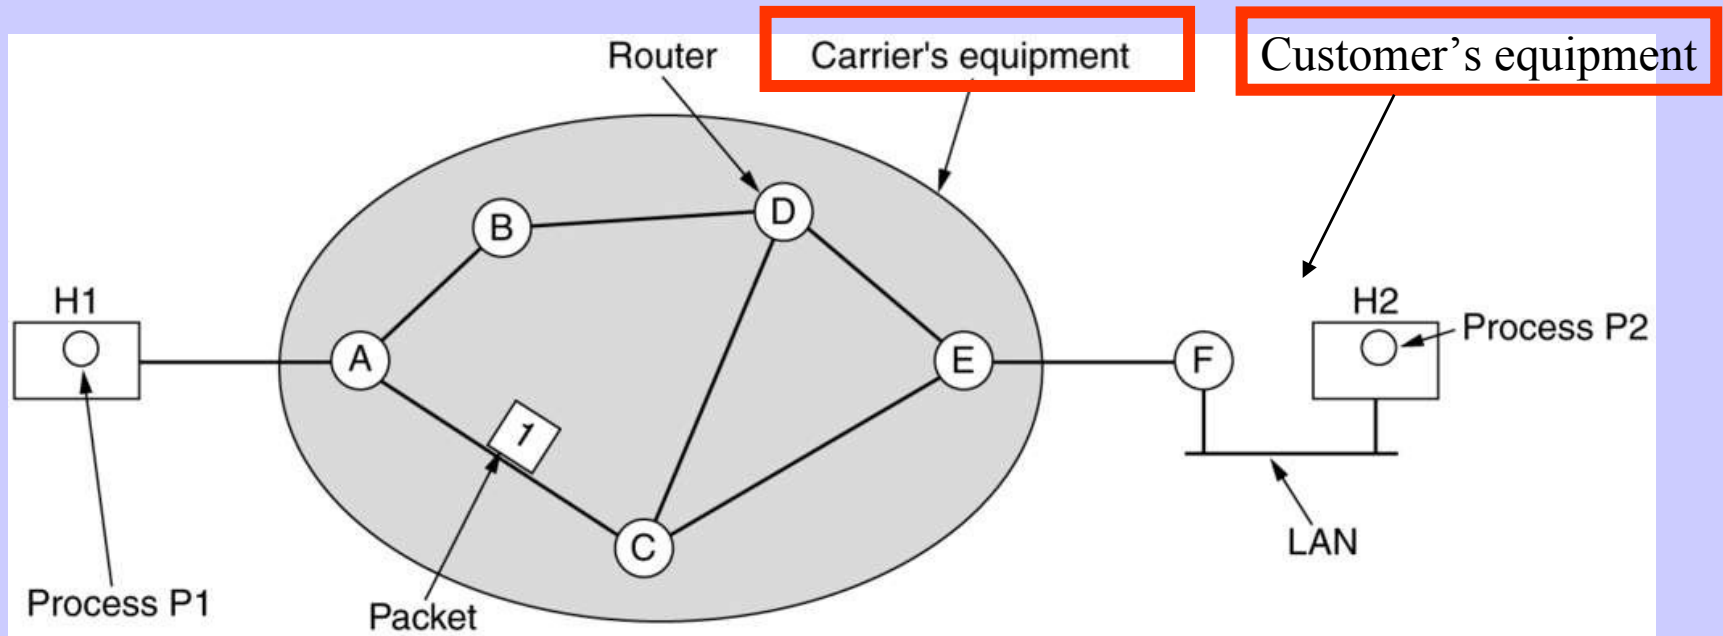


# 5.1 Network Layer Design Issues

- a) **Store-and-Forward Packet Switching**
- b) Services Provided to the Transport Layer
- c) Implementation of Connectionless Service
- d) Implementation of Connection-Oriented Service
- e) Comparison of Virtual-Circuit and Datagram Subnets




# Store-and-Forward Packet Switching



The environment of the network layer protocols.


- This equipment is used as follows:

a) A host with a packet to send transmits it to the nearest router, either on its own LAN or over a point-to-point link to the carrier. The packet is stored there until it has fully arrived so the checksum can be verified. Then it is forwarded to the next router along the path until it reaches the destination host, where it is delivered. This mechanism is **store-and-forward** packet switching, as we have seen in previous chapters.



# 5.1 Network Layer Design Issues

- a) Store-and-Forward Packet Switching
- b) **Services Provided to the Transport Layer**
- c) Implementation of Connectionless Service
- d) Implementation of Connection-Oriented Service
- e) Comparison of Virtual-Circuit and Datagram Subnets





# Services Provided to the Transport Layer



**What kind of services the network layer provides to the transport layer ?**


- The network layer services have been designed with the following goals:**
  - 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 present.**
  - 3. The network addresses made available to the transport layer should use a uniform numbering plan, even across**

**LANs and WANs.**

# One camp's view


The routers' job is moving packets around and nothing else. In their view, the subnet is inherently unreliable, no matter how it is designed. Therefore, the hosts should accept the fact that the network is **unreliable** and do error control (i.e., error detection and correction) and flow control themselves. So the network service should be **connectionless**.

The Internet offers connectionless network-layer service



# The other camp's view


The subnet should provide a **reliable, connection-oriented** service. In this view, quality of service is the dominant factor, and without connections in the subnet, quality of service is very difficult to achieve, especially for real-time traffic such as voice and video.



ATM networks offer connection-oriented network-layer service.

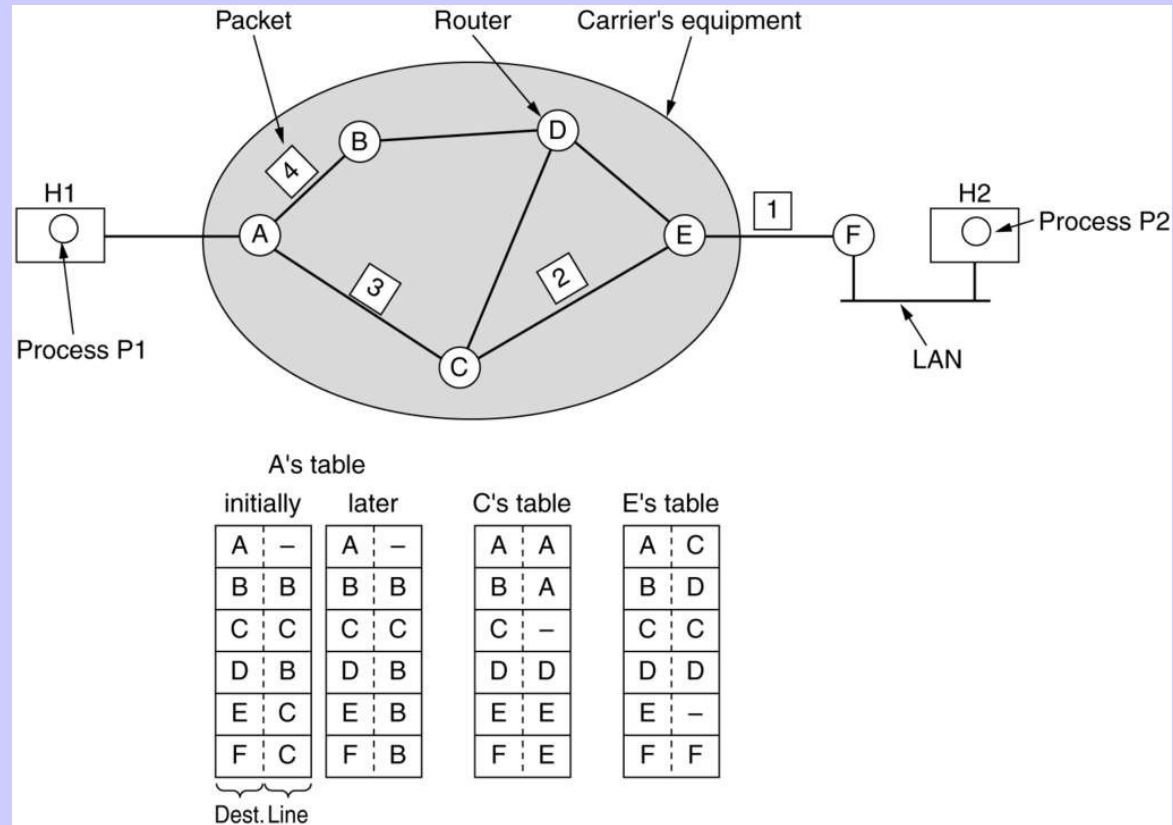
# 5.1 Network Layer Design Issues

- a) Store-and-Forward Packet Switching
- b) Services Provided to the Transport Layer
- c) **Implementation of Connectionless Service**
- d) Implementation of Connection-Oriented Service
- e) Comparison of Virtual-Circuit and Datagram Subnets



- a) Two different organizations are possible, depending on the type of service offered.**
- b) If connectionless service is offered, packets are injected into the subnet individually and routed independently of each other. No advance setup is needed. In this context, the packets are frequently called **datagrams** and the subnet is called a **datagram subnet**.**
- c) If connection-oriented service is used, a path from the source router to the destination router must be established before any data packets can be sent. This connection is called a **VC (virtual circuit)** and the subnet is called a **virtual-circuit subnet**.**

# Implementation of Connectionless Service




P346

**The question is:** a packet with a destination D arrives at router A. then which router will router A send this packet to?

# 5.1 Network Layer Design Issues

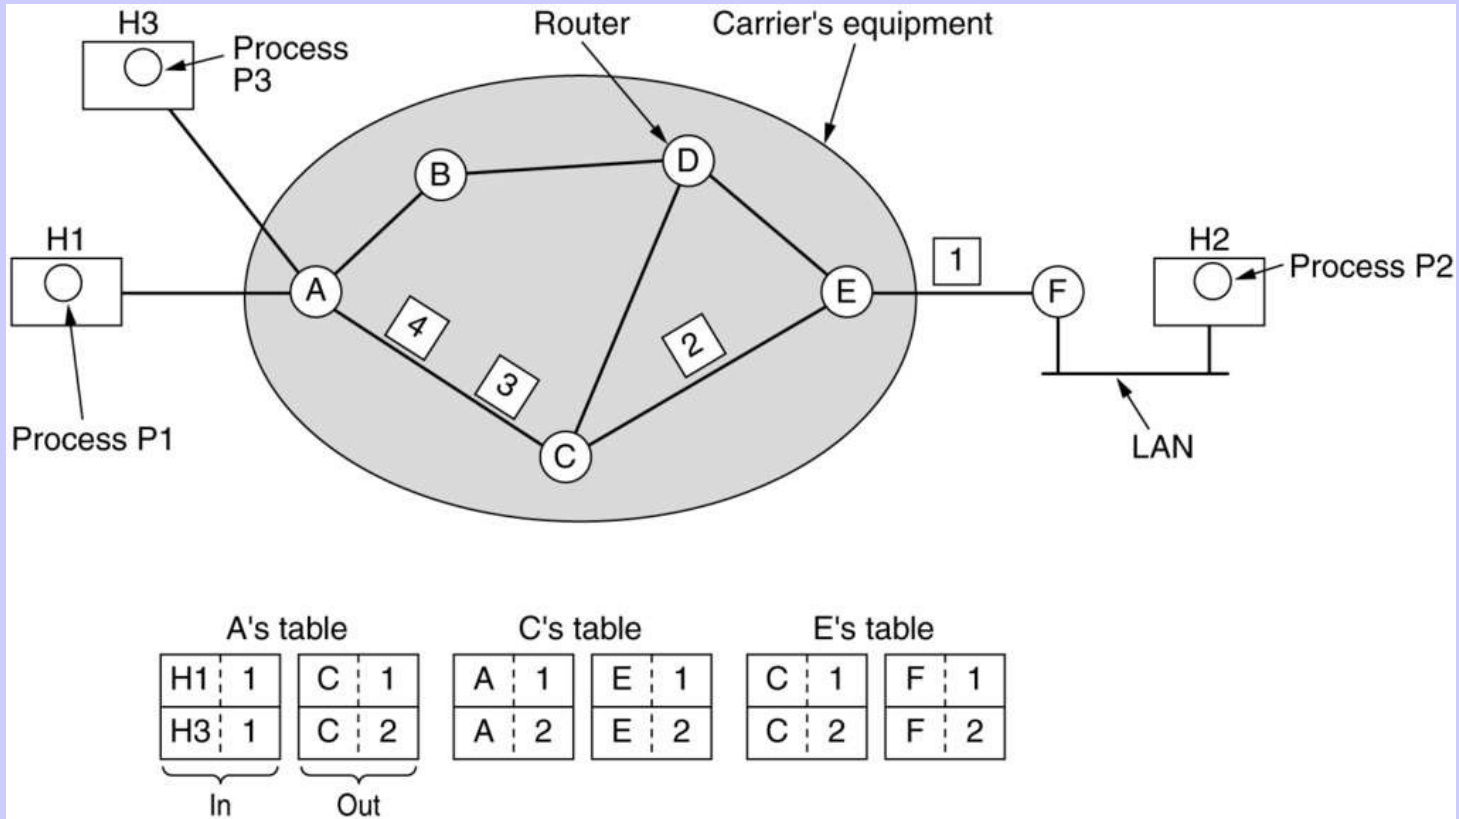
- a) Store-and-Forward Packet Switching
- b) Services Provided to the Transport Layer
- c) Implementation of Connectionless Service
- d) **Implementation of Connection-Oriented Service**
- e) Comparison of Virtual-Circuit and Datagram Subnets



- a) For connection-oriented service, we need a virtual-circuit subnet.
- b) The idea behind virtual circuits is to **avoid having to choose a new route for every packet sent**. Instead, when a connection is established, a route from the source machine to the destination machine is chosen as part of the connection setup and stored in tables inside the routers. **That route is used for all traffic flowing over the connection**, exactly the same way that the telephone system works. When the connection is released, the virtual circuit is also terminated.
- c) With connection-oriented service, each packet carries an **identifier** telling which virtual circuit it belongs to.




# Implementation of Connection-Oriented Service



Routing within a virtual-circuit subnet.

# 5.1 Network Layer Design Issues

- a) Store-and-Forward Packet Switching
- b) Services Provided to the Transport Layer
- c) Implementation of Connectionless Service
- d) Implementation of Connection-Oriented Service
- e) **Comparison of Virtual-Circuit and Datagram Subnets**




# Comparison of Virtual-Circuit and Datagram Subnets

<b>Issue</b>	<b>Datagram subnet</b>	<b>Virtual-circuit subnet</b>
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


- a) Inside the subnet, several trade-offs exist between virtual circuits and datagrams.**
- b) One trade-off is between router memory space and bandwidth.**
- **Virtual circuits allow packets to contain circuit numbers instead of full destination addresses. If the packets tend to be fairly short, a full destination address in every packet may represent a significant amount of overhead and hence, wasted bandwidth. The price paid for using virtual circuits internally is the table space within the routers. Depending upon the relative cost of communication circuits versus router memory, one or the other may be cheaper.**

a) **Another trade-off** is **setup time** versus **address parsing time**.

- **Using virtual circuits requires a setup phase, which takes time and consumes resources. However, figuring out what to do with a data packet in a virtual-circuit subnet is easy: the router just uses the circuit number to index into a table to find out where the packet goes. In a datagram subnet, a more complicated lookup procedure is required to locate the entry for the destination.**



- a) **Virtual circuits have some advantages in guaranteeing quality of service and avoiding congestion within the subnet.**
- b) **The loss of a communication line is fatal to virtual circuits using it but can be easily compensated for if datagrams are used. Datagrams also allow the routers to balance the traffic throughout the subnet, since routes can be changed partway through a long sequence of packet transmissions.**



# summary

- Store-and-Forward Packet Switching
- Services Provided to the Transport Layer
- Implementation of Connectionless Service
- Implementation of Connection-Oriented Service
- Comparison of Virtual-Circuit and Datagram Subnets

# Chapter 5 The Network Layer

## 5.1 Network Layer Design Issues

## 5.2 Routing Algorithms


## 5.6 The Network Layer in the Internet






# Description of Routing Algorithms

- 1 Definition:** The routing algorithm is that part of the network layer software responsible for deciding which output line an incoming packet should be transmitted on.



- 2 Properties of routing algorithm:** correctness, simplicity, robustness, stability, fairness, and optimality.


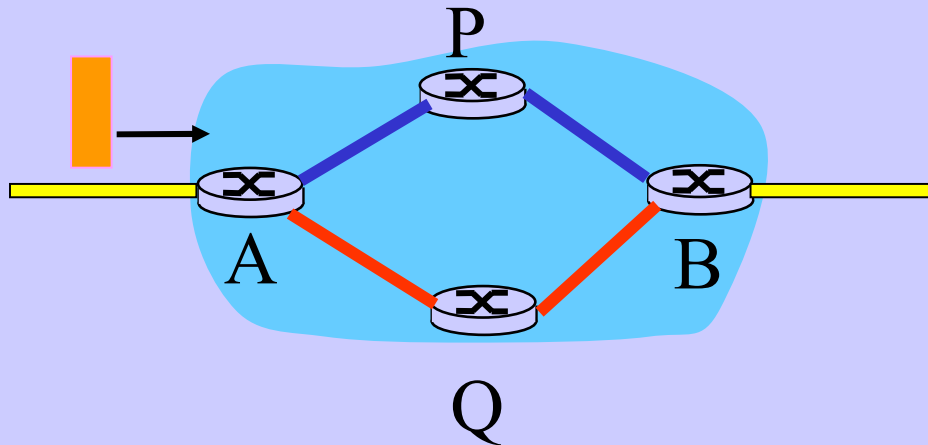


# Description of Routing Algorithms


- 1) Robustness:** Once a major network comes on the air, it may be expected to run continuously for years without system-wide failures. During that period there will be hardware and software failures of all kinds. Hosts, routers, and lines will fail repeatedly, and the topology will change many times. The routing algorithm should be able to cope with changes in the topology and traffic without requiring all jobs in all hosts to be aborted and the network to be rebooted every time some router crashes.

# Description of Routing Algorithms

- 2) **Stability**: It is also an important goal for the routing algorithm. There exist routing algorithms that never converge to equilibrium, no matter how long they run. A stable algorithm reaches equilibrium and stays there.



3) **Fairness and optimality** may sound obvious, but as it turns out, they are often contradictory goals.




i) There is enough traffic between A and A', between B and B', and between C and C' to saturate the horizontal links. To maximize the total flow, the X to X' traffic should be shut off altogether. Unfortunately, X and X' may not see it that way. Evidently, some compromise between global efficiency and fairness to individual connections is needed.

# Description of Routing Algorithms

3 Category of algorithm: ~~nonadaptive~~ and adaptive.

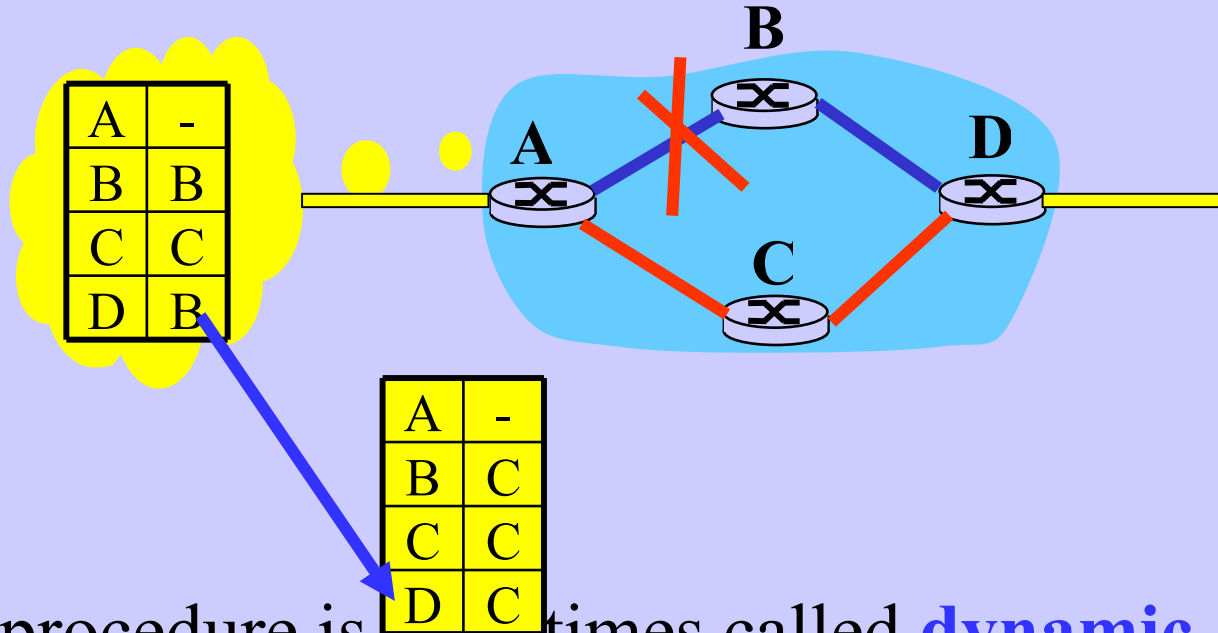
b) **Nonadaptive algorithms** do not base their routing decisions on measurements or estimates of the current traffic and topology. Instead, the choice of the route to use to get from I to J is computed in advance, off-line, and downloaded to the routers when the network is booted.



g) This procedure is sometimes called **static routing**.

# Description of Routing Algorithms


2) **Adaptive algorithms**, in contrast, change their routing decisions to reflect changes in the topology, and usually the traffic as well.



h) This procedure is sometimes called **dynamic routing**.


# 5.2 Routing Algorithms

- The Optimality Principle
- Shortest Path Routing
- Flooding
- Distance Vector Routing
- Link State Routing
- Hierarchical Routing
- Broadcast Routing



# 5.2.1 The Optimality Principle


- a) **The Optimality Principle:** 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.





# 5.2.1 The Optimality Principle

- a) The set of optimal routes from all sources to a given destination form a tree rooted at the destination. Such a tree is called a **sink tree**.
- b) Figure (a) A subnet. (b) A sink tree for router B.



# 5.2.1 The Optimality Principle


- a) **Note:** A sink tree is not necessarily unique; other trees with the same path lengths may exist.
- b) The goal of all routing algorithms is to discover and use the sink trees for all routers.

# 5.2.2 Shortest Path Routing

- a) A **technique** to study routing algorithms: The idea is to build a graph of the subnet, with each node of the graph representing a router and each arc of the graph representing a communication line (often called a link).
- b) To choose a route between a given pair of routers, the algorithm just finds the shortest path between them on the graph.
- c) One way of measuring path length is the number of **hops**. Another metric is the **geographic distance** in kilometers . Many other metrics are also possible. For example, each arc could be labeled with the **mean queuing and transmission delay** for some standard test packet as determined by hourly test runs.
- d) In the general case, the labels on the arcs could be computed as a **function** of the distance, bandwidth, average traffic, communication cost, mean queue length, measured delay, and other factors. By changing the weighting function, the algorithm would then compute the "shortest" path measured according to any one of a number of criteria or to a combination of criteria.

# 5.2.2 Shortest Path Routing







- a) **Dijkstra algorithm**: Each node is labeled (in parentheses) with its distance from the source node along the best known path. Initially, no paths are known, so all nodes are labeled with infinity. As the algorithm proceeds and paths are found, the labels may change, reflecting better paths. A label may be either **tentative** or **permanent**. Initially, all labels are tentative. When it is discovered that a label represents the shortest possible path from the source to that node, it is made permanent and never changed thereafter.



( 1930 年 5 月  
11 日 ~ 2002 年  
8 月 6 日 )

P353

Next steps?



- Figure. The first five steps used in computing the shortest path from A to D. The arrows indicate the working node.

The shortest path from A to D

ABEFHD

## 5.2.3 Flooding

- a) **Flooding algorithm**: every incoming packet is sent out on every outgoing line except the one it arrived on.
- b) Flooding obviously generates vast numbers of duplicate packets, in fact, an infinite number unless some measures are taken to damp the process.
- c) One such measure is to have a hop counter contained in the header of each packet, which is decremented at each hop, with the packet being discarded when the counter reaches zero.
- d) An alternative technique for damming the flood is to keep track of which packets have been flooded, to avoid sending them out a second time.
  - How to implement that? (P355)

## 5.2.3 Flooding

- a) A variation of flooding that is slightly more practical is **selective flooding**. In this algorithm the routers do not send every incoming packet out on every line, only on those lines that are going approximately in the right direction.
- b) Applications of flooding algorithm:
  1. military applications
  2. distributed database applications
  3. wireless networks
  4. as a metric against which other routing algorithms can be compared

# 5.2.4 Distance Vector Routing

- a) A dynamic routing algorithm
- b) **Distance vector routing algorithms** operate by having each router maintain a table (i.e, a vector) giving the best known distance to each destination and which line to use to get there. These tables are updated by exchanging information with the neighbors. (also named the distributed Bellman-Ford routing algorithm and the Ford-Fulkerson algorithm )



**Lester Randolph Ford**

NO PHOTO

**August 26, 1920**  
**~ March 19, 1984**

**September 23, 1927**  
**~**

**Delbert Ray Fulkerson**


NO PHOTO

**August 14, 1924**  
**~ January 10,**  
**1976**



# 5.2.4 Distance Vector Routing

- a) **Table content:** In distance vector routing, each router maintains a routing table indexed by, and containing one entry for, each router in the subnet. This entry contains two parts: the preferred outgoing line to use for that destination and an estimate of the time or distance to that destination.
- b) **Table updating method:** Assume that the router knows the delay to each of its neighbors. Once every  $T$  msec each router sends to each neighbor a list of its estimated delays to each destination. It also receives a similar list from each neighbor. Imagine that one of these tables has just come in from neighbor  $X$ , with  $X_i$  being  $X$ 's estimate of how long it takes to get to router  $i$ . If the router knows that the delay to  $X$  is  $m$  msec, it also knows that it can reach router  $i$  via  $X$  in  $X_i + m$  msec. By performing this calculation for each neighbor, a router can find out which estimate seems the best and use that estimate and the corresponding line in its new routing table. Note that the old routing table is not used in the calculation.



(a)

To	A	I	H	K	New estimated delay from J	
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	JI	JH	JK
delay is	8	10	12	6

Vectors received from J's four neighbors

New routing table for J	
8	A
20	A
28	I
20	H
17	I
30	I
18	H
12	H
10	I
0	-
6	K
15	K

(b)

Part (a) shows a subnet. The first four columns of part (b) show the delay vectors received from the neighbors of router J. Suppose that J has measured or estimated its delay to its neighbors, A, I, H, and K as 8, 10, 12, and 6 msec, respectively.

# 5.2.5 Link State Routing

- a) A dynamic routing algorithm
- b) The idea behind link state routing can be stated as five parts. Each router must do the following:
  1. Discover its neighbors and learn their network addresses.
  2. Measure the delay or cost to each of its neighbors.
  3. Construct a packet telling all it has just learned.
  4. Send this packet to all other routers.
  5. Compute the shortest path to every other router.
- In effect, the complete topology and all delays are experimentally measured and distributed to every router. Then Dijkstra's algorithm can be run to find the shortest path to every other router.

# 5.2.5 Link State Routing

## 1. Learning about the Neighbors

- It accomplishes this goal by sending a special HELLO packet on each point-to-point line. The router on the other end is expected to send back a reply telling who it is. These names must be globally unique because when a distant router later hears that three routers are all connected to F, it is essential that it can determine whether all three mean the same F.

## 2. Measuring Line Cost

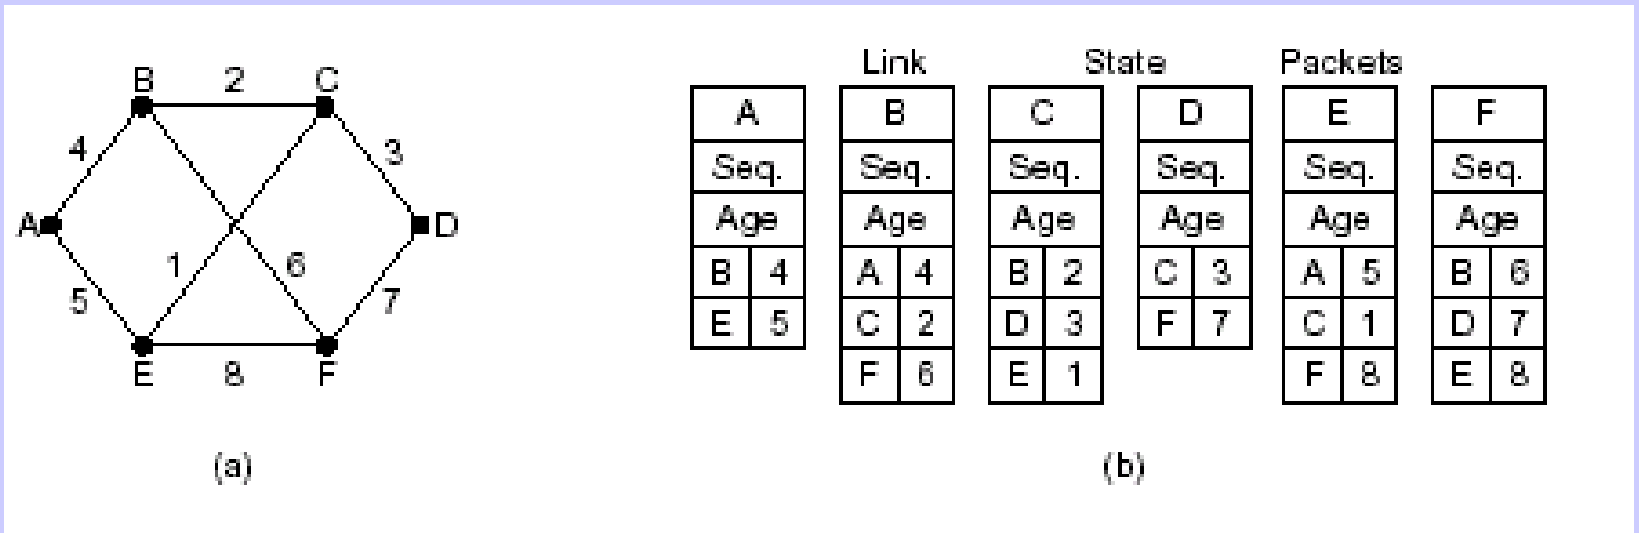
- The most direct way to determine this delay is to send over the line a special ECHO packet that the other side is required to send back immediately. By measuring the round-trip time and dividing it by two, the sending router can get a reasonable estimate of the delay. For even better results, the test can be conducted several times, and the average used. Of course, this method implicitly assumes the delays are symmetric, which may not always be the case.

# 5.2.5 Link State Routing

## 3. Building Link State Packets

- The packet starts with the identity of the sender, followed by a sequence number and age (to be described later), and a list of neighbors. For each neighbor, the delay to that neighbor is given.

(a) A subnet. (b) The link state packets for this subnet



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)

# 5.2.5 Link State

## Routing

- Building the link state packets is easy. The hard part is determining when to build them. One possibility is to build them periodically, that is, at regular intervals. Another possibility is to build them when some significant event occurs, such as a line or neighbor going down or coming back up again or changing its properties appreciably.

### 4. Distributing the Link State Packets

- c) **The basic distribution algorithm:** The fundamental idea is to use flooding to distribute the link state packets. To keep the flood in check, each packet contains a sequence number that is incremented for each new packet sent. Routers keep track of all the (source router, sequence) pairs they see. When a new link state packet comes in, it is checked against the list of packets already seen. If it is new, it is forwarded on all lines except the one it arrived on. If it is a duplicate, it is discarded. If a packet with a sequence number lower than the highest one seen so far ever arrives, it is rejected as being obsolete since the router has more recent data.

# 5.2.5 Link State

## Routing

- a) First problem with this algorithm:if the sequence numbers wrap around, confusion will reign. The solution here is to use a 32-bit sequence number. With one link state packet per second, it would take 137 years to wrap around, so this possibility can be ignored.
- b) Second problem :if a router ever crashes, it will lose track of its sequence number. If it starts again at 0, the next packet will be rejected as a duplicate.
- c) Third problem :if a sequence number is ever corrupted and 65,540 is received instead of 4 (a 1-bit error), packets 5 through 65,540 will be rejected as obsolete, since the current sequence number is thought to be 65,540.

# 5.2.5 Link State

## Routing

- a) The solution to all these problems is to include the age of each packet after the sequence number and decrement it once per second. When the age hits zero, the information from that router is discarded.

### 5. Computing the New Routes

- Once a router has accumulated a full set of link state packets, it can construct the entire subnet graph because every link is represented.
- Now Dijkstra's algorithm can be run locally to construct the shortest path to all possible destinations.

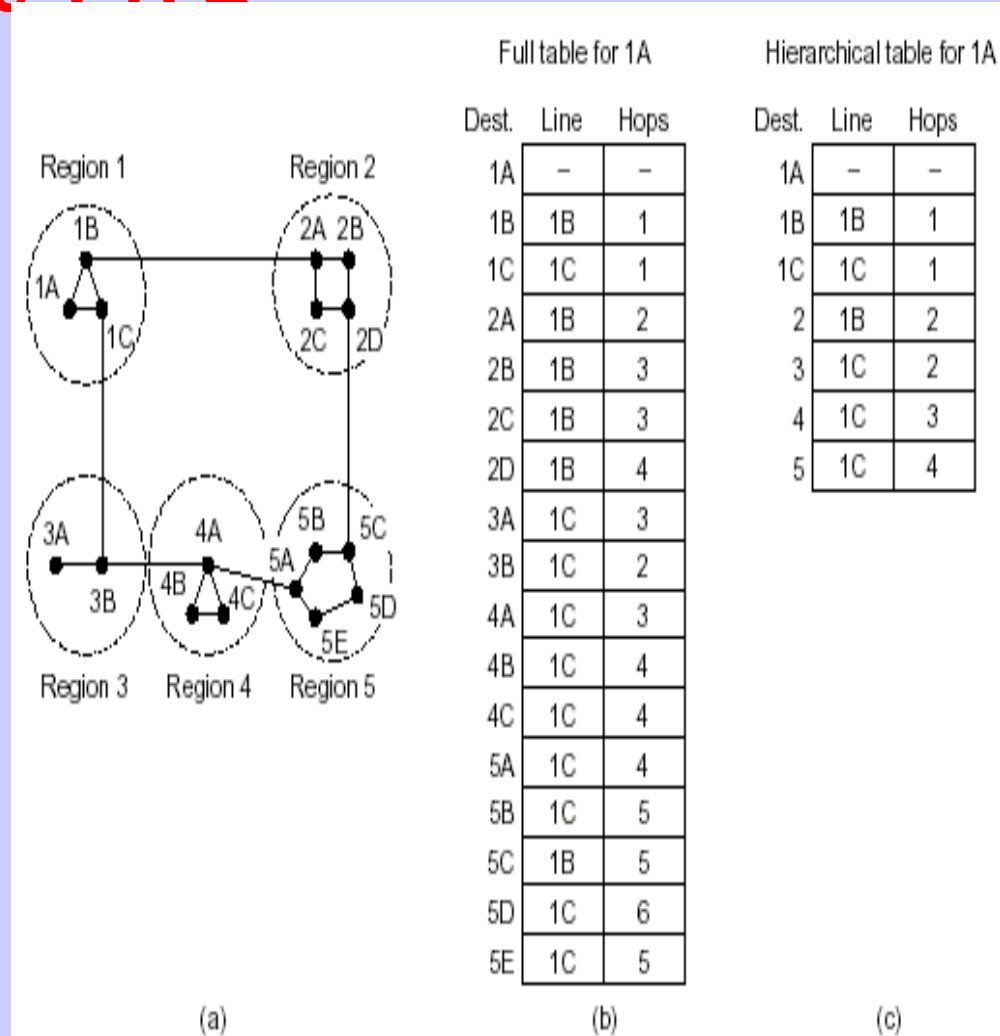


# 5.2.6 Hierarchical Routing

- The routers are divided into what we will call regions, with each router knowing all the details about how to route packets to destinations within its own region, but knowing nothing about the internal structure of other regions.
- For huge networks, a two-level hierarchy may be insufficient; it may be necessary to group the regions into clusters, the clusters into zones, the zones into groups, and so on, until we run out of names for aggregations.

# 5.2.6 Hierarchical Routing

- The full routing table for router 1A has 17 entries, as shown in (b). When routing is done hierarchically, as in (c), there are entries for all the local routers as before, but all other regions have been condensed into a single router, so all traffic for region 2 goes via the 1B -2A line, but the rest of the remote traffic goes via the 1C -3B line. Hierarchical routing has reduced the table from 17 to 7 entries.



# 5.2.6 Hierarchical Routing

- Unfortunately, these gains in space are not free. There is a penalty to be paid, and this penalty is in the form of increased path length. For example, the best route from 1A to 5C is via region 2, but with hierarchical routing all traffic to region 5 goes via region 3, because that is better for most destinations in region 5.

# 5.2.7 Broadcast Routing

- a) Sending a packet to all destinations simultaneously is called **broadcasting**.
- 2) The source simply sends a distinct packet to each destination. Not only is the method wasteful of bandwidth, but it also requires the source to have a complete list of all destinations.
- 3) Flooding.
  - The problem with flooding as a broadcast technique is that it generates too many packets and consumes too much bandwidth.

# 5.2.7 Broadcast Routing

1) Multi-destination routing.

- If this method is used, each packet contains either a list of destinations or a bit map indicating the desired destinations. When a packet arrives at a router, the router checks all the destinations to determine the set of output lines that will be needed. (An output line is needed if it is the best route to at least one of the destinations.) The router generates a new copy of the packet for each output line to be used and includes in each packet only those destinations that are to use the line. In effect, the destination set is partitioned among the output lines. After a sufficient number of hops, each packet will carry only one destination and can be treated as a normal packet.

# 5.2.7 Broadcast Routing


- 1) A fourth broadcast algorithm makes explicit use of the sink tree for the router initiating the broadcast—or any other convenient spanning tree for that matter.
- b) A **spanning tree** is a subset of the subnet that includes all the routers but contains no loops.
- c) If each router knows which of its lines belong to the spanning tree, it can copy an incoming broadcast packet onto all the spanning tree lines except the one it arrived on.

# 5.2.7 Broadcast Routing

- 1) Reverse path forwarding.
  - When a broadcast packet arrives at a router, the router checks to see if the packet arrived on the line that is normally used for sending packets to the source of the broadcast. If so, there is an excellent chance that the broadcast packet itself followed the best route from the router and is therefore the first copy to arrive at the router. This being the case, the router forwards copies of it onto all lines except the one it arrived on. If, however, the broadcast packet arrived on a line other than the preferred one for reaching the source, the packet is discarded as a likely duplicate.

# 5.2.7 Broadcast Routing

- how does the reverse path algorithm works? P369



Reverse path forwarding. (a) A subnet. (b) A sink tree. (c) The tree built by reverse path forwarding.



# Chapter 5 The Network Layer

**5.1 Network Layer Design Issues**


**5.2 Routing Algorithms**

**5.6 The Network Layer in the Internet**




# 5.6 The Network Layer in the Internet

- The IP Protocol
- IP Addresses
- Internet Control Protocols



# 5.6.1 The IP Protocol


- An IP datagram consists of a header part and a text part. The header has a 20-byte fixed part and a variable length optional part.



P<sub>434</sub>

# 5.6.2 IP Addresses

- Every host and router on the Internet has an IP address, which encodes its network number and host number.
- All IP addresses are 32 bits long. It is important to note that an IP address does not actually refer to a host. It really refers to a network interface, so if a host is on two networks, it must have two IP addresses.
- IP addresses were divided into the five categories



network mask

255.0.0.0

255.255.0.0

255.255.255.0


# 5.6.2 IP Addresses

- The values 0 and -1 (all 1s) have special meanings. The value 0 means this network or this host. The value of -1 is used as a broadcast address to mean all hosts on the indicated network.

0 0		This host		
0 0	...	0 0	Host	A host on this network
1 1				Broadcast on the local network
Network	1 1 1 1	...	1 1 1 1	Broadcast on a distant network
127	(Anything)			Loopback


# 5.6.2 IP Addresses

- All the hosts in a network must have the same network number. This property of IP addressing can cause problems as networks grow. For example.....
- The problem is the rule that a single class A, B, or C address refers to one network, not to a collection of LANs.
- The solution is to allow a network to be split into several parts for internal use but still act like a single network to the outside world.



# 5.6.2 IP Addresses

- To implement subnetting, the main router needs a subnet mask that indicates the split between network + subnet number and host.
- For example, if the university has a B address(130.50.0.0) and 35 departments, it could use a 6-bit subnet number and a 10-bit host number, allowing for up to 64 Ethernets, each with a maximum of 1022 hosts.



- The subnet mask can be written as 255.255.252.0. An alternative notation is /22 to indicate that the subnet mask is 22 bits long.

# 5.6.3 Internet Control Protocols

## 1、The Internet Control Message Protocol

- The operation of the Internet is monitored closely by the routers. When something unexpected occurs, the event is reported by the ICMP (Internet Control Message Protocol), which is also used to test the Internet.
- Each ICMP message type is encapsulated in an IP packet

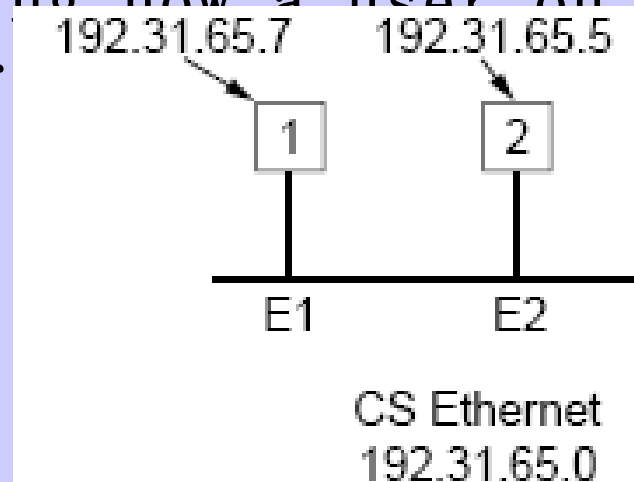
Message type	Description
Destination unreachable	Packet could not be delivered
Time exceeded	Time to live field hit 0
Parameter problem	Invalid header field
Source quench	Choke packet
Redirect	Teach a router about geography
Echo	Ask a machine if it is alive
Echo reply	Yes, I am alive
Timestamp request	Same as Echo request, but with timestamp
Timestamp reply	Same as Echo reply, but with timestamp



# 5.6.3 Internet Control Protocols

## 2、 ARP—The Address Resolution Protocol

- Most hosts at companies and universities are attached to a LAN by an interface board that only understands LAN addresses.
- The question : How do IP addresses get mapped onto data link layer addresses, such as Ethernet?
- Let us start out by seeing how a user on host 1 sends a packet to a user



# 5.6.3 Internet Control Protocols

- 1) The upper layer software on host 1 now builds a packet with 192.31.65.5 in the Destination address field and gives it to the IP software to transmit.
- 2) The IP software can look at the address and see that the destination is on its own network, but it needs some way to find the destination's Ethernet address.
  - Host 1 outputs a broadcast packet onto the Ethernet asking: Who owns IP address 192.31.65.5? The broadcast will arrive at every machine on Ethernet 192.31.65.0, and each one will check its IP address.
  - Host 2 alone will respond with its Ethernet address (E2). In this way host 1 learns that IP address 192.31.65.5 is on the host with Ethernet address E2.
  - The protocol used for asking this question and getting the reply is called ARP (Address Resolution Protocol).

# 5. 6. 3 Internet Control Protocols

- 3 ) The IP software on host 1 builds an Ethernet frame addressed to E2, puts the IP packet (addressed to 192.31.65.5) in the payload field, and dumps it onto the Ethernet.
- 4 ) The Ethernet board of host 2 detects this frame, recognizes it as a frame for itself, scoops it up, and causes an interrupt. The Ethernet driver extracts the IP packet from the payload and passes it to the IP software, which sees that it is correctly addressed and processes it.

# 5.6.3 Internet Control Protocols

3、RARP, BOOTP, and DHCP

- b) Given an Ethernet address, what is the corresponding IP address? In particular, this problem occurs when a diskless workstation is booted.
- c) The first solution devised was to use RARP (Reverse Address Resolution Protocol). This protocol allows a newly-booted workstation to broadcast its Ethernet address and say: My 48-bit Ethernet address is 14.04.05.18.01.25. Does anyone out there know my IP address? The RARP server sees this request, looks up the Ethernet address in its configuration files, and sends back the corresponding IP address.

# 5. 6. 3 Internet Control Protocols


- A disadvantage of RARP is that it uses a destination address of all 1s (limited broadcasting) to reach the RARP server. However, such broadcasts are not forwarded by routers, so a RARP server is needed on each network.
- Unlike RARP, BOOTP uses UDP messages, which are forwarded over routers. It also provides a diskless workstation with additional information, including the IP address of the file server holding the memory image, the IP address of the default router, and the subnet mask to use.
- A serious problem with BOOTP is that it requires manual configuration of tables mapping IP address to Ethernet address.


# 5. 6. 3 Internet Control Protocols

- DHCP allows both manual IP address assignment and automatic assignment.
- Like RARP and BOOTP, DHCP is based on the idea of a special server that assigns IP addresses to hosts asking for one. This server need not be on the same LAN as the requesting host.

# 5.6.3 Internet Control Protocols

- To find its IP address, a newly-booted machine broadcasts a DHCP DISCOVER packet. The DHCP relay agent on its LAN intercepts all DHCP broadcasts. When it finds a DHCP DISCOVER packet, it sends the packet as a unicast packet to the DHCP server, possibly on a distant network. The only piece of information the relay agent needs is the IP address of the DHCP server.





**Over!**

**Thank you very much!**



**robustness**

健壮性

**optimality**

**最优性**

**topology**

**拓扑结构**

**converge**

汇聚

**equilibrium**

平衡

**nonadaptive**

**adaptive**

非自适应

适应

自

**metric**

**参数，度量**

# hop 跳数

