1. Computer Networks & the Internet:

A network is a group of connected, communicating devices such as computers and printers. An internet.

- An internet (small i) is two or more networks that can communicate with each other.
- The most notable internet is Internet (capital I) composed of hundreds of thousands of interconnected networks so that a host on one network could communicate with a host on a second.

The public Internet is a world-wide **computer network**, i.e., a network that interconnects millions of computing devices throughout the world.

Computing Devices called **hosts** or **end systems** could be:

- traditional desktop PCs,
- Unix-based workstations,
- Servers that store and transmit information such as WWW pages and e-mail messages.
- Increasingly, non-traditional computing devices such as Web TVs, mobile computers, unusual devices (such as toasters) have been hooked up to the Internet.

1.1 Internets

- Consists of millions of connected computing devices (hosts = end systems) used for running network applications
- Communication links (fiber, copper, radio, satellite) are used to connect computing devices and networks with specific transmission rate (i.e. Different links can transmit data at different rates). Transmission rate is often called the link bandwidth, and is typically measured in bits/second. routers: forward packets (chunks of data)

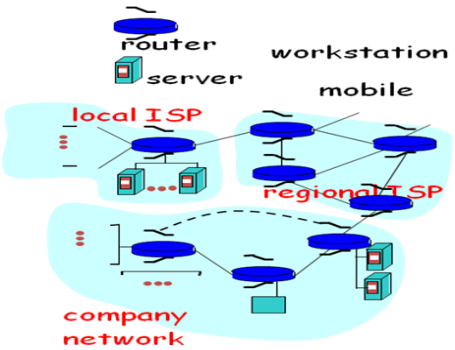

Why Networking?

- Distributed Software
  - Application : WEB, email, 3-tier appl., …
  - Database
  - Directory
- Resource Sharing
  - File, Software, Data, … (Network File System, File Transfer, …)
  - CPU, Memory, Peripherals, …
- Communication
  - Email, Chat, TV, Radio, Video Conference, Telephone, .
  - Virtual Terminal (Remote Login)

Internet Services

- Search Engines (Google)
- Email (Hotmail)
- Shopping (Amazon)
- Auctions (eBay)
- Chat (AOL)

Goals?

- Fast service (low latency)
- Service all users (scalability)
- Always available (fault tolerance)

1.2 Intranet (Intranet: access is denied from outside)

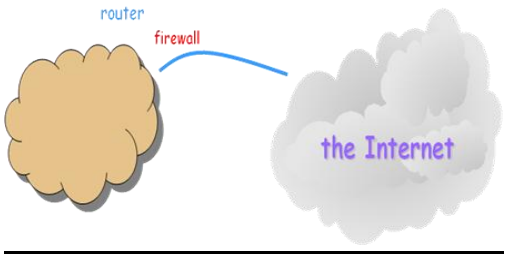

A private corporate network consisting of hosts, routers, and networks that use TCP/IP technology. An intranet may or may not connect to the global Internet.

1.3 Extranet

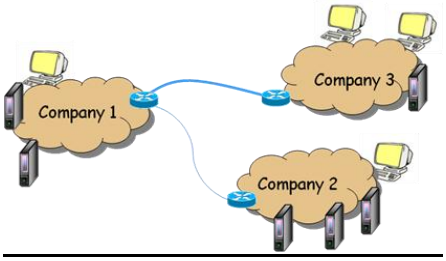

An internet of networks each of which is belong to individual company or organization

2.Network Structure

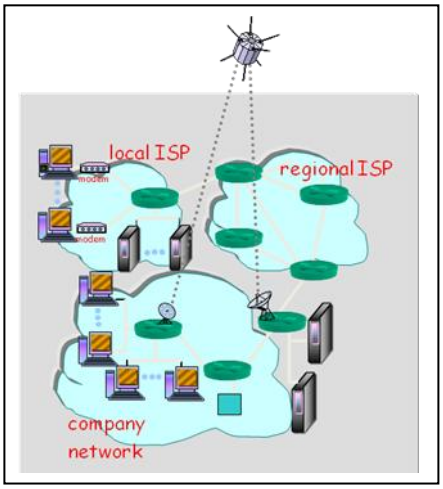

- *network edge*: applications and hosts
  a) end systems (hosts):
     - run application programs
     - e.g. Web, email
     - at "edge of network"
  b) client/server model
     - client host requests, receives service from always-on server
       e.g. Web browser/server; email client/server
  c) peer-peer model:
     - minimal (or no) use of dedicated servers
       e.g. Gnutella, KaZaA
- *network core:*
  - network of networks
    - access networks, physical media: communication links
  - routers: mesh of interconnected routers
    the fundamental question: how is data transferred through net?
    - circuit switching: dedicated circuit per call: telephone net
    - packet-switching: data sent through net in discrete "chunks"

Nodes are connected indirectly through switching devises (most popular are routers and /or link layer switches). Each packet has a route or path from source to destination. Each system access the internet through internet service provider (ISP) such as local telephone company, etc..

Communication networks can be classified based on the way in which the nodes exchange information:

2.1 Network Core

- **routers**: switches forwarding data

  the fundamental question: how is data transferred through net?

  - circuit switching: dedicated circuit per call: telephone net
  - packet-switching: data sent thru net in discrete "chunks"

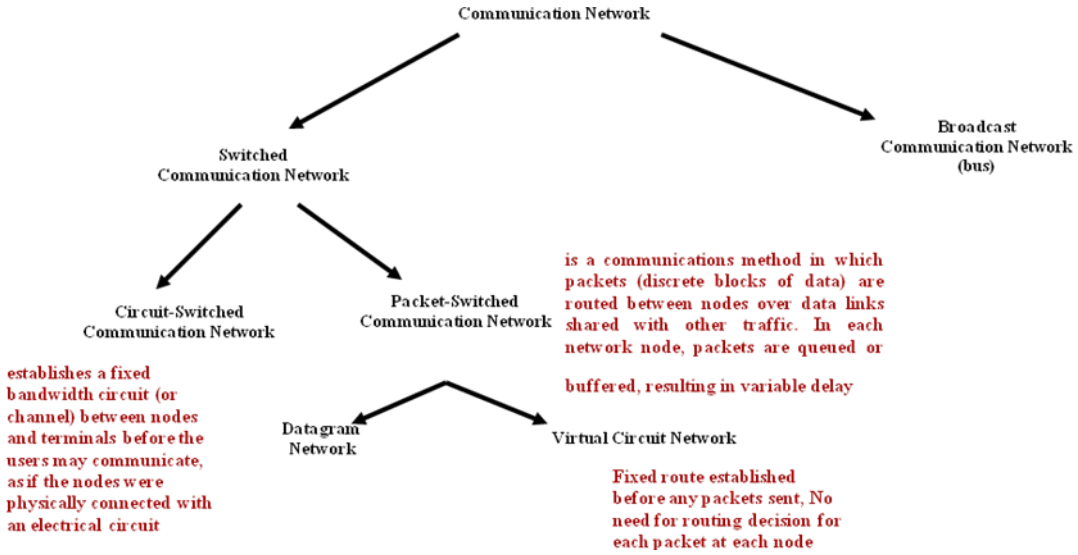

Circuit switching, the resources needed along a path (buffers, link bandwidth) to provide for communication between the end systems are *reserved* for the duration of the session. When two hosts desire to communicate, the network establishes a dedicated end-to-end circuit between two hosts. (Conference calls between more than two devices are, of course, also possible. In **circuit switching**, End-end resources reserved for "call"

- link bandwidth, switch capacity
- dedicated resources: no sharing
- circuit-like (guaranteed) performance
- call setup required

Packet-switched networks, these resources are not reserved; a session's messages use the resource on demand, and as a consequence, may have to wait (i.e., queue) for access to a communication link. In packet switch network, each end-end data stream divided into *packets*

- user A, B packets share network resources
- each packet uses full link bandwidth resources used as needed,
- Bandwidth division into "pieces", Dedicated allocation, Resource reservation are not allowed

Resource contention:

- Total resource demand can exceed amount available
- Congestion: packets queue, wait for link use
- Store and forward: packets move one hop at a time
  - transmit over link
  - wait turn at next link
  - Node receives complete packet before forwarding

Internet is a quintessential packet-switched network. Consider what happens when one host wants to send a packet to another host over:

- Circuit-switching, the packet is transmitted over a series of communication links.
- Packet-switching, the packet is sent into the network without reserving any bandwidth whatsoever. If one of the links is congested because other packets need to be transmitted over the link at the same time, then our packet will have to wait in a buffer at the sending side of the transmission line, and suffer a delay. The Internet

makes its best effort to deliver the data in a timely manner. But it does not make any guarantees.

Packet Switching vs Circuit Switching: Why?

- "reliability" – no congestion, in order data in circuit-switching
- packet switching: better sharing of bandwidth
- state, resources: packet switching has less state
  - advantage less control-plane processing resources along the way
  - More data plane (address lookup) processing
- Packet switching allows more users to use network!
- failure modes (routers/links down):
  - packet switching routing reconfigures sub-second timescale;
  - circuit-switching: more complex recovery – need to involve all (downstream) switches on path

Internet provides both *connection-oriented* (TCP) and *connectionless services* (UDP) to apps.

A) Virtual Circuit Networks: A virtual circuit (VC) consists of

1) A path (a series of links & packet switches) between the source and destination hosts
2) VC numbers, one number for each link along the path, and
3) Entries in VC-number translation tables in each packet switch along the path.

Once a VC is established between source and destination, packets can be sent with the appropriate VC numbers. Because a VC has a different VC number on each link, an intermediate packet switch must replace the VC number of each traversing packet with a new one.

The new VC number is obtained from the VC number translation table. To illustrate the concept, suppose host A requests that the network establish a VC between itself and host B. Suppose that the network chooses the path A - PS1 - The Network Core PS2 - B and assigns VC numbers 12, 22, 32 to the three links in this path. Then, when a packet as part of this VC leaves host A, the value in the VC number field is 12; when it leaves PS1, the value is 22; and when it leaves PS2, the value is 32. The numbers next to the links of PS1 are the **interface numbers**.

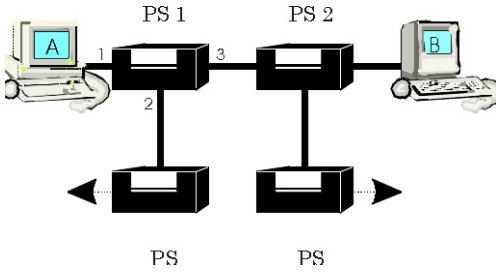

Each switch has a VC number translation table; for example, the VC number translation table in PS 1 might look something like this:

Incoming Interface	Incoming VC#	Outgoing Interface	Outgoing VC#
1	12	3	22
2	63	1	18
3	7	2	17
1	97	3	87

Whenever a new VC is established across a switch, an entry is added to the VC number table.

B) Datagram network: In a datagram network, each packet that traverses the network contains in its header the address of the destination. When a packet arrives at a packet switch in the network, the packet switch examines a portion of the packet's destination address and forwards the packet to an adjacent switch. More specifically, each packet switch has a routing table which maps destination addresses (or portions of the destination addresses) to an outbound link.

2.2 Network Edge: applications and hosts

Hosts or end systems are the computers that we use on a daily basis. They are referred to as "hosts" because they host (run) application-level programs such as a Web browser or server program, or an e-mail program. They are also referred to as "end systems" because they sit at the "edge" of the Internet.

Hosts are sometimes further divided into two categories: clients and servers.

In the so-called client-server model,

- A client program running on one end system requests and receives information from a server running on another end system.
- Since a client typically runs on one computer and the server runs on another computer, client/server Internet applications are, by definition, distributed applications.
- The client and the server interact with each other by communicating (i.e., sending each other message) over the Internet. At this level of abstraction, the routers, links and other "pieces" of the Internet serve as a "black box" that transfers messages between the distributed, communicating components of an Internet application.

The links, routers and other pieces of the Internet provide the means to transport these messages between the end system applications. The Internet provides two types of services to its applications: ***connectionless service*** and ***connection-oriented service.***

A) connection-oriented service (The Internet's connection-oriented service has a name -- TCP (Transmission Control Protocol); When an application uses the connection-oriented service, the client and the server (residing in different end systems) send control packets to each other before sending packets with real data (such as e-mail messages), this so-called ***handshaking*** procedure. Once the handshaking procedure is finished, a "connection" is said to be established between the two end systems. The Internet's connection oriented service comes with ***reliable data transfer***, ***flow control*** and ***congestion control***.

- ***Reliable data transfer***, mean that an application can rely on the connection to deliver all of its data without error and in the proper order. Reliability is achieved through the use of acknowledgments and retransmissions**.**

  Ex:

  When end system B receives a packet from A, it sends an acknowledgment; when end system A receives the acknowledgment, it knows that the corresponding packet has definitely been received. When end system A doesn't receive an acknowledgment, it assumes that the packet it sent was not received by B; it therefore retransmits the packet.

- ***Flow control*** makes sure that neither side of a connection overwhelms the other side by sending too many packets too fast (one side is faster than the other). The flow-control service forces the sending end system to reduce its rate whenever there is such a risk.

- ***Congestion control*** service helps prevent the Internet from entering a state of grid lock. ***When a router becomes congested, its buffers can overflow and packet loss can occur***.

In such circumstances, if every pair of communicating end systems continues to pump packets into the network as fast as they can, gridlock sets in and few packets are delivered to their destinations. The *Internet avoids* this problem by forcing end systems to diminish the rate at which they send packets into the network during periods of congestion. End systems are alerted to the existence of severe congestion when they stop receiving acknowledgments for the packets they have sent.

B) ***Connectionless service*** (UDP: User Datagram Protocol): There is no handshaking with the Internet's connectionless service. When one side of an application wants to send packets to another side of an application, the sending application simply sends the packets. Since there is no handshaking procedure prior to the transmission of the packets, data can be delivered faster. But there are no acknowledgments either, so a source never knows for sure which packets arrive at the destination. Moreover, the service makes no provision for flow control or congestion control.

- UDP - User Datagram Protocol [RFC 768]: Internet's connectionless service
  - unreliable data transfer
  - no flow control
  - no congestion control
- TCP service [RFC 793]: Internet's connection service
  - *reliable, in-order* byte-stream data transfer, when loss acknowledgements and retransmissions
  - *flow control:*
    - sender won't overwhelm receiver
  - *congestion control:*
    - senders "slow down sending rate" when network congested

App's using TCP:

- HTTP (WWW), FTP (file transfer), Telnet (remote login), SMTP (email)

App's using UDP:

- streaming media, teleconferencing, Internet telephony

3. Interconnecting LANs or WANs

LANs or WANs do not normally operate in isolation. They are connected to one another or to the Internet by using connecting devises. Connecting devices can operate in different layers of internet module (see figure 3.1). Three kinds of connecting devises will be discussed Hubs, Bridges, and Switches.

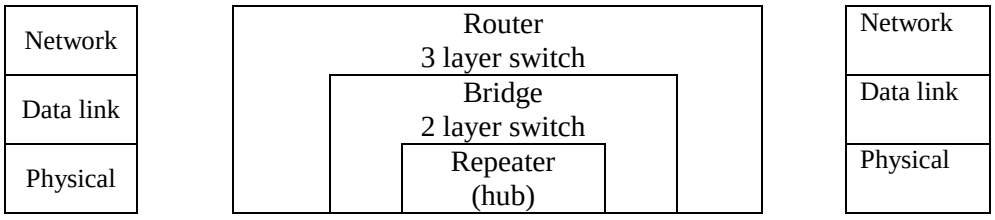

Fig. 3.1. Connecting devices.

3.1 HUB (repeater)

The simplest way to interconnect LANs is to use a hub (multi-port repeater). ***Hubs overcome the 10BASE5(***The *10* refers to its transmission speed of 10 Mbit/s, had a maximum length of 500 meters, maximum number of nodes that can be connected to a 10BASE5 segment is 100) ***Ethernet length restriction.*** Signals that carry information within a net can travel a fixed distance before attention endangers the integrity of the data. A repeater receives a signal and before it becomes too weak or corrupted, regenerates the original bit pattern. The repeater then

sends the refreshed signal. Repeater signal can extend the physical length of a LAN. Hubs are essentially repeaters, operating on bits. They are thus physical-layer devices. When a bit comes into a hub interface, the hub simply broadcasts the bit on all the other interfaces.

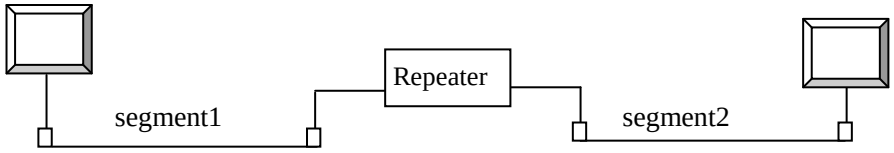

Repeaters actually connect 2 LAN segments (i.e. each connected LAN referred to as LAN segment). Repeaters forward every bit, it has no filtering capability. Repeater does not amplify the signal, it regenerate the signal. When it receives a weak or corrupted signal, it creates a copy, bit by bit, at the original strength. The repeater must be placed so that a signal reaches it before any noise changes its meaning or any of its bits. Hubs can be arranged in a hierarchy (or multi-tier design), with backbone hub at its top

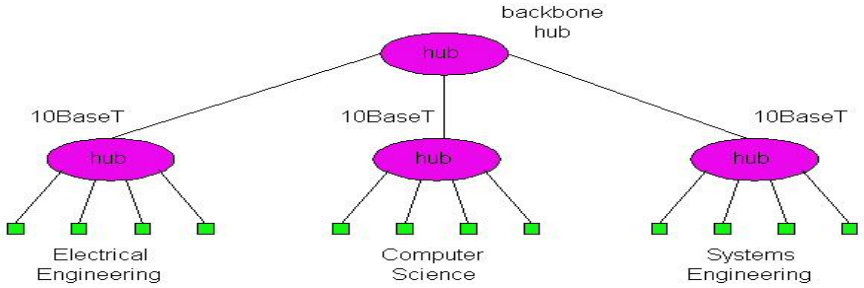

Fig 3.2 LAN: multi-tier hub design (two tiers or more), nodes connected to each hub is LAN segment.

Hubs do not isolate collision domains: node may collide with any node residing at any segment in LAN.

Hub Advantages:

- simple, inexpensive device
- Multi-tier provides graceful degradation: portions of the LAN continue to operate if one hub malfunctions
- extends maximum distance between node pairs (100m per Hub)

Hub limitations:

- single collision domain results in no increase in max throughput
    - multi-tier throughput same as single segment throughput
- Same collision domain and the maximum aggregate throughput are reduced to 10 Mbps. Before interconnecting the three departments, each departmental LAN had a maximum throughput of 10 Mbps, so that maximum aggregate throughput of the three LANs was 30 Mbps. But once the three LANs are interconnected with a hub, all of the hosts in the three departments belong to the same collision domain, and the maximum aggregate throughput is reduced to 10 Mbps.
- Cannot connect different Ethernet types (e.g., 10BaseT and 100BaseT) why? Since hubs are essentially repeaters and do not buffer frames, they cannot interconnect LAN segments operating at different rates.

- Each of the Ethernet technologies (10Base2, 10BaseT, 100BaseT, etc.) has restrictions on the maximum number of nodes that can be in a collision domain, the maximum distance between two hosts in a collision domain, and the maximum number of tiers that can be present in a multi-tier design. These restrictions constrain both the total number of hosts that connect to a multi-tier LAN as well as geographical reach of the multi-tier LAN.

3.2 Bridges (two layer switches)

Operates on both physical (regenerates the signal it receives), and data Link Layer devices (check the physical MAC address source and destination contained in the frame). It operates on Ethernet frames, examining frame header and selectively forwarding frame based on its destination.

- Bridge isolates (overcome) collision domains since it buffers frames. When frame is to be forwarded on segment, bridge uses CSMA/CD (Carrier Sense Multiple Access with Collision Detection) to access segment and transmit.
- Bridge has filtering capability, it can check the destination address of a frame and decide if the frame should be forwarded or dropped. If the frame should be forwarded, the decision must specify the port. Bridge has a table that maps address to ports.

Bridge advantages: Bridges can overcome many of the problems that plague hubs.

- Bridges permit inter-departmental communication while preserving isolated collision domains for each of the departments. Isolates collision domains resulting in higher total max throughput, and does not limit the number of nodes nor geographical coverage.
- Can connect different type Ethernet since it is a store and forward device bridges can interconnect different LAN technologies, including 10 Mbps and 100 Mbps Ethernets.
- There is no limit to how big a LAN can be when bridges are used to interconnect LAN segments: in theory, using bridges, it is possible to build a LAN that spans the entire globe.

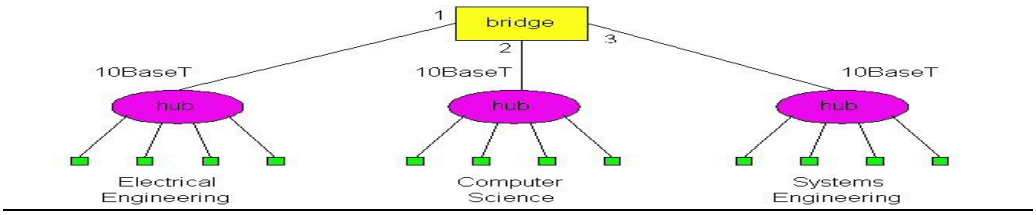

Fig. 3.3 Backbone Bridge (each LAN segment is now an isolated collision domain.)

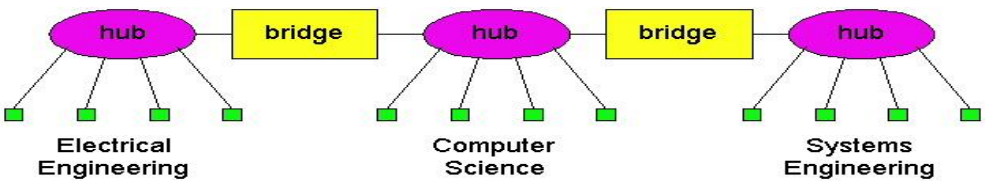

Fig 3.4 Interconnection without Backbone

Not recommended for two reasons:
- Single point of failure at Computer Science hub
- All traffic between EE and SE must path over CS segment

3.2.1 Bridge Forwarding and Filtering

Filtering is the ability to determine whether a frame should be forwarded to an interface or should just be dropped. When the frame should be forwarded, **forwarding** is the ability to determine which of the interfaces the frame should be directed to. Bridge filtering and forwarding are done with a **bridge table**. For each node on the LAN, the bridge table contains (1) the LAN address of the node, (2) the bridge interface that leads towards the node, (3) and the time at which the entry for the node was placed in the table. An example Table for the LAN in Figure 3.3 is shown below. We note here that the addressees used by bridges are physical addresses (not network addresses).

Bridge table 3.1 (for fig 3.3)

Address	**Interface**	**Time**
62-FE-F7-11-89-A3	**1**	9:32
7C-BA-B2-B4-91-10	**3**	9:36

To understand how bridge filtering and forwarding works, suppose a frame with destination address DD-DD-DD-DDDD-DD arrives to the bridge on interface x (in-port). The bridge indexes its table with the LAN address DD-DD-DD-DD-DDDD and finds the corresponding interface y (out-port).

a) If x equals y, then the frame is coming from a LAN segment that contains adapter DD-DD-DD-DD-DD-DD. There being no need to forward the frame to any of the other interfaces, the bridge performs the filtering function by discarding the frame.
b) If x does not equal y, then the frame needs to be routed to the LAN segment attached to interface y. The bridge performs its forwarding function by putting the frame in an output buffer that precedes interface y.

These simple rules allow a bridge to preserve separate collision domains for each of the different LAN segments connected to its interfaces. The rules also allow the nodes on different LAN segments to communicate.

Ex: Let's walk through these rules for the network in Figures 3.3 and its bridge table in 3.1.

a) Suppose that a frame with destination address 62-FE-F7-11-89-A3 arrives to the bridge from interface 1. The bridge examines its table and sees that the destination is on the LAN segment connected to interface 1 (i.e., the Electrical Engineering LAN). This means that the frame has already been broadcast on the LAN segment that contains the destination. The bridge therefore filters (i.e., discards) the frame.
b) Now suppose a frame with the same destination address arrives from interface 2. The bridge again examines its table and sees that the destination is the direction of interface 1; it therefore forwards the frame to the output buffer preceding interface 1.

It should be clear from this example that as long as the bridge table is complete and accurate, the bridge isolates the departmental collision domains while permitting the departments to communicate.

Recall that when a hub (or a repeater) forwards a frame onto a link, it just sends the bits onto the link without bothering to sense whether another transmission is currently taking place on the link. In contrast, when a bridge wants to forward a frame onto a link, it runs the CSMA/CD algorithm.

In particular, the bridge refrains from transmitting if it senses that some other node on the LAN segment is transmitting; furthermore, the bridge uses exponential back-off when one of its transmissions results in a collision. Thus bridge interfaces behave very much like node adapters. But technically speaking, they are *not* node adapters because neither a bridge nor its interfaces have LAN addresses. Recall that a node adapter always inserts its LAN address into the source

address of every frame it transmits. This statement is true for router adapters as well as host adapters. A bridge, on the other hand, does not change the source address of the frame.

3.2.2 Self-Learning

A bridge has the *property* of building its table automatically, dynamically and autonomously, without any intervention from a network administrator or from a configuration protocol. In other words, bridges are **self-learning.** This is accomplished as follows.

1. The bridge table is initially empty.
2. When a frame arrives on one of the interfaces and the frame's destination address is not in the table, then the bridge forwards copies of the frame to the output buffers of all of the other interfaces. (At each of these other interfaces, the frame accesses the LAN segment using CSMA/CD.)
3. For each frame received, the bridge stores in its table:
    a. The LAN address in the frame's *source address field,*
    b. The interface from which the frame arrived,
    c. The current time.

    In this manner the bridge records in its table the LAN segment on which the sending node resides. If every node in the LAN eventually sends a frame, then every node will eventually get recorded in the table.
4. When a frame arrives on one of the interfaces and the frame's destination address is in the table, then the bridge forwards the frame to the appropriate interface.
5. The bridge deletes an address in the table if no frames are received with that address as the source address after a period of time (the *aging time*). In this manner, if a PC is replaced by another PC (with a different adapter), the LAN address of the original PC will eventually be purged from the bridge table.

The self-learning property for the network in Figures 3.3 and its corresponding bridge table in table 3.2. Suppose at time 9:39 a frame with source address 01-12-23-34-45-56 arrives from interface 2. Suppose that this address is not in the bridge table. Then the bridge appends a new entry in the table.

Continuing with this same example, suppose that the aging time for this bridge is 60 minutes and no frames with source address 62-FE-F7-11-89-A3 arrive to the bridge between 9:32 and 10:32. Then at time 10:32 the bridge removes this address from its table. Table 3.2. Bridge learns about the location of adapter with address 01-12-23-34-45-56.

Address	Interface	Time
01-12-23-34-45-56	**2**	9:39
62-FE-F7-11-89-A3	**1**	9:32
7C-BA-B2-B4-91-10	**3**	9:36

Bridge Learning: example

Suppose C sends frame to D and D replies back with frame to C

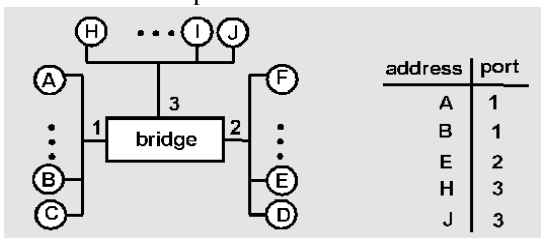

- C sends frame, bridge has no info about D, so floods to both LANs
  - bridge notes that C is on port 1
  - frame ignored on upper LAN
  - frame received by D

Bridge Learning: example

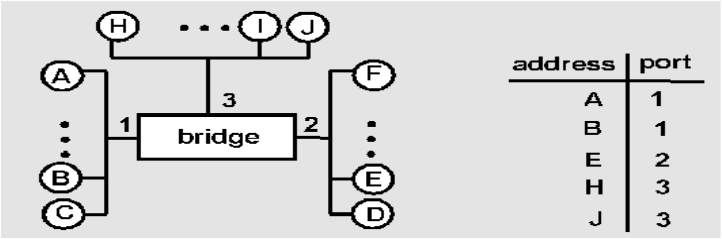

- D generates reply to C, sends
  - bridge sees frame from D
  - bridge notes that D is on interface 2
  - bridge knows C on interface 1, so *selectively* forwards frame out via interface 1

Bridges are **plug and play** also referred as **devices transparent bridges.** because they require absolutely no intervention from a network administrator or user. The administrator does not have to configure the bridge tables at the time of installation or when a host is removed from one of the LAN segments.

3.2.3 Spanning Tree

One of the problems with a pure hierarchical design for interconnected LAN segments is that if a hub or a bridge near the top of the hierarchy fails, then much (if not all) of the interconnected LAN will go down. For this reason it is desirable to build networks with multiple paths between LAN segments. An example of such a network is shown in Figure 3.5

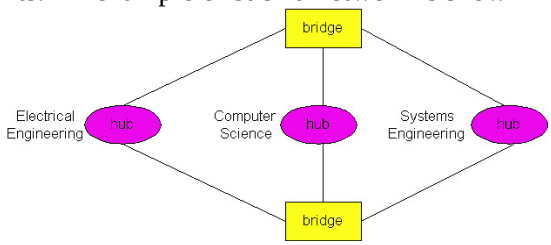

Figure 3.5: Interconnected LAN segments with redundant paths.

Multiple redundant paths between LAN segments (such as departmental LANs) can greatly improve fault tolerance. But, it ***has a side effect*** -- frames cycle and multiply within the interconnected LAN, thereby crashing the entire network. To see this,

1. suppose that the bridge tables in Figure 3.5 are empty, and a host in Electrical Engineering sends a frame to a host in Computer Science.
2. When the frame arrives to the Electrical Engineering hub, the hub will generate two copies of the frame and send one copy to each of the two bridges.
3. When a bridge receives the frame, it will generate two copies, send one copy to the Computer Science hub and the other copy to the Systems Engineering hub. Since both bridges do this, there will be four identical frames in the LAN. This multiplying of copies will continue indefinitely since the bridges do not know where the destination host resides. (To route the frame to the destination host in Computer Science, the destination host has to first generate a frame so that its address can be recorded in the

bridge tables.) The number of copies of the original frame grows exponentially fast, crashing the entire network.

To prevent the cycling and multiplying of frames, bridges use a spanning tree protocol. In the **spanning tree protocol**, bridges communicate with each other over the LANs in order to determine a spanning tree, that is, a subset of the original topology that has no loops. Once the bridges determine a spanning tree, the bridges disconnect appropriate interfaces in order to create the spanning tree out of the original topology.

For example, in Figure 3.5, a spanning tree is created by having the top bridge disconnect its interface to Electrical Engineering and the bottom bridge disconnect its interface to Systems Engineering. With the interfaces disconnected and the loops removed, frames will no longer cycle and multiply.

If, at some later time, one of links in the spanning tree fails, the bridges can reconnect the interfaces, run the spanning tree algorithm again, and determine a new set of interfaces that should be disconnected.

3.2.4 Bridges versus Routers

- both store-and-forward devices
  - routers: network layer devices (examine network layer headers)
  - bridges are Link Layer devices
- routers maintain routing tables, implement routing algorithms
- bridges maintain filtering tables, implement filtering, learning and spanning tree algorithms
- bridges do well in small (few hundred hosts) while routers used in large networks (thousands of hosts)

Bridges + and -

+Bridge operation is simpler requiring less processing

- Topologies are restricted with bridges: a spanning tree must be built to avoid cycles
- Bridges do not offer protection from broadcast storms (endless broadcasting by a host will be forwarded by a bridge)

Routers + and -

+ arbitrary topologies can be supported, cycling is limited by TTL counters (and good routing protocols)

+ provide firewall protection against broadcast storms

- require IP address configuration (not plug and play)

- require higher processing

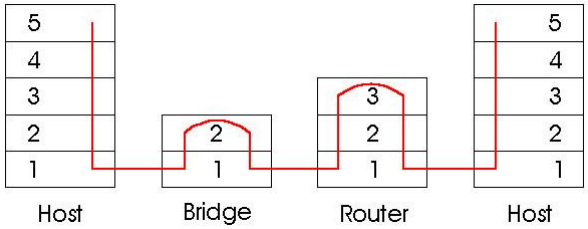

Figure 3.6: Packet processing and bridges, routers and hosts.

3.3 Ethernet Switches

Another interconnection device became widely available, namely, Ethernet switches. As do bridges, they forward and filter frames using LAN destination addresses, and they automatically

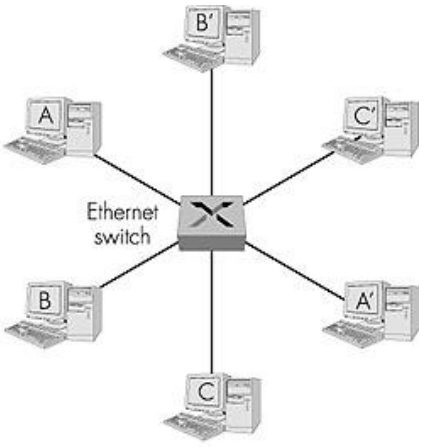

build routing tables using the source addresses in the traversing frames.

The most important difference between a bridge and switch is that bridges usually have a small number of interfaces (i. e., 2-4), whereas switches may have dozens of interfaces.

- layer 2 (frame) forwarding, filtering using LAN addresses
- Switching: A-to-B and A'-to-B' simultaneously, no collisions. If each host has a 10Mbps adapter card, then the aggregate throughput during the three simultaneous file transfers is 30 Mbps. If A and A' have 100 Mbps adapters and the remaining hosts have 10 Mbps adapters, then the aggregate throughput
- during the three simultaneous file transfers is 120 Mbps.
- large number of interfaces which generates a high aggregate forwarding rate through the switch fabric, therefore necessitating a high-performance design.
  Figure 3.7 Ethernet switch

- Switches can be purchased with various combinations of 10 Mbps, 100 Mbps and 1 Gbps interfaces. For example, you can purchase switches with four 100 Mbps interfaces and twenty 10 Mbps interfaces; or switches with four 100 Mbps interfaces and one 1 Gbps interface.
- Many switches also operate in a full-duplex mode; that is, they can send and receive.
- frames at the same time over the same interface.
- often: individual hosts, star-connected into switch
  - Ethernet, but no collisions!
  - Creates direct connections between hosts and the switch. When a host has a full-duplex direct connection to a switch, it can transmit
  - (and receive) frames at the full transmission rate of its adapter

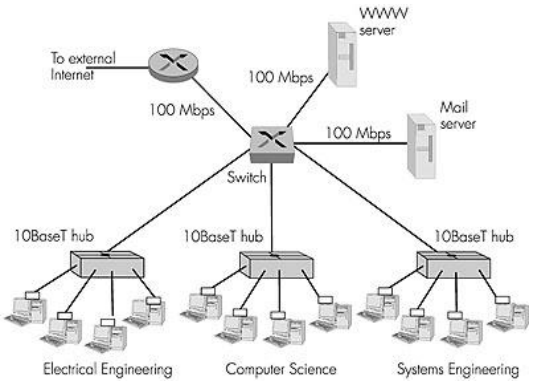

Figure 3.8:

- switches use **cut-through switching** rather than store-and-forward packet switching, used by routers and bridges
- cut-through switching: frame forwarded from input to output port without awaiting for assembly of entire frame
  - slight reduction in latency
- combinations of shared/dedicated, 10/100/1000 Mbps interfaces

3.4 Routers (three layer switches)

- Routers is a three-layer device. It operates in:
  - Physical layer: regenerate the signal it receives.
  - Data link layer: checks the (source and destination) contained in the packet.
  - Network layer: check s the network layer address (address in IP layer).
- Router can connects independent LANs or WANs to create the internetwork (internet).
- Router acts like a station on a network, but unlike any station, router has addresses on and links to two or more networks.
- The major differences between a router and repeater or bridge:
  - Router has physical and logical (IP) address for each interface.
  - A router works only on those packets in which the physical destination address matches the address of the interface at which the packet arrives.
  - A router changes the physical address of the packet (both source and destination) when it forward the packet.

  Ex:

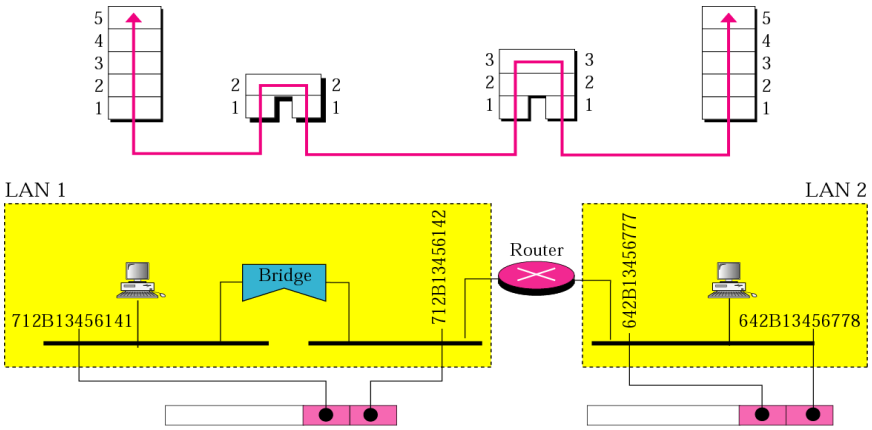

Figure 3.9: two LANs Separated by router

The router changes the source and destination address of the packet.

- When packet travel in the left LAN, its source physical address is the address of the sending station; its destination is the address of the router.
- When packet travel in the right LAN, its source physical address is the router address and its destination address is the address of the final destination station.

Comparison of the typical features of popular interconnection devices.

	hubs	**bridges**	**Routers**	**switches**
traffic isolation	*no*	*yes*	*Yes*	*yes*
plug and play	**yes**	*yes*	*No*	*yes*
optimal routing	*no*	*no*	*Yes*	*no*
cut-through	*yes*	*no*	*No*	*yes*

4 Access Networks and Physical Media

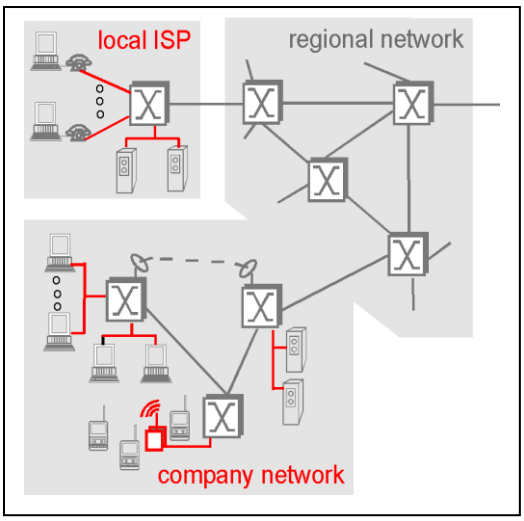

How to connect end systems to edge router?

Access networks can be loosely divided into three categories:

- Residential access nets: connecting a home end system into the network;
- Institutional access networks connecting school, company.
- Mobile access networks: connecting a mobile end system into the network

Keep in mind:

- Bandwidth (bits per second) of access network?
- Shared or dedicated?

4.1 Residential Access: point to point access

A residential access network connects a home end system (a PC, perhaps a Web TV or other residential system) to an edge router. Most common form of home access is using a **modem** over a POTS (plain old telephone system) dialup line to an Internet service provider (ISP). In this case, the "access network" is simply a point-to-point dialup link into an edge router (twisted-pair phone line)

- Dialup via modem
  - up to 56Kbps direct access to router (often less)
  - Can't surf and phone at same time: can't be "always on"

  Telephone lines, but can transmit at rates of up to about 8 Mbps from the ISP router to a home end system. The data rate in the reverse direction is less than 1 Mbps. i.e.
  - up to 1 Mbps upstream (today typically < 256 kbps)
  - up to 8 Mbps downstream (today typically < 1 Mbps)
- ADSL (asymmetric digital subscriber line): is conceptually similar to dialup modems. It is a modem technology running also over existing twisted pair.

  ADSL divides the communication link between the home the ISP into three non-overlapping frequency bands:
  - A high-speed downstream channel, in the 50 KHz to 1 MHz band;
  - A medium-speed upstream channel, in the 4 KHz to 50 KHz band;
  - Ordinary POTs two-way telephone channel, in the 0 to 4 KHz band.

One of the features of ADSL is that the service ***allows the user to make an ordinary telephone call, using the POTs channel, while simultaneously surfing the Web***.

- HFC (hybrid fiber coaxial cable) a cable head end station broadcasts through a distribution of coaxial cable and amplifiers to residences.
  - ✓ asymmetric: up to 30Mbps downstream, 2 Mbps upstream

  In particular, every packet sent by the head-end travels downstream on every link to every home; and every packet sent by a home travels on the upstream channel to the Head-end. ***HFC available via cable TV companies***

Network of cable and fiber attaches homes to ISP router, homes share access to router.

-For this reason, if several users are receiving different Internet videos on the downstream channel, actual rate at which ***each user receives its video will be significantly less than downstream rate***.

-On the other hand, if all the active users are Web surfing, then each of the users may actually receive Web pages at the full downstream rate, as a small collection of users will rarely receive a Web page at exactly the same time.

-Because the upstream channel is also shared, packets sent by two different homes at the same time will collide, which further decreases the effective upstream bandwidth.

ADSL is a point-to-point connection between the home and ISP, and therefore ***all the ADSL bandwidth is dedicated rather than shared***. Cable advocates, however, argue that a reasonably dimensioned HFC network provides higher bandwidths than ADSL**.**

Both ADSL, and HFC requires special modems, called cable modems

4.2 Enterprise Access Networks

- In enterprise access networks, a local area network (LAN) is used to connect an end system to an edge router.
- There are many different types of LAN technology. However, Ethernet technology is currently by far the most prevalent access technology in enterprise networks. **Ethernet** operates 10 Mbps or 100Mbps (and now even at 1 Gbps). Uses either twisted-pair copper wire or coaxial cable.
- The edge router is responsible for routing packets that have destinations outside of that LAN.
- Ethernet uses a shared medium, so that end user shares the transmission rate of the LAN. More recently, shared Ethernet technology has been migrating towards switched Ethernet technology.
- Switched Ethernet uses multiple coaxial cable or twisted pair Ethernet segments connected at a "switch" to allow the full bandwidth an Ethernet to be delivered to different users on the same LAN simultaneously.

4.3 Mobile Access Networks

- Use the radio spectrum to connect a mobile end system (e.g., a laptop PC or a PDA with a wireless modem) to a base station,
- This base station, in turn, is connected to an edge router of a data network.
- shared *wireless* access network connects end system to router via base station "access point"

5 Typical home network components:

- ADSL or cable modem
- router/firewall/NAT
- Ethernet
- wireless access point

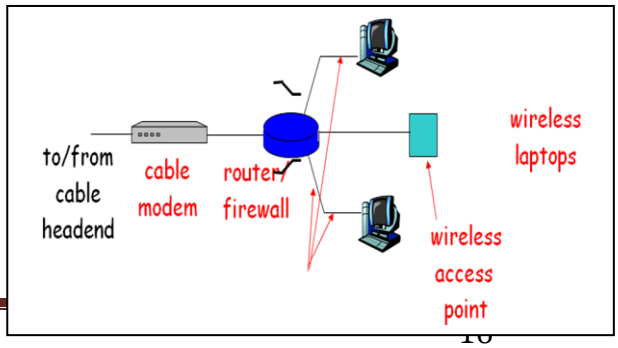

5.1 Physical Media

- **Bit**: propagates between transmitter/receiver pairs
- **physical link**: what lies between transmitter & receiver
- **guided media**: signals propagate in solid media: copper, fiber, coax
- **unguided media**: signals propagate freely, e.g., radio

Physical media include:

1. Twisted Pair (TP) two insulated copper wires.
2. Coaxial cable: two concentric copper conductors, bidirectional
    a. baseband: single channel on cable
    b. broadband: multiple channel on cable

3. Fiber optic cable: glass fiber carrying light pulses, each pulse is a bit
    a. high-speed operation:, high-speed point-to-point transmission (e.g., 5 Gps)
    b. low error rate: repeaters spaced far apart; immune to electromagnetic noise
4. Physical media: radio
    a. signal carried in electromagnetic spectrum
    b. no physical "wire"
    c. bidirectional
    d. propagation environment effects: reflection, obstruction by objects, interference

Radio link types:

- terrestrial microwave : ex. up to 45 Mbps channels
- LAN (e.g., Wifi), 2Mbps, 11Mbps
- wide-area (e.g., cellular) e.g. 3G: hundreds of kbps
- satellite up to 50Mbps channel (or multiple smaller channels)

6. Internet structure: network of networks

Internet structure roughly hierarchical,

- At center: "tier-1" ISPs, national/international coverage, treat each other as equals.

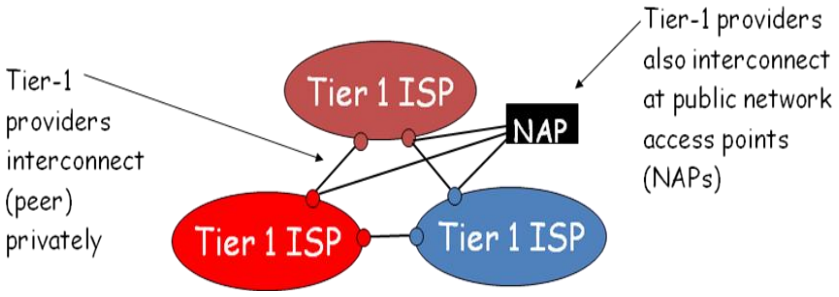

- "Tier-2" ISPs: smaller (often regional) ISPs, Connect to one or more tier-1 ISPs, possibly other tier-2 ISPs.

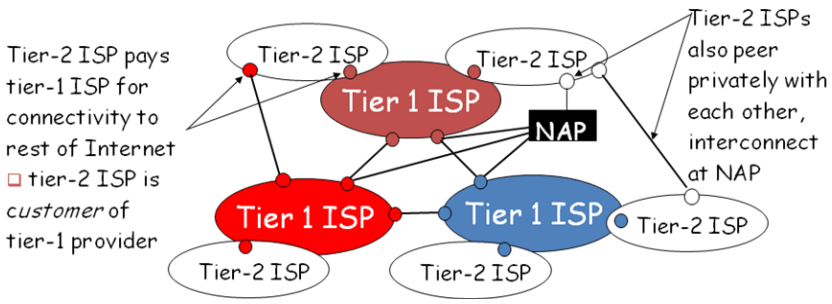

- "Tier-3" ISPs and local ISPs, last hop ("access") network (closest to end systems)

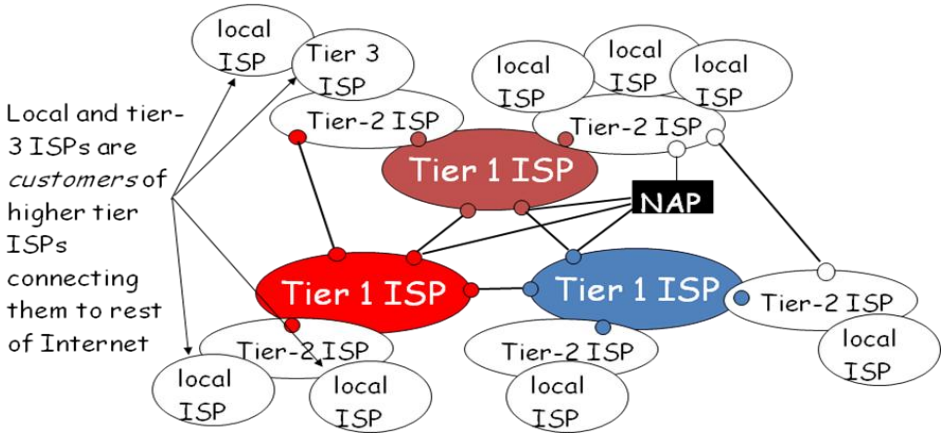

A Packet Passes Through Many Networks

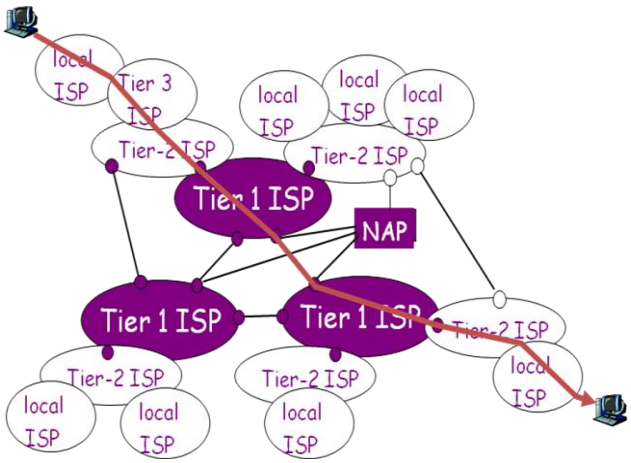

7. Protocol Layers and Their Service Models

A layered architecture allows us to discuss a well-defined, specific part of a large and complex system. ***Each layer implements a service via its own internal-layer actions, or relying on services provided by layer below.***

Why layering?

Dealing with complex systems:

- explicit structure allows identification, relationship of complex system's pieces, layered reference model for discussion.

- modularization eases maintenance, updating of system (change of implementation of layer's service transparent to rest of system e.g., change in gate procedure doesn't affect rest of system)

layering considered harmful?

- one layer may duplicate lower-layer functionality. For example, many protocol stacks provide error recovery on both a link basis and an end-to-end basis.
- A second potential drawback is that functionality at one layer may need information (e.g., a timestamp value) that is present only in another layer; this violates the goal of separation of layers.

7.1 Internet Protocol Stack

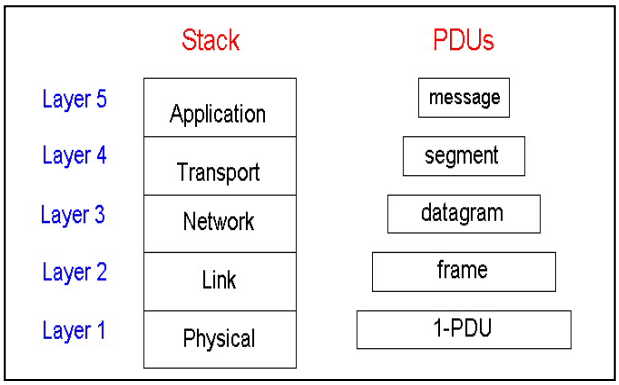

Application: supporting network applications, support may protocols:

- FTP, SMTP, STTP

- Transport: host-host data transfer (transporting application-layer messages between the client and server sides .transport protocols:
  - TCP provides a connection-oriented service to its applications. TCP also segments long messages into shorter segments and provides a congestion control mechanism.
  - UDP provides its applications a connectionless service,
- Network: routing of datagrams from source to destination
  - IP, routing protocols that determine the routes that datagrams take between sources and destinations. Network administrator can run any routing protocol desired.
- Link: data transfer between neighboring network elements. To move a packet from one node (host or packet switch) to the next node in the route, the network layer must rely on the services of the link layer. In particular, at each node IP passes the datagram to the link layer, which delivers the datagram to the next node along the route. At this next node, the link layer passes the IP datagram to the network layer. Ex: PPP, Ethernet

  As datagrams typically need to traverse several links to travel from source to destination, a datagram may be handled by different link-layer protocols at different links along its route. For example, a datagram may be handled by Ethernet on one link and then PPP on the next link.
- Physical: bits "on the wire"

Encapsulation