CT2108 – Networks and Data Communications 1

Dr Des Chambers

Introduction

1

Content

- Computer Networks vs. Distributed Systems
- Uses of computer networks
- Network Hardware
- Network Software
- Network Technologies

Computer Networks

- A Computer Network is a collection of autonomous devices interconnected by some type of network technology
- Networks come in many shapes and sizes
- The Internet is a network of networks
  - World Wide Web is not a physical network. It is a distributed Client-Server application that runs over the Internet

3

Two computers are said to be INTERCONNECTED if they are able to exchange information. The connection can be copper wire, optical fiber, wireless, etc…

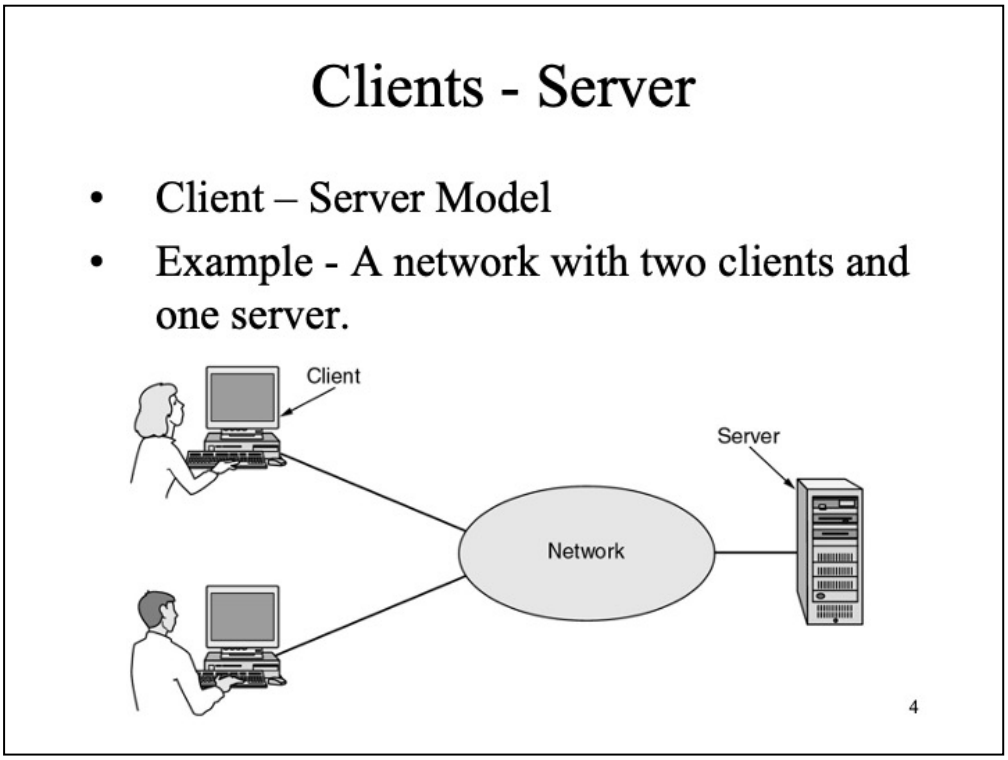

CLIENT – SERVER model is employed in most network applications. The Server is a powerful machine that can have multiple concurrent clients accessing its resources at the same time. Clients are usually simpler devices that run apps to interpret or display information provided by a server.

Client - Server

- The client-server model involves requests and replies.

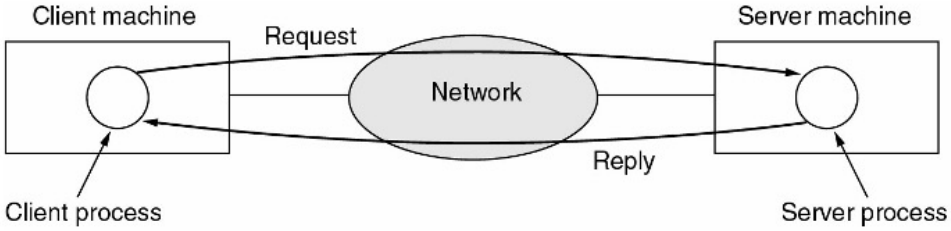

- Client – Server model employs at least two processes: one running on the server and one running on the client

5

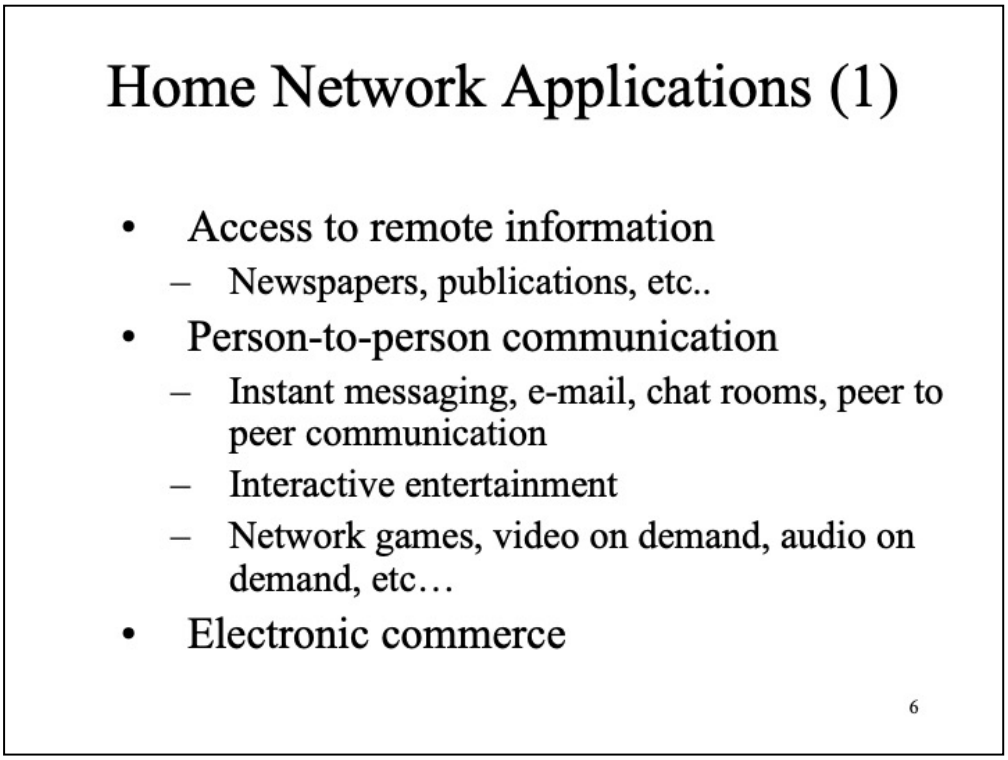

KEN OLSEN – president of Digital (second big computer vendor in the world, after IBM, in 1977) said " There is no reason for any individual to have a computer in his home". History proved otherwise, and Digital went out of business in the 1990s.

Home Network Applications (2)

- In peer-to-peer system there are no fixed clients and servers.

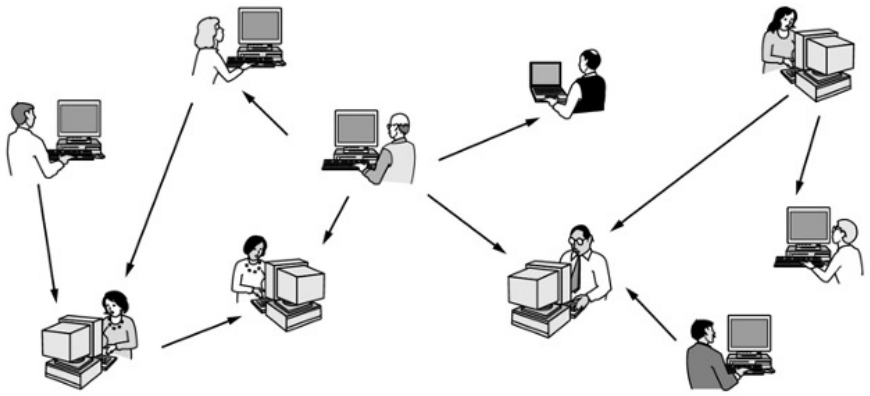

Newer peer to peer systems (like BitTorrent) don't have a centralized db. Lookup of the content comes from a local db-s maintained by each of the members. Besides content, each user maintains a list of other users as well.

Network Types

- Local Area Networks
- Metropolitan Area Networks
- Wide Area Networks
- Wireless Networks
- Home Networks
- The Internet

8

Network Types

- Classification of interconnected processors by scale.

Interprocessor distance	Processors located in same	Example
1 m	Square meter	Personal area network
10 m	Room	Local area network
100 m	Building	
1 km	Campus	
10 km	City	Metropolitan area network
100 km	Country	Wide area network
1000 km	Continent	
10,000 km	Planet	The Internet

9

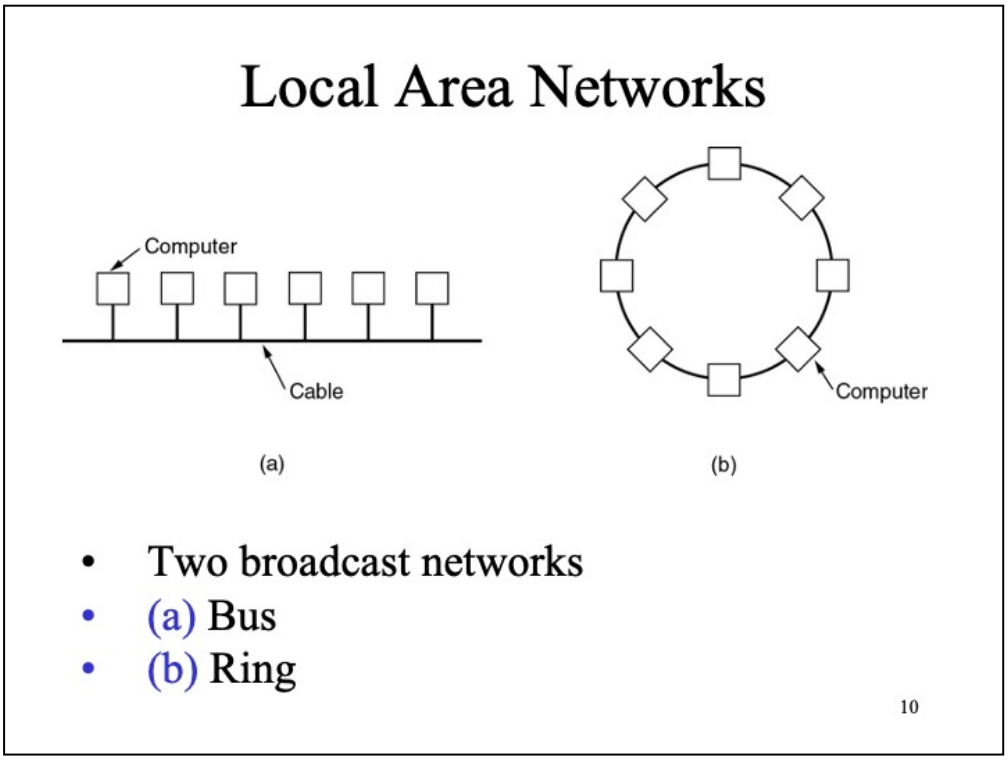

Local Area Networks are privately owned networks within a single building or camps, up to a few km in size. Are restricted in size, which means that worst case transmission time is bounded and known in advance.

LANs use very often same cable to which all the machines are attached. Speeds ranging from 100Mb/s to about 10Gb/s. Various **TOPOLOGIES** are possible for broadcast LANs: **BUS**, **STAR** (most popular here is Ethernet) and **RING** are mostly used.

ETHERNET – bus-based network, bus and/or star topology, broadcast decentralized network, usually operating at 100Mb/s to 10Gb/s. Computers in Ethernet can transmit whenever they want; if two packets collide, each computer just waits a random time and tries again later.

Metropolitan Area Networks

- A metropolitan area network based on cable TV.

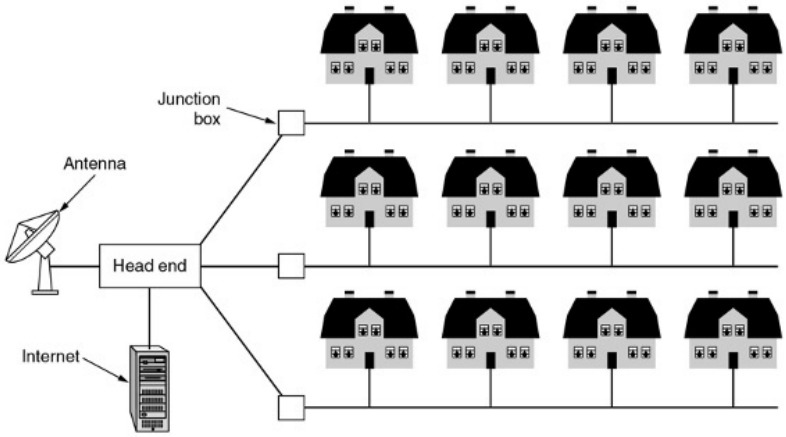

Metropolitan Area Network (MAN) covers a city. Best known example of a MAN is the cable TELEVISION NETWORK available in many cities. Until late 1990's they were intended for television only. After that, the cable providers realized that they could offer two-way Internet in the unused parts of the spectrum. This transformed the TV network in a metropolitan area network.

Wide Area Networks (1)

- Relation between hosts on LANs and the subnet.

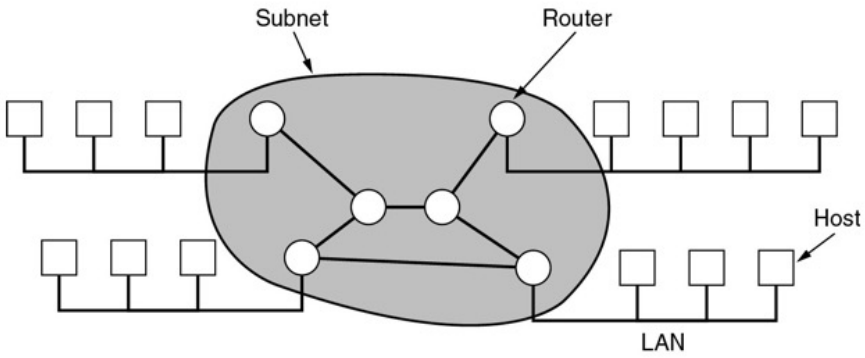

A Wide Area Network (WAN) spans over a large area, often a country or a continent. It contains a number of machines (called **HOSTS** in the networking context) that are connected by a communication **SUBNET**. The hosts are usually owned by people, while the subnet is owned by the telecom providers or Internet providers.

The job of the subnet is to carry messages from host to host. Separation of the pure communication aspects of the network (the subnet) from the application aspects (hosts) simplifies the complete network design.

Subnets contain two components:

•**TRANSMISSION LINES** – move bits between machines (copper, optical fiber, radio, etc)

•**SWITCHING COMPONENTS** – specialized computers that connect three or more transmission lines. When data comes on one of the lines, the switching element must choose an ongoing line to forward the data. **ROUTER** is the technical name for those switching elements.

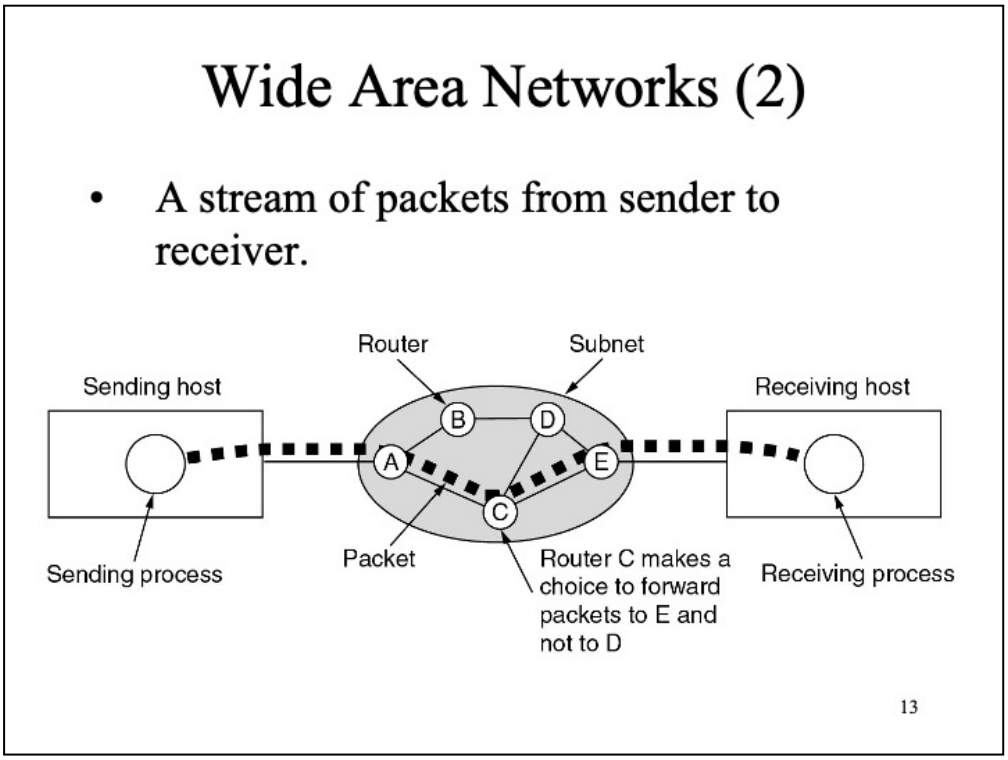

A **STORE-AND-FORWARD** or **PACKET SWITCHED** subnet is one where the packets are received entirely at intermediate routers, stored until some outgoing transmission line is free and then forwarded to the next router. When packets are small and all the same size, they are called **cells**.

When a process on a host wants to send a message to another host in the network, the sending host cuts the message into packets, each one caring some sort of sequence number. Those packets are then injected into the network, one at a time, in quick succession. The packets are delivered over the network and delivered to the receiving host, where they are reassembled and delivered to the receiving process.

In this figure, all the packets from sender to receiver followed same route ACE. In some subnets, the packets *must* follow always same path, in other subnets, the packets can follow different paths (they are routed separately). When the packet is getting to router A, the decision to follow path C or path B is made locally. This decision is made by A and how this decision is made is calling routing algorithm.

Wireless Networks

- (a) Bluetooth configuration
- (b) Wireless LAN

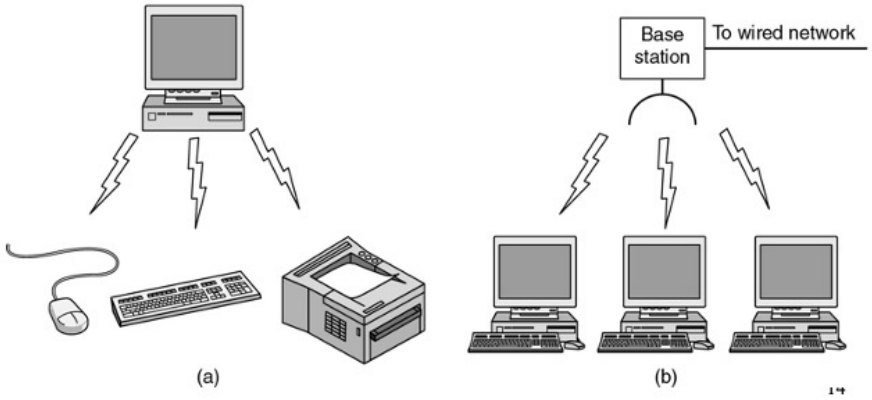

BLUETOOTH is an example of system interconnection network, and it refers to interconnecting computer components (monitor, mouse, keyboard, etc…). It is a master slave topology. The master tells the slaves what addresses to use, when they can broadcast, how long they can transmit, what frequencies they can use and other information.

WIRELESS LANs – each computer has a radio modem and antenna with which it can communicate with other systems. IEEE 802.11 is a basic standard for wireless LAN. A number of newer, derivate standards are in place now.

WIRELESS WANs – cellular phone networks: 3G/4G/5G

Network Software

- Protocol Hierarchies
- Design Issues for the Layers
- Connection-Oriented and Connectionless Services
- Service Primitives
- The Relationship of Services to Protocols

15

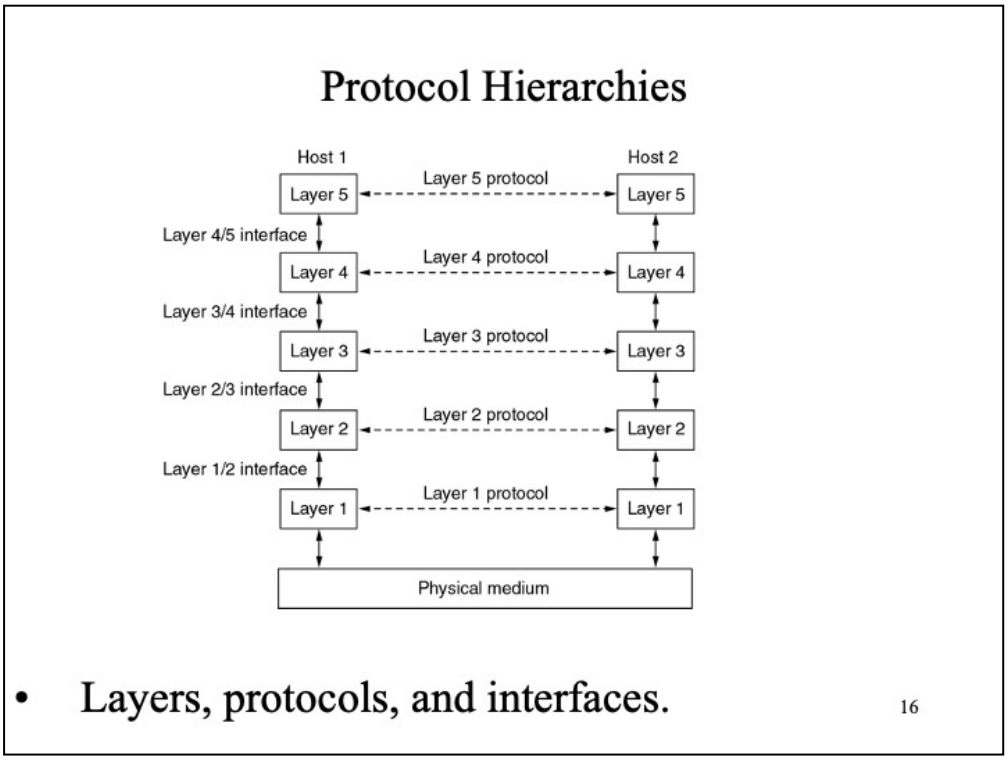

To reduce complexity of design, networks are organized as layer, each one build upon the one below it. The number of layers, the name of each layer, the contents and function of each layer differ from network to network.

The purpose of each layer is to create services for the layers above, hiding to those layers the details of how those services are actually implemented. The fundamental idea is that a particular piece of software (or even hardware) provides a service to its users but keeps the details of its internal state and algorithms hidden from them. Layer n on a machine carries a conversation with layer n on another machine. The rules and conventions used in this conversation are known as layer n protocol. In essence, a **PROTOCOL** is an agreement between the communicating parties on how communication is to proceed. The entities that implement the protocol at different layers level are called **PEERS**. It is peers that communicate using the protocol.

In reality, no data is directly transferred from layer n on one machine on layer n on the other machine. In effect, each layer passes data and control information to the layer below it, until the lowest layer is reached. Below Layer 1 is the **physical medium**, through the communication occurs. Between each pair of adjacent layers there is an interface. The interface defines which primitive operations and services the lower layer makes available to the upper one. Most difficult design issue is to define clean interfaces between layers. A set of layers and protocols is called a **NETWORK ARCHITECTURE** (it has to contain enough information to allow hardware and software engineers to design hardware and software that would obey the right protocol). A list of protocols used by certain systems, one protocol per layer, is called a **PROTOCOL STACK**.

Protocol Hierarchies (2)

- The philosopher-translator-secretary architecture.

17

Layer3 – two philosophers (peer processes), one speaks English and the other one speaks French. Since they have no common language, they engage a translator (peer processes at Layer 2). The translators, each contact a secretary (peer process Layer 1). Translators have agreed on a common language that both know (Dutch).

Note that each protocol is completely independent of the other ones. If at any time, the translators decide to change the language, all they have to do is to agree between each other. None of the interfaces with layer 3 or layer 1 will be changed. Similarly, the secretary could choose to use a different transmission medium (say e-mail), without disturbing or even informing the other layers.

Protocol Hierarchies (3)

- Example information flow supporting virtual communication in layer 5.

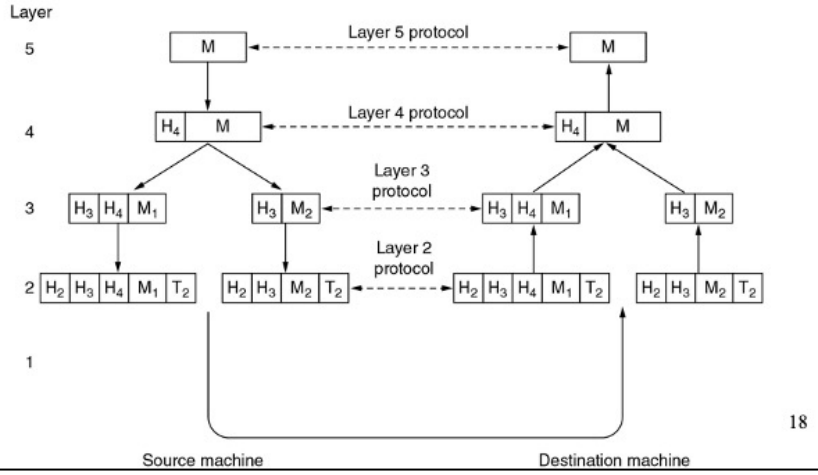

A message is produced by an application process running at layer 5. Message M is then given to layer 4 for transmission. Layer 4 puts a header in the front of the message to identify the message and pass the result to layer 3. The header includes control information, such as sequence numbers to allow layer 4 on destination to deliver messages in the right order, if the lower layers do not maintain sequence. In some layers, headers can also contain sizes, times and other control fields.

At layer 3 there is a limit on the size of the packet that can be transmitted. So layer 3 will break the incoming message into smaller parts, packets, adding header H3 corresponding to layer 3 on each packet. In this example, message M is split into M1 and M2.

Layer 3 decides which outgoing lines to use and passes the message to layer 2. Layer 2 adds not only a header to each piece, but also a trailer and gives the resulting units to layer 1 for physical transmission.

At the receiving machine, the message moves upwards, from layer to layer, with headers being stripped off as it progresses.

Design Issues for the Layers

- **Addressing**
  - consequence of having multiple destinations
- **Error Control**
  - The receiver should be able to inform sender which data was received correctly
- **Flow Control**
  - Keep sender from swamping slow receiver with data
  - Keep the sender from swamping with data slow networks
- **Multiplexing**
  - Use same communication channel for multiple, unrelated conversations
- **Routing**
  - When multiple paths between source and destination, one path must be chosen

19

Connection-Oriented and Connectionless Services

- Layers can offer two types of services to the layers above: connection oriented services and connection-less services
- Connection oriented services
  - Reliable message stream (sequence of pages)
  - Reliable byte stream (remote login, file transfer, etc..)
  - Unreliable connection (digitized voice or video)
- Connection-less
  - Datagram service (in analogy with telegram service)
  - Acknowledged datagram service
  - Request-reply service

20

CONNECTION ORIENTED SERVICE – modeled after the phone systems. The service users establishes a connection, uses the connection and then releases the connection. The main idea, is that the connection acts as a PIPE, at one end data is pushed and at the other end data is received. In most of the cases, the order is preserved. Sometime, during the connection establishment phase, a **NEGOTIATION** is employed (for establishing some parameters of the connection)

CONNECTION-LESS SERVICE – modeled after the postal system. Each message (letter) carries the full destination address, each one being routed through the system independent of the others. It is possible that the messages will arrive at the destination in out of order.

Each service is characterized by **QUALITY OF SERVICE**. Some services are reliable in the sense that they never loose data. Usually, reliability is implemented with acknowledgements from the receiver that it received data. This introduces overhead and delays in the communication, which sometime is OK, but sometime is not (in real time voice and video communication).

Services to Protocols Relationship

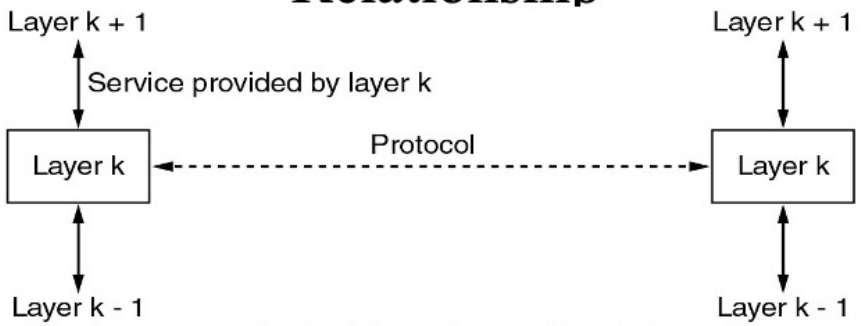

- Service – set of primitives (operations) that a layer provides to the layer above it; relate to the interfaces between layers
- Protocol – set of rules governing the format and meaning of packets exchanged by peer entities within a layer; relate to the packets that are sent between peer entities between different machines

21

Reference Models

- The OSI Reference Model
- The TCP/IP Reference Model
- A Comparison of OSI and TCP/IP
- A Critique of the OSI Model and Protocols
- A Critique of the TCP/IP Reference Model

22

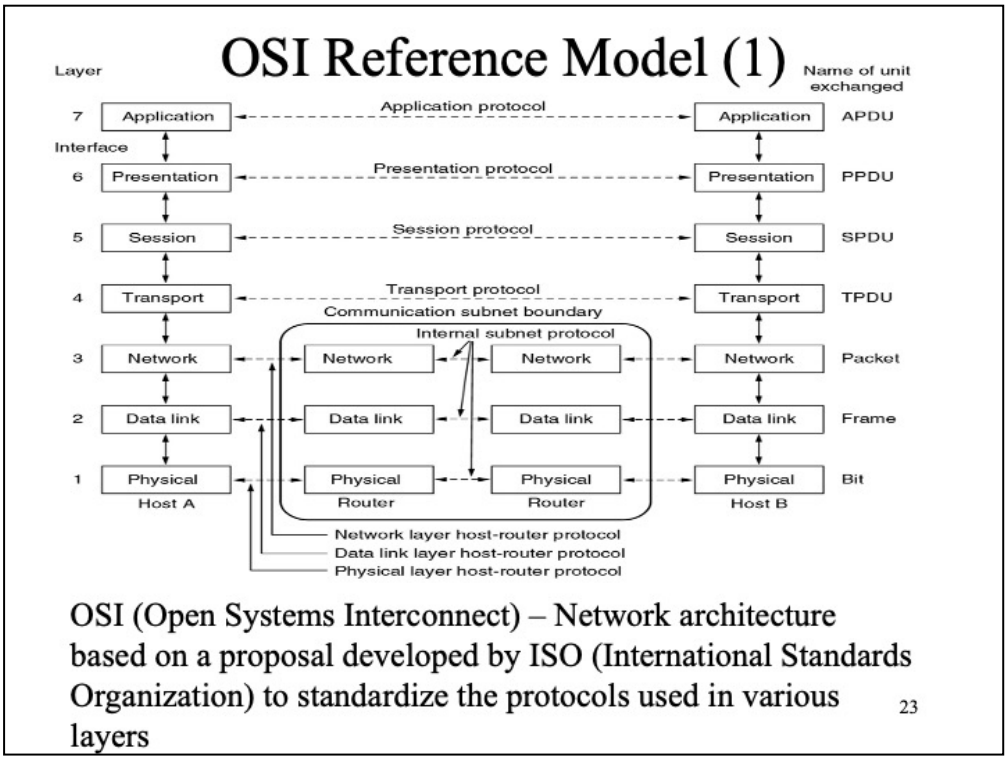

Design principles that lead to the seven-layer design are as follows:

1. A layer should be created where a **different abstraction is needed**
2. Each layer should perform a **well-defined function**
3. The function of each layer should be chosen with an eye toward defining **internationally standardized protocols**
4. The layer boundaries should be chosen to **minimize the information flow across the interfaces**
5. The number of layers should be large enough to avoid throwing together separate, **distinct functions** out of necessity and small enough to avoid inefficiency

OSI Reference Model (2)

- Physical Layer
  - Transmitting raw bits over communication channel
  - Typical questions that are addressed:
    - How many volts used to represent a "1" and how many for "0"
    - How many nanoseconds a bit last
    - Full duplex transmission or not (both directions)
    - How initial connection is established and how is torn down when both sides are finished
    - How many pins the network connector will have and what is each pin used for
  - Design issues – sending one bit "1" on one side has to get in the other side as "1" not as "0"
    - Mechanical, electrical and timing interfaces
    - Physical transmission medium

24

OSI Reference Model (3)

- Data Link Layer
  - Transform the raw transmission facility (offered by the physical layer) into a line that appears free of undetected transmission errors to the network layer
  - Design issues
    - **Error detection and correction**
      - The sender breaks up the input data into **data frames** (typically a few hundred or thousands bytes) and transmits the frames sequentially. If the service is reliable, the receiver has to confirm the correct receipt of each frame
    - **Flow control** – keep a fast transmitter drowning a slow receiver with data
      - Some traffic regulation mechanism is often needed to let the transmitter know how much buffer space the receiver has at the moment. Usually, this is integrated with the error handling mechanism
    - Broadcast networks have an additional issue in the data-link layer: how to control **access to the shared channel**. A special sub-layer of the data-link layer, the medium access control sub-layer deals with this problem

25

OSI Reference Model (4)

- Network Layer
  - Controls the operation of the subnet
  - Design issues
    - **Routing** - how packets are routed from source to destination
    - **Congestion control** – if too many packets are present in the subnet at the same time
    - Allow heterogeneous networks to be interconnected
    - In broadcast networks, the routing problem is thin or non existent

26

OSI Reference Model (5)

- Transport Layer
  - Accepts data from the above layer, split it into smaller units and pass them to the network layer. Ensures that those pieces arrive correctly at the other end
  - Determines what type of service to provide
    - Most popular type of transport connection is error free, point to point channel that delivers messages or bytes in the order in which they were sent
    - Other type of transport services: delivering datagrams with no guarantee about the order of delivery, broadcasting messages
  - It is a true end-to-end layer, all the way from source to the destination

27

The Transport layer has to perform its function in a way that isolates the upper layers from the inevitable changes in the hardware technology

Layers one to three are chained, while layers four to seven are END to END layers.

OSI Reference Model (6)

- Session Layer
  - Allows users on different machines to establish sessions between them
  - Sessions offer different services:
    - **Dialog control** – keeping track of those whose turn is to transmit
    - **Token management** – preventing two parties from attempting same critical operation at the same time
    - **Synchronization** – marking long transmissions to make sure they can be resumed from where they were when a crash happened

28

OSI Reference Model (7)

- Presentation Layer
  - It is not concerned with moving bits around, but with checking the syntax and semantics of data that is being moved by the layers below
  - In order to make it possible for computers with different data representations to communicate, the data structures to be exchanged can be defined in an abstract way, along with an standard encoding to be used on the wire.
  - It manages these abstract data structures and allows higher-level data structures to be defined and exchanged

29

OSI Reference Model (8)

- Application Layer
  - The application layer contains a number of different protocols and applications that are needed by the users.
  - A good example of widely used application protocol is HTTP (HyperText Transfer Protocol), which is the basis for Wide World Web distributed system.
  - When a browser wants a page, it sends the name of the page, to a web server, using HTTP protocol
  - Other application: FTP, e-Mail, news, etc…

30

TCP/IP Reference Model (1)

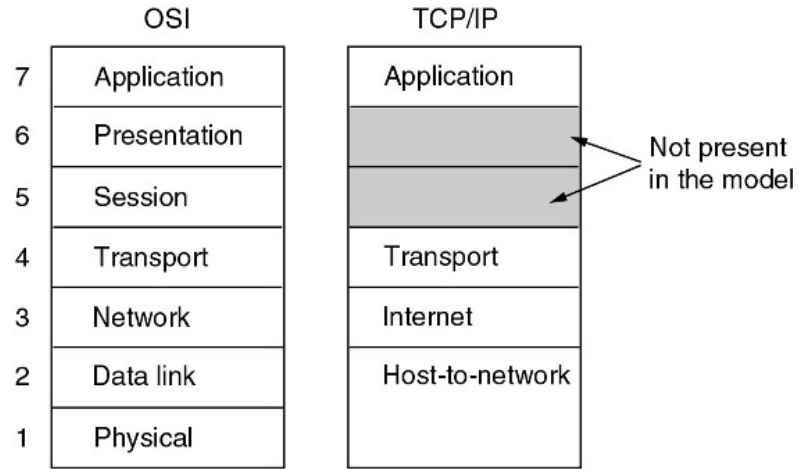

- Used by Internet, packet switching network (of networks) based on a connectionless internetwork layer

31

TCP/IP Reference Model (2)

- Internet Layer
  - Permits the hosts to inject packets into any network and have them travel independently to the destination (potentially using different paths or networks). The packets may arrive in a different order. It is the job of the higher layer to rearrange them
  - It defines an official packet format and protocol, called IP (**Internet Protocol**). The job of internet layer is to deliver IP packets where they want to go
    - Packet routing is one of the biggest issues
    - Avoiding congestion is another big issue

32

Analogy with the snail mail system

TCP/IP Reference Model (3)

- The Transport Layer
  - Designed to allow peer entities on the source and destination to carry on a conversation
  - Two end to end protocols: TCP and UDP
    - Transmission Control Protocol – end to end reliable connection oriented protocol that allows a byte stream originating from one machine to be delivered with no error on another machine in the Internet
    - User Datagram Protocol – unreliable connectionless protocol for applications that don't want TCP's sequencing flow control and want to provide theirs (or to apps that don't want connection overhead)

33

TCP/IP Reference Model (3)

- Application Layer – contains all the high-level protocols: http, file transfer, e-mail, domain name system, etc…

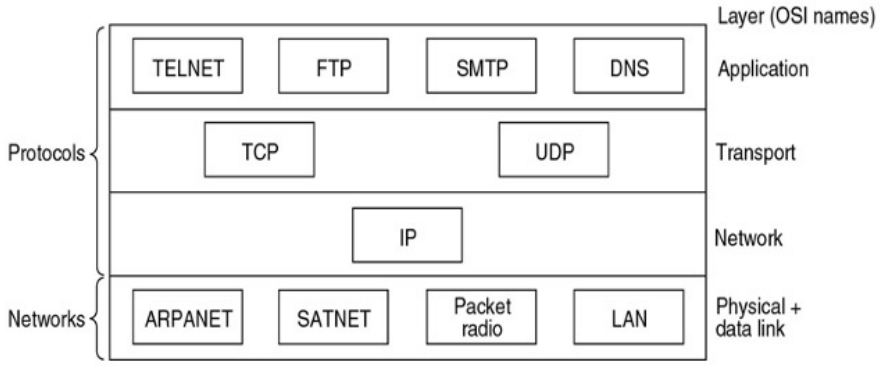

TCP/IP Reference Model (3)

- Host to Network Layer
  - Below the Internet Layer, in TCP/IP reference model is a great void
  - The model doesn't say much about it, except that the host has to connect to the network using some protocol, so it can send IP packets to it
  - This protocol is not defined and varies from host to host and from network to network

35

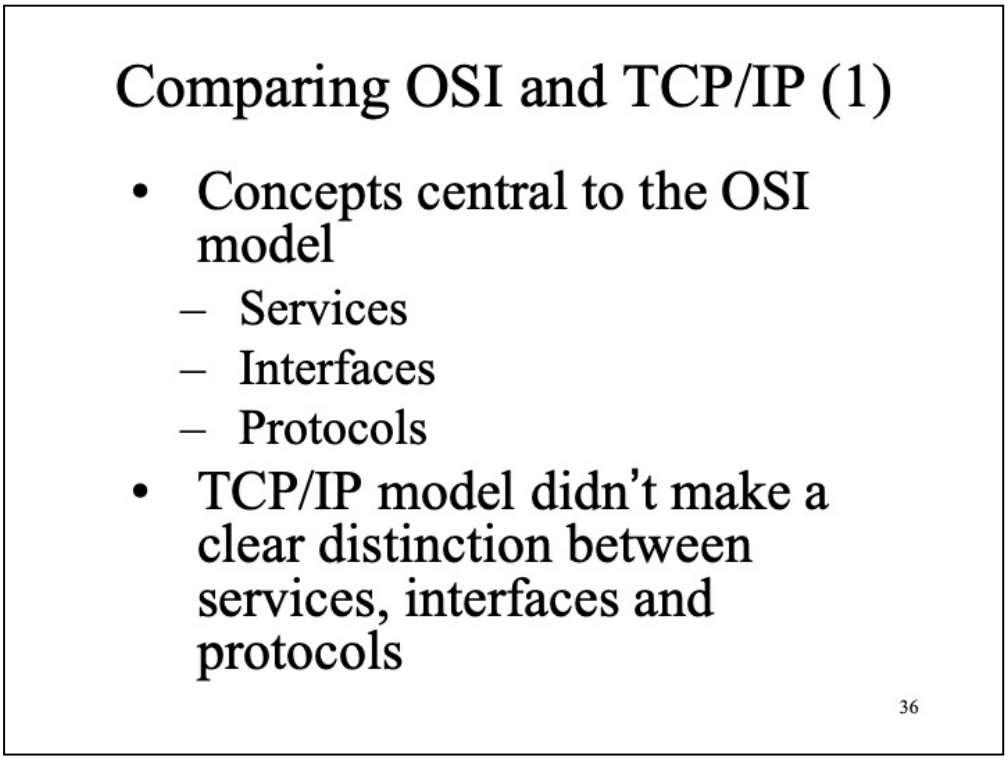

We are only comparing the models, not the corresponding protocol stacks. Both OSI and TCP/IP models have much in common. Both are based on a concept of a stack of independent protocols. Also, the functionality of the layers is somehow similar. **OSI** has three concepts that are central, and perhaps the biggest contribution of OSI was to make the explicit distinction between the three concepts:

SERVICES – tells what the layer does, not how entities above it access it, nor how the layer works. It defines the layer's semantics.

INTERFACES – tells the processes above it how to access it. It specifies what the parameters are and what result to expect. It says nothing about how the layer works inside.

PROTOCOLS – are peer layers own business. A layer can use any protocol if it gets the job done (provides the offered services). The protocols can change without affecting the software in the higher layer.

TCP/IP Model did not originally make a clear distinction between services, interfaces and protocols. People have tried to retrofit after the specifications, to make it look more like OSI. i.e., the only real services offered by the IP layer are: SEND_IP_PACKET and RECEIVE_IP_PACKET. Consequently, protocols in the OSI model are better hidden than in TCP/IP model and can be replaced a lot easier (as the technology changes), without disturbing the layers above.

Comparing OSI and TCP/IP (2)

- OSI reference was described **before** the protocols were invented
  - It means that the model is not biased towards a set of protocols, but is rather generic
  - Downside is that designers didn't have much experience (what function to put in which layer)
- TCP/IP protocols came first. The model was done just as a description of the protocols
  - The protocols did fit perfectly the model
  - The problem was that the model didn't fit any other protocols, useless to describe other types of networks than TCP/IP based.

37

Comparing OSI and TCP/IP (3)

- OSI has 7 layers, TCP/IP has only 4 layers
- Connection less vs. connection-oriented communication
  - OSI model – supports connection less and connection-oriented services at the network layer, while at transport layer supports only one type of connection-oriented service
  - TCP/IP model – supports only connection less services at the network layer, while at the transport layer offers both connection oriented and connection less services, giving the users a choice

38

A Critique of the OSI Model and Protocols

- Why OSI did not take over the world
  - Bad timing
  - Bad technology
  - Bad implementations
  - Bad politics

39

Bad Timing

- The apocalypse of the two elephants.

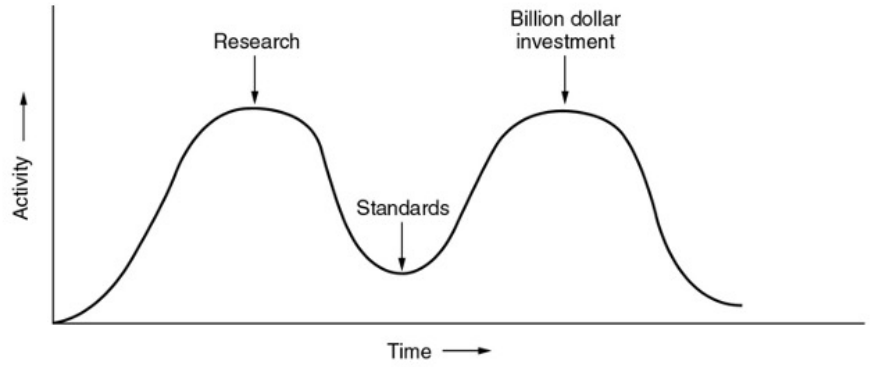

40

Standards have to be written between the research phase and investment phase. With OSI this didn't happen. Partially because the two phases were too close, partially because the competing TCP/IP model was already used in the research institutions, and companies already started to offer TCP/IP based products.

Bad Technology

- Two of the layers in OSI were nearly empty (session and presentation) while two others were overcrowded (network and data-link)
- The protocols and service definitions are very complex. They were almost incomprehensible.
- Some of the functions (such as addressing, flow control and error control) reappear again and again at different layers. This is unnecessary and inefficient

41

Bad Implementation

- Given the complexity of the protocols, the implementations were huge and inefficient
- People started to associate OSI with poor quality because the implementations were slow on the available equipment
- In contrast, one of the first implementations of TCP/IP was part of Berkeley UNIX and was quite good (also free). People begun to use it -> improvements -> large community -> more improvements

42

Bad Politics

- OSI was created by European Telecommunications ministries and USA Government – this has been seen (perceived) as a bunch of bureaucrats trying to push an inferior standard on the throats of researchers and scientists, that already had a working solution (TCP/IP)
- Of course, this was only partially true, but enough for TCP/IP model to get a lot of supporters

43

A Critique of the TCP/IP Reference Model

- Problems:
  - Service, interface, and protocol not distinguished
  - Not a general model
  - Host-to-network "layer" not really a layer
  - No mention of physical and data link layers
  - Minor protocols deeply entrenched, hard to replace

44

Hybrid Model

- The hybrid reference model to be used in this course

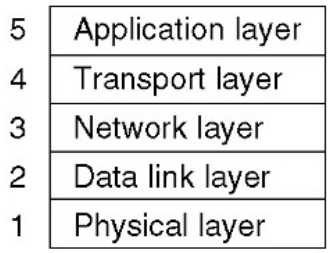

45

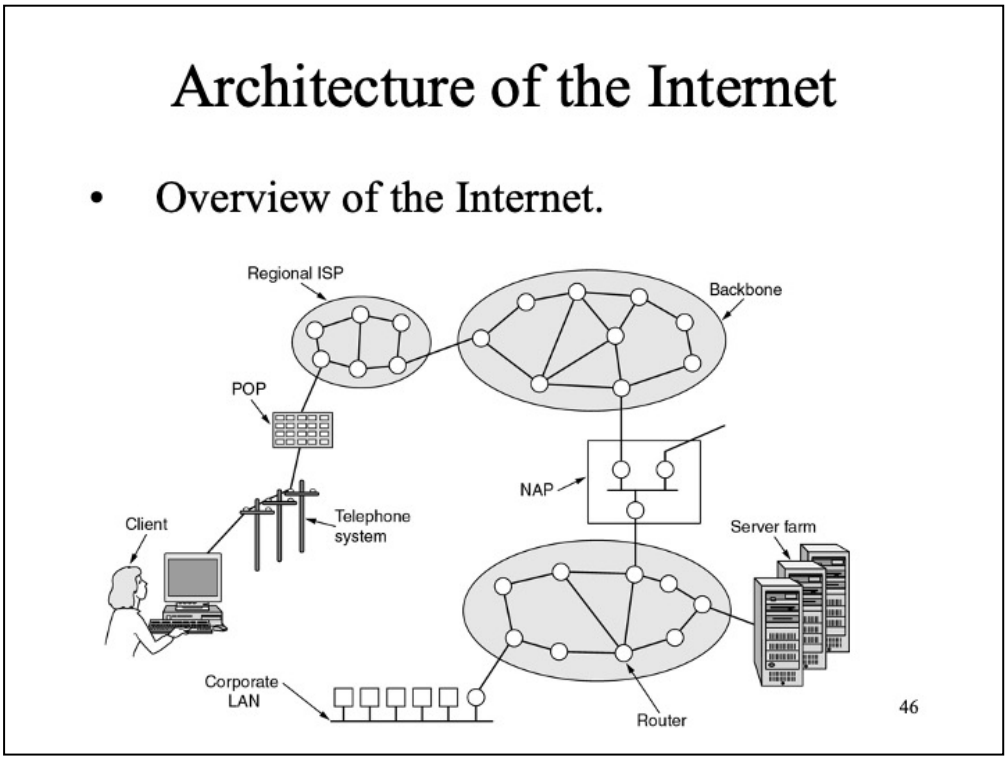

The **INTERNET** is not a network at all, but a vast collection of different networks that use common protocol and provide common services. Internet was not planed nor controlled by anyone. It all started back in late 1950's when DoD realized that all the communication is based on telephony systems with little or non redundancy. Therefore, a highly distributed, fault tolerant system needed to be in place. Late 1950's ARPA (Advanced Research Project Agency) was formed. In 1967 ARPA's interest turned into networking … ARPANET, first network of computers has been built. The subnet would consist of minicomputers (called **IMPs – INTERNET MESSAGE PROCESSORS**), interconnected by 56kb/s transmission lines. For reliability, each IMP would be connected to at least two another IMPs. The subnet was a datagram subnet, so if some lines and IMPs were destroyed, the messages could automatically be rerouted along the alternative paths.

The packets carried over the modem are transferred to the ISPs POP (**POINT OF PRESENCE**) where they are removed from the telephone system and injected into the ISP regional network. From this point on, the system is fully digital, and packet switched. The regional ISP consists of several routers, located in main cities where the ISP operates. The regional ISP subnet is connected to a backbone, that usually run between states or countries (even continents). The backbones are usually made of high bandwidth optical fiber lines, interconnected by routers. Interconnecting backbones often belonging to different competing ISPs.

Network Standardization

- Telecommunications World
  - ITU (International Telecommunication Union)
    - Radiocommunication Sector (ITU-R)
    - Telecommunications Standardization Sector (ITU-T)
    - Development Sector (ITU-D)
- International Standards World
  - ISO (International Standards Organization)
    - ANSI (American National Standards Institute), BSI (British Standards Institute), etc…
  - IEEE (Institute of Electrical and Electronics Engineers)
    - i.e. IEEE 802 group standardize LAN standards
- Internet Standards World
  - ITB (Internet Architecture Board)
    - RFC (Request for comments) for different standards
  - Got divided into:
    - IRTF (Internet Research Task Force)
    - ITEF (Internet Engineering Task Force)

47

References

- Andrew S. Tanenbaum – Computer Networks, ISBN 0-13-066102-3