Advanced Operating Systems

XI

Distributed Operating Systems

Communication

Prof. Muhammad Saeed

Layered Protocols

Open Systems Interconnection model (OSI model)

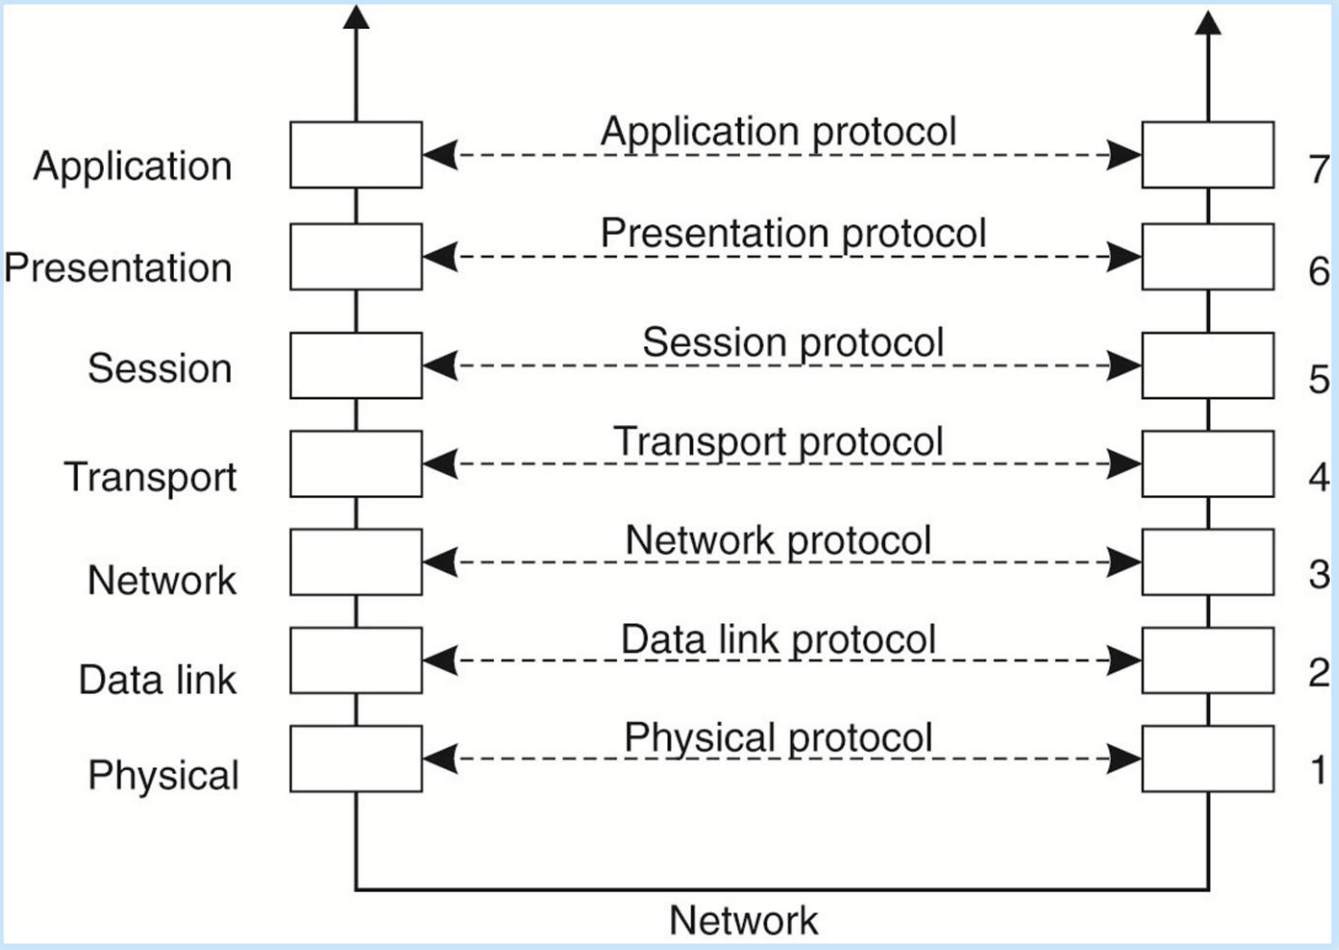

Layers, interfaces, and protocols in the OSI model

Layered Protocols

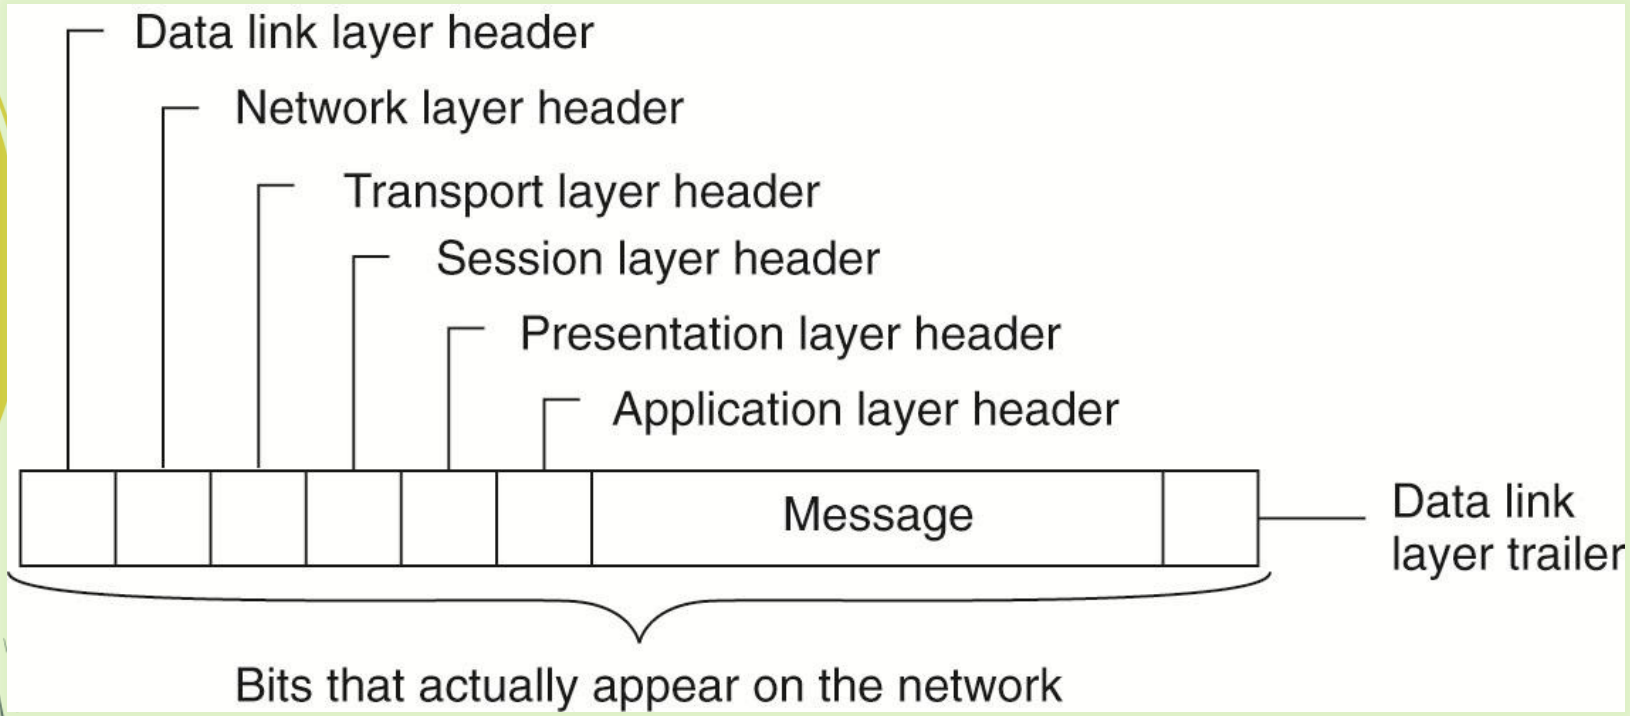

A typical message as it appears on the network.

Layered Protocols

Layer 1: physical layer

The physical layer defines electrical and physical specifications for devices. In particular, it defines the relationship between a device and a transmission medium, such as a copper or optical cable.

Layer 2: data link layer

The data link layer provides the functional and procedural means to transfer data between network entities and to detect and possibly correct errors that may occur in the physical layer.

Layer 3: network layer

The network layer provides the functional and procedural means of transferring variable length data sequences from a source host on one network to a destination host on a different network, while maintaining the quality of service requested by the transport layer (in contrast to the data link layer which connects hosts within the same network). The network layer performs network routing functions, and might also perform fragmentation and reassembly, and report delivery errors. Routers operate at this layer

Layered Protocols

Layer 4: transport layer

The transport layer provides transparent transfer of data between end users, providing reliable data transfer services to the upper layers. The transport layer controls the reliability of a given link through flow control, segmentation/desegmentation, and error control.

Layer 5: session layer

The session layer controls the dialogues (connections) between computers. It establishes, manages and terminates the connections between the local and remote application. It provides for full-duplex, half-duplex, or simplex operation, and establishes checkpointing, adjournment, termination, and restart procedures.

Layer 6: presentation layer

This layer provides independence from data representation (e.g., encryption) by translating between application and network formats. The presentation layer transforms data into the form that the application accepts.

Layer 7: application layer

The application layer is the OSI layer closest to the end user, which means that both the OSI application layer and the user interact directly with the software application. This layer interacts with software applications that implement a communicating component

Remote Procedure Call

Client and Server Stubs

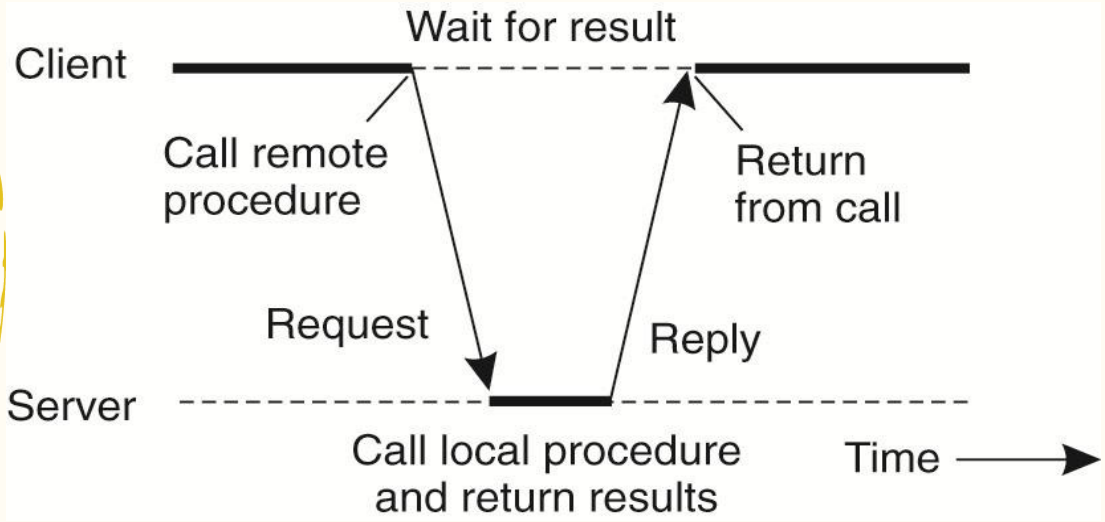

Client and Server Stubs

A **method stub** or simply **stub** in software development is a piece of code used to stand in for some other programming functionality. A stub may simulate the behavior of existing code (such as a procedure on a remote machine) or be a temporary substitute for yet-to-be-developed code. Stubs are therefore most useful in porting, distributed computing as well as general software development and testing.

Remote Procedure Call

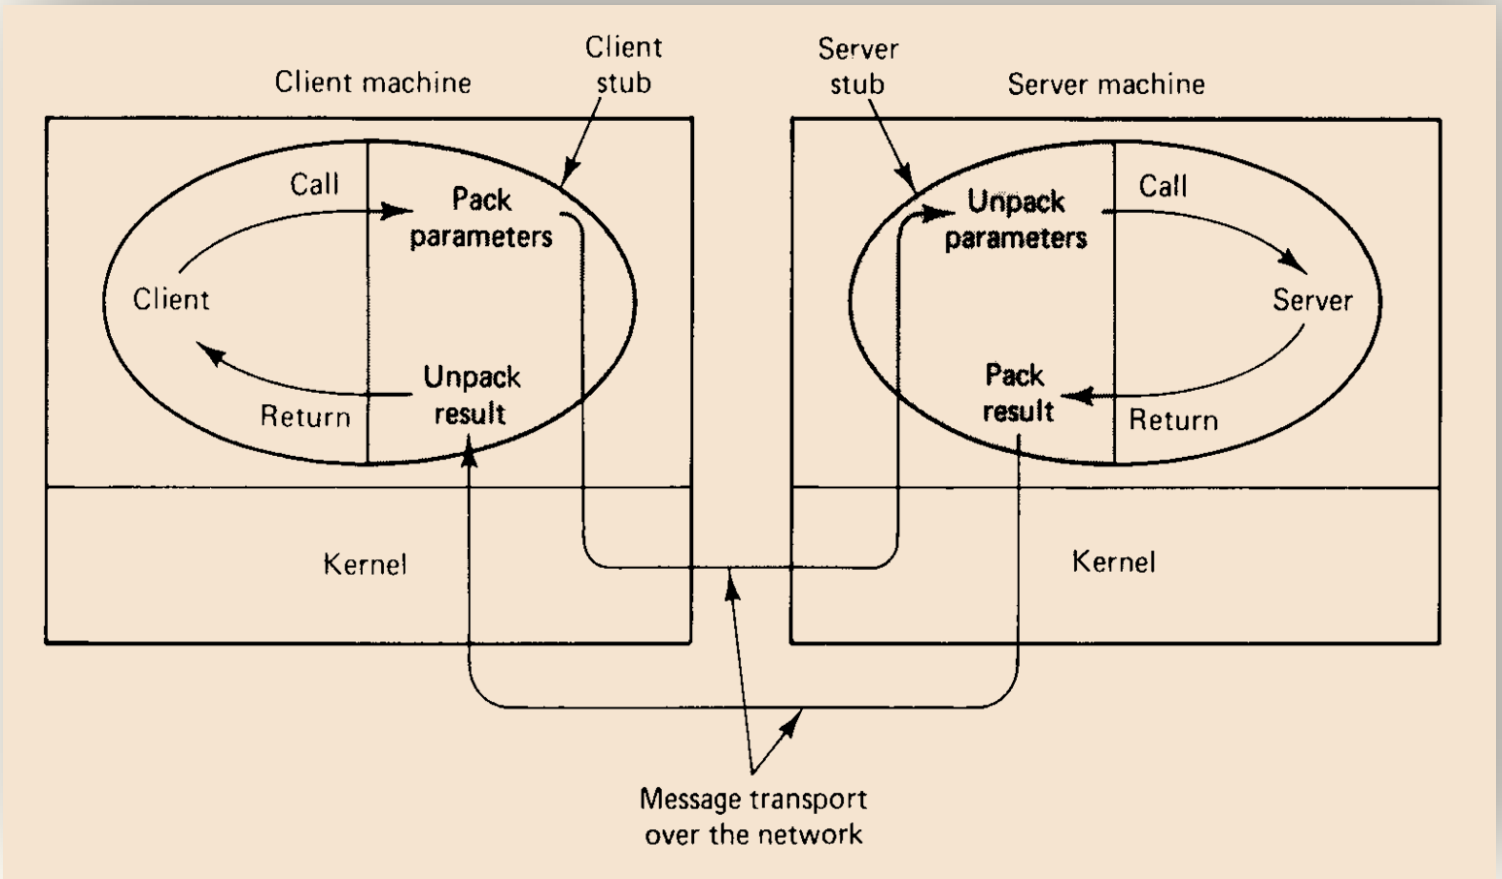

Remote Procedure Call

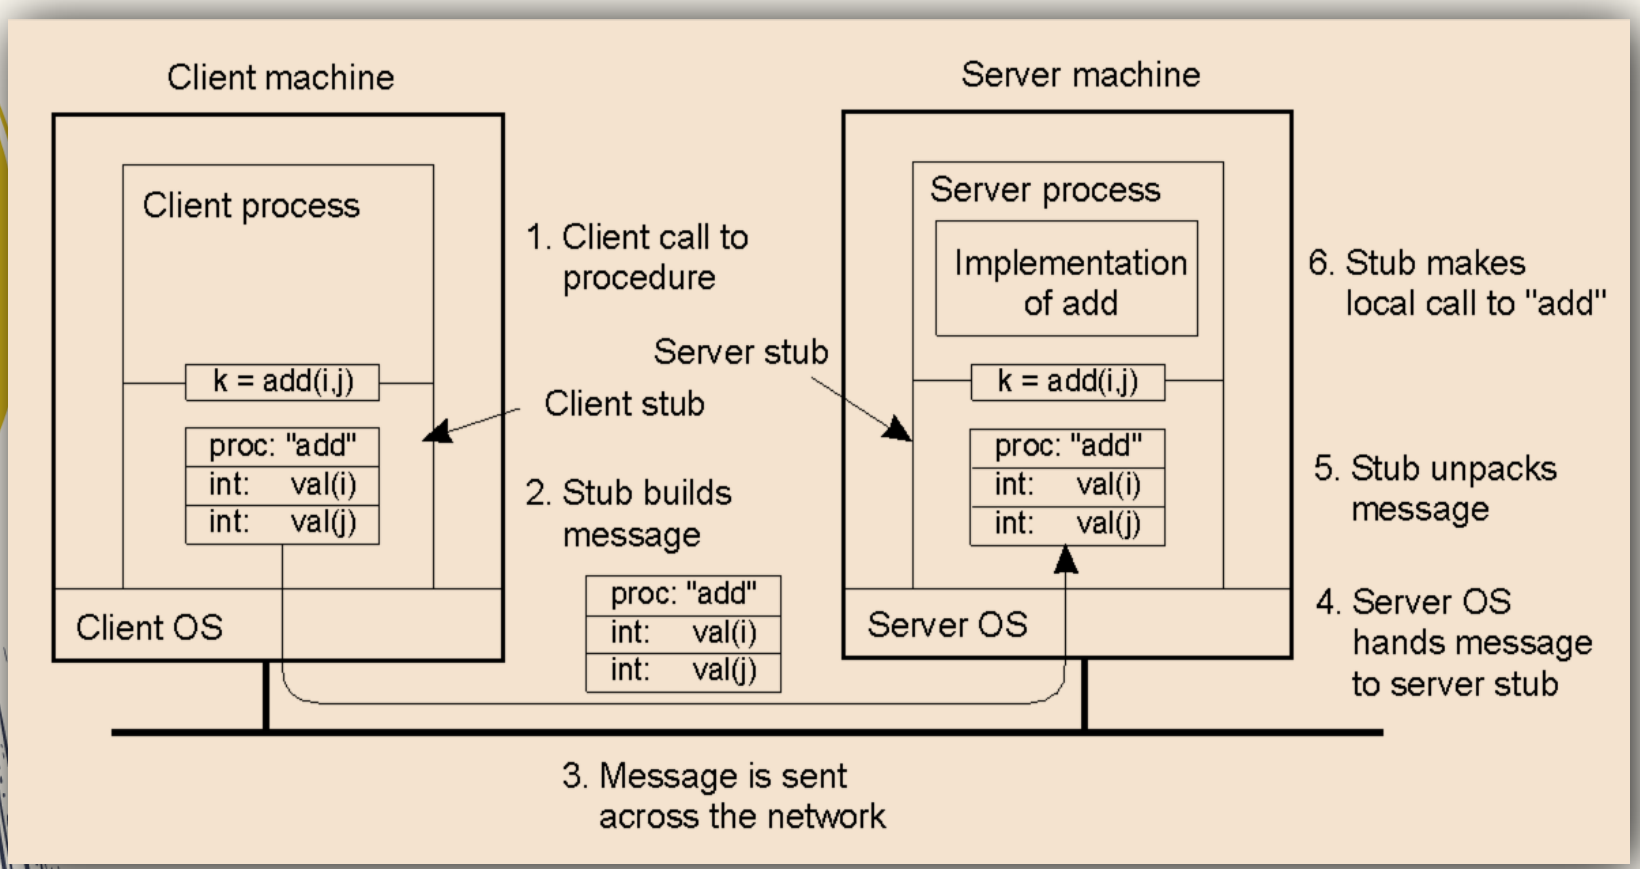

Remote Procedure Call

Steps for remote procedure call:

1. The client procedure calls the client stub in the normal way.
2. The client stub builds a message and calls the local operating system.
3. The client's OS sends the message to the remote OS.
4. The remote OS gives the message to the server stub.
5. The server stub unpacks the parameters and calls the server.
6. The server does the work and returns the result to the stub.
7. The server stub packs it in a message and calls its local OS.
8. The server's OS sends the message to the client's OS.
9. The client's OS gives the message to the client stub.
10. The stub unpacks the result and returns to the client.

Parameter Marshaling

Packet and Circuit Switching

Packet switching

Packet switching is a digital networking communications method that groups all transmitted data – regardless of content, type, or structure – into suitably sized blocks, called ***packets***. Packet switching features delivery of variable-bit-rate data streams (sequences of packets) over a shared network. When traversing network adapters, switches, routers and other network nodes, packets are buffered and queued, resulting in variable delay and throughput depending on the traffic load in the network.

Packet switching contrasts with another principal networking paradigm, **circuit switching**, a method which sets up a limited number of dedicated connections of constant bit rate and constant delay between nodes for exclusive use during the communication session. In case of traffic fees (as opposed to flat rate), for example in cellular communication services, circuit switching is characterized by a fee per time unit of connection time, even when no data is transferred, while packet switching is characterized by a fee per unit of information.

……..

..... Packet and Circuit Switching

..... Packet switching

Two major packet switching modes exist;

(1) **connectionless packet switching**
(datagram switching) and

(2) **connection-oriented packet switching**
(virtual circuit switching).

In the first case each packet includes complete addressing or routing information. The packets are routed individually, sometimes resulting in different paths and out-of-order delivery. In the second case a connection is defined and preallocated in each involved node during a connection phase before any packet is transferred. The packets include a connection identifier rather than address information, and are delivered in order.

Frames and Frame Relay

Frame

A **frame** is a digital data transmission unit or **data packet** that includes frame synchronization, i.e. a sequence of bits or symbols making it possible for the receiver to detect the beginning and end of the packet in the stream of symbols or bits.

Frame Relay is a standardized wide area network technology that specifies the physical and logical link layers of digital telecommunications channels using a packet switching methodology. Originally designed for transport across **Integrated Services Digital Network** (ISDN) infrastructure, it may be used today in the context of many other network interfaces.

Network providers commonly implement **Frame Relay for voice** (VoFR) and data as an encapsulation technique, used between **local area networks** (LANs) over a **wide area network**(WAN). Each end-user gets a private line (or leased line) to a Frame Relay node. The Frame Relay network handles the transmission over a frequently-changing path transparent to all end-users.

Frame Relay has become one of the most extensively-used WAN protocols. Its cheapness (compared to leased lines) provided one reason for its popularity. The extreme simplicity of configuring user equipment in a Frame Relay network offers another reason for Frame Relay's popularity.

……. Packet and Circuit Switching

Circuit switching

Circuit switching is a methodology of implementing a telecommunications network in which two network nodes establish a dedicated communications channel (circuit) through the network before the nodes may communicate. The circuit guarantees the full bandwidth of the channel and remains connected for the duration of the communication session. The circuit functions as if the nodes were physically connected as with an electrical circuit.

The defining example of a circuit-switched network is the early analog telephone network. When a call is made from one telephone to another, switches within the telephone exchanges create a continuous wire circuit between the two telephones, for as long as the call lasts.

Circuit switching contrasts with packet switching which divides the data to be transmitted into packets transmitted through the network independently. Packet switching shares available network bandwidth between multiple communication sessions.

……

..... Packet and Circuit Switching

..... Circuit switching

In circuit switching, the bit delay is constant during a connection, as opposed to packet switching, where packet queues may cause varying packet transfer delay. Each circuit cannot be used by other callers until the circuit is released and a new connection is set up. Even if no actual communication is taking place, the channel remains unavailable to other users. Channels that are available for new calls are said to be idle.

Virtual circuit switching is a packet switching technology that emulates circuit switching, in the sense that the connection is established before any packets are transferred, and packets are delivered in order.

While circuit switching is commonly used for connecting voice circuits, the concept of a dedicated path persisting between two communicating parties or nodes can be extended to signal content other than voice. Its advantage is that it provides for continuous transfer without the overhead associated with packets making maximal use of available bandwidth for that communication. The disadvantage is inflexibility; the connection and the bandwidth associated with it are reserved and unavailable for other uses.

..... Packet and Circuit Switching

Datagrams

A **datagram** is a basic transfer unit associated with a packet-switched network in which the delivery, arrival time, and order of arrival are not guaranteed by the network service.

Each datagram has two components, a **header** and a **data payload**. The header contains all the information sufficient for routing from the originating equipment to the destination without relying on prior exchanges between the equipment and the network. Headers may include source and destination addresses as well as a type field. The payload is the data to be transported. This process of nesting data payloads in a tagged header is called **encapsulation**.

The **Internet Protocol** defines standards for several types of datagrams. The term *datagram* is often considered synonymous to *packet* but there are some nuances. The term *datagram* is generally reserved for packets of an unreliable service that does not notify the user if delivery fails, while the term *packet* applies to any message formatted as a packet. For example, **Internet Protocol** (IP) provides an unreliable service and **UDP** over IP is also unreliable. That is why IP and UDP packets are generally called datagrams.

Connections

Connection-oriented (CO-mode communication)

Connection-oriented (**CO-mode communication**) is a data communication mode in telecommunications whereby the devices at the end points use a protocol to establish an end-to-end logical or physical connection before any data may be sent. In case of digital transmission, in-order delivery of a bit stream or byte stream is provided. Connection-oriented protocol services are often but not always *reliable* network services, that provide acknowledgment after successful delivery, and automatic repeat request functions in case of missing data or detected bit-errors.

Connectionless Communication

In packet switching networks, **CL-mode** or **connectionless communication** is a data transmission method in which each data packet carries information in a header record that contains a destination address sufficient to permit the independent delivery of the packet to its destination via the network.

A packet transmitted in a connectionless mode is frequently called a **datagram**. It has the advantage over a connection-oriented mode in that it has low overhead.

Addressing and Routing Methodologies

Multicast

To transmit a single message to a select group of recipients. A simple example of multicasting is sending an e-mail message to a mailing list. Teleconferencing and videoconferencing also use multicasting, but require more robust protocols and networks.

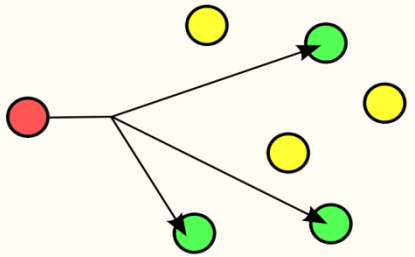

Broadcast

Broadcasting refers to transmitting a packet that will be received by every device on the network. In practice, the scope of the broadcast is limited to a broadcast domain. Not all network technologies support broadcast addressing. Broadcasting is largely confined to local area network (LAN) technologies, most notably Ethernet and Token Ring, where the performance impact of broadcasting is not as large as it would be in a wide area network (WAN). The successor to Internet Protocol Version 4 (IPv4), IPv6 also does not implement the broadcast method.

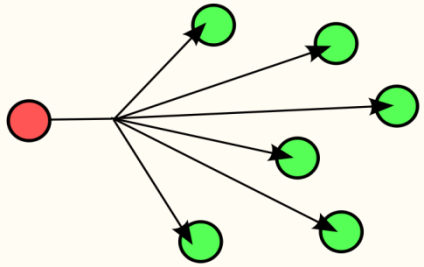

…… Addressing and Routing Methodologies

Unicast

In unicast addressing a host sends datagrams to another single host identified by a unique IP address. Unicast-based media servers open and provide a stream for each unique user.

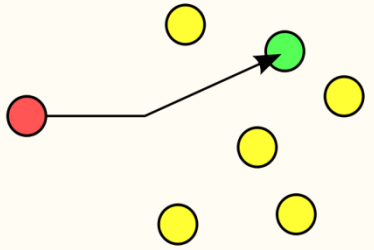

Anycast

Anycast is a Network addressing and routing methodology in which datagrams from a single sender are routed to the topologically nearest node in a group of potential receivers all identified by the same destination address.

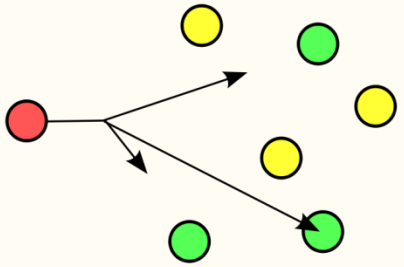

Geocast

Geocast refers to the delivery of information to a group of destinations in a network identified by their geographical locations. It is a specialized form of multicast addressing used by some routing protocols for mobile ad hoc networks.

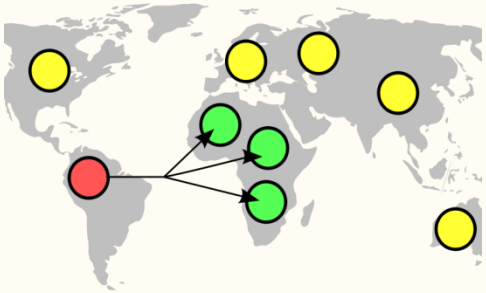

Sockets

Network Socket:

A **network socket** is an endpoint of an inter-process communication flow across a computer network. Today, most communication between computers is based on the Internet Protocol; therefore most network sockets are **Internet sockets**.

A **socket API** is an application programming interface (API), usually provided by the operating system, that allows application programs to control and use network sockets. Internet socket APIs are usually based on the **Berkeley sockets** standard(**BSD sockets**, is a computing library with an application programming interface (API) for **internet sockets** and **Unix domain sockets**, used for inter-process communication (IPC)).

A **socket address** is the combination of an IP address and a port number, much like one end of a telephone connection is the combination of a phone number and a particular extension. Based on this address, internet sockets deliver incoming data packets to the appropriate application process or thread.

..... Sockets

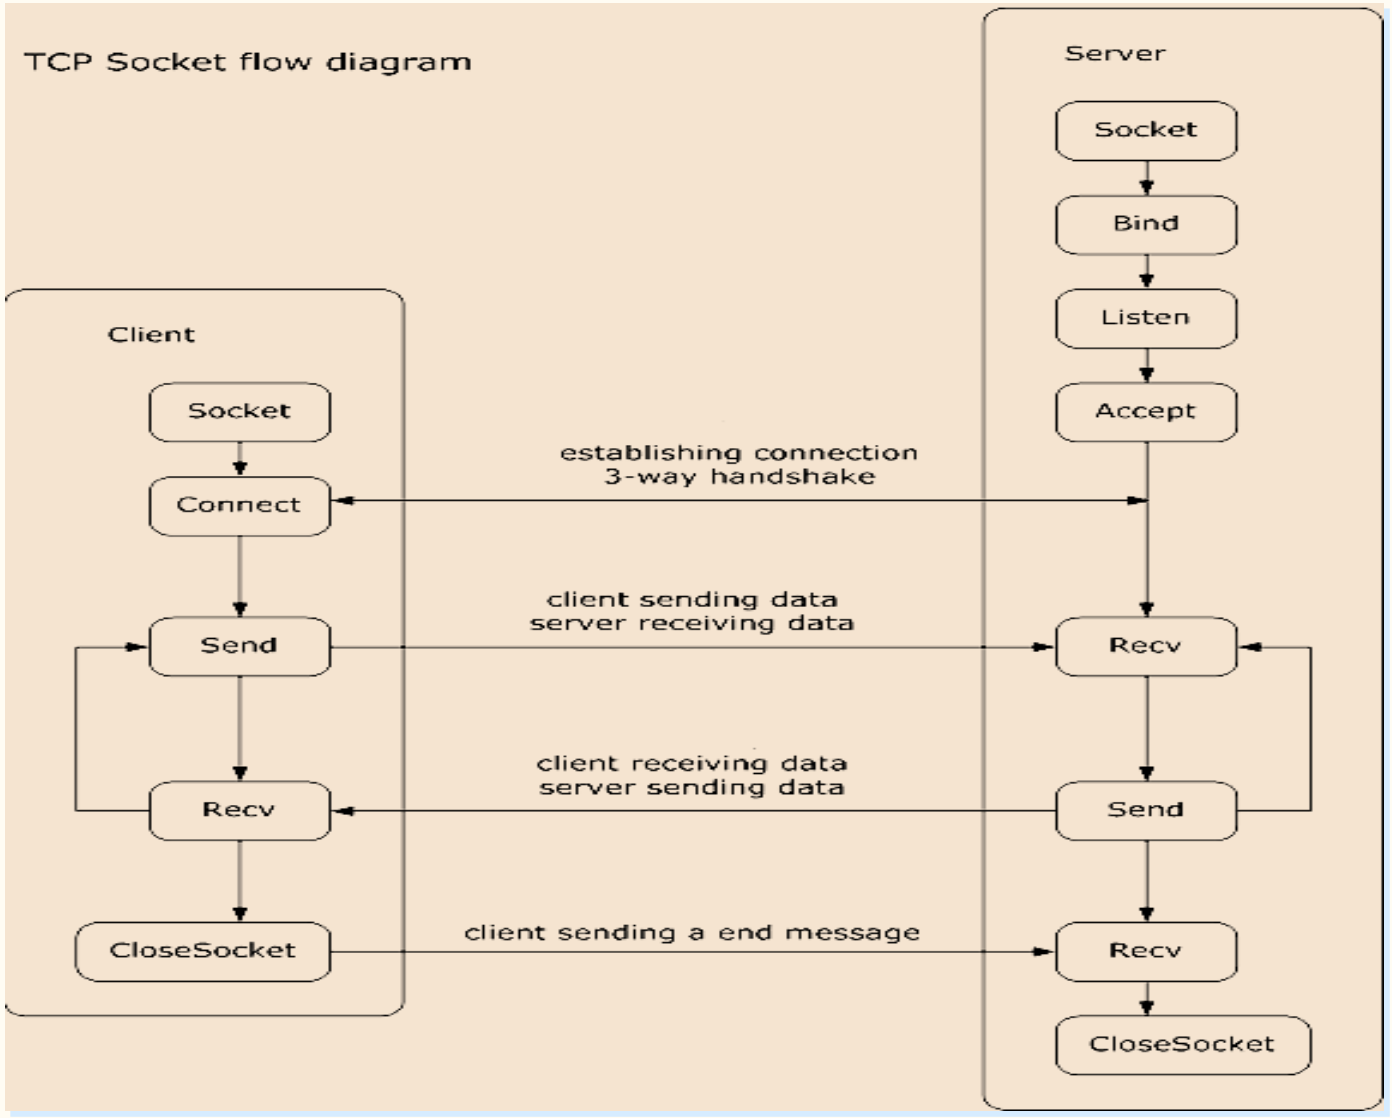
TCP Socket flow diagram
Server
Socket
Bind
Listen
Accept
Client
Socket
Connect
establishing connection
3-way handshake
Send
client sending data
server receiving data
Recv
Recv
client receiving data
server sending data
Send
CloseSocket
client sending a end message
Recv
CloseSocket

Ports

Network Ports

In computer networking a **port** is an application-specific or process-specific software construct serving as a communications endpoint in a computer's host operating system. A port is associated with an IP address of the host, as well as the type of protocol used for communication. The protocols that primarily use the ports are the Transport Layer protocols, such as the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP) of the Internet Protocol Suite. A port is identified for each address and protocol by a 16-bit number, commonly known as the **port number**. The port number completes the destination address for a communications session. Thus, different IP addresses or protocols may use the same port number for communication, e.g. on a given host or interface UDP and TCP may use the same port number, or on a host with two interfaces, both addresses are associated with a port having the same number.

A range of well-known ports is reserved by convention to identify specific service types on a host. In the client-server model of application architecture ports are used to provide a multiplexing service on each port number that network clients connect to for service initiation, after which communication is reestablished on other connection-specific port number.

Some of the slides: Courtesy by Tanenbaum & Van Steen, Distributed Systems: Principles and Paradigms, 2e, (c) 2007 Prentice-Hall, Inc.