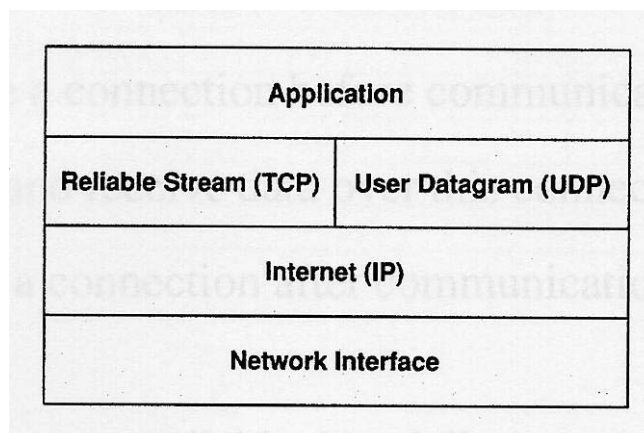


Chapter 3

TCP and UDP

(Recommended reading: Comer, Vol 1, Chapters 12, 13)

- In the TCP/IP protocol suite, there are two major transport protocols: transmission control protocol (TCP) and user datagram protocol (UDP).
- Role of TCP and UDP in the TCP/IP reference model:

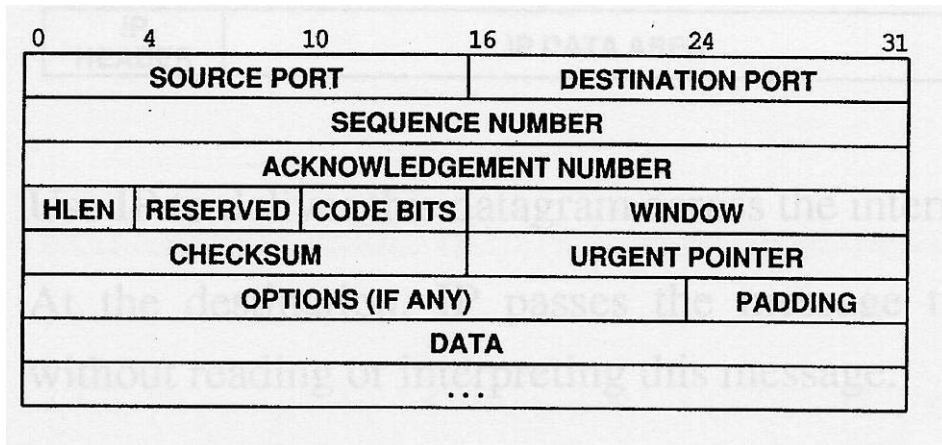


- Transmission Control Protocol (TCP)
 - Basic Features
 - TCP provides connection-oriented communication (virtual circuit connection, like telephone communication). It manages a point-to-point and full-duplex connection for an application between two computers:
 - creates a connection before communication;
 - sends and receives data over this connection;
 - » Application data is considered to be a continuous unstructured stream of bits.
 - » Application programs must understand stream contents and boundaries of records.
 - » piggybacking permitted
 - closes a connection after communication.
 - TCP guarantees reliable data delivery.
 - The TCP recipient will receive data in a correct order without data loss or error.

- Connection abstraction
 - TCP allows **multiple** application programs on a given machine to communicate with programs in other machines.
 - A virtual circuit connection consists of two endpoints.
 - Each endpoint is a pair of integers (host, port)
 - Host = IP address
 - A port is a positive integer on the host, usually identifying an application program, e.g.,
 - » daytime: 13
 - » FTP: 20 and 21
 - » SMTP: 25
 - » HTTP: 80
 - Functionalities of a port
 - » If a process attempts to extract data from a port before any data arrives, the OS temporarily stops (blocks) the process until data arrives; once the data arrives, the OS passes the data to the process and unblocks it.
 - » Ports are buffered.
 - » Multiple packets arriving at a port are being queued up; the packets are being extracted on a first-in-first-out (FIFO) basis.

- Examples:
 - The endpoints (18.9.0.32, 1184) and (128.10.2.3, 53) form a connection.
 - The endpoints (128.2.254.139, 1184) and (128.10.2.3, 53) form another connection.
 - A given TCP port number can be shared by multiple connections on the same machine.
- A connection is setup once during the virtual circuit initialization. Subsequent data transfer uses the connection object, not the source and destination IP addresses to specify the transfer.

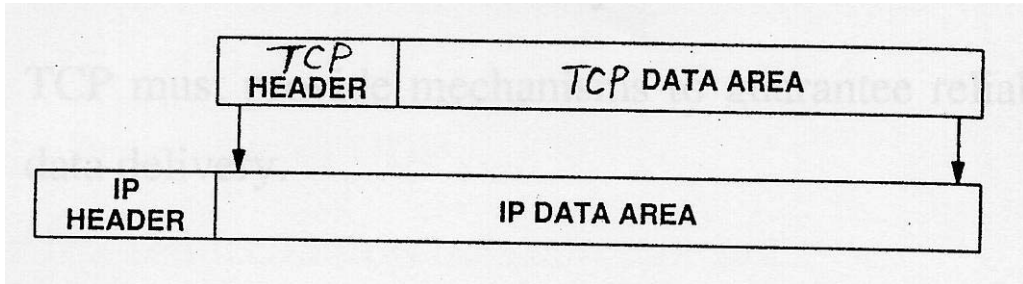
- TCP segment and its delivery
 - In TCP, the basic data unit is called a segment.
 - Format of a TCP segment:



- Header
 - The header is composed of a 20-byte fixed part and an optional part with a variable length. The total size of the header (in 32-bit words) is specified in HLEN.
- Data
 - The data can have a variable size, which can be up to $65535 - 20 - 20 = 65515$ bytes.

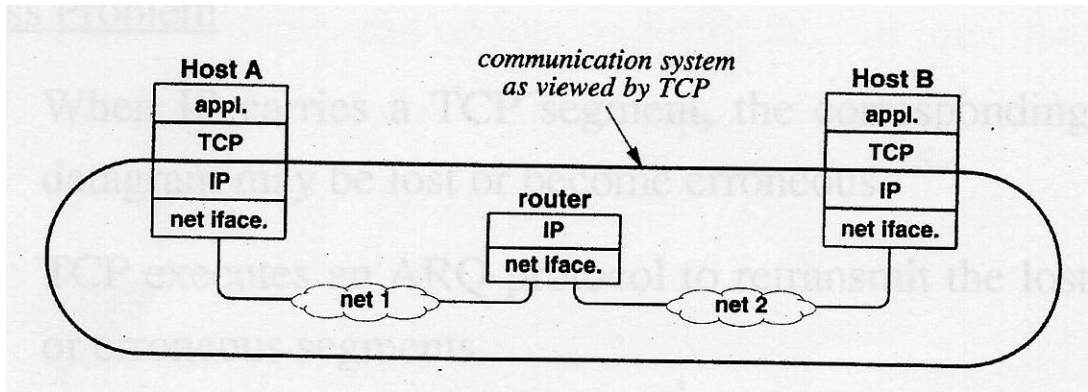
- **CODE BITS**
 - Control bits to indicate end of stream, acknowledgement field being valid, connection reset, urgent pointer field being valid, etc.
- **URGENT POINTER**
 - for out-of-band data (urgent data)
 - e.g., user sending keyboard interrupts or aborts to a remote login session
- **WINDOW**
 - During piggybacking, how much data a receiver is willing to accept.

- TCP uses IP to carry segments as follows:
 - Encapsulate each TCP segment in an IP datagram.



- Use IP to deliver this datagram across the internet.
- At the destination, IP passes the message to TCP without reading or interpreting this message.
- Automatic buffering: even if the application program generates the stream one octet at a time, the TCP software usually collects enough data to fill a reasonably large datagram before transmission.

- TCP views the underlying communication system as a logical connection:



- Providing Reliable Data Delivery
 - TCP uses the unreliable IP to carry data.
 - TCP must provide mechanisms to guarantee reliable data delivery.
 - » Sliding window protocol; selective reject (SREJ)
 - » positive acknowledgement with retransmission
 - Out-of-Order Problem
 - At the receiver side, IP may receive the datagrams in a wrong order.
 - After IP passes the TCP segments to TCP, TCP reorders these segments into the correct order. (Also called stream delivery service.)

- Loss Problem
 - When IP carries a TCP segment, the corresponding datagram may be lost or become erroneous.
 - TCP executes an Automatic Repeat Request (ARQ) protocol to retransmit the lost or erroneous segments.
- end-to-end flow control via positive acknowledgement
 - acknowledgement of segments received
 - How many additional octets of data the receiver is prepare to receive (receiver's current buffer size).
- Recall: in Sliding Window protocol with Selective Reject scheme,
 - number of bits in sequence number = 32
 - maximum window size (number of unacknowledged data units) = 2^{31}
 - In TCP, each data unit is one octet, **not** a segment or a packet.

- User Datagram Protocol (UDP)
 - UDP is another popular transport protocol.
 - Basic Properties of UDP
 - UDP is a connectionless transport protocol.
 - A UDP application sends messages without establishing and then closing a connection.
 - UDP has a smaller overhead than TCP, especially when the total size of the messages is small.
 - UDP does not guarantee reliable data delivery.
 - UDP messages can be lost or duplicated, or they may arrive out of order; and they can arrive faster than the receiver can process them.
 - The application programmers using UDP have to consider and tackle these issues themselves.
 - UDP has no mechanism for flow control.
 - Difference between UDP and IP
 - UDP distinguishes among applications within a given host via the destination port number; an IP datagram only identifies a destination host via the IP address.

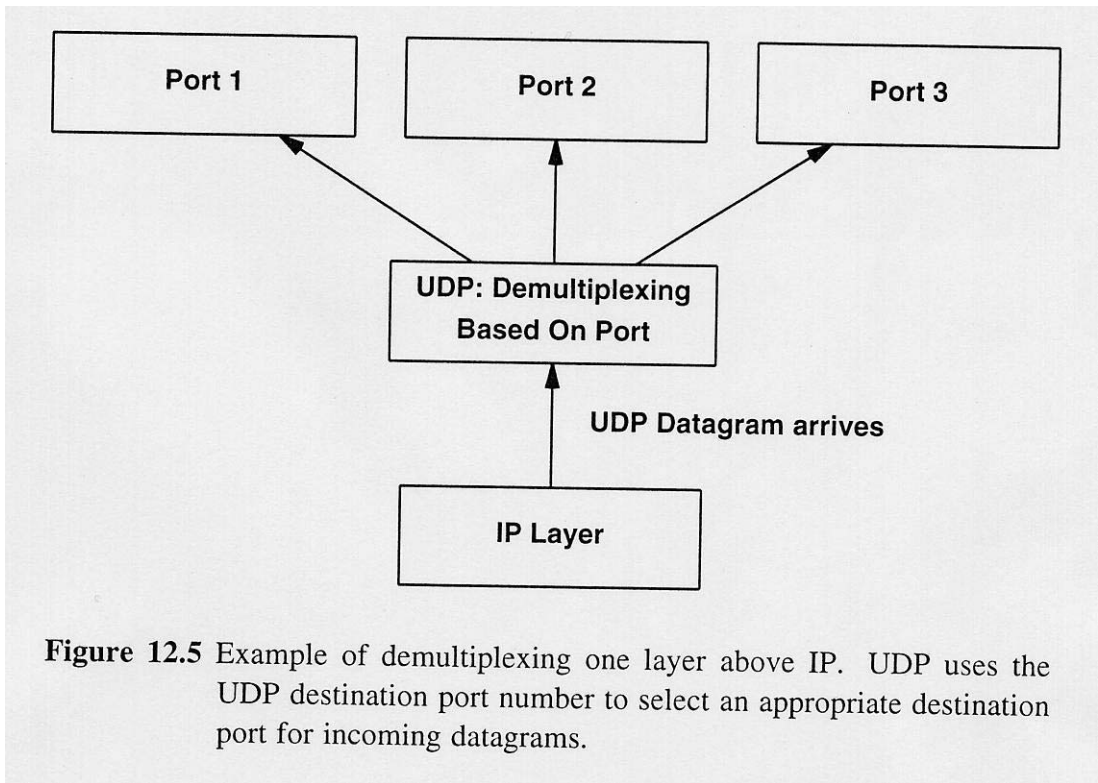
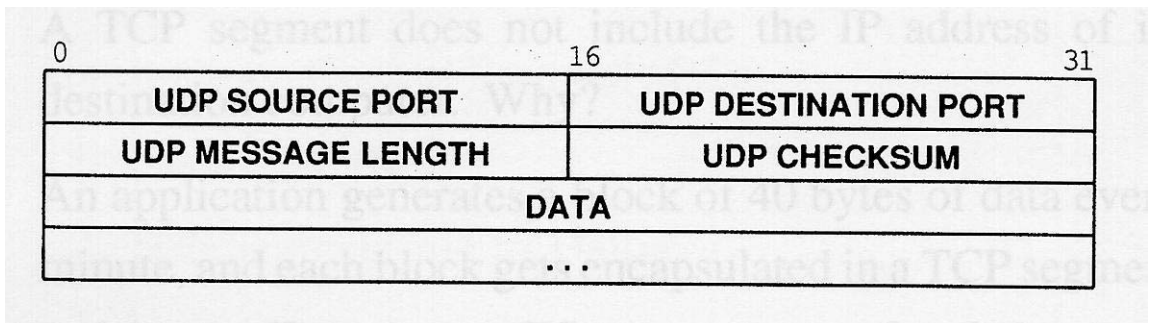


Figure 12.5 Example of demultiplexing one layer above IP. UDP uses the UDP destination port number to select an appropriate destination port for incoming datagrams.

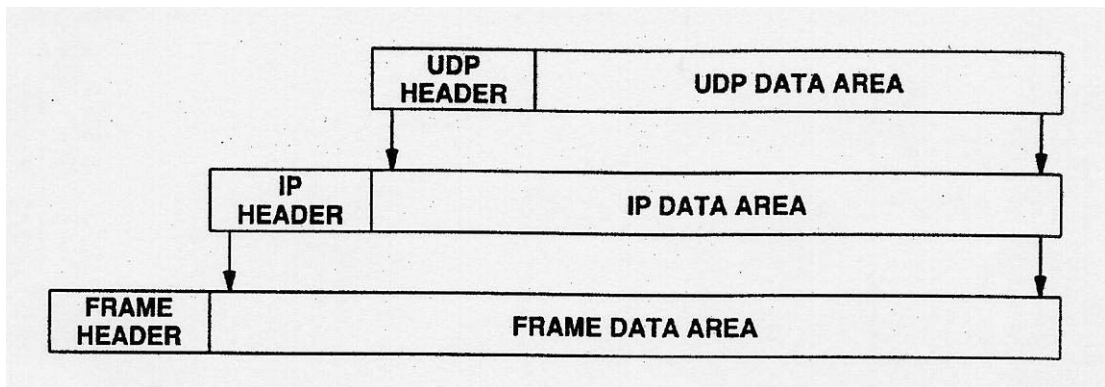
– UDP Segment and Its Transport

- Format of a UDP segment:



- The header consists of 8 bytes.
- The data can have a variable size.

- UDP uses IP to carry UDP segments:
 - Encapsulate each UDP message in an IP datagram, and use IP to deliver this datagram across the internet.
 - When this datagram is delivered across any specific network, it is further encapsulated:



- Application programmers using UDP often ignore the problems of unreliability, message loss, duplication, delay, out-of-order delivery, flooding, and etc.
- These programs are often tested using highly reliable, low-delay LANs.
- Many application programs that rely on UDP work well locally, but fail dramatically in a larger TCP/IP internet.

- Tutorial Problems

- What is the relationship between TCP and IP?
- A TCP segment does not include the IP address of its destination computer. Why?
- When do duplicated segments occur in TCP?
- An application generates a block of 40 bytes of data every minute, and each block gets encapsulated in a TCP segment and then an IP datagram. What percentage of each datagram will be the actual data?
- Compare and contrast TCP and UDP.
- State some applications in which UDP is more suitable than TCP.
- For WWW browsing, would you choose TCP or UDP?