

Network Programming

Dr. Thayer Hayajneh

Computer Engineering Department

Transport Layer

UDP Protocol

1

Transport Layer

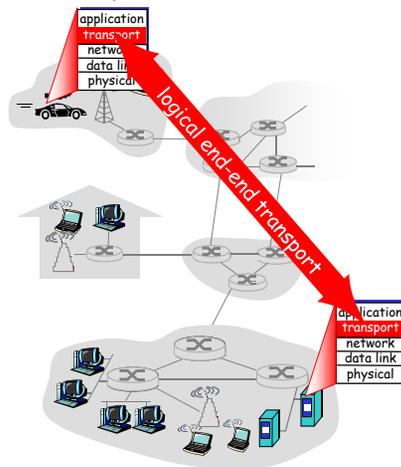
Our goals:

- understand principles behind transport layer services:
 - multiplexing/demultiplexing
 - reliable data transfer
 - flow control
 - congestion control
- learn about transport layer protocols in the Internet:
 - UDP: connectionless transport
 - TCP: connection-oriented transport
 - TCP congestion control

2

Transport services and protocols

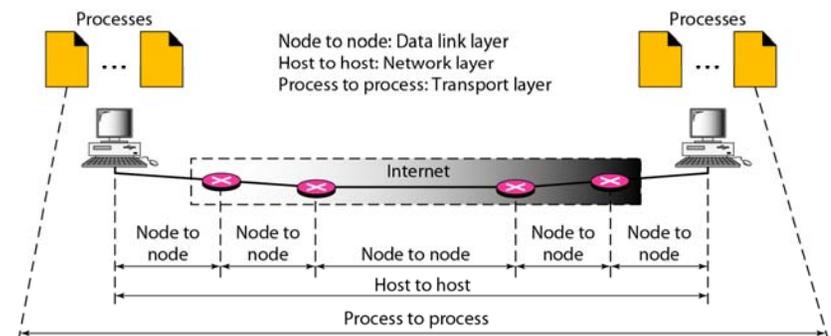
- provide *logical communication* between app processes running on different hosts
- transport protocols run in end systems
 - send side: breaks app messages into *segments*, passes to network layer
 - rcv side: reassembles segments into messages, passes to app layer
- more than one transport protocol available to apps
 - Internet: TCP and UDP



3

Transport vs. network layer

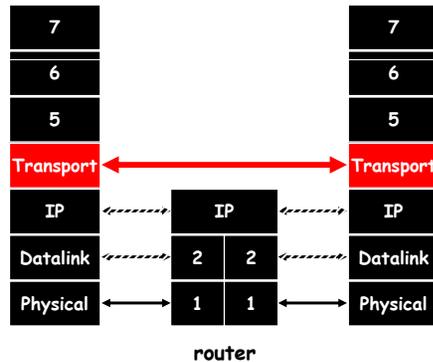
- network layer*: logical communication between hosts
- transport layer*: logical communication between processes
 - relies on, enhances, network layer services



4

Transport Protocols

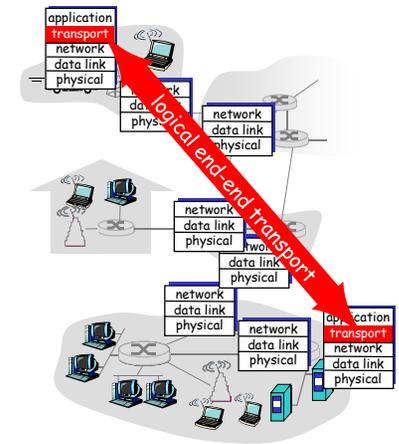
- Lowest level end-to-end protocol.
 - Header generated by sender is interpreted only by the destination
 - Routers view transport header as part of the payload



5

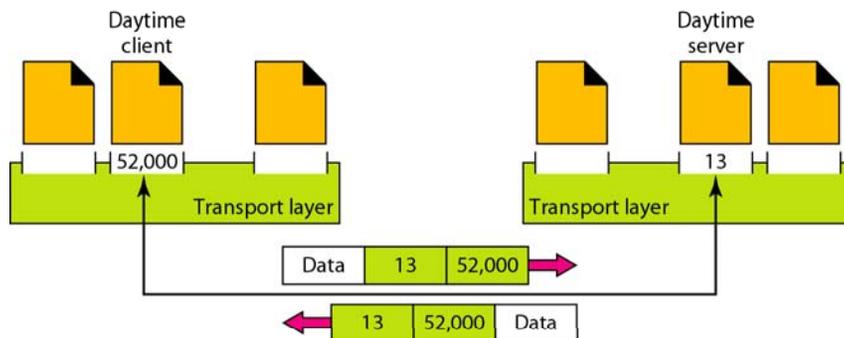
Internet transport-layer protocols

- reliable, in-order delivery (TCP)
 - congestion control
 - flow control
 - connection setup
- unreliable, unordered delivery: UDP
 - no-frills extension of "best-effort" IP
- services not available:
 - delay guarantees
 - bandwidth guarantees



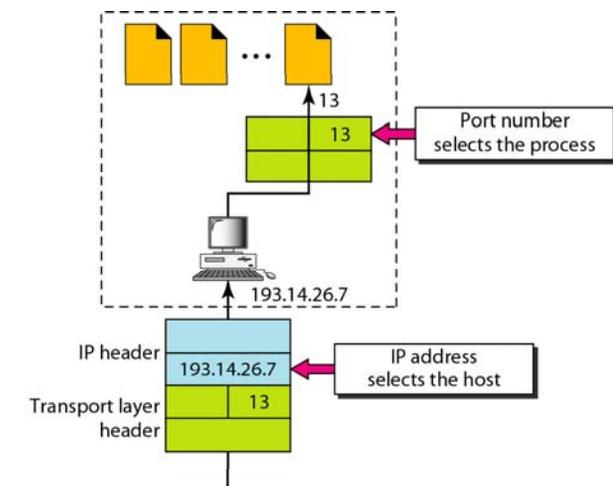
6

Port numbers 16 bits



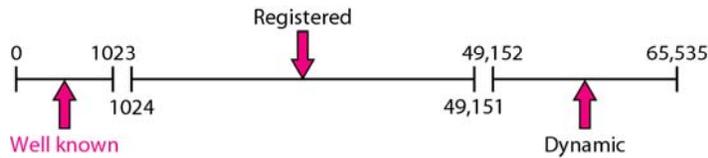
7

IP addresses versus port numbers



8

IANA ranges



9

Port Numbers in UNIX

In UNIX, the well-known ports are stored in a file called `/etc/services`. Each line in this file gives the name of the server and the well-known port number. We can use the `grep` utility to extract the line corresponding to the desired application. The following shows the port for FTP. Note that FTP can use port 21 with either UDP or TCP.

```
$ grep ftp /etc/services
ftp      21/tcp
ftp      21/udp
```

10

Socket address



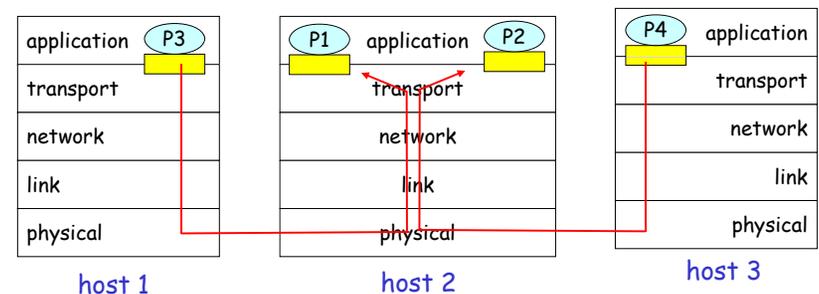
11

Multiplexing/demultiplexing

Demultiplexing at rcv host:
delivering received segments to correct socket

Multiplexing at send host:
gathering data from multiple sockets, enveloping data with header (later used for demultiplexing)

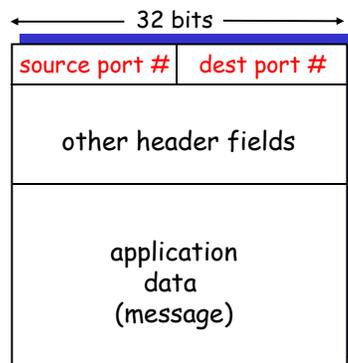
■ = socket ○ = process



12

How demultiplexing works

- host receives IP datagrams
 - each datagram has source IP address, destination IP address
 - each datagram carries 1 transport-layer segment
 - each segment has source, destination port number
- host uses IP addresses & port numbers to direct segment to appropriate socket



TCP/UDP segment format

Connectionless demultiplexing

- Create sockets with port numbers:

```
DatagramSocket mySocket1 = new
    DatagramSocket( );
DatagramSocket mySocket2 = new
    DatagramSocket(12535);
```

- UDP socket identified by two-tuple:
(dest IP address, dest port number)

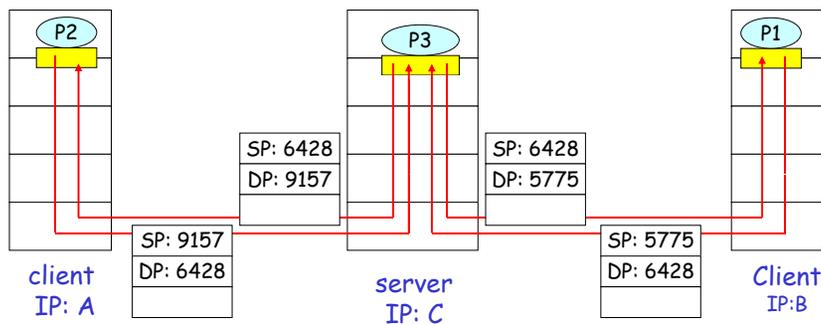
- When host receives UDP segment:

- checks destination port number in segment
- directs UDP segment to socket with that port number

- IP datagrams with different source IP addresses and/or source port numbers directed to same socket

Connectionless demux (cont)

```
DatagramSocket serverSocket = new DatagramSocket(6428);
```



SP provides "return address"

Connection-oriented demux

- TCP socket identified by 4-tuple:
 - source IP address
 - source port number
 - dest IP address
 - dest port number
- recv host uses all four values to direct segment to appropriate socket

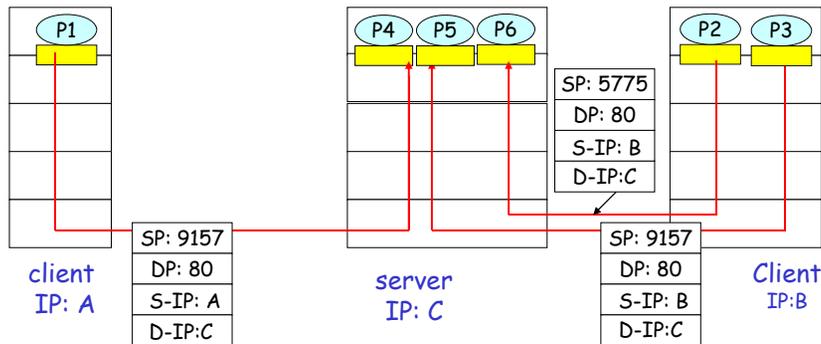
- Server host may support many simultaneous TCP sockets:

- each socket identified by its own 4-tuple

- Web servers have different sockets for each connecting client

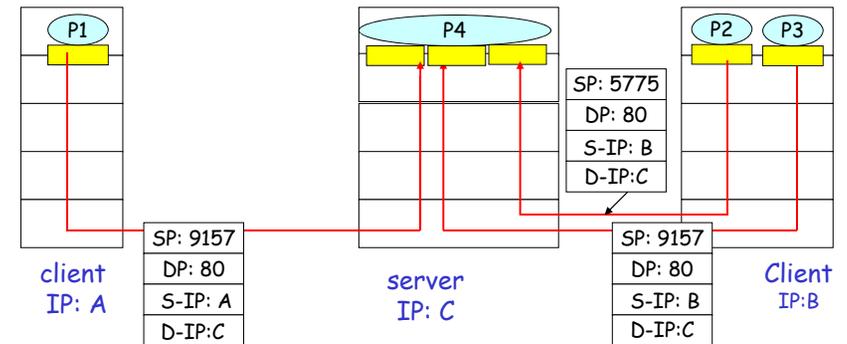
- non-persistent HTTP will have different socket for each request

Connection-oriented demux (cont)



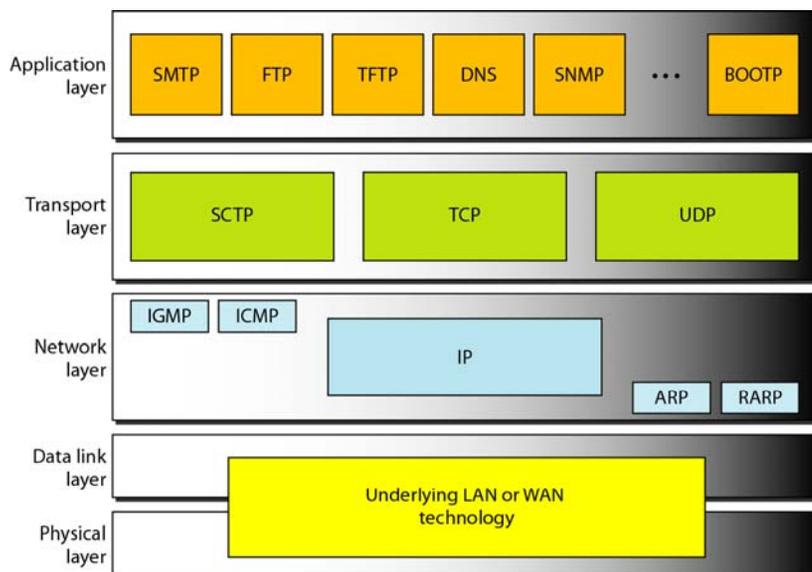
17

Connection-oriented demux: Threaded Web Server



18

Position of UDP, TCP, and SCTP in TCP/IP suite



19

OSI Model

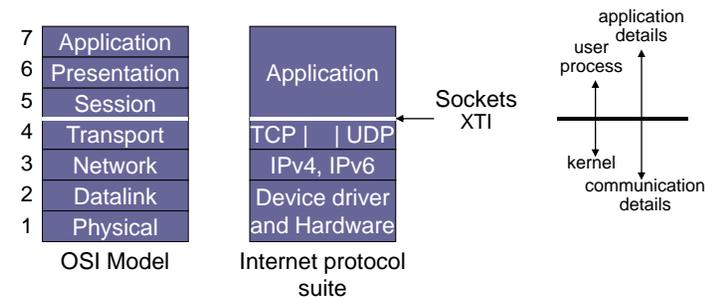


Figure 1.14 Layers on OSI model and Internet protocol suite

- Why do both sockets and XTI provide the interface from the upper three layers of the OSI model into the transport layer?
 - First, the upper three layers handle all the details of the application and the lower four layers handle all the communication details.
 - Second, the upper three layers are called a user process while the lower four layers are provided as part of the operating system kernel.

20

UDP: User Datagram Protocol [RFC 768]

- "no frills," "bare bones" Internet transport protocol
- "best effort" service, UDP segments may be:
 - lost
 - delivered out of order to app
- **connectionless:**
 - no handshaking between UDP sender, receiver
 - each UDP segment handled independently of others

21

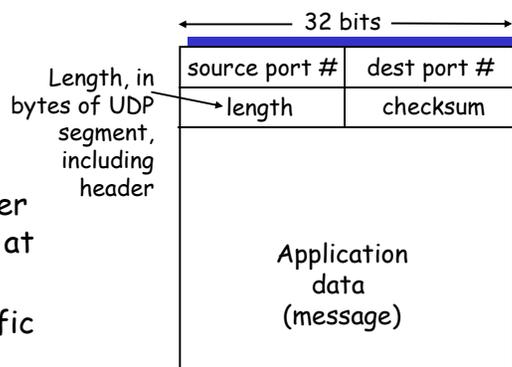
Why is there a UDP?

- no connection establishment (which can add delay)
- simple: no connection state at sender, receiver
- small segment header
- no congestion control: UDP can blast away as fast as desired
- often used for streaming multimedia apps
 - loss tolerant
 - rate sensitive
- other UDP uses → DNS and SNMP
- Suitable for multicasting

22

UDP: more

- reliable transfer over UDP: add reliability at application layer
 - application-specific error recovery!



UDP segment format

23

UDP checksum

Goal: detect "errors" (e.g., flipped bits) in transmitted segment

Sender:

- treat segment contents as sequence of 16-bit integers
- checksum: addition (1's complement sum) of segment contents
- sender puts checksum value into UDP checksum field

Receiver:

- compute checksum of received segment
- check if computed checksum equals checksum field value:
 - NO - error detected
 - YES - no error detected. *But maybe errors nonetheless? More later*
-

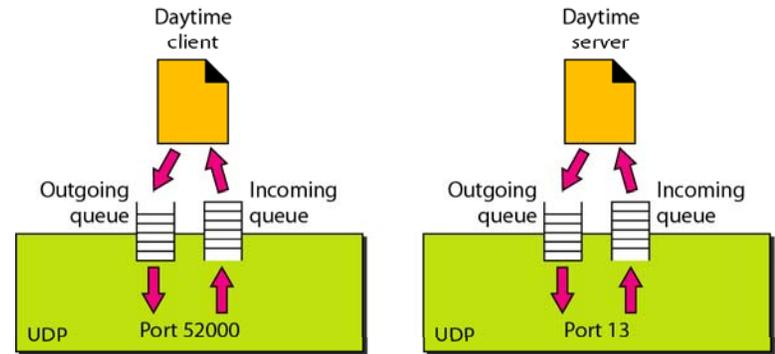
24

Internet Checksum Example

- Note
 - When adding numbers, a carryout from the most significant bit needs to be added to the result
- Example: add two 16-bit integers

	1 1 1 0 0 1 1 0 0 1 1 0 0 1 1 0
	1 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1
wraparound	1 1 0 1 1 1 0 1 1 1 0 1 1 1 0 1 1
sum	1 0 1 1 1 0 1 1 1 0 1 1 1 1 0 0
checksum	0 1 0 0 0 1 0 0 0 1 0 0 0 0 1 1

Queues in UDP



UDP output

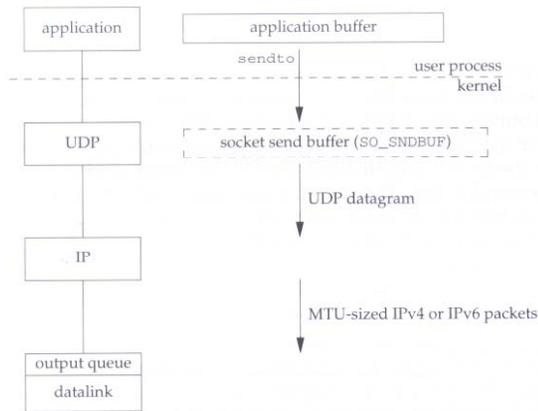


Figure 2.12 Steps and buffers involved when application writes to a UDP socket.

UDP is suitable for:

- A process that requires simple request-response communication with little concern for flow and error control - FTP
- A process with internal flow and error control mechanisms - TFTP
- Multicasting
- Management processes - SNMP
- Some rout updating protocols - RIP