

Transport Layer

EECS 3214

Slides courtesy of J.F. Kurose and K.W. Ross, All Rights Reserved

29-Jan-18

1-1

Chapter 3: Transport Layer

our goals:

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

Transport Layer 3-2

Chapter 3 outline

3.1 transport-layer services

3.2 multiplexing and demultiplexing

3.3 connectionless transport: UDP

3.4 principles of reliable data transfer

3.5 connection-oriented transport: TCP

- segment structure
- reliable data transfer
- flow control
- connection management


3.6 principles of congestion control

3.7 TCP congestion control

Transport Layer 3-3

Transport services and protocols

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



Transport Layer 3-4

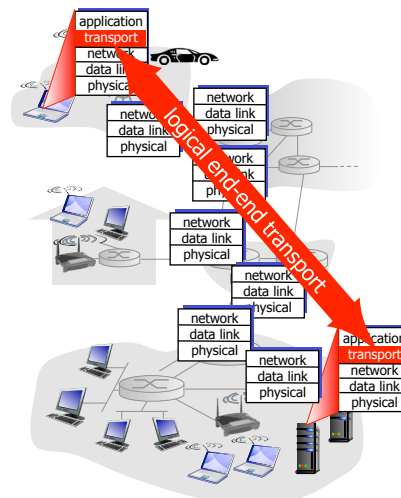
Transport vs. network layer

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

Transport Layer 3-5

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



Transport Layer 3-6

Chapter 3 outline

3.1 transport-layer services

3.2 multiplexing and demultiplexing

3.3 connectionless transport: UDP

3.4 principles of reliable data transfer

3.5 connection-oriented transport: TCP

- segment structure
- reliable data transfer
- flow control
- connection management

3.6 principles of congestion control

3.7 TCP congestion control

Transport Layer 3-7

Terminology

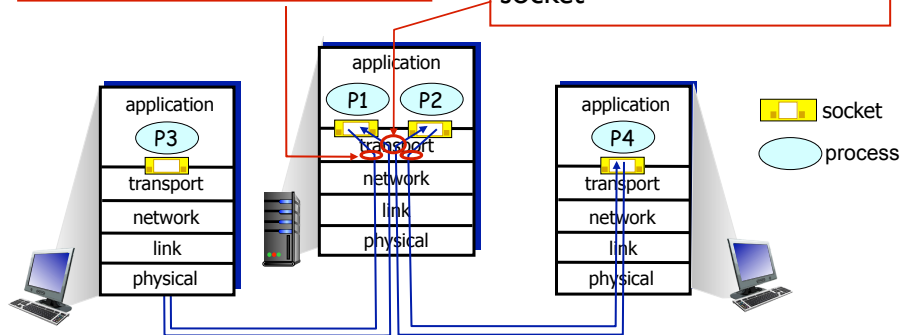
- Port: external end-point at a node
- **Socket**: internal end-point of local inter-process communication
- In Internet protocols, **socket address** = IP address + port number
- Berkeley sockets: API for Internet sockets (Unix)
 - socket represented as a file descriptor (**socket descriptor**)
 - `socket ()`: creates a new socket and gives it a socket descriptor
 - `bind ()`: associates a socket descriptor with a socket address
- An Internet socket is characterized by at least
 - socket address (IP address + port number)
 - protocol (TCP, UDP)
 - TCP port 1234 and UDP port 1234 are distinct sockets

Transport Layer 3-8

Multiplexing/demultiplexing

multiplexing at sender:
handle data from multiple sockets, add transport header (later used for demultiplexing)

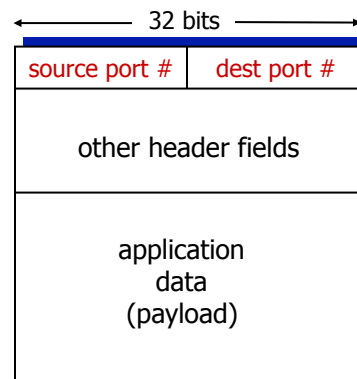
demultiplexing at receiver:
use header info to deliver received segments to correct socket



Transport Layer 3-9

How demultiplexing works

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



TCP/UDP segment format

Transport Layer 3-10

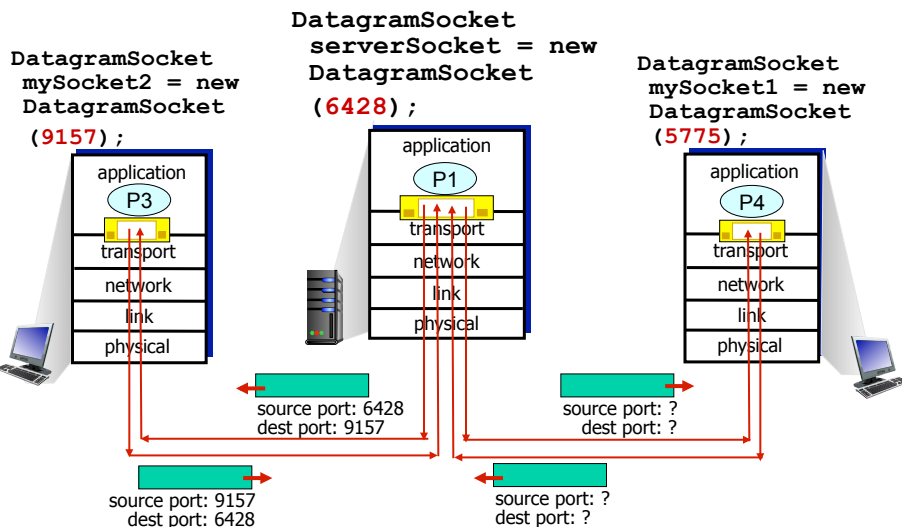
Connectionless demultiplexing

- *recall:* created socket has host-local port #:


```
DatagramSocket mySocket1
= new DatagramSocket(12534);
```
 - *recall:* when creating datagram to send into UDP socket, must specify
 - destination IP address
 - destination port #
 - when host receives UDP segment:
 - checks destination port # in segment
 - directs UDP segment to socket with that port #
- IP datagrams with *same dest. port #*, but different source IP addresses and/or source port numbers will be directed to *same socket* at dest

Transport Layer 3-11

Connectionless demux: example



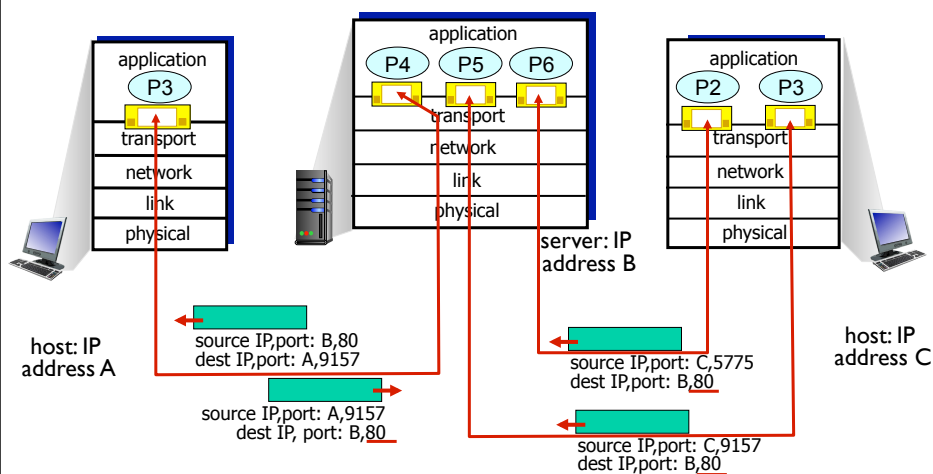
Transport Layer 3-12

Connection-oriented demux

- TCP socket identified by 4-tuple:
 - source IP address
 - source port number
 - dest IP address
 - dest port number
- demux: receiver 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

Transport Layer 3-13

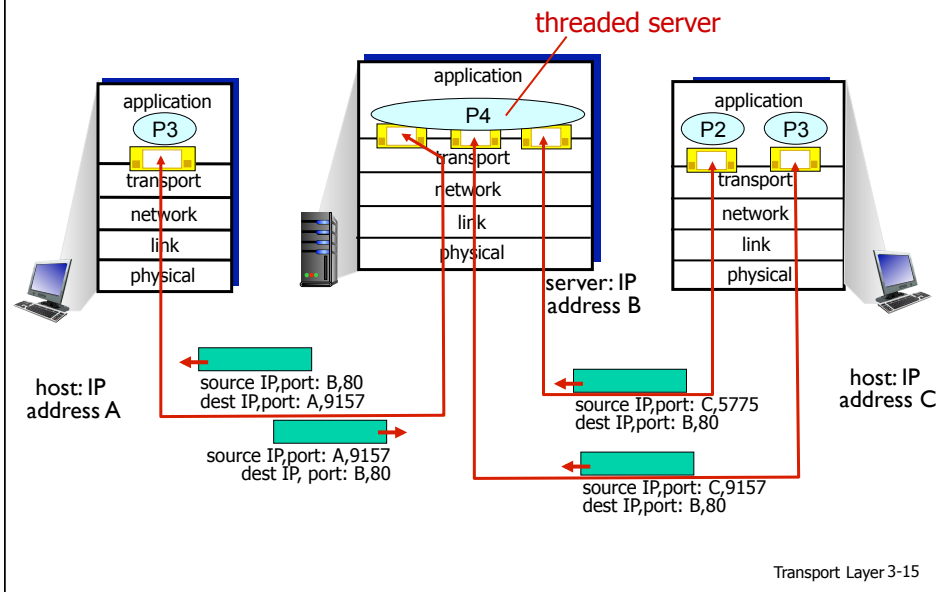
Connection-oriented demux: example



three segments, all destined to IP address: B,
dest port: 80 are demultiplexed to *different* sockets

Transport Layer 3-14

Connection-oriented demux: example



Chapter 3 outline

3.1 transport-layer services

3.2 multiplexing and demultiplexing

3.3 connectionless transport: UDP

3.4 principles of reliable data transfer

3.5 connection-oriented transport: TCP

- segment structure
- reliable data transfer
- flow control
- connection management

3.6 principles of congestion control

3.7 TCP congestion control

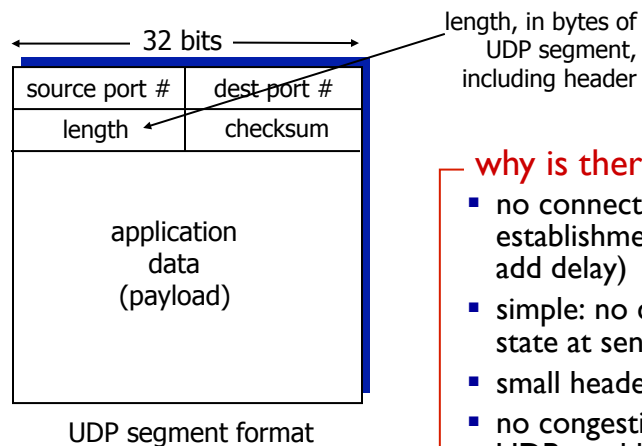
Transport Layer 3-16

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
- UDP use:
 - streaming multimedia apps (loss tolerant, rate sensitive)
 - DNS
 - SNMP
- reliable transfer over UDP:
 - add reliability at application layer
 - application-specific error recovery!

Transport Layer 3-17

UDP: segment header



why is there a UDP?

- no connection establishment (which can add delay)
- simple: no connection state at sender, receiver
- small header size (8 bytes)
- no congestion control: UDP can blast away as fast as desired

Transport Layer 3-18

UDP checksum

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

sender:

- treat segment contents, including header fields, as sequence of 16-bit integers
- checksum: addition (one'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
-

Transport Layer 3-19

Internet checksum: example

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
<hr/>																
wraparound	1	1	0	1	1	1	0	1	1	1	0	1	1	1	0	1
<hr/>																
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

Note: when adding numbers, a carryout from the most significant bit needs to be added to the result

* Check out the online interactive exercises for more examples: http://gaia.cs.umass.edu/kurose_ross/interactive/

Transport Layer 3-20

Chapter 3 outline

3.1 transport-layer services

3.2 multiplexing and demultiplexing

3.3 connectionless transport: UDP

3.4 principles of reliable data transfer (reading assignment)

3.5 connection-oriented transport: TCP

- segment structure
- reliable data transfer
- flow control
- connection management

3.6 principles of congestion control

3.7 TCP congestion control

Transport Layer 3-21

Chapter 3: summary

■ principles behind transport layer services:

- multiplexing, demultiplexing
- reliable data transfer
- flow control
- congestion control

■ instantiation, implementation in the Internet

- UDP
- TCP

next:

- leaving the network “edge” (application, transport layers)
- into the network “core”
- two network layer chapters:
 - data plane
 - control plane

Transport Layer 3-22