



# Transport Layer

## - Socket -

Kyunghan Lee

Networked Computing Lab (NXC Lab)

Department of Electrical and Computer Engineering

Seoul National University

<https://nxc.snu.ac.kr>

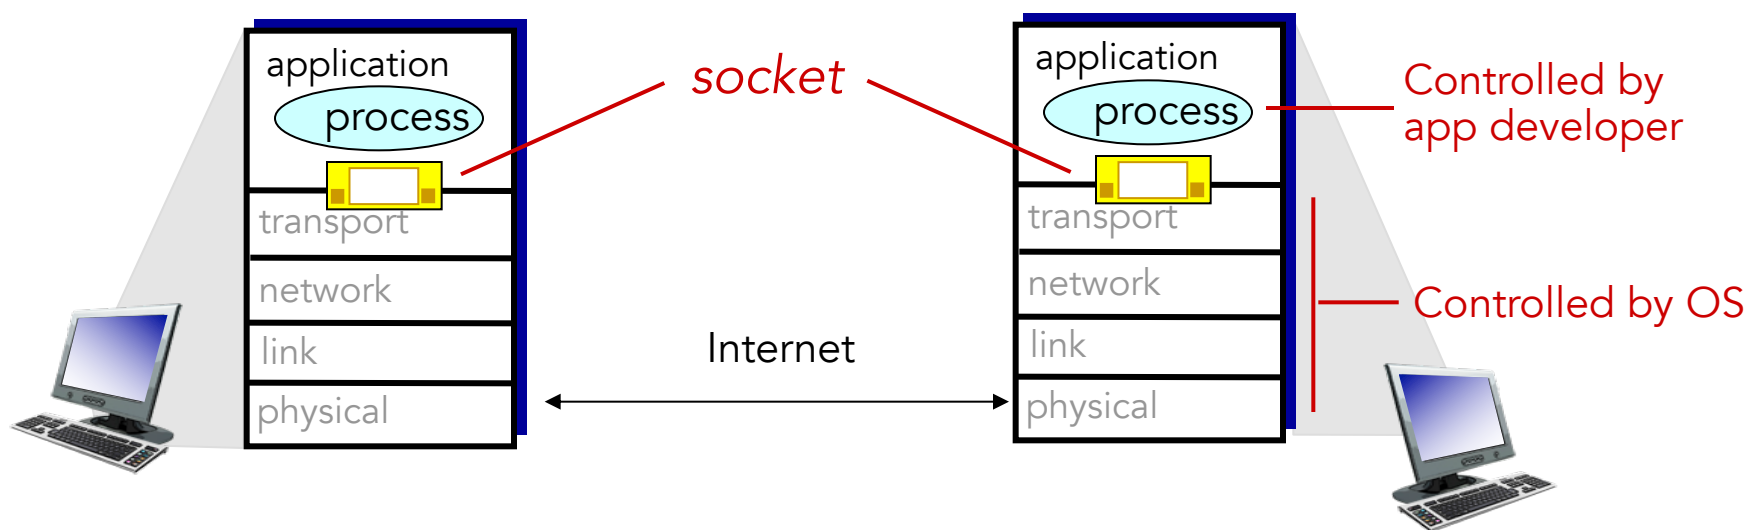
[kyunghanlee@snu.ac.kr](mailto:kyunghanlee@snu.ac.kr)



# Socket programming

*Goal:* learn how to build client/server applications that communicate using sockets

*Socket:* door between application process and end-to-end-transport protocol



# Socket programming

*Two socket types for two transport services:*

- *UDP*: unreliable datagram
- *TCP*: reliable, byte stream-oriented

*Application Example:*

1. client reads a line of characters (data) from its keyboard and sends data to server
2. server receives the data and converts characters to uppercase
3. server sends modified data to client
4. client receives modified data and displays line on its screen



# Socket programming

## UDP: no “connection” between client & server

- ❑ no handshaking before sending data
- ❑ sender explicitly attaches IP destination address and port # to each packet
- ❑ receiver extracts sender IP address and port# from received packet

## UDP: transmitted data may be lost or received out-of-order

## Application viewpoint:

- ❑ UDP provides *unreliable* transfer of groups of bytes (“datagrams”) between client and server



# Client/server socket interaction: UDP

## server (running on server IP)

create socket, port= x:  
`serverSocket =  
socket(AF_INET,SOCK_DGRAM)`

↓  
read datagram from  
`serverSocket`

↓  
write reply to  
`serverSocket`  
specifying  
client address,  
port number

## client

create socket:

`clientSocket =  
socket(AF_INET,SOCK_DGRAM)`

↓  
Create datagram with server IP and  
port=x; send datagram via  
`clientSocket`

↓  
read datagram from  
`clientSocket`

↓  
close  
`clientSocket`

# Socket programming with TCP

## Client must contact server

- server process must first be running
- server must have created socket (door) that welcomes client's contact

## Client contacts server by:

- Creating TCP socket, specifying IP address, port number of server process
- *when client creates socket:* client TCP establishes connection to server TCP

- When contacted by client, *server TCP creates new socket* for server process to communicate with that particular client
  - allows server to talk with multiple clients
  - source port numbers used to distinguish clients

## Application viewpoint:

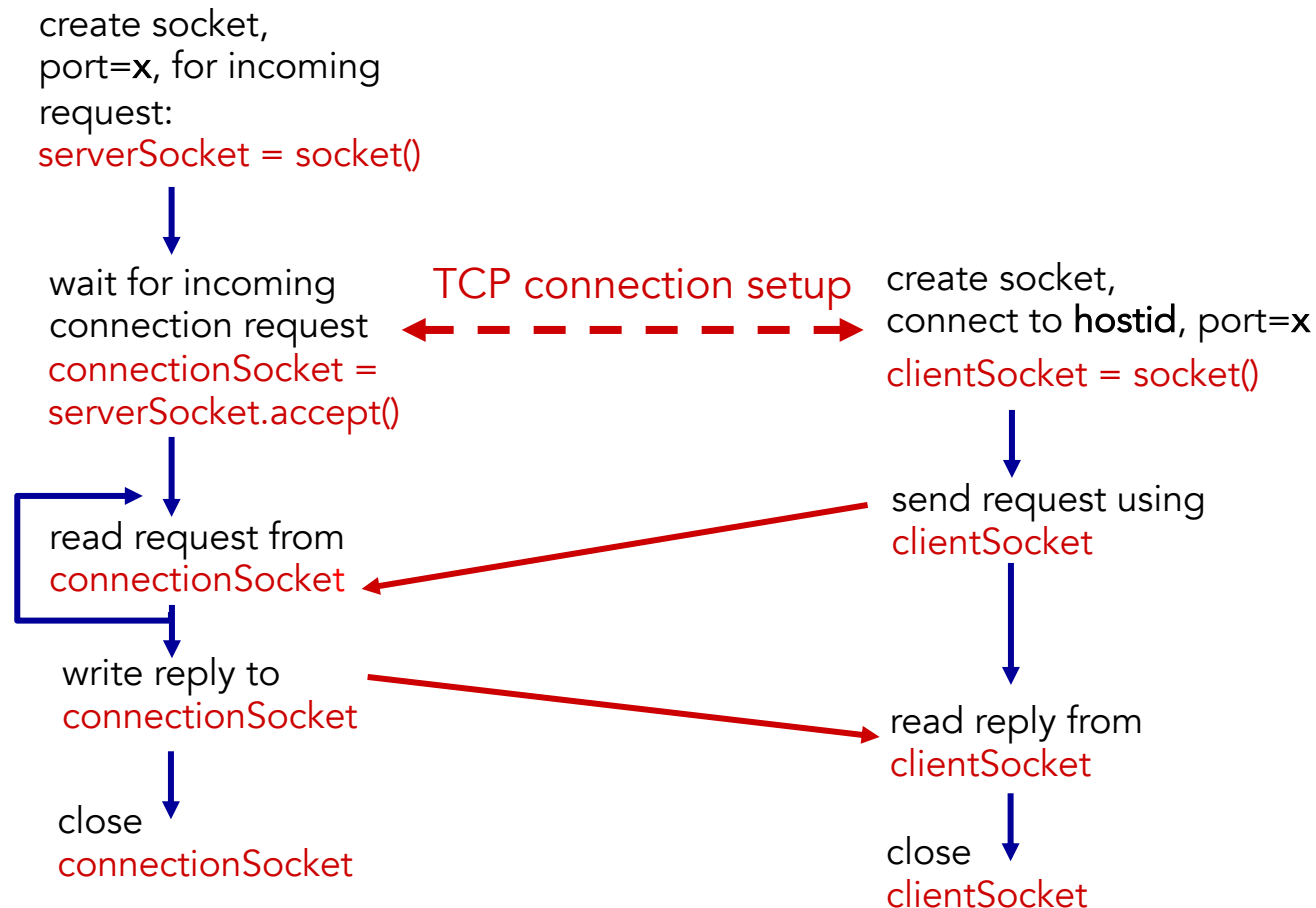
TCP provides reliable, in-order byte-stream transfer (“pipe”) between client and server



# Client/server socket interaction: TCP

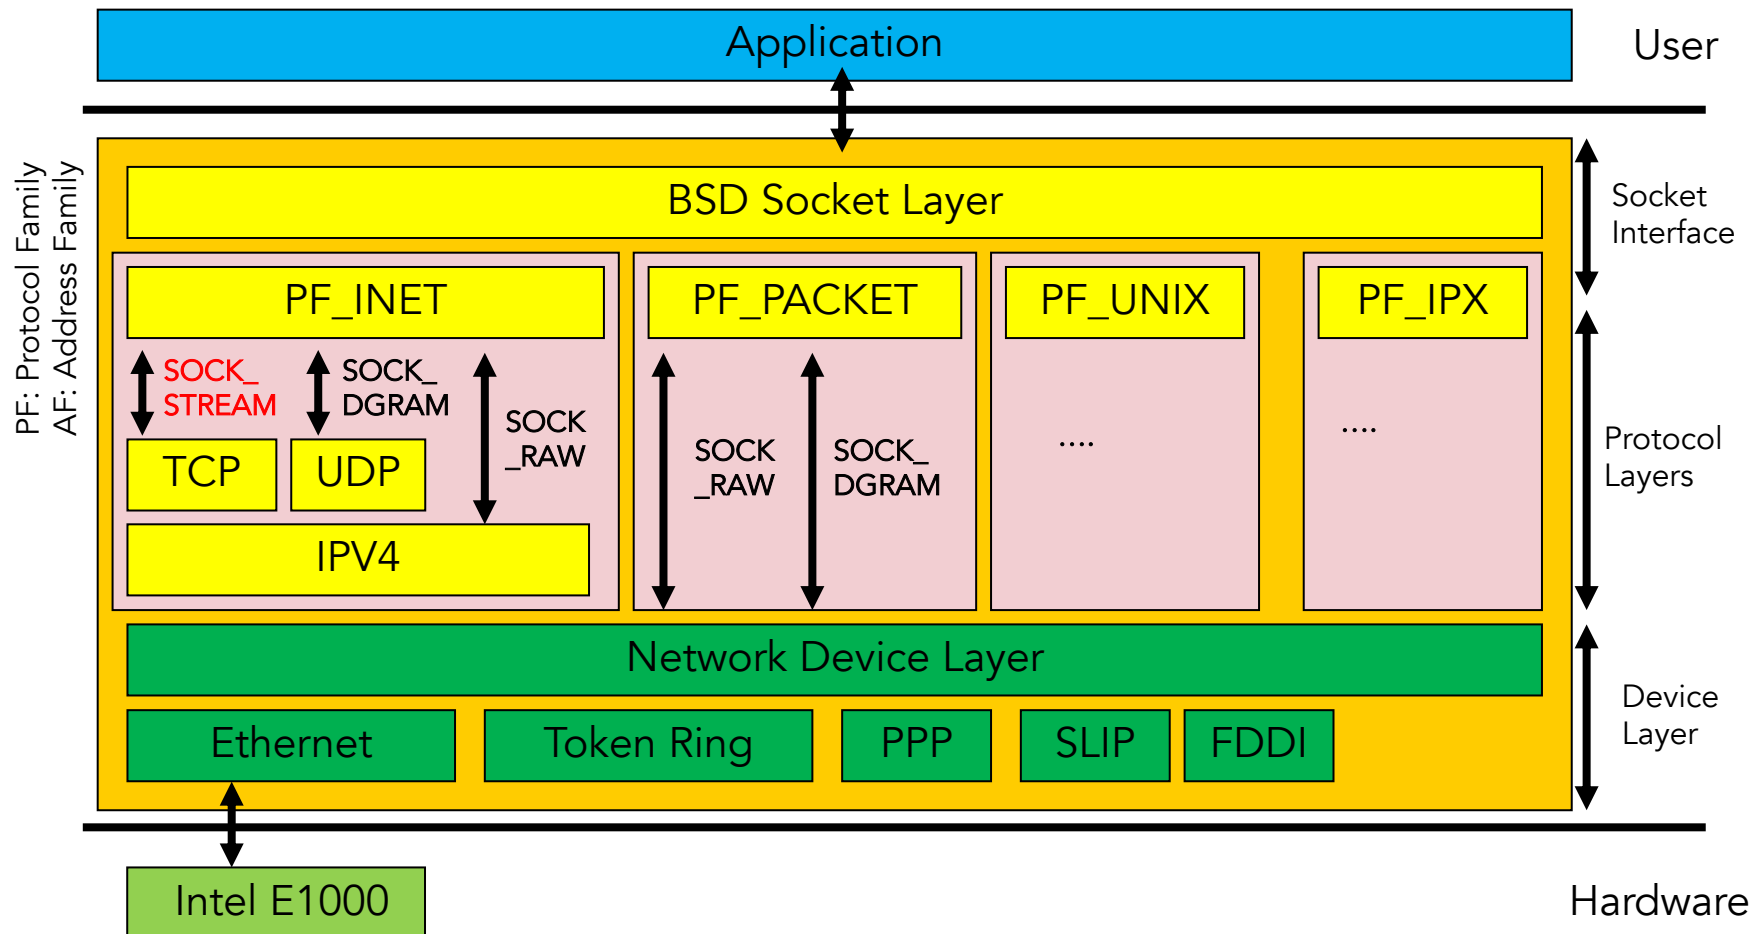
## server (running on host ID)

## client



# Socket: Application-TCP Interaction

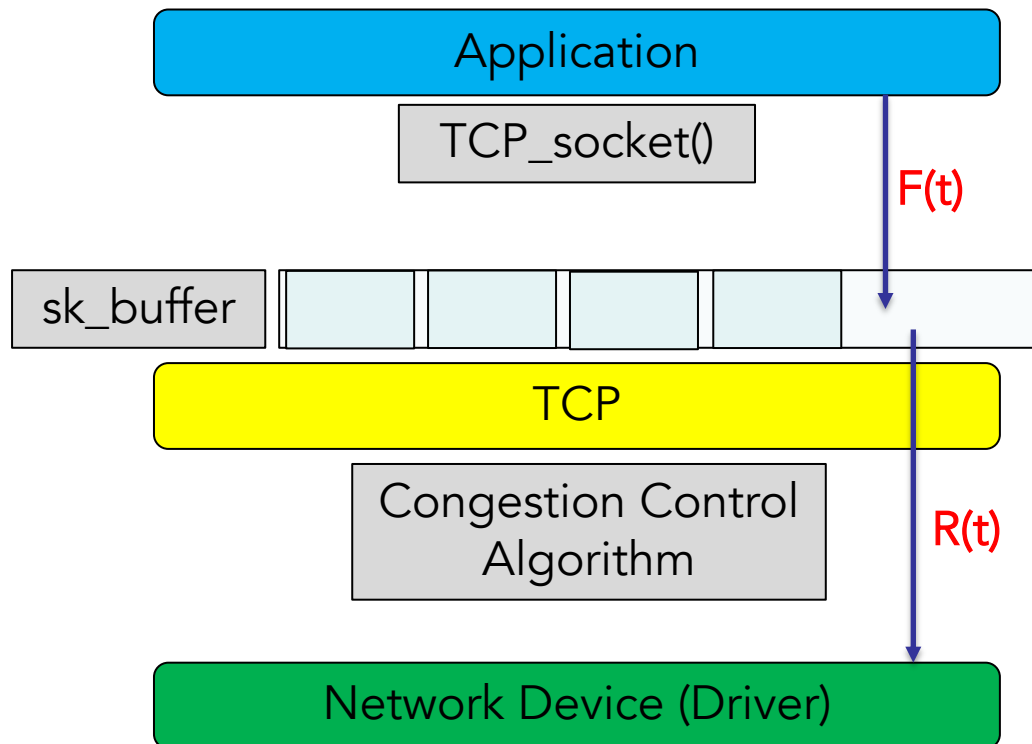
- Linux transport-layer implementation





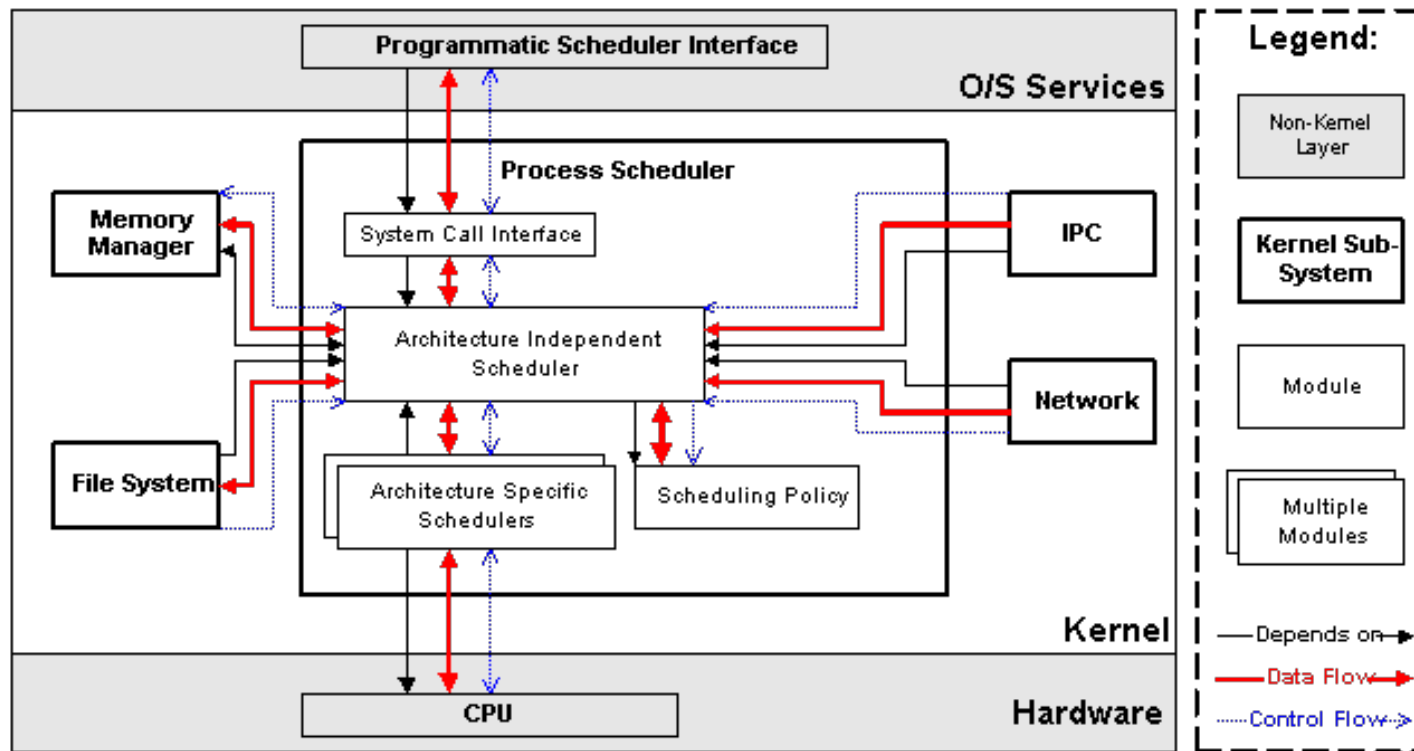
# Latency from Socket

- Transport-layer implementation with socket buffer
  - Imperfect synchronization between  $F(t)$  and  $R(t)$  in action




# Latency from Socket

- Process scheduling latency
  - Determined by Linux Kernel's CFS scheduler
  - Order of some milliseconds


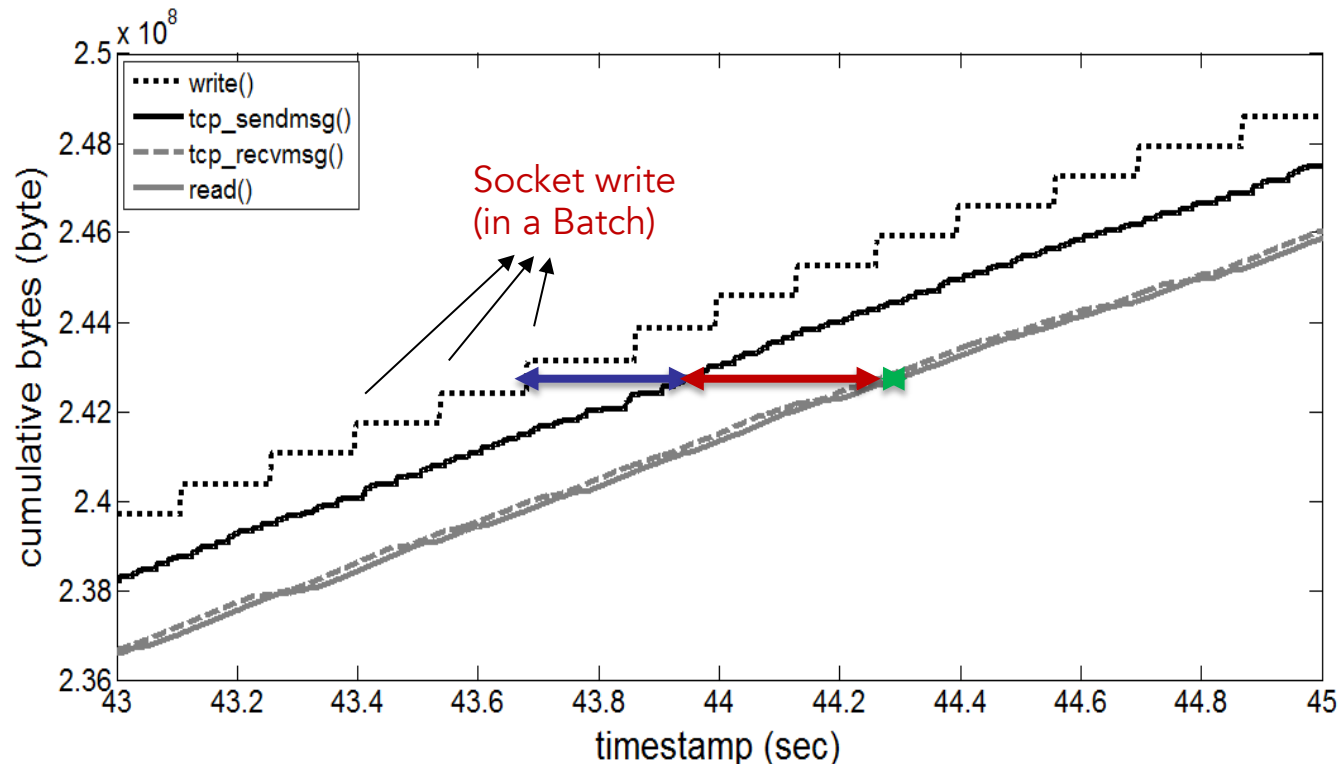


# Latency from Socket

- D1: App to TCP
- D2: TCP to TCP
- D3: TCP to App



D: App to App Latency





# Transport Layer

Kyunghan Lee

Networked Computing Lab (NXC Lab)

Department of Electrical and Computer Engineering

Seoul National University

<https://nxc.snu.ac.kr>

[kyunghanlee@snu.ac.kr](mailto:kyunghanlee@snu.ac.kr)



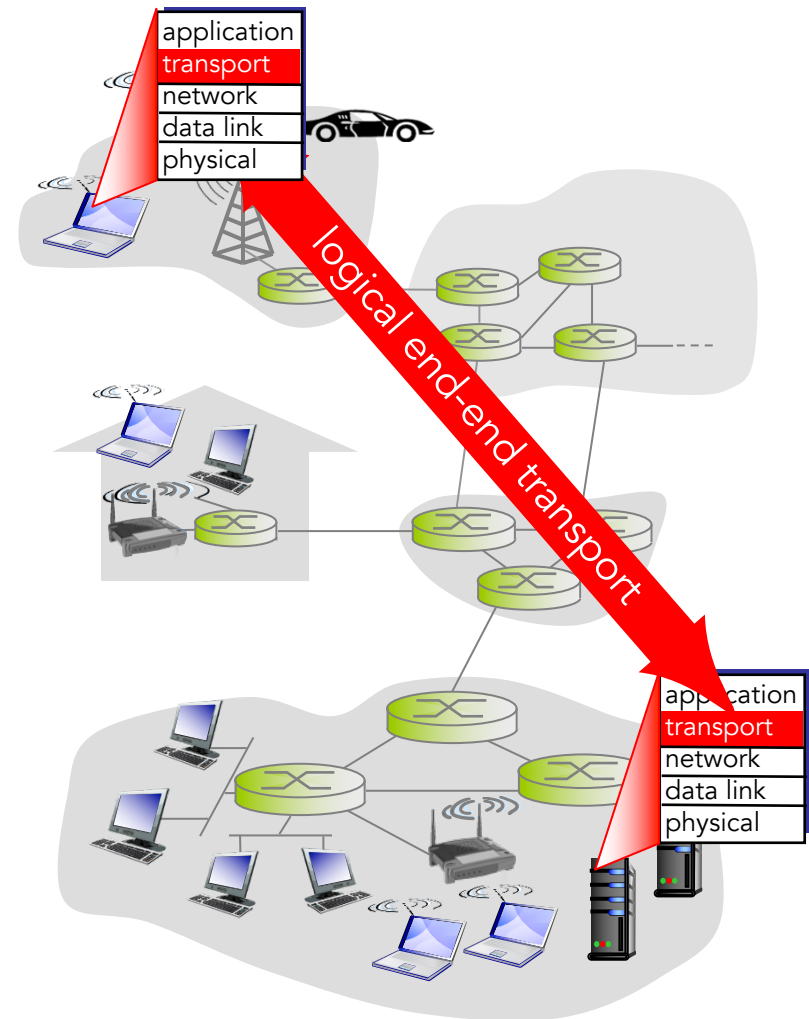
# Objectives

- understand principles behind transport layer services:
  - multiplexing and 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



# 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



# Transport vs. Network layer

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

## *household analogy:*

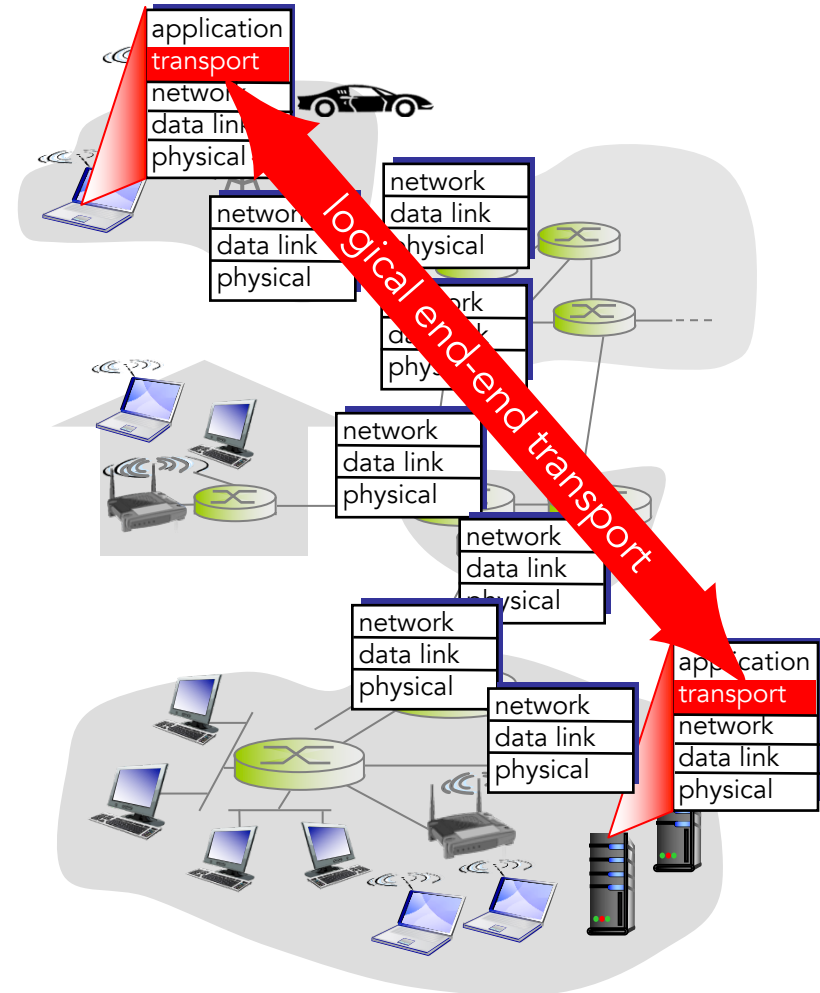
*12 kids in Ann's house sending letters to 12 kids in Bill's house:*

- hosts = houses
- processes = kids
- app messages = letters in envelopes
- transport protocol = Ann and Bill who demux to in-house siblings
- network-layer protocol = postal service



# 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





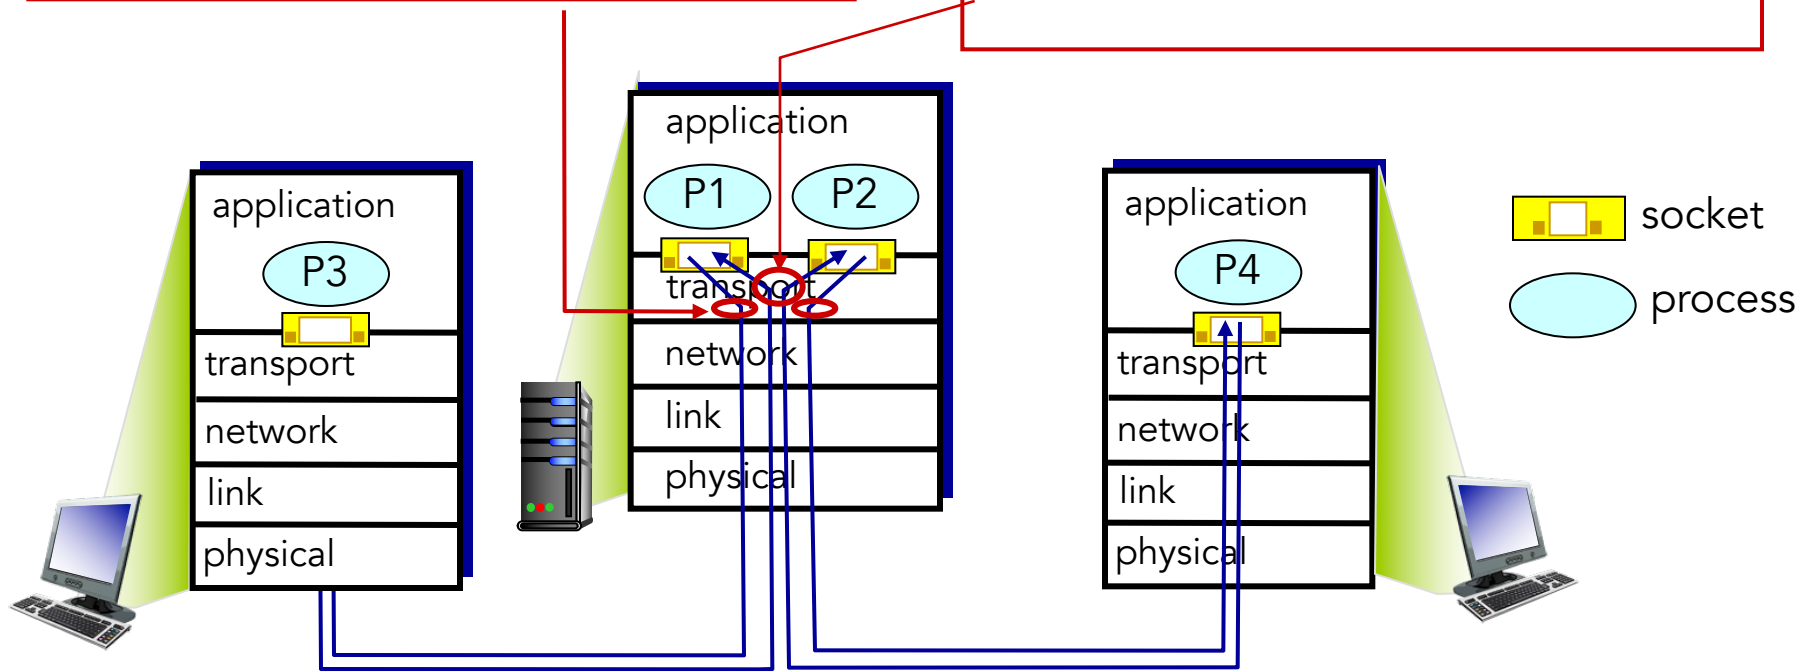
# Multiplexing/Demultiplexing

## *multiplexing at sender:*

handle data from multiple sockets, add transport header (later used for demultiplexing)

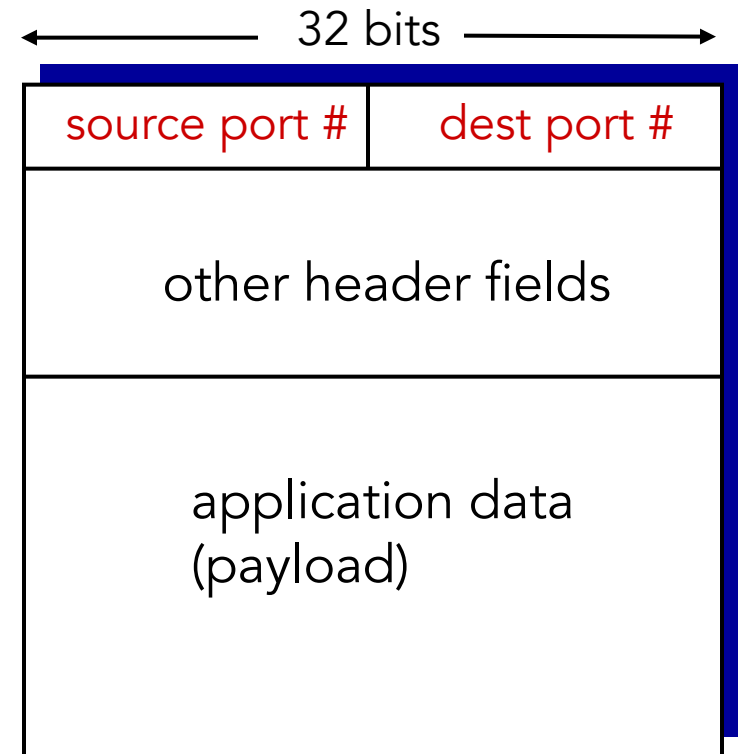
## *demultiplexing at receiver:*

use header info to deliver received segments to correct socket



# 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 & port numbers* to direct segment to appropriate socket



TCP/UDP segment format

# Connectionless demultiplexing

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

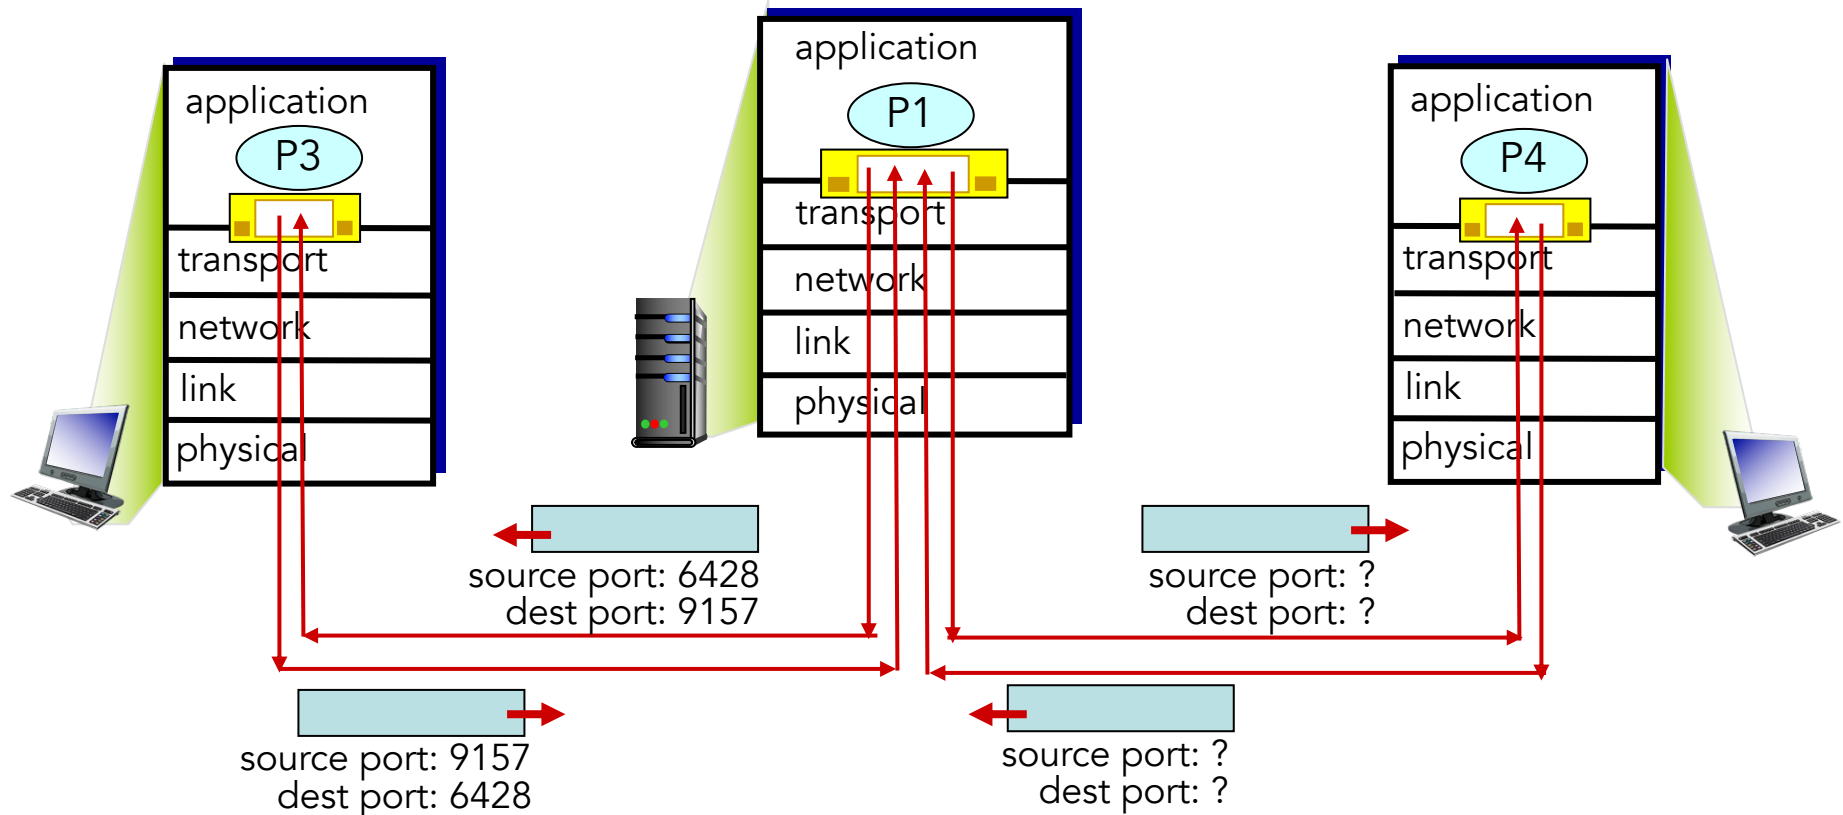
```
DatagramSocket mySocket1  
= new DatagramSocket(12534);
```
  - *recall*: when creating datagram to send into UDP socket, must specify
    - destination IP address
    - destination port #
- 
- when host receives UDP segment:
    - checks destination port # in segment
    - directs UDP segment to 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.

# Connectionless demux: Example

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

```
DatagramSocket mySocket2
= new DatagramSocket(9157);
```

```
DatagramSocket mySocket1
= new DatagramSocket(5775);
```

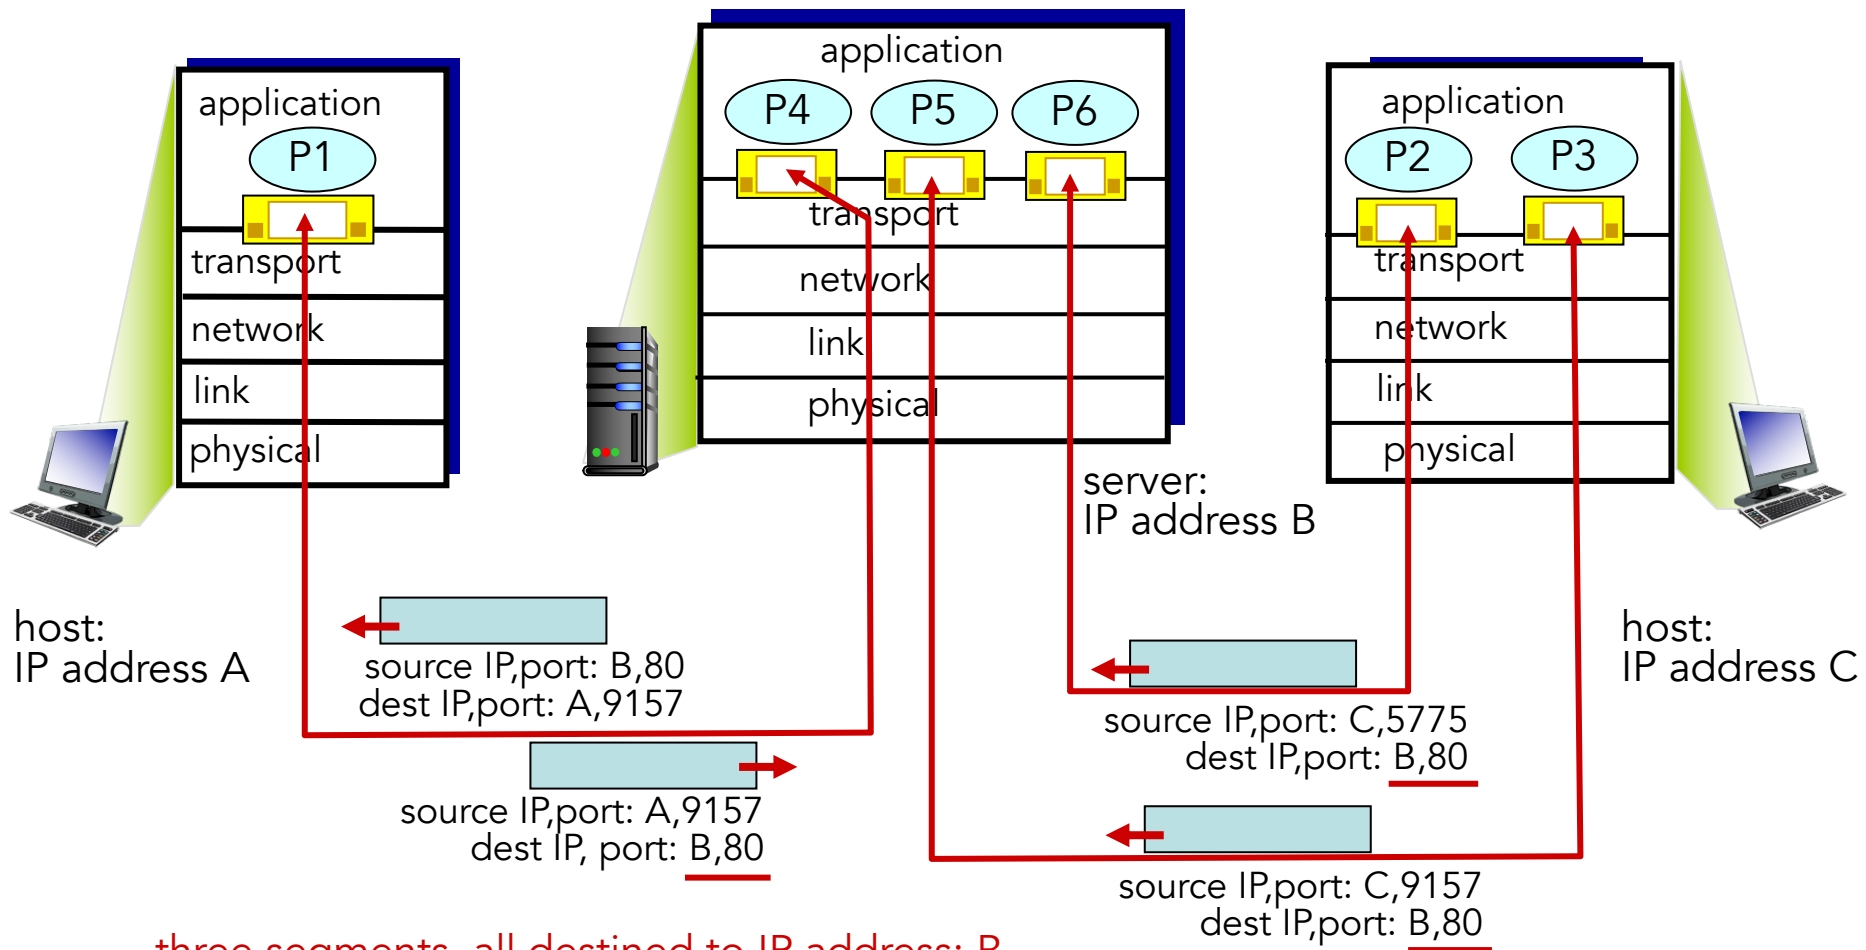


# Connection-oriented demux

- 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

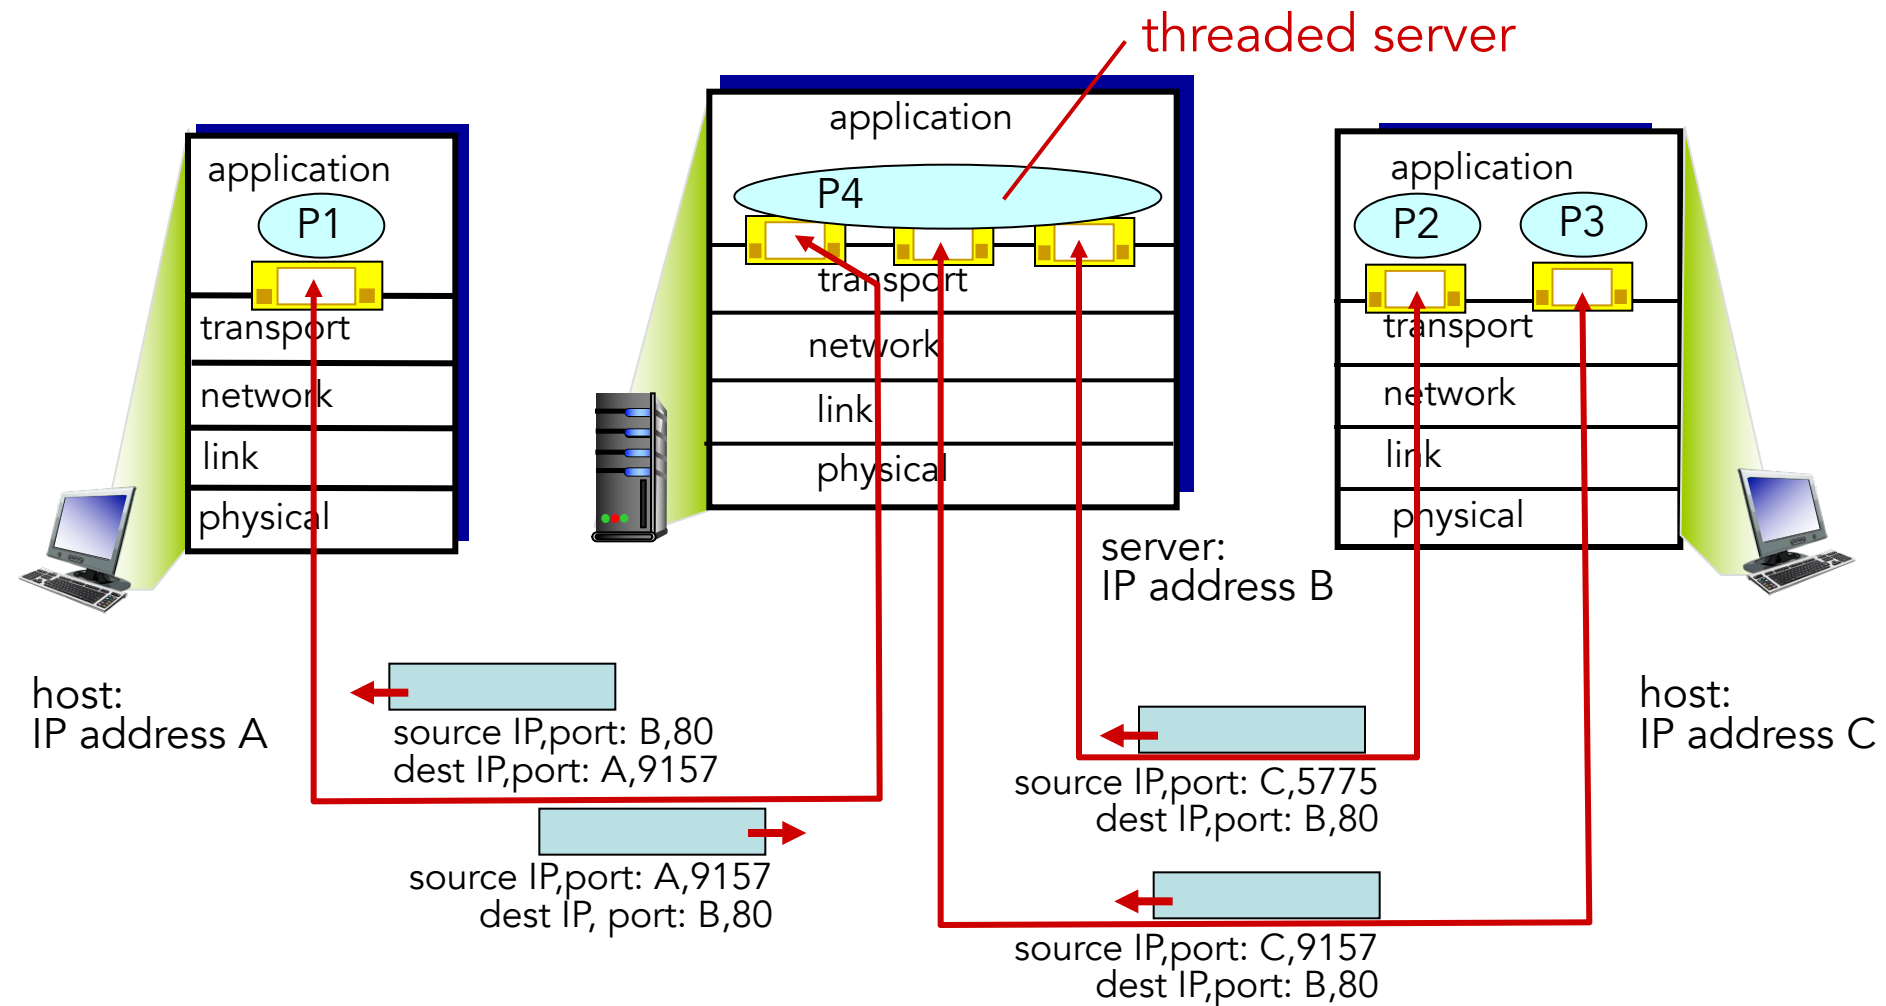


# Connection-oriented demux: example



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

# Connection-oriented demux: example



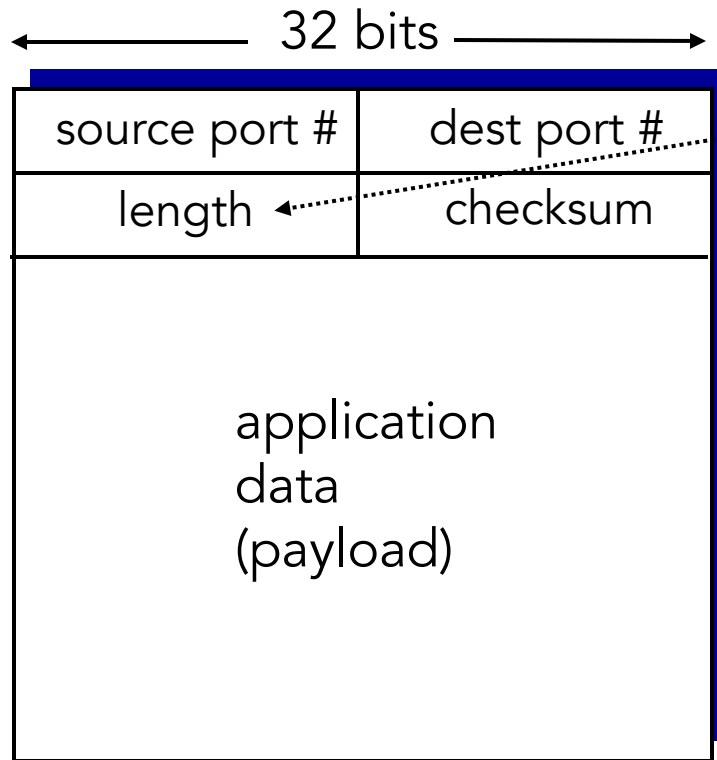
# UDP: User Datagram Protocol [RFC 768]

- “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!





# UDP: segment header



length, in bytes of UDP segment, including header

UDP segment format

## why is there a UDP?

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

# 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 ....*




# 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	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

# Reading Assignment #3 – Chapter

Quiz #2: November 2<sup>nd</sup> (4~5 questions)



## Chapter 3 Transport Layer

<b>Chapter 3</b>	<b>Transport Layer</b>	<b>215</b>
3.1	Introduction and Transport-Layer Services	216
3.1.1	Relationship Between Transport and Network Layers	216
3.1.2	Overview of the Transport Layer in the Internet	219
3.2	Multiplexing and Demultiplexing	221
3.3	Connectionless Transport: UDP	228
3.3.1	UDP Segment Structure	232
3.3.2	UDP Checksum	232
3.4	Principles of Reliable Data Transfer	234
3.4.1	Building a Reliable Data Transfer Protocol	236
3.4.2	Pipelined Reliable Data Transfer Protocols	245
3.4.3	Go-Back-N (GBN)	249
3.4.4	Selective Repeat (SR)	254
3.5	Connection-Oriented Transport: TCP	261
3.5.1	The TCP Connection	261
3.5.2	TCP Segment Structure	264
3.5.3	Round-Trip Time Estimation and Timeout	269
3.5.4	Reliable Data Transfer	272
3.5.5	Flow Control	280
3.5.6	TCP Connection Management	283
3.6	Principles of Congestion Control	289
3.6.1	The Causes and the Costs of Congestion	289
3.6.2	Approaches to Congestion Control	296
3.7	TCP Congestion Control	297
3.7.1	Fairness	307
3.7.2	Explicit Congestion Notification (ECN): Network-assisted Congestion Control	310
3.8	Summary	312
	Homework Problems and Questions	314
	Programming Assignments	329
	Wireshark Labs: Exploring TCP, UDP	330
	Interview: Van Jacobson	331