Introduction to Networking and the Internet

Where are we?

- Function calls, system calls, threads and processes

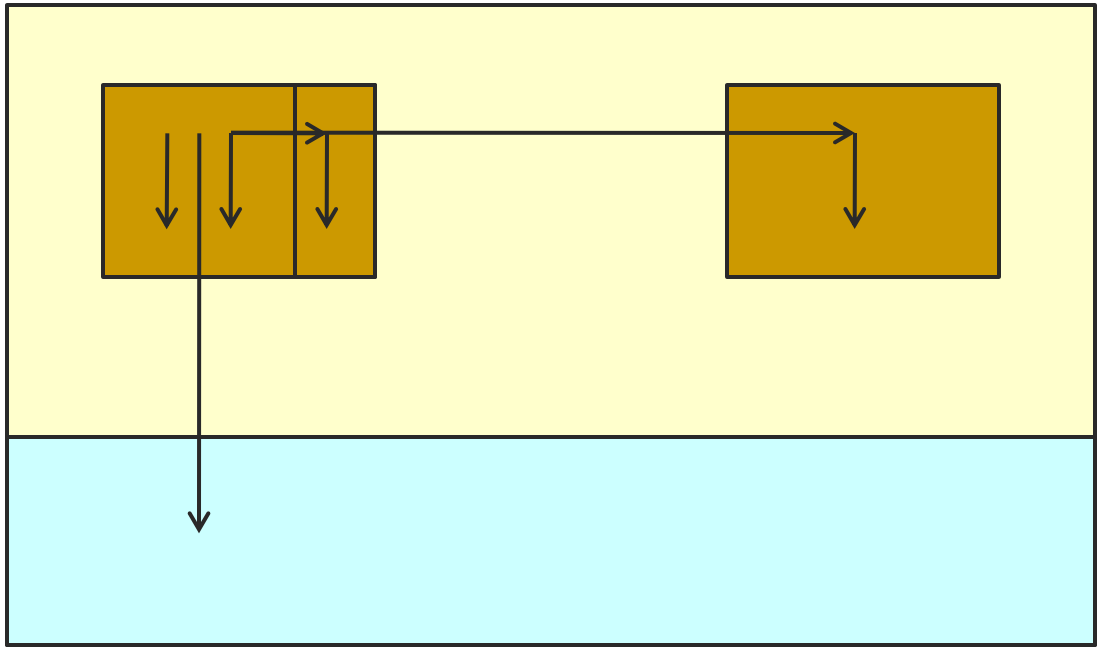

What's next?

- Networked communication and distributed applications

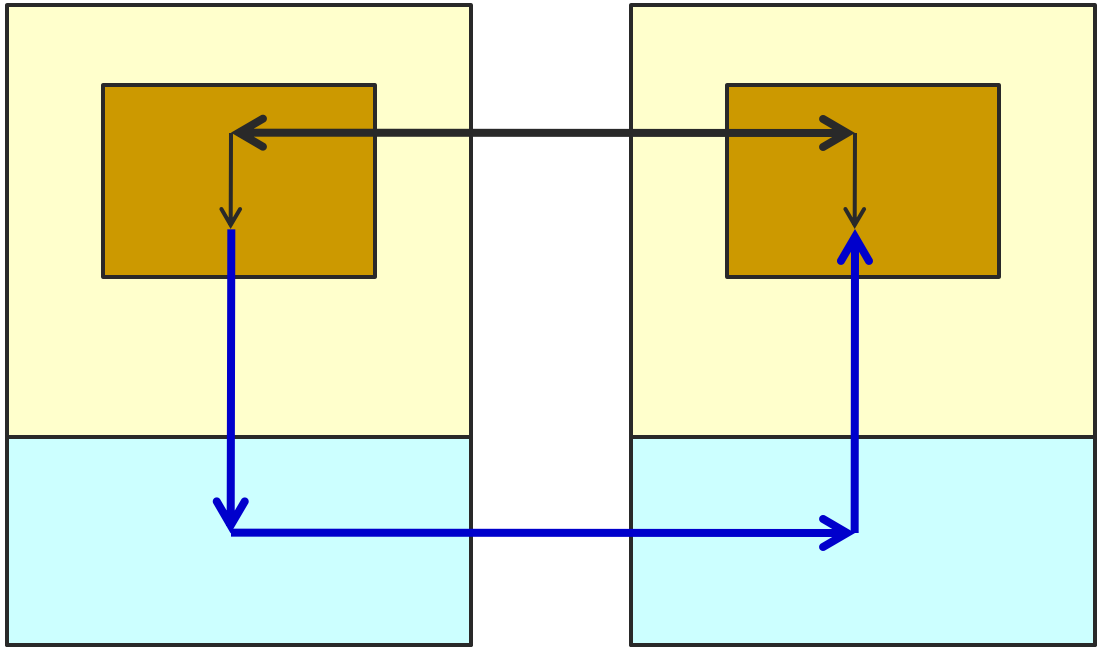

Introduction

- What is the Internet?
- Network edge
- What is a protocol?
- Protocol layers, service models

What is the Internet?

- Communication infrastructure
  - Enables distributed applications
  - Web, VoIP, email, games, e-commerce, file sharing
- Communication services
  - Provided to applications
  - Reliable data delivery from source to destination
  - “best effort” (unreliable) data delivery

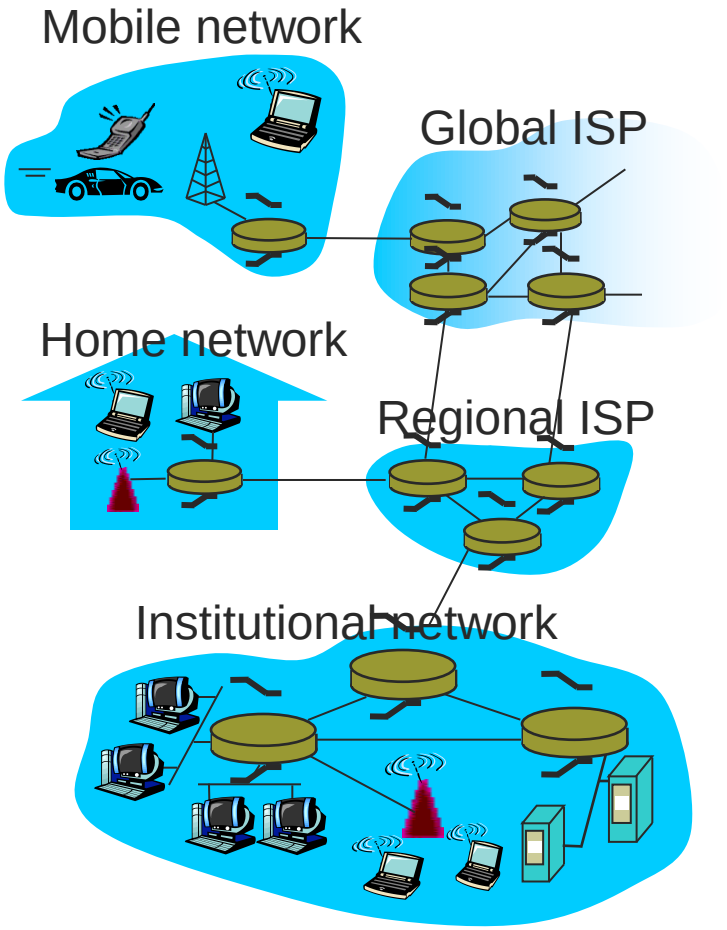

Network Service

- Goal
  - Transfer data between end systems
- Support For Common Services
  - Simplify the role of applications
  - Hide the complexity of the network
  - Semantics and interface depend on applications

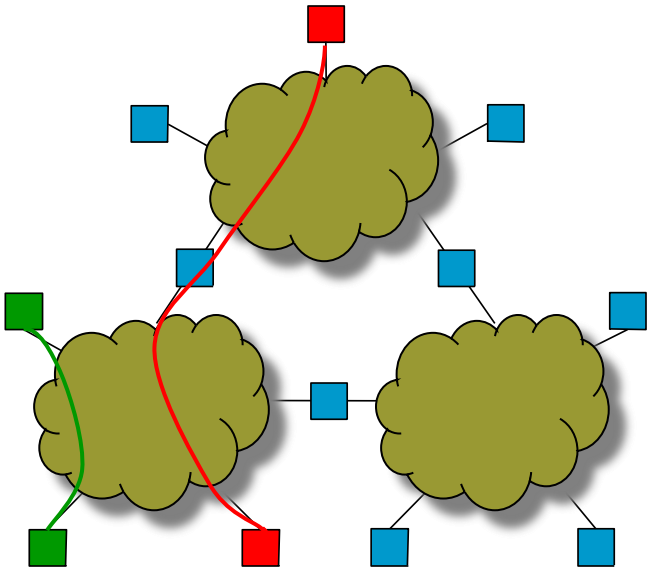

Example: Sending a Letter

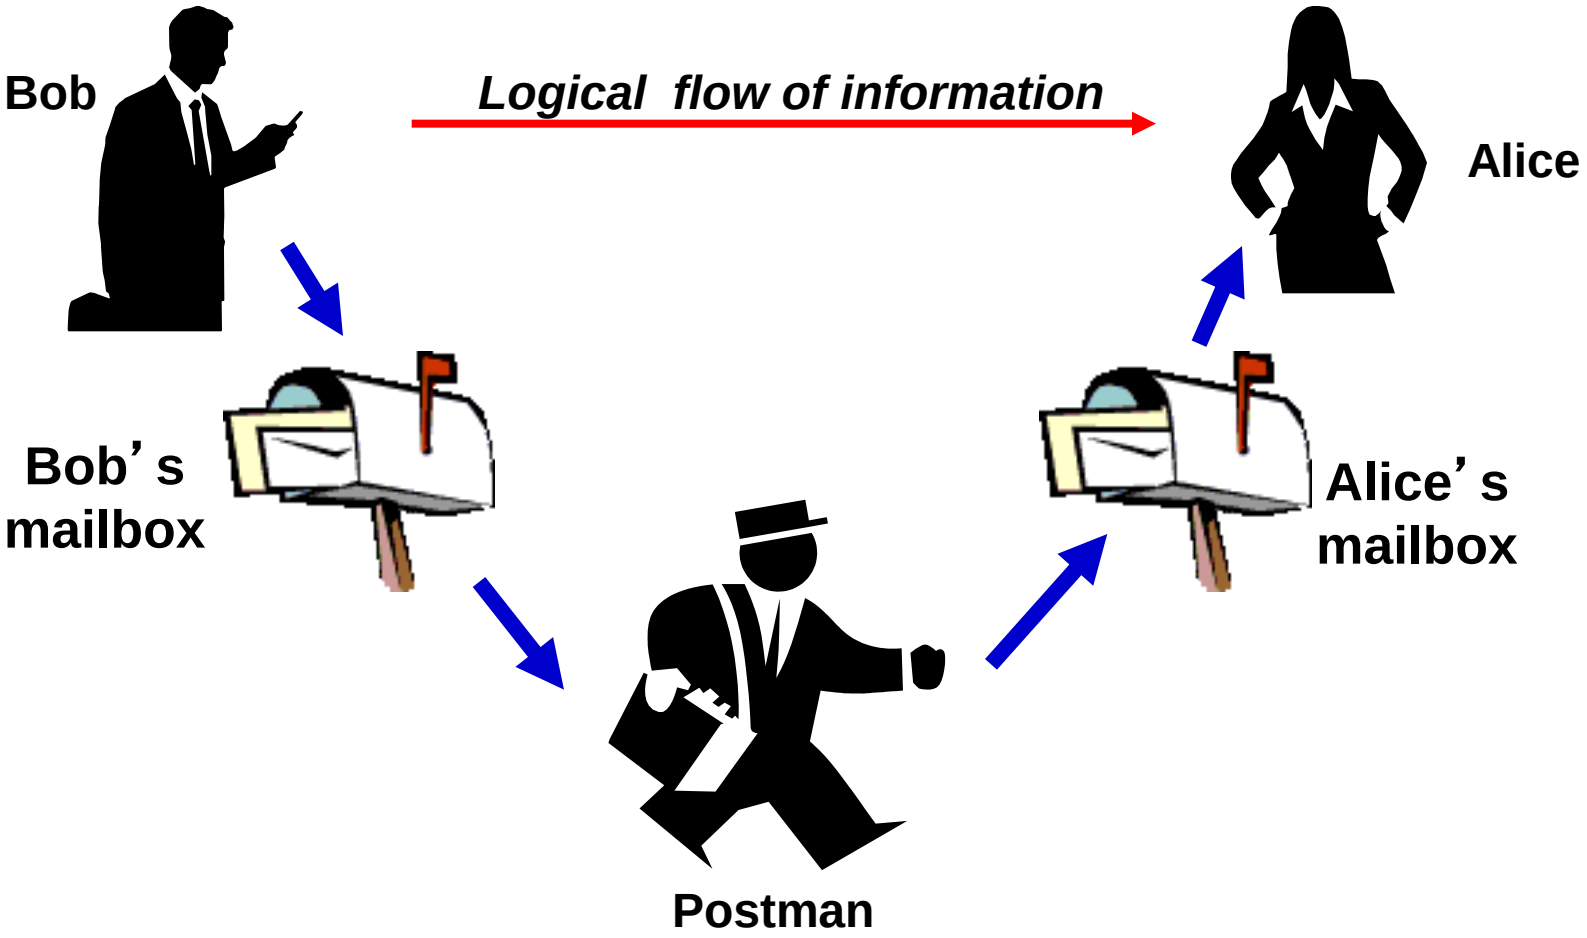

Services

- Unconfirmed service

- Acknowledged service

Channels

- Channel
  - The abstraction for application-level communication
- Idea
  - Turn host-to-host connectivity into process-to-process communication

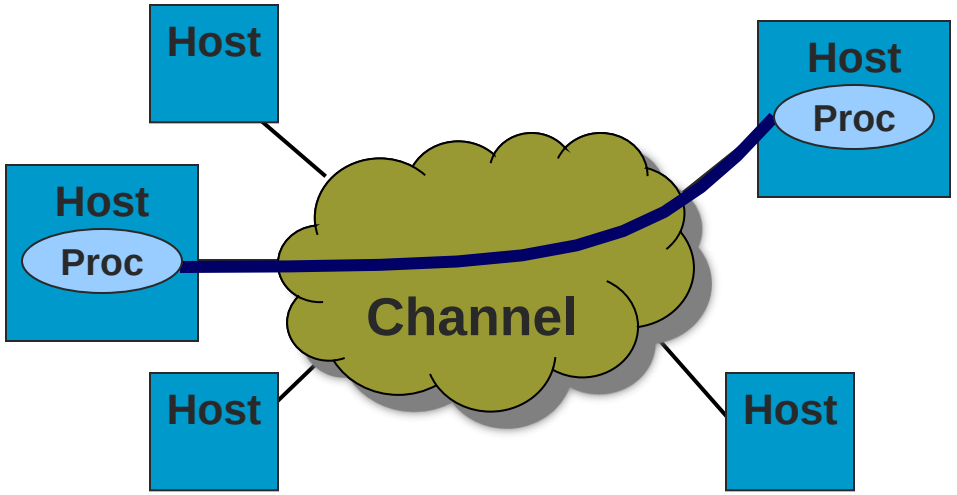

Looks like IPC!

Networked Communication Challenges

- Networked communication ≠ IPC
- Problems typically masked by communication channel abstractions
  - Bit errors (electrical interference)
  - Packet errors (congestion)
  - Link/node failures
  - Message delays
  - Out-of-order delivery
  - Eavesdropping
- Goal
  - Fill the gap between what applications expect and what the underlying technology provides

Network Architecture

- Networks are complex!
- Many “pieces”
  - Hosts
  - Routers
  - Links of various media
  - Applications
  - Protocols
  - Hardware, software
- Question
  - Is there any hope of organizing structure of network?

Abstraction through Layering

- Abstract system into layers:
  - Decompose the problem of building a network into manageable components
    - Each layer provides some functionality
  - Modular design provides flexibility
    - Modify layer independently
    - Allows alternative abstractions

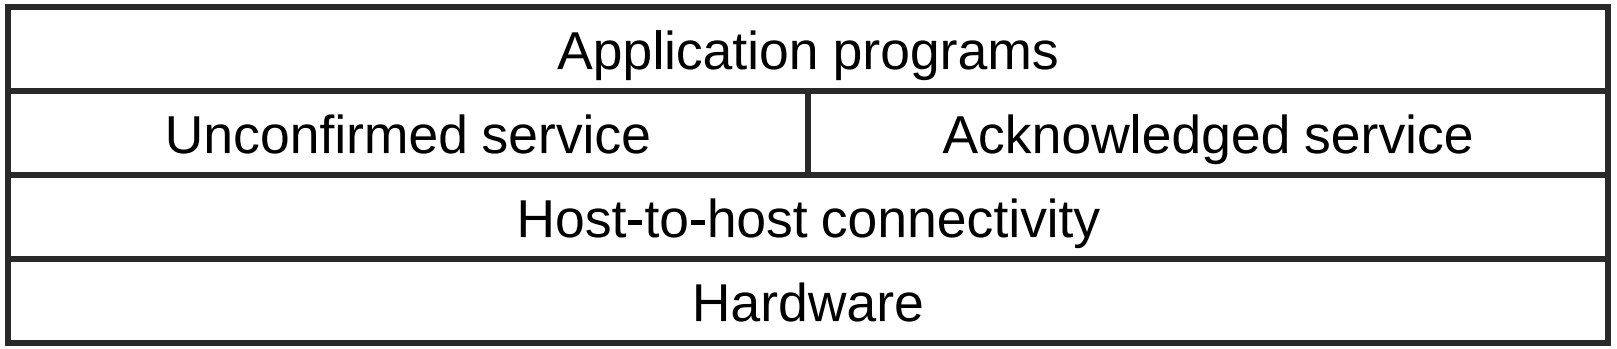

Layering Example: Air Travel

- Layers
  - Each layer implements a service
  - Via its own internal-layer actions
  - Relying on services provided by layer below

Why layering?

- Complexity
  - Explicit structure allows identification, relationship of complex system's pieces
- Modularity
  - Eases maintenance, updating of system
    - Change of implementation of layer's service transparent to rest of system
    - e.g., change in gate procedure doesn't affect rest of system

Protocol: Language of communication across hosts

- Defines structure of communication between two instances of a layer (on two hosts)
- Protocols are defined by
  - Specific msgs sent
  - Specific actions taken when msgs received, or other events
- Protocols define
  - Format
  - Order of msgs sent and received among network entities
  - Actions taken on msg transmission, receipt

What is a Protocol?

- Human protocols
  - "what's the time?"
  - "I have a question"
  - Introductions

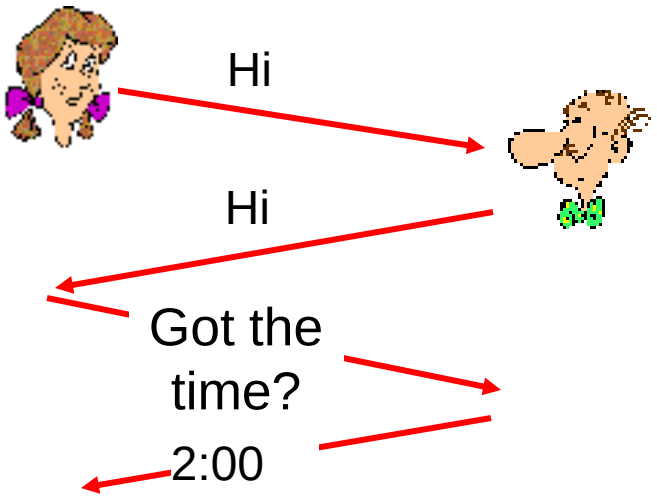

time

- Network protocols
  - Machines rather than humans
  - All internet communication is governed by protocols

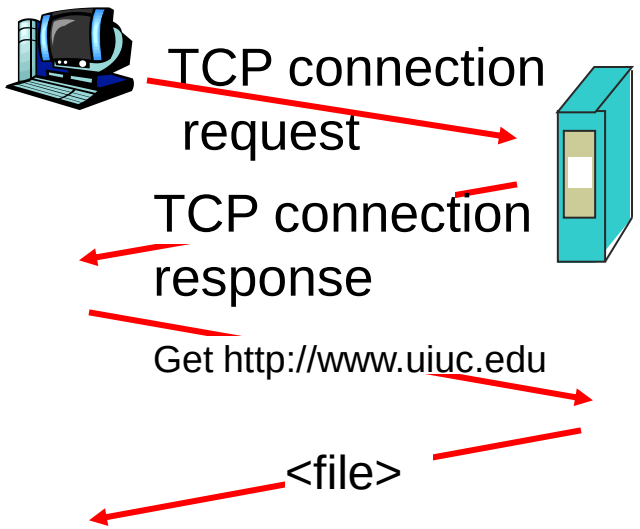

Layering Concepts

- Encapsulation
  - Higher layer protocols create messages and send them via the lower layer protocols
  - These messages are treated as data by the lower-level protocol
  - Higher-layer protocol adds its own control information in the form of headers or trailers
- Multiplexing and Demultiplexing
  - Use protocol keys in the header to determine correct upper-layer protocol

Internet Protocol Stack

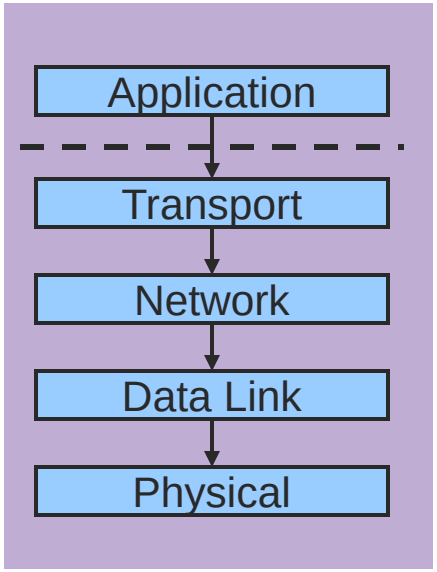

- Application: Application specific protocols
- Transport: Process-to-process channel
- Network: Host-to-host packet delivery
- Data Link: Framing of data bits
- Physical: Transmission of raw bits

Network Packet

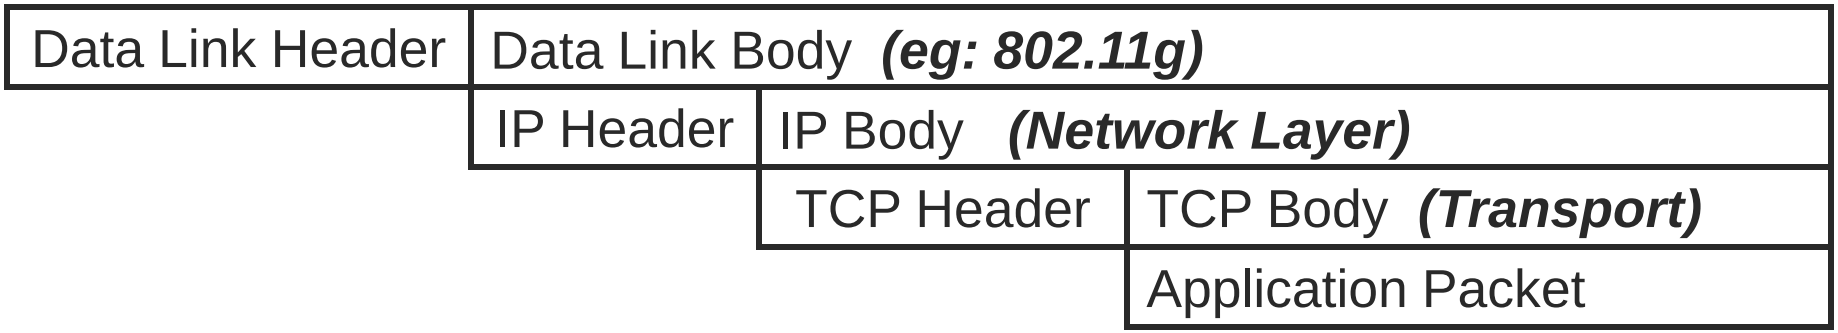

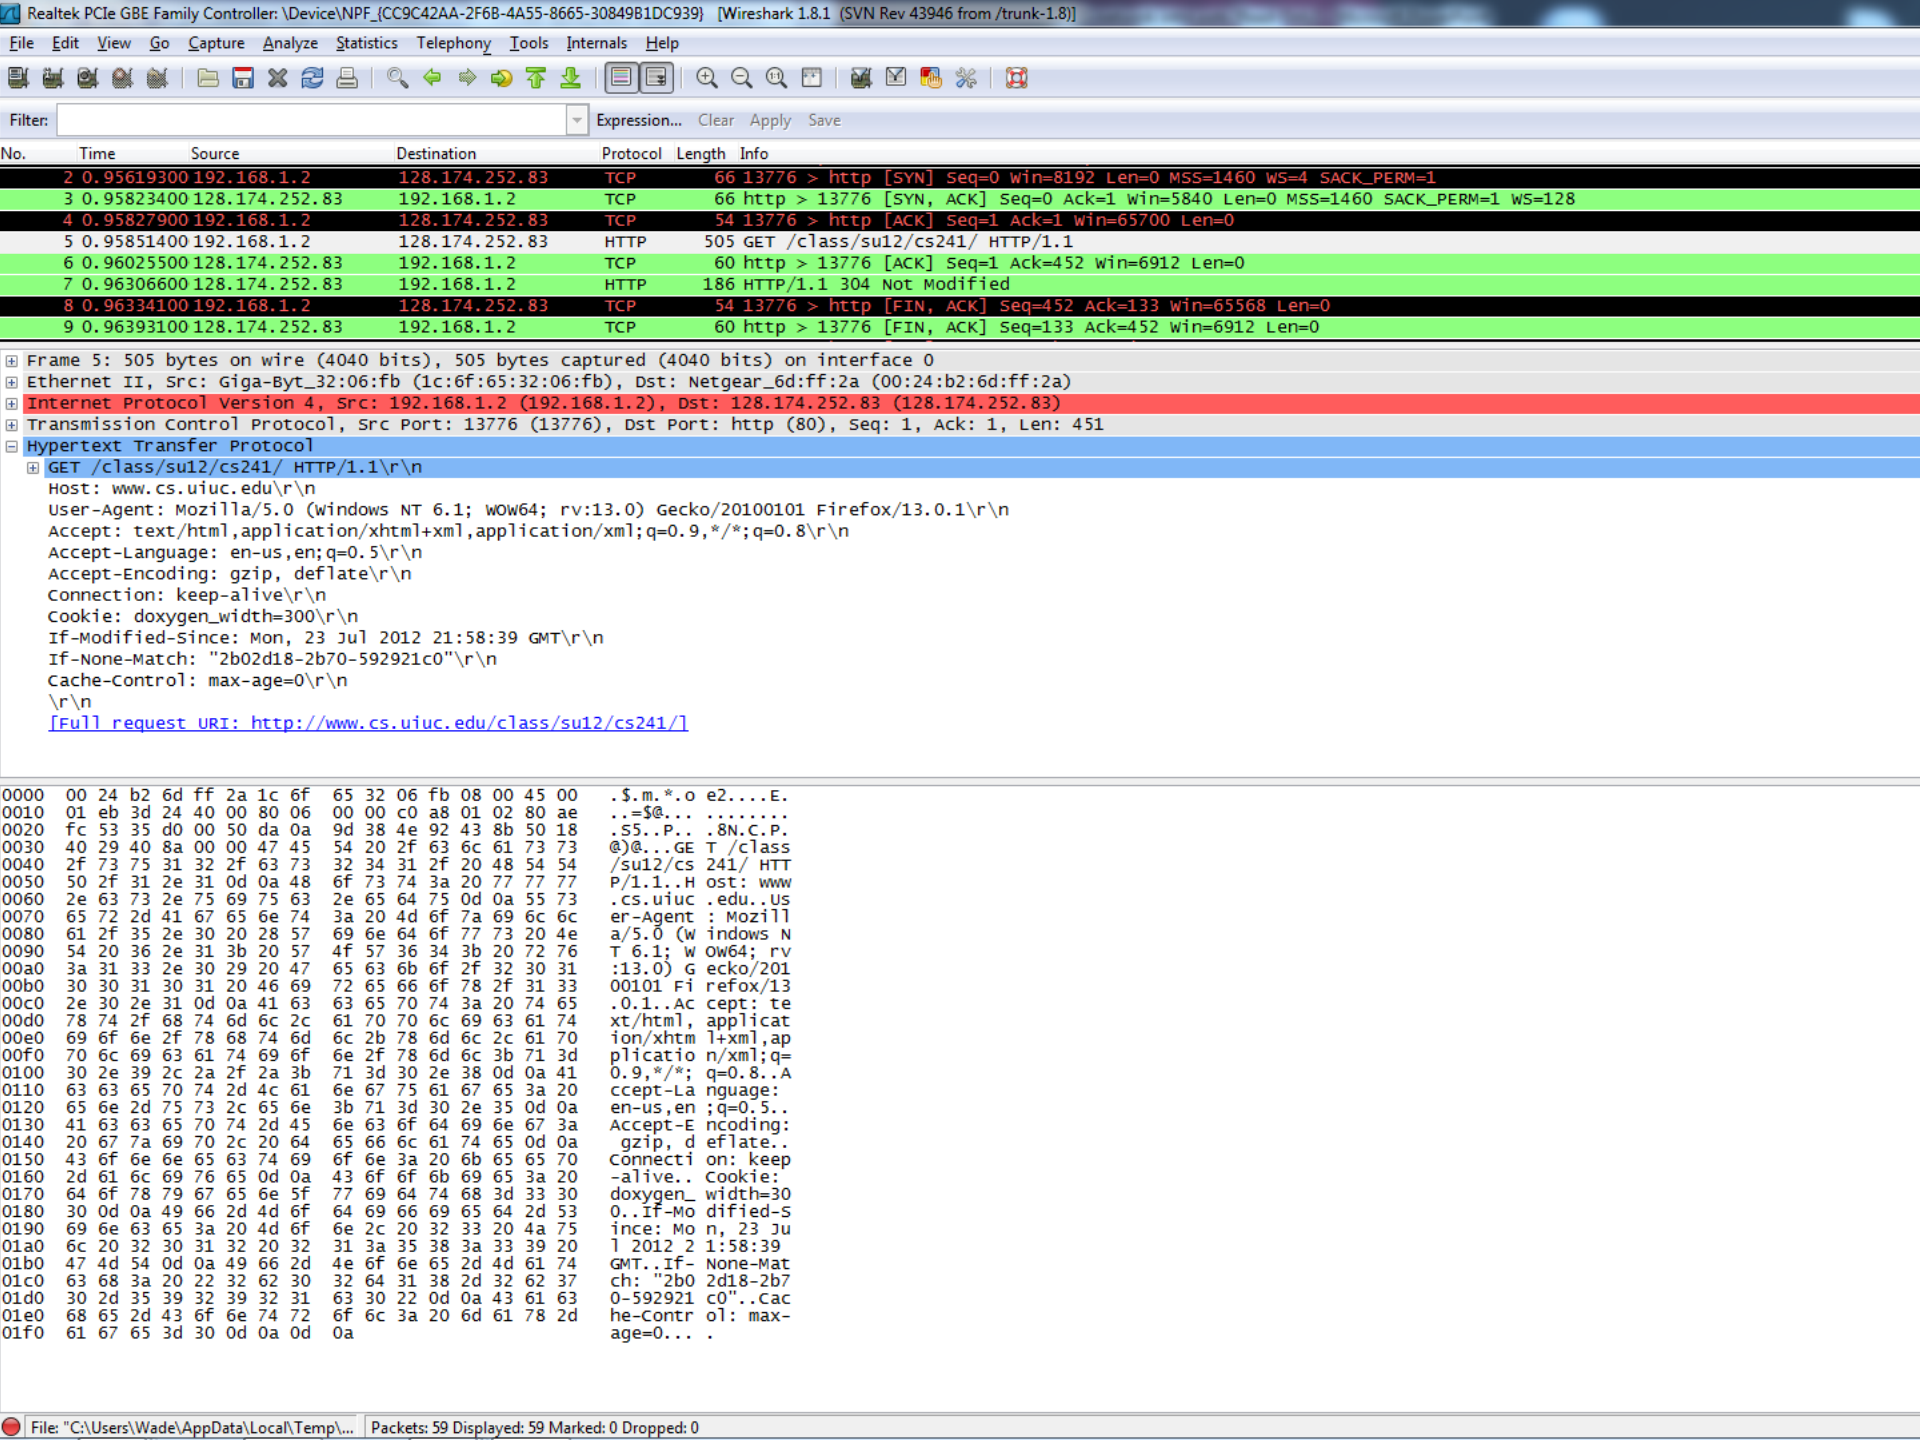
Realtek PCIe GBE Family Controller: \Device\NPF_{CC9C42AA-2F6B-4A55-8665-30849B1DC939} [Wireshark 1.8.1 (SVN Rev 43946 from /trunk-1.8)]
File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help
Filter: Expression... Clear Apply Save
No. Time Source Destination Protocol Length Info
2 0.95619300 192.168.1.2 128.174.252.83 TCP 66 13776 > http [SYN] Seq=0 Win=8192 Len=0 MSS=1460 WS=4 SACK_PERM=1
3 0.95823400 128.174.252.83 192.168.1.2 TCP 66 http > 13776 [SYN, ACK] Seq=0 Ack=1 Win=5840 Len=0 MSS=1460 SACK_PERM=1 WS=128
4 0.95827900 192.168.1.2 128.174.252.83 TCP 54 13776 > http [ACK] Seq=1 Ack=1 Win=65700 Len=0
5 0.95851400 192.168.1.2 128.174.252.83 HTTP 505 GET /class/su12/cs241/ HTTP/1.1
6 0.96025500 128.174.252.83 192.168.1.2 TCP 60 http > 13776 [ACK] Seq=1 Ack=452 Win=6912 Len=0
7 0.96306600 128.174.252.83 192.168.1.2 HTTP 186 HTTP/1.1 304 Not Modified
8 0.96334100 192.168.1.2 128.174.252.83 TCP 54 13776 > http [FIN, ACK] Seq=452 Ack=133 Win=65568 Len=0
9 0.96393100 128.174.252.83 192.168.1.2 TCP 60 http > 13776 [FIN, ACK] Seq=133 Ack=452 Win=6912 Len=0
Frame 5: 505 bytes on wire (4040 bits), 505 bytes captured (4040 bits) on interface 0
Ethernet II, Src: Giga-Byt_32:06:fb (1c:6f:65:32:06:fb), Dst: Netgear_6d:ff:2a (00:24:b2:6d:ff:2a)
Internet Protocol Version 4, Src: 192.168.1.2 (192.168.1.2), Dst: 128.174.252.83 (128.174.252.83)
Transmission Control Protocol, Src Port: 13776 (13776), Dst Port: http (80), Seq: 1, Ack: 1, Len: 451
Hypertext Transfer Protocol
GET /class/su12/cs241/ HTTP/1.1\r\n
Host: www.cs.uiuc.edu\r\n
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:13.0) Gecko/20100101 Firefox/13.0.1\r\n
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8\r\n
Accept-Language: en-us,en;q=0.5\r\n
Accept-Encoding: gzip, deflate\r\n
Connection: keep-alive\r\n
Cookie: doxygen_width=300\r\n
If-Modified-Since: Mon, 23 Jul 2012 21:58:39 GMT\r\n
If-None-Match: "2b02d18-2b70-592921c0"\r\n
Cache-Control: max-age=0\r\n
\r\n
[Full request URI: http://www.cs.uiuc.edu/class/su12/cs241/]
0000 00 24 b2 6d ff 2a 1c 6f 65 32 06 fb 08 00 45 00 .$.m.*.o e2....E.
0010 01 eb 3d 24 40 00 80 06 00 00 c0 a8 01 02 80 ae ..=$@... ....
0020 fc 53 35 d0 00 50 da 0a 9d 38 4e 92 43 8b 50 18 .S5..P.. .8N.C.P.
0030 40 29 40 8a 00 00 47 45 54 20 2f 63 6c 61 73 73 @)@...GE T /class
0040 2f 73 75 31 32 2f 63 73 32 34 31 2f 20 48 54 54 /su12/cs 241/ HTT
0050 50 2f 31 2e 31 0d 0a 48 6f 73 74 3a 20 77 77 77 P/1.1..H ost: www
0060 2e 63 73 2e 75 69 75 63 2e 65 64 75 0d 0a 55 73 .cs.uiuc .edu..Us
0070 65 72 2d 41 67 65 6e 74 3a 20 4d 6f 7a 69 6c 6c er-Agent : Mozill
0080 61 2f 35 2e 30 20 28 57 69 6e 64 6f 77 73 20 4e a/5.0 (W indows N
0090 54 20 36 2e 31 3b 20 57 4f 57 36 34 3b 20 72 76 T 6.1; W OW64; rv
00a0 3a 31 33 2e 30 29 20 47 65 63 6b 6f 2f 32 30 31 :13.0) G ecko/201
00b0 30 30 31 30 31 20 46 69 72 65 66 6f 78 2f 31 33 00101 Fi refox/13
00c0 2e 30 2e 31 0d 0a 41 63 63 65 70 74 3a 20 74 65 .0.1..Ac cept: te
00d0 78 74 2f 68 74 6d 6c 2c 61 70 70 6c 69 63 61 74 xt/html, applicat
00e0 69 6f 6e 2f 78 68 74 6d 6c 2b 78 6d 6c 2c 61 70 ion/xhtm l+xml,ap
00f0 70 6c 69 63 61 74 69 6f 6e 2f 78 6d 6c 3b 71 3d plicatio n/xml;q=
0100 30 2e 39 2c 2a 2f 2a 3b 71 3d 30 2e 38 0d 0a 41 0.9,*/*; q=0.8..A
0110 63 63 65 70 74 2d 4c 61 6e 67 75 61 67 65 3a 20 ccept-La nguage:
0120 65 6e 2d 75 73 2c 65 6e 3b 71 3d 30 2e 35 0d 0a en-us,en ;q=0.5..
0130 41 63 63 65 70 74 2d 45 6e 63 6f 64 69 6e 67 3a Accept-E ncoding:
0140 20 67 7a 69 70 2c 20 64 65 66 6c 61 74 65 0d 0a gzip, d eflate..
0150 43 6f 6e 6e 65 63 74 69 6f 6e 3a 20 6b 65 65 70 Connecti on: keep
0160 2d 61 6c 69 76 65 0d 0a 43 6f 6f 6b 69 65 3a 20 -alive.. Cookie:
0170 64 6f 78 79 67 65 6e 5f 77 69 64 74 68 3d 33 30 doxygen_ width=30
0180 30 0d 0a 49 66 2d 4d 6f 64 69 66 69 65 64 2d 53 0..If-Mo dified-S
0190 69 6e 63 65 3a 20 4d 6f 6e 2c 20 32 33 20 4a 75 ince: Mo n, 23 Ju
01a0 6c 20 32 30 31 32 20 32 31 3a 35 38 3a 33 39 20 l 2012 2 1:58:39
01b0 47 4d 54 0d 0a 49 66 2d 4e 6f 6e 65 2d 4d 61 74 GMT..If- None-Mat
01c0 63 68 3a 20 22 32 62 30 32 64 31 38 2d 32 62 37 ch: "2b0 2d18-2b7
01d0 30 2d 35 39 32 39 32 31 63 30 22 0d 0a 43 61 63 0-592921 c0"..Cac
01e0 68 65 2d 43 6f 6e 74 72 6f 6c 3a 20 6d 61 78 2d he-Contr ol: max-
01f0 61 67 65 3d 30 0d 0a 0d 0a age=0... .
File: "C:\Users\Wade\AppData\Local\Temp\... Packets: 59 Displayed: 59 Marked: 0 Dropped: 0

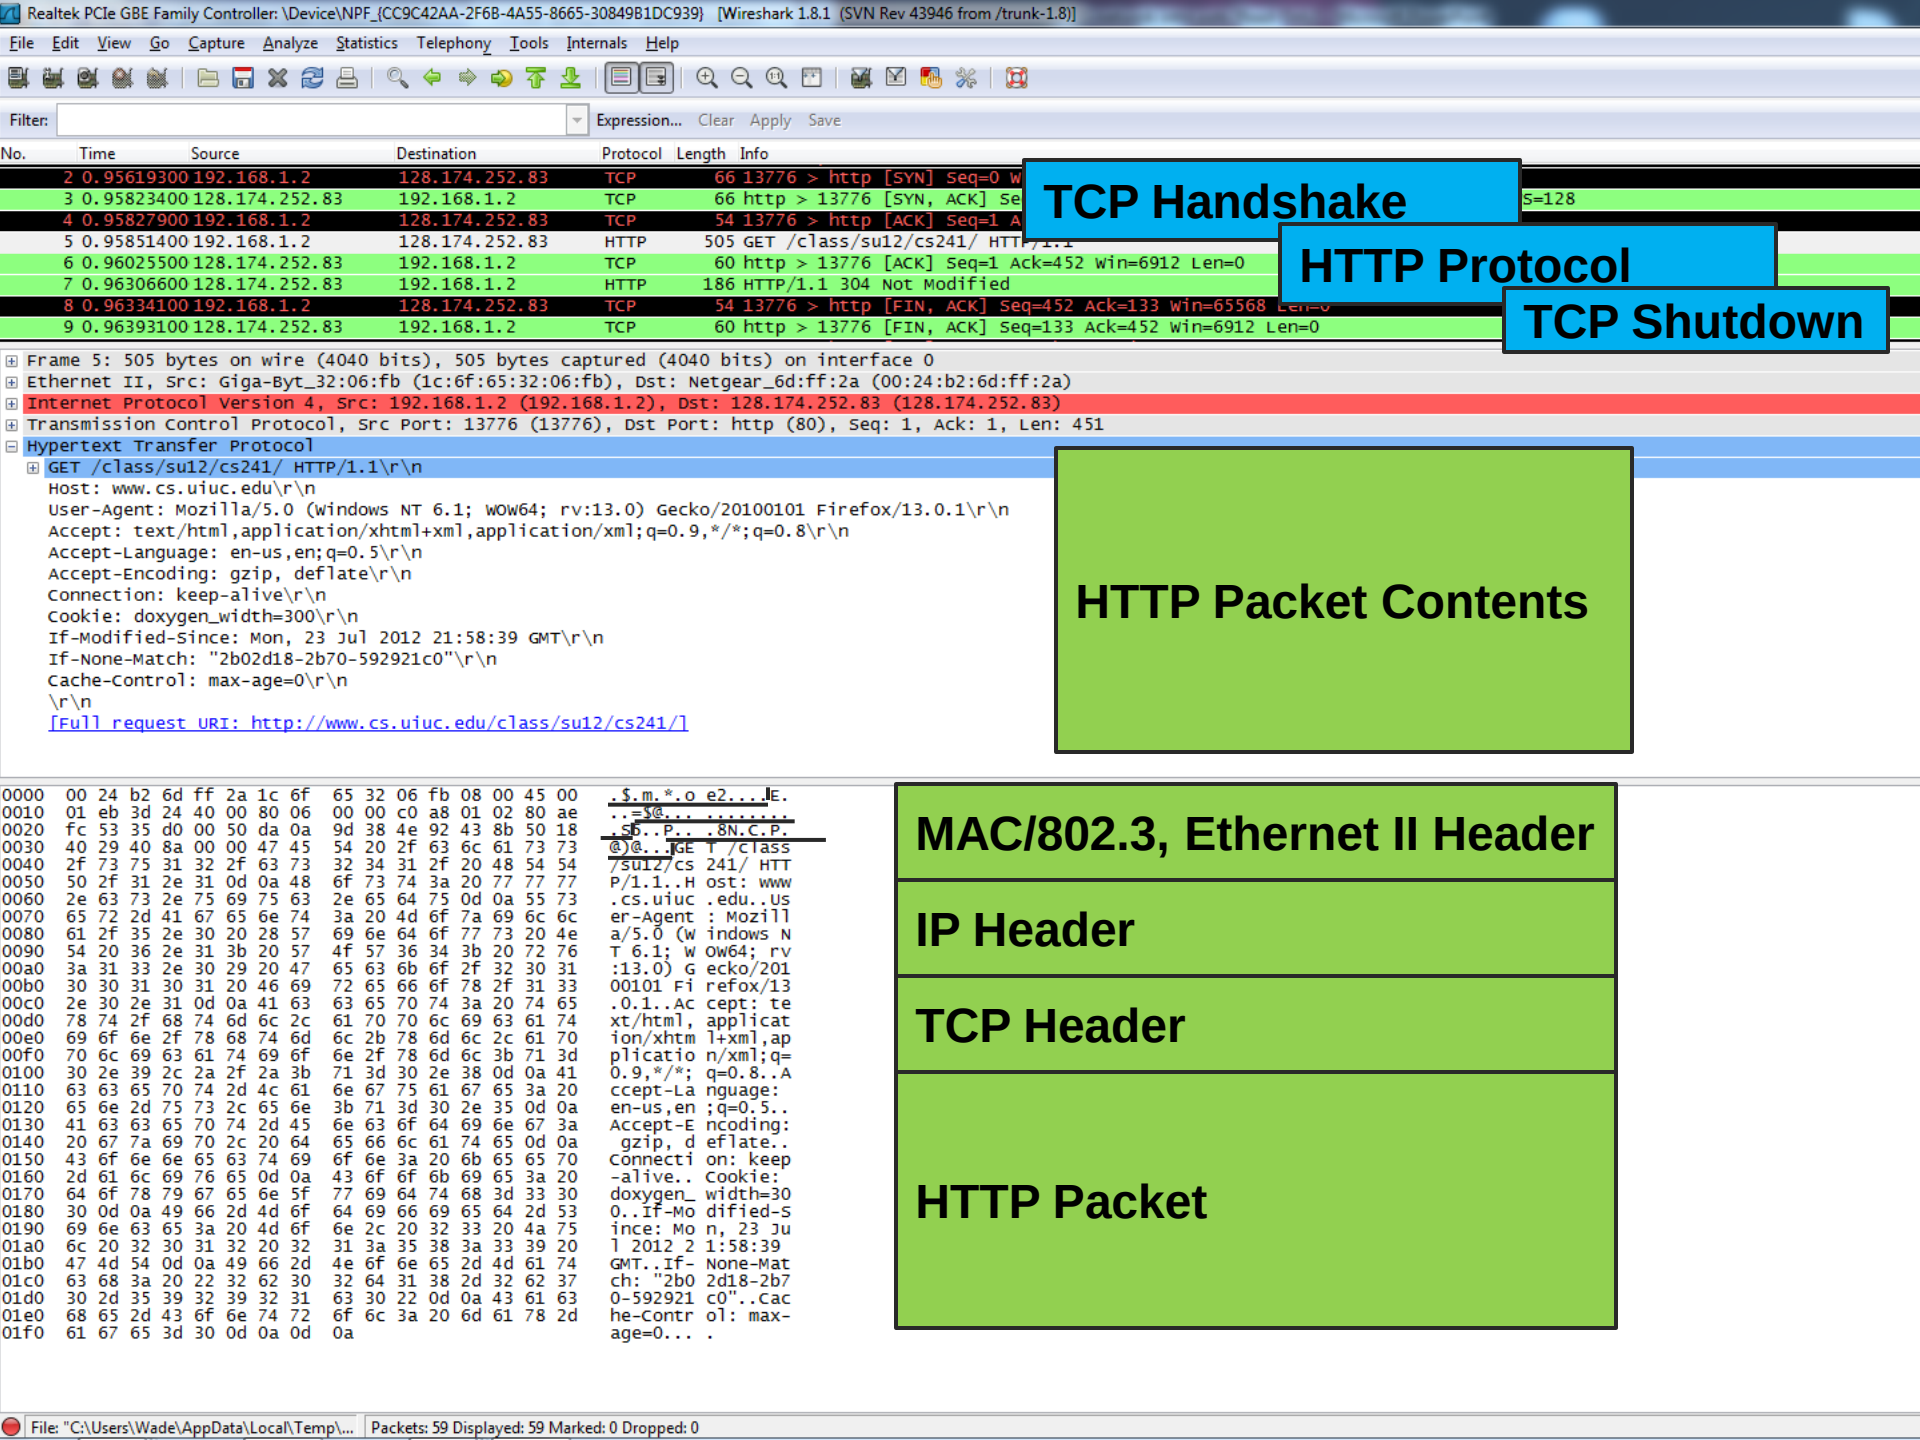
Realtek PCIe GBE Family Controller: \Device\NPF_{CC9C42AA-2F6B-4A55-8665-30849B1DC939} [Wireshark 1.8.1 (SVN Rev 43946 from /trunk-1.8)]
File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help
Filter: Expression... Clear Apply Save
No. Time Source Destination Protocol Length Info
2 0.95619300 192.168.1.2 128.174.252.83 TCP 66 13776 > http [SYN] Seq=0 W
3 0.95823400 128.174.252.83 192.168.1.2 TCP 66 http > 13776 [SYN, ACK] Se S=128
4 0.95827900 192.168.1.2 128.174.252.83 TCP 54 13776 > http [ACK] Seq=1 A
5 0.95851400 192.168.1.2 128.174.252.83 HTTP 505 GET /class/su12/cs241/ HTTP/1.1
6 0.96025500 128.174.252.83 192.168.1.2 TCP 60 http > 13776 [ACK] Seq=1 Ack=452 Win=6912 Len=0
7 0.96306600 128.174.252.83 192.168.1.2 HTTP 186 HTTP/1.1 304 Not Modified
8 0.96334100 192.168.1.2 128.174.252.83 TCP 54 13776 > http [FIN, ACK] Seq=452 Ack=133 Win=65568
9 0.96393100 128.174.252.83 192.168.1.2 TCP 60 http > 13776 [FIN, ACK] Seq=133 Ack=452 Win=6912 Len=0
TCP Handshake
HTTP Protocol
TCP Shutdown
Frame 5: 505 bytes on wire (4040 bits), 505 bytes captured (4040 bits) on interface 0
Ethernet II, Src: Giga-Byt_32:06:fb (1c:6f:65:32:06:fb), Dst: Netgear_6d:ff:2a (00:24:b2:6d:ff:2a)
Internet Protocol Version 4, Src: 192.168.1.2 (192.168.1.2), Dst: 128.174.252.83 (128.174.252.83)
Transmission Control Protocol, Src Port: 13776 (13776), Dst Port: http (80), Seq: 1, Ack: 1, Len: 451
Hypertext Transfer Protocol
GET /class/su12/cs241/ HTTP/1.1\r\n
Host: www.cs.uiuc.edu\r\n
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:13.0) Gecko/20100101 Firefox/13.0.1\r\n
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8\r\n
Accept-Language: en-us,en;q=0.5\r\n
Accept-Encoding: gzip, deflate\r\n
Connection: keep-alive\r\n
Cookie: doxygen_width=300\r\n
If-Modified-Since: Mon, 23 Jul 2012 21:58:39 GMT\r\n
If-None-Match: "2b02d18-2b70-592921c0"\r\n
Cache-Control: max-age=0\r\n
\r\n
[Full request URI: http://www.cs.uiuc.edu/class/su12/cs241/]
HTTP Packet Contents
0000 00 24 b2 6d ff 2a 1c 6f 65 32 06 fb 08 00 45 00 .$.m.*.o e2....E.
0010 01 eb 3d 24 40 00 80 06 00 00 c0 a8 01 02 80 ae ..=$@... ....
0020 fc 53 35 d0 00 50 da 0a 9d 38 4e 92 43 8b 50 18 .S5..P.. .8N.C.P.
0030 40 29 40 8a 00 00 47 45 54 20 2f 63 6c 61 73 73 @)@...GE T /class
0040 2f 73 75 31 32 2f 63 73 32 34 31 2f 20 48 54 54 /su12/cs 241/ HTT
0050 50 2f 31 2e 31 0d 0a 48 6f 73 74 3a 20 77 77 77 P/1.1..H ost: www
0060 2e 63 73 2e 75 69 75 63 2e 65 64 75 0d 0a 55 73 .cs.uiuc .edu..Us
0070 65 72 2d 41 67 65 6e 74 3a 20 4d 6f 7a 69 6c 6c er-Agent : Mozill
0080 61 2f 35 2e 30 20 28 57 69 6e 64 6f 77 73 20 4e a/5.0 (W indows N
0090 54 20 36 2e 31 3b 20 57 4f 57 36 34 3b 20 72 76 T 6.1; W OW64; rv
00a0 3a 31 33 2e 30 29 20 47 65 63 6b 6f 2f 32 30 31 :13.0) G ecko/201
00b0 30 30 31 30 31 20 46 69 72 65 66 6f 78 2f 31 33 00101 Fi refox/13
00c0 2e 30 2e 31 0d 0a 41 63 63 65 70 74 3a 20 74 65 .0.1..Ac cept: te
00d0 78 74 2f 68 74 6d 6c 2c 61 70 70 6c 69 63 61 74 xt/html, applicat
00e0 69 6f 6e 2f 78 68 74 6d 6c 2b 78 6d 6c 2c 61 70 ion/xhtm l+xml,ap
00f0 70 6c 69 63 61 74 69 6f 6e 2f 78 6d 6c 3b 71 3d plicatio n/xml;q=
0100 30 2e 39 2c 2a 2f 2a 3b 71 3d 30 2e 38 0d 0a 41 0.9,*/*; q=0.8..A
0110 63 63 65 70 74 2d 4c 61 6e 67 75 61 67 65 3a 20 ccept-La nguage:
0120 65 6e 2d 75 73 2c 65 6e 3b 71 3d 30 2e 35 0d 0a en-us,en ;q=0.5..
0130 41 63 63 65 70 74 2d 45 6e 63 6f 64 69 6e 67 3a Accept-E ncoding:
0140 20 67 7a 69 70 2c 20 64 65 66 6c 61 74 65 0d 0a gzip, d eflate..
0150 43 6f 6e 6e 65 63 74 69 6f 6e 3a 20 6b 65 65 70 Connecti on: keep
0160 2d 61 6c 69 76 65 0d 0a 43 6f 6f 6b 69 65 3a 20 -alive.. Cookie:
0170 64 6f 78 79 67 65 6e 5f 77 69 64 74 68 3d 33 30 doxygen_ width=30
0180 30 0d 0a 49 66 2d 4d 6f 64 69 66 69 65 64 2d 53 0..If-Mo dified-S
0190 69 6e 63 65 3a 20 4d 6f 6e 2c 20 32 33 20 4a 75 ince: Mo n, 23 Ju
01a0 6c 20 32 30 31 32 20 32 31 3a 35 38 3a 33 39 20 l 2012 2 1:58:39
01b0 47 4d 54 0d 0a 49 66 2d 4e 6f 6e 65 2d 4d 61 74 GMT..If- None-Mat
01c0 63 68 3a 20 22 32 62 30 32 64 31 38 2d 32 62 37 ch: "2b0 2d18-2b7
01d0 30 2d 35 39 32 39 32 31 63 30 22 0d 0a 43 61 63 0-592921 c0"..Cac
01e0 68 65 2d 43 6f 6e 74 72 6f 6c 3a 20 6d 61 78 2d he-Contr ol: max-
01f0 61 67 65 3d 30 0d 0a 0d 0a age=0... .
MAC/802.3, Ethernet II Header
IP Header
TCP Header
HTTP Packet
File: "C:\Users\Wade\AppData\Local\Temp\... Packets: 59 Displayed: 59 Marked: 0 Dropped: 0

Transport Layer

- Provide logical communication between application processes running on different hosts
- Transport protocols run in end systems
  - Send side
    - Break application messages into segments
    - Pass to network layer
  - Receive side
    - Reassemble segments into messages
    - Pass to application layer
- More than one transport protocol available to applications
  - Internet: TCP and UDP

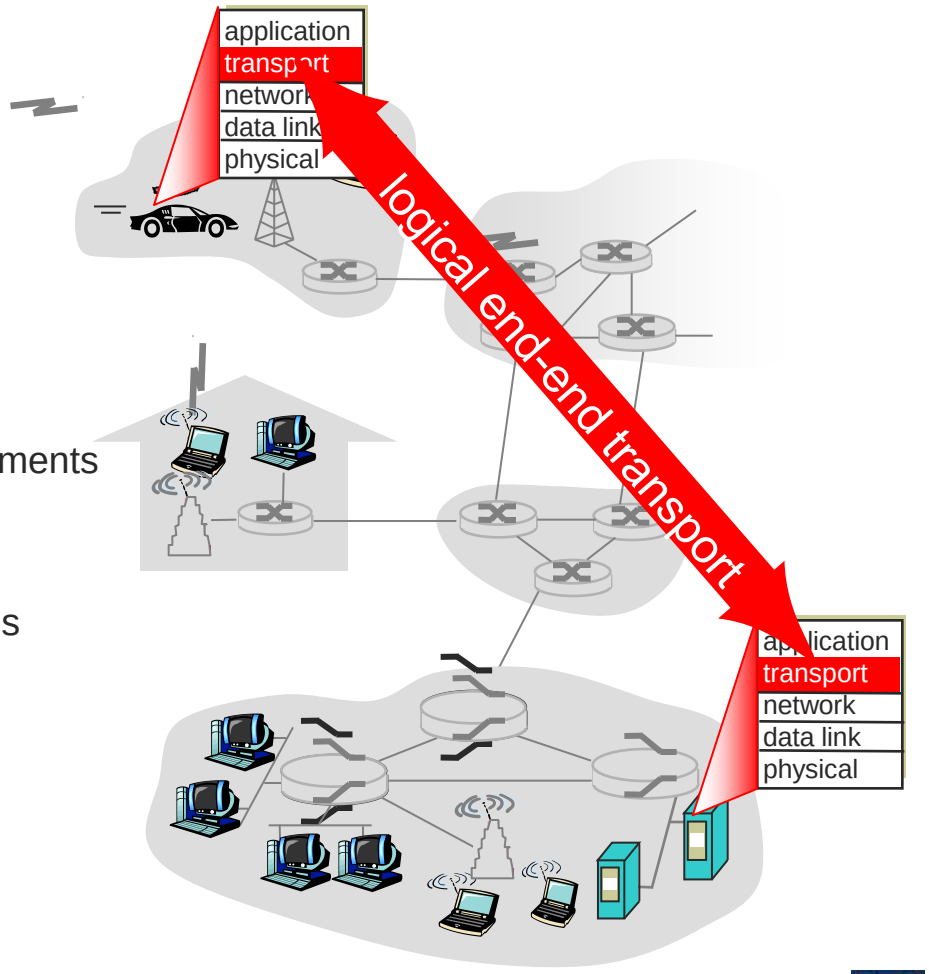

Transport vs. Network Layer

- Transport layer
  - Logical communication between processes
  - Relies on, enhances, network layer services
- Network layer
  - Logical communication between hosts

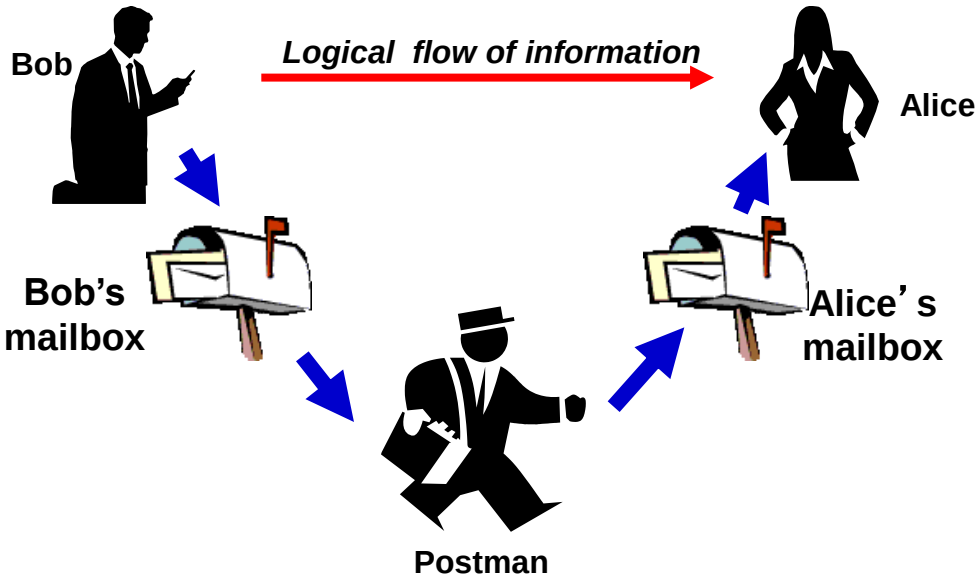

Internet Architecture – Hourglass Design

- Features
  - Hourglass shape – IP is the focal point

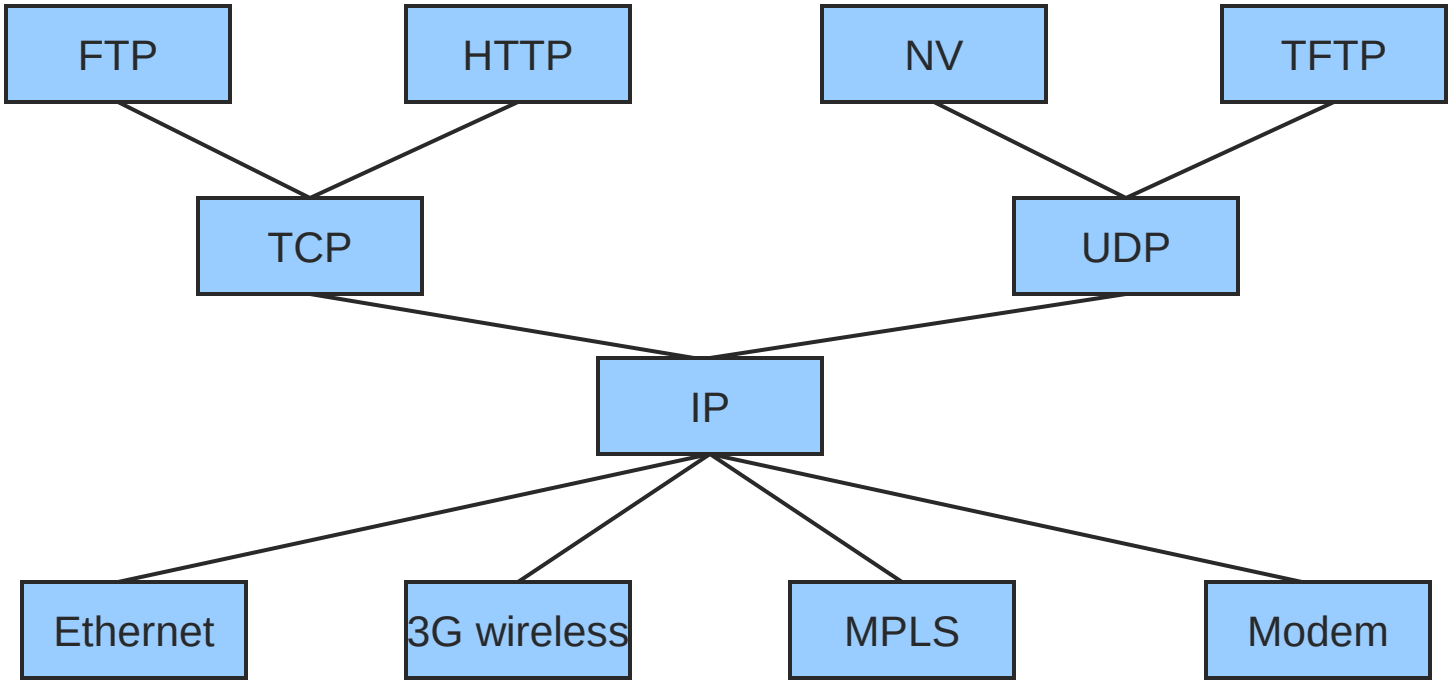

Network Applications

Creating a Network Application

- Write programs that
  - Run on (different) end systems
  - Communicate over network
    - e.g., web server software communicates with browser software
- No need to write software for network-core devices
  - Network-core devices do not run user applications

Client-server Architecture

- Server
  - Always-on host
  - Well-known IP address
- Clients
  - Communicate with server
  - May be intermittently connected
  - May have dynamic IP addresses
  - Do not communicate directly with each other

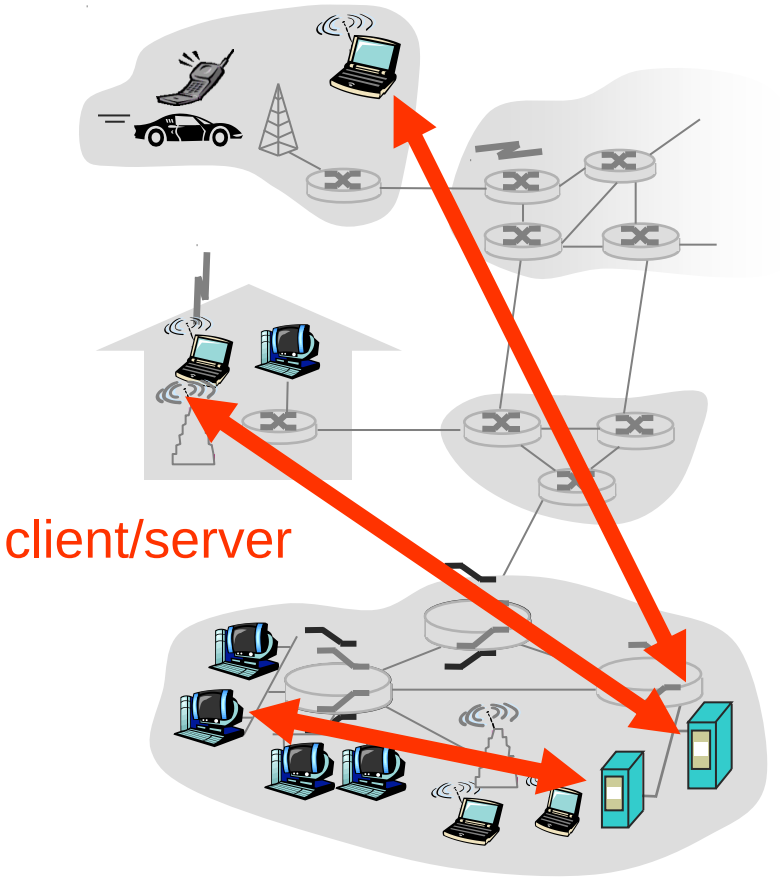

P2P Architecture

- No always-on server
- Arbitrary end systems directly communicate
- Peers are intermittently connected and change IP addresses

- Highly scalable but difficult to manage

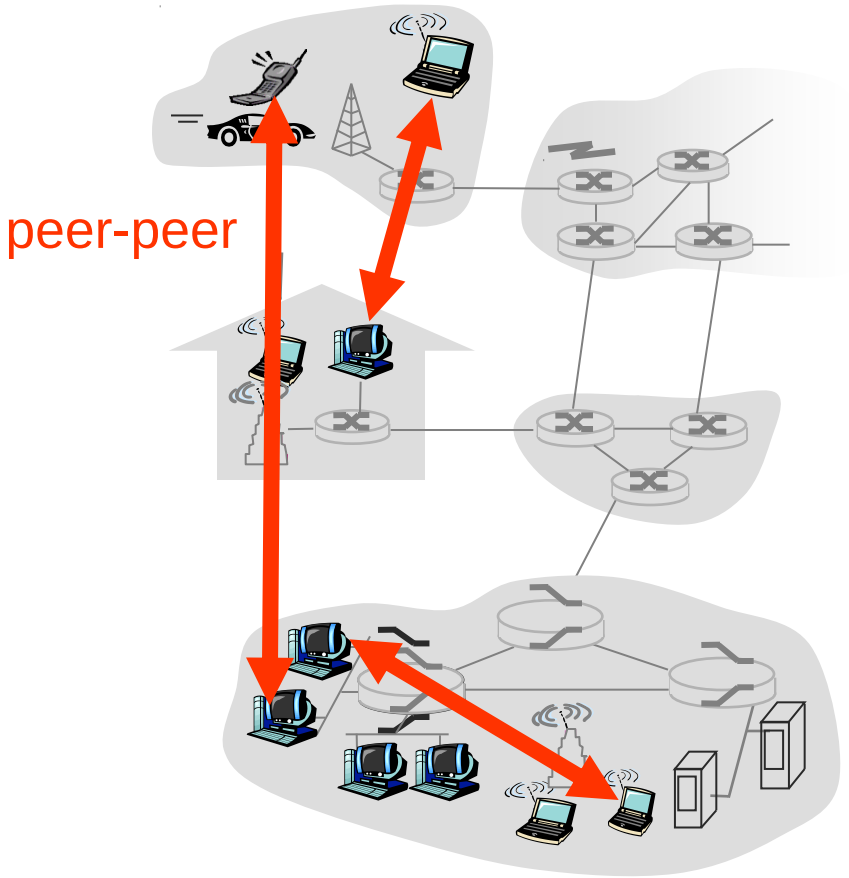

Hybrid Client-server and P2P

- Skype
  - Voice-over-IP P2P application
  - Centralized server: finding address of remote party
  - Client-client connection: direct (not through server)
- Instant messaging
  - Chatting between two users is P2P
  - Centralized service: client presence detection/location
  - User registers its IP address with central server when it comes online
  - User contacts central server to find IP addresses of buddies

Addressing Processes

- Receiving messages
  - Process must have identifier
  - Host device has unique 32-bit IP address
- Question
  - Does the IP address of host suffice for identifying the process?
  - Answer: No, many processes can be running on same host
- Process Identifier
  - Include both IP address and port number associated with process on host
- Example port numbers
  - HTTP server: 80
  - Mail server: 25

Sockets

- Process sends/receives messages to/from its socket
  - Analogous to a door
  - Sending process shoves messages out the door
  - Transport infrastructure on other side of door brings message to socket at receiving process

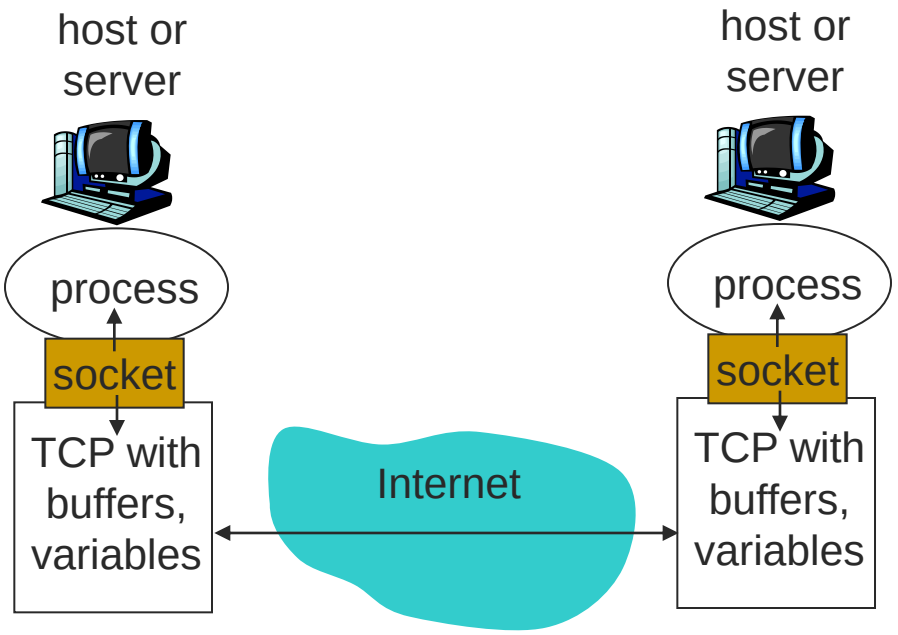

Sockets

- API
  - Choice of transport protocol
  - Ability to set a few parameters

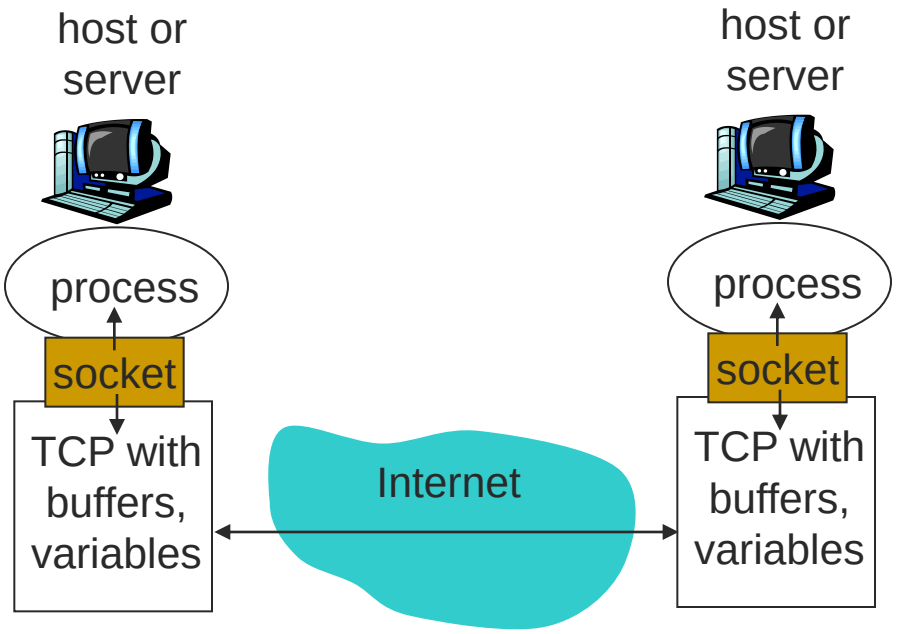

Transport Services

- Data loss
  - Some applications (e.g., audio) can tolerate some loss
  - Other apps (e.g., file transfer, telnet) require 100% reliability
- Timing
  - Some applications (e.g., IP phones, interactive games) require low delay to be "effective"
- Throughput
  - Some applications (e.g., multimedia) have a minimum throughput to be "effective"
  - other applications ("elastic apps") make use of whatever throughput they get
- Security
  - Encryption, data integrity, …

Internet Transport Protocols

TCP

- Connection-oriented
  - setup required between client and server
- Reliable transport
- Flow control
  - Won't overwhelm receiver
- Congestion control
  - Won't overwhelm network
- Does not provide
  - Timing, throughput guarantees, security

UDP

- Unreliable data transfer
- Does not provide
  - Connection setup, reliability, flow control, congestion control, timing, throughput guarantee, or security
- Question
  - Why bother? Why is there a UDP?