

Lecture 18: Networks

CS213 – Intro to Computer Systems
Branden Ghen a – Winter 2026

Some slides borrowed from:
Bryant and O'Hallaron (CMU) and Eleftherios Kosmas (UCSD)

Administrivia

- SETI Lab
 - Due tonight, but it's okay to use slip days or late days
 - Feel free to ask questions on Piazza
- Office hours
 - Normal today and Friday
 - No normal office hours next week
 - We'll try to put something together on Tuesday/Wednesday

Exam reminders

- Midterm Exam 3 details
 - Thursday, March 19 at noon in lecture room (Tech LR3)
 - 75-minute exam, same style as last time
 - Covers lectures 10-18 (Memory Security to Networks)
 - Particular emphasis:
 - Security
 - Caches and Cache Performance
 - Virtual Memory
 - Concurrency
- **One 8.5" x 11" sheets of paper** with notes on front and back
- Prior exams posted on Canvas

Today's Goals

- How do computers communicate?
- What choices does the Internet make about how to communicate?
- Explore Internet communication mechanisms: TCP and HTTP
- See CS340 for more details!
 - This is an 80-minute speedrun of a full quarter-long course

Outline


- **Computer Networks**
 - **Topology**
 - Inter-network communication
- The Internet
 - Protocol choices
- Sockets
 - System calls for communicating with other computers
- The Web

Computer networks

- A network is a connection of two or more computers
 - Definition includes hardware and software components
- Networks might serve a variety of regions
 - Local Area Network (LAN)
 - Connects computers in some small area
 - Home, building, campus
 - Wide Area Network (WAN)
 - Connects computers over distant areas
 - City, country, world
- Could be wired or wireless

OSI model of communication layers

- Transport
 - How to form connections between computers
 - Example: TCP and UDP
- Network
 - How to send packets between networks
 - Example: IP
- Data Link
 - How to send frames of data
 - Example: Ethernet, WiFi
- Physical
 - How to send individual bits
 - Example: Ethernet, WiFi



How do we connect computers?

- First question of a network: how are the computers actually connected?
- There are various choices here with different implications


How to connect: point-to-point networks

- How do we connect computers in a network?
 - This is a question of “network topology”
 - Simple option: just connect them directly
- Problem: what if I find a third computer?



How to connect: bus networks

- Connect everything to one wire in parallel
 - Actually a “multidrop bus”
 - Scales pretty well to many computers
- Problem: which computer gets to transmit when?
 - Simultaneous transmissions conflict




How to connect: ring networks

- Connect everything with point-to-point connections
 - Connect the last computer back to the first computer
- Problem: what if a computer stops sending?




How to connect: star networks

- Connect to a hub with point-to-point connections
 - Hub connects all computers
 - Hub is a simple computer with one job: transfer communications between computers
 - Hopefully more reliable than any of the computers




Common home networks

- Example of a star topology network
 - Your home router is the hub connecting multiple computers
 - Could be wired (Ethernet) or wireless (WiFi)




Building a network: local area network

- Smallest division of a network
- **Hosts** are connected to **router** (wired or wireless)
- Operation (same for Ethernet or WiFi)
 - Each “adapter” has a unique 48-bit address (MAC Address)
 - e.g. 00:16:ea:e3:54:e6
 - Hosts send bits to router in chunks called **frames**
 - Routers are aware of other connected hosts (by ID) and forward message
- Story: how to connect a kindle to airport WiFi




Building a network: internets

- Multiple local area networks (LANs) can be physically connected
- The connected networks are called an *internet* (short for internetworking)
 - Networks might be running totally separate protocols: Ethernet, WiFi, ...
 - Modems translate between network types: e.g., Ethernet <-> Cable



Structure of an internet



- Ad hoc interconnection of networks
 - No particular topology
 - Possibly vastly different link capabilities
- Send packets from source to destination by hopping through networks

Outline

- **Computer Networks**

- Topology
- **Inter-network communication**

- The Internet

- Protocol choices

- Sockets

- System calls for communicating with other computers

- The Web

Managing inter-network communication with a protocol


- Now we have the ability to form and connect networks
 - But how do we transmit data between them?
 - Especially given that they differ in capabilities and design
- Solution: *protocol* software running on each host AND router
 - The protocol is a set of rules that governs how hosts and routers should cooperate when they transfer data from network to network
 - Smooths out the differences between the different networks

The role of an inter-network protocol


- Provides a **naming scheme**
 - We cannot rely on any particular lower-layer address
 - Defines a uniform format for *host addresses*
 - Each host (and router) is assigned at least one of these internet addresses that uniquely identifies it
- Provides a **delivery mechanism**
 - Defines a standard transfer unit: packet
 - A packet consists of a header and a payload
 - Header: Metadata such as packet size, source and destination addresses
 - Payload: data bits sent from the source host

Protocols are “layered”

- Headers for each layer of communication wrap data
 - Data is wrapped with header for the network to make a packet
 - i.e., bytes are added to the start/end of it
 - Packet is wrapped with header for the link to make a frame




Transmitting data between networks



Other inter-network issues

- We are glossing over a number of important questions:
 - What if different networks have different maximum frame sizes? (segmentation)
 - How do routers know *where* to forward frames?
 - How are routers informed when the network topology changes?
 - What if packets get lost?
- These (and other) questions are addressed by the area of systems known as **computer networking**
 - CS340 – Intro to Computer Networks

Break + xkcd








Outline

- Computer Networks
 - Topology
 - Inter-network communication
- **The Internet**
 - **Protocol choices**
- Sockets
 - System calls for communicating with other computers
- The Web




The ARPANET in December 1969









TeleGeography Submarine Cable Map

The [Submarine Cable Map](#) is a free resource from TeleGeography. Data contained in this map is drawn from the [Global Bandwidth Research Service](#) and is updated on a regular basis.

To learn more about TeleGeography or this map please [click here](#).



Sponsored in part by Huawei Marine

Feedback [t](#) [f](#) [github](#)

[Q Search](#)

[Submarine Cable List](#)

Seattle, WA, United States

[Email link](#)


[Cables](#)

[Arctic Fibre](#)

The global Internet

- Most famous example of an internet (uppercase to distinguish)
- Based on the TCP/IP protocol family
 - **IP** (Internet Protocol)
 - Provides a *naming scheme* and unreliable *delivery of packets* from **host-to-host**
 - **UDP** (Unreliable Datagram Protocol)
 - Uses IP to provide *unreliable data delivery* from **process-to-process**
 - **TCP** (Transmission Control Protocol)
 - Uses IP to provide *reliable data delivery* from **process-to-process**
- Accessed via a mix of Unix file I/O and the **sockets** interface

Hardware and software organization of an Internet application



A programmer's view of the internet

1. Hosts are mapped to a set of 32-bit **IP addresses**
 - 129.105.5.212
2. The set of IP addresses is mapped to a set of identifiers called Internet **domain names**
 - 129.105.5.212 is mapped to moore.wot.eecs.northwestern.edu
3. A process on one Internet host can communicate with a process on another Internet host over a **connection**

A programmer's view of the internet

1. Hosts are mapped to a set of 32-bit **IP addresses**

- 129.105.5.212

2. The set of IP addresses is mapped to a set of identifiers called Internet **domain names**

- 129.105.5.212 is mapped to moore.wot.eecs.northwestern.edu

3. A process on one Internet host can communicate with a process on another Internet host over a **connection**

1. IP addresses

- 32-bit IP addresses are stored in an **IP address struct**
 - IP addresses are always stored in memory in *network byte order* (big-endian)
 - Remember: most computers use little-endian 😓
 - True in general for any integer transferred in a packet header from one machine to another
 - E.g., the port number used to identify an Internet connection

```
/* Internet address structure */
struct in_addr {
    uint32_t    s_addr; /* network byte order (big-endian) */
};
```

- By convention, each byte in a 32-bit IP address is represented by its decimal value and separated by a period
 - IP address: 0x816905D4 = 129.105.7.212

How many hosts can there be?


- IP addresses are 32-bits

How many hosts can there be?

- IP addresses are 32-bits
 - $2^{32} = 4.3$ billion hosts
- **However**, some addresses are reserved for special purposes
 - Private networks (10.x.x.x, 192.168.x.x, ...)
 - Within a computer (127.x.x.x)
 - Various testing or reserved purposes
- 588 million are reserved
 - So 3.7 billion **publicly accessible** hosts
- Typically, a home router has **one** publicly accessible IP address
 - All devices within the home are on a private network
 - Phones on cellular data are also within a private network

Aside: IPv4 and IPv6

- The original Internet Protocol, with its 32-bit addresses, is known as *Internet Protocol Version 4* (IPv4)
- 1996: Internet Engineering Task Force (IETF) introduced *Internet Protocol Version 6* (IPv6) with 128-bit addresses
 - Intended as the successor to IPv4
- Majority of Internet traffic still carried by IPv4 (we'll focus on it today)



IPv6 traffic at Google

2026: 46%

2025: 48%

2024: 44%

2023: 41%

Northwestern IPv4 ranges

- Format: `ip.addr/NUM`
- **NUM** is the number of meaningful bits, from the most-significant bit
 - 129.105.0.0/16 means the most-significant 16 bits are meaningful
i.e., 129.105.x.x
- Remaining bits can be used for unique addresses
 - $2^{(32-\text{NUM})}$ machines

Public IPv4 Address Ranges

Network	Description
129.105.0.0/16	Northwestern Administrative Mainly - Evanston
165.124.0.0/16	Northwestern Administrative Evanston and Chicago
165.124.188.0/22	Northwestern-Qatar1, all others (except NU-Qatar2) on main NU campus
165.124.236.0/22	Northwestern-Qatar2, all others (except NU-Qatar1) on main NU campus
192.5.143.0/24	Peering point to point links, various other usage
192.26.86.0/24	ABF Rubloff 4th floor
192.26.87.0/24	Starlight mgmt network
192.31.155.0/24	Northwestern Affiliate
192.31.253.0/24	Utility
199.74.64.0/18	Northwestern dorms
199.249.165.0/24	Northwestern Affiliate
199.249.166.0/24	Northwestern Affiliate
199.249.167.0/24	Northwestern Affiliate
199.249.168.0/24	Northwestern Affiliate

A programmer's view of the internet

1. Hosts are mapped to a set of 32-bit **IP addresses**


- 129.105.5.212

2. The set of IP addresses is mapped to a set of identifiers called Internet **domain names**

- 129.105.5.212 is mapped to moore.wot.eecs.northwestern.edu

3. A process on one Internet host can communicate with a process on another Internet host over a **connection**

2. Internet domain names



Top-level domain names

Second-level domain names

Third-level domain names
and onwards...

Note: Northwestern owns 129.105.0.0/16 and 165.124.0.0/16

Some top-level domain names (TLDs)

.space	.store	.stream	.studio
.study	.style	.supplies	.supply
.support	.surf	.surgery	.sydney
.systems	.taipei	.tattoo	.tax
.taxi	.team	.tech	.technology
.tennis	.theater	.theatre	.tienda
.tips	.tires	.tirol	.today
.tokyo	.tools	.top	.tours
.town	.toys	.trade	.trading
.training	.tube	.university	.uno
.vacations	.vegas	.ventures	.versicherung
.vet	.viajes	.video	.villas
.vin	.vip	.vision	.vlaanderen
.vodka	.vote	.voting	.voto
.voyage	.wales	.wang	.watch
.webcam	.website	.wed	.wedding
.whoswho	.wien	.wiki	.win
.wine	.work	.works	.world
.wtf	.在线	.移动	.онлайн
.сайт	.сайт	.opr	.中文网
.संगठन	.机构	.みんな	.游戏
.企业	.xyz	.yoga	.yokohama
.zone			

Domain Naming System (DNS)

- The Internet maintains a mapping between IP addresses and domain names in a huge worldwide distributed database called **DNS**
- Conceptually, programmers can view the DNS database as a collection of millions of **host entries**
 - Each host entry defines the mapping between a set of domain names and IP addresses
- A special name: **localhost**
 - Refers back to the computer being used (IP address 127.0.0.1)

DNS lookups

- The entire DNS mapping is rather large
 - Billions of entries at 10s of bytes in size = 10s-100s of GB
- DNS servers maintain the lists and requests are made to those servers to translate into an IP address
 - DNS entry can then be cached on the local client!!
- DNS servers have well-known IP addresses
 - Google: 8.8.8.8 and 8.8.4.4
 - Cloudflare: 1.1.1.1 and 1.0.0.1
 - Comcast: 75.75.75.75

A programmer's view of the internet

1. Hosts are mapped to a set of 32-bit **IP addresses**

- 129.105.7.30

2. The set of IP addresses is mapped to a set of identifiers called Internet **domain names**

- 129.105.7.30 is mapped to moore.wot.eecs.northwestern.edu

3. A process on one Internet host can communicate with a process on another Internet host over a **connection**

3. Internet connections


- Clients and servers communicate by sending streams of bytes over **TCP connections**. Each connection is:
 - Point-to-point: connects a pair of processes.
 - Full-duplex: data can flow in both directions at the same time,
 - Reliable: stream of bytes sent by the source is eventually received by the destination in the same order it was sent.
- There are other communication mechanisms, like UDP
 - But we'll focus on TCP today

Connections are identified by IP address and Port


- Connection endpoints are called “sockets”
 - Socket address is an **IPAddress:port** pair
- A **port** is a 16-bit integer that identifies a process:
 - Ephemeral port
 - Assigned automatically by client kernel when client makes a connection request
 - Usually 1000 and up
 - Well-known port
 - Associated with some service provided by a server

Anatomy of a connection

- A connection is uniquely identified by socket addresses of endpoints
 - Socket address is an IP_address:Port pair
 - Connection: (client_address:client_port, server_address:server_port)



Ports are used to identify services to the kernel



Well-known service ports and names

- Popular services have permanently assigned *well-known ports and corresponding well-known service names*:
 - echo servers: echo 7
 - ftp servers: ftp 21
 - ssh servers: ssh 22
 - email servers: smtp 25
 - Web servers: http 80
- Mappings between well-known ports and services are listed in `/etc/services` on Linux

Break + Thinking

- What are the steps for viewing a website?
 1. You enter a domain name for the website
 2. ???

Break + Thinking

- What are the steps for viewing a website?
 1. You enter a domain name for the website
 2. Computer looks up domain name to get IP Address
 3. Computer sends request to IP_address:80
 4. Computer gets back data, which it renders into a website

ALL the layers

- A 'famous' interview question
 - "What happens when you type google.com into your browser's address bar and press enter?"
 - <https://github.com/alex/what-happens-when> (11 pages!)
 - Keyboard events
 - Parsing URL
 - DNS lookup
 - Opening socket
 - HTTP protocol
 - HTML parsing
 - GPU rendering

Outline

- Computer Networks
 - Topology
 - Inter-network communication
- The Internet
 - Protocol choices
- **Sockets**
 - **System calls for communicating with other computers**
- The Web

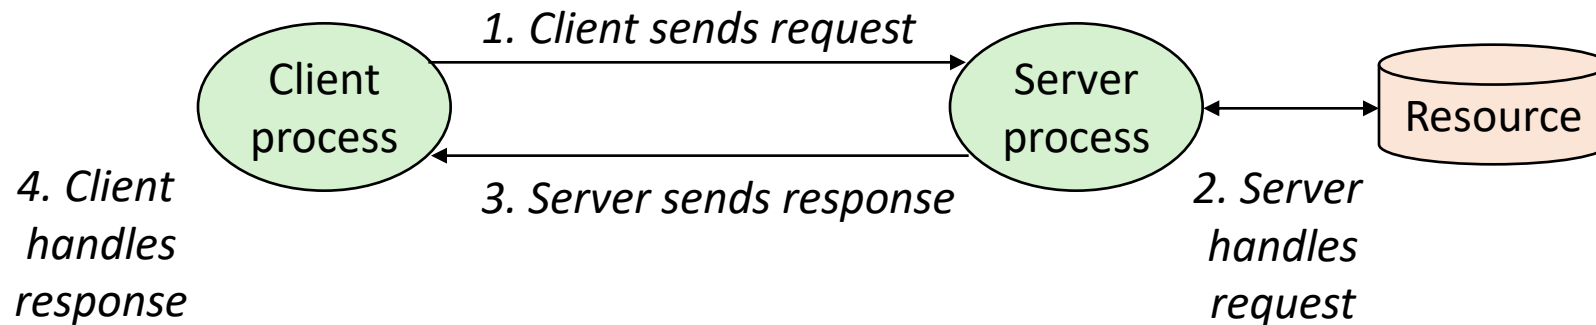
What is a socket?

- To an application, a socket is a file descriptor that lets the application read/write from/to the network
 - Connects a process on one host to a process on another
- Clients and servers communicate with each other by reading from and writing to socket descriptors
- The main distinction between regular file I/O and socket I/O is how the application “opens” the socket descriptors



Client-Server transaction design pattern

- Most network applications are based on the client-server model:
 - A **server** process and one or more **client** processes
 - Server manages some **resource**
 - Server provides **service** by manipulating resource for clients
 - Server activated by request from client



*Note: clients and servers are processes running on hosts
(can be the same or different hosts)*

Socket communication flow

Server


Client

1. Initialize
socket()

2. Create
connection

3. Transfer
data

4. Close



Socket communication flow

Server


Client

1. Initialize
socket()

2. Create
connection

3. Transfer
data

4. Close



Initializing a socket

- Both sides create a socket
 - Identifies the type of communication: AF_INET for internet communication
 - Alternatives are AF_INET6, AF_BLUETOOTH, etc.
- Server needs to do more work here
 - Pick an IP address and Port to connect to
 - Start listening for incoming requests

Socket communication flow

Server

Client

1. Initialize
socket()

socket()

bind()

listen()

socket()

2. Create
connection

accept()

connect()

3. Transfer
data

read()

write()

write()

read()

4. Close

close()

close()

Creating a connection

- Client “connects” to a server somewhere
 - Provides a server IP address and Port
 - OS will pick a random outgoing Port on the client
- Server waits until a connection arrives
 - For performance, modern servers often create a worker thread to handle each incoming connection
- Both sides get a “file descriptor” for the connection
 - Looks very similar to file I/O from a system call perspective

Socket communication flow

Server


Client

1. Initialize
socket()

2. Create
connection

3. Transfer
data


4. Close



Socket-specific send/receive calls

- The same as read/write for files except with additional flags
 - Blocks until data is completed
- `ssize_t recv(int sockfd, void *buf, size_t len, int flags);`
 - Receive data into an array (equivalent to read)
- `ssize_t send(int sockfd, const void *buf, size_t len, int flags);`
 - Send data from an array (equivalent to write)
- Flags: various networking-specific configurations
 - Wait for all requested data, or Never wait for data
 - Request confirmation of send, Hint that there will be more data after this
 - Etc.

Socket communication flow



Example client and server code

- Adapted from: <https://www.geeksforgeeks.org/tcp-server-client-implementation-in-c/>
 - Warning: this code is written with terrible style
 - But is a good socket example
- IPv6 server example:
<https://www.tack.ch/unix/network/sockets/udpv6.shtml>
- Sidebar: should you write your own server in C?
 - That's just asking for trouble (see buffer overflow attacks)
 - But if you're a major web company and then you might need to for performance reasons
 - Generally: use some existing server library instead

Outline


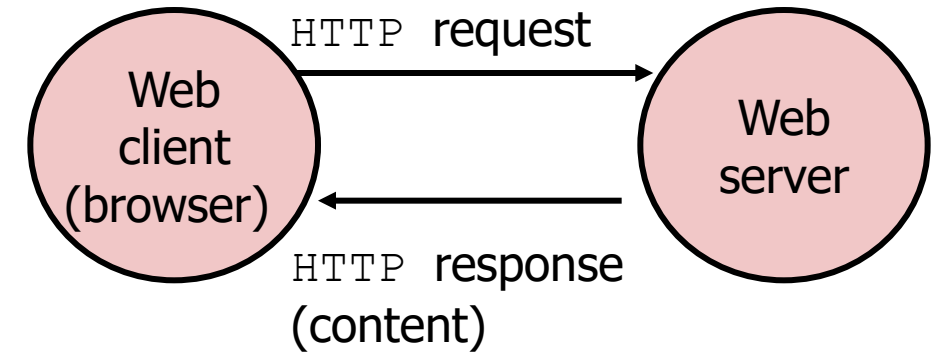
- Computer Networks
 - Topology
 - Inter-network communication
- The Internet
 - Protocol choices
- Sockets
 - System calls for communicating with other computers
- **The Web**

Internet versus the World Wide Web (WWW)

- Internet allows for communication to occur between computers all over the world
- The Web is a system for sharing file content across the Internet
 - Uses HTTP (HyperText Transfer Protocol) for communication
 - Identifies the types of files so they can be interpreted
 - Images
 - Plaintext
 - HTML (HyperText Markup Language)
 - A web browser is an application that can request and interpret web content

Web server basics

- Clients and servers communicate using HyperText Transfer Protocol (HTTP)
 1. Client and server establish TCP connection
 2. Client requests content
 3. Server responds with requested content
 4. Client and server close connection (eventually)



URLs and how they are used

- Unique name for a file: URL (Universal Resource Locator)
- Example URL: `http://www.northwestern.edu:80/index.html`
- Clients use prefix (`http://www.northwestern.edu:80`) to infer:
 - What kind (protocol) of server to contact (HTTP)
 - Where the server is (`www.northwestern.edu`)
 - What port it is listening on (80)
- Servers use suffix (`/index.html`) to:
 - Find file on file system
 - Initial `"/` in suffix denotes home directory for requested content.
 - Minimal suffix is `"/`, which server expands to configured default filename (usually, `index.html`)
 - Alternatively determine what dynamic content to serve

HTTP requests

- HTTP requests are sent from a client to a server
 - It's in plain text, so it's human readable
- Request line: **<method> <uri> <version>**
 - **<method>** is one of **GET**, **POST**, **PUT**, **DELETE**, etc.
 - **<uri>** path to the data you want to get (URI: Uniform Resource Identifier)
 - **<version>** is HTTP version of request (**HTTP/1.0** or **HTTP/1.1**)
- Request headers: **<header name>: <header data>**
 - Provide additional information to the server
 - Appended after the request line

Example HTTP requests

- `GET /index.html HTTP/1.0\r\n\r\n`
 - Method: GET
 - URI: /index.html
 - Version: Http/1.0
- More complex example:

Request

```
POST / HTTP/1.1
```

```
Host: example.com/listener
```

```
User-Agent: curl/8.6.0
```

```
Accept: */*
```

```
Content-Type: application/json
```

```
Content-Length: 345
```

```
{  
  "data": "ABC123"  
}
```

← Request headers

← Representation headers

HTTP responses

- HTTP response is sent back from the server to the client
- Response line: **<version> <status code> <status msg>**
 - **<version>** is HTTP version of the response
 - **<status code>** is numeric status
 - **<status msg>** is corresponding English text
 - 200 OK Request was handled without error
 - 301 Moved Provide alternate URL
 - 404 Not found Server couldn't find the file
- Response headers: **<header name>: <header data>**
 - Provide additional information about response
 - **Content-Type:** MIME type of content in response body
 - **Content-Length:** Length of content in response body

Example HTTP responses

- HTTP/1.1 200 OK\r\nContent-Length: 13\r\nContent-Type: text/html\r\n\r\nHello World!
 - Version: HTTP/1.1
 - Status code/message: 200 OK
 - Headers: Content-Length: 12 Content-Type: text/html
 - Data: "Hello World!"
- More complex example

Response

HTTP/1.1 200 OK

Server: Apache

Date: Fri, 21 Jun 2024 12:52:39 GMT

Cache-Control: public, max-age=3600

```
Content-Type: text/html
```

ETag: "abc123"

Last-Modified: Thu, 20 Jun 2024 11:30:00 GMT

```
<!DOCTYPE html>
```

```
<html lang="en"
```

(more data)

← Response headers

← Representation headers

Web content

- Web servers return *content* to clients
 - *content*: a sequence of bytes with an associated MIME type (Multipurpose Internet Mail Extensions) which describe how to interpret the data
- Example MIME types
 - `text/html` HTML document
 - `text/plain` Unformatted text
 - `image/gif` Binary image encoded in GIF format
 - `image/png` Binary image encoded in PNG format
 - `image/jpeg` Binary image encoded in JPEG format


List of MIME types: <http://www.iana.org/assignments/media-types/media-types.xhtml>

Static and dynamic content

- Web content is associated with a filename that is managed by the server
- The content returned in HTTP responses can be **static** or **dynamic**
 - Static content: content stored in files and retrieved in response to an HTTP request
 - Examples: HTML files, images, audio clips
 - Request identifies which content file
 - Dynamic content: content produced on-the-fly in response to an HTTP request
 - Example: content produced by a program executed by the server on behalf of the client
 - Request identifies name of content desired from executable code

The Internet is made of layers

- Each layer handles its own issues and abstracts details away
- Ethernet/WiFi sends bits and packets
- IP identifies which computer to send packets to
- TCP creates a reliable connection to a process on that computer
- HTTP allows file data to be requested and interpreted



CS340 – Intro to Computer Networking

- Computer-to-Computer Communication
 - OSI Model
 - IPv4, IPv6, and Routing
 - TCP, UDP
 - Application Protocols: HTTP, FTP
 - Sockets and Web Servers
- A systems class that feels more *different* from CS213 rather than more similar
- Usually taught twice a year

Outline

- Computer Networks
 - Topology
 - Inter-network communication
- The Internet
 - Protocol choices
- Sockets
 - System calls for communicating with other computers
- The Web