CSSE132
Introduction to Computer Systems

28 : Internetworking

April 25, 2013

Today

- **Networks**
- **internet (lowercase i)**
- **Internet (uppercase i)**
  - History and evolution
  - IP addresses and domain names
  - Sockets

A Client-Server Transaction

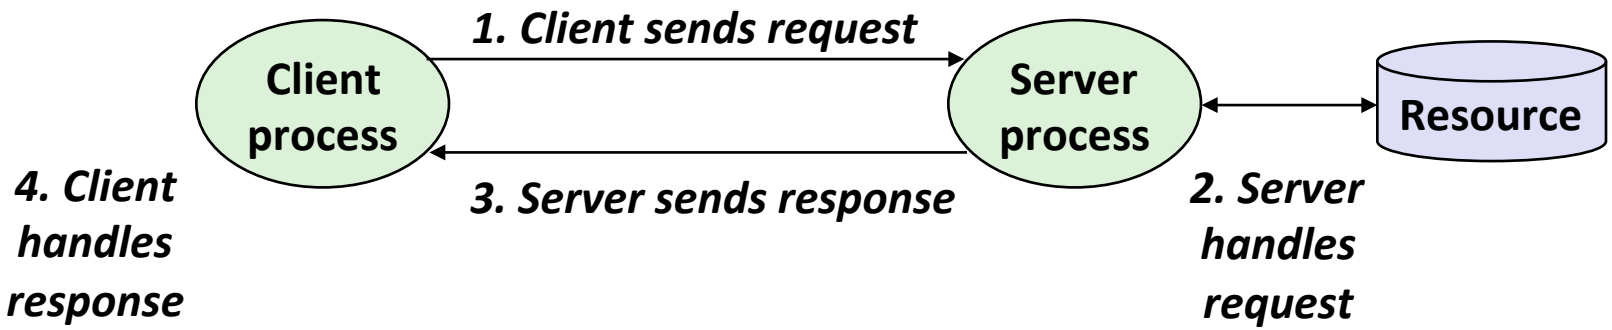

Note: clients and servers are processes running on hosts (can be the same or different hosts)

- **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 (vending machine analogy)

Hardware Organization of a Network Host

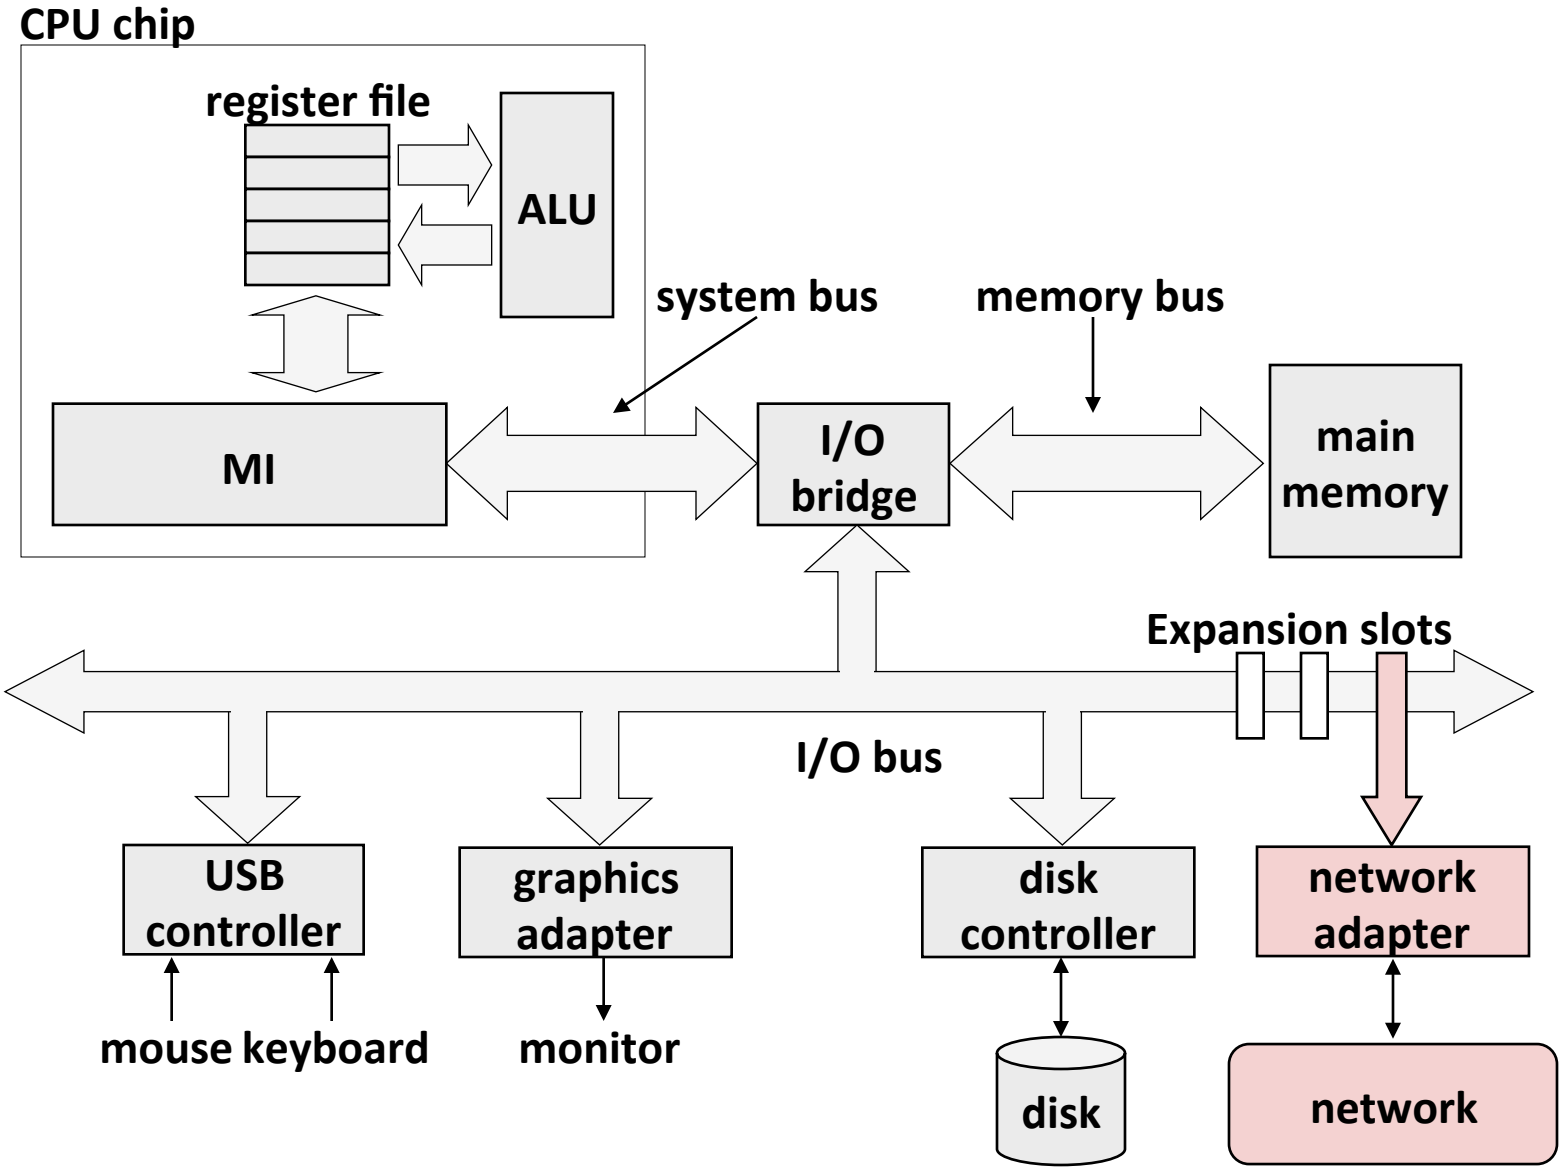

Computer Networks

- **A *network* is a hierarchical system of boxes and wires organized by geographical proximity**
  - SAN (System Area Network) spans cluster or machine room
    - Switched Ethernet, Quadrics QSW, …
  - LAN (Local Area Network) spans a building or campus
    - Ethernet is most prominent example
  - WAN (Wide Area Network) spans country or world
    - Typically high-speed point-to-point phone lines
- **An *internetwork* (*internet*) is an interconnected set of networks**
  - The Global IP Internet (uppercase “I”) is the most famous example of an internet (lowercase “i”)
- **Let’s see how an internet is built from the ground up**

Today

- **Networks**
- **internet (lowercase i)**
- **Internet (uppercase i)**
  - History and evolution
  - IP addresses and domain names
  - Sockets

Switches, routers and subnets

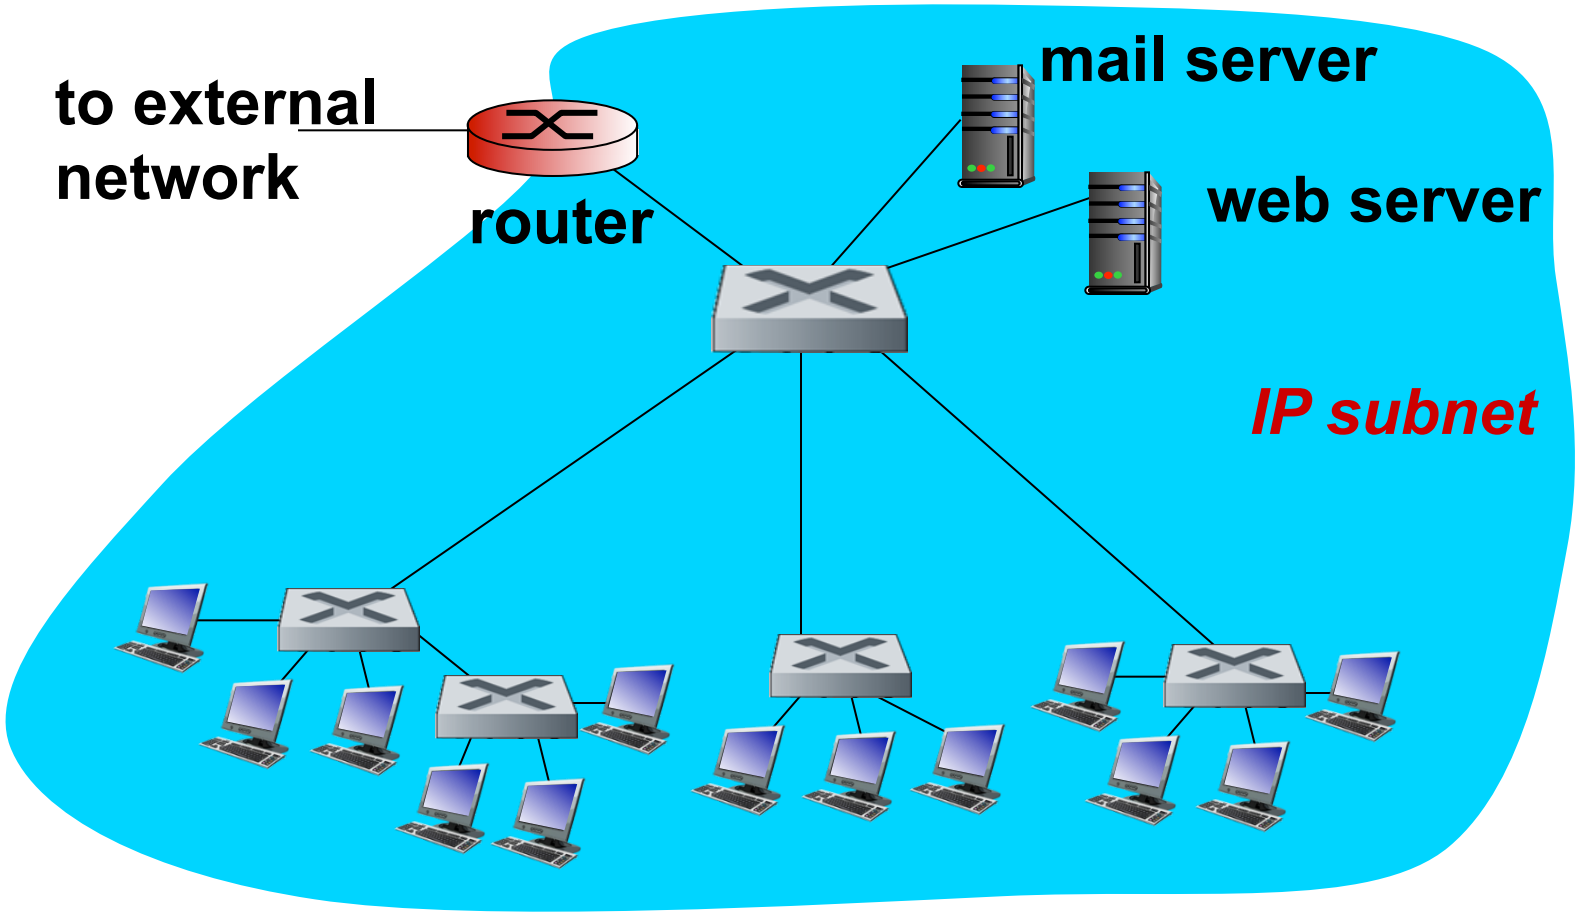

Logical Structure of an internet

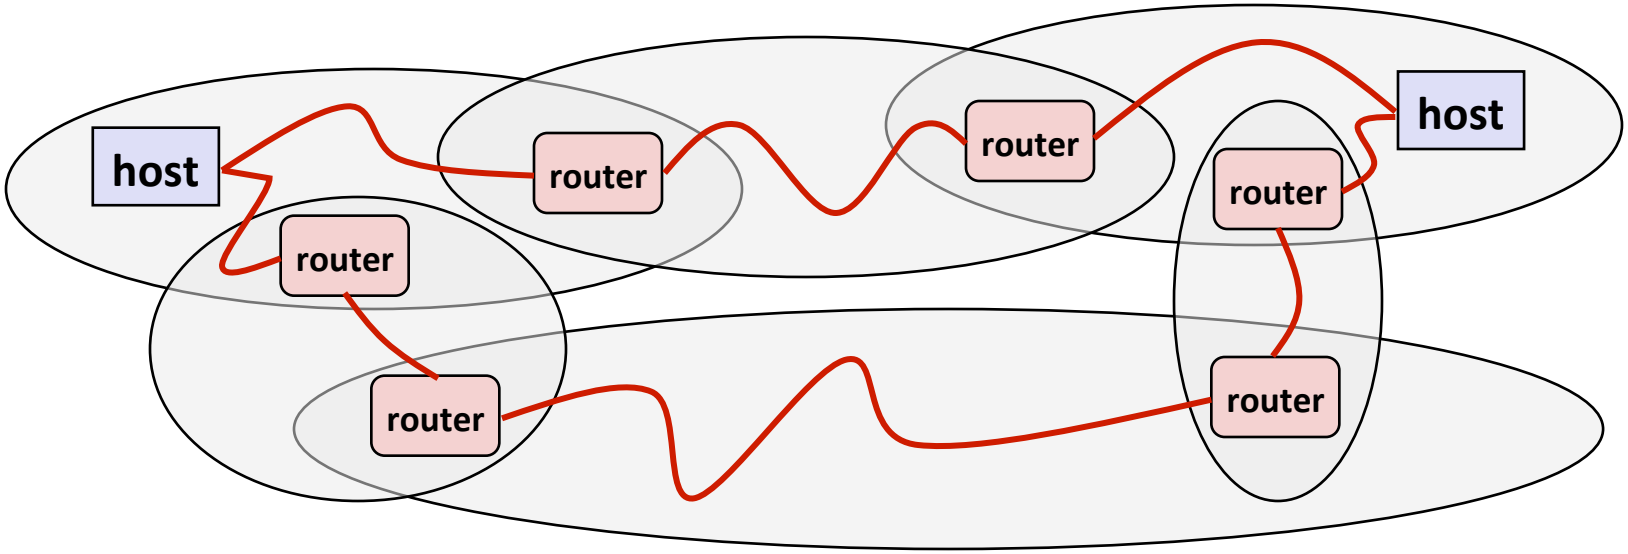

- **Ad hoc interconnection of networks**
  - No particular topology
  - Vastly different router & link capacities
- **Send packets from source to destination by hopping through networks**
  - Router forms bridge from one network to another
  - Different packets may take different routes

The Notion of an internet Protocol

- **How is it possible to send bits across incompatible LANs and WANs?**

- **Solution:**
  - protocol software running on each host and router
  - smooths out the differences between the different networks

- **Implements an internet protocol (i.e., set of rules)**
  - governs how hosts and routers should cooperate when they transfer data from network to network
  - TCP/IP is the protocol for the global IP Internet

What Does an internet Protocol Do?

- **Provides a naming scheme**
  - An internet protocol 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**
  - An internet protocol defines a standard transfer unit (***packet***)
  - Packet consists of ***header*** and ***payload***
    - Header: contains info such as packet size, source and destination addresses
    - Payload: contains data bits sent from source host

Internet protocol stack

- ***application:*** **supporting network applications**
  - FTP, SMTP, HTTP
- ***transport:*** **process-process data transfer**
  - TCP, UDP
- ***network:*** **routing of datagrams from source to destination**
  - IP, routing protocols
- ***link:*** **data transfer between neighboring network elements**
  - Ethernet, 802.111 (WiFi), PPP
- ***physical:*** **bits "on the wire"**

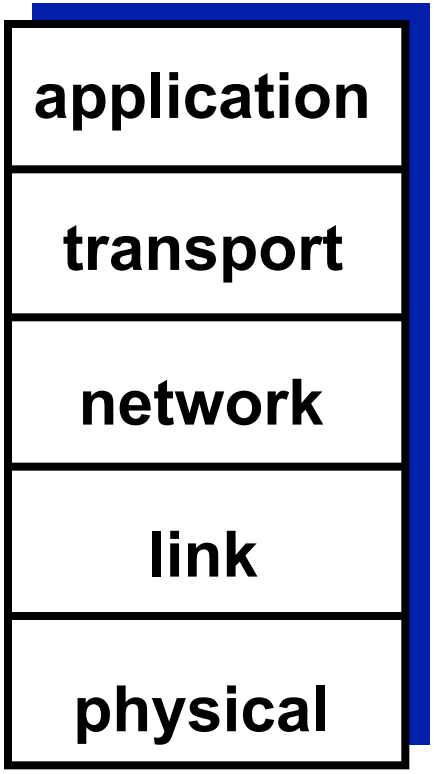

Other Issues

- **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***

Today

- **Networks**
- **internet (lowercase i)**
- **Internet (uppercase i)**
  - History and evolution
  - IP addresses and domain names
  - Sockets

Global IP Internet

- **Most famous example of an internet**

- **Based on the TCP/IP protocol family**
  - IP (Internet protocol) :
    - Provides *basic naming scheme* and unreliable *delivery capability* of packets (datagrams) from host-to-host
  - UDP (Unreliable Datagram Protocol)
    - Uses IP to provide unreliable datagram delivery from *process-to-process*
  - TCP (Transmission Control Protocol)
    - Uses IP to provide *reliable* byte streams from process-to-process over connections

- **Accessed via a mix of Unix file I/O and functions from the *sockets interface***

Hardware and Software Organization of an Internet Application

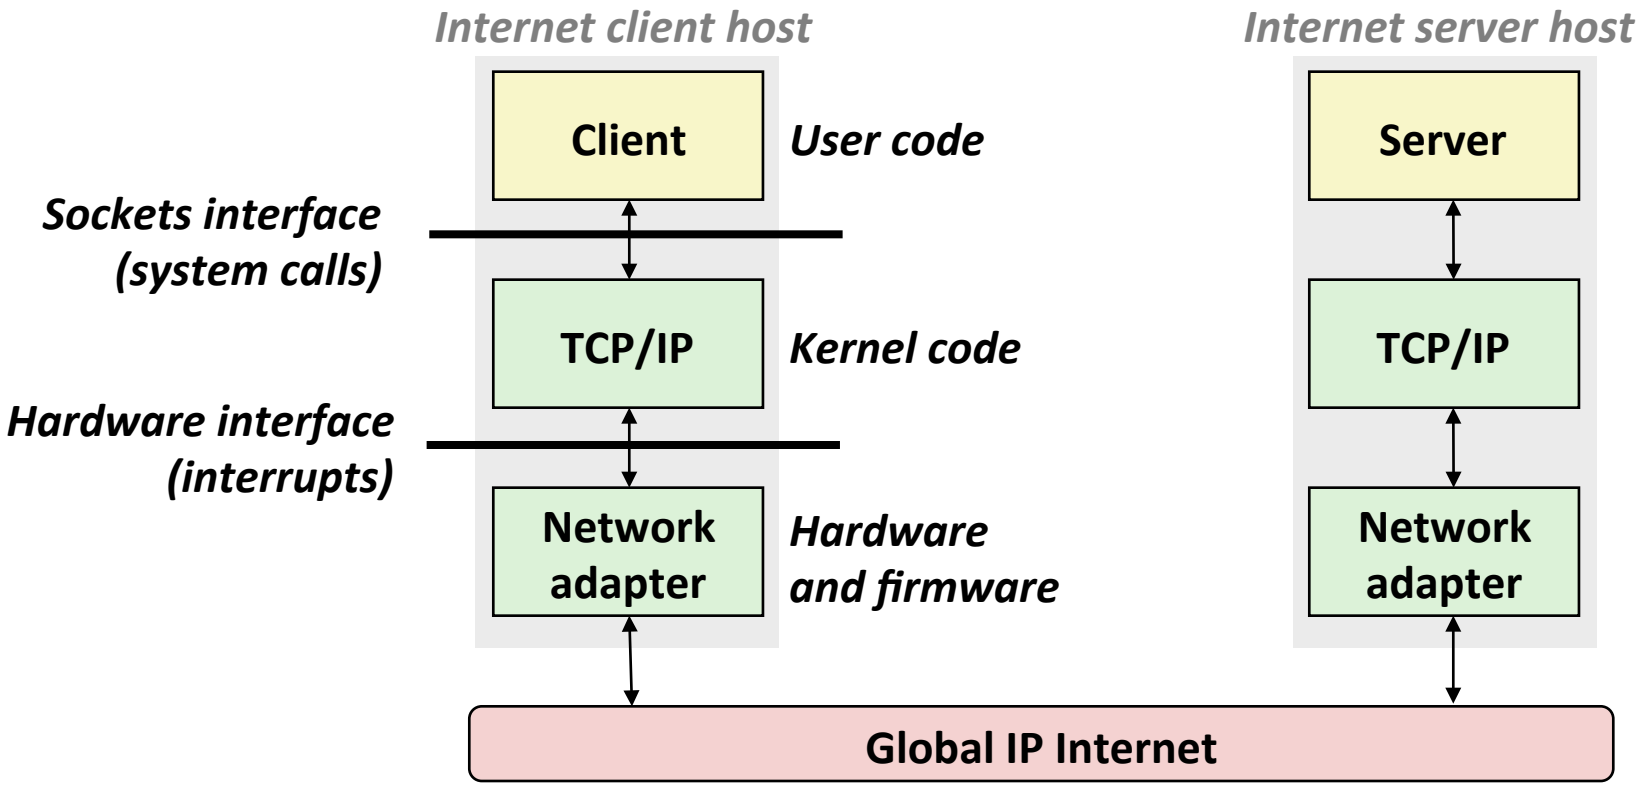

History

- Launched as ARPAnet in 1969
- Switched to TCP/IP in 1983
- DNS (naming service) invented in 1985
- Web invented in 1989

Naming and Communicating on the Internet

- **Original Idea**
  - Every node on Internet would have unique IP address
    - Everyone would be able to talk directly to everyone
  - No secrecy or authentication
    - Messages visible to routers and hosts on same LAN
    - Possible to forge source field in packet header

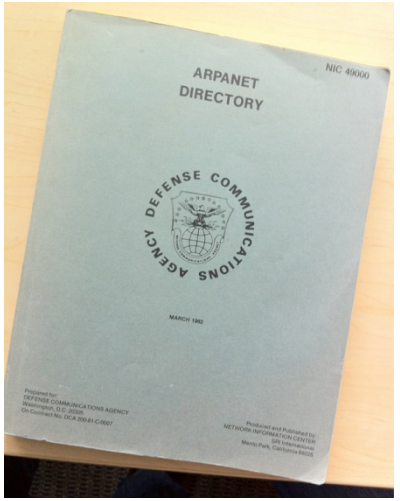

- **Shortcomings**
  - There aren't enough IP addresses available (2^{32} possible)
  - Don't want everyone to have access or knowledge of all other hosts
  - Security issues mandate secrecy & authentication
  - Hard to remember numeric addresses

Evolution of Internet: Naming

- **Dynamic address assignment**
  - Most hosts don't need to have known address
    - Only those functioning as servers
  - **DHCP** (Dynamic Host Configuration Protocol)
    - Local ISP assigns address for temporary use

- **Example:**
  - ifconfig (Unix)
  - ipconfig (Windows)

Evolution of Internet: NAT

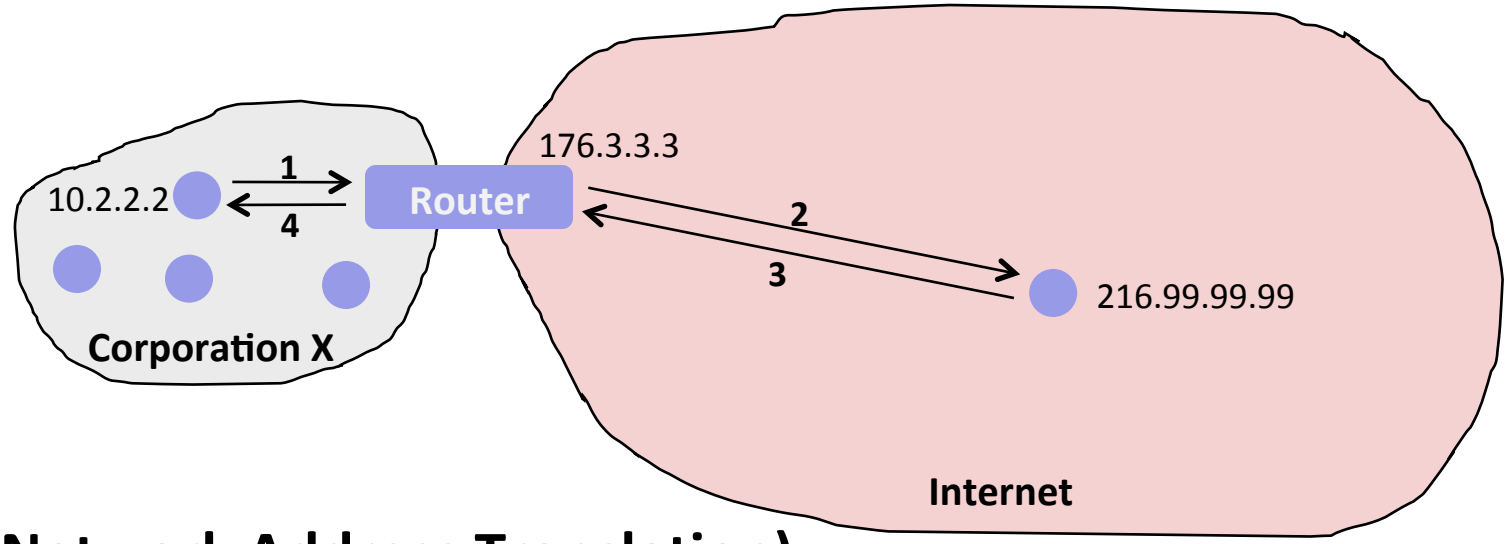

- **NAT (Network Address Translation)**
  - Hides organizations nodes from rest of Internet
  - Use local IP addresses within organization
  - For external service, provides proxy service
    1. Client request: src=10.2.2.2, dest=216.99.99.99
    2. Router forwards: src=176.3.3.3, dest=216.99.99.99
    3. Server responds: src=216.99.99.99, dest=176.3.3.3
    4. Router forwards response: src=216.99.99.99, dest=10.2.2.2

Virtual Private Networks

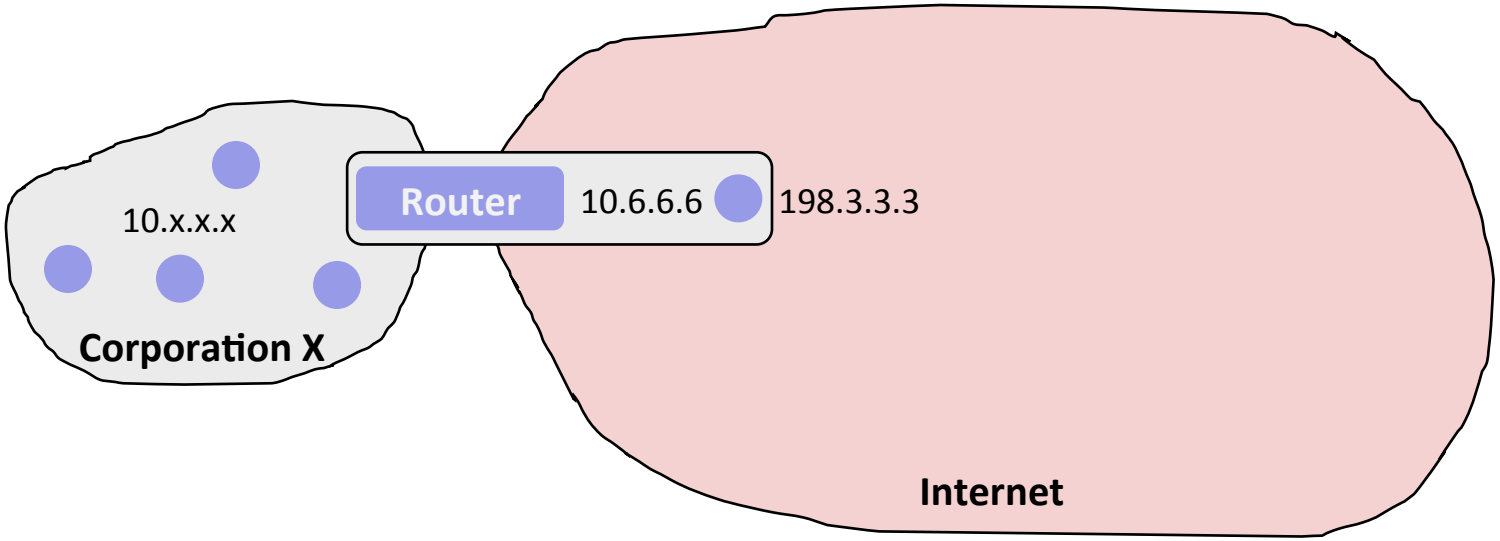

- **Supporting road warrior**
  - Employee working remotely with assigned IP address 198.3.3.3
  - Wants to appear to rest of corporation as if working internally
    - From address 10.6.6.6
    - Gives access to internal services (e.g., ability to send mail)
- **Virtual Private Network (VPN)**
  - Overlays private network on top of regular Internet

A Programmer's View of the Internet

- **Hosts are mapped to a set of 32-bit *IP addresses***
  - 128.2.203.179

- **The set of IP addresses is mapped to a set of identifiers called Internet *domain names***
  - 128.2.203.179 is mapped to www.cs.cmu.edu

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

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 byte order)
  - 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 {
    unsigned int s_addr; /* network byte order (big-endian) */
};
```

Useful network byte-order conversion functions ("l" = 32 bits, "s" = 16 bits)

`htonl`: convert uint32_t from host to network byte order
`htons`: convert uint16_t from host to network byte order
`ntohl`: convert uint32_t from network to host byte order
`ntohs`: convert uint16_t from network to host byte order

Dotted Decimal Notation

- **By convention, each byte in a 32-bit IP address is represented by its decimal value and separated by a period**
  - IP address: `0x8002C2F2 = 128.2.194.242`

- **Functions for converting between binary IP addresses and dotted decimal strings:**
  - **`inet_aton`:** dotted decimal string → IP address in network byte order
  - **`inet_ntoa`:** IP address in network byte order → dotted decimal string

  - "n" denotes network representation
  - "a" denotes application representation

Internet Domain Names

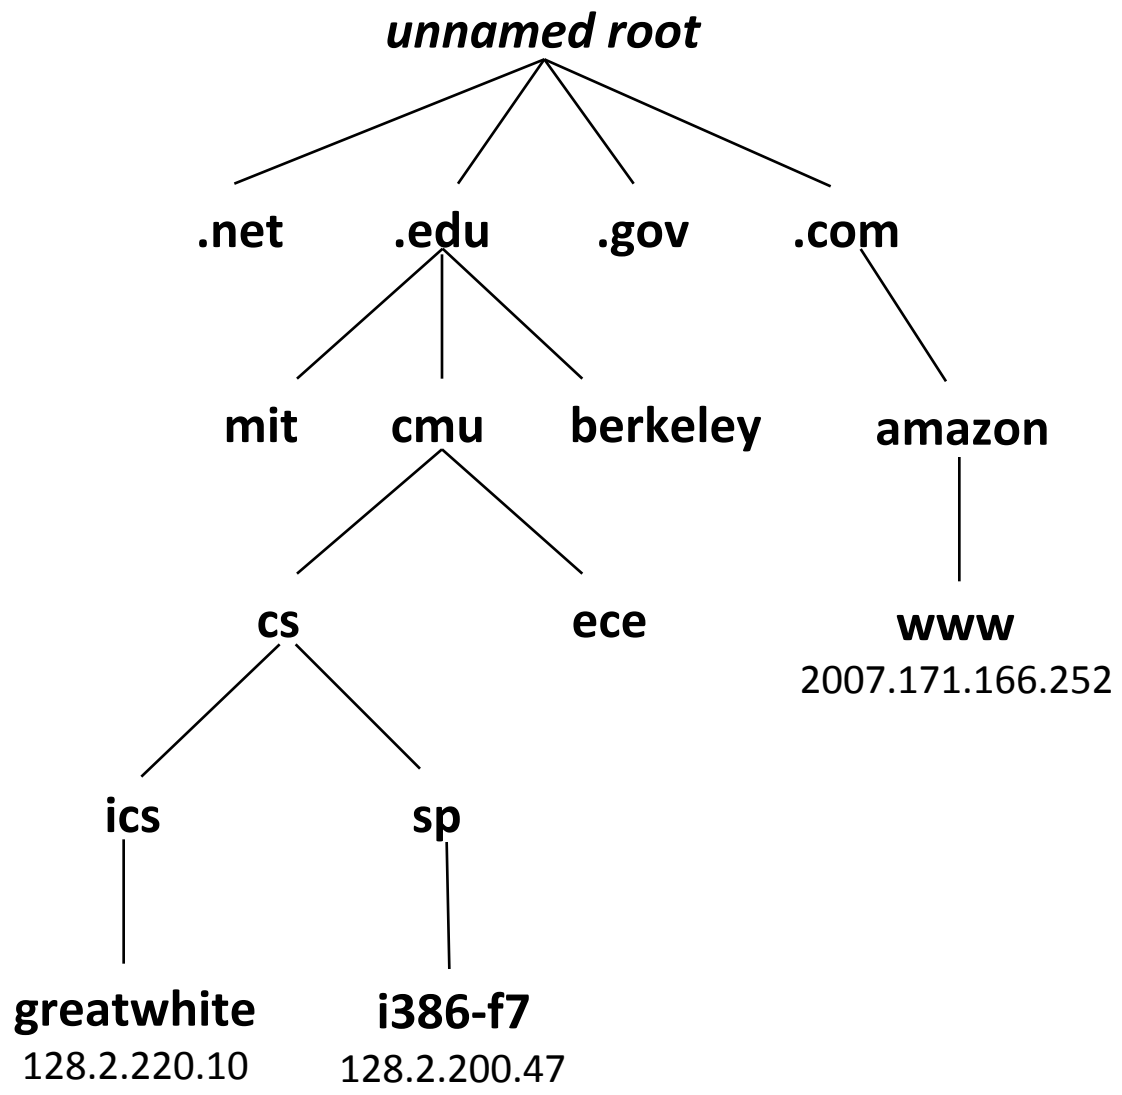

First-level domain names

Second-level domain names

Third-level domain names

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 entry structures*:

```
/* DNS host entry structure */
struct hostent {
   char    *h_name;        /* official domain name of host */
   char    **h_aliases;    /* null-terminated array of domain names */
   int     h_addrtype;     /* host address type (AF_INET) */
   int     h_length;       /* length of an address, in bytes */
   char    **h_addr_list; /* null-terminated array of in_addr structs
*/
};
```

- **Functions for retrieving host entries from DNS:**
  - **`gethostbyname`**: query key is a DNS domain name.
  - **`gethostbyaddr`**: query key is an IP address.

Properties of DNS Host Entries

- **Each host entry is an equivalence class of domain names and IP addresses**
- **Each host has a locally defined domain name `localhost` which always maps to the *loopback address* `127.0.0.1`**
- **Different kinds of mappings are possible:**
  - Simple case: one-to-one mapping between domain name and IP address:
    - **`greatwhile.ics.cs.cmu.edu`** maps to **`128.2.220.10`**
  - Multiple domain names mapped to the same IP address:
    - **`eecs.mit.edu`** and **`cs.mit.edu`** both map to **`18.62.1.6`**
  - Multiple domain names mapped to multiple IP addresses:
    - **`google.com`** maps to multiple IP addresses
  - Some valid domain names don't map to any IP address:
    - for example: **`ics.cs.cmu.edu`**

Internet Connections

- **Clients and servers communicate by sending streams of bytes over *connections*:**
  - Point-to-point, full-duplex (2-way communication), and reliable.

- ***A socket* is an endpoint of a connection**
  - Socket address is an `IPaddress:port` pair

- **A *port* is a 16-bit integer that identifies a process:**
  - ***Ephemeral port*:** Assigned automatically on client when client makes a connection request
  - ***Well-known port:*** Associated with some service provided by a server (e.g., port 80 is associated with Web servers)

- **A connection is uniquely identified by the socket addresses of its endpoints (*socket pair*)**
  - `(cliaddr:cliport, servaddr:servport)`

Putting it all Together: Anatomy of an Internet Connection

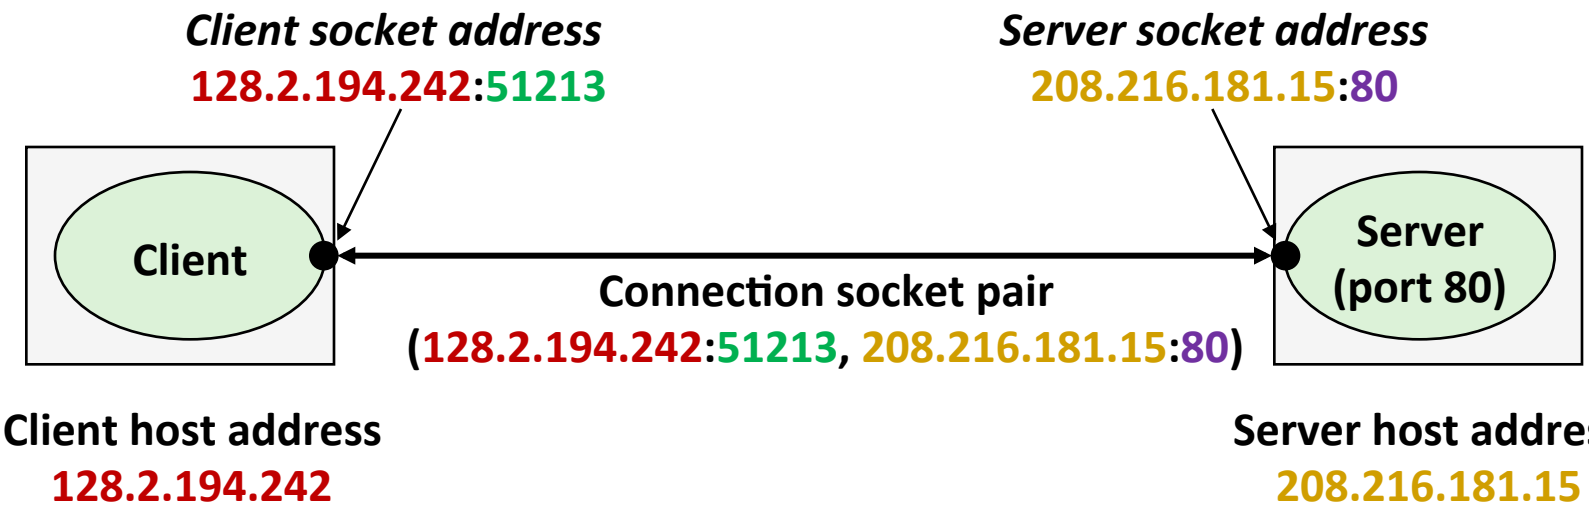

Next Time

- How to use the sockets interface to establish Internet connections between clients and servers
- How to use Unix I/O to copy data from one host to another over an Internet connection