



# Internet Fundamentals

Jeffrey Shafer


Some slides adapted from CMU 15.213 slides

# 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”)

# Ethernet Segment

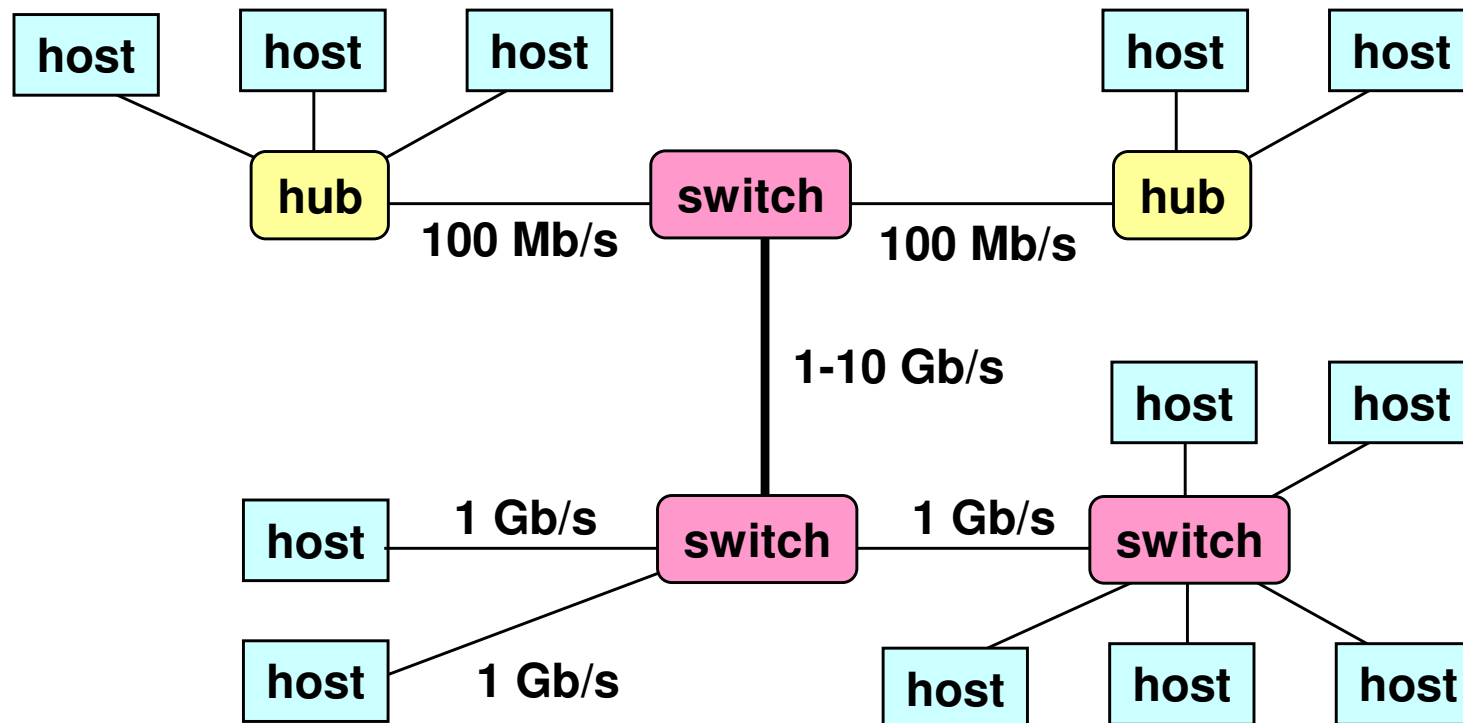
- Ethernet segment consists of a collection of hosts connected by wires to a hub



- Operation
  - Each Ethernet adapter has a unique 48-bit address
  - Hosts send bits to any other host in chunks called frames
  - Hub slavishly copies each bit from each port to every other port
    - Every host sees every bit


# Switched Ethernet Segment

- Spans room, building, or campus
- Switches (bridges)
  - Learn which hosts are reachable from which ports
  - Selectively copies frames from port to port



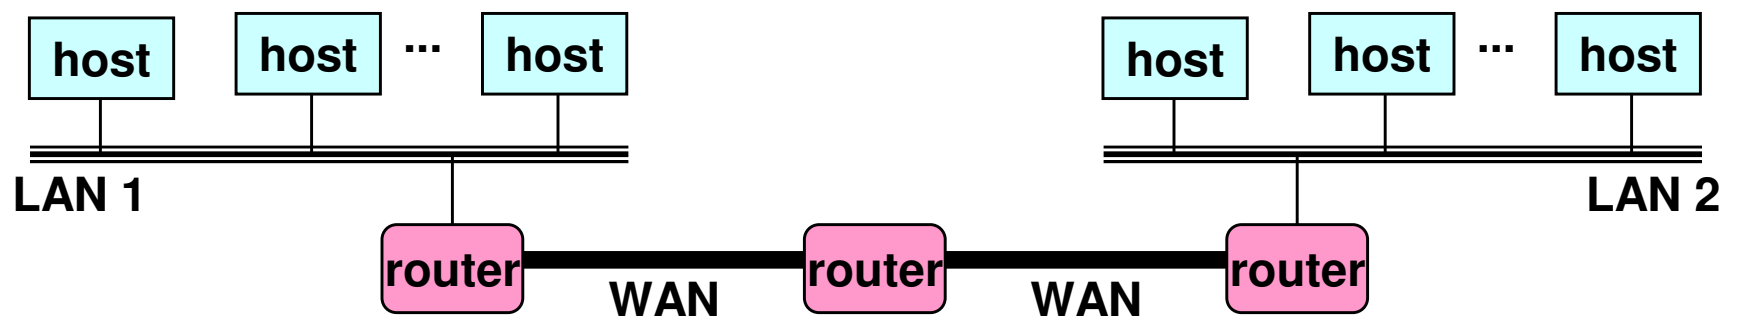
# Conceptual View of LANs

- For simplicity, hubs, switches, and wires are often shown as a collection of hosts attached to a single wire:



# Next Level: internets

- Multiple incompatible LANs can be physically connected by specialized computers called routers
- The connected networks are called an internet



**LAN 1 and LAN 2 might be completely different, totally incompatible LANs (e.g., Ethernet and WiFi, 802.11\*, T1-links, DSL, ...)**



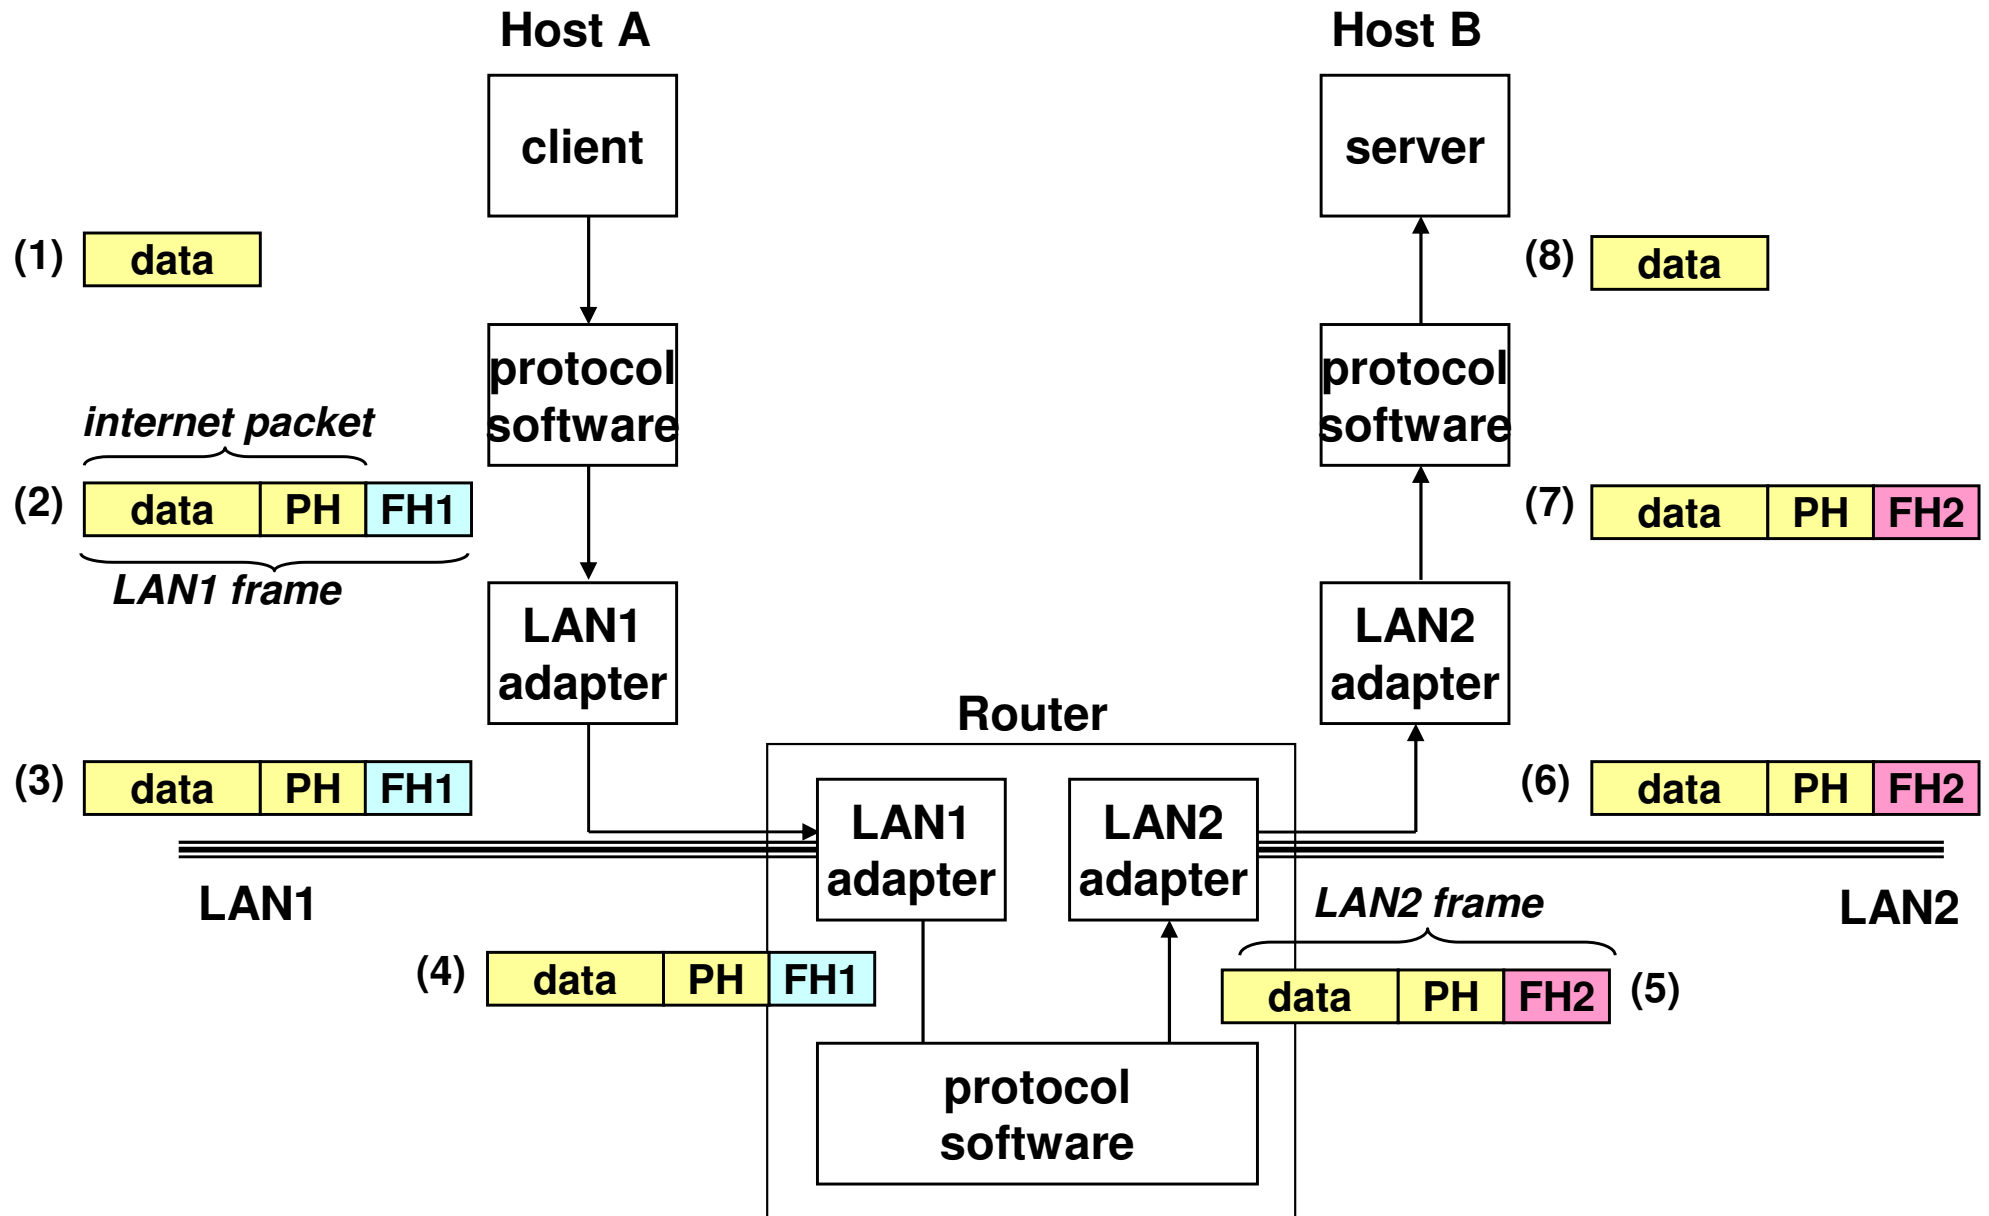
# The Notion of an internet Protocol

- How is it possible to send bits across incompatible LANs and WANs?
- Solution: protocol software running on each host and router smoothes out the differences between the different networks
- Implements an internet protocol (i.e., set of rules) that 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

# Transferring Data Over an internet



# Issues

- What if different networks have different maximum frame sizes? (segmentation)
- What if the router is connected to more than two LANs?
- What if the router is connected to other routers?
- What if packets get lost? (buffer overflows, errors, failures, etc.)

# Global IP Internet

- Most famous example of an internet
- Characteristics
  - Each packet is individually routed
  - No time guarantee for delivery
  - No guarantee of delivery in sequence
  - No guarantee of delivery at all!
  - No guarantee of integrity of data
  - Packets can be fragmented
  - Packets may be duplicated

# Layering in the Internet

- Application Layer
  - Provides application-specific functions (end-to-end reliability, data interpretation, etc.)
- Transport Layer
  - Provides reliable (if desired), in-sequence (if desired) delivery of data on behalf of application
- Network Layer
  - Provides “best-effort”, but unreliable, delivery of datagrams
- Link Layer
  - Carries data over (usually) point-to-point links between hosts and routers; or between routers and routers

# Protocols

## ■ TCP/IP Protocol Family

### □ Network Layer

#### ■ IP (Internet protocol) :

- Provides basic naming scheme and unreliable delivery capability of packets (datagrams) from host-to-host

### □ Transport Layer

#### ■ UDP (User Datagram Protocol)

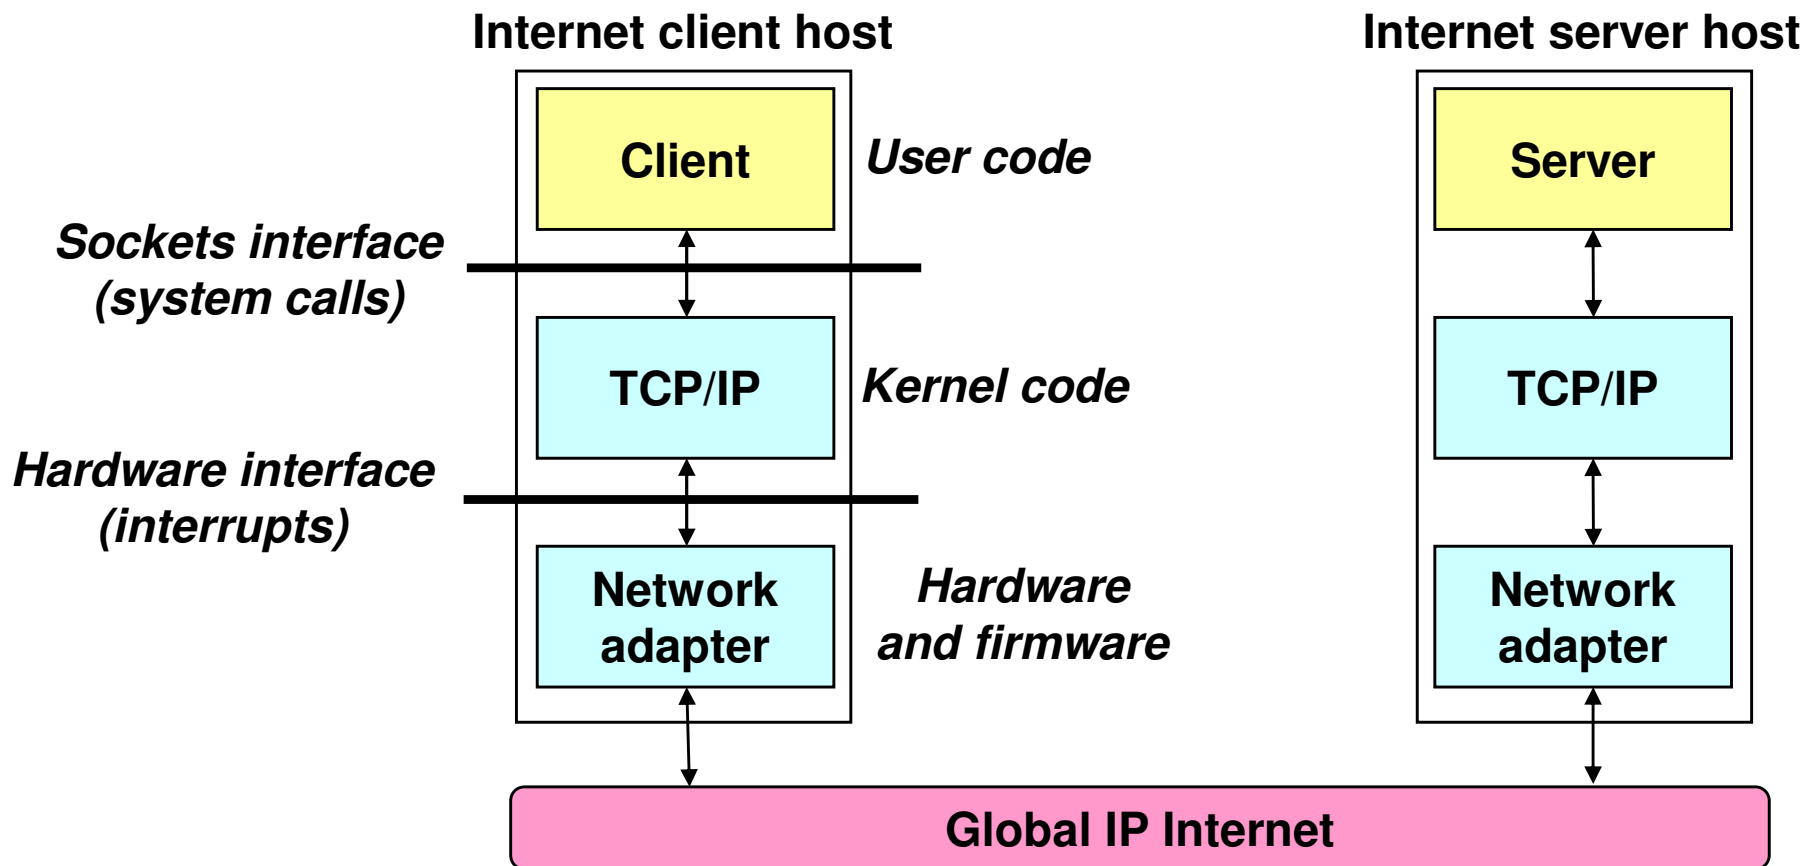
- Uses IP to provide unreliable datagram delivery from process-to-process

#### ■ TCP (Transmission Control Protocol)

- Uses IP to provide reliable byte streams from process-to-process over connections

## ■ Newer transport layer protocols exist or are being developed (e.g., SCTP)

# Organization of an Internet Application





# A Programmer's View of the Internet

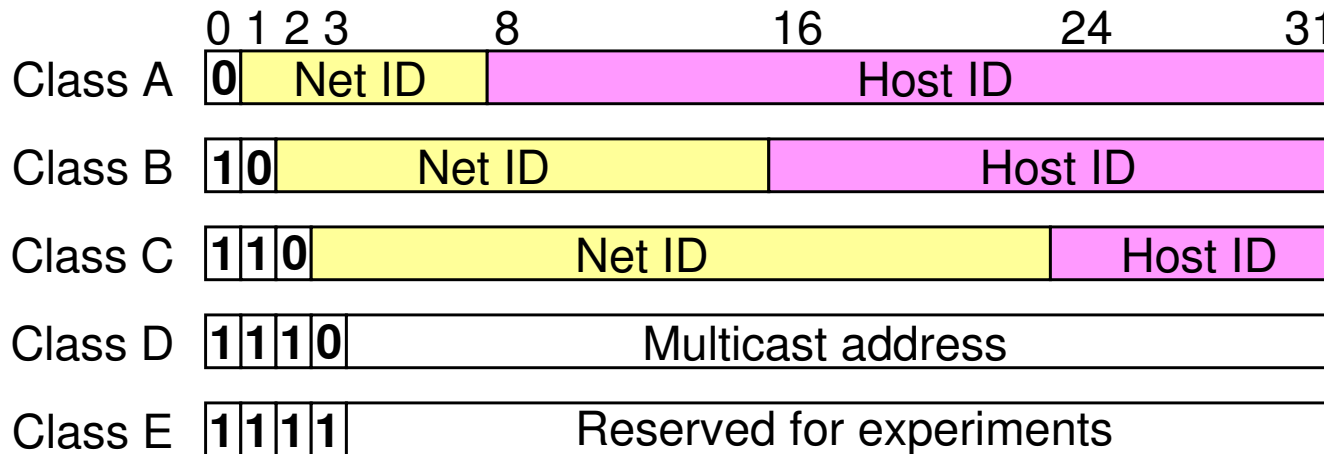
1. Hosts are mapped to a set of 32-bit IP addresses
  - 128.42.1.125
2. A set of identifiers called Internet domain names are mapped to the set of IP addresses for convenience
  - www.cs.rice.edu is mapped to 128.42.1.125
3. A process on one Internet host can communicate with a process on another Internet host over a connection

# 1. IP Addresses

- IP version 4 addresses are 32-bits
- 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 a TCP connection
- IP addresses are usually displayed in dotted decimal notation
  - Each byte represented by decimal value
  - Bytes are separated by a period
  - IP address  $0x8002C2F2 = 128.2.194.242$

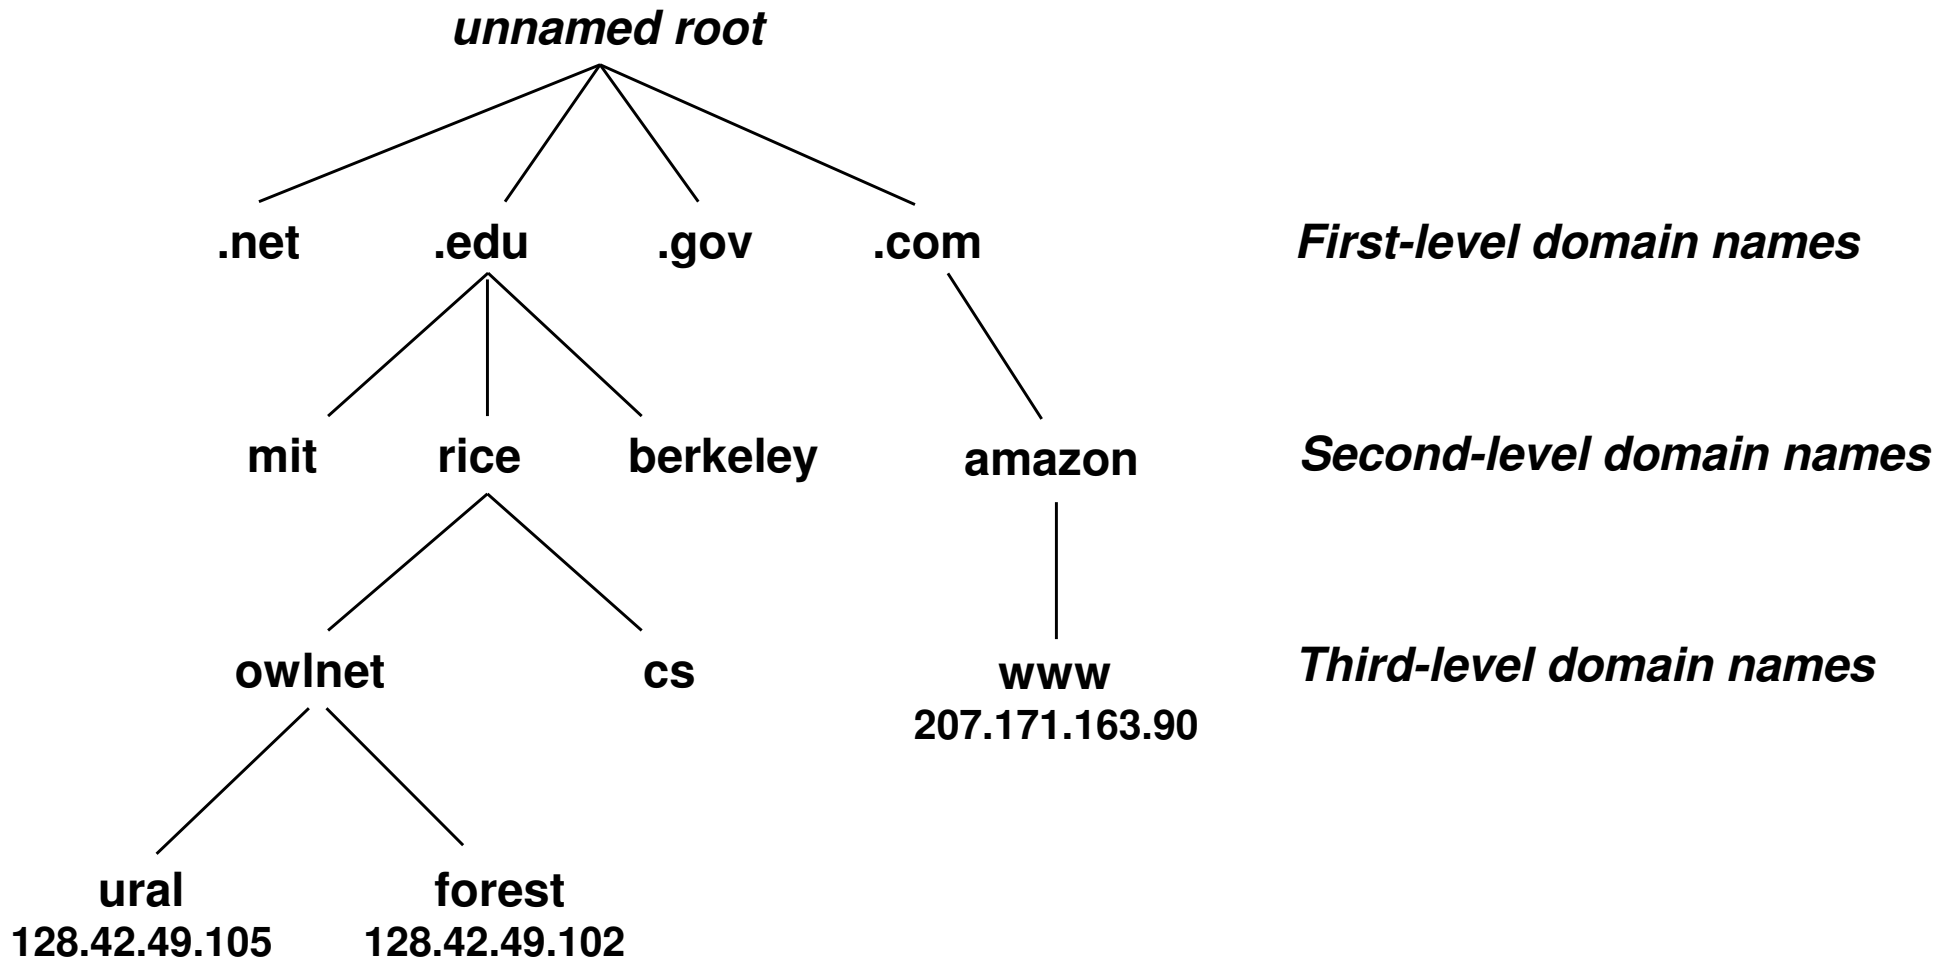
# IP Address Structure

- IP (V4) Address space divided into classes:



- Loop-back address: 127.0.0.1
- Unrouted (private) IP addresses:
  - 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16
- Dynamic IP addresses (DHCP)

# 2. Internet 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 */
};
```


# 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: 1 domain name maps to one IP address:
    - `forest.owlnet.rice.edu` maps to `128.42.49.102`
  - Multiple domain names mapped to the same IP address:
    - `www.cs.rice.edu`, `eecs.rice.edu`, and `bianca.cs.rice.edu` all map to `128.42.1.125`
  - Multiple domain names mapped to multiple IP addresses:
    - `aol.com` and `www.aol.com` map to multiple IP addresses


# 3. Internet Connections

- Clients and servers communicate by sending streams of bytes over TCP connections:
  - Point-to-point, full-duplex (2-way communication), and reliable
- A socket is an endpoint of a connection
  - Socket address is an IP address, 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)

# Anatomy of an Internet Connection



# Transiting the Internet



# Reliability

- Packets can get lost for a variety of reasons
  - Buffer overflow
  - Failures
  - Etc...
- IP is not reliable, so the “network” does not need to do anything
- It is up to the communication endpoints to ensure reliable delivery, if necessary
  - TCP
  - Application-level reliability protocols

# Segmentation

- Networks along the path may be different types
  - Wireless
  - Ethernet
  - ...
- What if they have different maximum frame sizes?
  - Routers must segment IP packets!

# Dynamic Routing

- How do routers know where to forward frames?
  - Must find a route initially
  - Routes may change over time
    - Failures
    - Congestion
    - Etc.
- Dynamic routing protocols
  - OSPF
  - BGP

# Routers

- Deal mostly with network (e.g., IP) and link (e.g., Ethernet) layers
  - Best-effort, unreliable delivery
  - Point-to-point communication with other routers/hosts
- More sophisticated network systems have emerged that blur the boundaries
  - Route traffic
  - Have knowledge of transport (e.g., TCP) and application (e.g., HTTP) layers



# Next

- Ethernet – link layer protocol
- Internet Protocol – network layer protocol
- IP Routing
- ... and NetFPGA architecture  
(so we can start projects!)