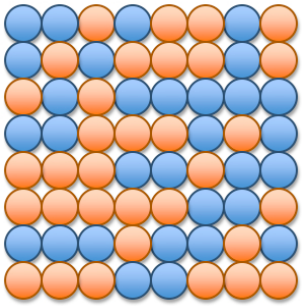

UNIT 13A

The Internet: Fundamentals

What is the Internet?

- The Internet is a system to deliver data (bits) from one computational device to another.
- No one entity controls/owns the Internet.
- The Internet is governed by protocols and standards that are commonly agreed to by developers of network software and applications.

ARPANET to Internet

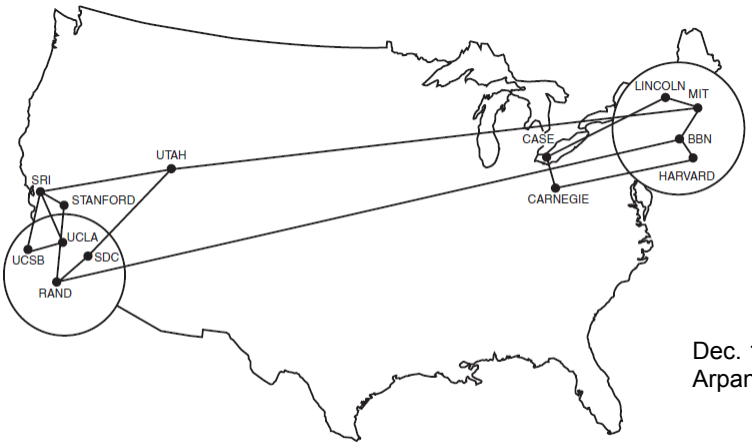

Dec. 1970
Arpanet

Source: Heart, F., McKenzie, A., McQuillian, J., and Walden, D., ARPANET Completion Report, Bolt, Beranek and Newman, Burlington, MA, January 4, 1978.

ARPANET to Internet

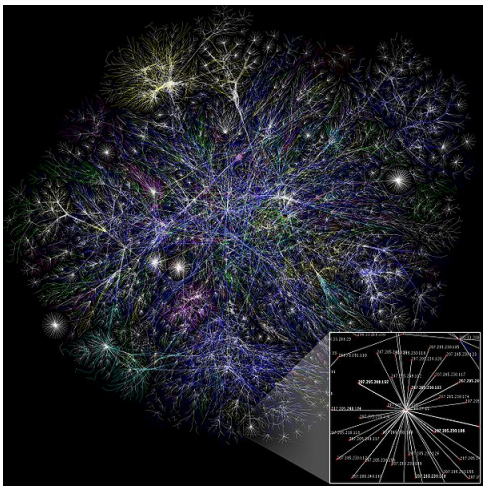

2000's
Internet Map
(small section)

Structure of the Internet

- Core
  - routers
  - gateways
  - Internet Service Providers (ISP's)
  - domain name servers
- Edges
  - individual users
  - private networks

Circuit Switching

- Two network nodes (e.g. phones) establish a dedicated connection via one or more switching stations.

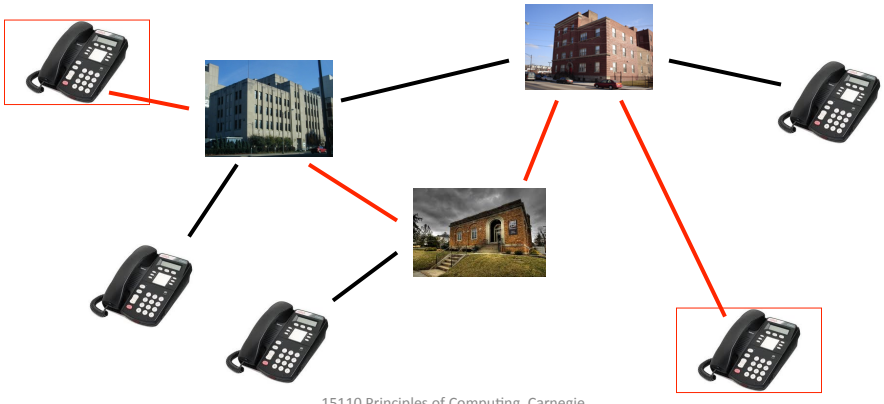

Packet Switching

- Two network nodes (e.g. computers) send messages by breaking the message up into small packets and sending each packet on to the network with a serial number and a destination address.
- Routers in the network use a buffer (queue) to hold packets until they can be routed toward their destination.
- Packets may be received at the destination in any order and may get lost and retransmitted. Serial numbers are used to put packets back into order at the destination.

Packet Switching

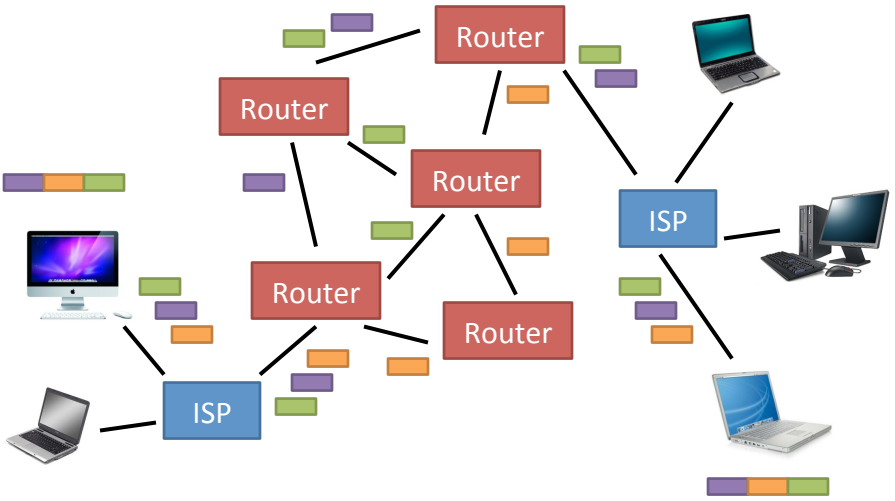

IP Addresses (IPv4)

- Computers on the internet are assigned an IP Address

  ____ . ____ . ____ . ____

  Four numbers between 0 and 255, inclusive.

  Example: 128.2.13.163
- This means that each part of the address is an 8-bit value, and an IP address is 32 bits.

  → supports up to 2^{32} computers on the network at the same time
- ISPs can reassign IP addresses dynamically.

Network Address Translation (NAT)

- To accommodate more users on the Internet, NAT is used.
- The gateway assigns an additional code called a port for each user. Packets are tagged with the port.
- The gateway knows where to route the messages on the private network, but all messages from that private network share the same single IP address.

Network Address Translation (NAT)

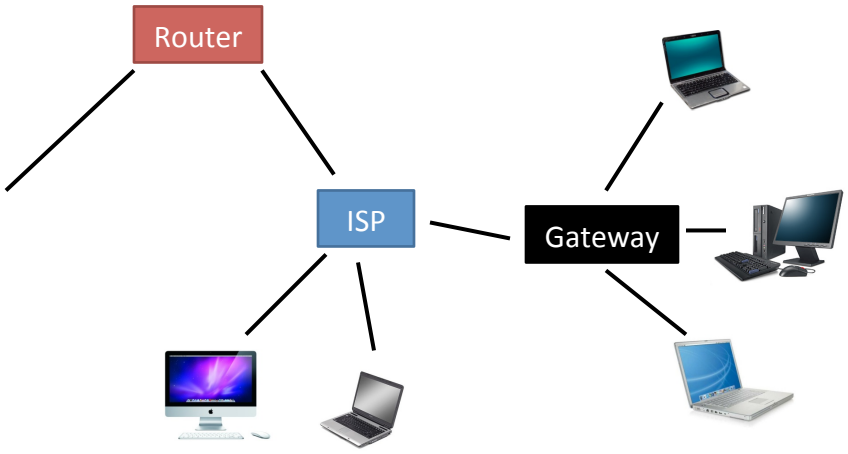

IPv4 Address Assignment

- The original IPv4 had several classes of addresses:
  - Class A **0** + 7-bit network + 24-bit address
    Accommodates up to 2^{24} unique IP addresses in a company or location.
  - Class B **10** + 14 bit network + 16-bit address
    Accommodates up to 2^{16} unique IP addresses in a company or location.
  - Class C **110** + 21-bit network + 8-bit addr
    Accommodates up to 2^{8} unique IP addresses in a company or location.

IPv4 Address Assignment

- In 1993, the Internet switched to classless internet-domain routing. In this scheme, the network part is an arbitrary length prefix of the address, such as 10.10.1.32/27, which has a 27-bit network part and a 5-bit address part (so there can only be 32 machines on that network).
- IPv6 also follows classless routing, but the standard subnetwork size is 64-bits (which allows using the MAC address manufactured into each ethernet card as the local part. Normally 48-bit prefixes of IPv6 are assigned to individual organizations, allowing each organization to have a 65,535 subnetworks with up to 2^{64} machines per subnetwork.

New IP (IPv6)

- IPv6 uses 128-bit addresses
  → supports 2^{128} unique computer addresses
  → = 3.4 X 10^{38}
- Allows for many more devices (cell phones, video game machines, appliances, automobiles, etc.)
- Designed to deal with the approaching use of all available addresses in IPv4.

ISPs

- An Internet Service Provider (ISP) is a company that provides access for users to the Internet.
  - AT&T, Comcast, EarthLink, Verizon, etc.
  - access can be provided via copper cable, wireless transmission, fiber optic cable, etc.
  - In rural areas, an ISP may be a company providing Internet services by satellite.
  - Universities (like CMU) and big companies (like Google and Microsoft) are their own ISPs.

Internet Protocol (IP)

- A *protocol* is a standard for communicating messages between networked computers.
- An IP address in each packet determine the intended destination of the packet.
- A domain name server translates machine names to equivalent IP addresses to make it easier for users to indicate message destinations.
  - Example: www.cnn.com, unix.andrew.cmu.edu, employees.verizon.net

Transfer Control Protocol (TCP)

- TCP is the main protocol used on the Internet to transmit messages using packets.
  - used for the web, email and file transfer
- TCP can detect if a packet is lost, delivered out of order or duplicated.
- TCP is optimized for accurate delivery rather than timely delivery.
  - For streaming data, other protocols are used (e.g. UDP) where packet loss is not as critical.

TCP and "Handshaking"

- The process of two parties determine that each has received the other's transmission correctly is called *handshaking*.
  - Alice sends several packets to Bob using TCP.
  - Each packet includes parity information so Bob can check its accuracy.
  - When Bob receives a packet, if it is ok, Bob sends an acknowledgement packet back to Alice.
  - If Bob is missing a packet, he can send a request for a retransmission of the packet.
  - If Alice doesn't get an acknowledgement within a set period of time, she can retransmit the packet.

Routers

- Routers are considered to be very simple devices whose sole purpose is to route data traffic.
- The end-to-end principle in the Internet
- Routers only implement IP by routing packets. It is up to the end units to run the more involved TCP to check for transmission errors, omissions and duplications.

Fault Tolerance

- The Internet is subject to faults at individual nodes. The protocols are designed to allow data traffic to be rerouted if nodes go down or become too overloaded.
  - World Trade Center Attack (9/11/2001)
  - New Orleans & Hurricane Katrina (2005)
  - Henchung earthquake (2006)
  - Qatar Internet blackout (2008)

Communications Protocol

- agreement between communicating parties
  - syntax: how are the messages' contents organized?
  - semantics: what do the messages mean?
  - synchronization: when are messages sent?
- sometimes standardized as a "Request for Comments (RFC)" by the Internet Engineering Task Force (IETF)

Layers of the TCP/IP Reference Model

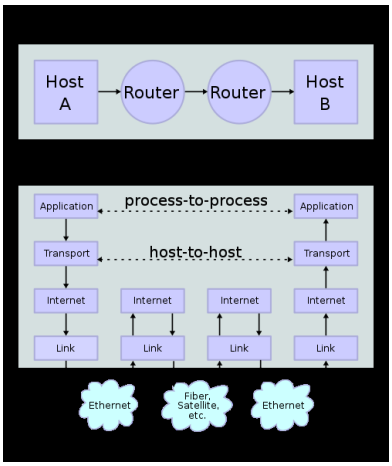

wikipedia.org/wiki/File:IP_stack_connections.svg

Layers of the TCP/IP Reference Model

- Application Layer
  - Handles requests from the user for data on the Internet.
- Transport Layer
  - Handles splitting messages into packets for delivery.
- Internet Layer
  - Handles the task of sending packets across one or more networks.
- Link Layer
  - Handles the physical transfer and reception of bits.

TCP/UDP vs IP

- **Internet Protocol (IP)**
  - delivers packets to IP address
  - best effort delivery
- **User Datagram Protocol (UDP)**
  - delivers packets to port at IP address
  - port identifies
  - still best effort delivery
- **Transmission Control Protocol (TCP)**
  - creates a reliable bi-directional stream (source address/port and destination address/port)
    - acknowledgements, resend, reassembly in correct order,
    - error detection
    - connection must be opened and closed (established with three-way handshake)
    - flow/congestion control

Examples of Protocols

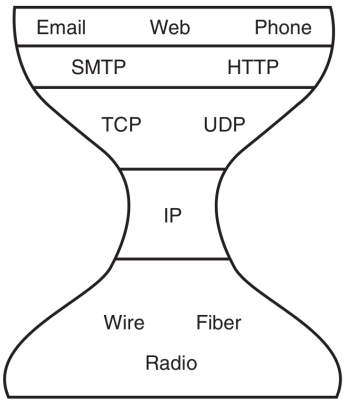

- Application
  - Hypertext Transfer Protocol (HTTP)
  - Simple Mail Transfer Protocol (SMTP)
  - Domain Name System (DNS)
    - XYZ.com → **w.x.y.z**
  - Secure Shell (SSH) Protocol
    - ssh unix.andrew.cmu.edu
  - Voice Over IP (Phone calls)
    - Session Initiation Protocol (SIP)
    - Real-time Transport Protocol (RTP)

Examples of Protocols

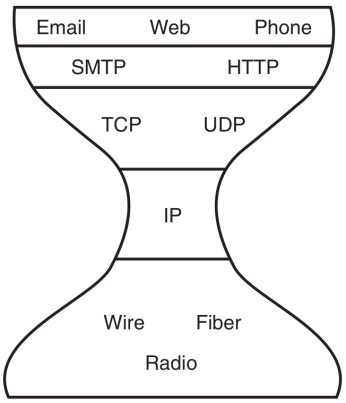

- Transport
  - Transmission Control Protocol (TCP)
  - User Datagram Protocol (UDP)
- Internet
  - IPv4
  - IPv6

Examples of Protocols

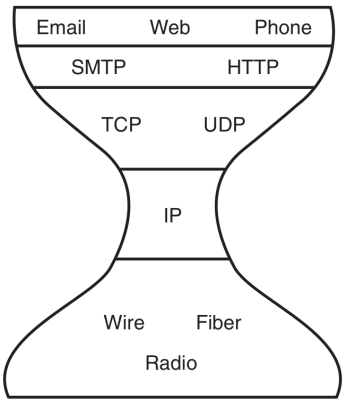

- Link
  - 1000BaseT — Gigabit Ethernet
  - Data Over Cable Service Interface Specification — Cable Modems
  - Long Term Evolution (LTE) – 4G cell phone
  - 802.11N (Wi-Fi) — Wireless Ethernet

Net Neutrality

- Should different kinds of packets be treated differently?
- The principle of net neutrality advocates no restrictions by ISPs or governments on consumers' access to networks that participate in the Internet.

Net neutrality means simply that all like Internet content must be treated alike and move at the same speed over the network. The owners of the Internet's wires cannot discriminate. This is the simple but brilliant "end-to-end" design of the Internet that has made it such a powerful force for economic and social good.
– Lawrence Lessig and Robert W. McChesney (Washington Post, June 8, 2006)