Background & Review of Communication Fundamentals

CS 161/194-1
Anthony D. Joseph
September 9, 2005

Outline

- Communications Network Taxonomy
  - Packet Networks
- The Internet
- Transport Layer: UDP/IP, TCP/IP
- Network Service Examples
- P2P applications

Many slides courtesy of EE 122 (Stoica/Katz)

What is a Communication Network? (End-system Centric View)

- Network offers one basic service: move information
  - Bird, fire, messenger, truck, telegraph, telephone, Internet …
- What distinguish different types of networks?
  - The services they provide, security, …
- What distinguish the services?
  - Latency
  - Bandwidth
  - Loss rate
  - Number of end systems
  - Service interface (how to invoke the service?)
  - Others
    - Reliability, unicast vs. multicast, real-time...
- What are the security issues?
  - Authentication, privacy, anonymity, integrity, …

What is a Communication Network? (Infrastructure Centric View)

- Communication medium: electron, photon
- Network components:
  - Links – carry bits from 1 place to 1 or more: fiber, copper, wireless,…
  - Interfaces – attach devices to links
  - Switches/routers – interconnect links: electronic/optic, crossbar/Banyan
  - Hosts – communication endpoints: PCs, PDAs, cell phones, toasters
- Protocols – rules governing communication between nodes
  - TCP/IP, ATM, MPLS, SONET, Ethernet, X.25
- Applications: Web browser, X Windows, FTP, ...
- Low-level security issues:
  - Authentication, privacy, integrity, …

Network Components (Examples)

Links

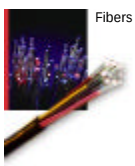

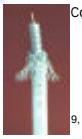

Interfaces

Ethernet card

Wireless card

Switches/routers

Large router

Taxonomy of Communication Networks

- Communication networks can be classified based on the way in which the nodes exchange information:

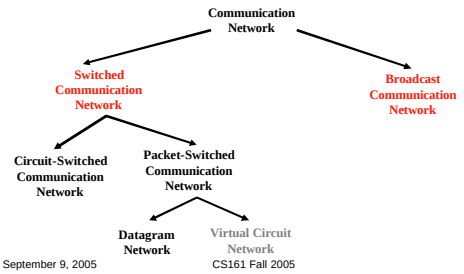

Broadcast vs. Switched Communication Networks

- Broadcast Communication Networks
  - Information transmitted by any node is received by every other node in the network
    - Examples: usually in LANs (non-switched Ethernet, WiFi)
- Switched Communication Networks
  - Information transmitted to a sub-set of designated nodes
    - Examples: WANs (Telephony Network, Internet), switched Ethernet
  - Problem: how to forward information to intended node(s)?
    - Done by special nodes (e.g., routers, switches) executing routing protocols
    - Can the routing process be subverted?

Taxonomy of Communication Networks

- Communication networks can be classified based on the way in which the nodes exchange information:

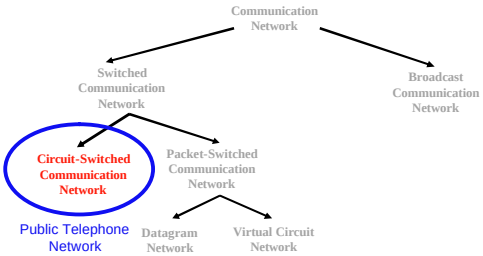

Taxonomy of Communication Networks

- Communication networks can be classified based on the way in which the nodes exchange information:

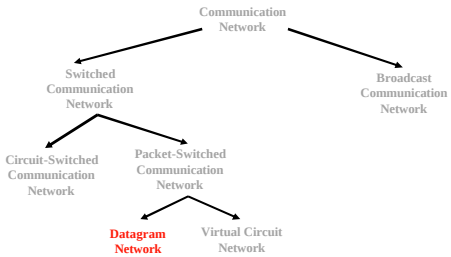

Datagram Packet Switching

- Each packet is independently switched
  - Each packet header contains destination address
- No resources are pre-allocated (reserved) in advance
- Example: IP networks

Timing of Datagram Packet Switching

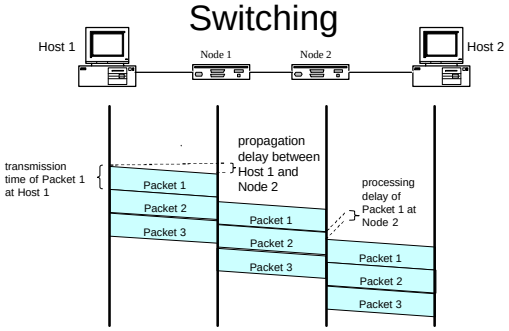

Datagram Packet Switching

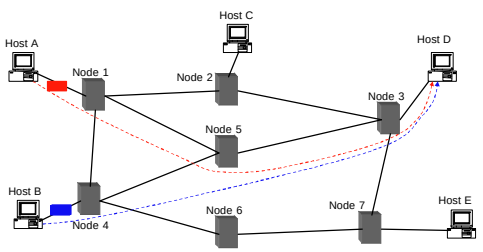

Outline

- Communications Network Taxonomy
  - Packet Networks
- The Internet
- Transport Layer: UDP/IP, TCP/IP
- Network Service Examples
- P2P applications

Many slides courtesy of EE 122 (Stoica/Katz)

The Internet

- Global scale, general purpose, heterogeneous-technologies, public, computer network
- Internet Protocol
  - Open standard: Internet Engineering Task Force (IETF) as standard body (http://www.ietf.org)
  - Technical basis for other types of networks
    - Intranet: enterprise IP network
- Developed by the research community

History of the Internet

- 68-70's: started as a research project, 56 kbps, initially 4 nodes (UCLA, UCSB, SRI, Utah) then < 100 computers
- 80-83: TCP/IP, DNS; ARPANET and MILNET split
- 85-86: NSF builds NSFNET as backbone, links 6 Supercomputer centers, 1.5 Mbps, 10,000 computers
- 87-90: link regional networks, NSI (NASA), ESNet (DOE), DARTnet, TWBNet (DARPA), 100,000 computers
- 90-92: NSFNET moves to 45 Mbps, 16 mid-level networks
- 94: NSF backbone dismantled, multiple private backbones; Introduction of Commercial Internet
- Today: backbones run at 10 Gbps, close to 320M computers in 150 countries

Network "Cloud"

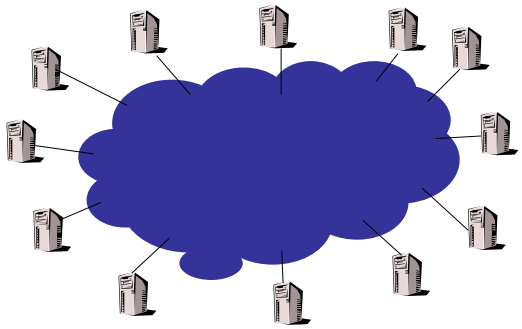

Regional Nets + Backbone

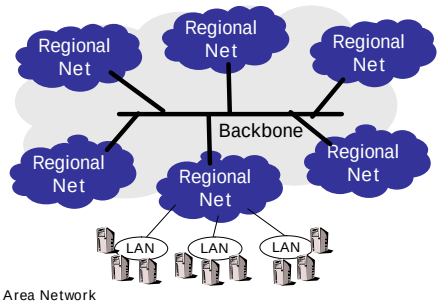

LAN: Local Area Network

Backbones + NAPs + ISPs

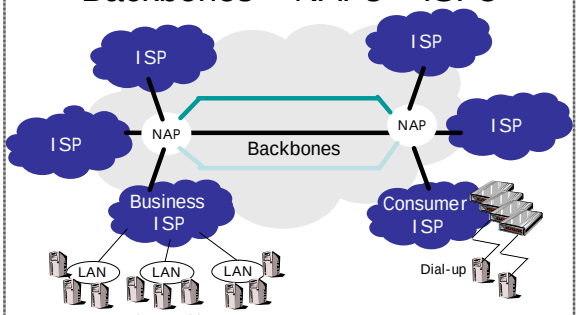

ISP: Internet Service Provider
NAP: Network Access Point

Core Networks + Access Networks

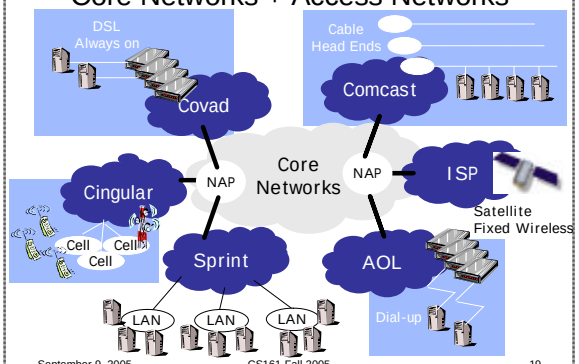

Computers Inside the Core

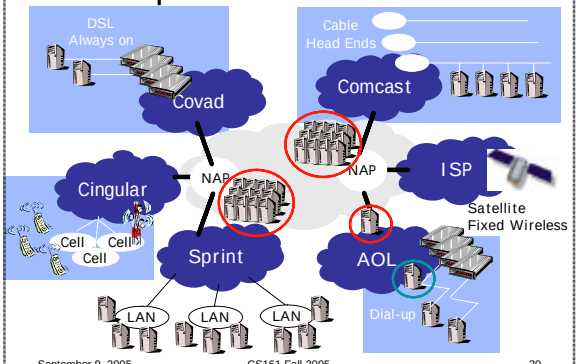

Internet Protocol Layers

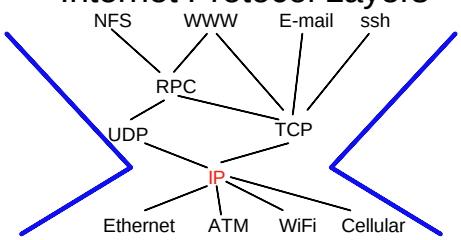

- Many different network technologies
- IP was invented to glue them together
  - n translations, not n x n!
  - Minimal requirements (datagram)
  - "IP over everything"

Services Provided by the Internet

- Shared access to computing resources
  - telnet (1970's), ssh (1990's)
- Shared access to data/files
  - FTP, NFS, AFS (1980's), CIFS (late 90's)
- Communication medium over which people interact
  - email (1980's), on-line chat rooms, instant messaging (1990's)
  - audio, video, Voice-over-IP (1990's, early 00's)
    - replacing telephone network?
- Medium for information dissemination
  - USENET (1980's)
  - WWW (1990's)
    - replacing newspaper, magazine?
  - Audio, video (late 90's, early 00's)
    - replacing radio, TV?
  - File sharing (late 90's, early 00's)

Addressing

- Every Internet host has an IP address
  - e.g., 67.114.133.15
- Packets include destination address
  - Network is responsible for routing packet to address
- Host-view:

IP-centric View

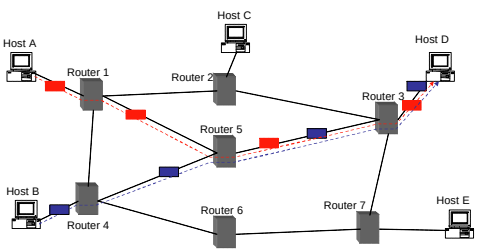

Routing

- Routers have "routing tables"
  - Tables mapping each destination with an outgoing link
  - Requires that routing table is highly compressible!
  - Implications for address assignment, mobility, etc.
- Routing decisions made packet-by-packet
  - Routers keep no connection state
- Question: Why have the network do routing?
  - Why not the hosts?
  - Compare delivery-by-hand to FedEx

Internet Service

- "Best-Effort" service
  - No guarantees about packet delivery
  - Hosts must cope with loss, delay, reordering, duplication
- Why not guarantee no loss and low delay?
- IP packets are addressed to a host
  - How to decide which application gets which packets?
- Need a transport layer!

Outline

- Communications Network Taxonomy
  - Packet Networks
- The Internet
- Transport Layer: UDP/IP, TCP/IP
- Network Service Examples
- P2P applications

Many slides courtesy of EE 122 (Stoica/Katz)

Transport Layer

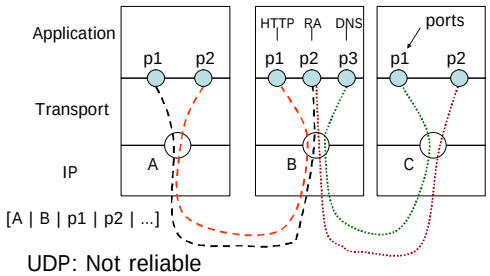

UDP: Not reliable

TCP: Ordered, reliable, well-paced

Ports

- Need to decide which application gets which packets
- Solution: map each socket to a *port*
- Client must know server's port
- Separate 16-bit port address space for UDP and TCP
  - (src_IP, src_port, dst_IP, dst_port) uniquely identifies TCP connection
- *Well known ports* (0-1023): everyone agrees which services run on these ports
  - e.g., ssh:22, http:80
  - on UNIX, must be root to gain access to these ports (why?)
- Ephemeral ports (most 1024-65535): given to clients
  - e.g. chat client gets one of these

Headers

- IP header ➔ used for IP routing, fragmentation, error detection…
- UDP header ➔ used for multiplexing/demultiplexing, error detection
- TCP header ➔ used for multiplexing/demultiplexing, flow and congestion control

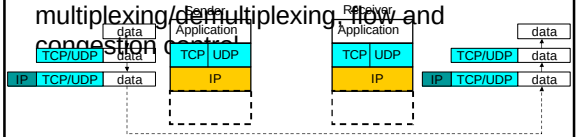

UDP

- User Datagram Protocol
- Minimalist transport protocol
- Same best-effort service model as IP
- Messages up to 64KB
- “Fire and Forget”
- Provides multiplexing/demultiplexing to IP
- Does not provide flow and congestion control
- Application examples: video/audio streaming, VoIP

UDP Service & Header

- Service:
  - Send datagram from (IPa, Port 1) to (IPb, Port 2)
  - Service is unreliable, but error detection possible
- Header:

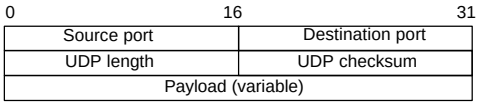

•UDP length is UDP packet length (including UDP header and payload, but not IP header)
•Optional UDP checksum is over UDP packet

TCP

- Transmission Control Protocol
- Reliable, in-order, and at most once delivery
- Messages can be of arbitrary length
- Provides multiplexing/demultiplexing to IP
- Provides congestion control and avoidance
- Application examples: file transfer, chat, P2P

TCP Service

1) Open connection
2) Reliable byte stream transfer from (IPa, TCP Port1) to (IPb, TCP Port2)
   - Indication if connection fails: Reset
3) Close connection

Timing Diagram

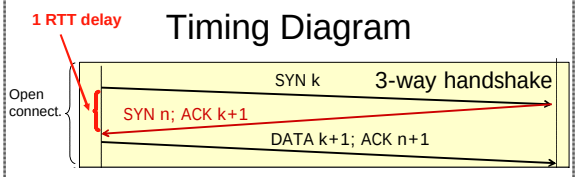

Outline

- Communications Network Taxonomy
  - Packet Networks
- The Internet
- Transport Layer: UDP/IP, TCP/IP
- Network Service Examples
- P2P applications

Many slides courtesy of EE 122 (Stoica/Katz)

Domain Name Service (DNS)

- Humans/applications use machine names
  - e.g., www.cs.berkeley.edu
- Network (IP) uses IP addresses
  - e.g., 67.114.112.23
- DNS translates between the two
  - An overlay service in its own right
  - Global distribution of name-to-IP address mappings—a kind of content distribution system as well
  - Unsung hero of the Internet

File Transfer (FTP, SCP, etc.)

Get file from soup.cs.berkeley.edu

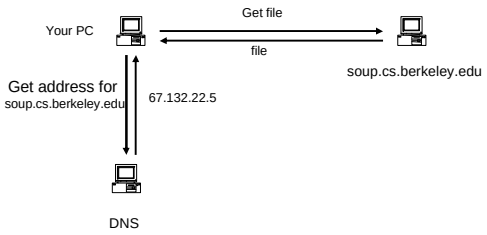

Question

- Why isn't the network in this picture?
  - Network just delivers (or not) packets to their destination
  - It plays no other role in application
- Obvious concept now, but novel at the time
  - Makes it both harder and easier for applications
  - Hosts more complex, applications less efficient
  - Long-term flexibility
- Security issues are hidden
  - Ex: Broadcast vs. switched

Email

Email message exchange is similar to previous example, except

- Exchange is between mail servers
- DNS gives name of mail server for domain

Web

Get www.icir.org/file.html

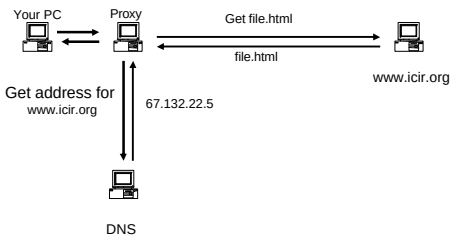

Outline

- Communications Network Taxonomy
  - Packet Networks
- The Internet
- Transport Layer: UDP/IP, TCP/IP
- Network Service Examples
- P2P applications

Many slides courtesy of EE 122 (Stoica/Katz)

Peer-to-Peer Networking (P2P)

- A killer application: Naptser
  - Free music over the Internet
- Key idea: share the storage *and* bandwidth of individual (home) users

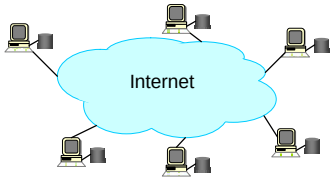

P2P Model

- Each user stores a subset of files
- Each user has access (can download) files from all users in the system

Challenges

- Find where a particular file is stored
  - Napster used central servers to store index
- Scaling to hundreds of millions
- Machines come and go

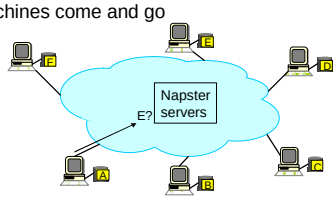

P2P Services in the Internet: Napster, Gnutella, BitTorent, …

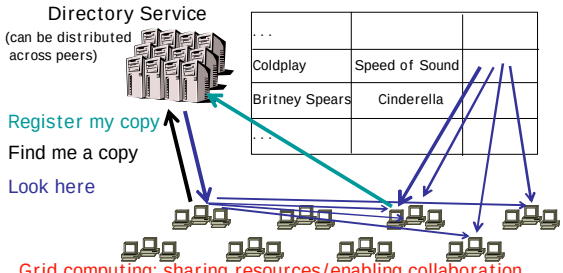

Grid computing: sharing resources/enabling collaboration