CS 268: Graduate Computer Networks – Spring 2003

- Instructor:
  - Ion Stoica (istoica@cs.berkeley.edu, 645 Soda Hall)
- Lecture time: TT, 12:30-2:00 pm
- Place: 310 Soda Hall
- Office hour: Tu, 2 - 3 pm

Overview

- Administrative trivia
- Overview and history of the Internet
- A Taxonomy of Communication Networks
- Router Architecture in Packet-Switching Networks

Administrative Trivia's

- Course Web page:
  - http://inst.eecs.berkeley.edu/~cs268/sp03
  - Check it periodically to get the latest information
- Deadline means deadline
  - Unless otherwise specified, it means 12:20pm on the date (10 minutes before lecture)
  - Special circumstances should be brought to my attention way ahead of deadlines
- Exams are close-book

Goals of this Course

- Understand how the Internet works
- Get familiar with current Internet research efforts
- Understand solutions in context
  - Goals
  - Assumptions
- Appreciate what is good research
  - Problem selection
  - Solution & research methodology
  - Presentation
- Apply what you learned in a class project

What Do You Need To Do?

- A research-oriented class project
- Two exams
- Paper reading

Research Project

- Investigate new ideas and solutions in a class research project
  - Define the problem
  - Execute the research
  - Work with your partner
  - Write up and present your research
- Ideally, best projects will become conference papers (e.g., SIGCOMM, INFOCOM, MOBICOM)

Research Project: Steps

- I'll distribute a list of projects
  - You can either choose one of these projects or come up with your own
- Pick your project, partner, and submit a one page proposal describing:
  - The problem you are solving
  - Your plan of attack with milestones and dates
  - Any special resources you may need
- A midterm presentation of your progress (five minutes)
- Final project presentation (ten minutes) + poster session
- Submit project papers

Paper Reviews

- Goal: synthesize main ideas and concepts in the papers
- Number: up to two papers per class
- Length: no more than half page per paper
- Content
  - Main points intended by the author
  - Points you particularly liked/disliked
  - Other comments (writing, conclusions…)
- Submission:
  - Submit each review via e-mail before 12:20 pm on lecture day
  - See class web page for details

Grading

Term project	50%
Final exam	15%
Midterm exam	15%
Class participation	10%
Paper reviews	10%

- This is a graduate networking class: more important is what you realize/learn than the grade

Enrollment Policy

- Graduate students get highest priority
- Among other students, priority given to those who
  - Have backgrounds in networking, operating systems
  - Have relatively light course load
- Procedure of enrollment for undergraduate students
  - Be officially on the waiting list
  - Send me an email with URL that has pointers to
    - Your resume
    - A short statement of relevant courses (textbook, university, grade) and experiences
    - Other courses you are taking this semester

Send the Following Information

- Please send me (istoica@cs.berkeley.edu) an e-mail with the subject "cs**268 registration**" and the following information:
  - Last and first name
  - Student ID
  - Your department
  - Preferred email address
  - URL of your home page
- Please indicate explicitly if we can add you to the on-line web page that lists each student enrolled in the class (only your name and URL will be made publicly available here).

Overview

- Administrative trivia
- Overview and history of the Internet
- A Taxonomy of Communication Networks

What is a Communication Network? (End system view)

- Network offers a service: move information
  - Bird, fire, messenger, truck, telegraph, telephone, Internet …
  - Another example, transportation service: move objects
    - horse, train, truck, airplane ...
- What distinguish different types of networks?
  - The services they provide
- What distinguish the services?
  - Latency
  - Bandwidth
  - Loss rate
  - Number of end systems
  - Service interface (how to invoke?)
  - Other details
    - Reliability, unicast vs. multicast, real-time, message vs. byte ...

What is a Communication Network? (Infrastructure Centric View)

- Electrons and photons as communication medium
- Links: fiber, copper, satellite, …
- Switches: electronic/optical, crossbar/Banyan
- Protocols: TCP/IP, ATM, MPLS, SONET, Ethernet, PPP, X.25, FrameRelay, AppleTalk, IPX, SNA
- Functionalities: routing, error control, congestion control, Quality of Service (QoS)
- Applications: FTP, WEB, X windows, ...

Types of Networks

- Geographical distance
  - Local Area Networks (LAN): Ethernet, Token ring, FDDI
  - Metropolitan Area Networks (MAN): DQDB, SMDS
  - Wide Area Networks (WAN): X.25, ATM, frame relay
  - Caveat: LAN, MAN, WAN may mean different things
    - service, network technology, networks
- Information type
  - Data networks vs. telecommunication networks
- Application type
  - Special purpose networks: airline reservation network, banking network, credit card network, telephony
  - General purpose network: Internet

Types of Networks

- Right to use
  - private: enterprise networks
  - public: telephony network, Internet
- Ownership of protocols
  - proprietary: SNA
  - open: IP
- Technologies
  - terrestrial vs. satellite
  - wired vs. wireless
- Protocols
  - IP, AppleTalk, SNA

The Internet

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

History of the Internet

- 70's: started as a research project, 56 kbps, < 100 computers
- 80-83: 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
- Today: backbones run at 10 Gbps, 10s millions computers in 150 countries

Time Line of the Internet

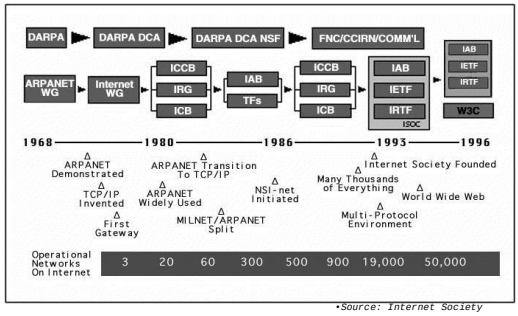

•Source: Internet Society

Growth of the Internet

- **Number of Hosts on the Internet:**

Aug. 1981	213
Oct. 1984	1,024
Dec. 1987	28,174
Oct. 1990	313,000
Oct. 1993	2,056,000
Apr. 1995	5,706,000
Jul. 1997	19,540,000
Jul. 1999	56,218,000
Jul. 2001	125,888,197
Jul. 2002	162,128,493

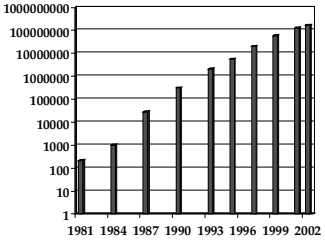

Recent Growth (1991-2002)

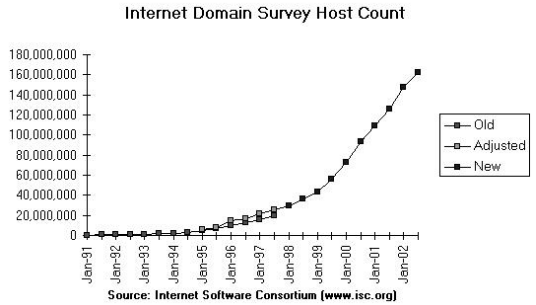

Who is Who in the Internet ?

- **Internet Engineering Task Force (IETF):** The IETF is the protocol engineering and development arm of the Internet. Subdivided into many working groups, which specify Request For Comments or RFCs.
- **IRTF (Internet Research Task Force):** The Internet Research Task Force is a composed of a number of focused, long-term and small Research Groups.
- **Internet Architecture Board (IAB)**: The IAB is responsible for defining the overall architecture of the Internet, providing guidance and broad direction to the IETF.
- **The Internet Engineering Steering Group (IESG)**: The IESG is responsible for technical management of IETF activities and the Internet standards process. Standards. Composed of the Area Directors of the IETF working groups.

Internet Standardization Process

- All standards of the Internet are published as RFC (Request for Comments). But not all RFCs are Internet Standards !
  - available: http://www.ietf.org
- A typical (but not only) way of standardization is:
  - Internet Drafts
  - RFC
  - Proposed Standard
  - Draft Standard (requires 2 working implementation)
  - Internet Standard (declared by IAB)
- David Clark, MIT, 1992: "We reject: kings, presidents, and voting. We believe in: rough consensus and running code."

Services Provided by the Internet

- Shared access to computing resources
  - Telnet (1970's)
- Shared access to data/files
  - FTP, NFS, AFS (1980's)
- Communication medium over which people interact
  - Email (1980's), on-line chat rooms, instant messaging (1990's)
  - Audio, video (1990's)
    - Replacing telephone network?
- A medium for information dissemination
  - USENET (1980's)
  - WWW (1990's)
    - Replacing newspaper, magazine?
  - Audio, video (2000's)
    - Replacing radio, CD, TV?

Internet Physical Infrastructure

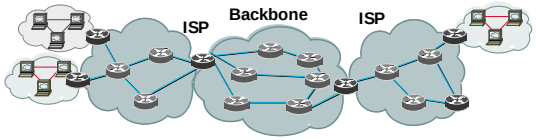

- Residential Access
  - Modem
  - DSL
  - Cable modem
  - Satellite
- Enterprise/ISP access, Backbone transmission
  - T1/T3, DS-1 DS-3
  - OC-3, OC-12
  - ATM vs. SONET, vs. WDM
- Campus network
  - Ethernet, ATM
- Internet Service Providers
  - access, regional, backbone
  - Point of Presence (POP)
  - Network Access Point (NAP)

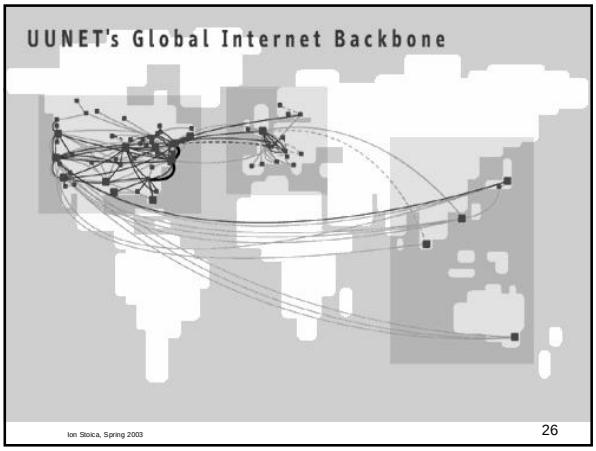

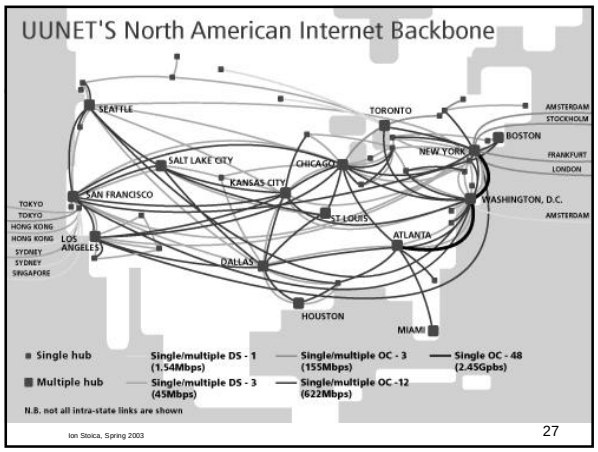

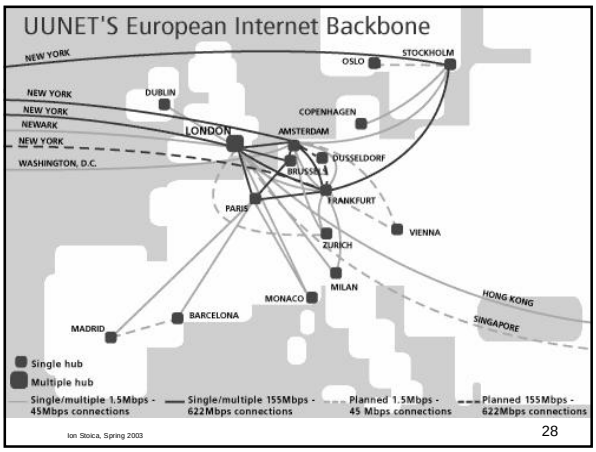

Overview

- Administrative trivia
- Overview and history of the Internet
- A Taxonomy of Communication Networks

A Taxonomy of Communication Networks

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

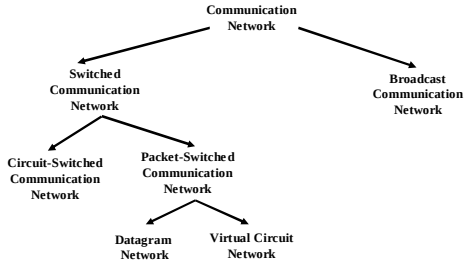

Broadcast vs. Switched Communication Networks

- Broadcast communication networks
  - information transmitted by any node is received by every other node in the network
    - examples: usually in LANs (Ethernet, Wavelan)
  - Problem: coordinate the access of all nodes to the shared communication medium (Multiple Access Problem)
- Switched communication networks
  - information is transmitted to a sub-set of designated nodes
    - examples: WANs (Telephony Network, Internet)
  - Problem: how to forward information to intended node(s)
    - this is done by special nodes (e.g., routers, switches) running routing protocols

A Taxonomy of Communication Networks

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

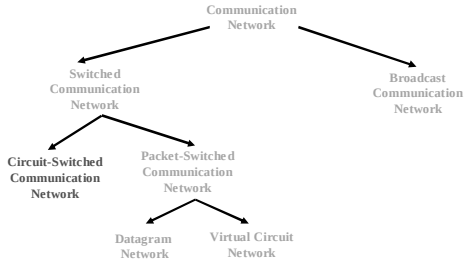

Circuit Switching

- Three phases
  1. circuit establishment
  2. data transfer
  3. circuit termination
- If circuit not available: “Busy signal”
- Examples
  - Telephone networks
  - ISDN (Integrated Services Digital Networks)

Timing in Circuit Switching

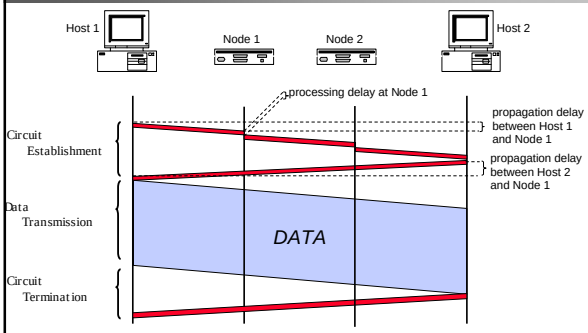

Circuit Switching

- A node (switch) in a circuit switching network

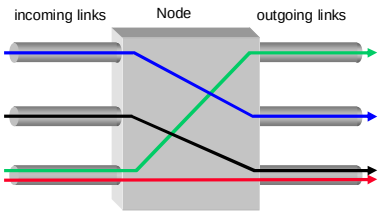

Circuit Switching: Multiplexing/Demultiplexing

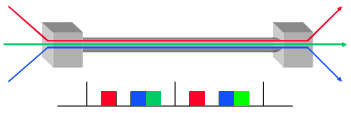

- Time divided in frames and frames divided in slots
- Relative slot position inside a frame determines which conversation the data belongs to
- Needs synchronization between sender and receiver
- In case of non-permanent conversations
  - Needs to dynamic bind a slot to a conservation
  - How to do this?

A Taxonomy of Communication Networks

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

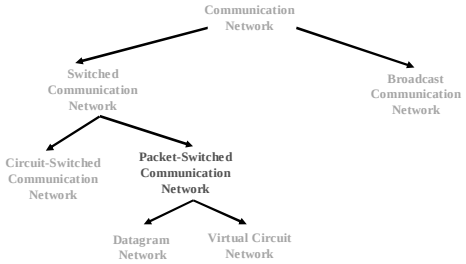

Packet Switching

- Data are sent as formatted bit-sequences, so-called packets.
- Packets have the following structure:

  - Header and Trailer carry control information (e.g., destination address, check sum)
- Each packet is passed through the network from node to node along some path (**Routing**)
- At each node the entire packet is received, stored briefly, and then forwarded to the next node (**Store-and-Forward Networks**)
- Typically no capacity is allocated for packets

Packet Switching

- A node in a packet switching network

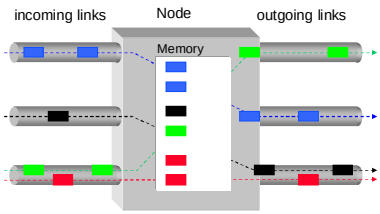

Packet Switching: Multiplexing/Demultiplexing

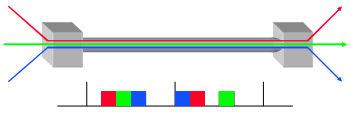

- Data from any conversation can be transmitted at any given time
- How to tell them apart?
  - use meta-data (header) to describe data

A Taxonomy of Communication Networks

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

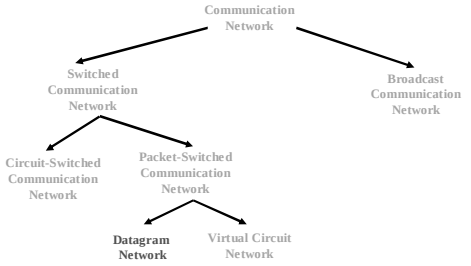

Datagram Packet Switching

- 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

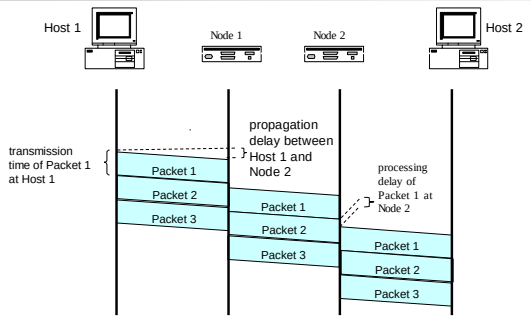

Datagram Packet Switching

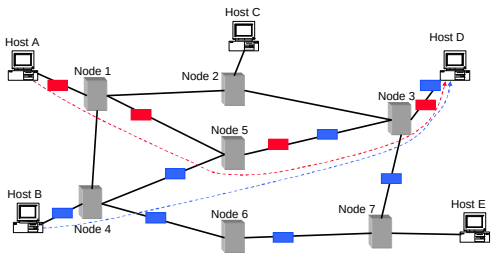

A Taxonomy of Communication Networks

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

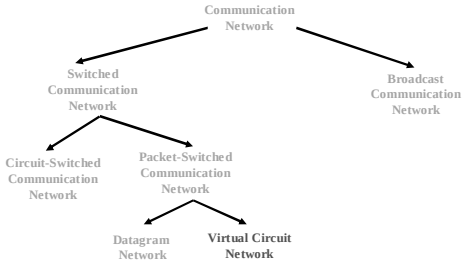

Virtual-Circuit Packet Switching

- Hybrid of circuit switching and packet switching
  - data is transmitted as packets
  - all packets from one packet stream are sent along a pre-established path (=virtual circuit)
- Guarantees in-sequence delivery of packets
- **However**: Packets from different virtual circuits may be interleaved
- Example: ATM networks

Virtual-Circuit Packet Switching

- Communication with virtual circuits takes place in three phases
  1. VC establishment
  2. data transfer
  3. VC disconnect
- Note: packet headers don't need to contain the full destination address of the packet

Timing of Datagram Packet Switching

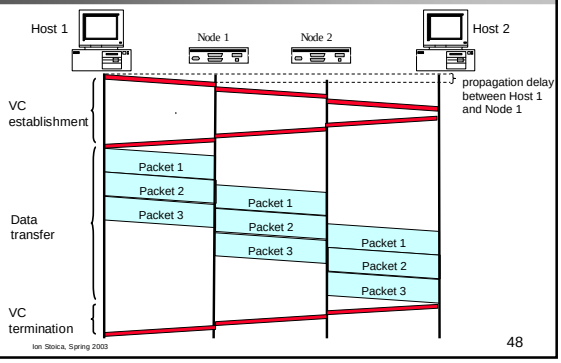

Datagram Packet Switching

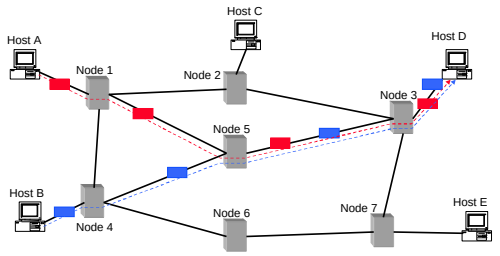

Packet-Switching vs. Circuit-Switching

- Most important advantage of packet-switching over circuit switching: Ability to exploit statistical multiplexing:
  - efficient bandwidth usage; ratio between peek and average rate is 3:1 for audio, and 15:1 for data traffic
- However, packet-switching needs to deal with congestion:
  - more complex routers
  - harder to provide good network services (e.g., delay and bandwidth guarantees)
- In practice they are combined:
  - IP over SONET, IP over Frame Relay

Summary

- Course administrative trivia
- Internet history and trivia
- Rest of the course a lot more technical and (hopefully) exciting