CS162
Operating Systems and Systems Programming
Lecture 8

Introduction to Networking, Packet Switching

September 26, 2011
Anthony D. Joseph and Ion Stoica
http://inst.eecs.berkeley.edu/~cs162

Review: Example: What's in a Search Query?

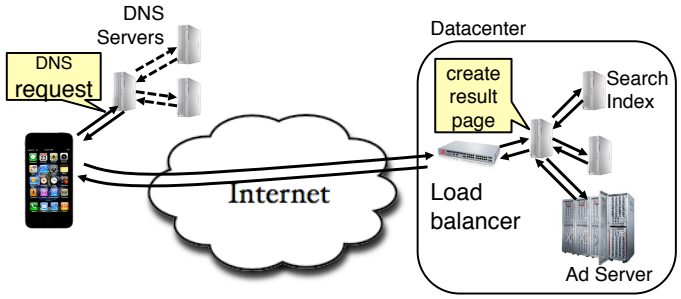

- Complex interaction of multiple components in multiple administrative domains

Goals for Today

- Communication network taxonomy
  - Circuit switching
  - Packet switching
- Statistical multiplexing

Note: Some slides and/or pictures in the following are adapted from notes by Vern Paxson, and Randy Katz.

What Global (non-digital) Communication Network Do You Use Every Day?

Roughly speaking, how does it work?

Postal System

What's Another Such Network That You Use Every Day?

Examples of Other Information Networks

Sending messages by birds

Fire signals

Other Network Examples

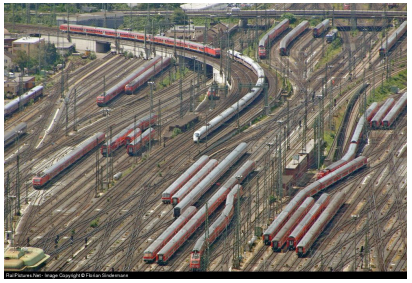

Taxonomy of Communication Networks

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

Communication Network

Taxonomy of Communication Networks

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

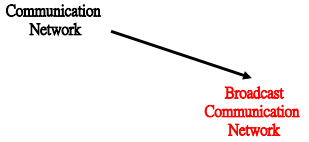

Broadcast Communication Networks

- Information transmitted by any node is received by every other node in the network
  - Examples?
  - Usually in LANs (*Local Area* Networks)
    - » E.g., Ethernet (classical), WiFi
    - » E.g., lecture!
- What problems does this raise?
- Problem #1: limited range
- Problem #2: privacy of communication
- Problem #3: coordinating access to the shared communication medium (*Multiple Access Problem*)

Taxonomy of Communication Networks

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

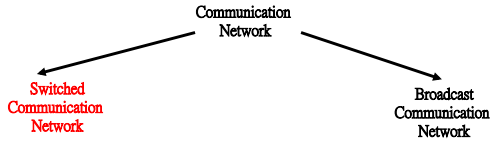

Switched Communication Networks

- Information transmitted along a path of intermediary *nodes* ("routers" or "switches")
- Basic issue: how the routers figure out the next hop along the path

Taxonomy of Communication Networks

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

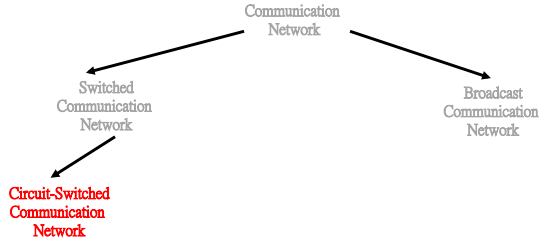

Circuit Switching (e.g., Phone Network)

- Establish: source creates circuit to destination
  - Nodes along the path store connection info
  - Nodes generally reserve resources for the connection
  - If circuit not available: "Busy signal"
- Transfer: source sends data over the circuit
  - No destination address, since nodes know path
- Teardown: source tears down circuit when done

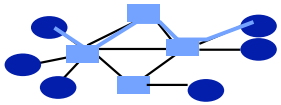

Circuit Switching

- Node (switch) in a circuit switching network

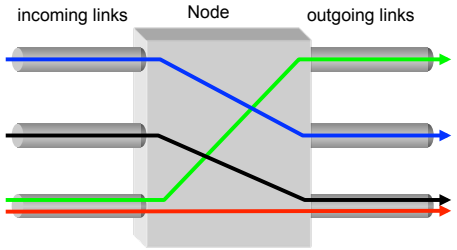

How does the node connect the incoming link to the outgoing?

Telephone Network

- Alexander Graham Bell
  - 1876: Demonstrates the telephone at US Centenary Exhibition in Philadelphia

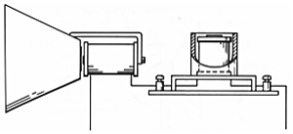

Circuit Switching With Human Operator

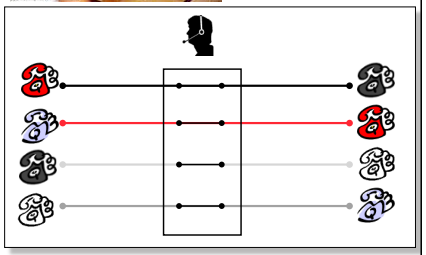

Telephone Network

- Almon Brown Strowger (1839 - 1902)
  - 1889: Invents the "girl-less, cuss-less" telephone system -- the *mechanical switching system*

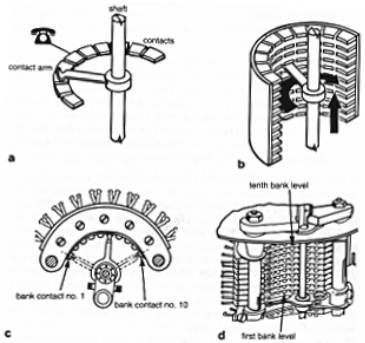

Timing in Circuit Switching

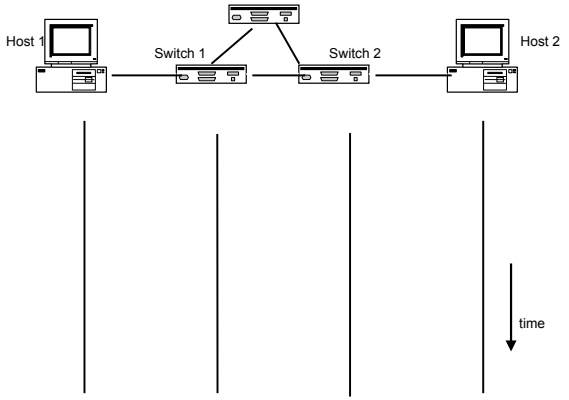

Timing in Circuit Switching

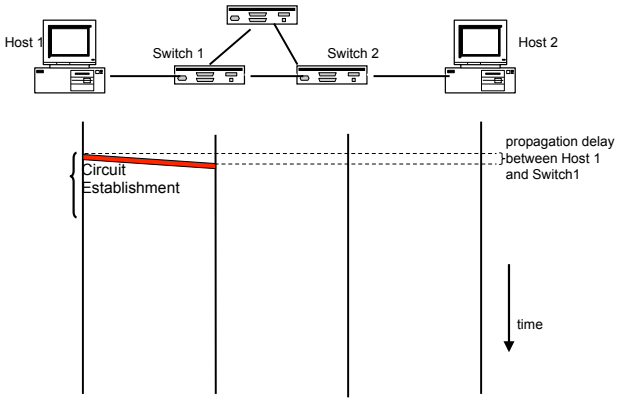

Timing in Circuit Switching

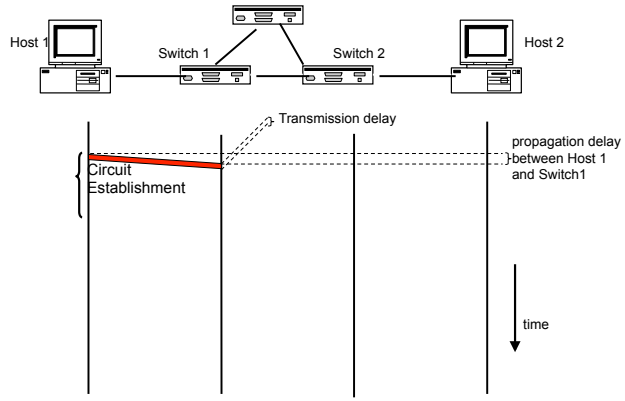

Timing in Circuit Switching

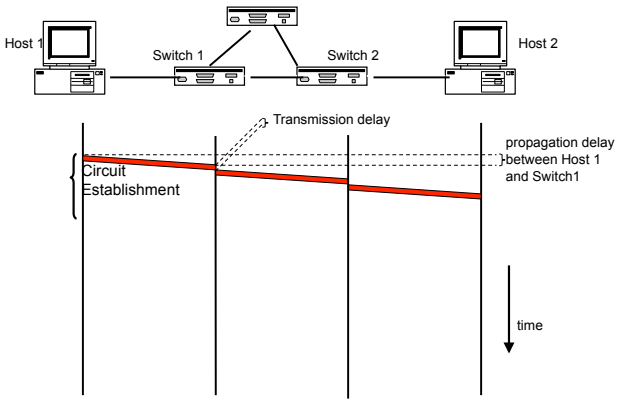

Timing in Circuit Switching

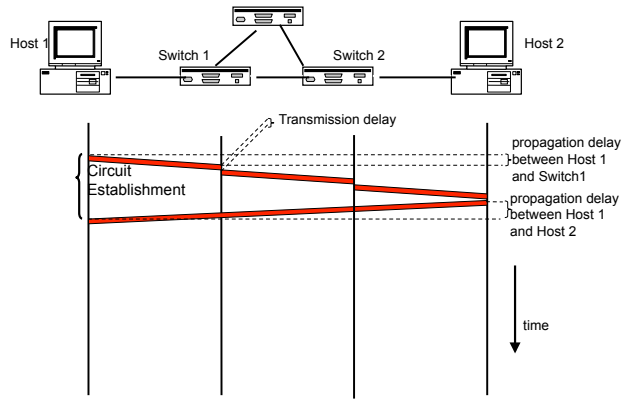

Timing in Circuit Switching

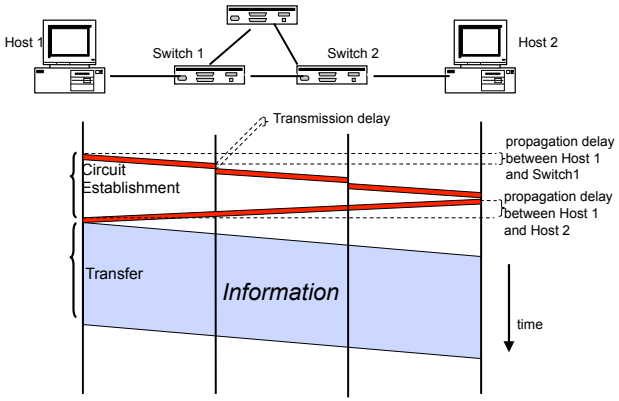

Timing in Circuit Switching

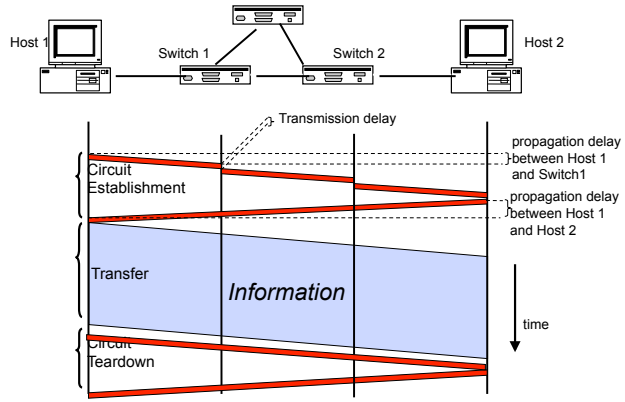

Circuit Switching

- Node (switch) in a circuit switching network

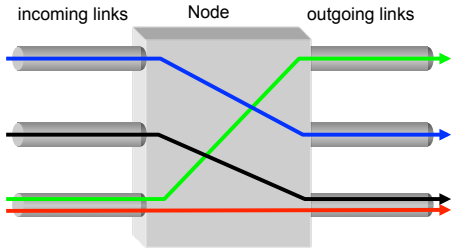

How do the black and red circuits share the outgoing link?

Circuit Switching: *Multiplexing* a Link

- Time-division
  - Each circuit allocated certain time slots
- Frequency-division
  - Each circuit allocated certain frequencies

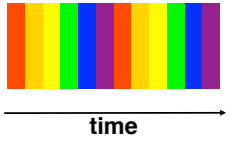

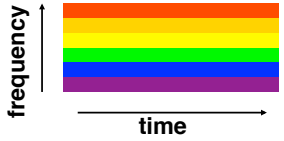

Time-Division Multiplexing/ Demultiplexing

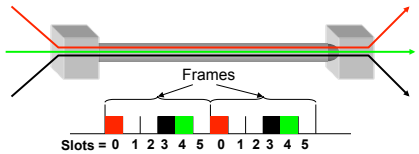

- Time divided into frames; frames into slots
- Relative slot position inside a frame determines to which circuit data belongs
  - E.g., slot 0 belongs to red circuit
- Requires synchronization between sender and receiver—surprisingly non-trivial!
- In case of non-permanent circuits
  - Need to dynamically bind a slot to a circuit
  - How to do this?
- If sender does not send data **the circuit's capacity is lost!**

5min Break

Taxonomy of Communication Networks

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

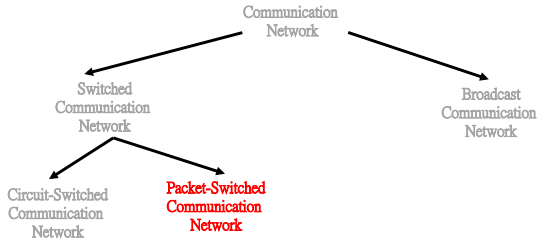

Packet Switching

- Data sent as chunks of formatted bit-sequences (**Packets**)
- Packets have following structure:

  - » Header and Trailer carry control information (e.g., destination address, checksum)
- Each packet traverses the network from node to node along some path (**routing**) based on header info
- Usually, once a node receives the entire packet, it stores it (hopefully briefly) and then forwards it to the next node (**Store-and-Forward Networks**)

Packet Switching

- Node in a packet switching network

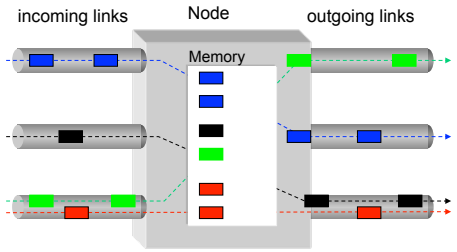

Packet Switching: Multiplexing/ Demultiplexing

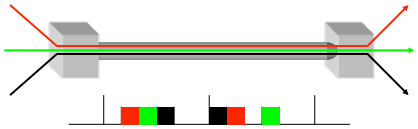

- Data from any flow can be transmitted at any given time
  - Single flow can use *the entire link capacity* if it is alone
- How to tell them apart?
  - Use meta-data (header) to describe data
- Note: for packet switching we use flow (instead of circuit) to denote packets sent by a sender to a receiver

Taxonomy of Communication Networks

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

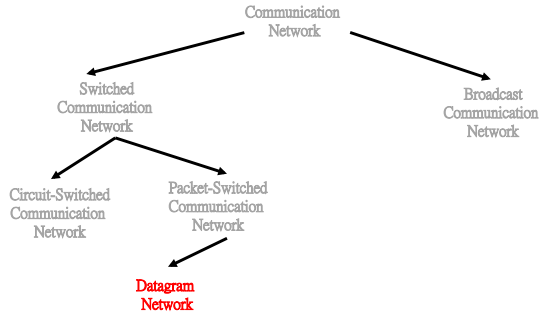

Datagram Packet Switching

- Each packet is independently switched
  - Each packet header contains full destination address
- No resources are pre-allocated (reserved) in advance
- Leverages "statistical multiplexing" (or *stat-muxing*)
  - Essentially: "chances are good that packets from different flows won't all arrive at the same time, so we can get by without enough capacity for all of them at their peak transmission rate"
  - *Assuming independence of traffic sources*, can compute probability that there is enough capacity
- Example: IP networks; postal system

Statistical Multiplexing

- Use trivial time-driven simulation to illustrate statistical multiplexing
- Probabilistically generate the bandwidth of a flow at each time unit, e.g.,
  - With probability 0.2, bandwidth is 6
  - With probability 0.8, bandwidth is 1
- Average bandwidth, avg=0.2*6 + 0.8*1 = 2
- peak/avg = 6/2 = 3

One Flow

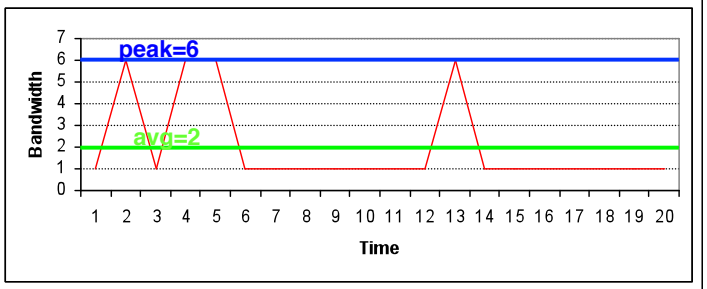

peak / avg = 3

Two Flows

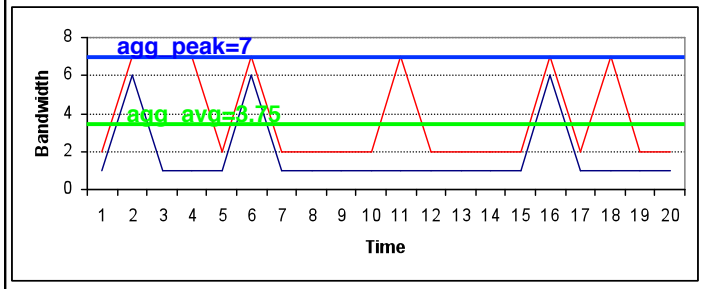

agg_peak / agg_avg = 7/3.75 = 1.86
(agg_avg = average of aggregate bandwidth)
(agg_peak=maximum value of aggregate bandwidth)

50 Flows

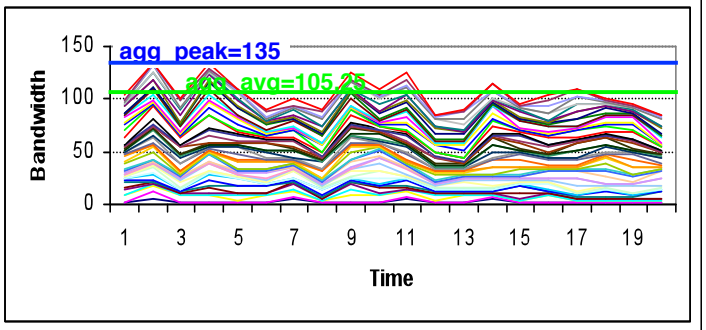

agg_peak / agg_avg = 7/3.75 = 135/105.25 = 1.28

Statistical Multiplexing (cont'd)

- As number of flows increases, agg_peak/agg_avg decreases
  - For 1000 flows, peak/avg = 2125/2009=1.057
- Q: What does this mean?
- A: Multiplexing a large enough number of flows "eliminates" burstiness
  - *Use average bandwidth to provision capacity, instead of peak bandwidth*
  - E.g., For 1000 flows
    - » Average of aggregate bandwidth = 2,000
    - » Sum of bandwidth peaks = 6,000

Timing of Datagram Packet Switching

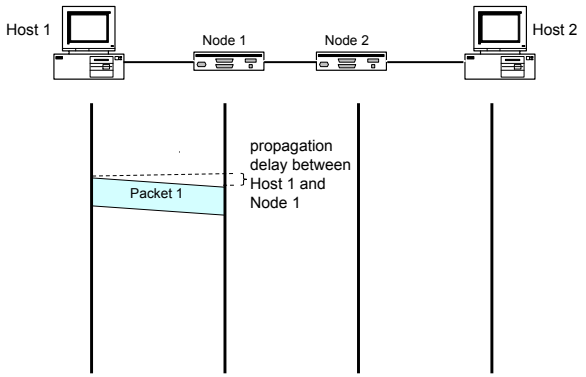

Timing of Datagram Packet Switching

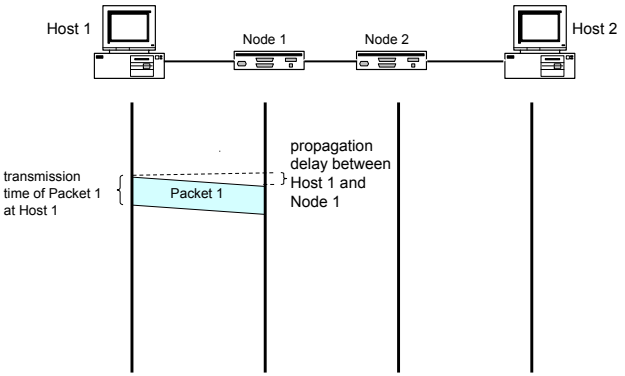

Timing of Datagram Packet Switching

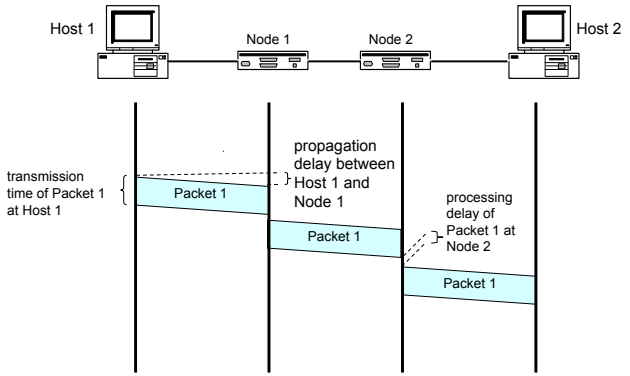

Timing of Datagram Packet Switching

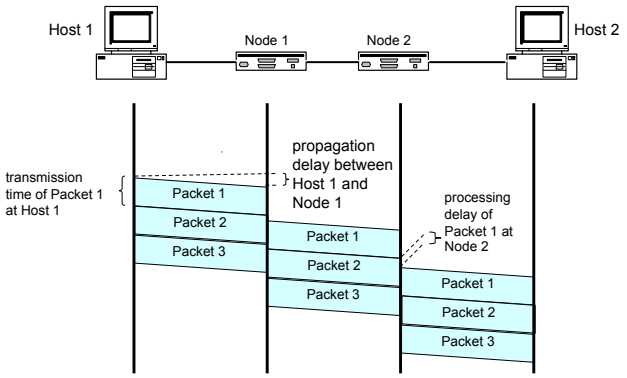

Datagram Packet Switching

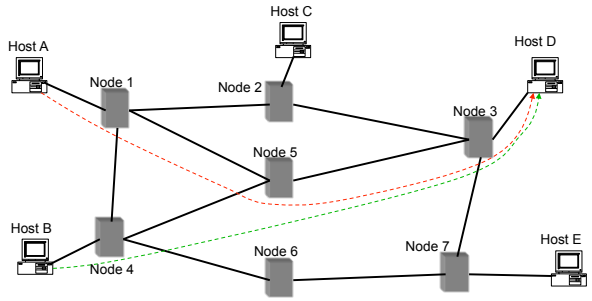

Datagram Packet Switching

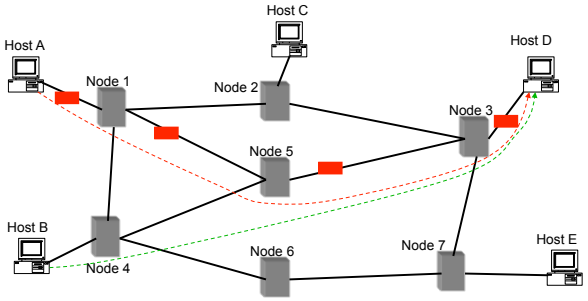

Datagram Packet Switching

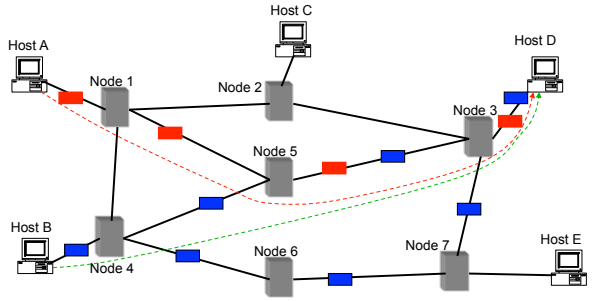

Taxonomy of Communication Networks

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

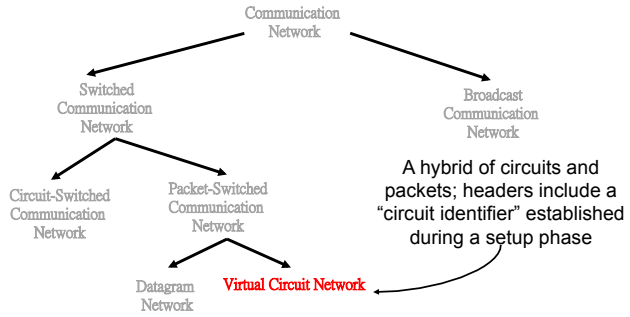

Advantages of Circuit Switching

- **Guaranteed bandwidth**
  - Predictable communication performance
- **Simple abstraction**
  - Reliable communication channel between hosts
  - No worries about lost or out-of-order packets
- **Simple forwarding**
  - Forwarding based on time slot or frequency
  - No need to inspect a packet header
- **Low per-packet overhead**
  - Forwarding based on time slot or frequency
  - No IP (and TCP/UDP) header on each packet

Disadvantages of Circuit Switching

- **Wasted bandwidth**
  - Bursty traffic leads to idle connection during silent period
  - Unable to achieve gains from "statistical multiplexing"
- **Blocked connections**
  - Connection refused when resources are not sufficient
  - Unable to offer "okay" service to everybody
- **Connection set-up delay**
  - No communication until the connection is set up
  - Unable to avoid extra latency for small data transfers
- **Network state**
  - Network nodes must store per-connection information
  - Unable to avoid per-connection storage and state
  - This makes failures more disruptive!

Packet-Switching vs. Circuit-Switching

- Critical advantage of packet-switching over circuit switching: Exploitation of statistical multiplexing
- Another: since routers don't know about individual flows, when a router or link fails, it's: Easy to fail over to a different path
- A third: easier for different parties to link their networks together because they're **not** promising to reserve resources for one another
- However, packet-switching must handle congestion:
  - More complex routers
  - Harder to provide good network services (e.g., delay and bandwidth guarantees)

Summary

- Packet switching:
  - Store and forward: a packet is stored at routers before being forwarded
  - Each packet can take a different path
  - No resource reservation: leverage statistical multiplexing
- Circuit switching:
  - Set-up path in advance
  - Reserve resources for each connection
- Statistically multiplexing:
  - Peak aggregate bandwidth much lower than sum of the peak of individual connections