High-Speed Networking: Trends and Issues

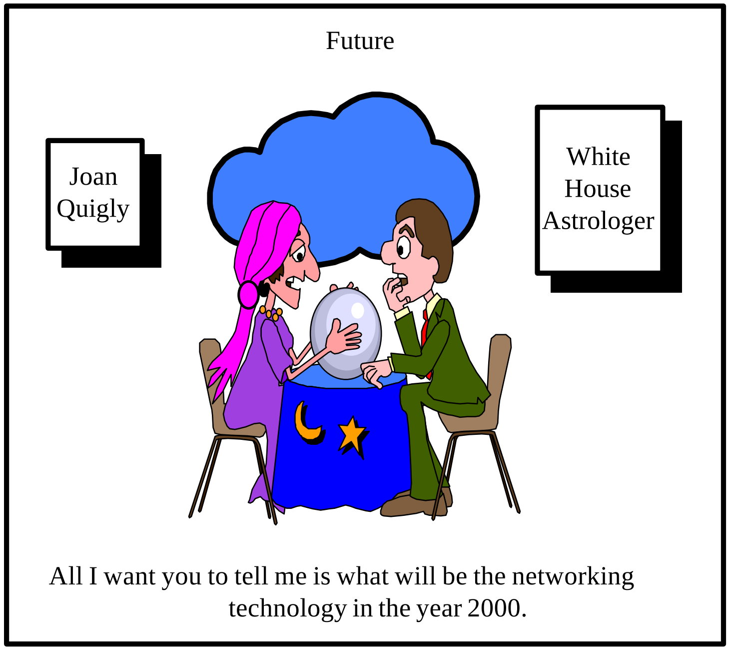

Raj Jain

Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
http://www.cse.wustl.edu/~jain/

Future

Joan Quigly

White House Astrologer

All I want you to tell me is what will be the networking technology in the year 2000.

- Industry trends
- High-speed network design
- A Simple rule of thumb
- Trends in traffic
- Trends in network topology

Trend: Telecommunication and Networking

- From computerization of telephone traffic switching to telephonization of computer traffic switching.

Trend: Networking is Critical

- Communication more critical than computing
  ⇒ Bus performance vs ALU speed
  ⇒ I/O performance vs SPECMarks
- User Location:
  - 1960: Computer room	1970: Terminal room
  - 1980: Desktop	1990: Mobile
- System Extent:
  - 1980: 1 Node within 10 m
  - 1990: 100 nodes within 10 km

- Last 10 years: Individual computing
  Next 10 years: Cooperative computing
- Past: Corporate networks
  Future:
  - Intercorporate networks
  - National Info Infrastructures
  - International Info Infrastructures

Life Cycles of Technologies

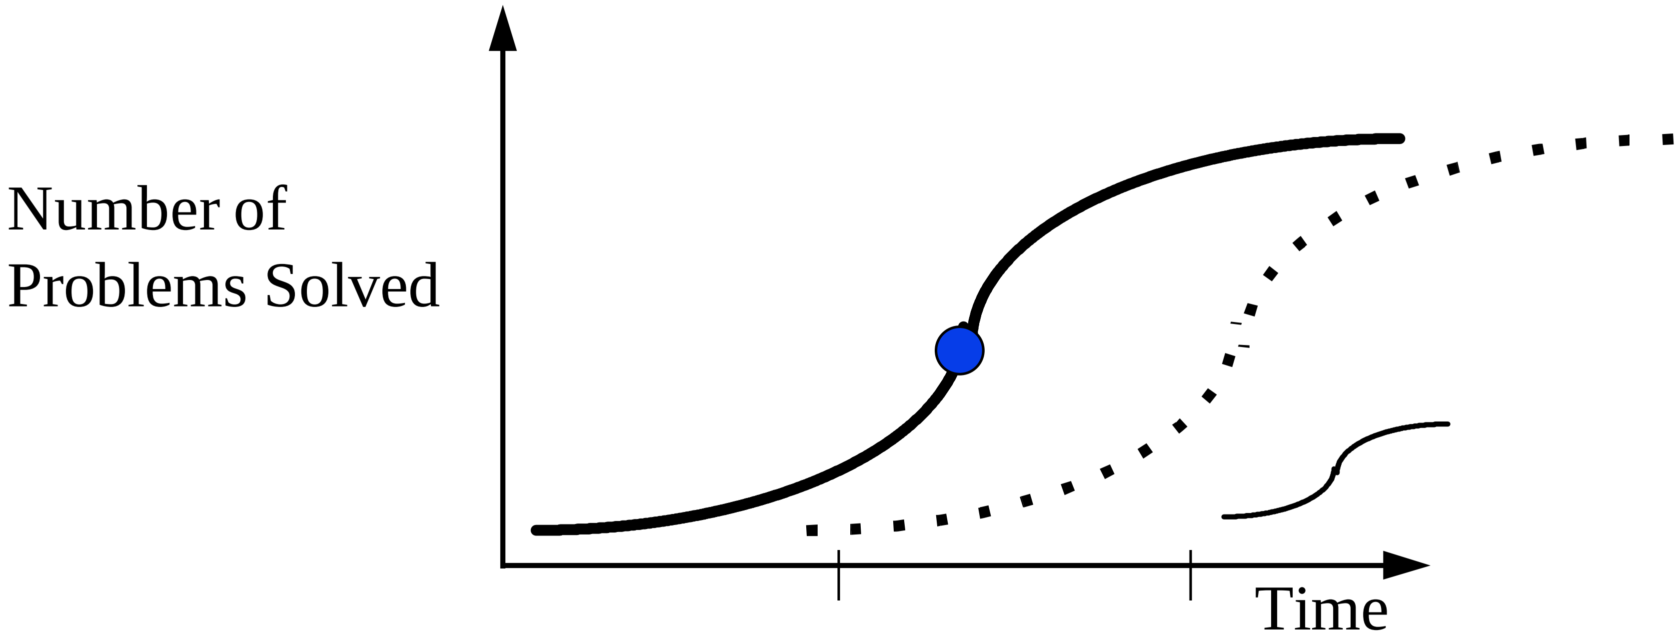

Life Cycles of Networking Technology

Number of Hosts
Bytes per Hosts
Number of Networks
MIPS
Memory Size
Storage

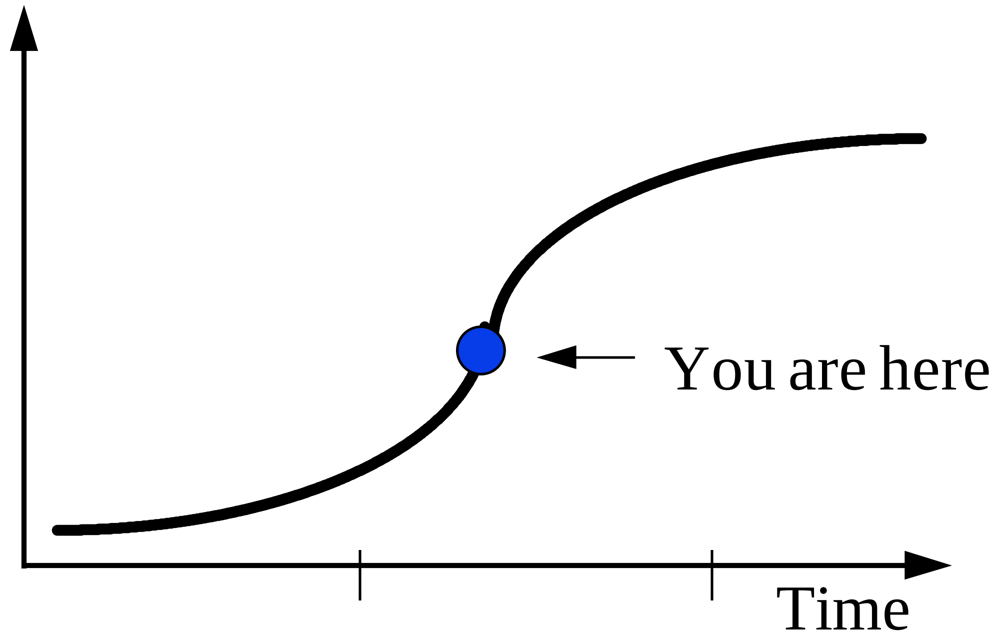

Trend: Exponential Growth

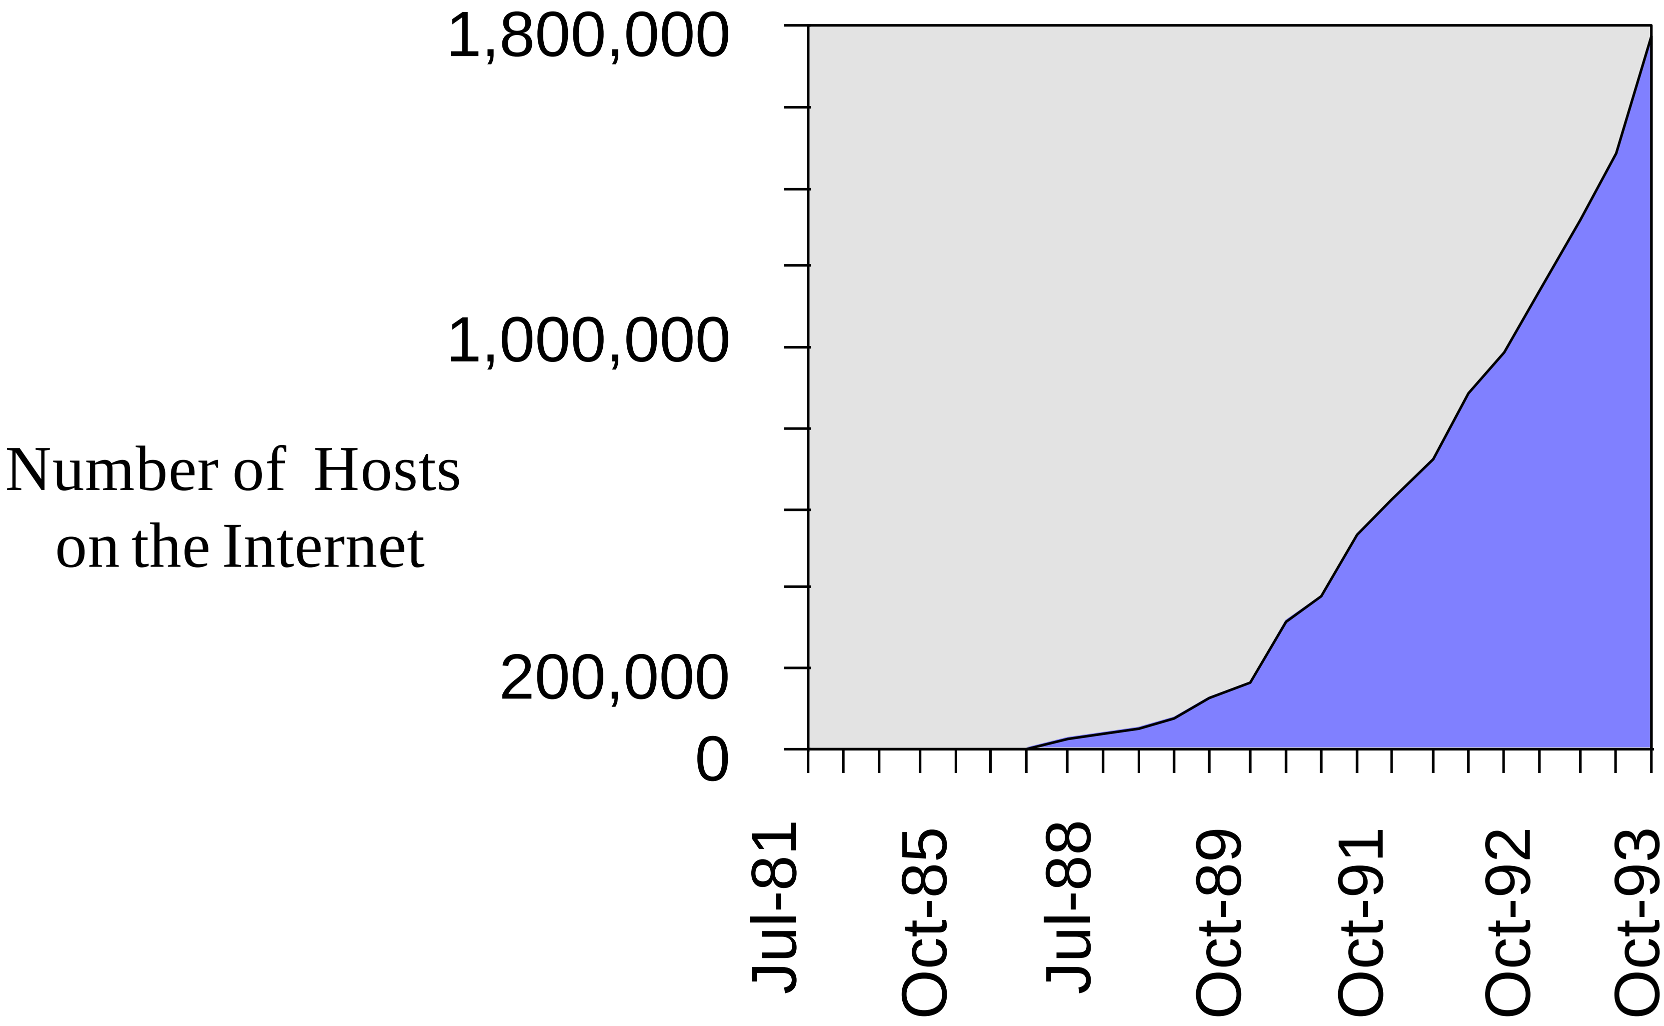

Trend: Standardization

- Religion must be forgotten
  ⇒ Improve on someone else's ideas as naturally as yours
- Can't succeed alone
  ⇒ Innovation + Technology partnerships
- To impact: Participate in standardization
  Publication is too late and insufficient
- Vertical vs horizontal specialization
  ⇒ Switch, router, host, applications

Routers
Switches
Hubs

Challenge: Economy of Scale

- Technology is far ahead of the applications.
  Invention is becoming the mother of necessity.
  We have high speed fibers, but not enough video traffic.
- Low-cost is the primary motivator. Not necessity.
  ⇒ Buyer's market (Like $99 airline tickets to Bahamas.)
  Why? vs Why not?
- Ten 100-MIPS computer are cheaper than one 1000-MIPS computer ⇒ Parallel computing, not supercomputing
- Ethernet was and still is cheaper than 10 one-Mbps links.
- No FDDI if it is 10 times as expensive as Ethernet.
  10/100 Ethernet adapters = $50 over 10 Mbps
- Q: Given ATM or 100 Mbps Ethernet at the same cost, which network will you buy?
  A: Ethernet. Proven Technology.

Challenge: Performance

Application Designers	Video Coding, FTP
Protocol Architects/Implementers	TCP/IP, UDP
O/S Architects/Implementers	UNIX, DOS
CPU, Memory, Disk Designers	Pentium, Alpha
LAN Interface Designers	Adapters
Media Access (LAN) Architects	FDDI, ATM
Optic Device Designers	Fibers, Lasers

- Faster link ≠ Faster applications
- Need to consider trends of all layers

Protocol Design: Key Parameters

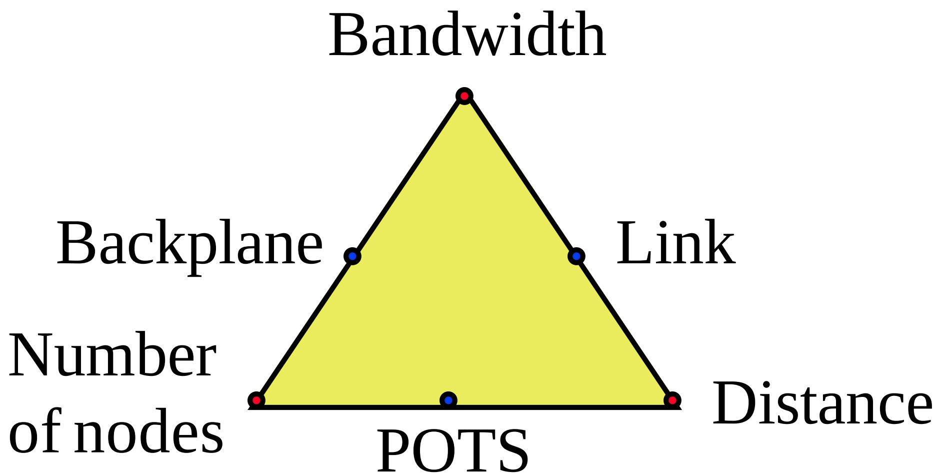

- Two out of three is trivial
- Potential: 25 THz on a single fiber
- 1992 Records: 5 Gbps over 15,000 km
  10 Gbps over 11,000 km
  10 Gbps over 4500 km fiber
- Borderless society ⇒ Increasing distances
  Increasing # of nodes

The Magic Word: α

Performance Fundamentals

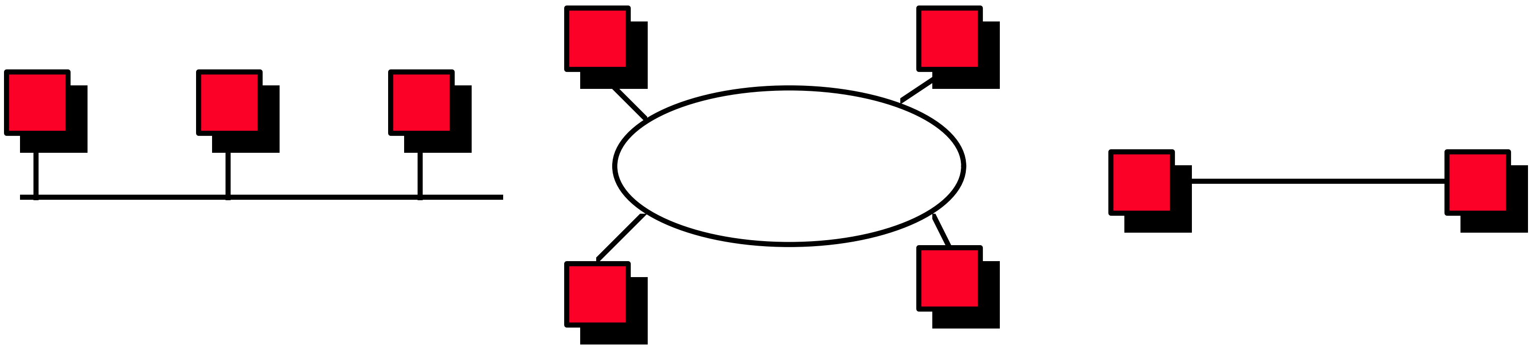

- Efficiency = Maximum throughput/Media bandwidth
- Efficiency is a decreasing function of α

  = Propagation delay /Transmission time

  = (Distance/Speed of light)/(Transmission size/Bits/sec)

  = Distance*Bits/sec/(Speed of light)(Transmission size)
- Bit rate-distance-transmission size tradeoff.
- Most people cannot visualize bit rate but can see distance easily.

Lessons

- For any given access method: the throughput (or efficiency) goes down as either the bit rate is increased, distance is increased, or frame size is decreased.
- If you scale the bit rate and packet size by the same factor, all tilizations, delays remain same.
- If you increase the bit rate by a factor of 10 but decrease the distance by a factor of 10, ff remains same.
- If you increase the bit rate by a factor of 10 but increase the frame size by a factor of 10, ff remains same.
- Designing a high-speed network is somewhat similar to designing a l-speed long-distance network.

Networking to Mars

- Distance*speed = constant
- 1 Gb/s between Boston and San Francisco is similar to 56 kb/s to Mars
- Earth-Mars Distance/Boston-SF Distance
  $= 49 \times 10^6$ Miles/3128 Miles = 1 Gb/s/56 kb/s
- Rule of Thumb: Don't do on a high-speed network, what you wouldn't do on a network to Mars.

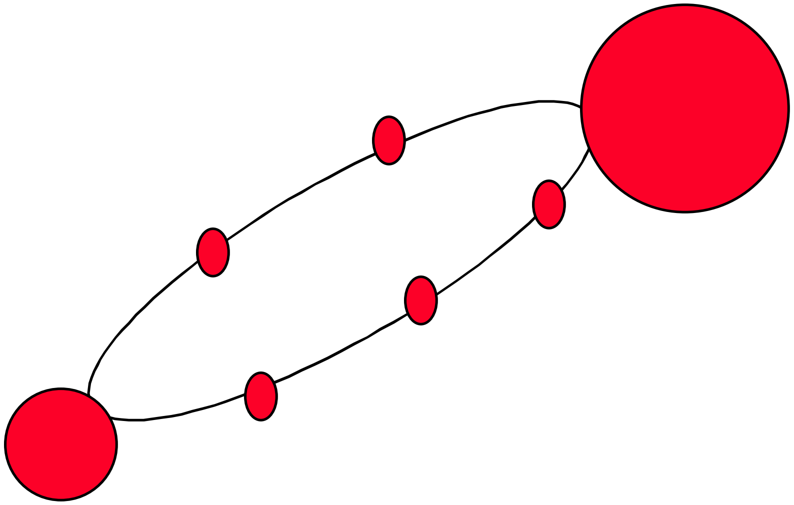

What You Wouldn't Do on A Network to Mars?

Media Access:

- Transmit and wait to hear others (e.g., Ethernet)
- Hold token while your frame goes around the ring (e.g., IEEE 802.5)
- Hold entire path while using only a part of it (e.g., FDDI) $\Rightarrow$ Spatial reuse

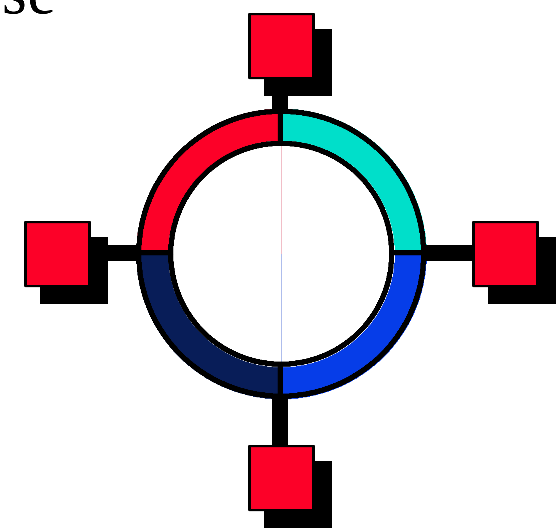

What You Wouldn't Do on A Network to Mars?

Transport or logical-link layer:

- Drop all packets if one is lost
  ⇒ (out-of-order caching)
- Retransmit all packets when just one is lost
  ⇒ (Selective retransmission)
- Wait for a packet to be resent to you if it is lost
  ⇒ (Forward Error Correction)
- Wait until last minute to order
  ⇒ (Anticipation, prefetching)
- Wait for a three-way (or two-way) handshake before sending first byte ⇒ (Implicit handshake)
- Summary: Minimize delay vs maximize throughput
  ⇒ Generation gap

Trends in Applications

- Little Voice
- AT&T: 125 to 130 M calls/day @ 5 min/call 64 kbps = 28.8 Gbps = 1/1000 of one fiber
- 200 Million X 24 hr/day X 64 kbps = 12.8 Tbps
- Survey of 1750 businesses:

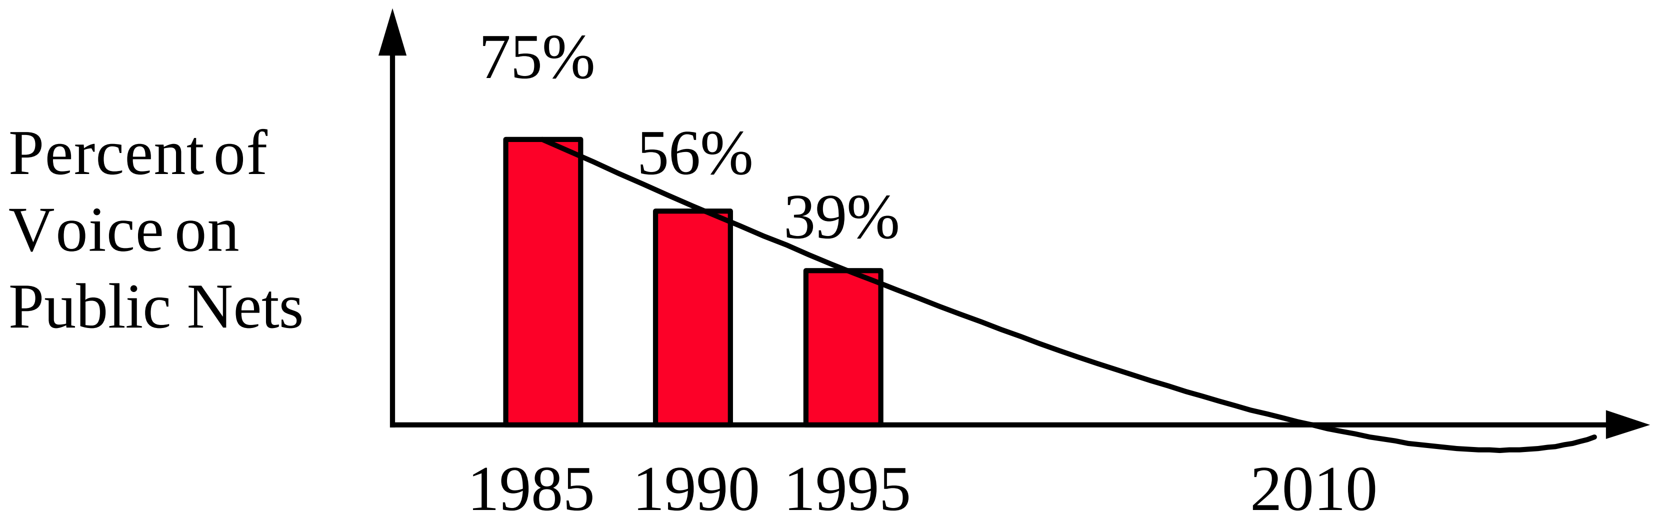

◆Ref: IEEE Spectrum, August 1992, p 19.

Video Characteristics

- Size: 1 Hr uncompressed HDTV= 540 GB = $150/sec
- 1 Hr compressed HDTV= 9 GB = $2.5/sec
  ⇒ Needs to be compressed for storage
  ⇒ Variable bit rate
- Holding time: At 1 Gbps:
  - 10 Mb image = 10 ms
  - 1 hour compressed VHS movie =10 secs or less
    ⇒ Bursty short-lived traffic

Electro-optic Bottleneck

- Bandwidth of fiber = 25 THz/window
- Bandwidth of electronics = 1-10 Gbps
- Switching bottleneck ⇒ Optical switching ⇒ All-optical networks
- Switches more expensive than media: Less switches and more links
- Higher connectivity, less hops
- Distributed media shared switching (like WANs) and not distributed
- switching shared media (like LANs)

Switching vs Transmission

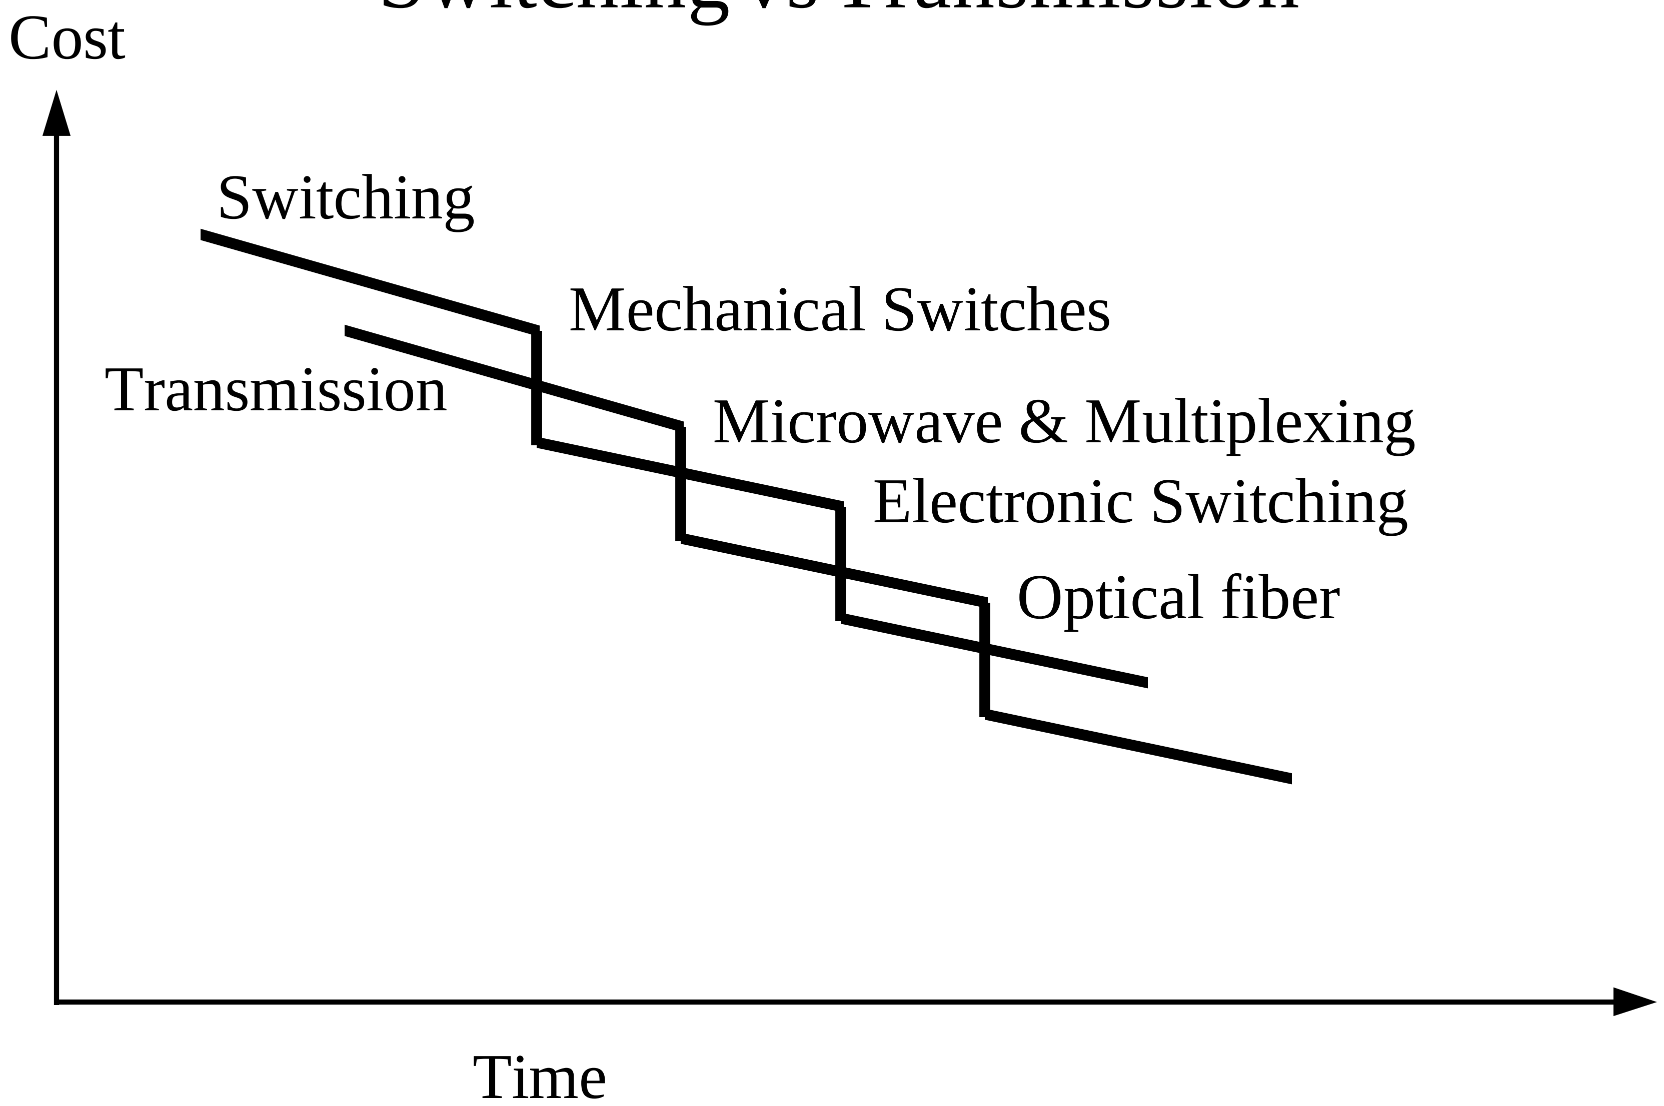

Shared Media vs Shared Switches

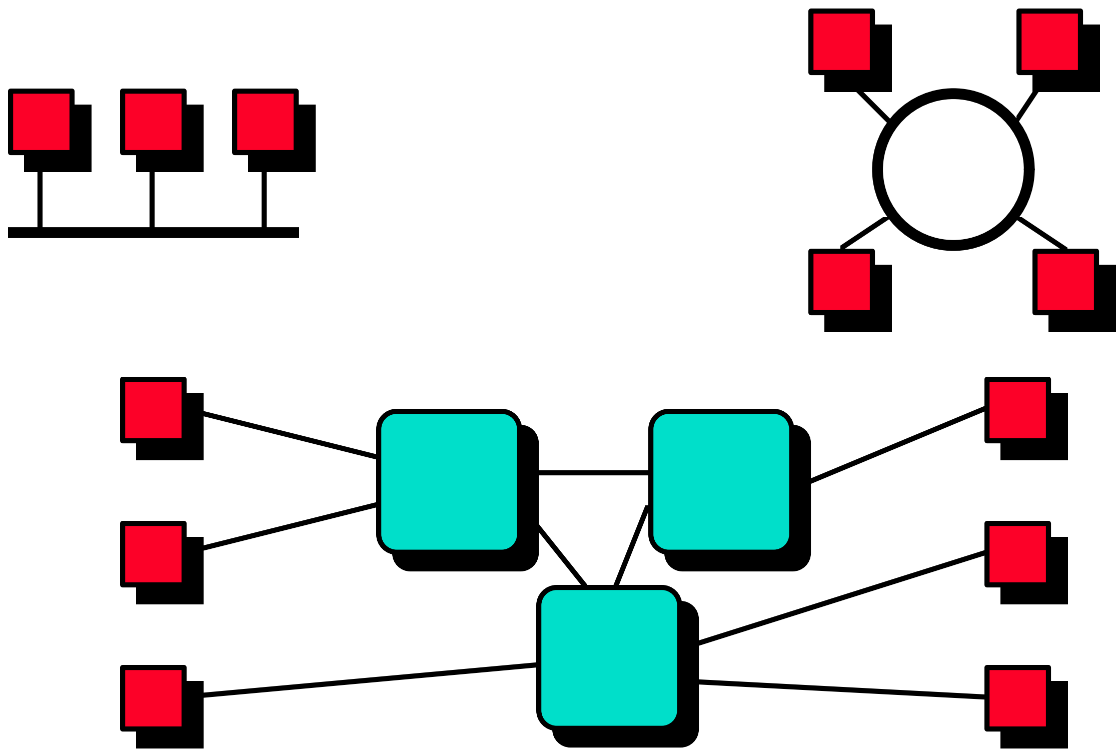

- Variable bandwidth/station
- Cost $\propto$ bandwidth
- Incremental upgradability
- Natural spatial reuse

Summary

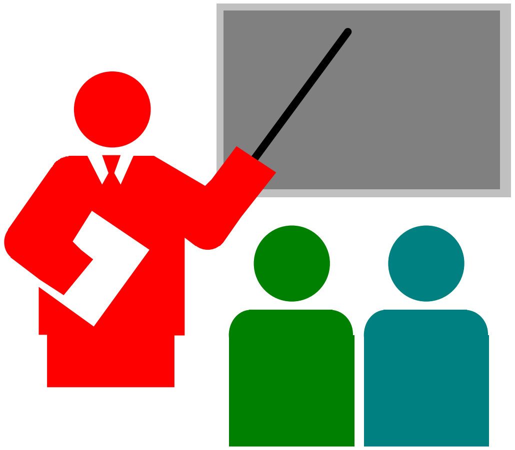

- High-speed links iff economy of scale.
- Bursty, short holding time traffic.
- Shared-switch distributed-media. No shared-media access.
- Speed-distance-transmission size tradeoff ⇒ Don't do on a high-speed network what you wouldn't do on a network to Mars.

Further Reading

- Craig Partridge, *Gigabit Networking*, Addison-Wesley, 1993.
- D. Clark and D. Tennenhouse, "Architectural Considerations for a New Generation of Protocols," Proc. SIGCOMM-90, pp. 200-208, August 1990..
- IEEE Journal on Selected Areas in Communications, Special Issue on High-Speed Computer Network Interfaces, February 1993.
- C. Partridge, "How slow is One Gigabit Per Second," Computer Communications Review, January 1990
- D. Clark, et al, "An Analysis of TCP Processing Overhead," IEEE Communications Magazine, June 1989.
- Craig Partridge, "Building Gigabit Network Interfaces," ConneXions, 1993.

References (Cont)

- R. Watson and S. Mamrak, “Gaining Efficiency in Transport Services by Appropriate Design and Implementation Choices,” ACM Transactions on Computer Systems, May 1987
- P. Druschel, et al, “Network Subsystem Design,” IEEE Network, July 1993.
- J. Smith and C. Traw, “Giving Applications Access to Gb/s Network,” IEEE Network, July 1993.