Throughput

- *throughput:* rate (bits/time unit) at which bits transferred between sender/receiver
  - *instantaneous:* rate at given point in time
  - *average:* rate over longer period of time

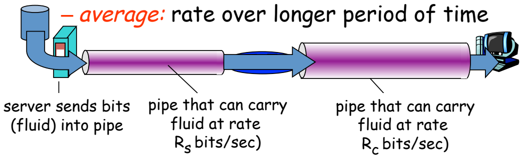

Throughput (more)

- $R_s < R_c$ What is average end-end throughput?

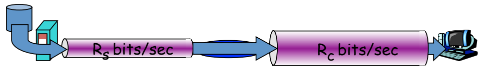

- $R_s > R_c$ What is average end-end throughput?

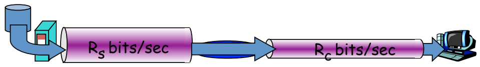

bottleneck link

link on end-end path that constrains end-end throughput

Throughput: Internet scenario

- per-connection end-end throughput: min (R_c,R_s,R/10)
- in practice: R_c or R_s is often bottleneck

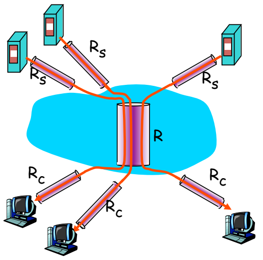

10 connections (fairly) share backbone bottleneck link R bits/sec

Protocol "Layers"

Networks are complex!

- many "pieces":
  - hosts
  - routers
  - links of various media
  - applications
  - protocols
  - hardware, software

Question:

Is there any hope of *organizing* structure of network?

Or at least our discussion of networks?

Distributed implementation of layer functionality

Departing airport

- ticket (purchase)
- baggage (check)
- gates (load)
- runway takeoff
- airplane routing

arriving airport

- ticket (complain)
- baggage (claim)
- gates (unload)
- runway landing
- airplane routing

intermediate air traffic sites

- airplane routing
- airplane routing
- airplane routing

Why layering?

Dealing with complex systems:

- explicit structure allows identification, relationship of complex system's pieces
  - layered reference model for discussion
- modularization eases maintenance, updating of system
  - change of implementation of layer's service transparent to rest of system
  - e.g., change in gate procedure doesn't affect rest of system
- layering considered harmful?

Internet protocol stack

- application: supporting network applications
  - FTP, SMTP, STTP
- transport: host-host data transfer
  - TCP, UDP
- network: routing of datagrams from source to destination
  - IP, routing protocols
- link: data transfer between neighboring network elements
  - PPP, Ethernet
- physical: bits "on the wire"

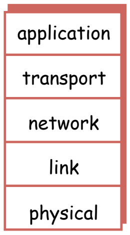

ISO/OSI reference model

- presentation: allow applications to interpret meaning of data, e.g., encryption, compression, machine-specific conventions
- *session:* synchronization, checkpointing, recovery of data exchange
- Internet stack "missing" these layers!
  - these services, *if needed,* must be implemented in application
  - needed?

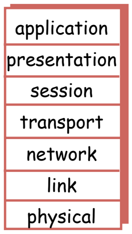

Layering: logical communication

Each layer:

- distributed
- "entities" implement layer functions at each node
- entities perform actions, exchange messages with peers

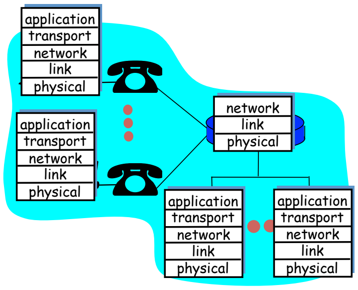

Layering: *logical* communication

E.g.: transport

- take data from app
- add addressing, reliability check info to form "datagram"
- send datagram to peer
- wait for peer to ack receipt
- analogy: post office

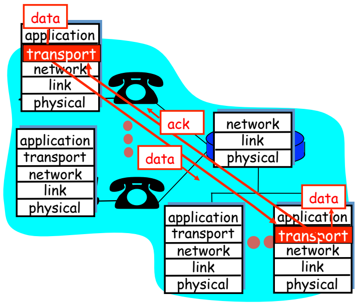

Layering: physical communication

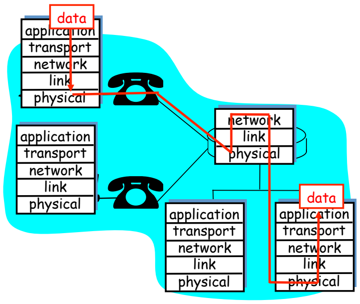

Protocol layering and data

Each layer takes data from above

- adds header information to create new data unit
- passes new data unit to layer below

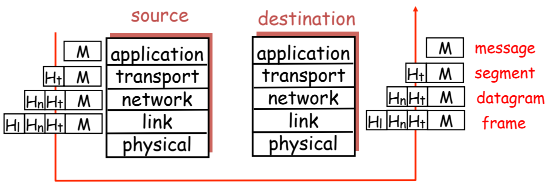

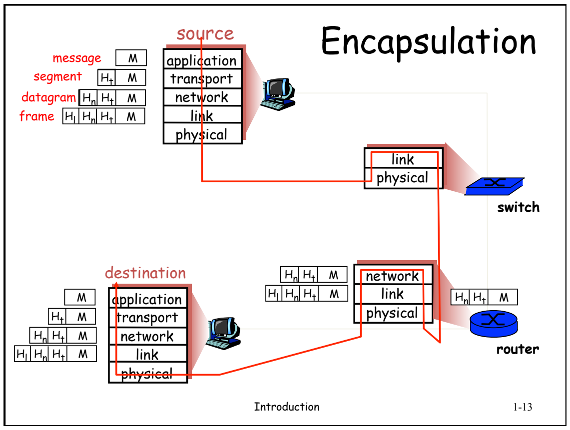

Encapsulation

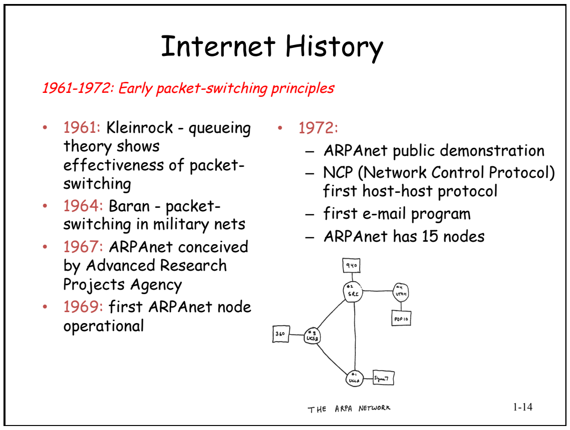

Internet History

1961-1972: Early packet-switching principles

- 1961: Kleinrock - queueing theory shows effectiveness of packet-switching
- 1964: Baran - packet-switching in military nets
- 1967: ARPAnet conceived by Advanced Research Projects Agency
- 1969: first ARPAnet node operational
- 1972:
  - ARPAnet public demonstration
  - NCP (Network Control Protocol) first host-host protocol
  - first e-mail program
  - ARPAnet has 15 nodes

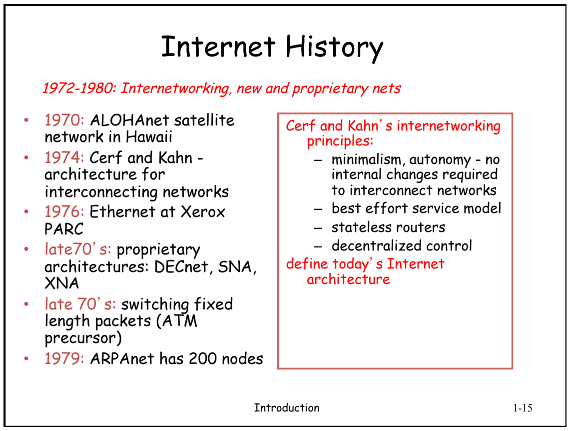

Internet History

1972-1980: Internetworking, new and proprietary nets

- 1970: ALOHAnet satellite network in Hawaii
- 1974: Cerf and Kahn - architecture for interconnecting networks
- 1976: Ethernet at Xerox PARC
- late70's: proprietary architectures: DECnet, SNA, XNA
- late 70's: switching fixed length packets (ATM precursor)
- 1979: ARPAnet has 200 nodes

Cerf and Kahn's internetworking principles:
- minimalism, autonomy - no internal changes required to interconnect networks
- best effort service model
- stateless routers
- decentralized control

define today's Internet architecture

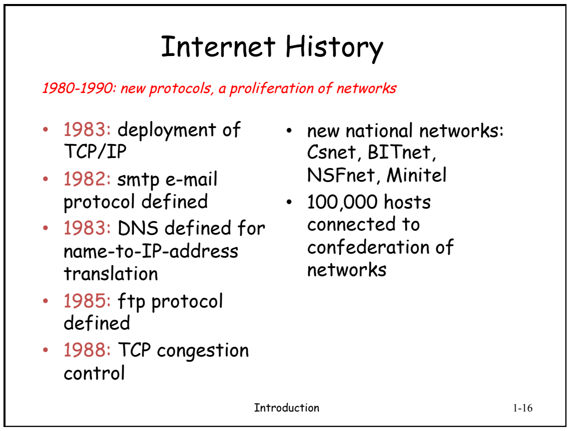

Internet History

1980-1990: new protocols, a proliferation of networks

- 1983: deployment of TCP/IP
- 1982: smtp e-mail protocol defined
- 1983: DNS defined for name-to-IP-address translation
- 1985: ftp protocol defined
- 1988: TCP congestion control
- new national networks: Csnet, BITnet, NSFnet, Minitel
- 100,000 hosts connected to confederation of networks

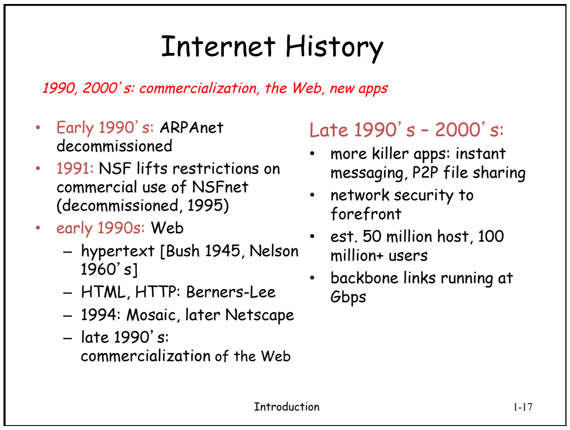

Internet History

1990, 2000's: commercialization, the Web, new apps

- Early 1990's: ARPAnet decommissioned
- 1991: NSF lifts restrictions on commercial use of NSFnet (decommissioned, 1995)
- early 1990s: Web
  - hypertext [Bush 1945, Nelson 1960's]
  - HTML, HTTP: Berners-Lee
  - 1994: Mosaic, later Netscape
  - late 1990's: commercialization of the Web

Late 1990's – 2000's:
- more killer apps: instant messaging, P2P file sharing
- network security to forefront
- est. 50 million host, 100 million+ users
- backbone links running at Gbps

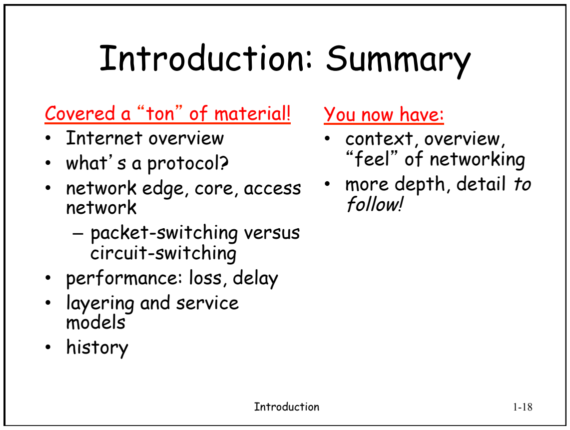

Introduction: Summary

Covered a "ton" of material!
- Internet overview
- what's a protocol?
- network edge, core, access network
  - packet-switching versus circuit-switching
- performance: loss, delay
- layering and service models
- history

You now have:
- context, overview, "feel" of networking
- more depth, detail *to follow!*

Quiz #1:

1. Consider an application that transmits data at a steady rate. Also, when such an application starts, it will continue running for a relatively long period of time. Answer the following questions:
   a) Would a packet-switched network or a circuit-switched network be more appropriate for the application? Why?
   b) Suppose that a packet-switched network is used and the only traffic in this network comes from such applications as described above. Furthermore, assume that the sum of the application data rates is less than the capacities of each and every link. Is some form of congestion control needed? Why?
2. Suppose hosts, A and B, are separated by 10,000 kilometers and are connected by a direct link of R= 1 Mbps. Suppose the propagation speed over the link is 2.5 * 10^8 meters/sec. Consider sending a file of 400,000 bits from A to B. Suppose the file is sent continuously as one big message. What is the maximum number of bits that will be in the link at any given time?

Answers to Problem 1.

a) A circuit-switched network would be well suited to the application described, because the application involves long sessions with predictable smooth bandwidth requirements. Since the transmission rate is known and not bursty, bandwidth can be reserved for each application session circuit with no significant waste. In addition, we need not worry greatly about the overhead costs of setting up and tearing down a circuit connection, which are amortized over the lengthy duration of a typical application session.

b) Given such generous link capacities, the network needs no congestion control mechanism. In the worst (most potentially congested) case, all the applications simultaneously transmit over one or more particular network links. However, since each link offers sufficient bandwidth to handle the sum of all of the applications' data rates, no congestion (very little queueing) will occur.

Answers to Problem 2.

40,000 bits

Tprop = 0.1/2.5 sec = 0.04 sec

Data transmitted in 0.04 sec

0.04 * 1 Mbps = 40,000 bits