Part I: Introduction

Chapter goal:

- get context, overview, "feel" of networking
- more depth, detail *later* in course
- approach:
  - descriptive
  - use Internet as example

Overview:

- what's the Internet
- what's a protocol?
- network edge
- network core
- access net, physical media
- performance: loss, delay
- protocol layers, service models
- backbones, NAPs, ISPs
- history
- ATM network

What's the Internet: "nuts and bolts" view

- millions of connected computing devices: *hosts, end-systems*
  - pc's workstations, servers
  - PDA's phones, toasters

  running *network apps*
- *communication links*
  - fiber, copper, radio, satellite
- *routers:* forward packets (chunks) of data thru network

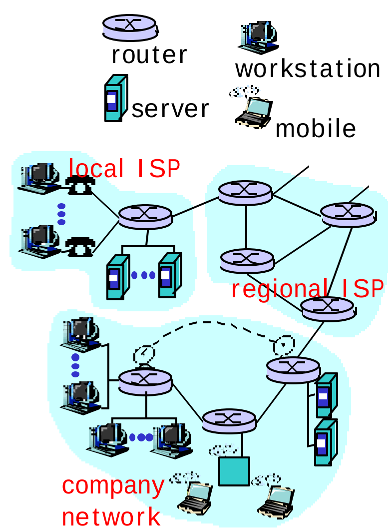

What's the Internet: "nuts and bolts" view

- *protocols*: control sending, receiving of msgs
  - e.g., TCP, IP, HTTP, FTP, PPP
- *Internet:* "network of networks"
  - loosely hierarchical
  - public Internet versus private intranet
- Internet standards
  - RFC: Request for comments
  - IETF: Internet Engineering Task Force

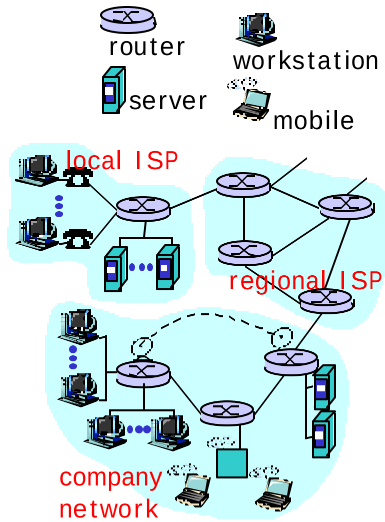

What's the Internet: a service view

- communication *infrastructure* enables distributed applications:
  - WWW, email, games, e-commerce, database., voting,
  - more?
- communication services provided:
  - connectionless
  - connection-oriented
- cyberspace [Gibson]:
  "a consensual hallucination experienced daily by billions of operators, in every nation, ..."

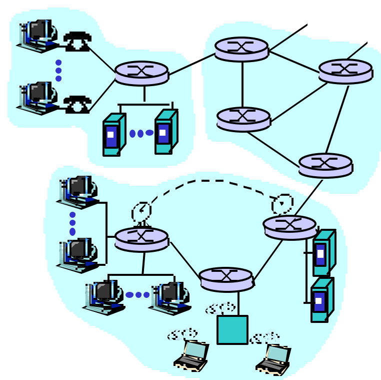

What's a protocol?

human protocols:

- "what's the time?"
- "I have a question"
- introductions

… specific msgs sent

… specific actions taken when msgs received, or other events

network protocols:

- machines rather than humans
- all communication activity in Internet governed by protocols

protocols define format, order of msgs sent and received among network entities, and actions taken on msg transmission, receipt

What's a protocol?

a human protocol and a computer network protocol:

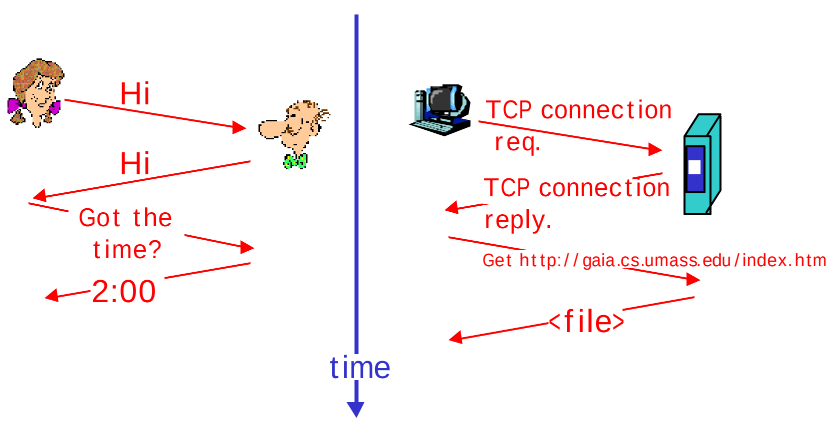

Q: Other human protocol?

A closer look at network structure:

- network edge: applications and hosts
- network core:
  - routers
  - network of networks
- access networks, physical media: communication links

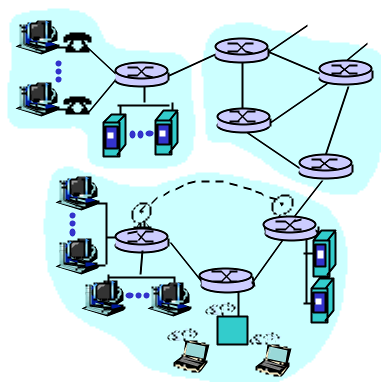

The network edge:

- end systems (hosts):
  - run application programs
  - e.g., WWW, email
  - at "edge of network"
- client/server model
  - client host requests, receives service from server
  - e.g., WWW client (browser)/ server; email client/server
- peer-peer model:
  - host interaction symmetric
  - e.g.: teleconferencing

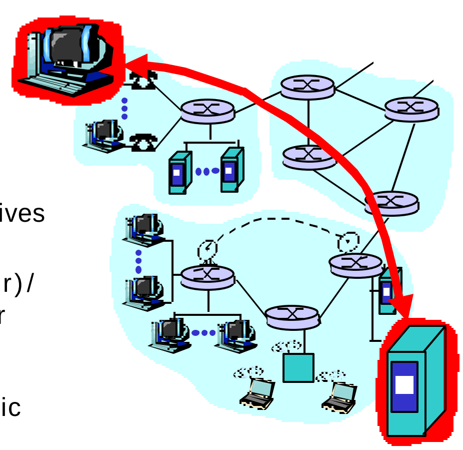

Network edge: connection-oriented service

Goal: data transfer between end sys.

- *handshaking:* setup (prepare for) data transfer ahead of time
  - Hello, hello back human protocol
  - *set up "state"* in two communicating hosts
- TCP - Transmission Control Protocol
  - Internet's connection-oriented service

TCP service [RFC 793]

- *reliable, in-order* byte-stream data transfer
  - loss: acknowledgements and retransmissions
- *flow control:*
  - sender won't overwhelm receiver
- *congestion control:*
  - senders "slow down sending rate" when network congested

Network edge: connectionless service

Goal: data transfer between end systems

- same as before!

- UDP - User Datagram Protocol [RFC 768]: Internet's connectionless service
  - unreliable data transfer
  - no flow control
  - no congestion control

App's using TCP:

- HTTP (WWW), FTP (file transfer), Telnet (remote login), SMTP (email)

App's using UDP:

- streaming media, teleconferencing, Internet telephony

The Network Core

- mesh of interconnected routers
- *the* fundamental question: how is data transferred through net?
  - circuit switching: dedicated circuit per call: telephone net
  - packet-switching: data sent thru net in discrete “chunks”

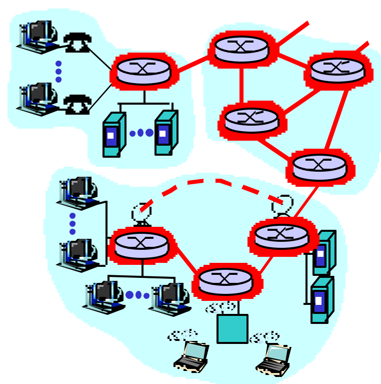

Network Core: Circuit Switching

End-end resources reserved for “call”

- link bandwidth, switch capacity
- dedicated resources: no sharing
- circuit-like (guaranteed) performance
- call setup required

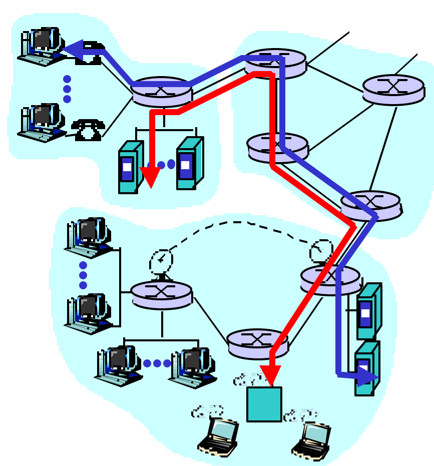

Network Core: Circuit Switching (Multiplexing)

network resources (e.g., bandwidth) divided into "pieces"

- pieces allocated to calls
- resource piece *idle* if not used by owning call *(no sharing)*
- dividing link bandwidth into "pieces"
  - frequency division
  - time division

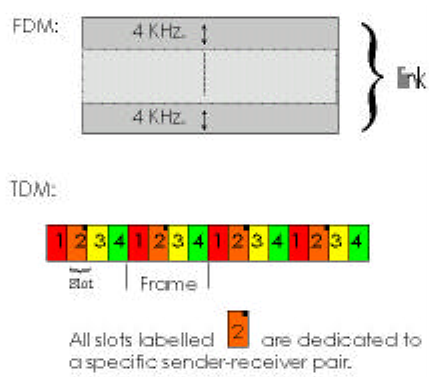

Network Core: Packet Switching

each end-end data stream divided into *packets*

- user A, B packets *share* network resources
- each packet uses full link bandwidth
- resources used *as needed*,

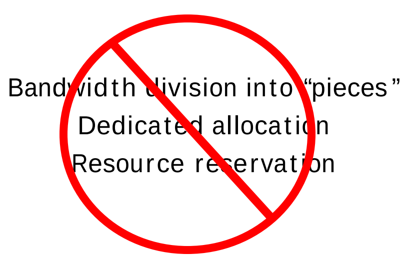

resource contention:

- aggregate resource demand can exceed amount available
- congestion: packets queue, wait for link use
- store and forward: packets move one hop at a time
  - transmit over link
  - wait turn at next link

Network Core: Packet Switching

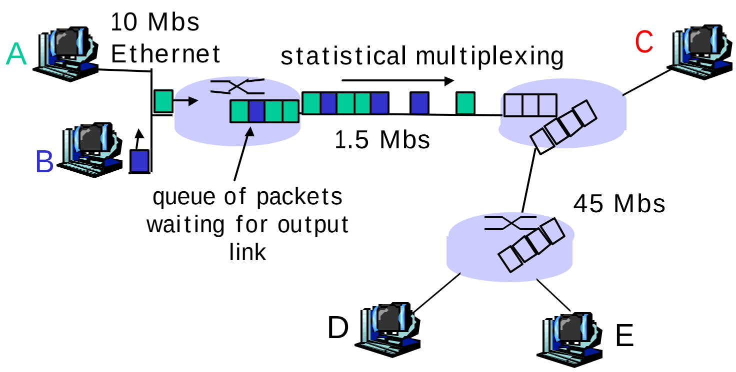

Packet-switching versus circuit switching: human restaurant analogy

- other human analogies?

Network Core: Packet Switching

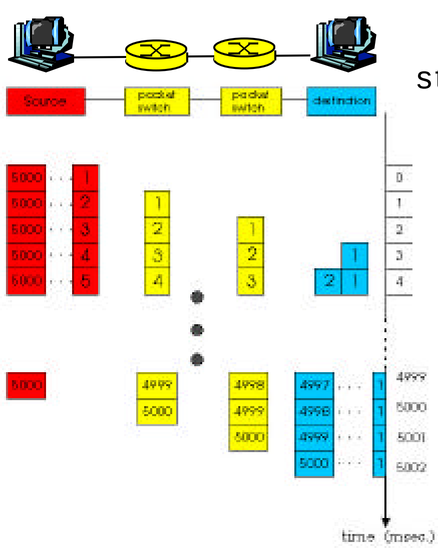

Packet-switching:
store and forward behavior

Packet switching versus circuit switching

Packet switching allows more users to use network!

- 1 Mbit link
- each user:
  - 100Kbps when “active”
  - active 10% of time
- circuit-switching:
  - 10 users
- packet switching:
  - with 35 users, probability > 10 active less that .004

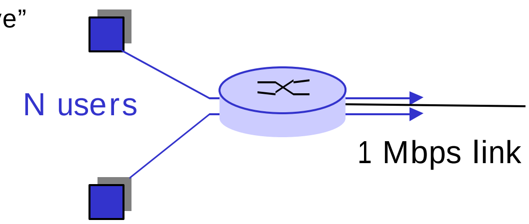

Packet switching versus circuit switching

Is packet switching a “slam dunk winner?”

- Great for bursty data
  - resource sharing
  - no call setup
- Excessive congestion: packet delay and loss
  - protocols needed for reliable data transfer, congestion control
- Q: How to provide circuit-like behavior?
  - bandwidth guarantees needed for audio/video apps

still an unsolved problem (chapter 6)

Packet-switched networks: routing

- *Goal:* move packets among routers from source to destination
  - we'll study several path selection algorithms (chapter 4)
- datagram network:
  - *destination address* determines next hop
  - routes may change during session
  - analogy: driving, asking directions
- virtual circuit network:
  - each packet carries tag (virtual circuit ID), tag determines next hop
  - fixed path determined at *call setup time*, remains fixed thru call
  - routers maintain per-call state

Network Taxonomy

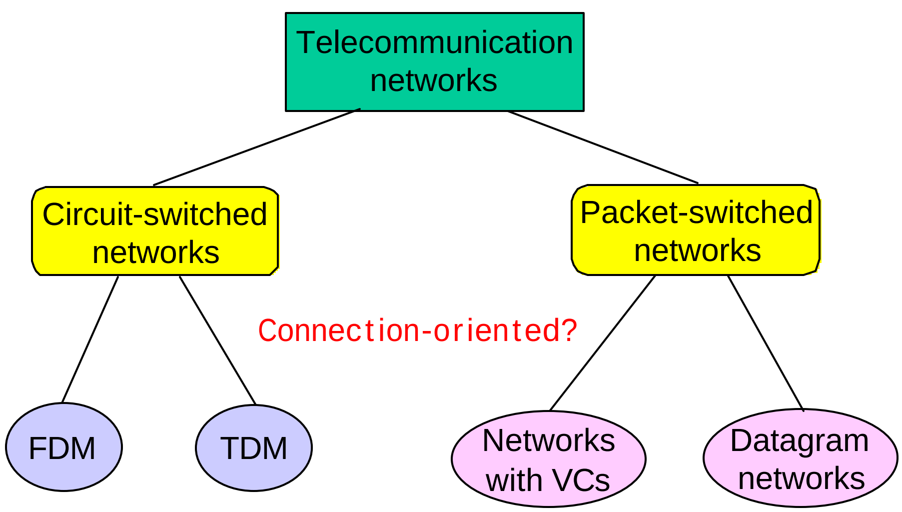

Access networks and physical media

Q: How to connection end systems to edge router?

- residential access nets
- institutional access networks (school, company)
- mobile access networks

Keep in mind:

- bandwidth (bits per second) of access network?
- shared or dedicated?

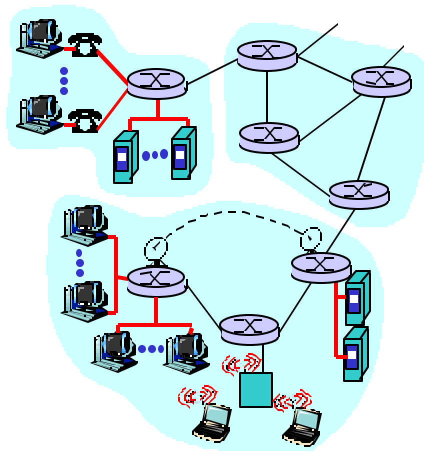

Residential access: point to point access

- Dialup via modem
  - up to 56Kbps direct access to router (conceptually)
- ISDN: intergrated services digital network: 128Kbps all-digital connect to router
- ADSL: asymmetric digital subscriber line
  - up to 1 Mbps home-to-router
  - up to 8 Mbps router-to-home
  - DSL/ADSL deployment: SBC, Covad

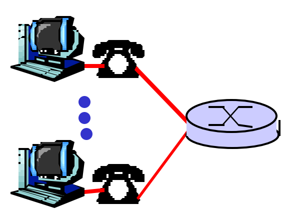

Residential access: cable modems

- HFC: hybrid fiber coax
  - asymmetric: up to 10Mbps upstream, 1 Mbps downstream
- network of cable and fiber attaches homes to ISP router
  - shared access to router among home
  - issues: congestion, dimensioning
- deployment: available via cable companies, e.g., Road Runner, @Home, MediaOne

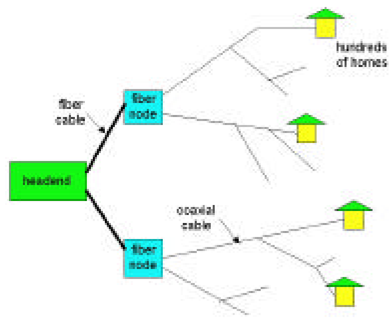

Institutional access: local area networks

- company/univ local area network (LAN) connects end system to edge router
- Ethernet:
  - shared or dedicated cable connects end system and router
  - 10 Mbs, 100Mbps, Gigabit Ethernet
- deployment: institutions, home LANs soon
- LANs: chapter 5

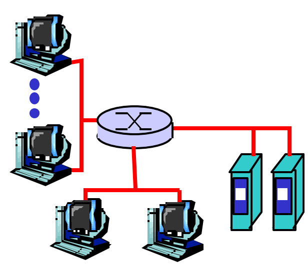

Wireless access networks

- shared *wireless* access network connects end system to router
- wireless LANs:
  - radio spectrum replaces wire
  - e.g., Lucent Wavelan 10 Mbps
- wider-area wireless access
  - CDPD: wireless access to ISP router via cellular network

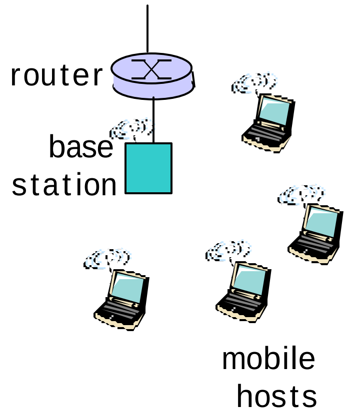

Physical Media

- physical link: transmitted data bit propagates across link
- guided media:
  - signals propagate in solid media: copper, fiber
- unguided media:
  - signals propagate freelye.g., radio

Twisted Pair (TP)

- two insulated copper wires
  - Category 3: traditional phone wires, 10 Mbps ethernet
  - Category 5 TP: 100Mbps ethernet

Physical Media: coax, fiber

Coaxial cable:

- wire (signal carrier) within a wire (shield)
  - baseband: single channel on cable
  - broadband: multiple channel on cable
- bidirectional
- common use in 10Mbs Ethernet

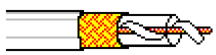

Fiber optic cable:

- glass fiber carrying light pulses
- high-speed operation:
  - 100Mbps/1Gbps Ethernet
  - high-speed point-to-point transmission (e.g., 5 Gps)
- low error rate

Qwest Fiber Networks

Physical media: radio

- signal carried in electromagnetic spectrum
- no physical “wire”
- bidirectional
- propagation environment effects:
  - reflection
  - obstruction by objects
  - interference

Radio link types:

- microwave
  - e.g. up to 45 Mbps channels
- LAN (e.g., waveLAN)
  - 2Mbps, 11Mbps
- wide-area (e.g., cellular)
  - e.g. CDPD, 10’s Kbps
- satellite
  - up to 50Mbps channel (or multiple smaller channels)
  - 270 Msec end-end delay
  - geosynchronous versus LEOS

Delay in packet-switched networks

packets experience delay on end-to-end path

- four sources of delay at each hop

- nodal processing:
  - check bit errors
  - determine output link
- queueing
  - time waiting at output link for transmission
  - depends on congestion level of router

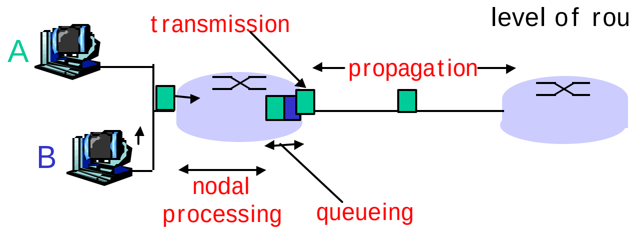

Delay in packet-switched networks

Transmission delay:

- R=link bandwidth (bps)
- L=packet length (bits)
- time to send bits into link = L/R

Propagation delay:

- d = length of physical link
- s = propagation speed in medium (~$2x10^8$ m/sec)
- propagation delay = d/s

Note: s and R are *very* different quantitites!

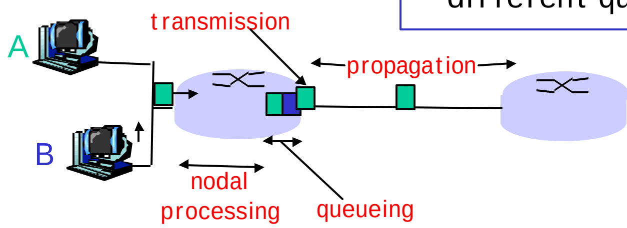

Queueing delay (revisited)

- R=link bandwidth (bps)
- L=packet length (bits)
- a=average packet arrival rate

traffic intensity = La/R

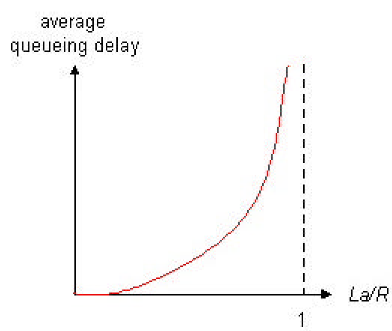

- La/R ~ 0: average queueing delay small
- La/R -> 1: delays become large
- La/R > 1: more "work" arriving than can be serviced, average delay infinite!

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?

Organization of air travel

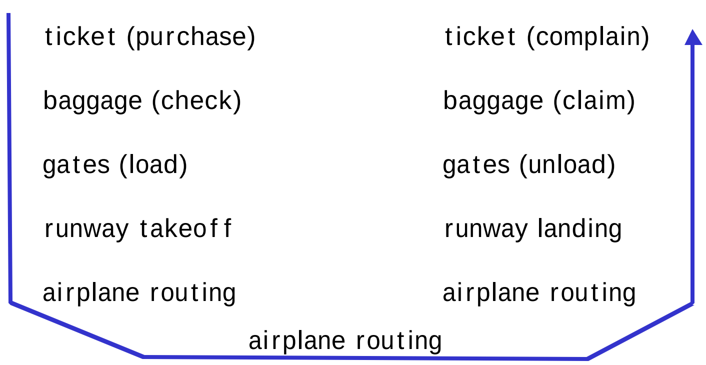

- a series of steps

Organization of air travel: a different view

ticket (purchase)	ticket (complain)
baggage (check)	baggage (claim)
gates (load)	gates (unload)
runway takeoff	runway landing
airplane routing	airplane routing
airplane routing	

Layers: each layer implements a service
- via its own internal-layer actions
- relying on services provided by layer below

Layered air travel: services

Counter-to-counter delivery of person+bags
baggage-claim-to-baggage-claim delivery
people transfer: loading gate to arrival gate
runway-to-runway delivery of plane
airplane routing from source to destination

Distributed implementation of layer functionality

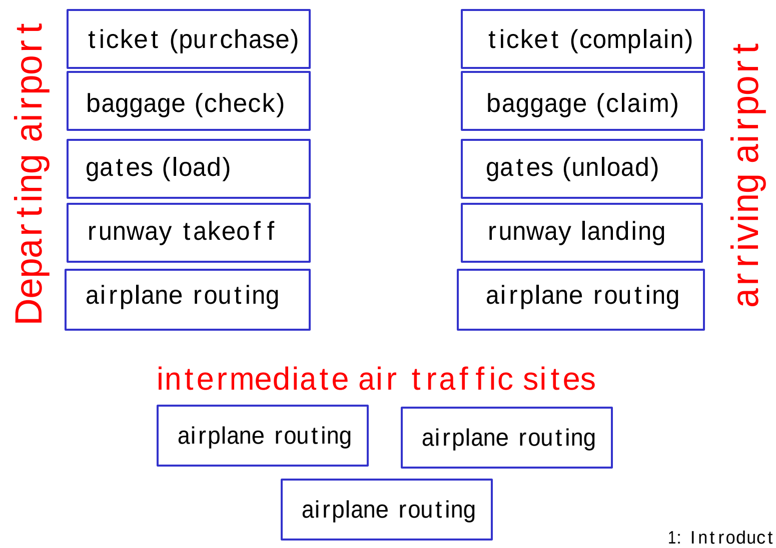

Why layering?

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?

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

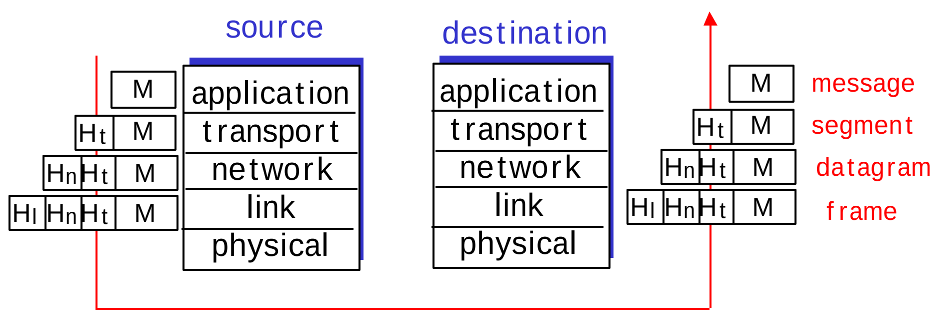

The OSI model

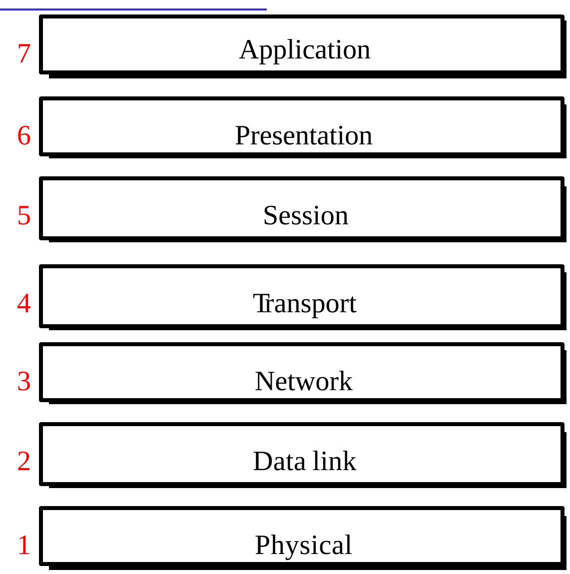

From: *TCP/IP Protocol Suite, McGraw-Hill, 2000*

OSI Layers

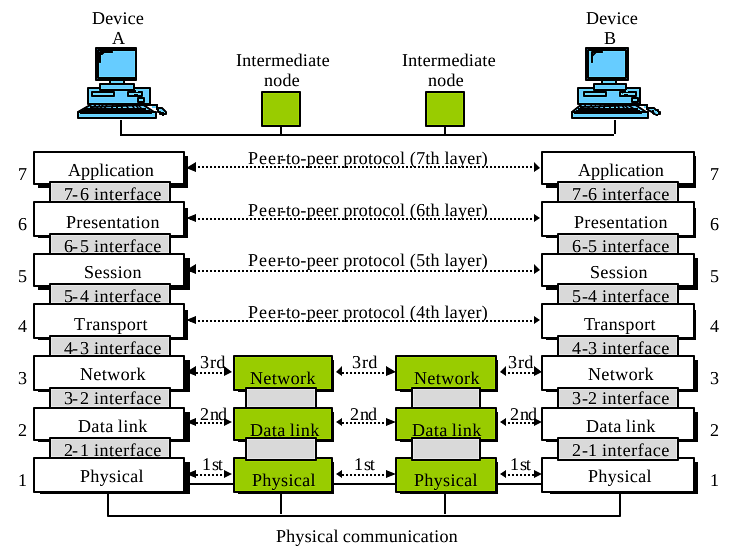

From: *TCP/IP Protocol Suite, McGraw-Hill, 2000*

TCP/IP and OSI model

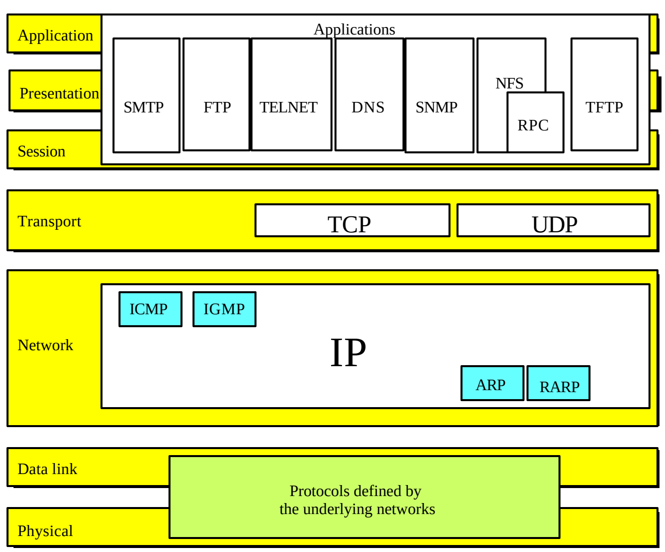

From: *TCP/IP Protocol Suite, McGraw-Hill, 2000*

Internet protocol stack

- application: supporting network applications
  - ftp, smtp, http
- 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”

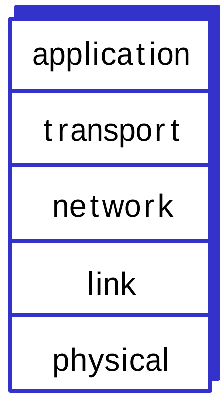

Layering: logical communication

Each layer:

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

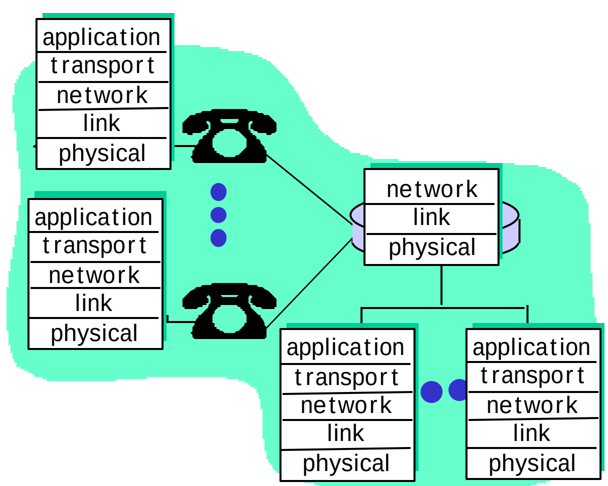

Layering: *logical* communication

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

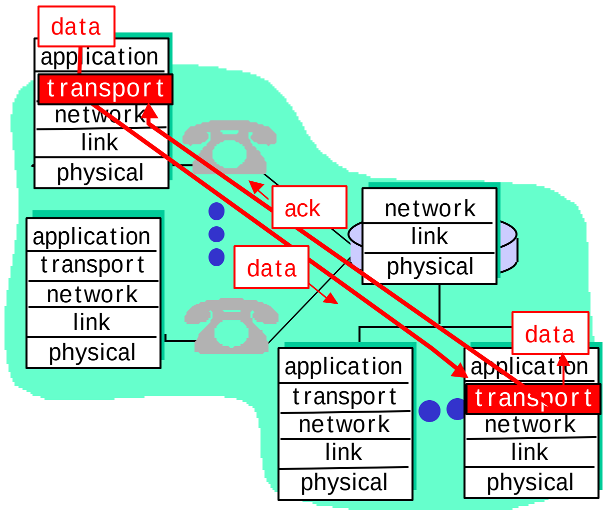

Layering: physical communication

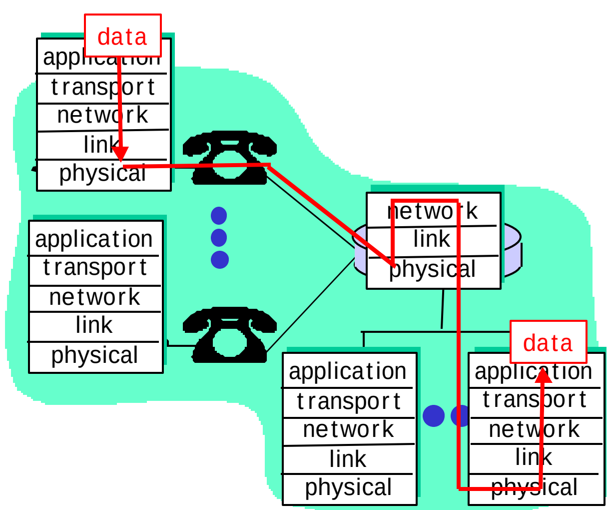

Internet structure: network of networks

- roughly hierarchical
- national/international backbone providers (NBPs)
  - e.g. BBN/GTE, Sprint, AT&T, IBM, UUNet
  - interconnect (peer) with each other privately, or at public Network Access Point (NAPs)
- regional ISPs
  - connect into NBPs
- local ISP, company
  - connect into regional ISPs

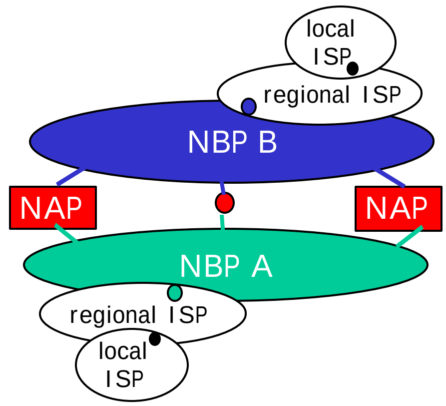

National Backbone Provider

e.g. BBN/GTE US backbone network

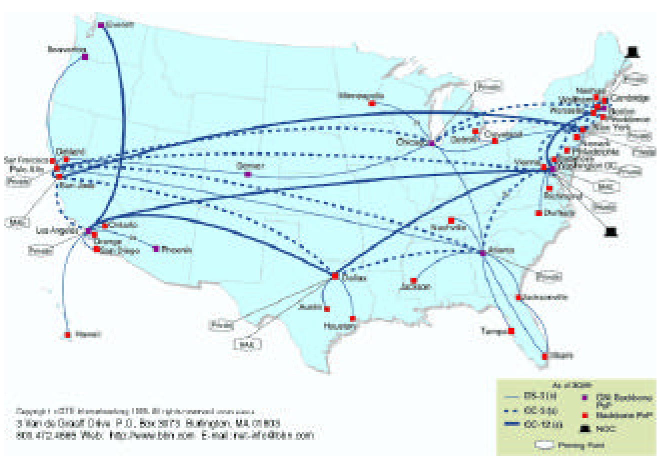

Other Backbone - vBNS

vBNS Backbone Network Map

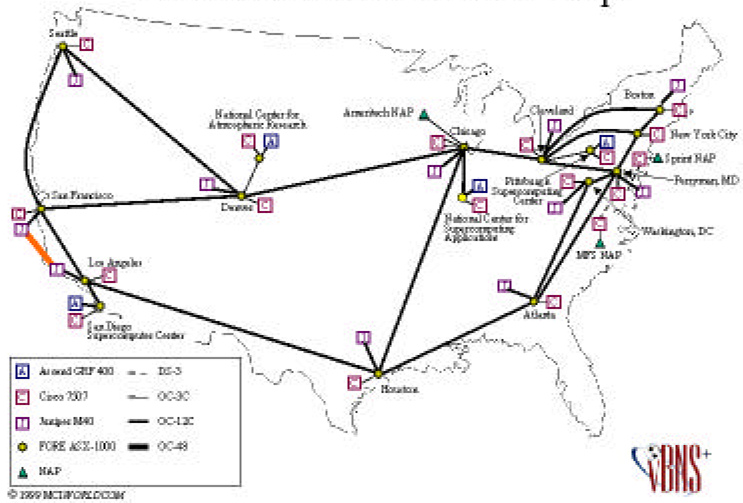

Other Backbone - vBNS (cont'd)

vBNS POS Backbone 2Q'00

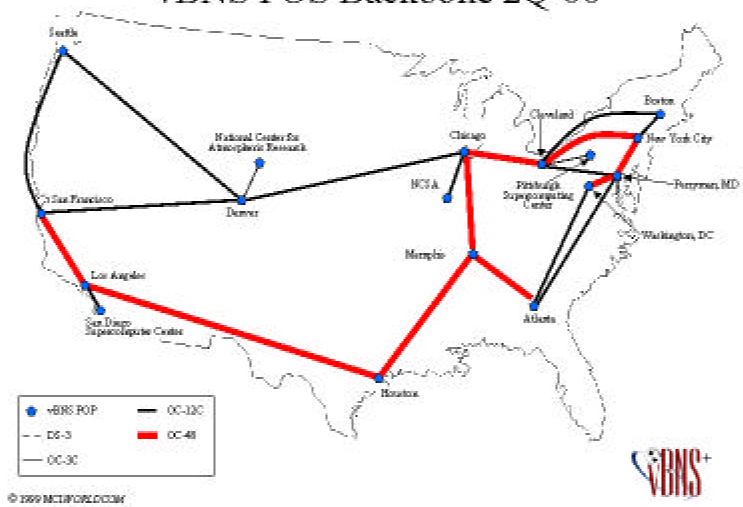

Internet 2

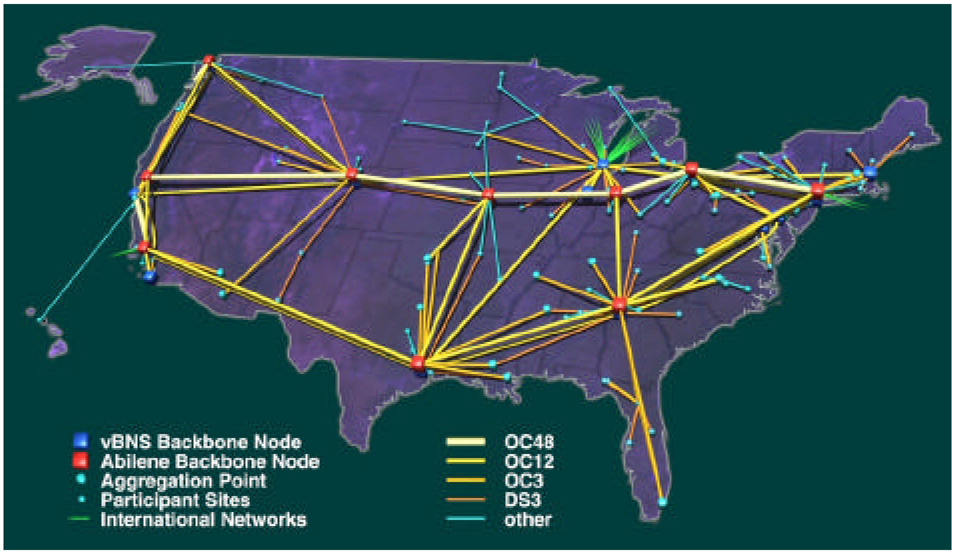

Internet History

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 Reearch Projects Agency
- 1969: first ARPAnet node operational
- 1972:
  - ARPAnet demonstrated publicly
  - 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
- 1973: Metcalfe's PhD thesis proposes Ethernet
- 1974: Cerf and Kahn - architecture for interconnecting networks
- 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's: commercialization, the WWW

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

Late 1990's:

- est. 50 million computers on Internet
- est. 100 million+ users
- backbone links runnning at 1 Gbps

ATM: Asynchronous Transfer Mode nets

Internet:

- today's *de facto* standard for global data networking

1980's:

- telco's develop ATM: competing network standard for carrying high-speed voice/data
- standards bodies:
  - ATM Forum
  - ITU

ATM principles:

- small (48 byte payload, 5 byte header) fixed length *cells* (like packets)
  - fast switching
  - small size good for voice
- virtual-circuit network: switches maintain state for each "call"
- well-defined interface between "network" and "user" (think of telephone company)

ATM layers

- ATM Adaptation Layer (AAL): interface to upper layers
  - end-system
  - segmentation/reassembly
- ATM Layer: cell switching
- Physical

Where's the application?

- ATM: lower layer
- functionality only
- IP-over ATM: later

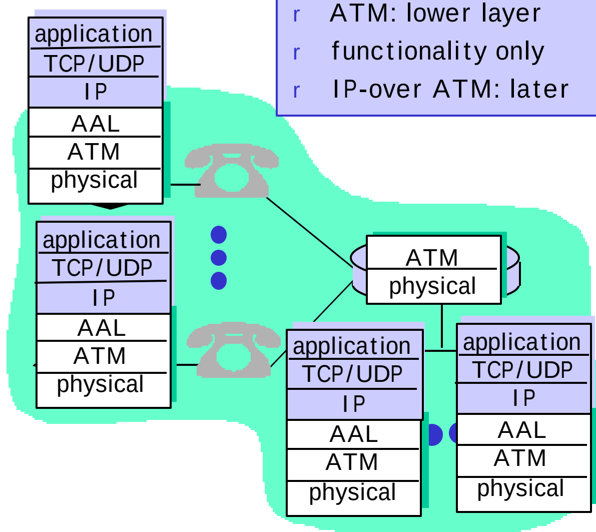

Chapter 1: Summary

Covered a "ton" of material!

- Internet overview
- what's a protocol?
- network edge, core, access network
- performance: loss, delay
- layering and service models
- backbones, NAPs, ISPs
- history
- ATM network

You now hopefully have:

- context, overview, "feel" of networking
- more depth, detail *later* in course