Chapter I: Introduction

Course on Computer Communication and Networks, CTH/GU

The slides are adaptation of the slides made available by the authors of the course's main textbook:

Computer Networking: A Top Down Approach , 5th edition.
Jim Kurose, Keith Ross
Addison-Wesley, July 2007.

Chapter I: Introduction

The slides are adaptation of the slides made available by the authors of the course'smain textbook

Overview:

- what's the Internet
- types of service
- ways of information transfer, routing, performance, delays, loss

---

- protocol layers, service models
- access net, physical media
- backbones, NAPs, ISPs
- (history)
- quick look into ATM networks

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

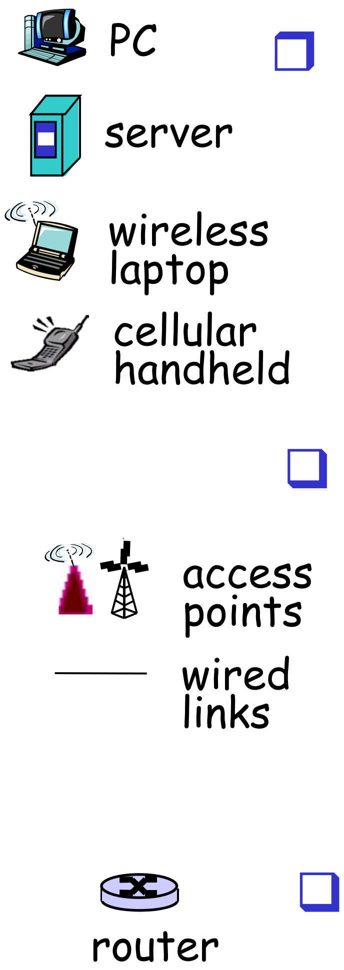

- millions of connected computing devices: *hosts = end systems*
  - running *network apps*
- *communication links*
  - fiber, copper, radio, satellite
  - transmission rate = *bandwidth*
- *routers:* forward packets (chunks of data)

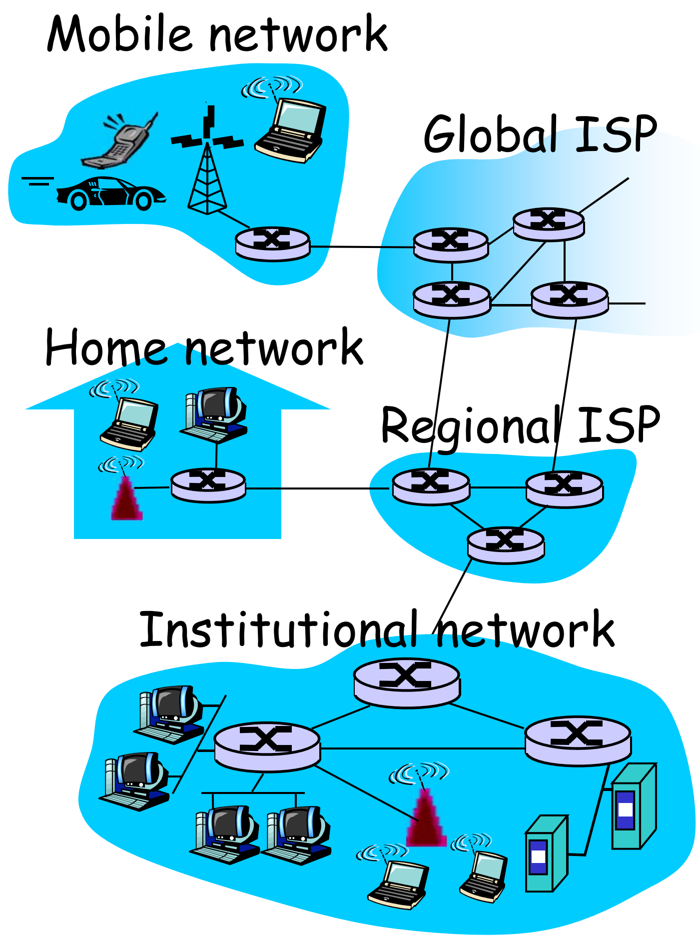

What's the Internet: "nuts and bolts" view

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

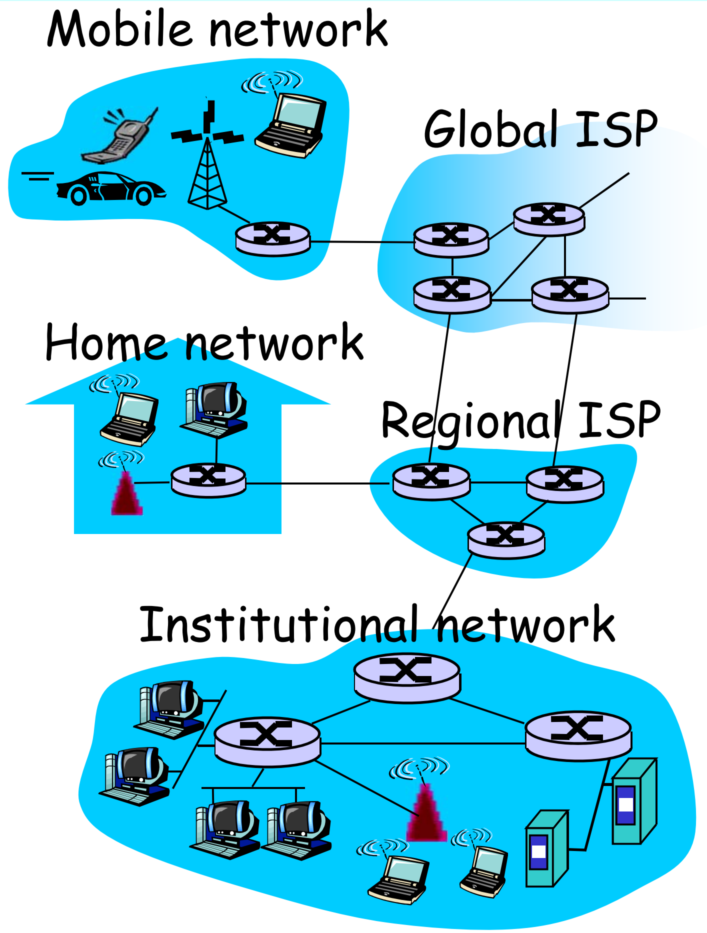

What's the Internet: a service view

- communication *infrastructure* enables distributed applications:
  - Web, VoIP, email, games, e-commerce, file sharing
- communication services provided to apps:
  - reliable data delivery from source to destination
  - "best effort" (unreliable) data delivery

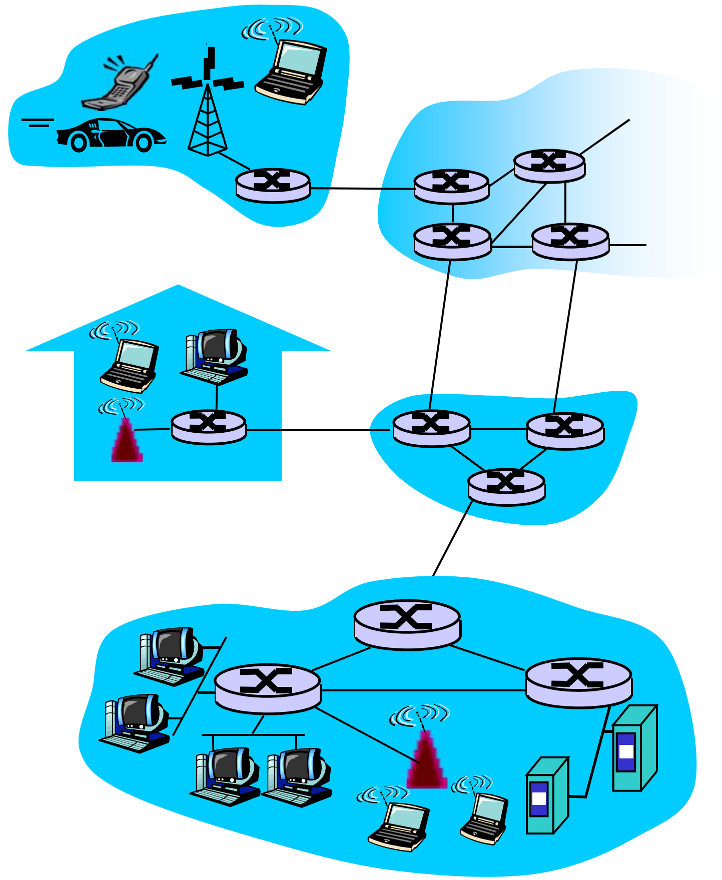

A closer look at network structure:

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

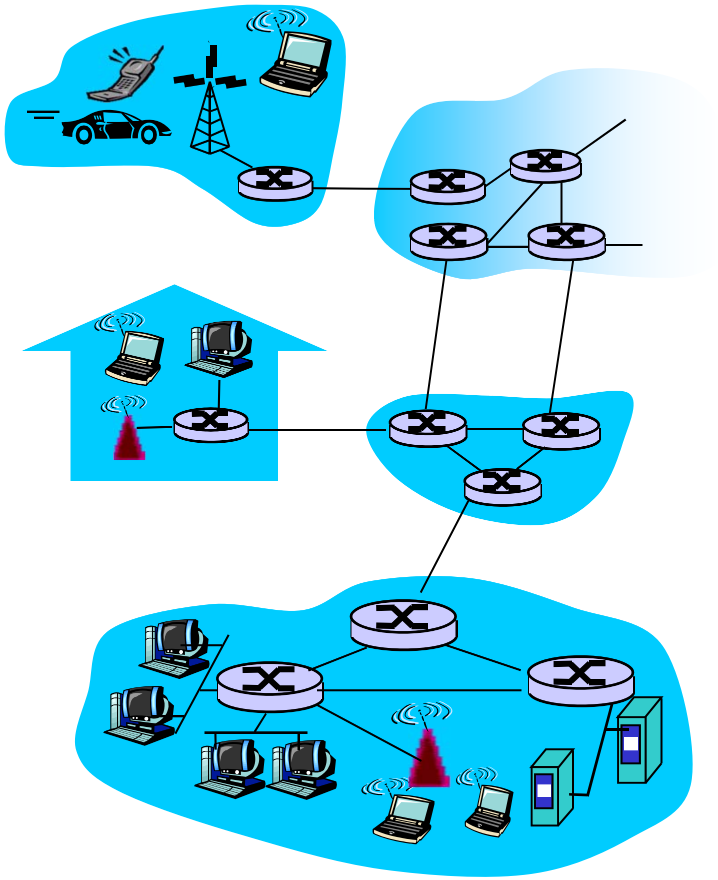

The network edge:

- end systems (hosts):
  - run application programs e.g. Web, email at "edge of network"
- client/server model
  - e.g. Web browser/server;
- peer-peer model:
  - e.g. Skype, BitTorrent

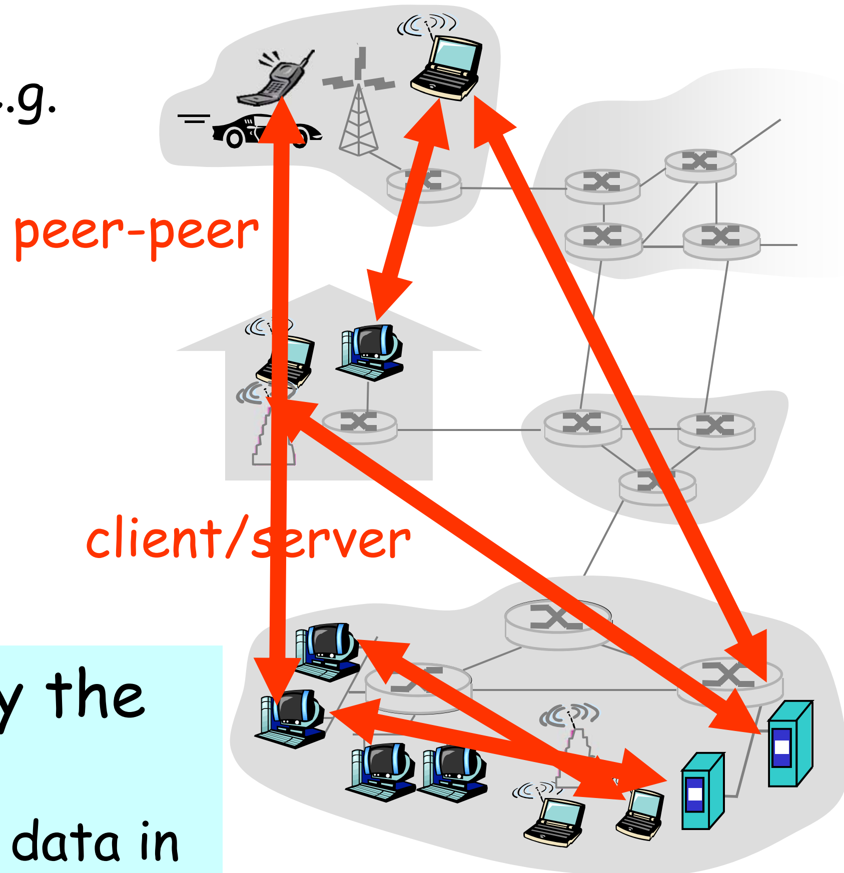

types of service offered by the network to applications:

connection-oriented: deliver data in the order they are sent

connectionless: delivery of data in arbitrary order

The Network Core

- mesh of interconnected routers
- fundamental question: how is data transferred through net?
- packet-switching: data sent thru net in discrete "chunks"
    - We will contrast with circuit switching: dedicated circuit per call: "classic"phone net

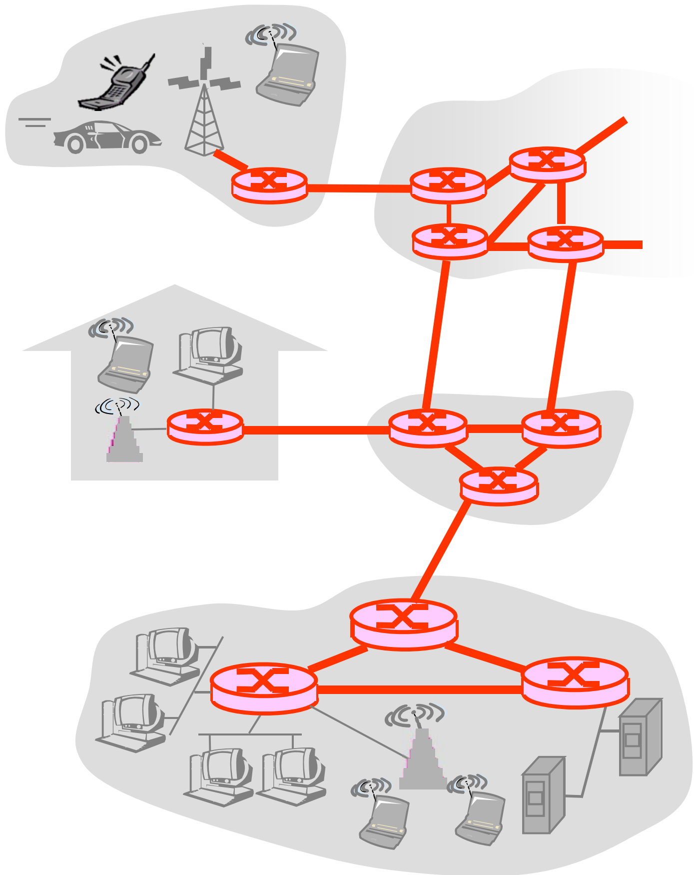

Network Core: Packet Switching

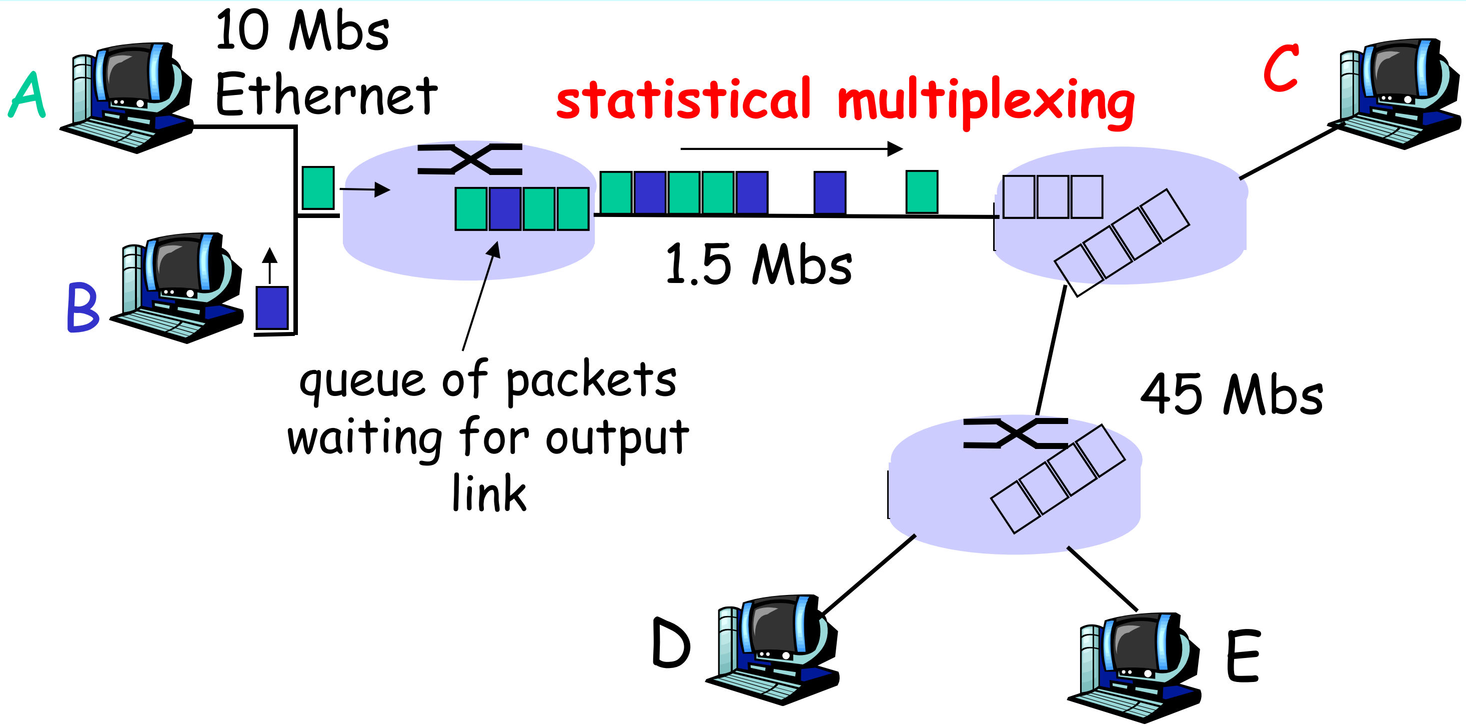

Network Core: Packet Switching

each end-end data stream divided into *packets*

- packets *share* network resources
- resources used *as needed*

store and forward:

- packets move one hop at a time
  - transmit over link
  - wait turn at next link
- http://www.youtube.com/watch?v=O7CuFlM4V54
  - Nice animation; disregard theterms used in narration; they do not follow exact protocolspecifications

resource contention:

- aggregate resource demand (bandwidth) can exceed amount available
- congestion: packets queue, wait for link use

Delay in packet-switched networks

packets experience delay on end-to-end path

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

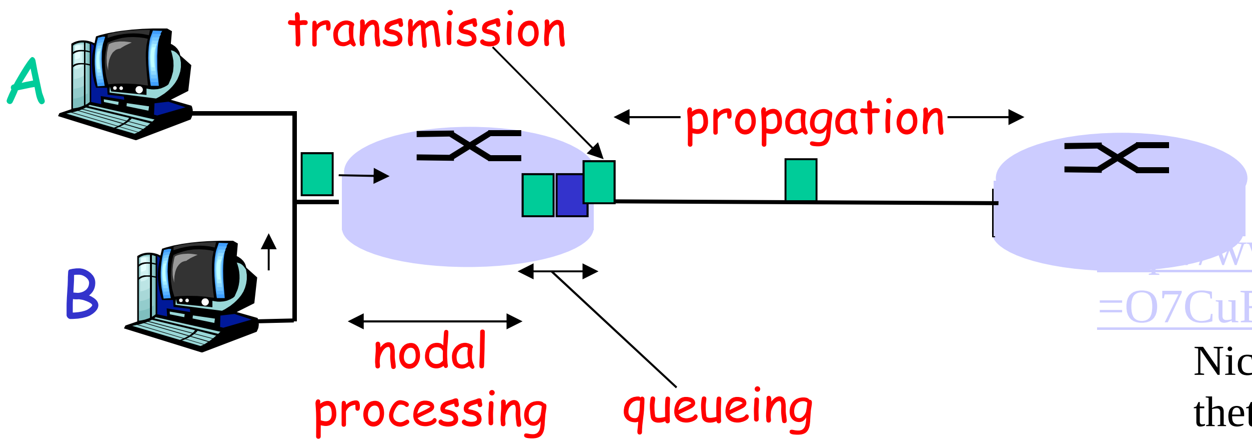

http://www.youtube.com/watch?v=O7CuFlM4V54

Nice animation; disregard theterms used in narration; they do not follow exact protocolspecifications

Delay in packet-switched networks

3. Transmission delay:

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

4. Propagation delay:

- d = length of physical link
- s = propagation speed in medium (~2×10^8 m/sec)
- propagation delay = d/s

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

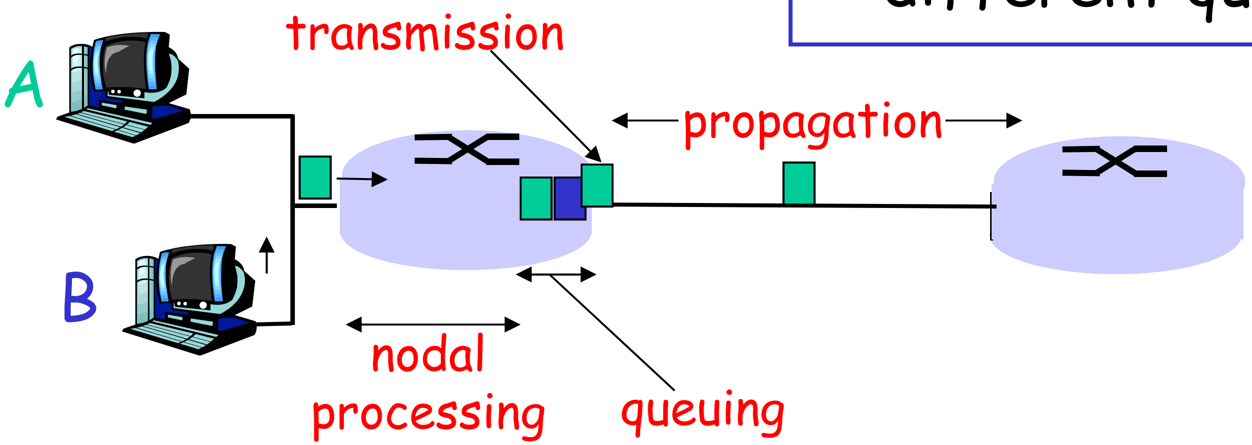

Circuit, message, packet switching

- store and forward behavior + other delays' visualization *(fig. from "Computer Networks" by A Tanenbaum,)*

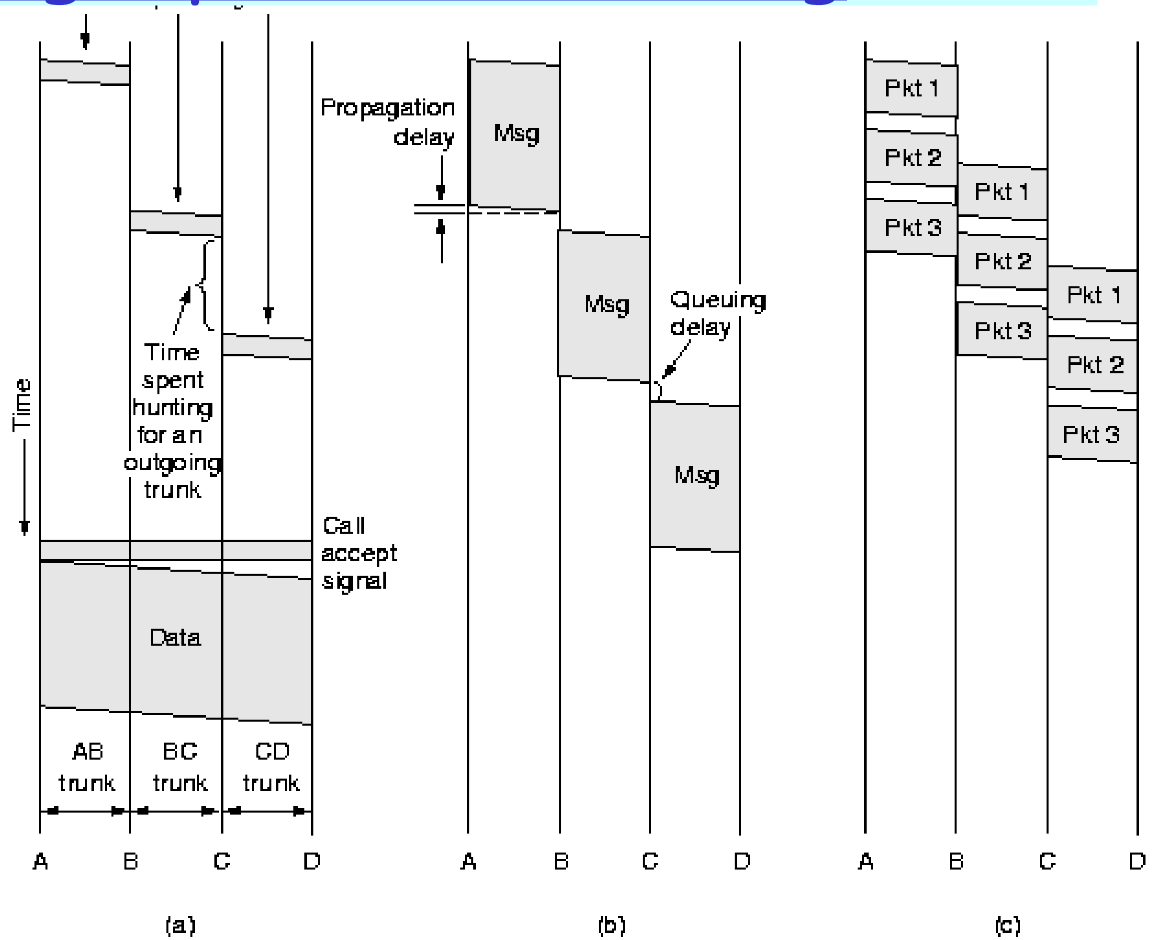

Fig. 2-35. Timing of events in (a) circuit switching, (b) message switching, (c) packet switching.

Network Core: Circuit Switching

End-end resources reserved/dedicated for "call"

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

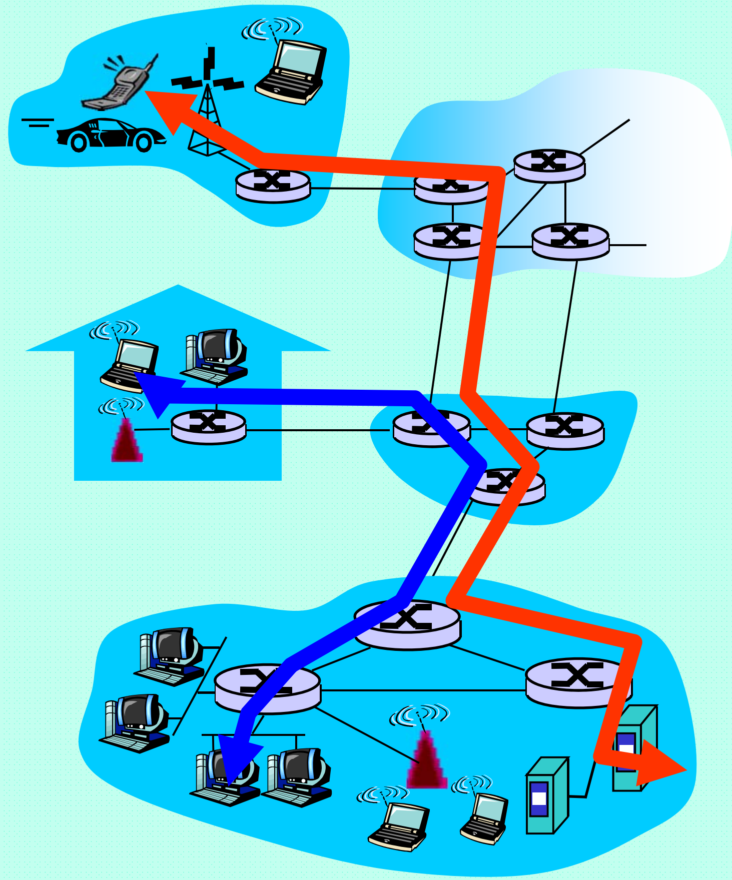

Packet switching versus "classical" circuit switching

Packet switching allows more users to use the network!

- 1 Mbit link
- each user:
  - 100Kbps when "active"
  - active 10% of time (bursty behaviour)
- circuit-switching:
  - 10 users
- packet switching:
  - with 35 users, probability > 10 active less than 0.0004 (⇒ almost all of the time same queuing behaviour as circuit switching)

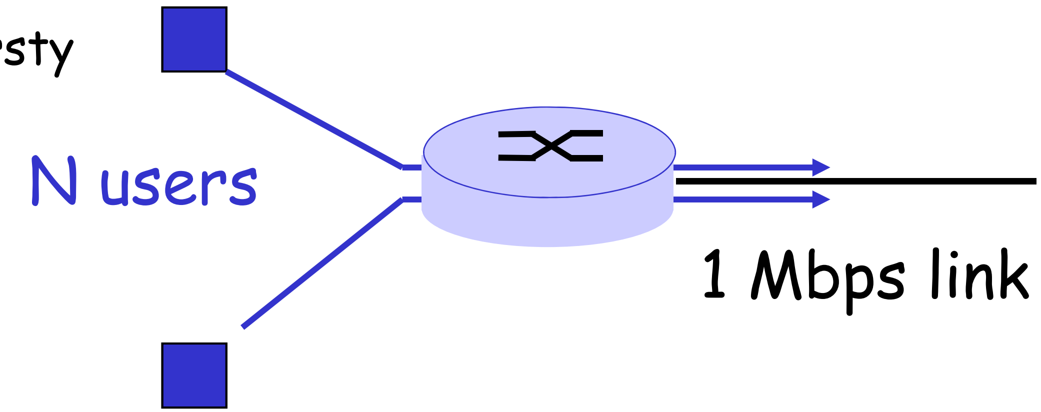

Queueing delay (revisited) ...

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

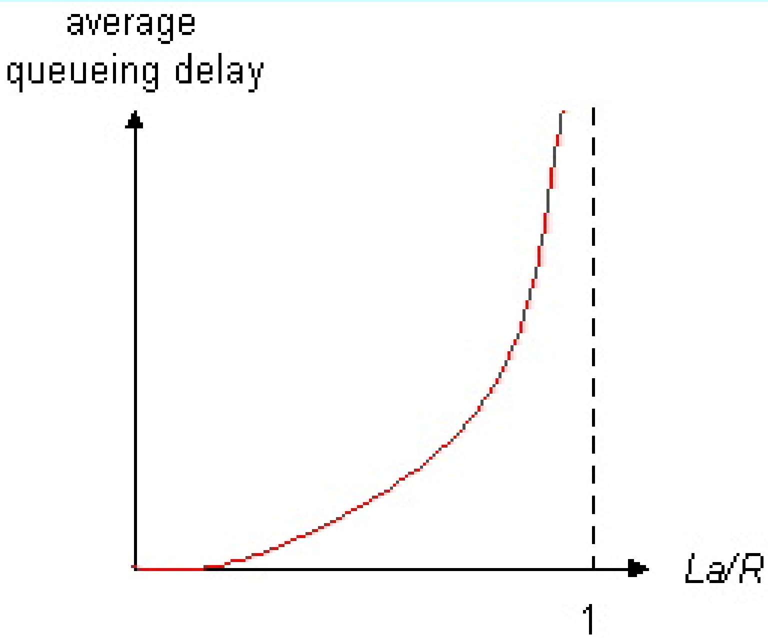

traffic intensity = La/R

- 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! Queues may grow unlimited, packets can be lost

… "Real" Internet delays and routes (1)…

- What do "real" Internet delay & loss look like?
- Traceroute program: provides delay measurement from source to router along end-end Internet path towards destination. For all *i:*
  - sends three packets that will reach router *i* on path towards destination
  - router *i* will return packets to sender
  - sender times interval between transmission and reply.

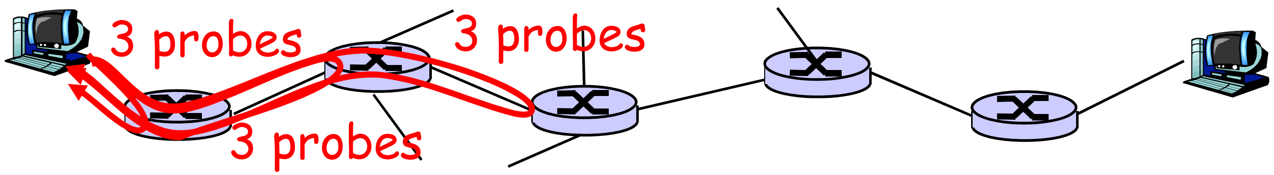

…"Real" Internet delays and routes (2)…

traceroute: gaia.cs.umass.edu to www.eurecom.fr

Three delay measurements from gaia.cs.umass.edu to cs-gw.cs.umass.edu

1 cs-gw (128.119.240.254) 1 ms 1 ms 2 ms
2 border1-rt-fa5-1-0.gw.umass.edu (128.119.3.145) 1 ms 1 ms 2 ms
3 cht-vbns.gw.umass.edu (128.119.3.130) 6 ms 5 ms 5 ms
4 jn1-at1-0-0-19.wor.vbns.net (204.147.132.129) 16 ms 11 ms 13 ms
5 jn1-so7-0-0-0.wae.vbns.net (204.147.136.136) 21 ms 18 ms 18 ms
6 abilene-vbns.abilene.ucaid.edu (198.32.11.9) 22 ms 18 ms 22 ms
7 nycm-wash.abilene.ucaid.edu (198.32.8.46) 22 ms 22 ms 22 ms
8 62.40.103.253 (62.40.103.253) 104 ms 109 ms 106 ms
9 de2-1.de1.de.geant.net (62.40.96.129) 109 ms 102 ms 104 ms
10 de.fr1.fr.geant.net (62.40.96.50) 113 ms 121 ms 114 ms
11 renater-gw.fr1.fr.geant.net (62.40.103.54) 112 ms 114 ms 112 ms
12 nio-n2.cssi.renater.fr (193.51.206.13) 111 ms 114 ms 116 ms
13 nice.cssi.renater.fr (195.220.98.102) 123 ms 125 ms 124 ms
14 r3t2-nice.cssi.renater.fr (195.220.98.110) 126 ms 126 ms 124 ms
15 eurecom-valbonne.r3t2.ft.net (193.48.50.54) 135 ms 128 ms 133 ms
16 194.214.211.25 (194.214.211.25) 126 ms 128 ms 126 ms
17 * * *
18 * * *
19 fantasia.eurecom.fr (193.55.113.142) 132 ms 128 ms 136 ms

trans-oceanic link

* means no reponse (probe lost, router not replying)

Packet switching properties

- Good: Great for bursty data
  - resource sharing
  - no call setup
- Not so good: Excessive congestion: packet delay and loss
  - protocols needed for reliable data transfer, congestion control
- http://www.youtube.com/watch?v=Dq1zpiDN9k4&feature=related
- Q: How to provide circuit-like behavior?
  - bandwidth guarantees needed for audio/video apps
  - Some routing policies can help (cf next slide)

Packet-switched networks: routing

- *Goal:* move packets among routers from source to destination
  - Challenge 1: path selection algorithms
  - Challenge2: Important design issue:
    - **datagram network**:
      - *destination address* determines next hop
      - routes may change during session
    - **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

Virtual circuits:

"source-to-dest path behaves almost like telephone circuit"

- call setup, teardown for each call *before* data can flow
  - signaling protocols to setup, maintain teardown VC (ATM, frame-relay, X.25; not in IP)
- each packet carries VC identifier (not destination host)
- *every* router maintains "state" for each passing connection
- resources (bandwidth, buffers) may be *allocated* to VC

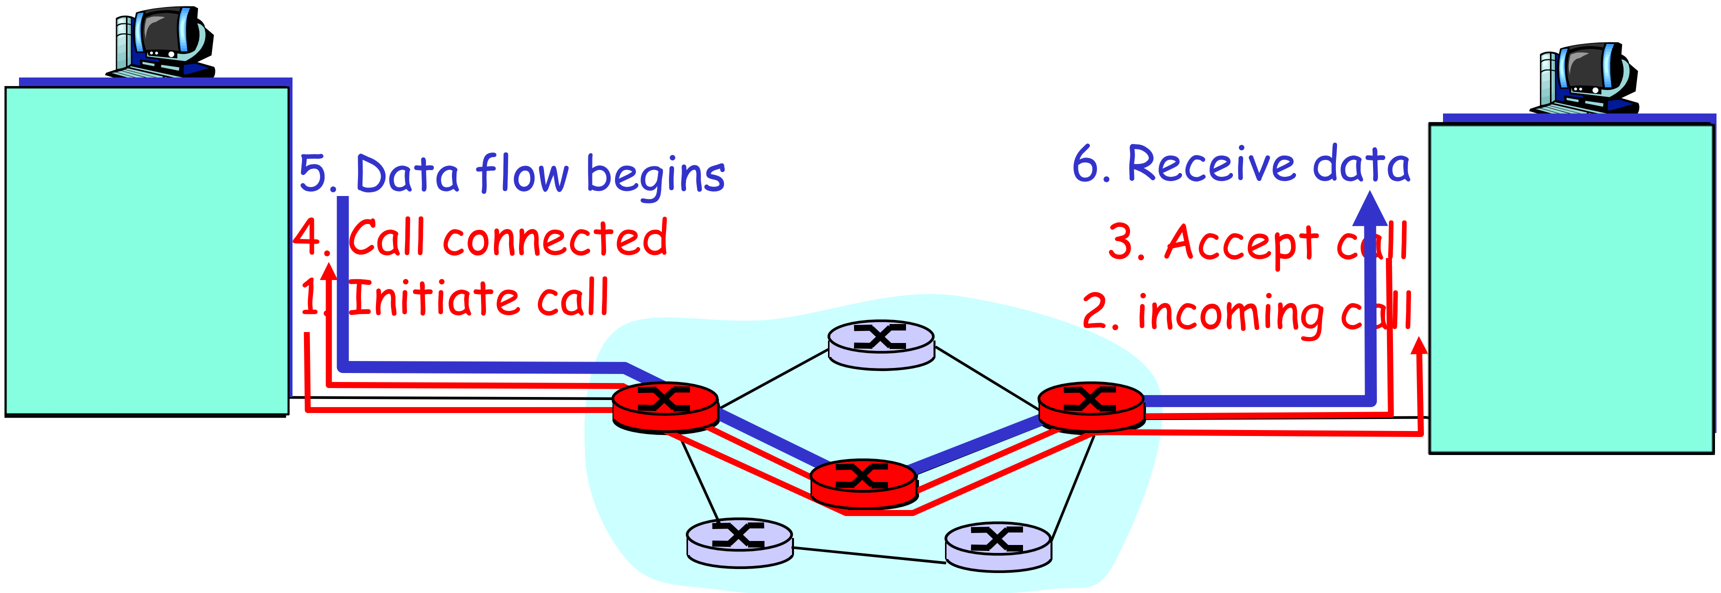

Network Taxonomy

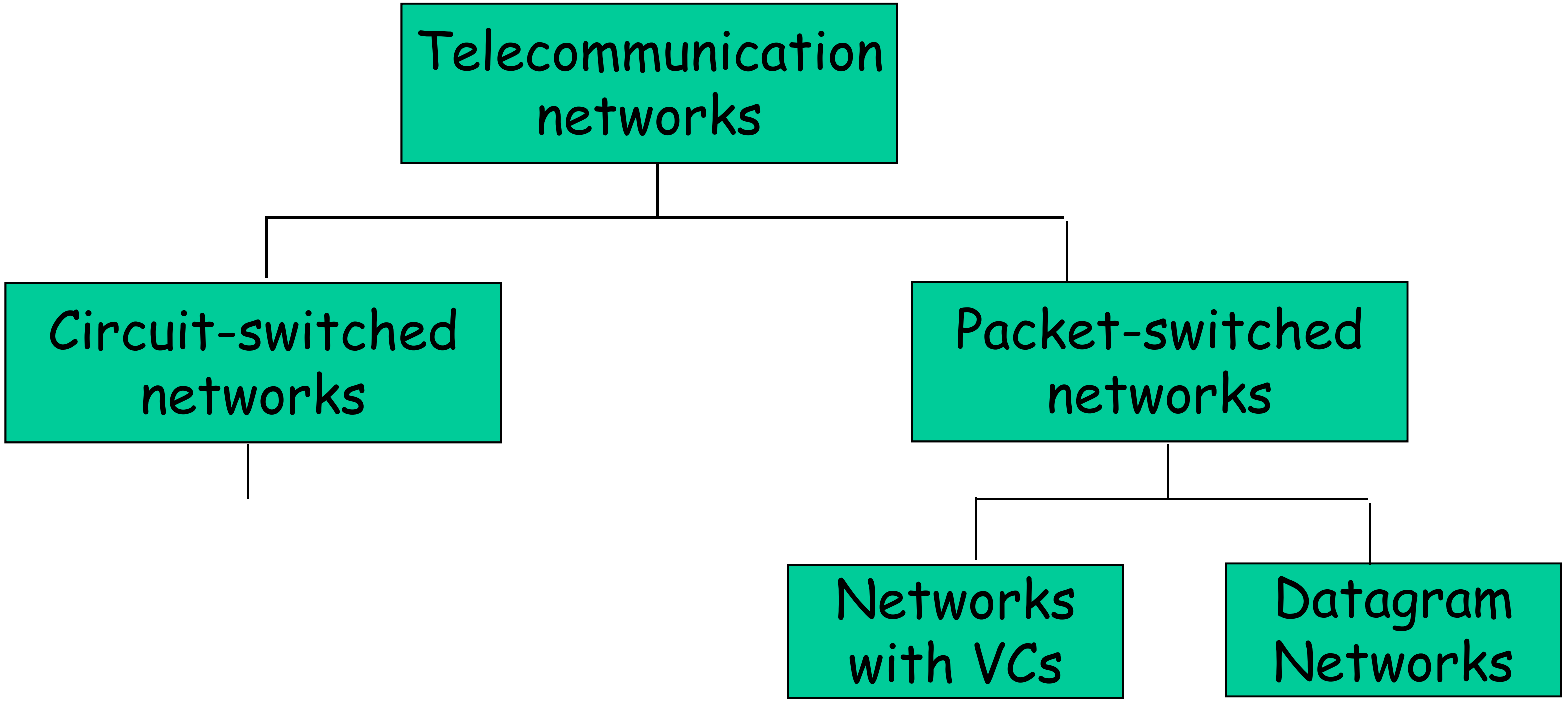

- Datagram network cannot be charecterized either connection-oriented or connectionless.
  - Internet provides both connection-oriented (TCP) and connectionless services (UDP) to apps.

Packet loss

- queue (aka buffer) preceding link has finite capacity
- packet arriving to full queue dropped (aka lost)
- lost packet may be retransmitted by previous node, by source end system, or not at all

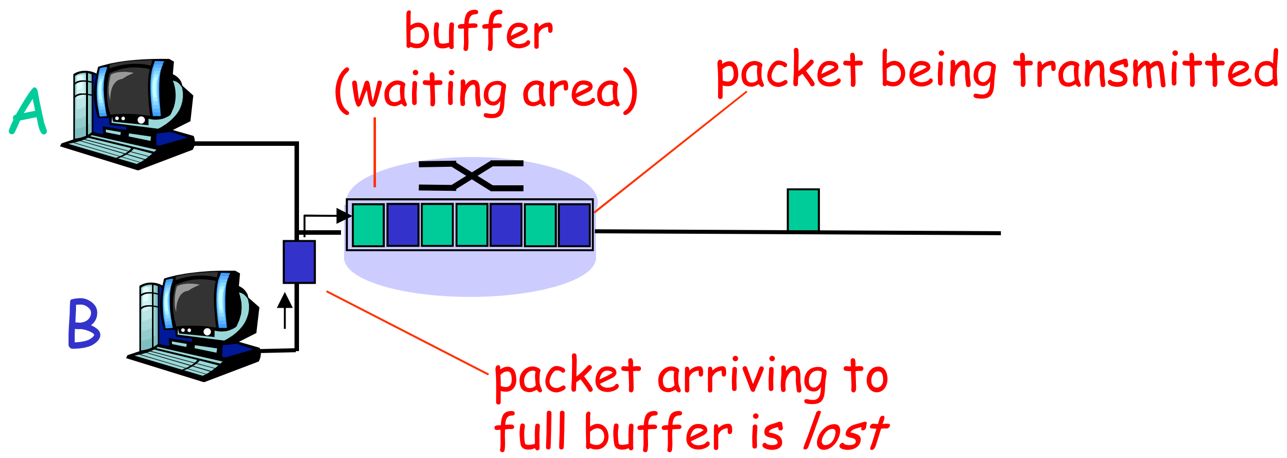

Throughput

- *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

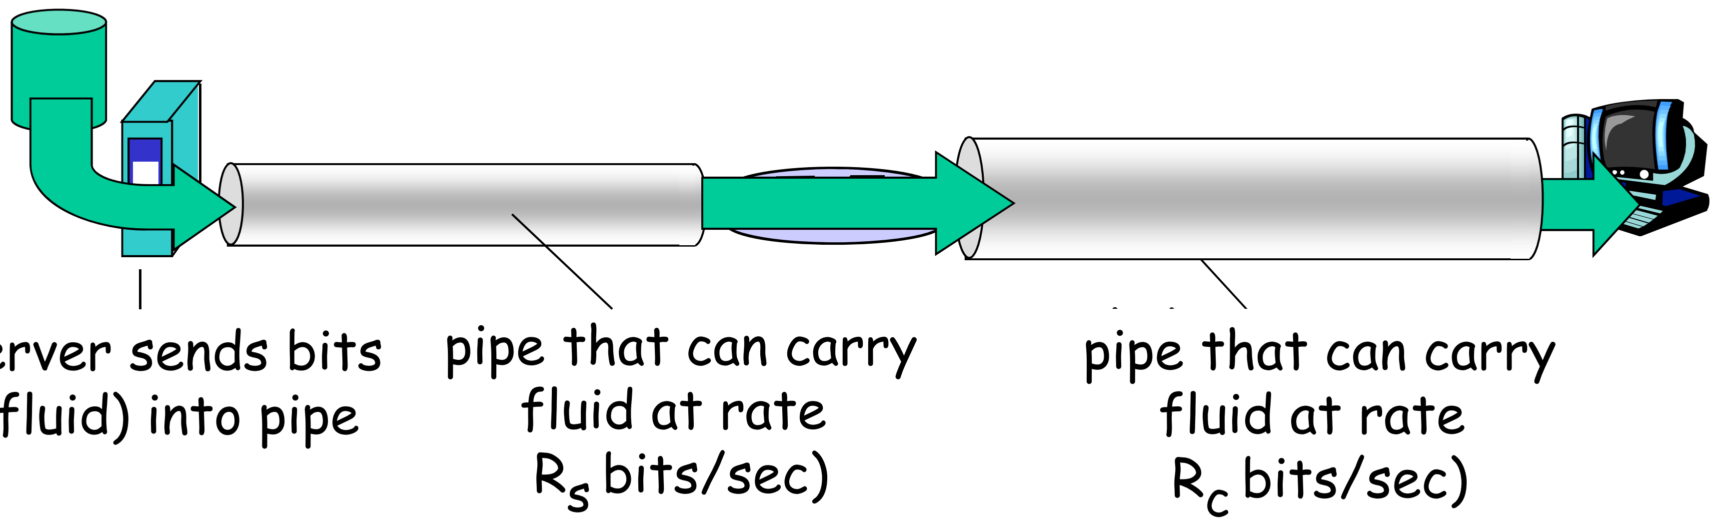

Throughput (more)

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

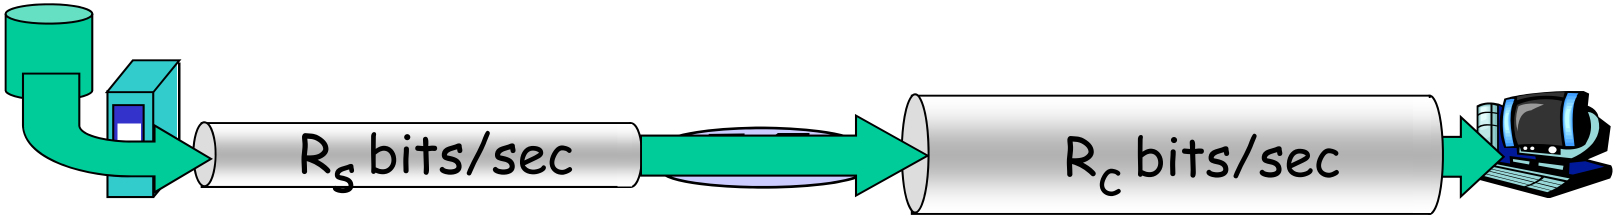

- $R_s > R_c$ What is average end-end throughput?

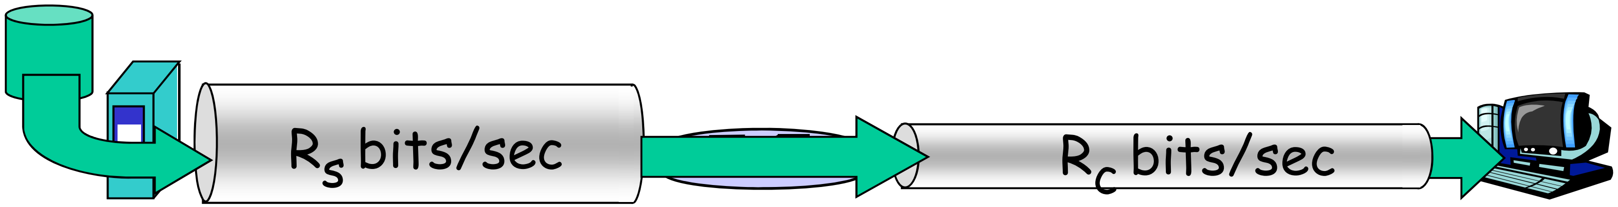

bottleneck link

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 (if fair))
- in practice: R_c or R_s is often bottleneck

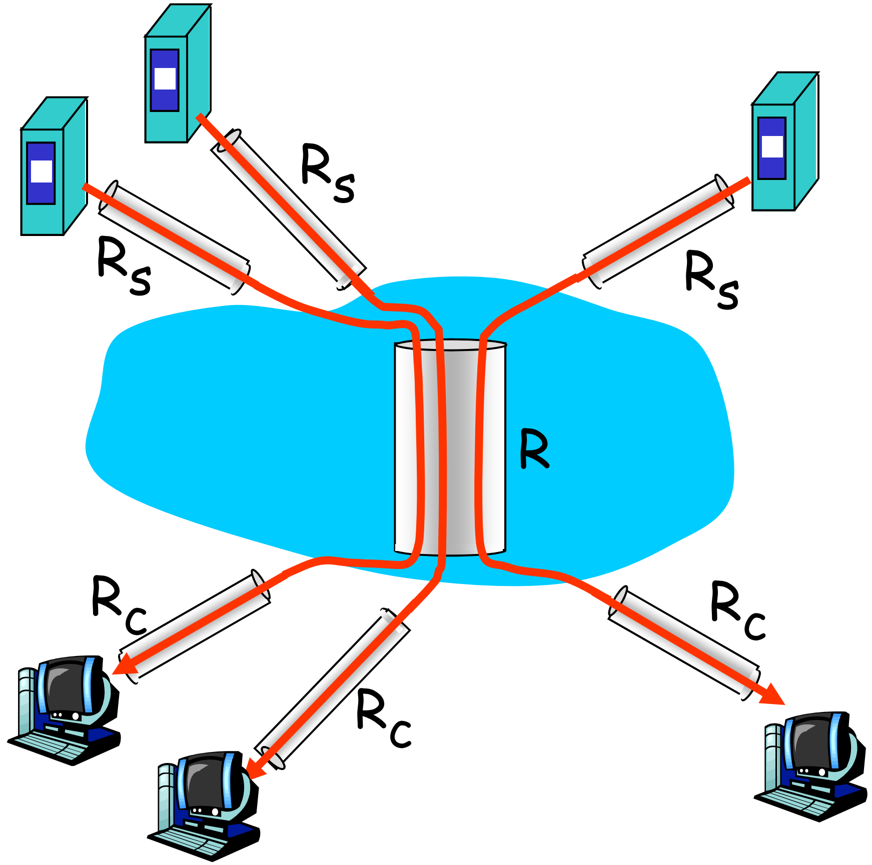

10 connections (fairly) share backbone bottleneck link R bits/sec

Access networks and physical media

Access networks and physical media

Q: How to connect end systems to edge router?

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

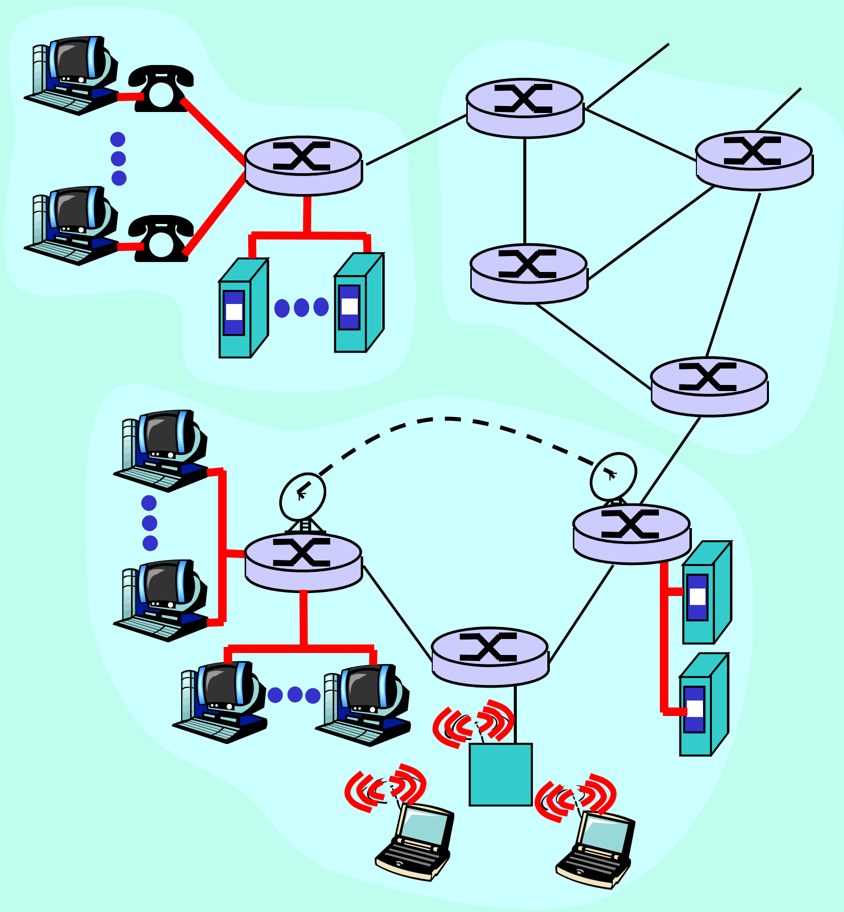

Keep in mind:

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

Dial-up Modem

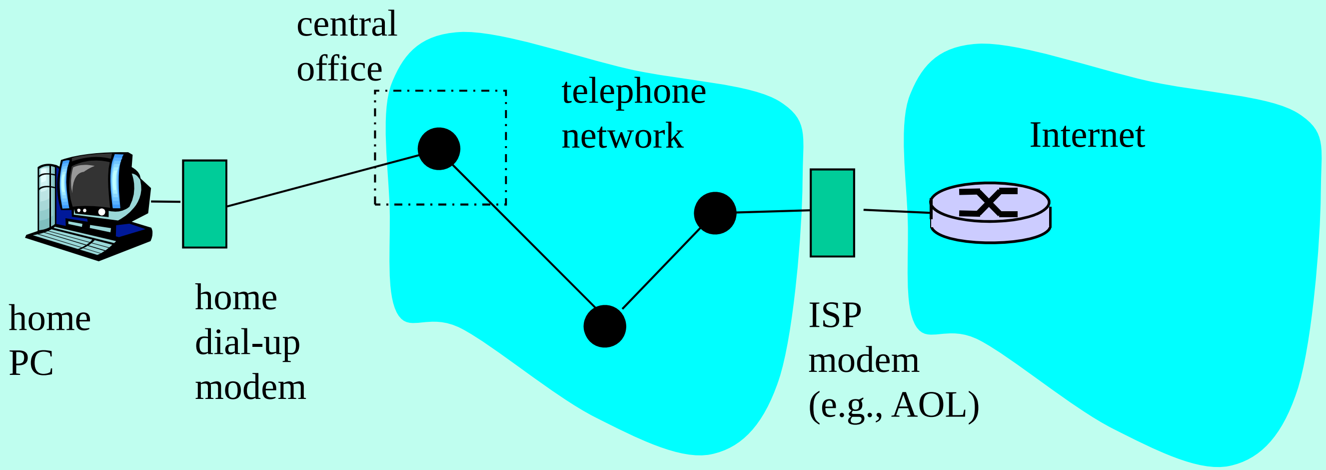

- Uses existing telephony infrastructure
  - Home is connected to central office
- up to 56Kbps direct access to router (often less)
- Can't surf and phone at same time: not "always on"

Digital Subscriber Line (DSL)

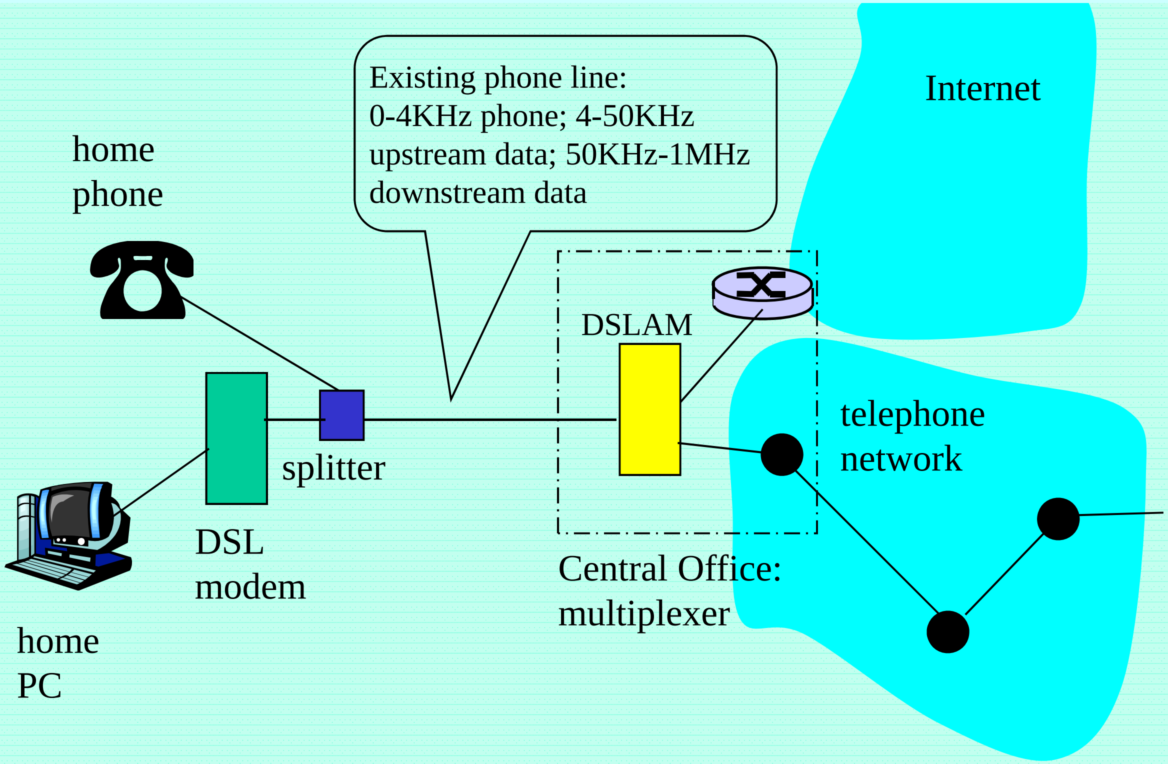

- Also uses existing telephone infrastruture
- Commonly up to 1 Mbps upstream (more typically < 256 kbps)
- Commonly up to 8 Mbps downstream (more typically < 1 Mbps)
- dedicated physical line to telephone central office

Residential access: cable modems

- Instead uses cable TV infrastructure
- HFC: hybrid fiber coax
  - asymmetric: <30Mbps downstream, 2 Mbps upstream
- network of cable and fiber attaches homes to ISP router
  - homes share access to router
  - unlike DSL, which has dedicated access

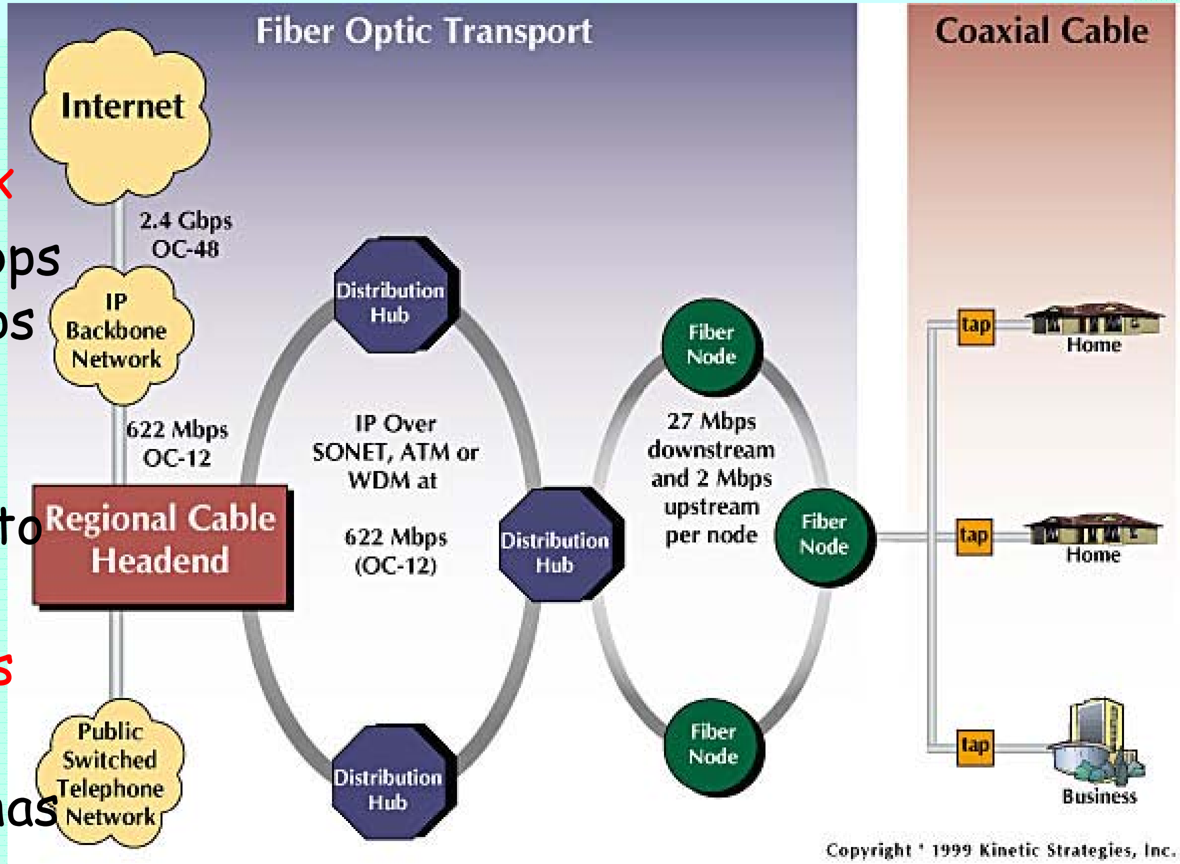

Diagram: http://www.cabledatacomnews.com/cmic/diagram.html

Cable Network Architecture: Overview

Typically 500 to 5,000 homes

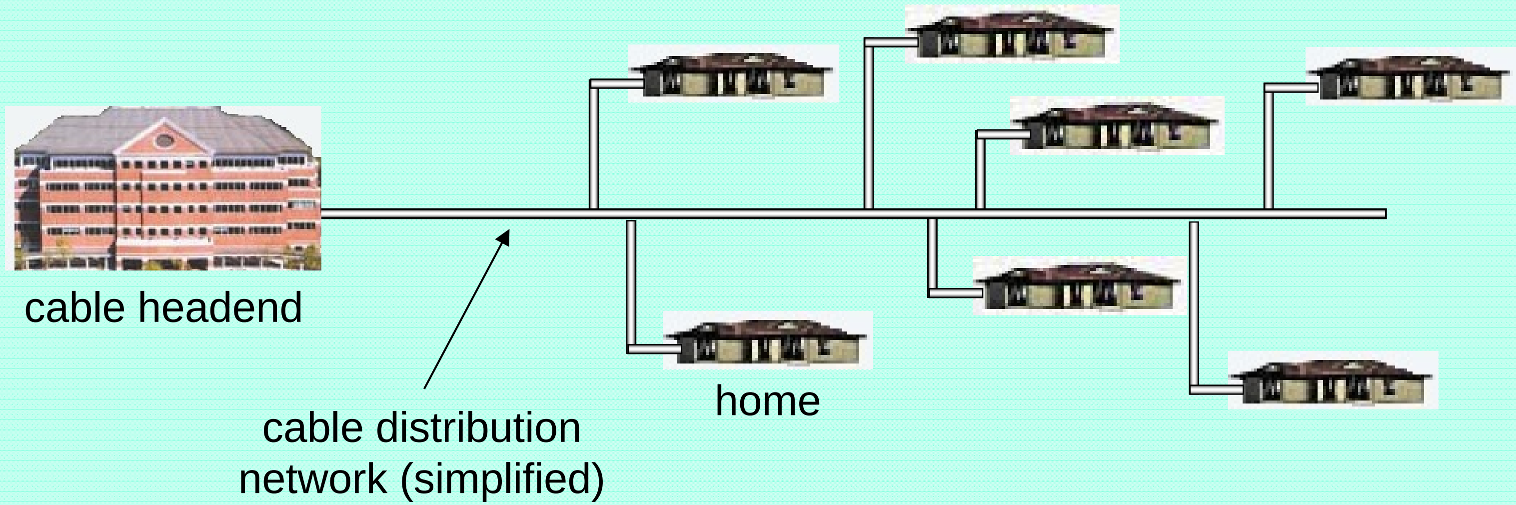

Cable Network Architecture: Overview

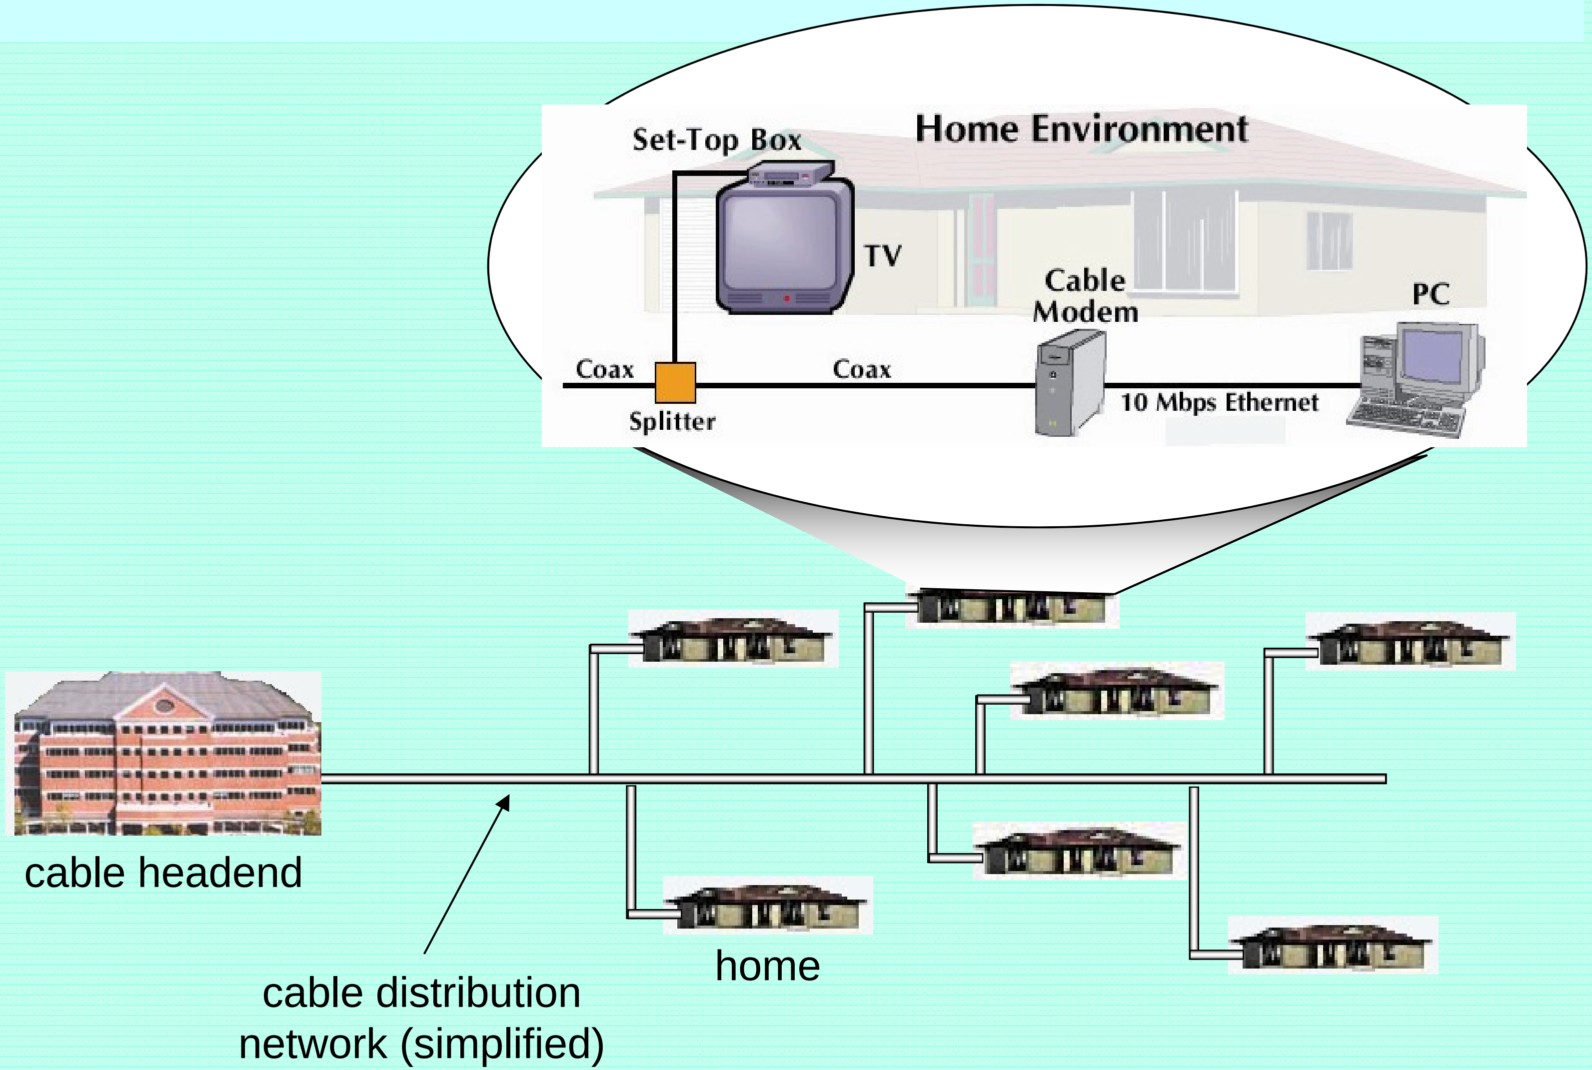

Cable Network Architecture: Overview

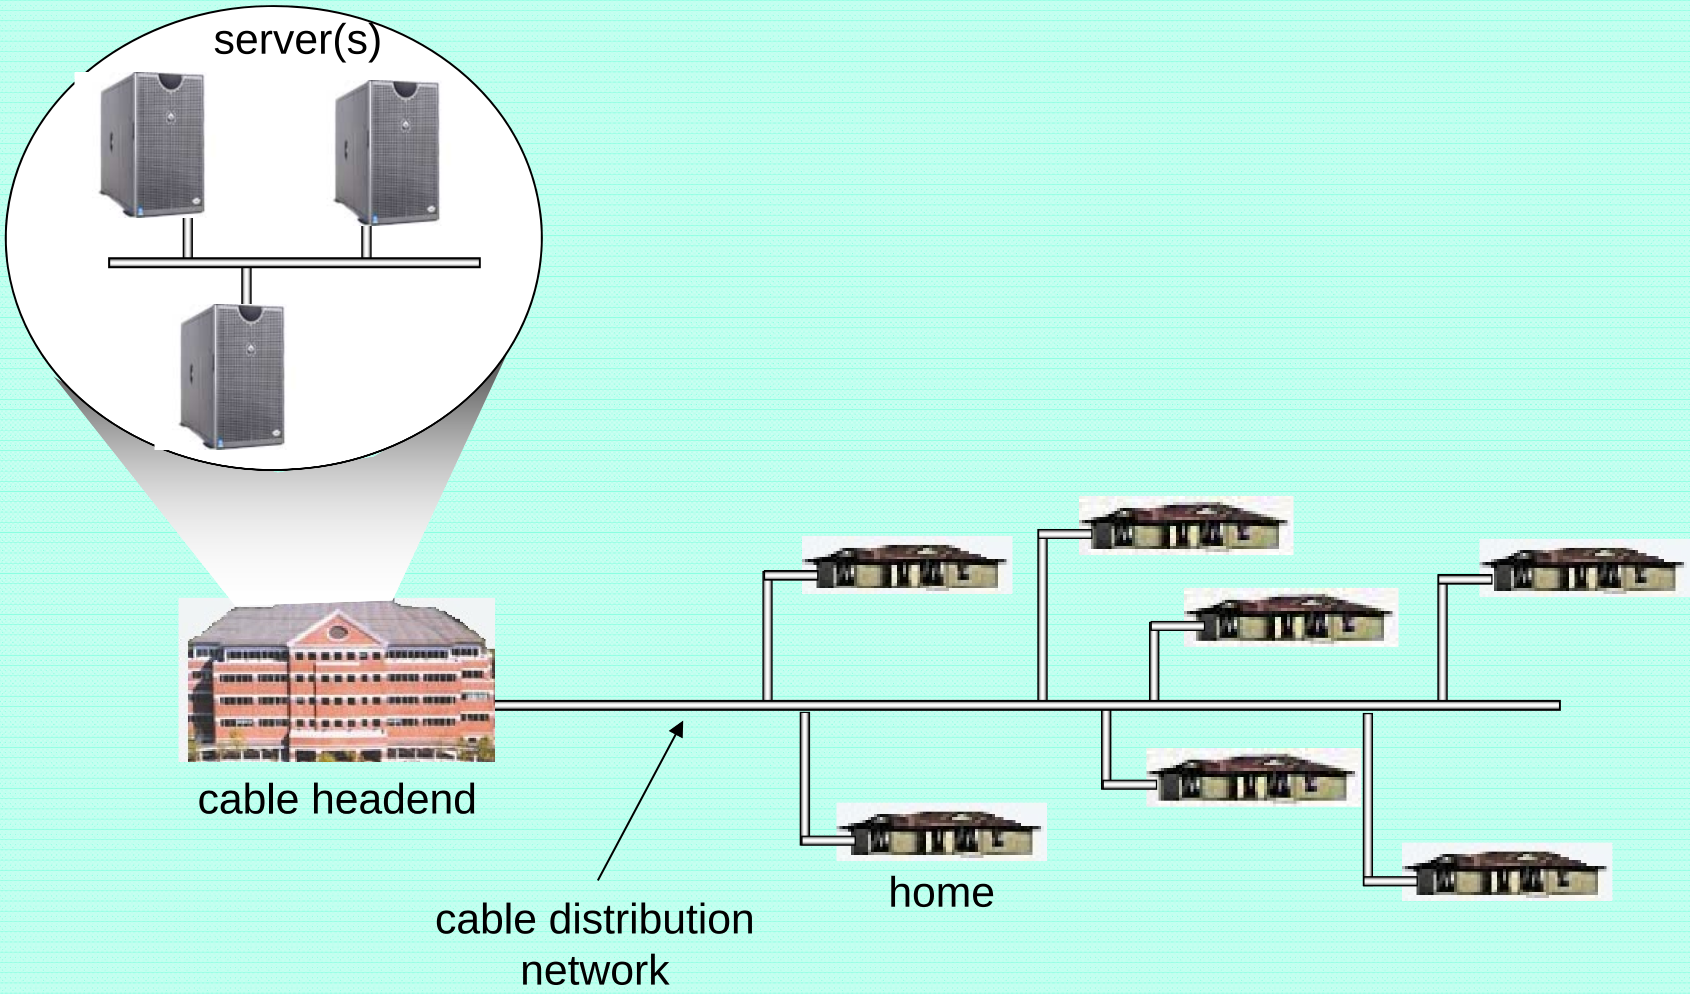

Cable Network Architecture: Overview

FDM:

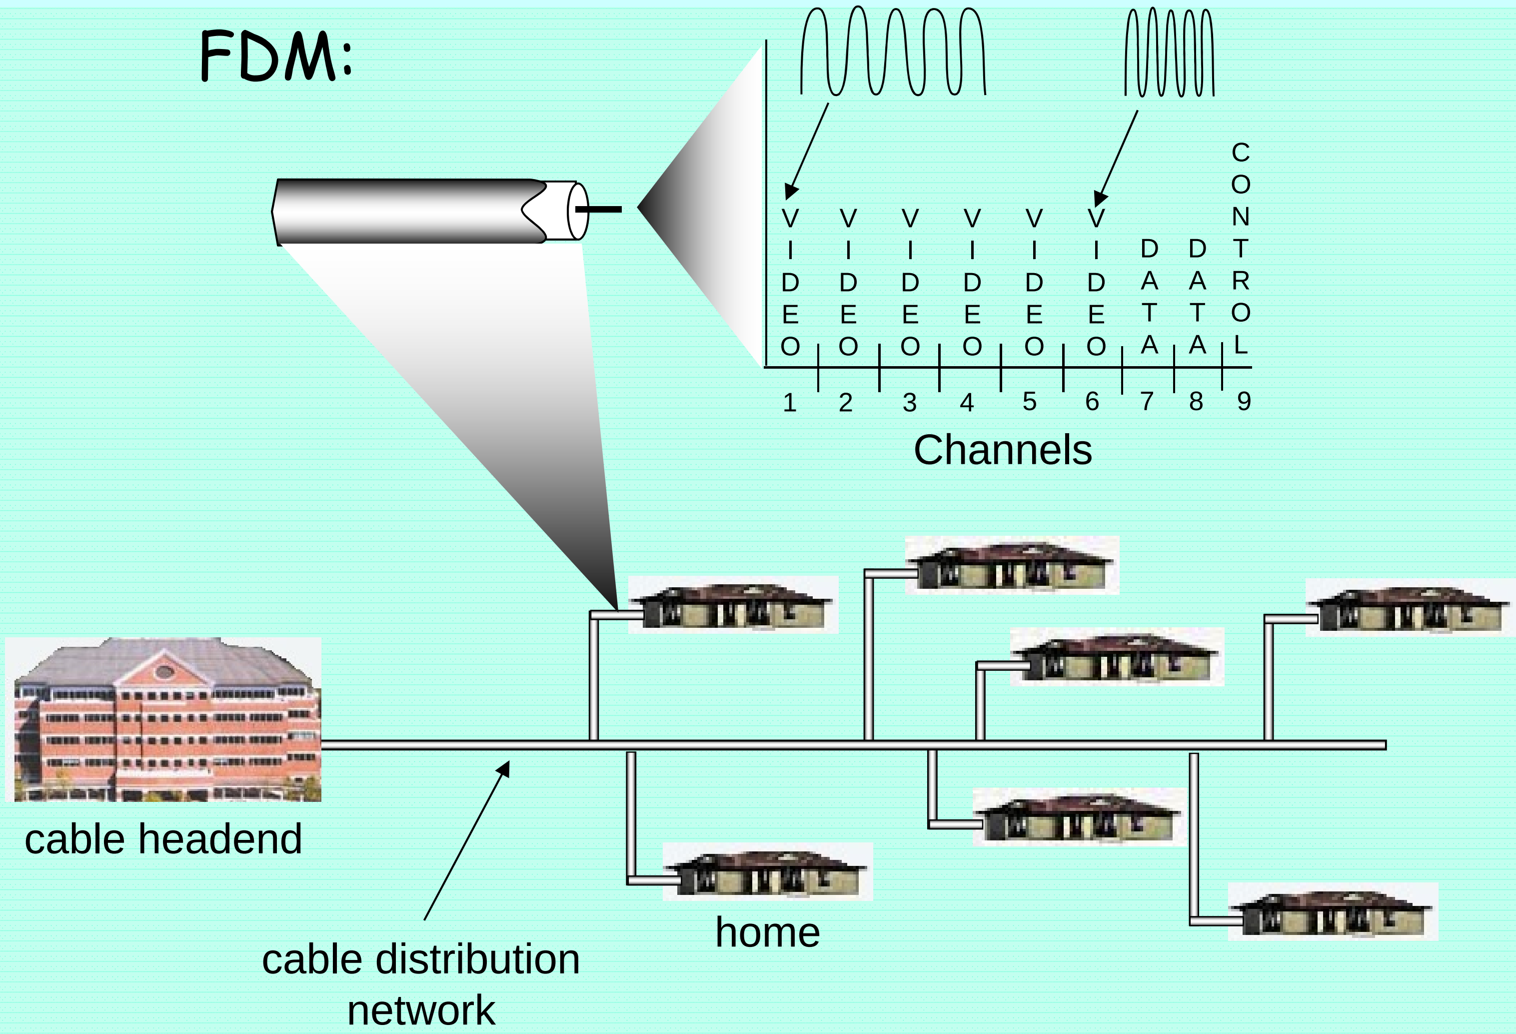

Fiber to the Home

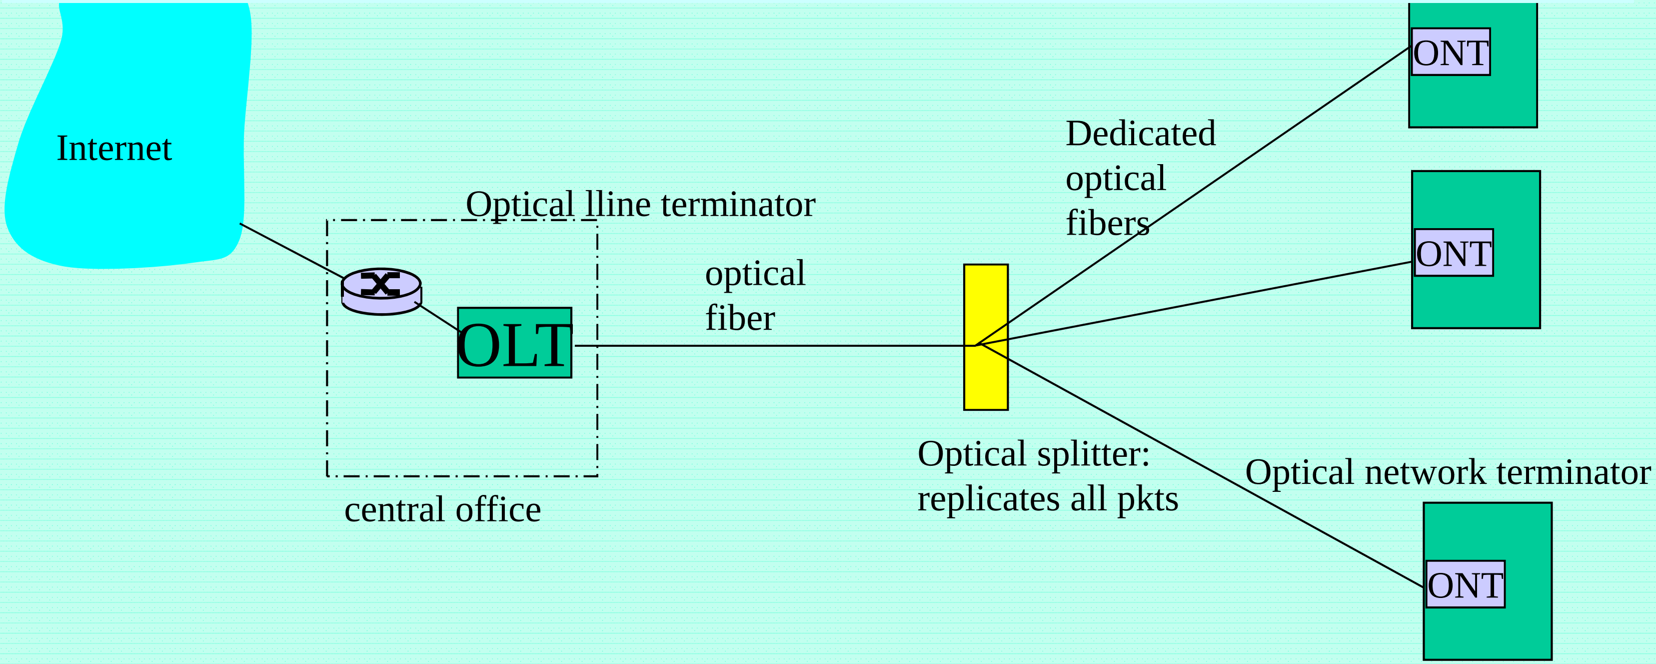

- Optical links from central office to the home
- Two competing optical technologies:
  - Passive Optical network (PON:)
  - Active Optical Network (AON: essentially switched Ethernet, as in institutional access -next)
- Much higher Internet rates; fiber also carries television and phone services

Institutional access: local area networks

- company/univ local area network (LAN) connects end system to edge router
- E.g. Ethernet:
  - shared or dedicated cable connects end system and router (usually switched now)
  - 10 Mbs, 100Mbps, Gigabit Ethernet
- deployment: institutions, home LANs

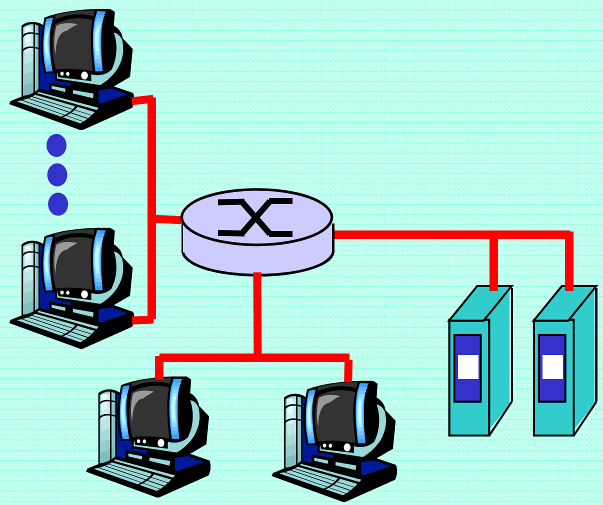

Wireless access networks

- shared *wireless* access network connects end system to router
  - via base station aka "access point"
- wireless LANs:
  - 802.11b/g (WiFi): 11 or 54 Mbps
- wider-area wireless access
  - provided by telco operator
  - ~1Mbps over cellular system
  - next up (?): WiMAX (10's Mbps) over wide area

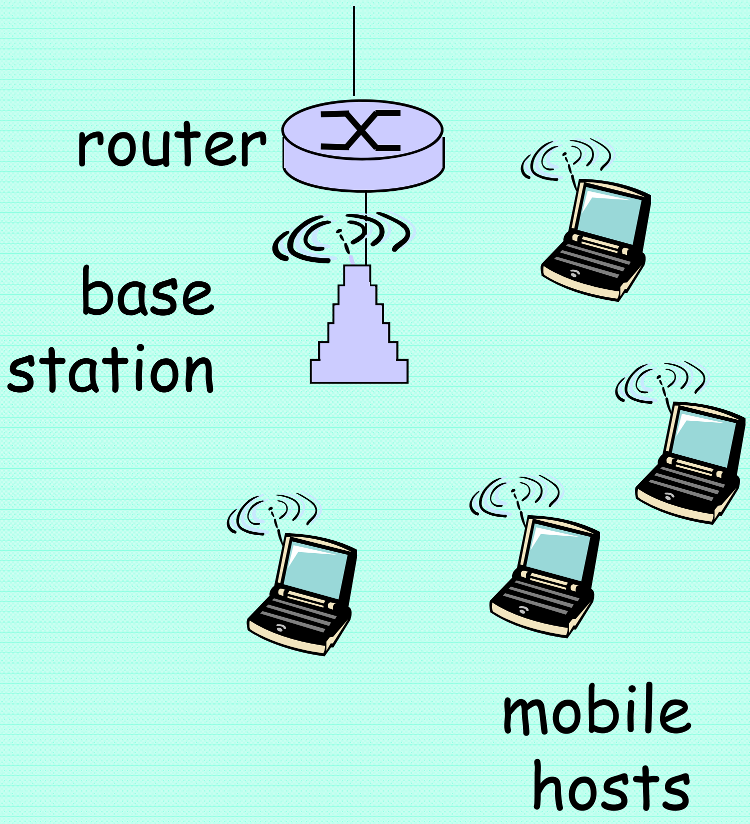

Home networks

Typical home network components:

- DSL or cable modem
- router/firewall/NAT
- Ethernet
- wireless access point

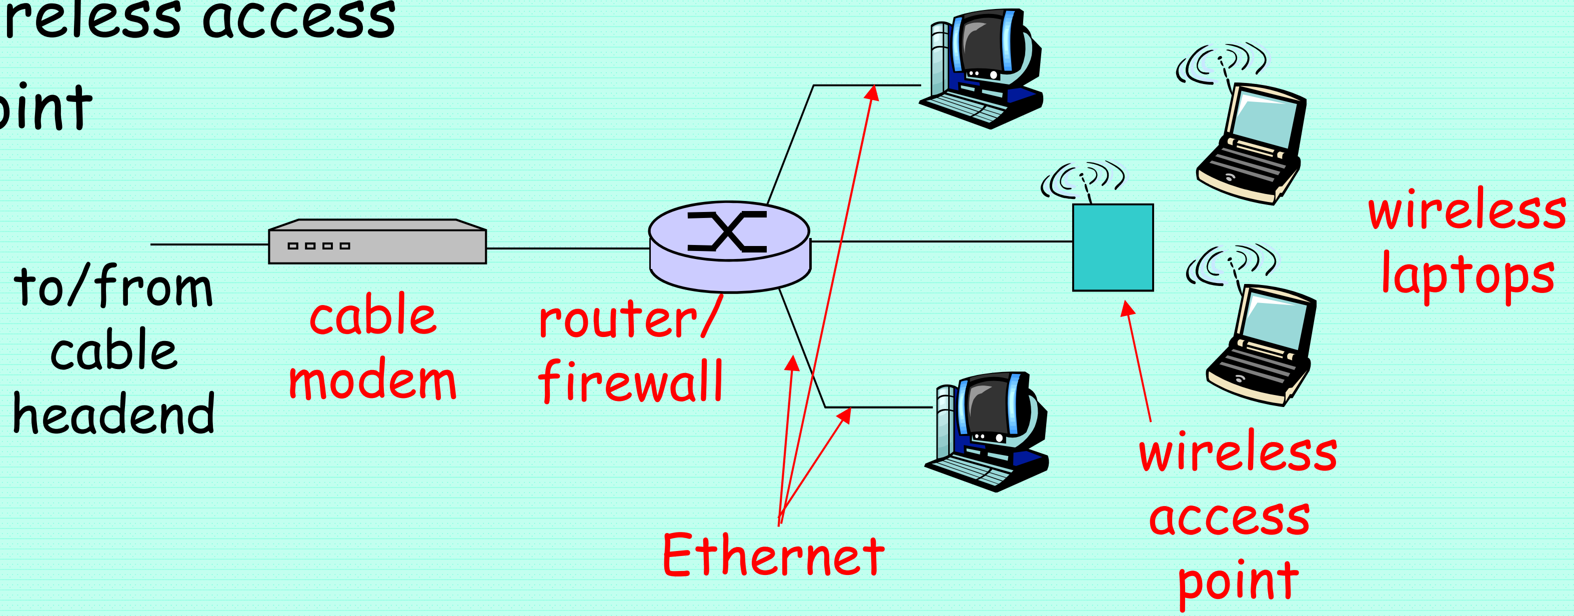

Physical Media

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

Physical media: wireless

- signal carried in electromagnetic spectrum
- Omnidirectional: signal spreads, can be received by many antennas
- Directional: antennas communicate with focused el-magnetic beams and must be aligned (requires higher frequency ranges)
- propagation environment effects:
  - reflection
  - obstruction by objects
  - interference

Properties: Attenuation, Multipath propagation

Signal can fade with distance, can get obstructed, can take many different paths between sender and receiver due to reflection, scattering, diffraction

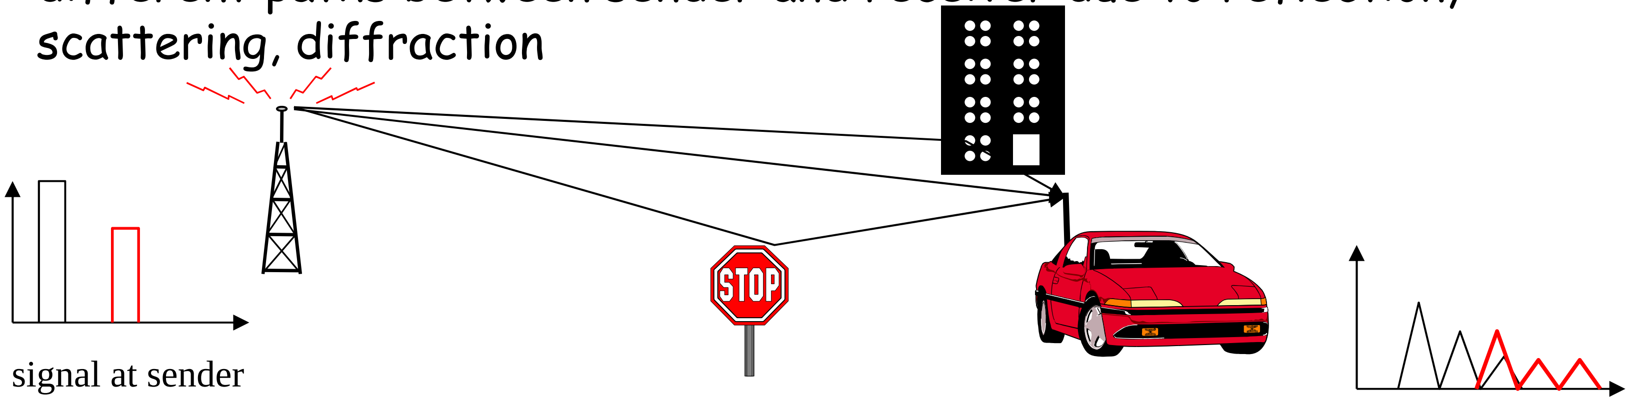

signal at sender

signal at receiver

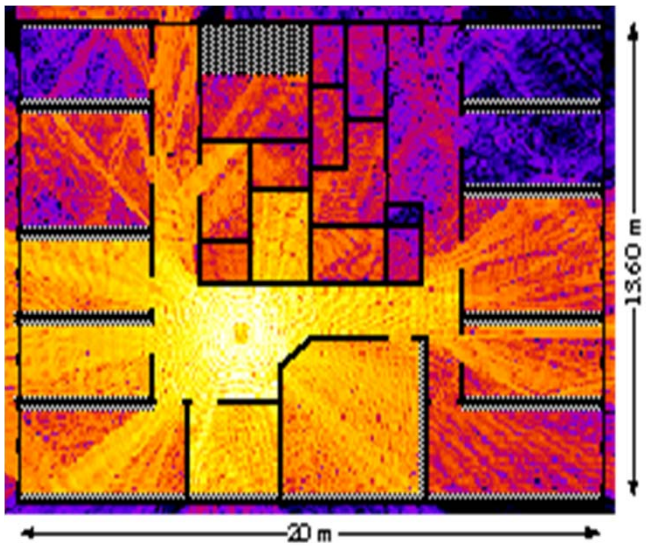

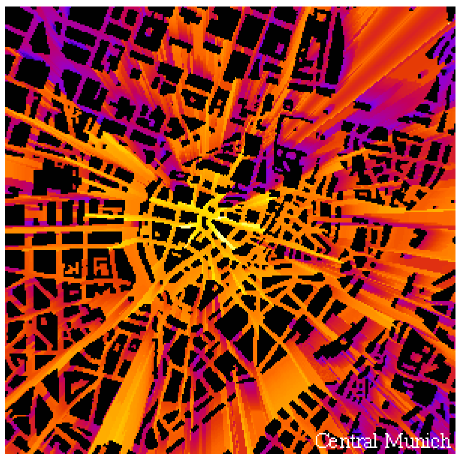

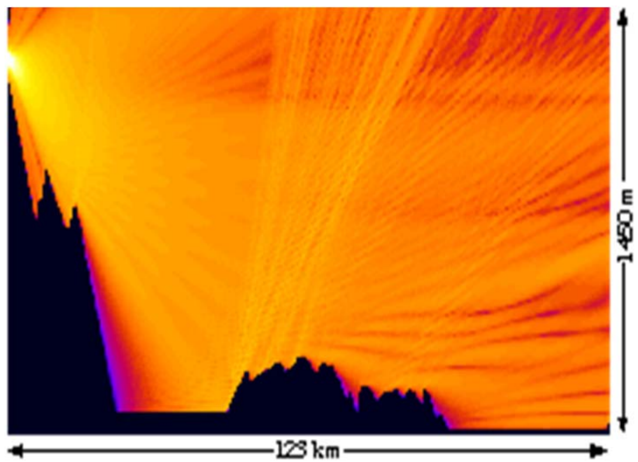

Physical Media: Twisted pair

Twisted Pair (TP)

- two insulated copper wires
  - Category 3: traditional phone wires, 10 Mbps Ethernet
  - Category 5 TP: more twists, higher insulation: 100Mbps Ethernet

Physical Media: coax, fiber

Coaxial cable:

- wire (signal carrier) within a wire (shield)
  - baseband: single channel on cable (common use in 10Mbs Ethernet)
  - broadband: multiple channels on cable (FDM; commonly used for cable TV)

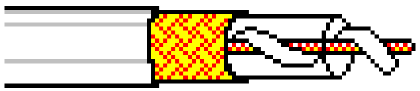

Fiber optic cable:

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

Back to Layers-discussion

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

Why layering?

Dealing with complex systems:

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

Terminology: Protocols, Interfaces

- Each layer offers **services** to the upper layers (shielding from the details how the services are implemented)
  - service interface: across layers in same host
- Layer n on a host carries a conversation with layer n on another host (data are not sent directly)
  - host-to-host interface: defines messages exchanged with peer entity
- Interfaces must be clean
  - min info exchange
  - make it simple for protocol replacements
- Network architecture (set of layers, interfaces) vs protocol stack (protocol implementation)

What's a protocol?

a human protocol and a computer network protocol:

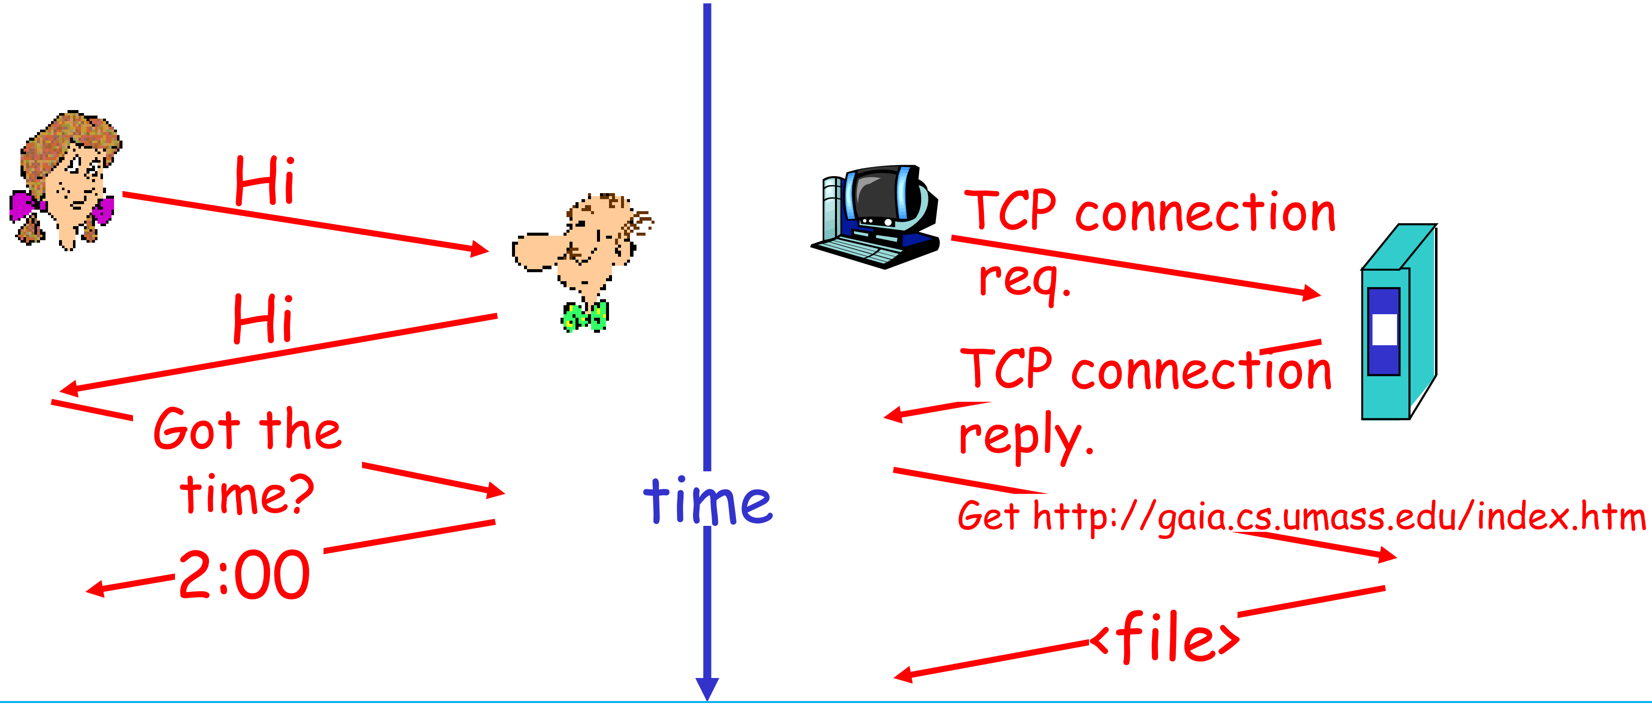

host-to-host interface: defines messages exchanged with peer entity:

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

The OSI Reference Model

- ISO (International Standards Organization) defines the OSI (Open Systems Inerconnect) model to help vendors create interoperable network implementation

- Reduce the problem into smaller and more manageable problems: 7 layers
  - a layer should be created where a different level of abstraction is needed; each layer should perform a well defined function)
  - The function of each layer should be chosen with an eye toward defining internationally standardized protocols

- ``X dot" series (X.25, X. 400, X.500) OSI model implementation (protocol stack)

Internet protocol stack

- application: ftp, smtp, http, etc
- transport: 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”

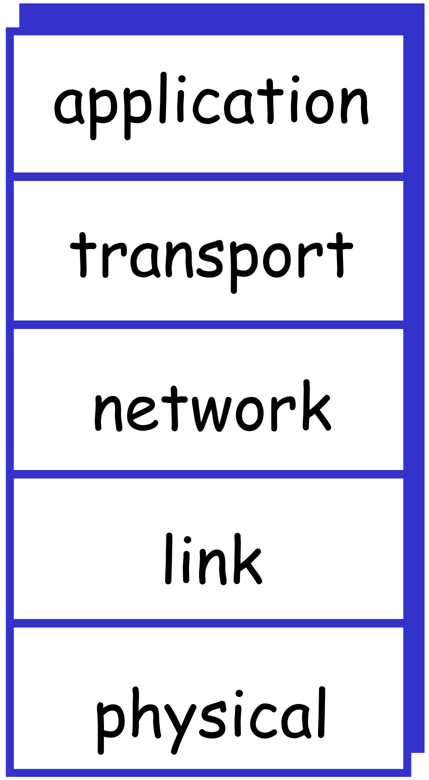

Internet protocol stack

- Architecture simple but not as good as OSI's
  - no clear distinction between interface-design and implementations;
  - hard to re-implement certain layers
- Successful protocol suite (de-facto standard)
  - was there when needed (OSI implementations were too complicated)
  - freely distributed with UNIX

Layering: logical communication

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

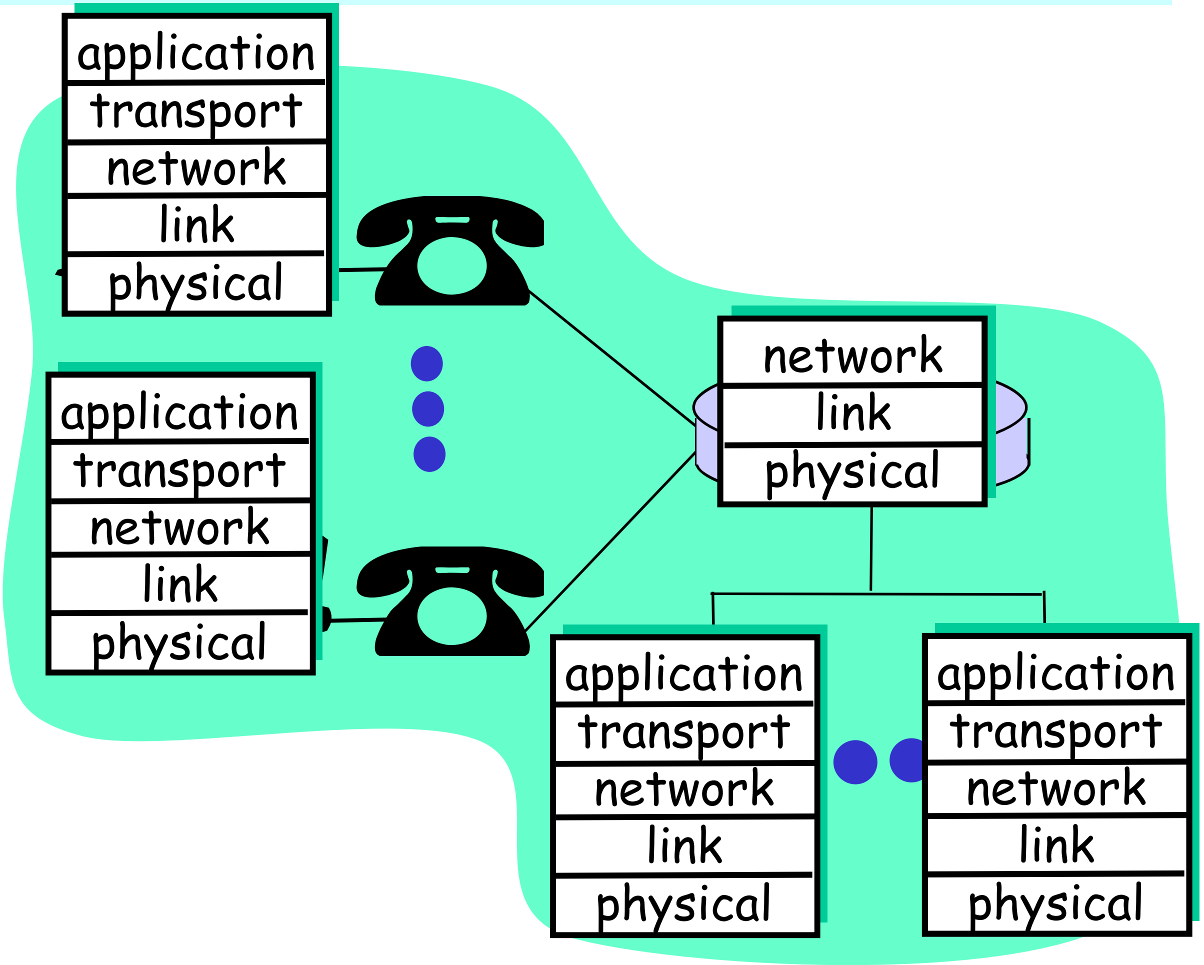

Layering: *logical* communication

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

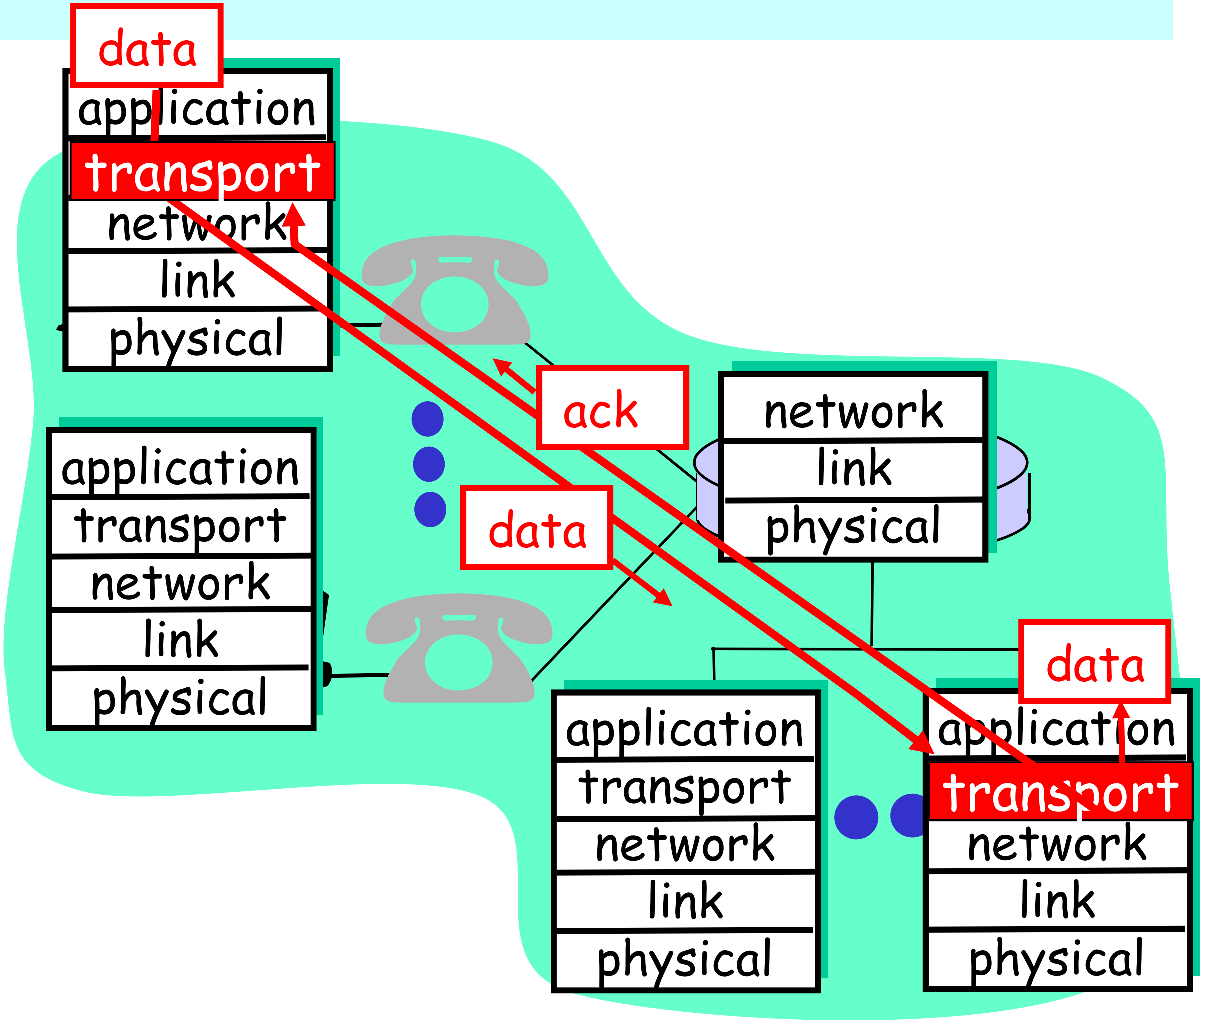

Layering: physical communication

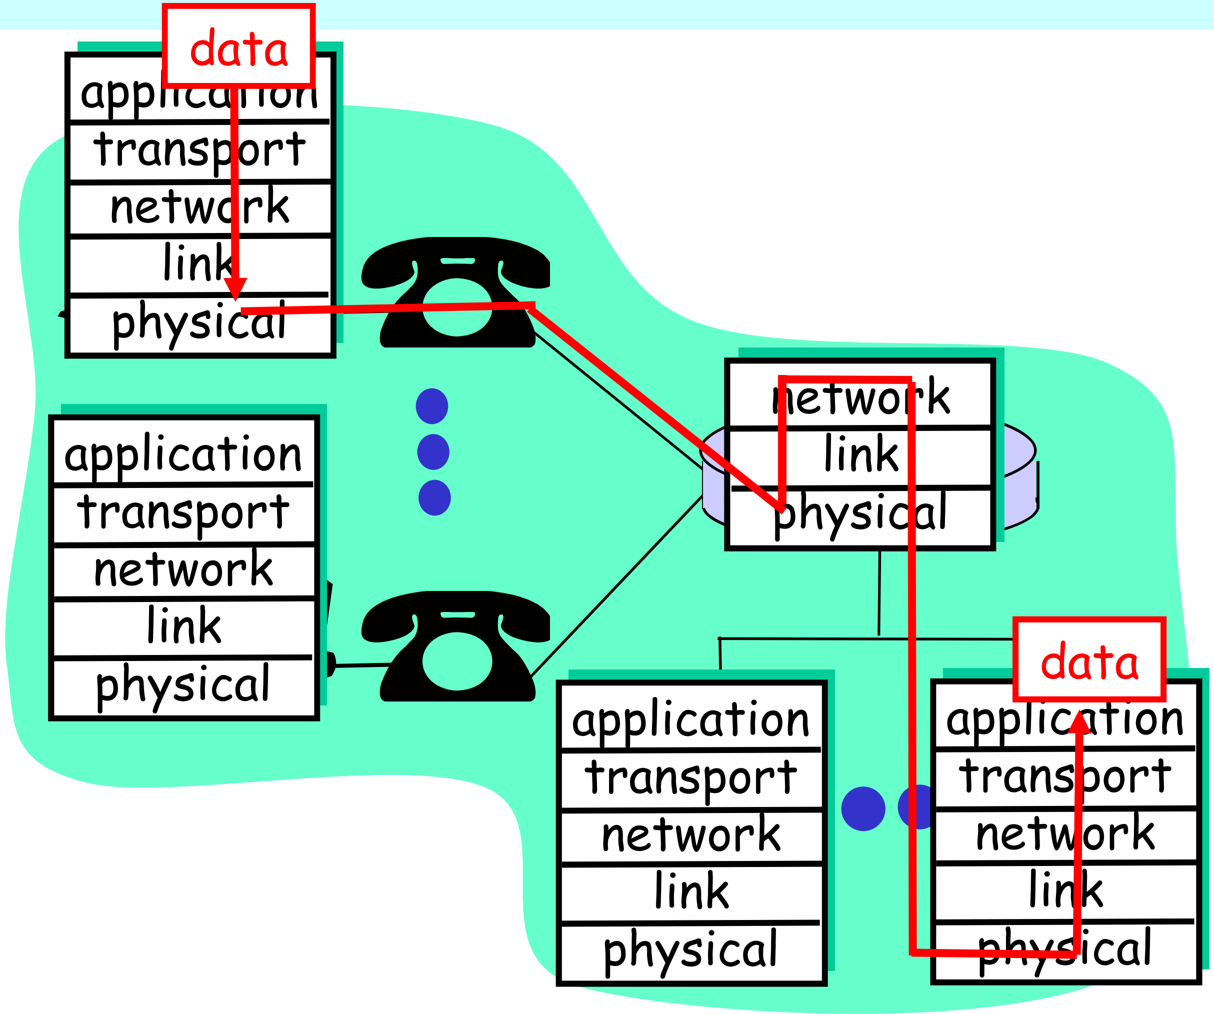

Protocol layering and data

Each layer takes data from above

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

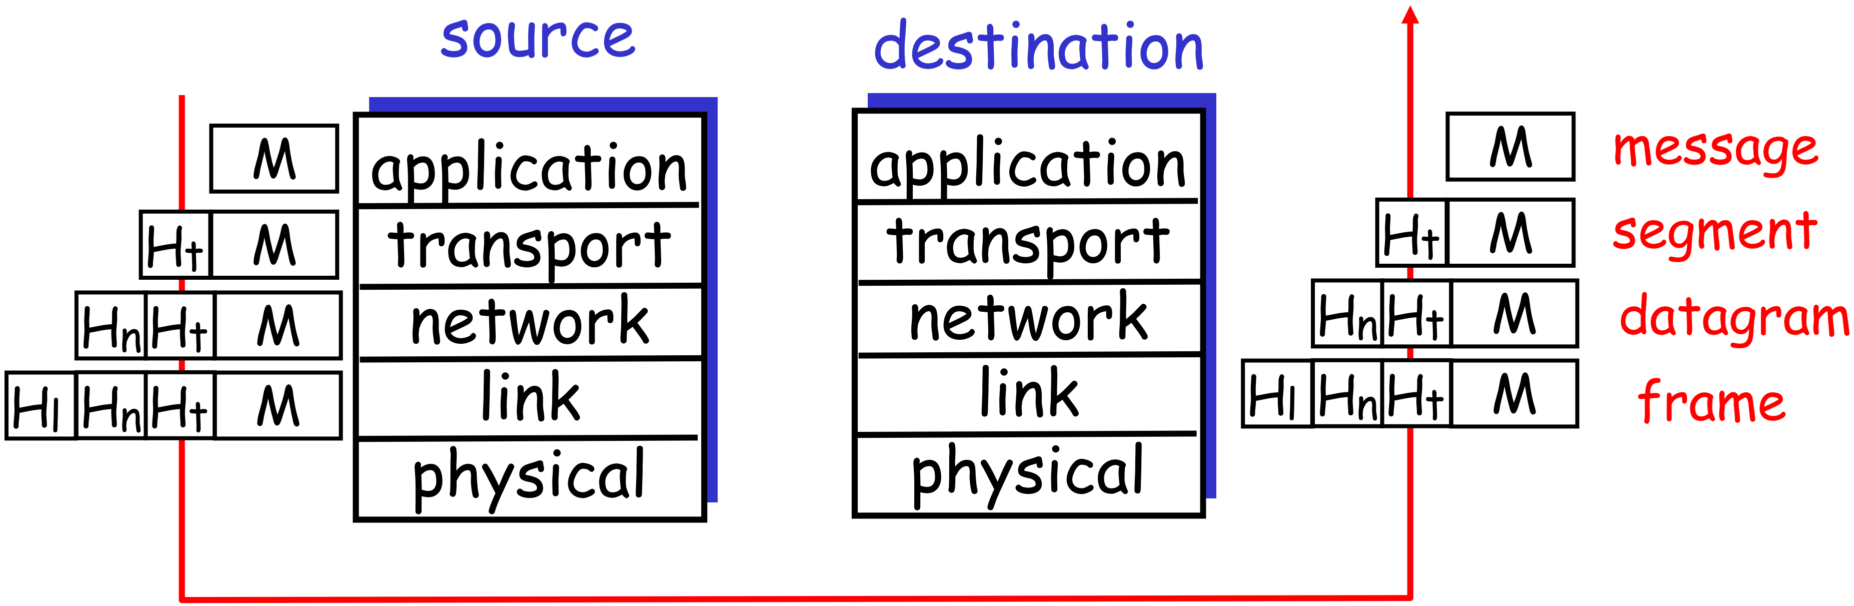

Internet structure: network of networks

- roughly hierarchical
- national/international backbone providers (NBPs)- tier 1 providers
  - e.g. BBN/GTE, Sprint, AT&T, IBM, UUNet/Verizon, TeliaSonera
  - interconnect (peer) with each other privately, or at public Network Access Point (NAPs: routers or NWs of routers)
- regional ISPs, tier 2 providers
  - connect into NBPs; e.g. Tele2
- local ISP, company
  - connect into regional ISPs, e.g. ComHem, Bredband2, Spray.se, …

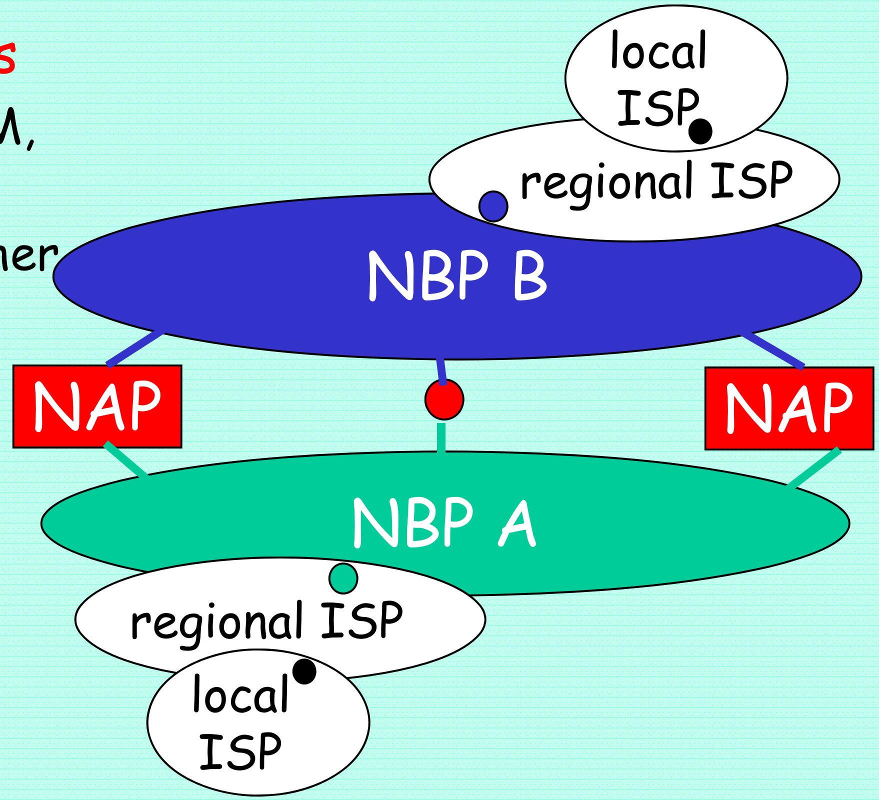

Internet structure: network of networks

- "Tier-2" ISPs: smaller (often regional) ISPs
  - Connect to one or more tier-1 ISPs, possibly other tier-2 ISPs

Tier-2 ISP pays tier-1 ISP for connectivity to rest of Internet

- tier-2 ISP is *customer* of tier-1 provider

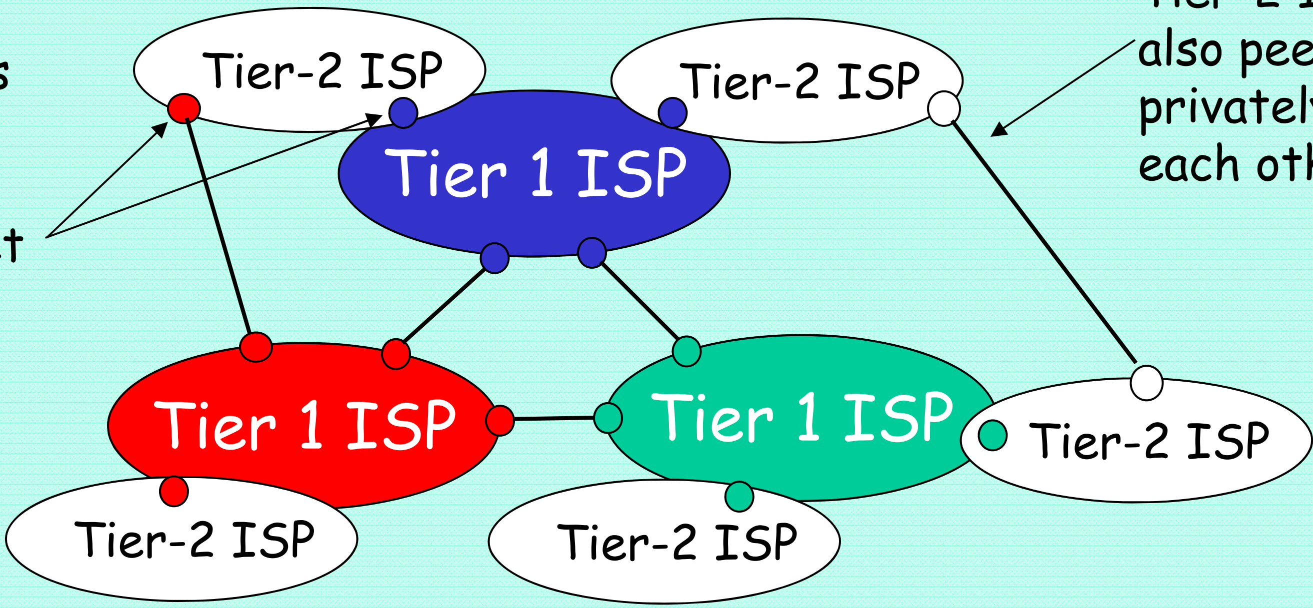

Tier-2 ISPs also peer privately with each other.

Internet structure: network of networks

- "Tier-3" ISPs and local ISPs
  - last hop ("access") network (closest to end systems)

Local and tier-3 ISPs are *customers* of higher tier ISPs connecting them to rest of Internet

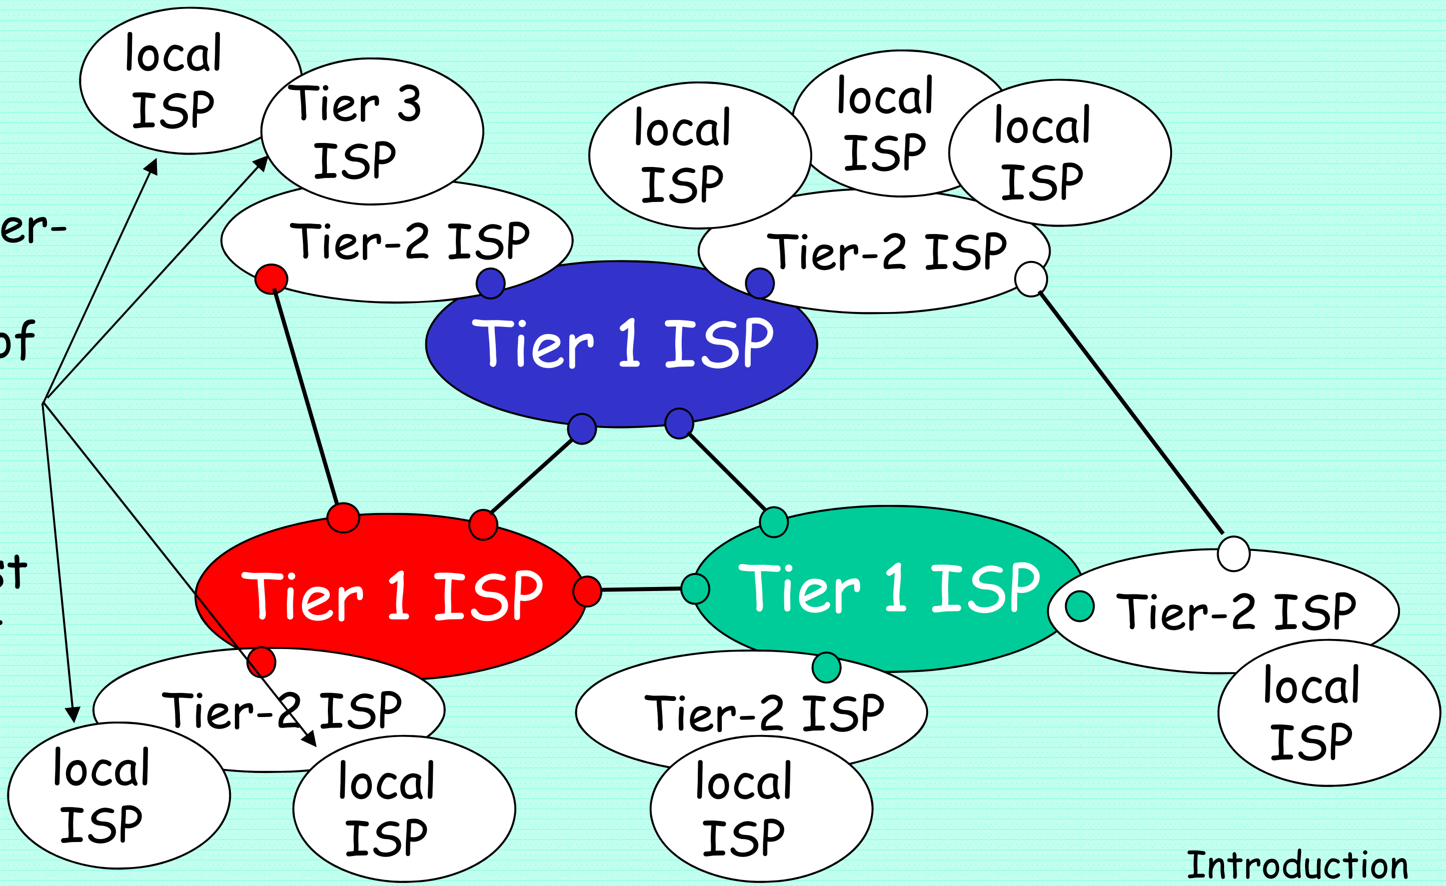

Internet structure: network of networks

- a packet passes through many networks!

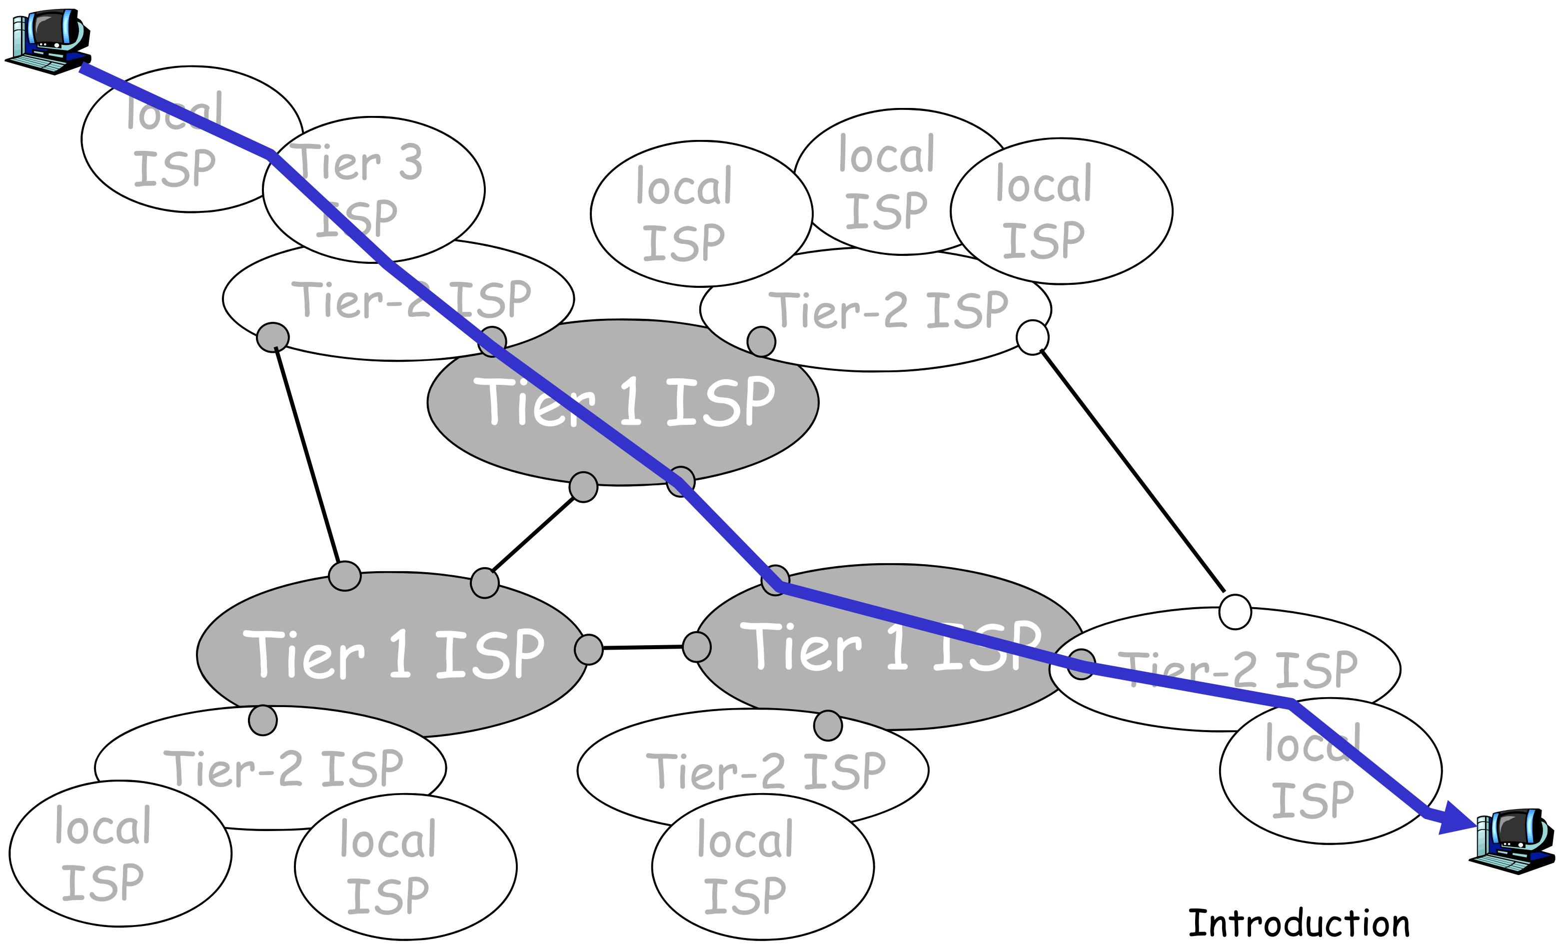

Recommended Reading:

- Internet History in the book: interesting and fun!

Security prelude

Network Security

- The field of network security is about:
  - how bad guys can attack computer networks
  - how we can defend networks against attacks
  - how to design architectures that are immune to attacks
- Internet not originally designed with (much) security in mind
  - *original vision:* "a group of mutually trusting users attached to a transparent network" ☺
  - Internet protocol designers playing "catch-up"
  - Security considerations in all layers!

Bad guys can put malware into hosts via Internet

- Malware can get in host from a virus, worm, or trojan horse.
- Spyware malware can record keystrokes, web sites visited, upload info to collection site.
- Infected host can be enrolled in a botnet, used for spam and DDoS attacks.
- Malware is often self-replicating: from an infected host, seeks entry into other hosts

Bad guys can put malware into hosts via Internet

- Trojan horse
  - Hidden part of some otherwise useful software
  - Today often on a Web page (Active-X, plugin)
- Virus
  - infection by receiving object (e.g., e-mail attachment), actively executing
  - self-replicating: propagate itself to other hosts, users
- Worm:
  - infection by passively receiving object that gets itself executed
  - self- replicating: propagates to other hosts, users

Sapphire Worm: aggregate scans/sec
in first 5 minutes of outbreak (CAIDA, UWisc data)

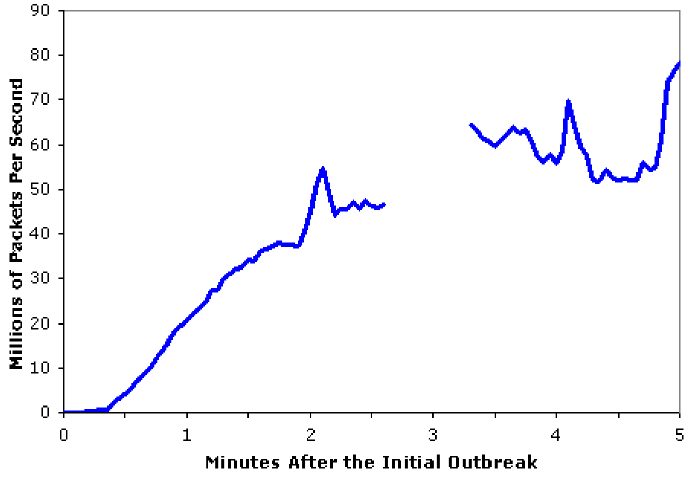

Bad guys can attack servers and network infrastructure

- Denial of service (DoS): attackers make resources (server, bandwidth) unavailable to legitimate traffic by overwhelming resource with bogus traffic

1. select target
2. break into hosts around the network (see botnet)
3. send packets toward target from compromised hosts

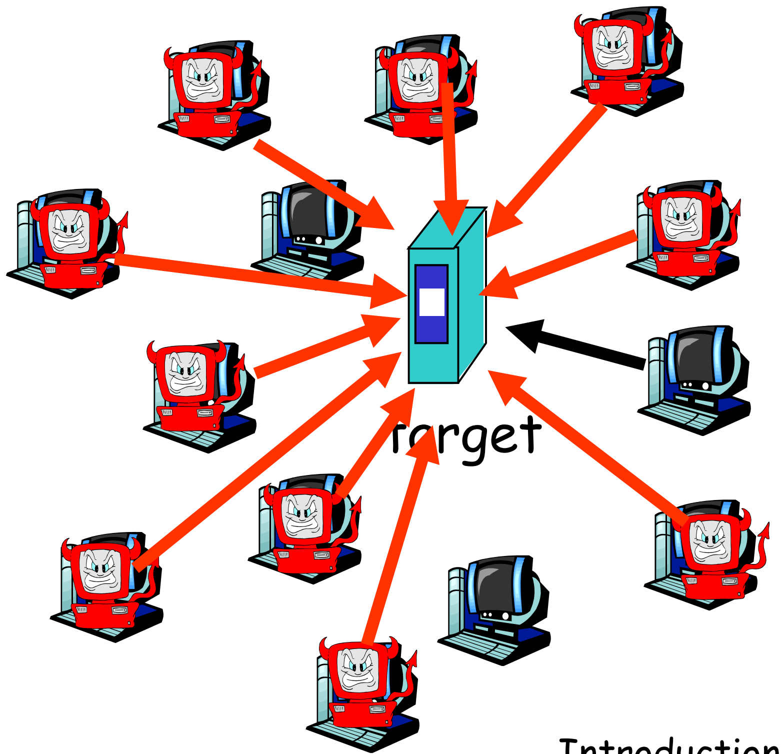

The bad guys can sniff packets

Packet sniffing:

- broadcast media (shared Ethernet, wireless)
- promiscuous network interface reads/records all packets (e.g., including passwords!) passing by

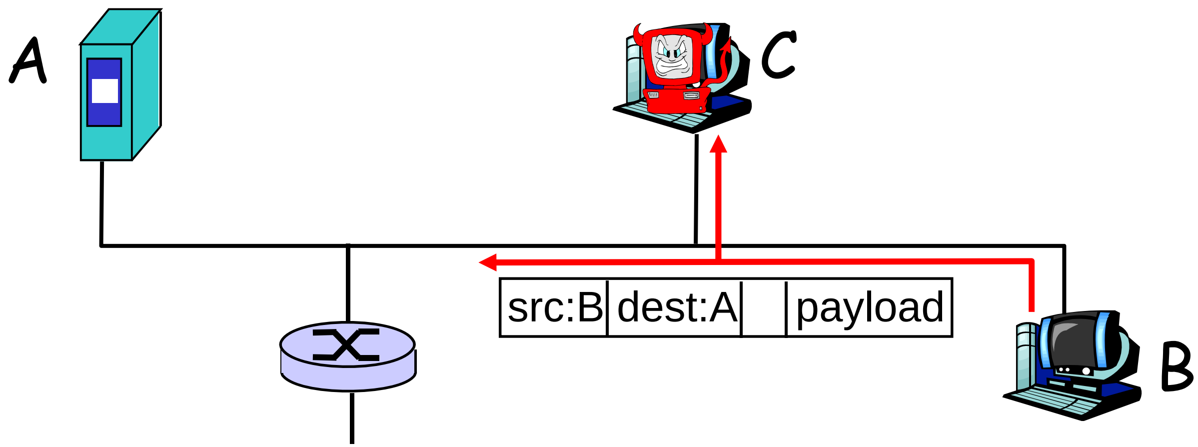

- Wireshark software used for end-of-chapter labs is a (free) packet-sniffer

The bad guys can use false source addresses

- *IP spoofing:* send packet with false source address

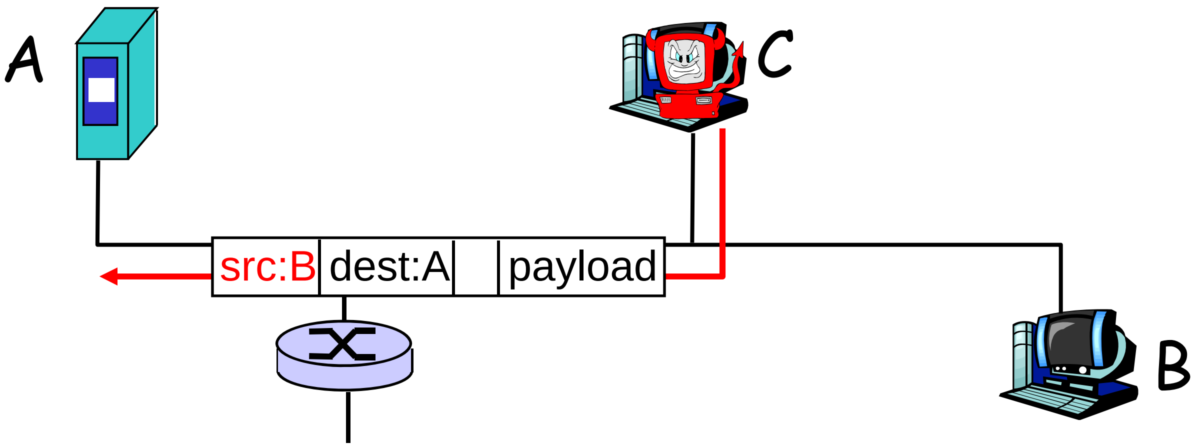

The bad guys can record and playback

- *record-and-playback*: sniff sensitive info (e.g., password), and use later
  - password holder *is* that user from system point of view

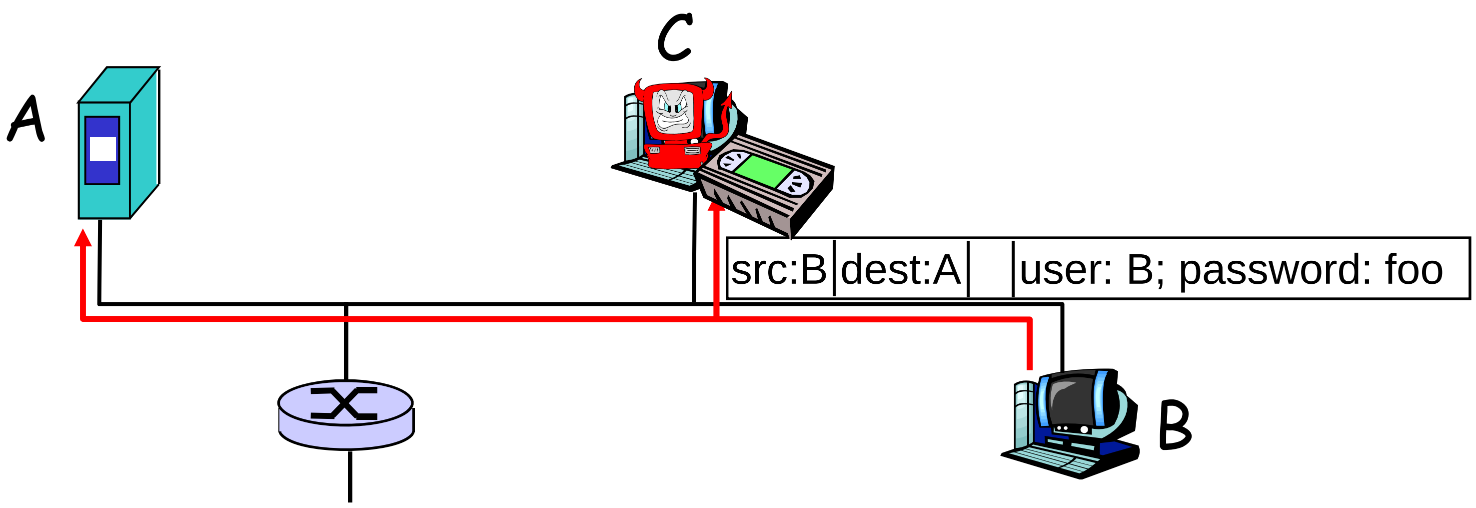

Chapter 1: Summary

Covered a "ton" of material!

- what's the Internet
- what's a protocol?
- network edge (types of service)
- network core (ways of transfer, routing, performance, delays, loss)
- access net, physical media
- protocol layers, service models
- backbones, NAPs, ISPs
- (history)
- Security concerns
- quick look into ATM networks (historical and service/resource-related perspective)

You now hopefully have:

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

Extra notes

Network Core: Circuit Switching

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

Circuit Switching: FDM and TDM

Example:

4 users

FDM

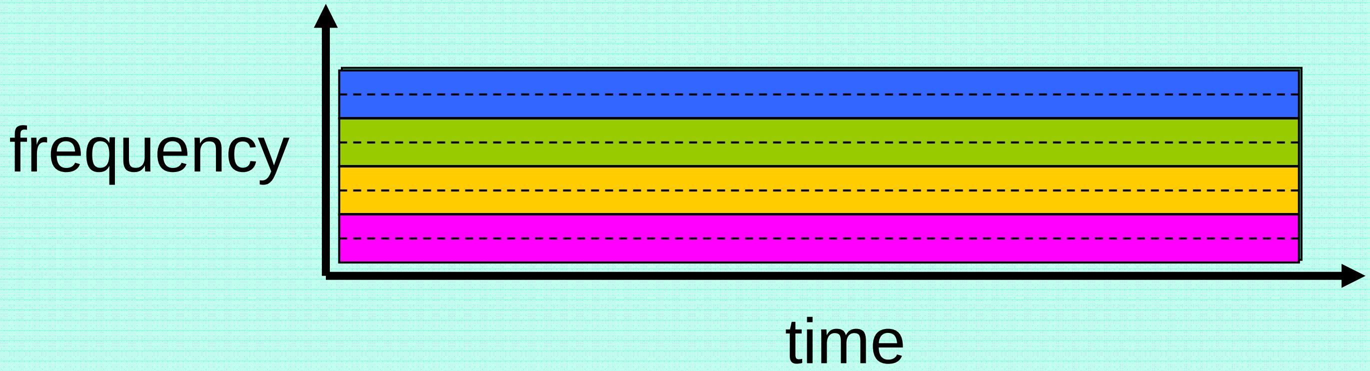

TDM

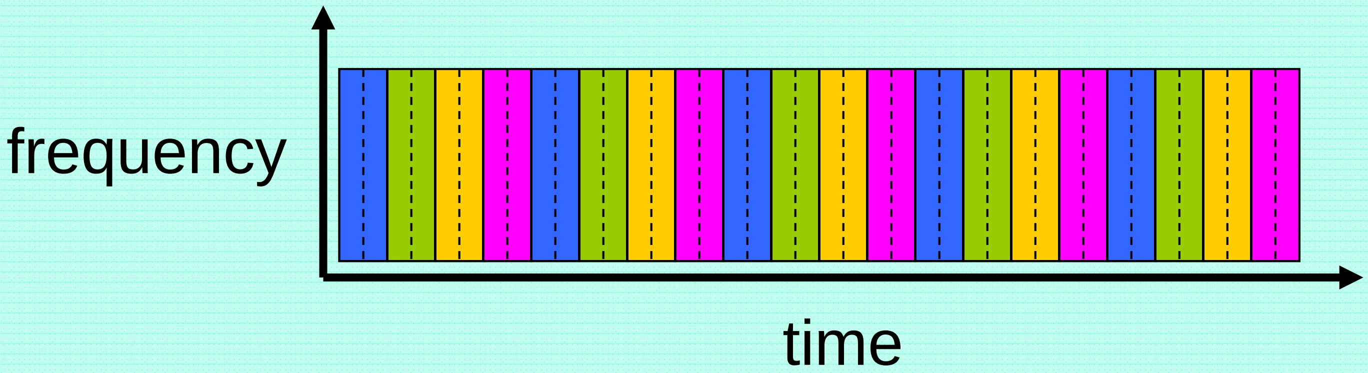

ATM Networking What/why was that?

(paved MPLS networking - Multiprotocol label switchng):

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 (transport-layer-like functionality)
  - end-system
  - segmentation/re-assembly
- ATM Layer: cell switching (network-layer-type functionality)
- Physical

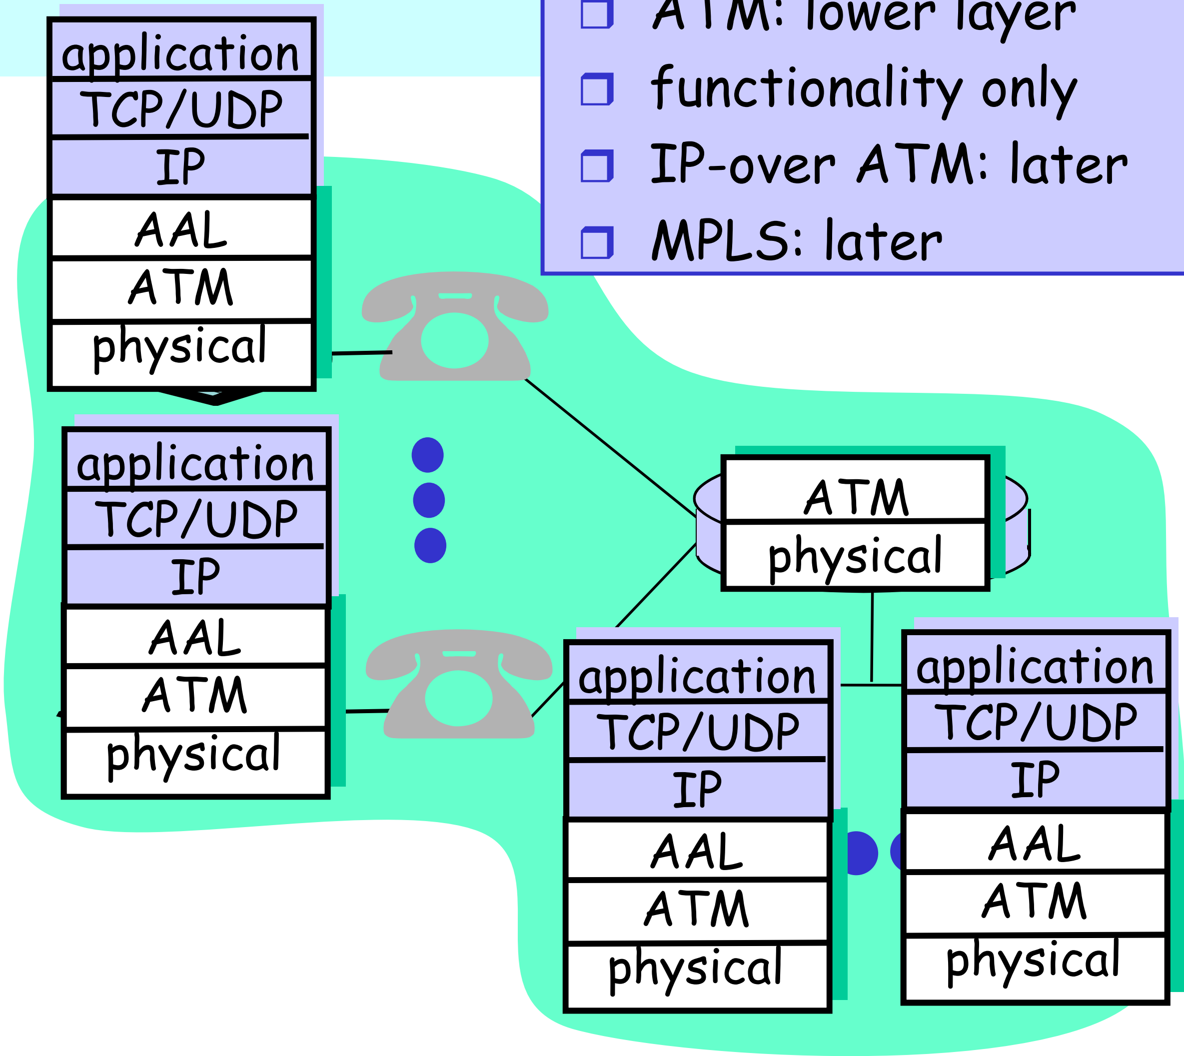

Where's the application?

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