UMBC

AN HONORS UNIVERSITY IN MARYLAND

IS 450/IS 650–
Data Communications and Networks

Introduction
(Chapter 1)

Nirmalya Roy
Department of Information Systems
University of Maryland Baltimore County

Chapter 1: Introduction

Our goal:

- get “feel” and terminology
- more depth, detail *later* in course
- approach:
  - use Internet as example

Overview:

- what’s the Internet
- what’s a protocol?
- network edge
- access net, physical media
- network core
- Internet/ISP structure
- performance: loss, delay
- protocol layers, service models
- network modeling

Chapter 1: Roadmap

1.1 What *is* the Internet?

1.2 Network edge

1.3 Network access and physical media

1.4 Network core

1.5 Internet structure and ISPs

1.6 Delay & loss in packet-switched networks

1.7 Protocol layers, service models

1.8 History

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
  - Different transmission rates
- *Packet switches*
  - *Routers/link layer switches:* forward packets (chunks of data)

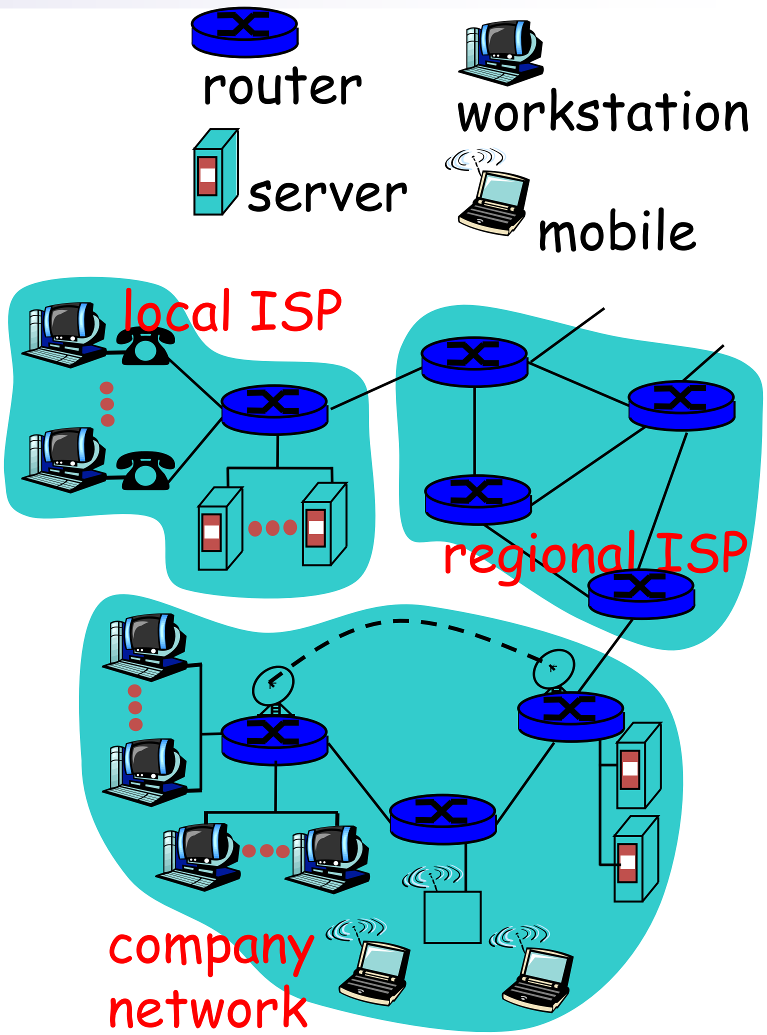

“Cool” Internet Appliances

IP picture frame
http://www.ceiva.com/

Web-enabled toaster +
weather forecaster

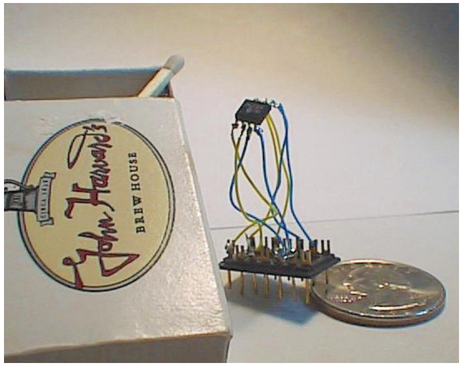

World’s smallest web server
http://www-ccs.cs.umass.edu/~shri/iPic.html

Internet phones

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

- *protocols* coordinate communication
  - Who gets to transmit?
  - What path to take?
  - What message format?
  - e.g., HTTP, FTP, PPP, TCP, IP

- *Internet:* "network of networks"
  - loosely hierarchical
  - public Internet vs. private intranet

- Internet standards
  - RFC: Request for comments
  - IETF: Internet Engineering Task Force

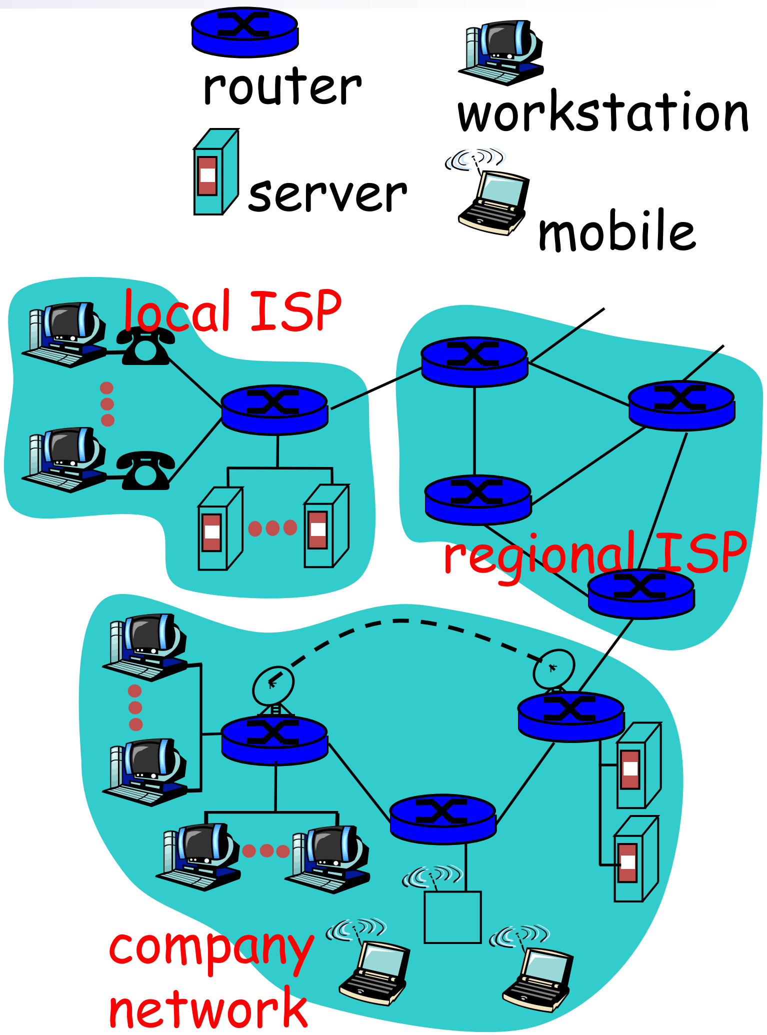

What's the Internet: A Service View

- communication *infrastructure* enables distributed applications:
  - Web, email, games, e-commerce, file sharing
- communication services provided to apps:
  - Different end systems
  - Internet API
  - Connectionless unreliable
  - connection-oriented reliable

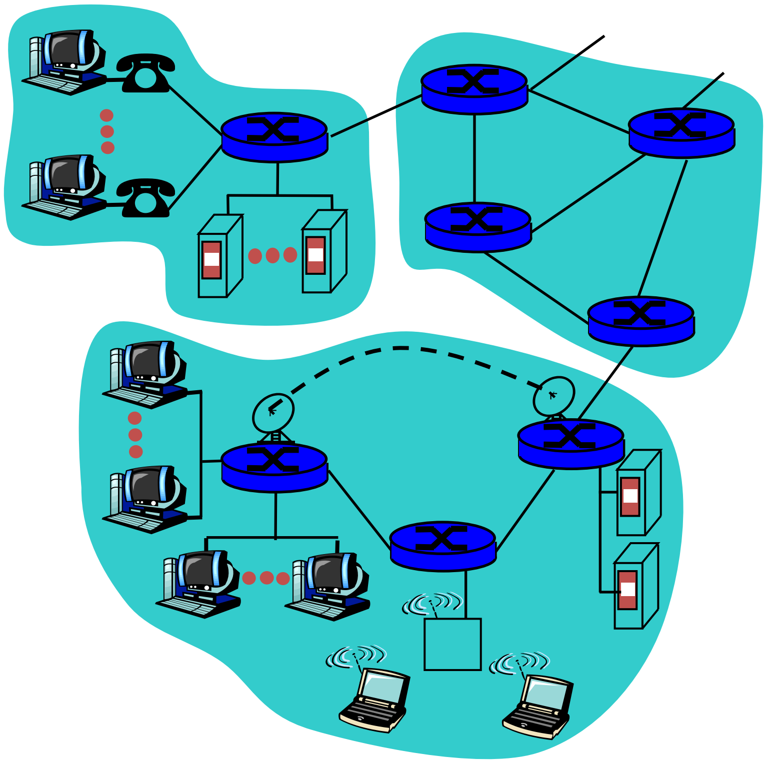

Think of an analogy of this in real life services

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 coordinated 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:

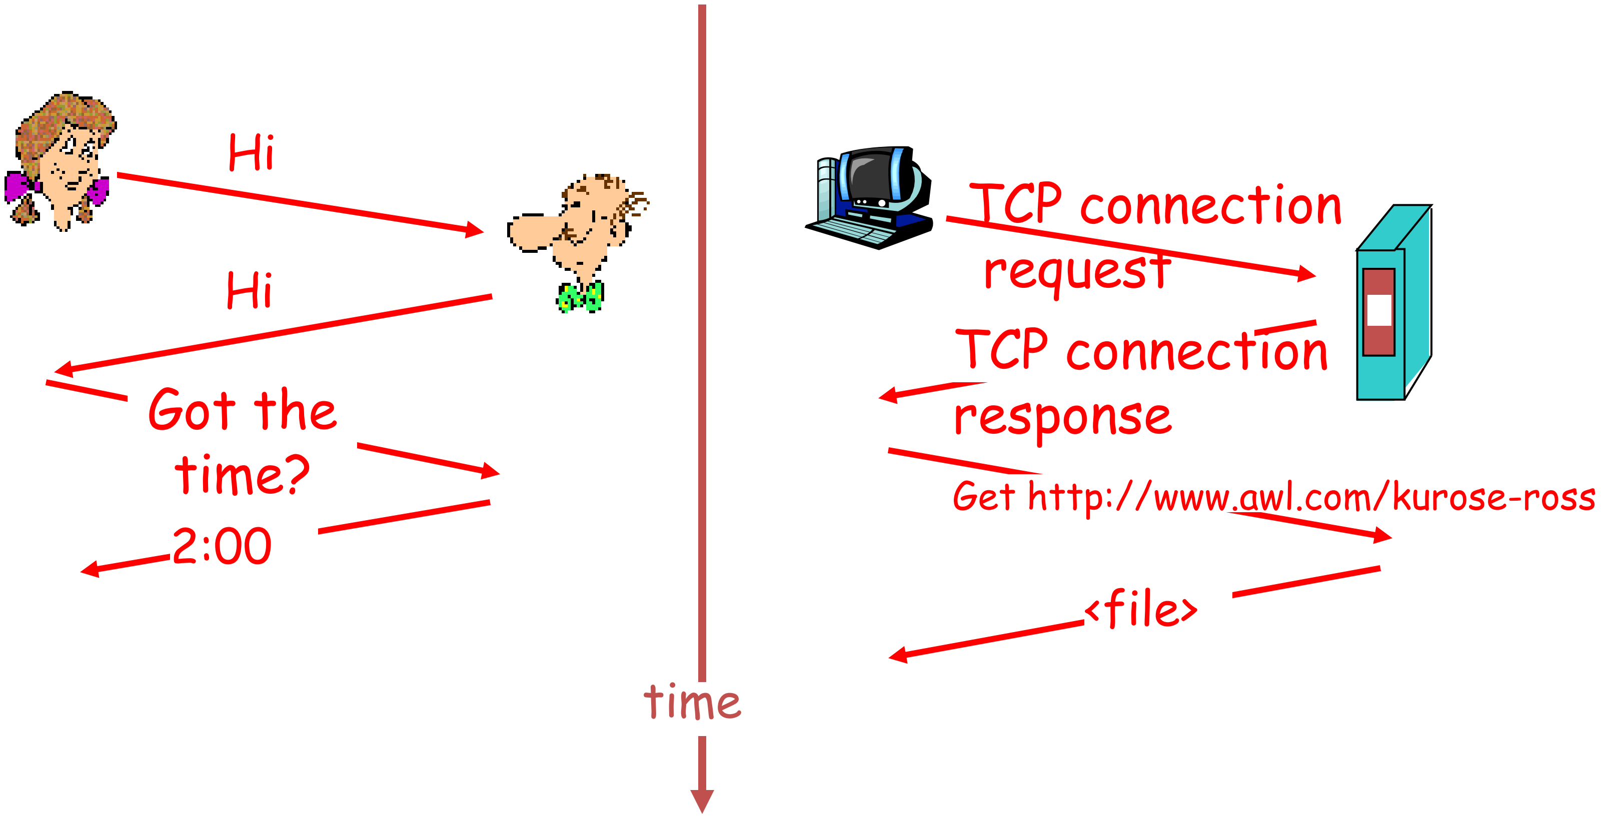

This one is trivial. Can you think of a more complex case?

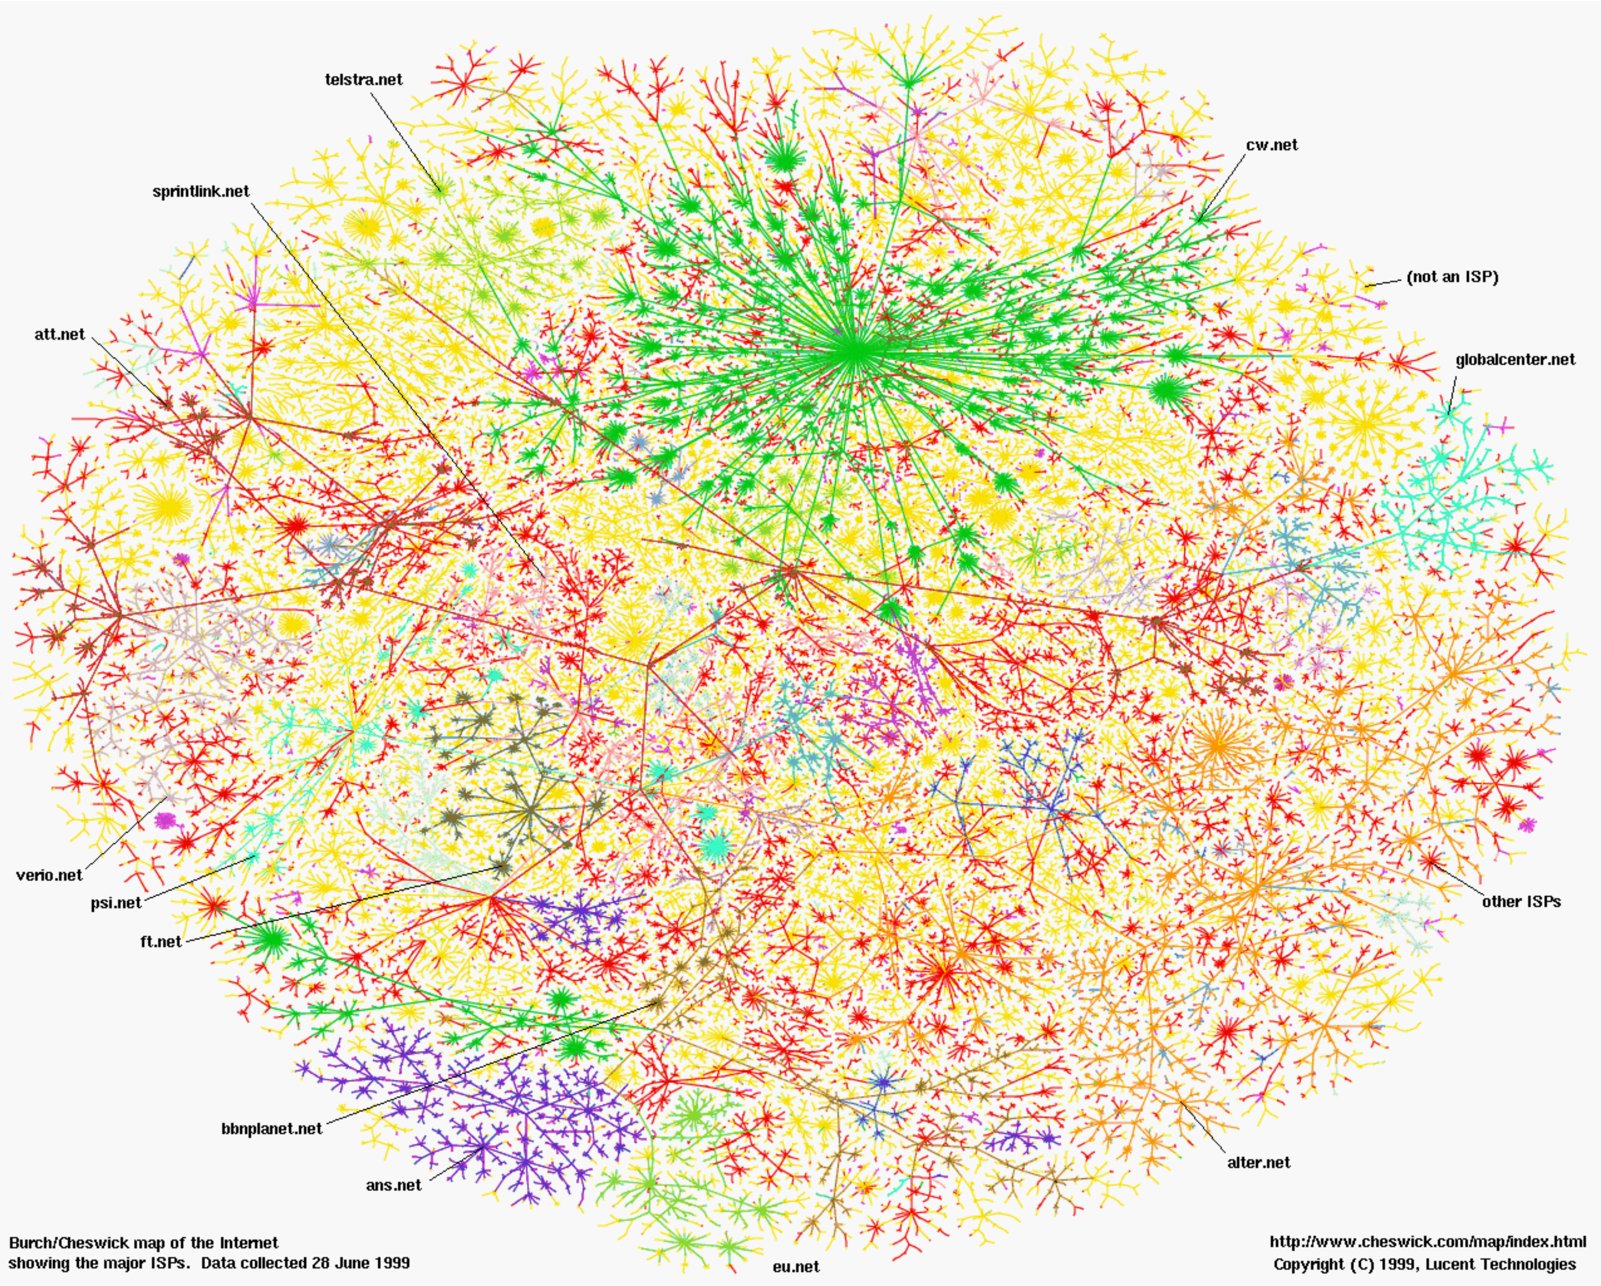
telstra.net
cw.net
sprintlink.net
(not an ISP)
att.net
globalcenter.net
verio.net
psi.net
ft.net
other ISPs
bbnplanet.net
ans.net
alter.net
Burch/Cheswick map of the Internet
showing the major ISPs. Data collected 28 June 1999
eu.net
http://www.cheswick.com/map/index.html
Copyright (C) 1999, Lucent Technologies

Chapter 1: Roadmap

1.1 What *is* the Internet?

1.2 Network edge

1.3 Network access and physical media

1.4 Network core

1.5 Internet structure and ISPs

1.6 Delay & loss in packet-switched networks

1.7 Protocol layers, service models

1.8 History

A Closer Look at Network Structure

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

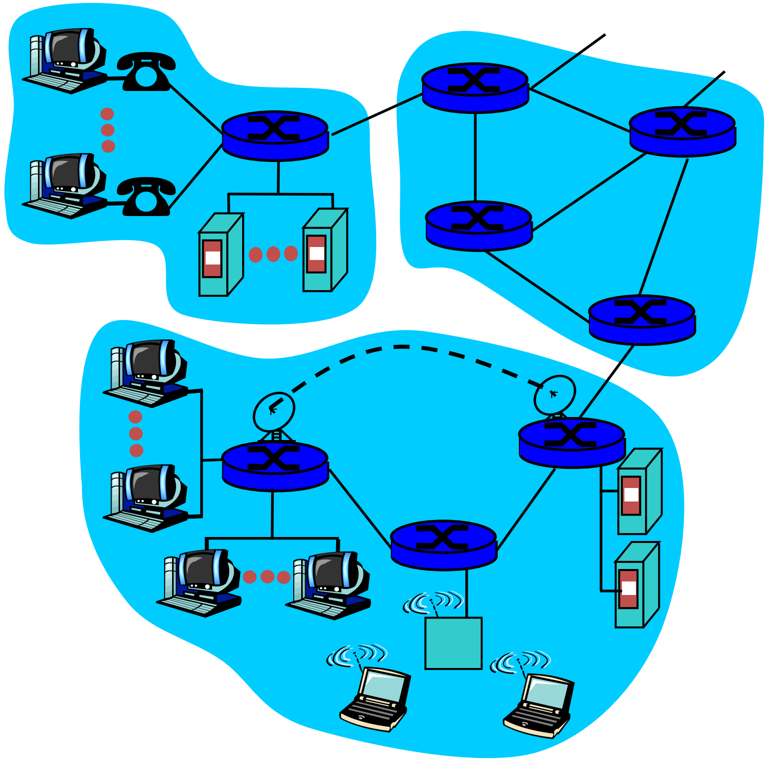

The Network Edge

- end systems (hosts):
  - run application programs
  - e.g. Web, email
- client/server model
  - client host requests, receives service from always-on server
  - e.g. Web browser/server; email client/server
- peer-peer model:
  - minimal use of dedicated servers
  - e.g. Skype, BitTorrent, KaZaA

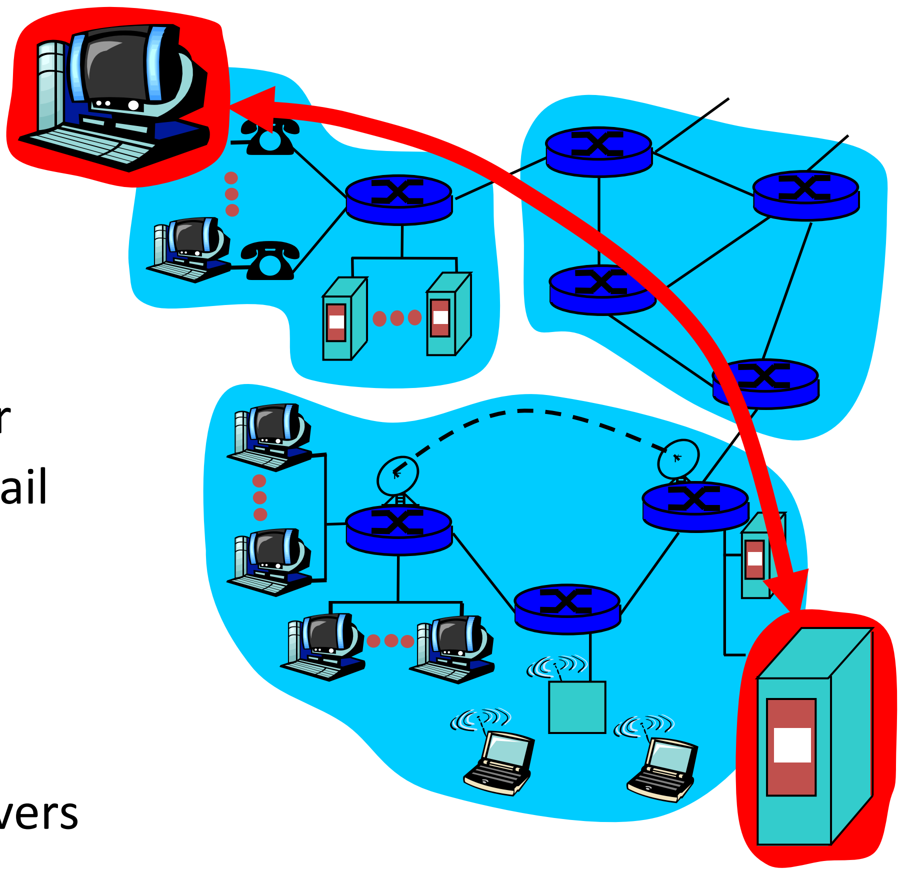

Network Edge: Connection-oriented Service

Goal: data transfer between end systems

- ***Connection:*** prepare for data transfer ahead of time
  - Request / Respond
  - *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

… like buying flight tickets for the full international trip

Network Edge: Connectionless Service

Goal: data transfer between end systems

- same as before!

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

App's using TCP:

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

App's using UDP:

- streaming media, teleconferencing, DNS, Internet telephony

… like buying separate flight tickets for each flight segment

Chapter 1: Roadmap

1.1 What *is* the Internet?

1.2 Network edge

1.3 Network access and physical media

1.4 Network core

1.5 Internet structure and ISPs

1.6 Delay & loss in packet-switched networks

1.7 Protocol layers, service models

1.8 History

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?

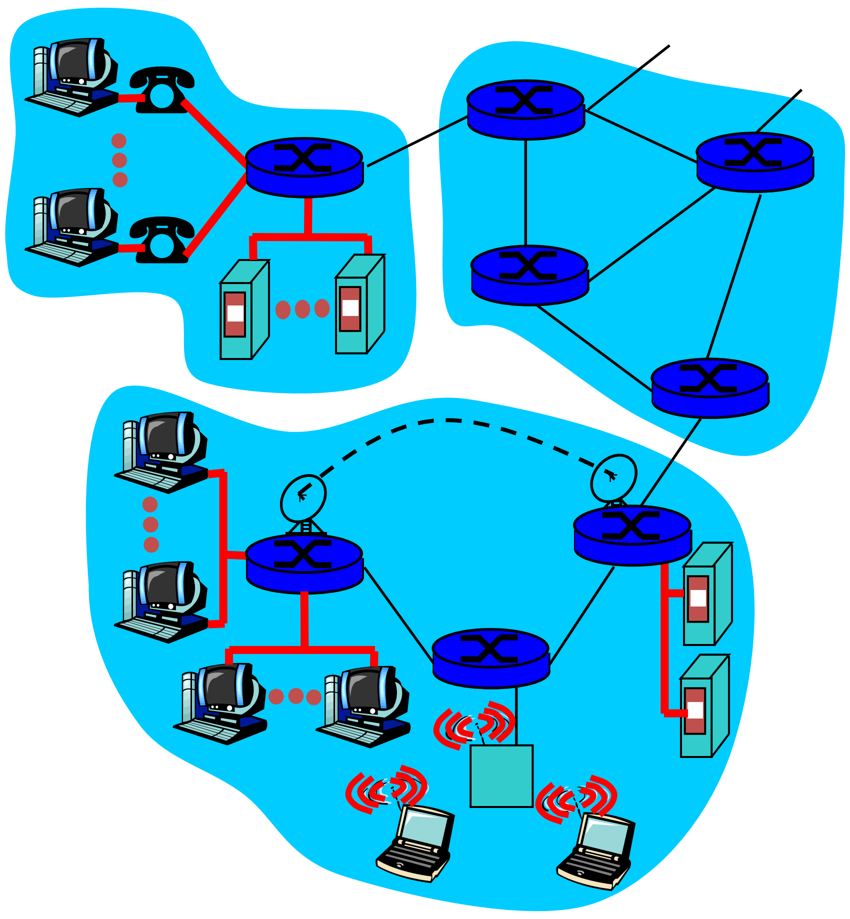

Residential Access: Point to Point Access

- Dialup via modem
  - up to 56Kbps direct access to router (often less)
  - Can't surf and phone at same time: can't be "always on"

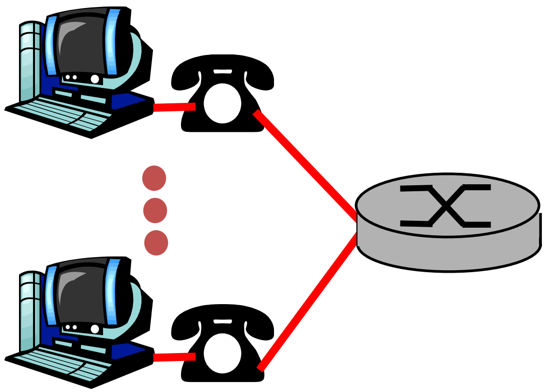

- ADSL: asymmetric digital subscriber line
  - up to 1 Mbps upstream (today typically < 256 kbps)
  - up to 8 Mbps downstream (today typically < 1 Mbps)
  - FDM: 50 kHz - 1 MHz for downstream

    4 kHz - 50 kHz for upstream

    0 kHz - 4 kHz for ordinary telephone

Residential Access : Digital Subscriber Line (DSL)

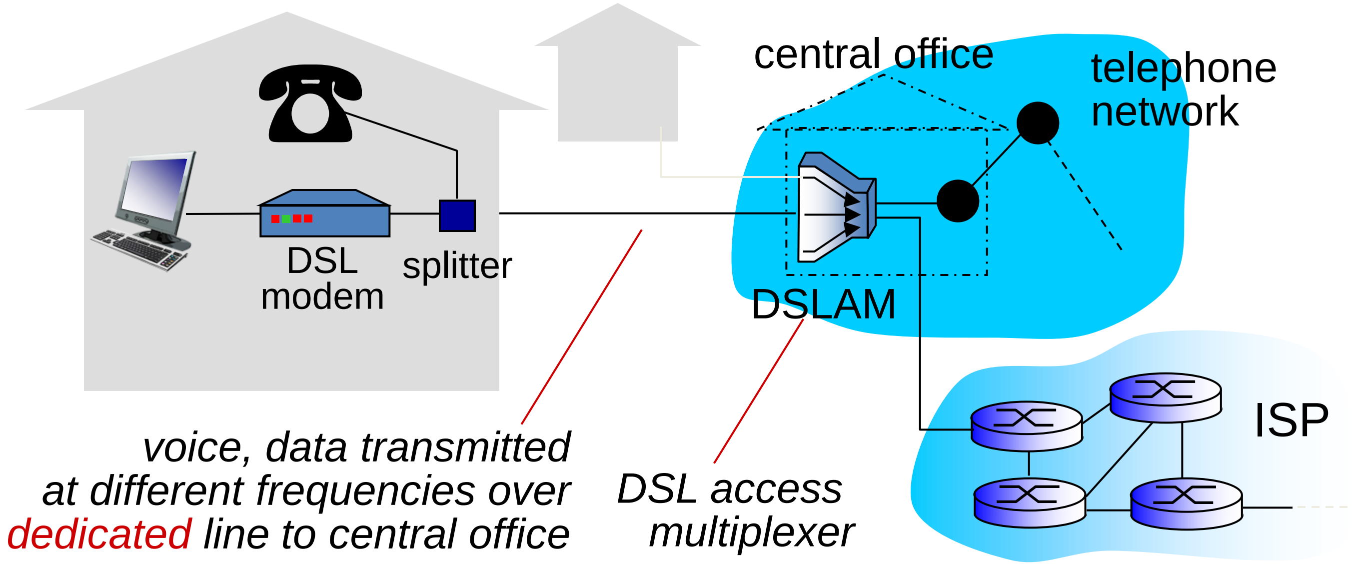

- use *existing* telephone line to central office DSLAM
  - data over DSL phone line goes to Internet
  - voice over DSL phone line goes to telephone net
  - < 2.5 Mbps upstream transmission rate (typically < 1 Mbps)
  - < 24 Mbps downstream transmission rate (typically < 10 Mbps)

Residential Access: Cable Modems

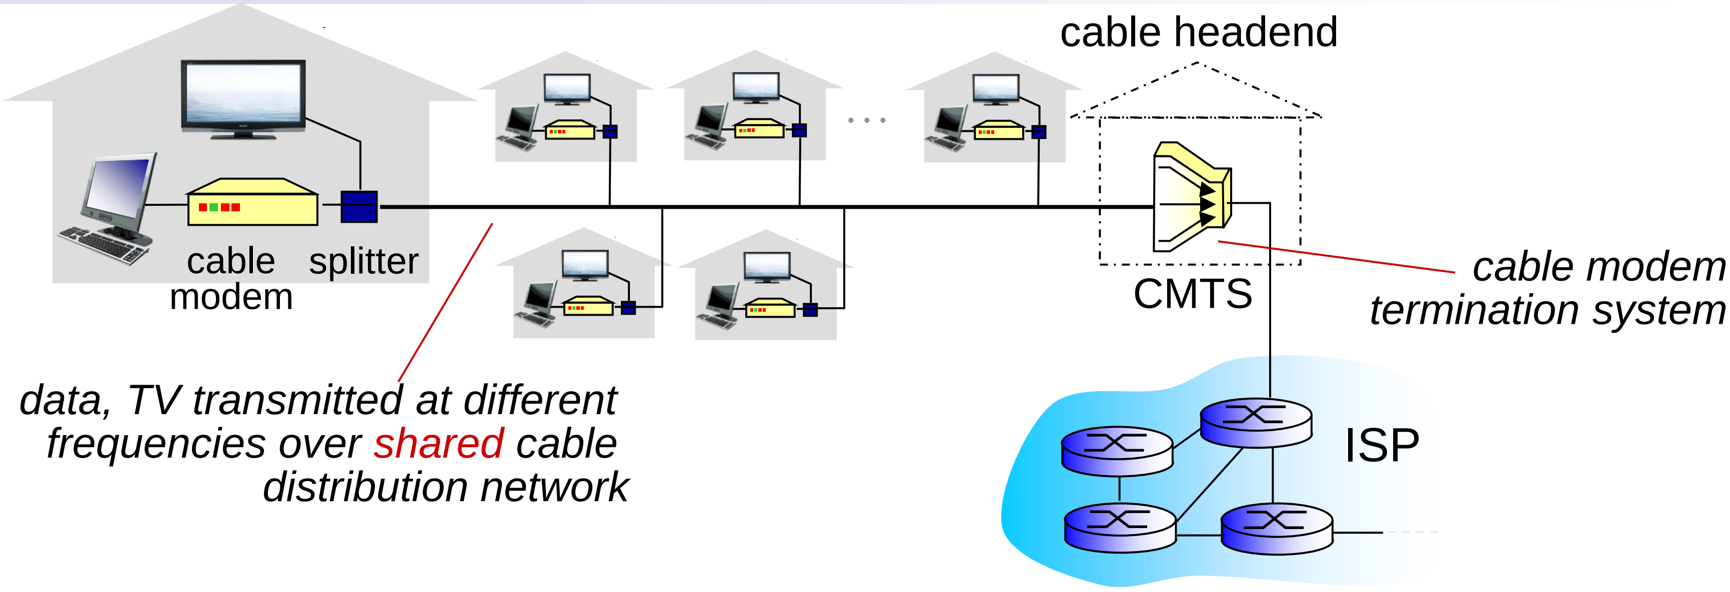

- HFC: hybrid fiber coax
  - asymmetric: up to 30Mbps downstream transmission rate, 2 Mbps upstream transmission rate
- network of cable, fiber attaches homes to ISP router
  - homes *share access network* to cable headend
  - unlike DSL, which has dedicated access to central office

Residential Access: Cable Modems

- HFC: hybrid fiber coax
  - asymmetric: up to 30Mbps downstream, 2 Mbps upstream
- network of cable and fiber attaches home to ISP router
  - homes share access to router
- deployment: available via cable TV companies

Residential Access: Cable Modems

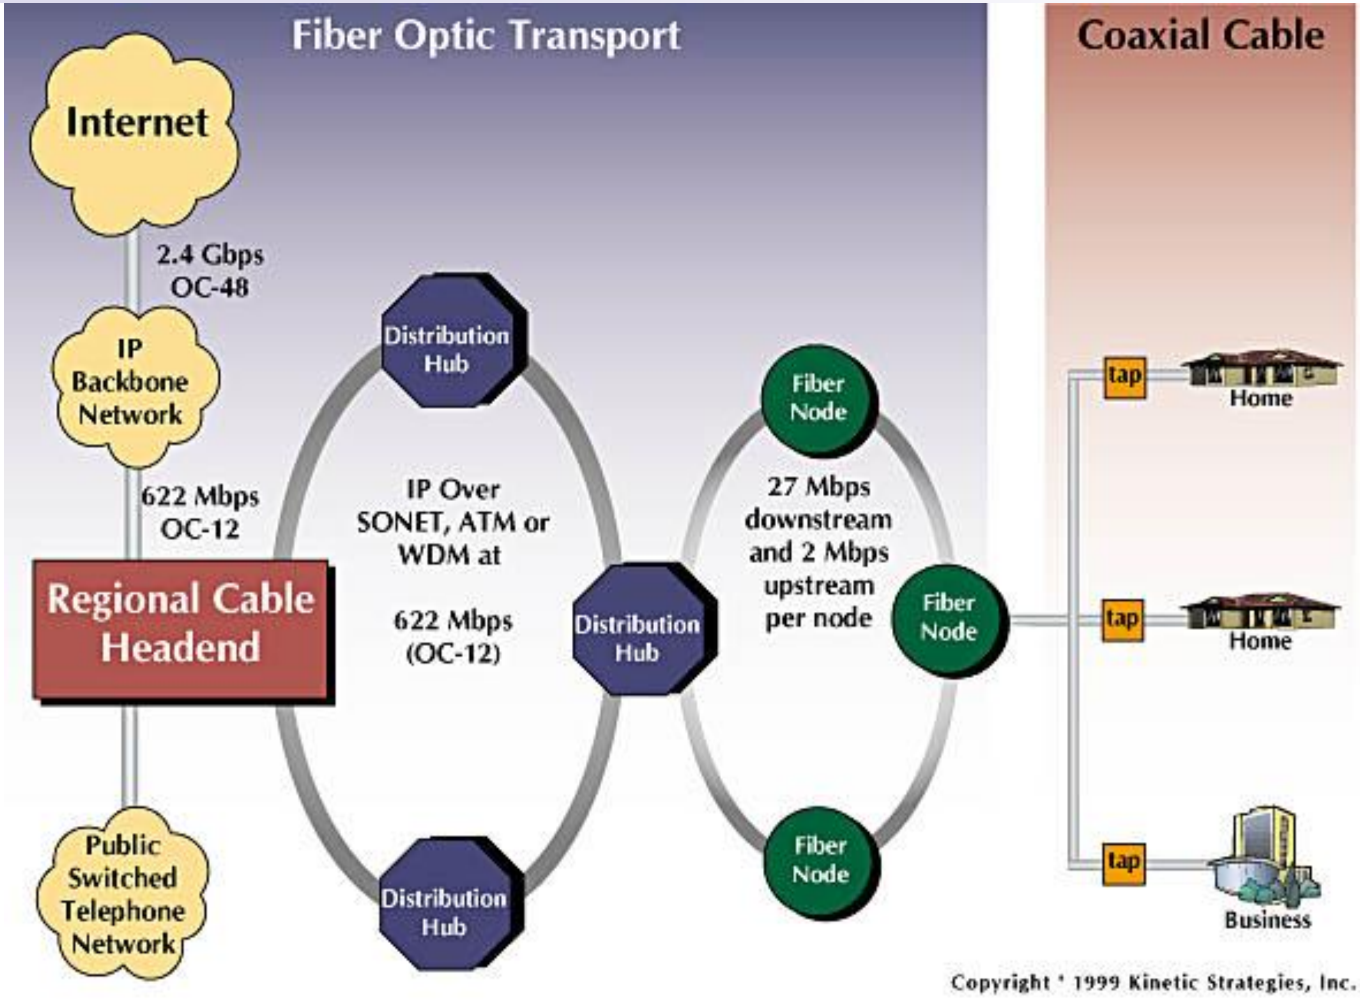

Diagram: http://www.cabledatacomnews.com/cmic/diagram.html

Cable Network Architecture: Overview

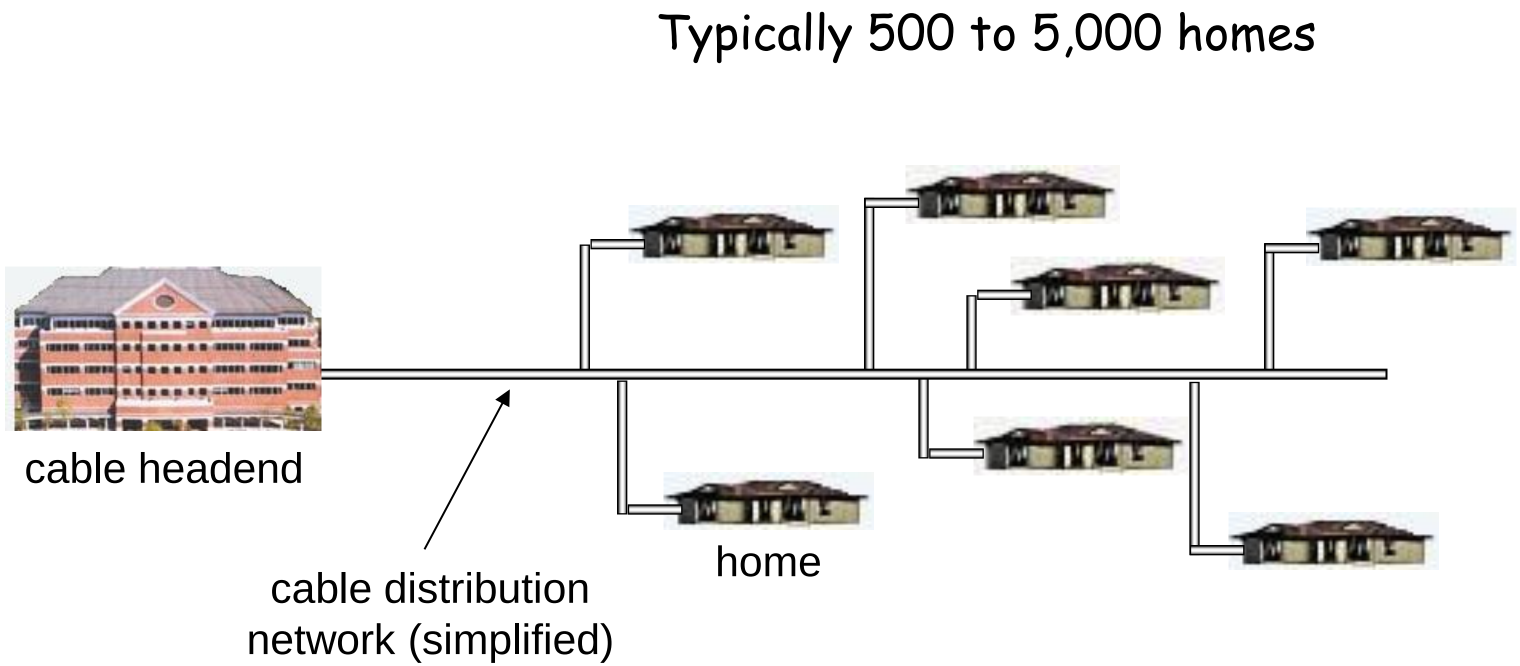

Cable Network Architecture: Overview

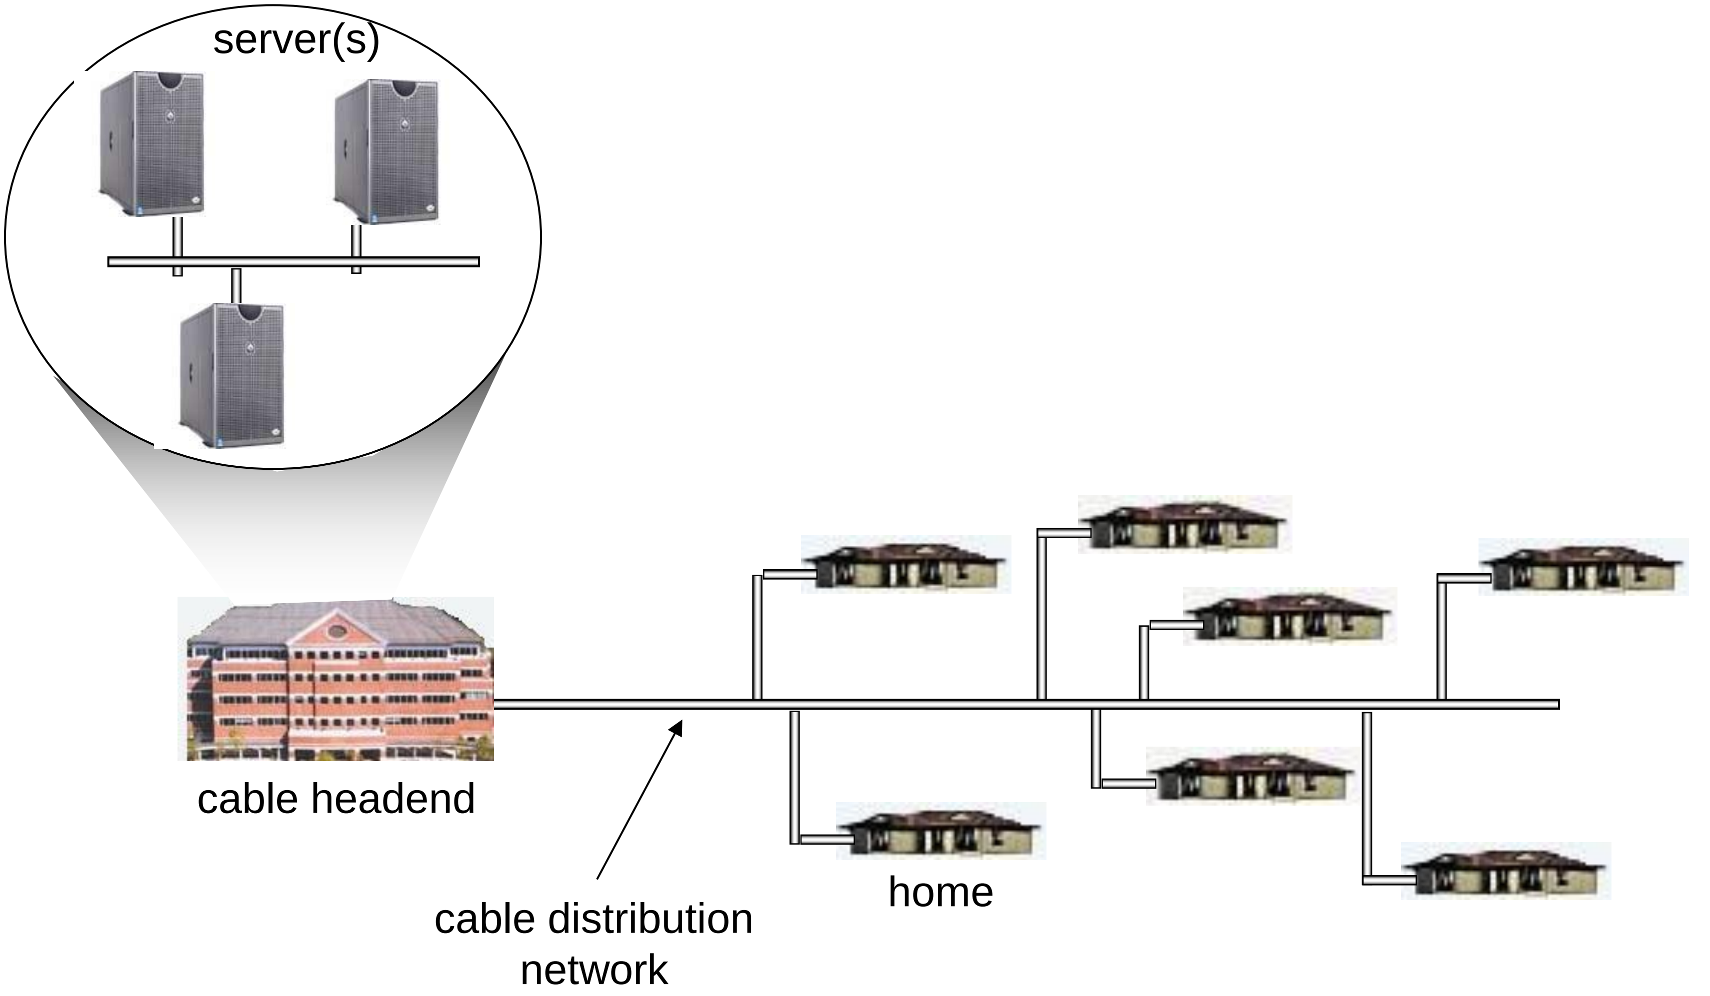

Cable Network Architecture: Overview

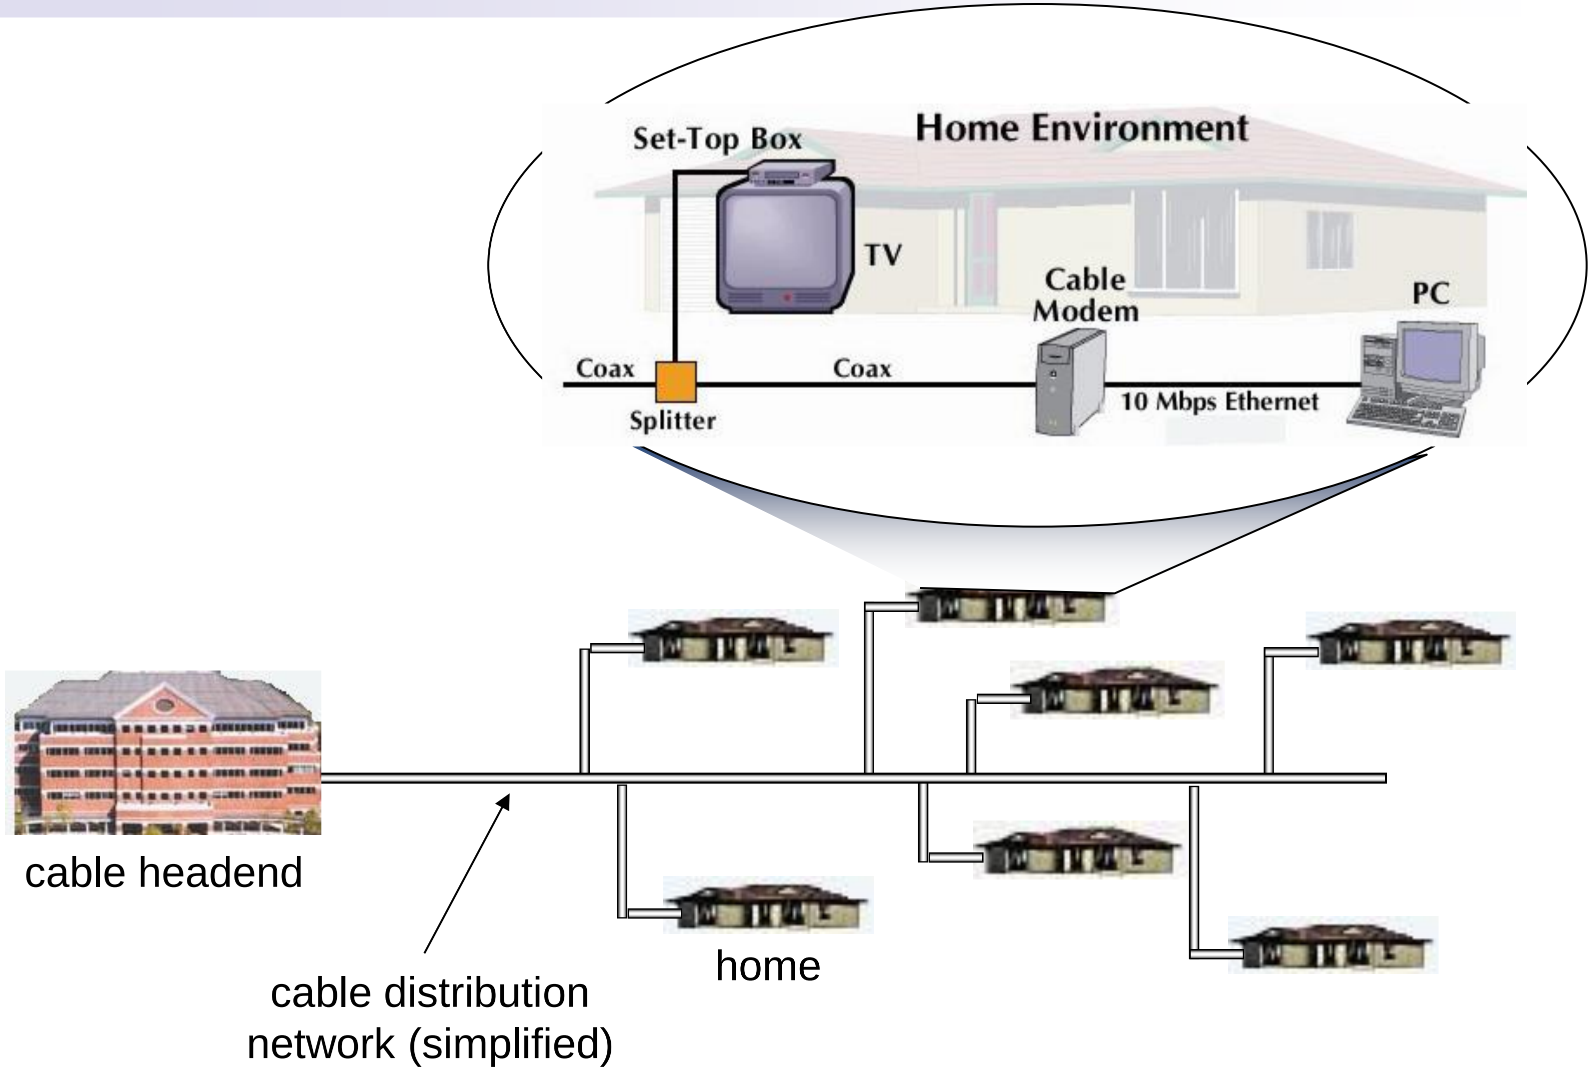

Cable Network Architecture: Overview

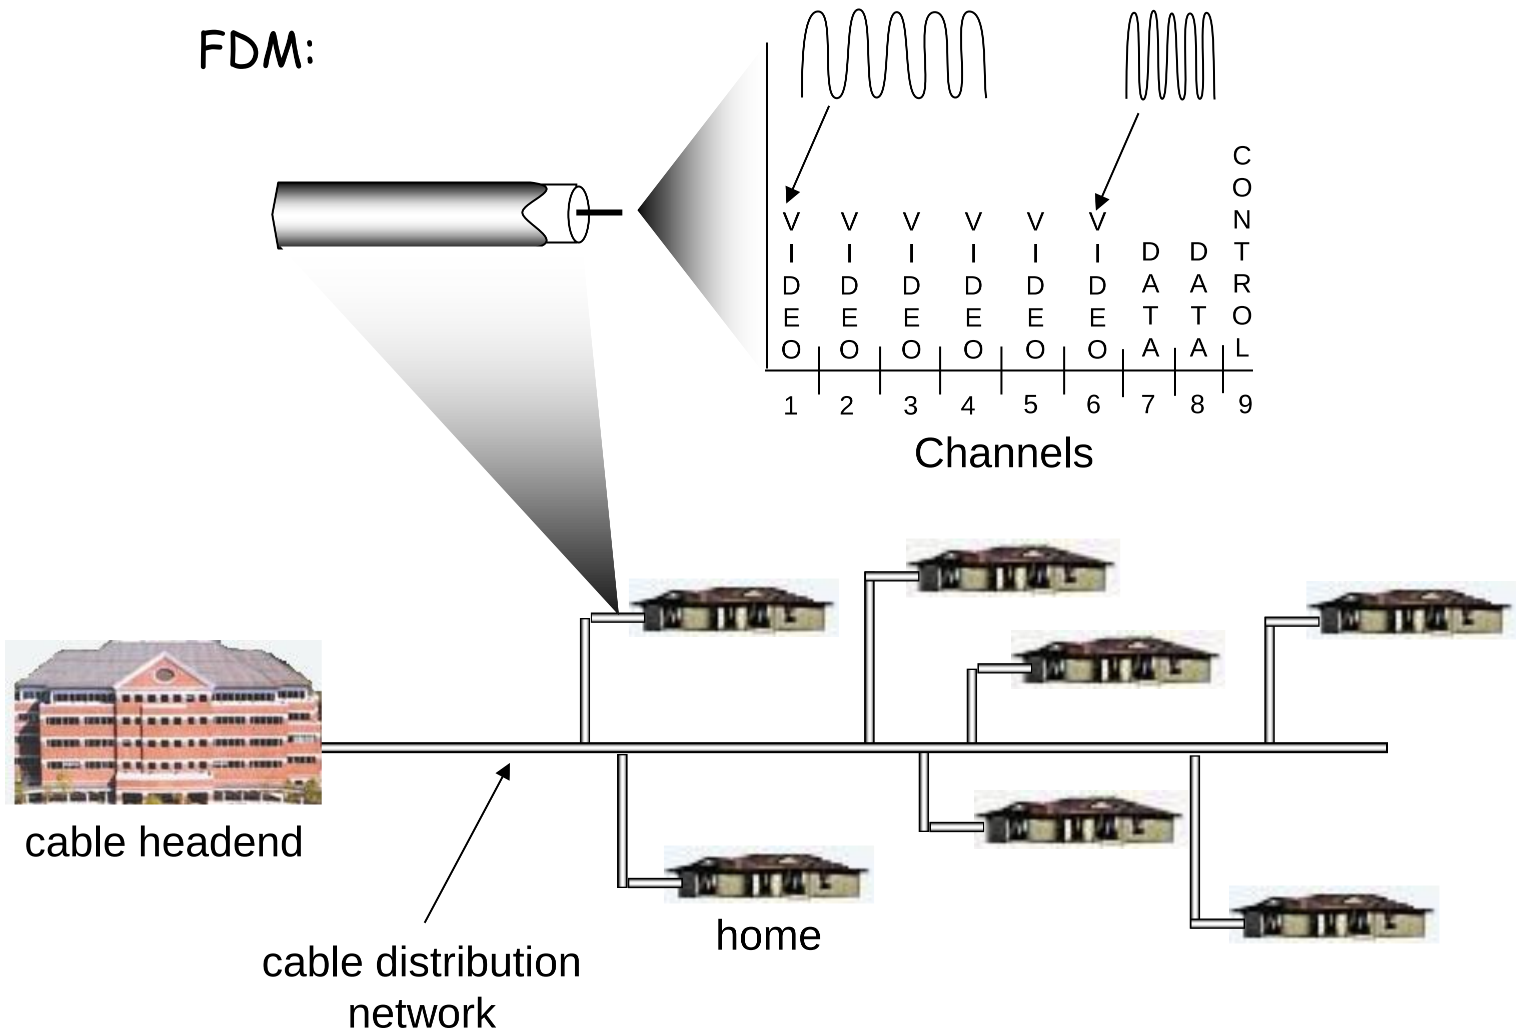

DSL vs Cable Modem

- DSL is point to point

  Thus data rate does not reduce when neighbor uses his/her DSL

- But, DSL uses twisted-pair, and transmission technology cannot support more than ~10Mbps

- Cable Modems share the pipe to the cable headend.

  Thus, your data rate can reduce when neighbors are surfing concurrently

- However, fibre optic lines have significantly higher data rate (fat pipe)

  Even if other users, data rate may still be higher

The debate / competition continues …

Company Access: Local Area Networks

- company/univ local area network (LAN) connects end system to edge router
- Ethernet:
  - shared or dedicated link connects end system and router
  - 10 Mbs, 100Mbps, Gigabit Ethernet
- LANs: Chapter 5

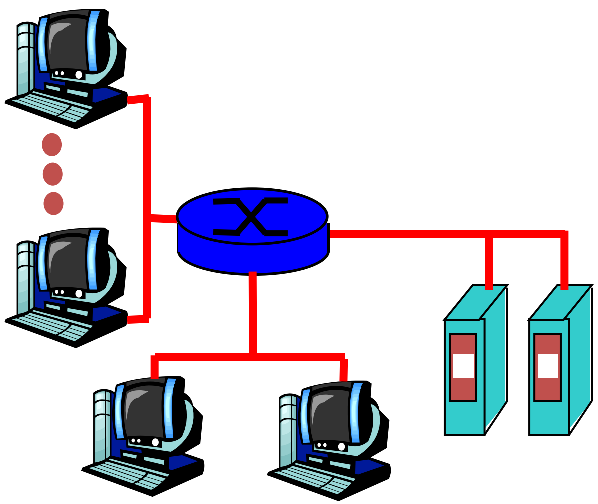

Wireless Access Networks

- 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
  - 3G ~ 384 kbps
    - Will it happen??
  - GPRS (General packet radio service) in Europe/US, LTE ~ 10 Mbps

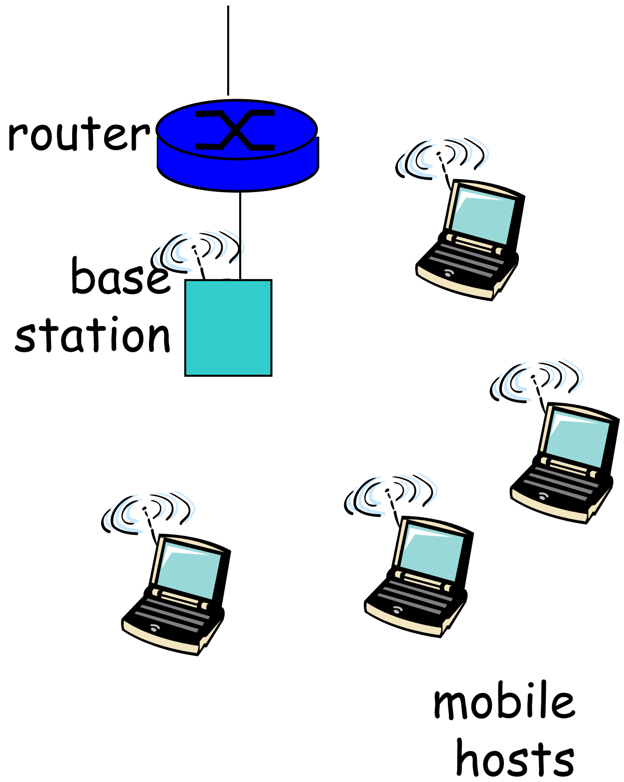

Home Networks

Typical home network components:

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

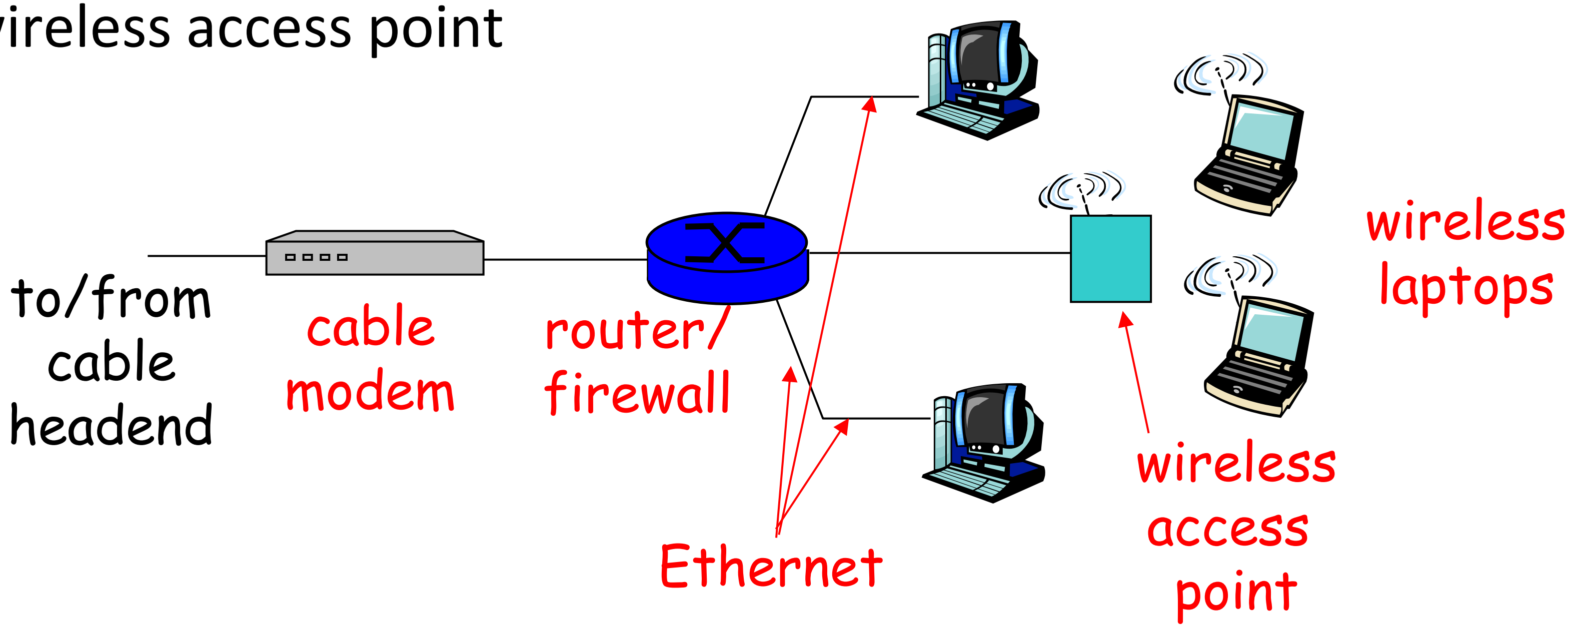

Chapter 1: Roadmap

1.1 What *is* the Internet?

1.2 Network edge

1.3 Network access and physical media

1.4 Network core

1.5 Internet structure and ISPs

1.6 Delay & loss in packet-switched networks

1.7 Protocol layers, service models

1.8 History

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"
  - Forwarding table and routing protocols

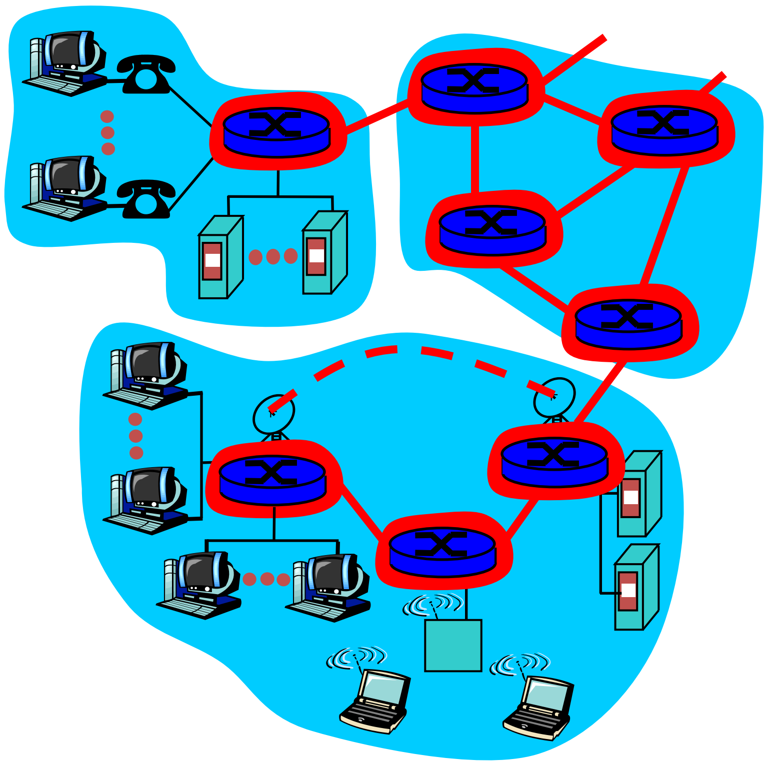

Network Core: Circuit Switching

End-end resources reserved for "call"

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

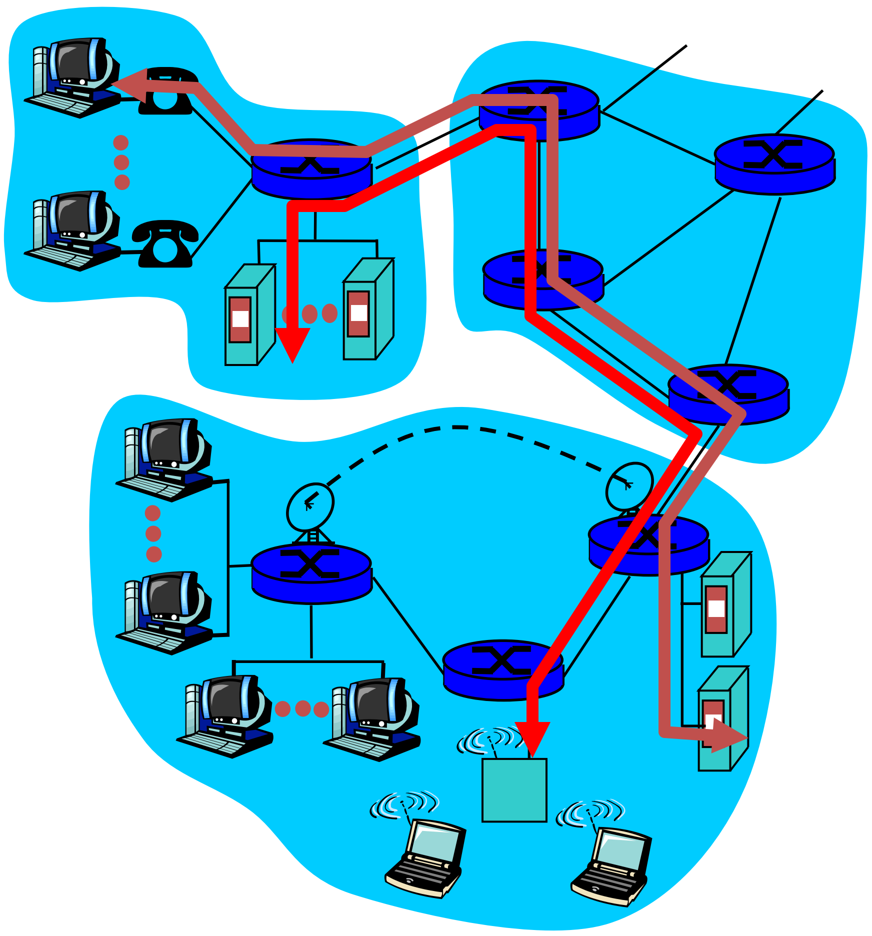

Network Core: Circuit Switching

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

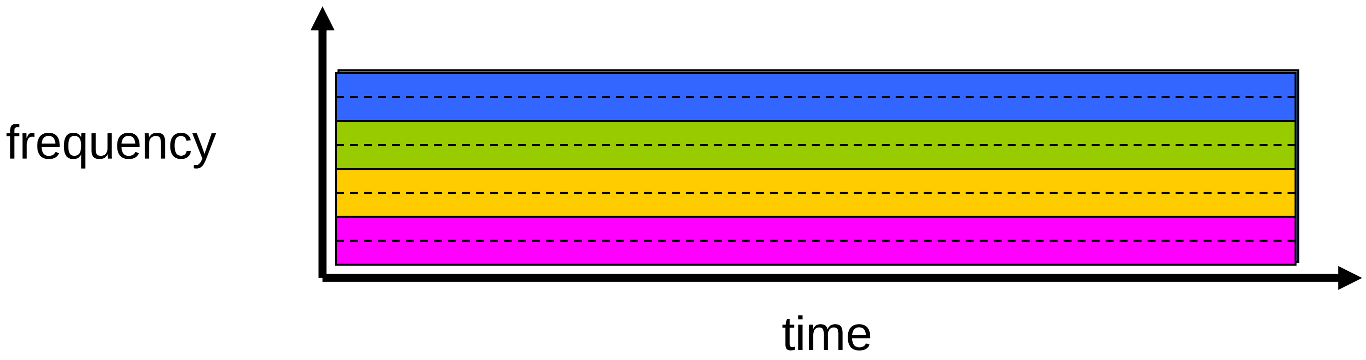

TDM

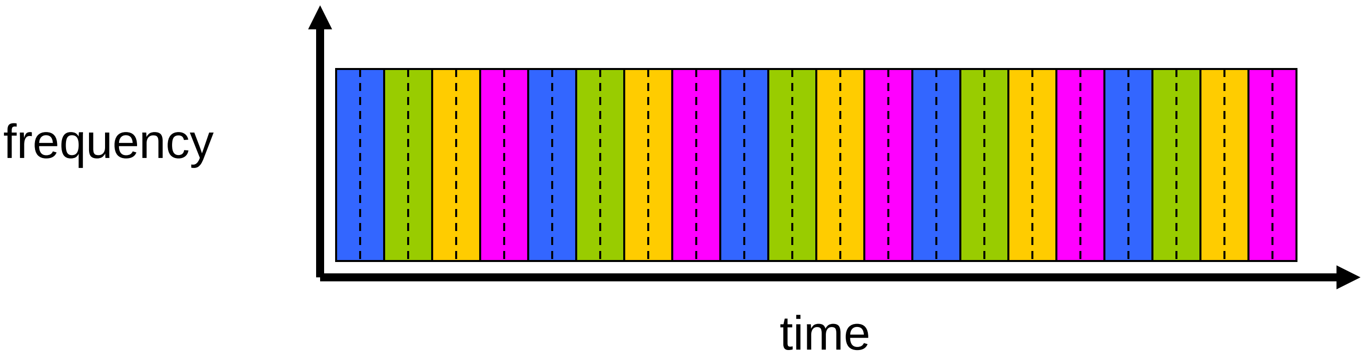

FDM vs TDM

- What are the tradeoffs?
  - Advantage and disadvantage of dividing frequency ?
  - Advantage and disadvantage of dividing time ?

Numerical example

- How long does it take to send a file of 640,000 bits from host A to host B over a circuit-switched network?
  - All links are 1.536 Mbps
  - Each link uses TDM with 24 slots/sec
  - 500 msec to establish end-to-end circuit

Let's work it out!

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*

Bandwidth division into "pieces"
Dedicated allocation
Resource reservation

resource contention:

- aggregate resource demand can exceed amount available
  - Packets queue up
- store and forward: packets move one hop at a time
  - Node receives complete packet before forwarding

Packet Switching: Statistical Multiplexing

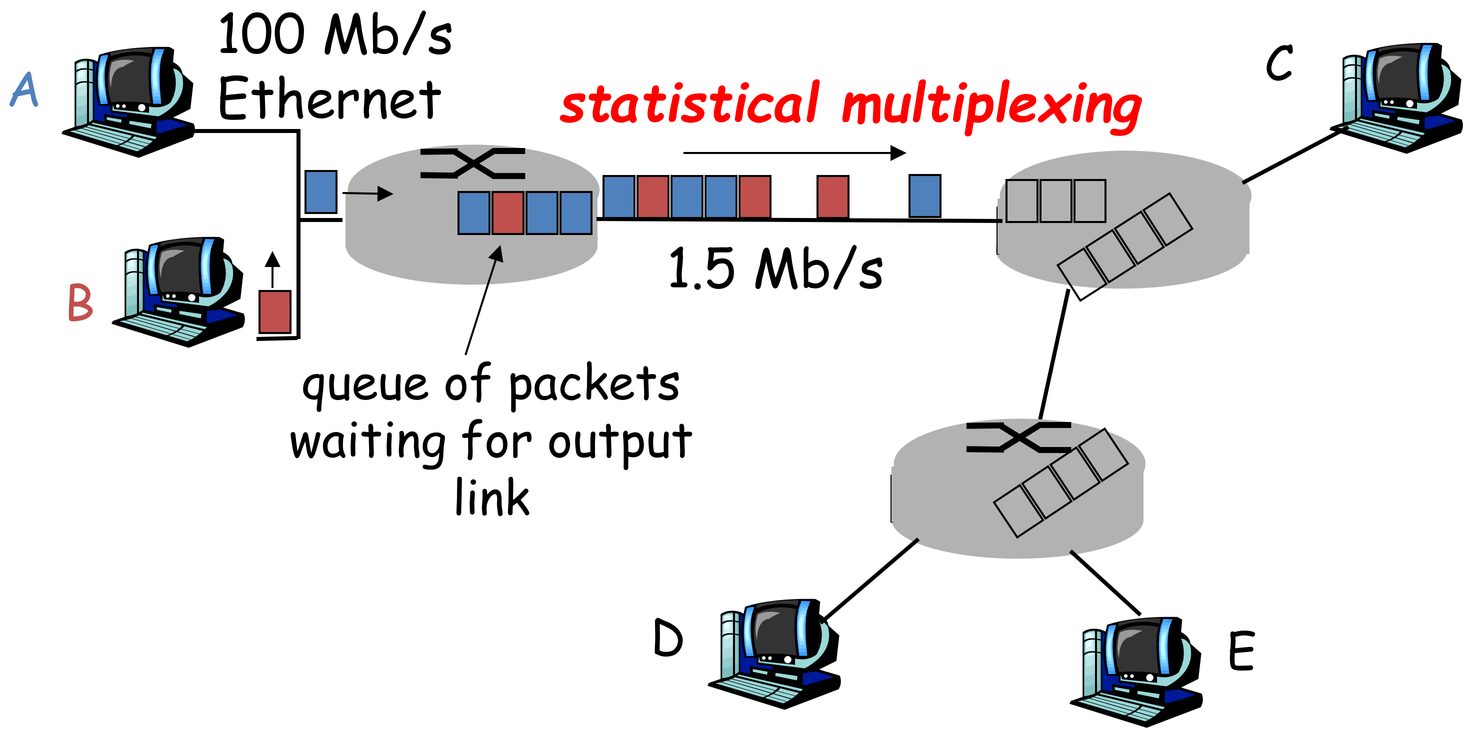

Sequence of A & B packets does not have fixed pattern, shared on demand ➔ ***statistical multiplexing***.

TDM: each host gets same slot in revolving TDM frame.

Packet-switching: store-and-forward

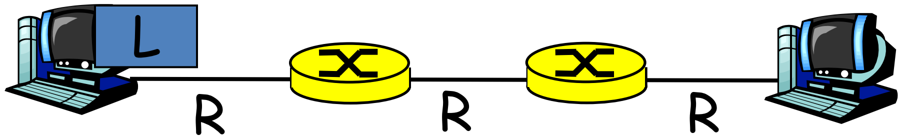

- Takes L/R seconds to transmit (push out) packet of L bits on to link of R bps
- Entire packet must arrive at router before it can be transmitted on next link: *store and forward*
- delay = 3L/R (assuming zero propagation delay)

Example:

- L = 7.5 Mbits
- R = 1.5 Mbps
- delay = 15 sec

more on delay shortly …

Packet-switched networks: forwarding

- *Goal:* move packets through routers from source to dest.
  - we'll study several path selection (routing) algorithms (chap 4)

- datagram network:
  - *destination address* in packet determines next hop
  - routes may change during session
  - analogy: driving, asking directions

- virtual circuit network:
  - 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*
  - (analogy: air trains in airports)

Compare

Thoughts on tradeoffs between packet switching and circuit switching?

Which one would you take?

Under what circumstances?

Why?

Packet switching versus Circuit switching

Packet switching allows more users to use network!

- problem: 1 Mbps link
- each user:
  - 100 kbps when "active"
  - active 10% of time

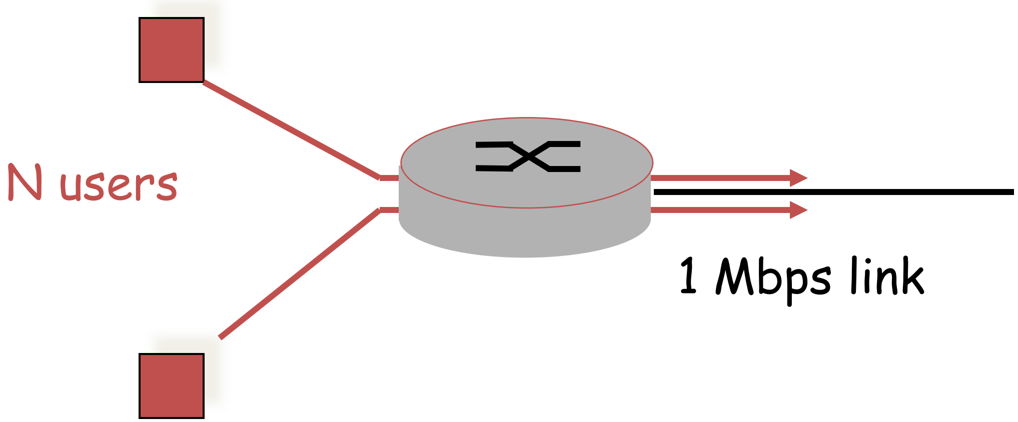

- circuit-switching:
  - 10 users
- packet switching (ps):
  - with 35 users,

probability > 10 active users is less than 0.0004

Q: how did we get value 0.0004?
Get performance of circuit switching with 3 times more users in case of PS

Factorials

- Denoted: n!
- Read: "n factorial"
- Definition:
  - n! = 1 if n = 0
  - = n (n – 1)! If n > 0
- $n! < n^n$
- How many different ways of arranging n distinct object into a sequence (called permutation of those objects)? n!

Combinations

- What if order *doesn't* matter?
- In poker, the following two hands are equivalent:
  - A♦, 5♥, 7♣, 10♠, K♠
  - K♠, 10♠, 7♣, 5♥, A♦
- The number of *r*-combinations of a set with *n* elements, where *n* is non-negative and 0≤*r*≤*n* is:

$$ {}^{n}C_{r} = C(n,r) = \frac{n!}{r!(n-r)!} $$

Binomial Distribution

- Binomial probability distributions allow us to deal with circumstances in which the outcomes belong to two relevant categories such as
  - success/failure or
  - acceptable/defective or
  - active/passive etc

Binomial Probability Formula

$$P(r) = {}^{n}C_{r}\, p^{r}(1-p)^{n-r} = \frac{n!}{r!(n-r)!}\, p^{r} q^{n-r}$$

for $r = 0, 1, 2, \ldots, n$

where

n = number of trials

r = number of successes among n trials

p = probability of success in any one trial

q = probability of failure in any one trial ($q = 1 - p$)

Problem on Circuit and Packet switching

- Suppose users share a 15 Mbps link. Also suppose each user requires 1 Mbps when transmitting, but each user transmit only 10% time.

a) When circuit switching is used, how many users can be supported?

b) Suppose there are 30 users. Find the probability that any given time, exactly 20 users are transmitting simultaneously. (Hint: Use the binomial distribution)

Packet switching versus Circuit switching

Is packet switching a "slam dunk winner?"

- Great for absorbing bursty data from individual sources
  - resource sharing (due to diversity)
  - simpler, no call setup

Why?

- Excessive congestion: packet delay and loss
  - protocols needed for reliability, congestion control

- Q: How to provide circuit-like behavior?
  - bandwidth guarantees needed for audio/video apps
  - still unsolved (chapter 7)

Network Taxonomy

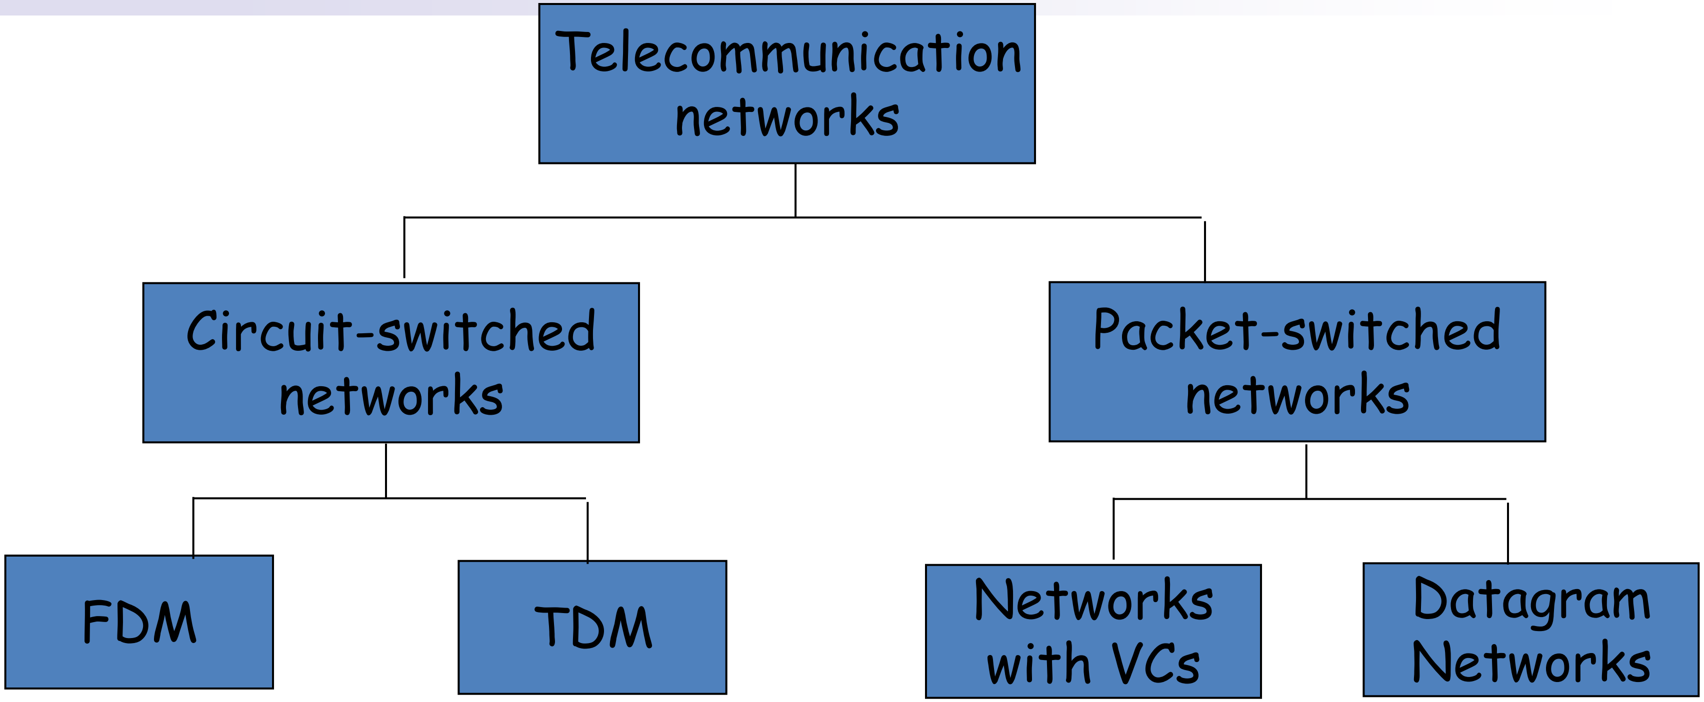

- Datagram network is *not* either connection-oriented or connectionless.
- Internet provides both connection-oriented (TCP) and connectionless services (UDP) to apps.
- Datagram service is a service provided by IP. It is a best effort, unreliable, message delivery service.

Chapter 1: Roadmap

1.1 What *is* the Internet?

1.2 Network edge

1.3 Network access and physical media

1.4 Network core

1.5 Internet structure and ISPs

1.6 Delay & loss in packet-switched networks

1.7 Protocol layers, service models

1.8 History

Internet Structure: Network of Networks

- roughly hierarchical
- at center: "tier-1" ISPs (e.g., MCI, Sprint, AT&T, Cable and Wireless), national/international coverage
  - treat each other as equals

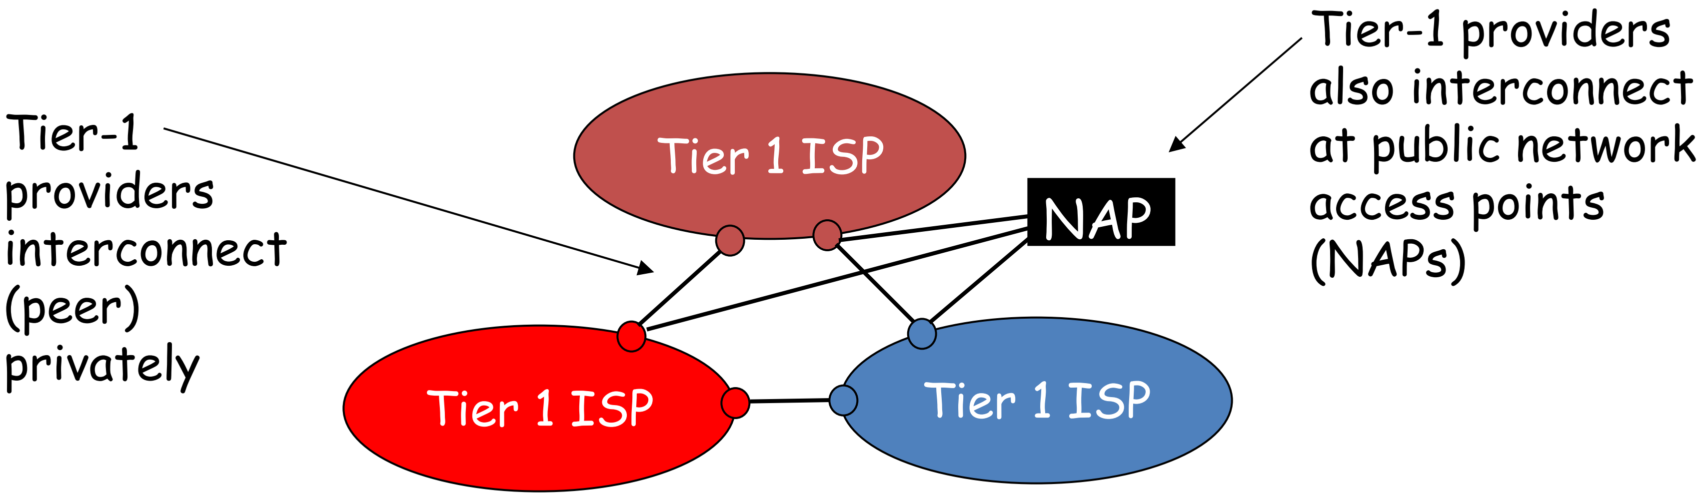

Tier-1 ISP: e.g., Sprint

Sprint US backbone network

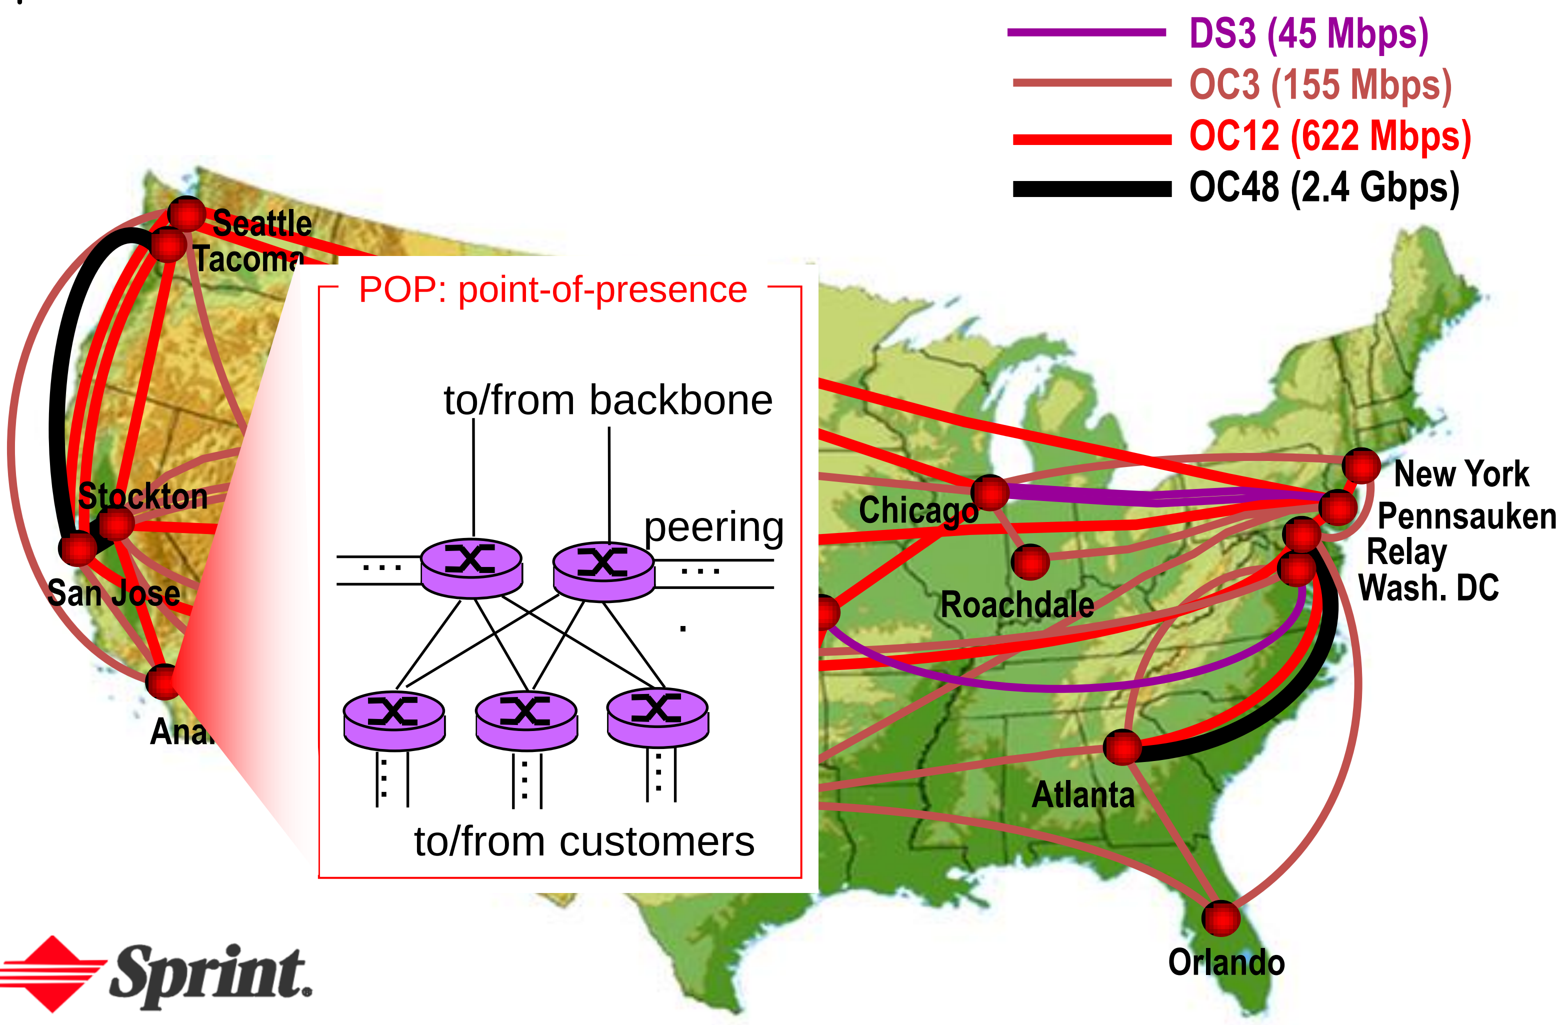

Internet Structure: Network of Networks

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

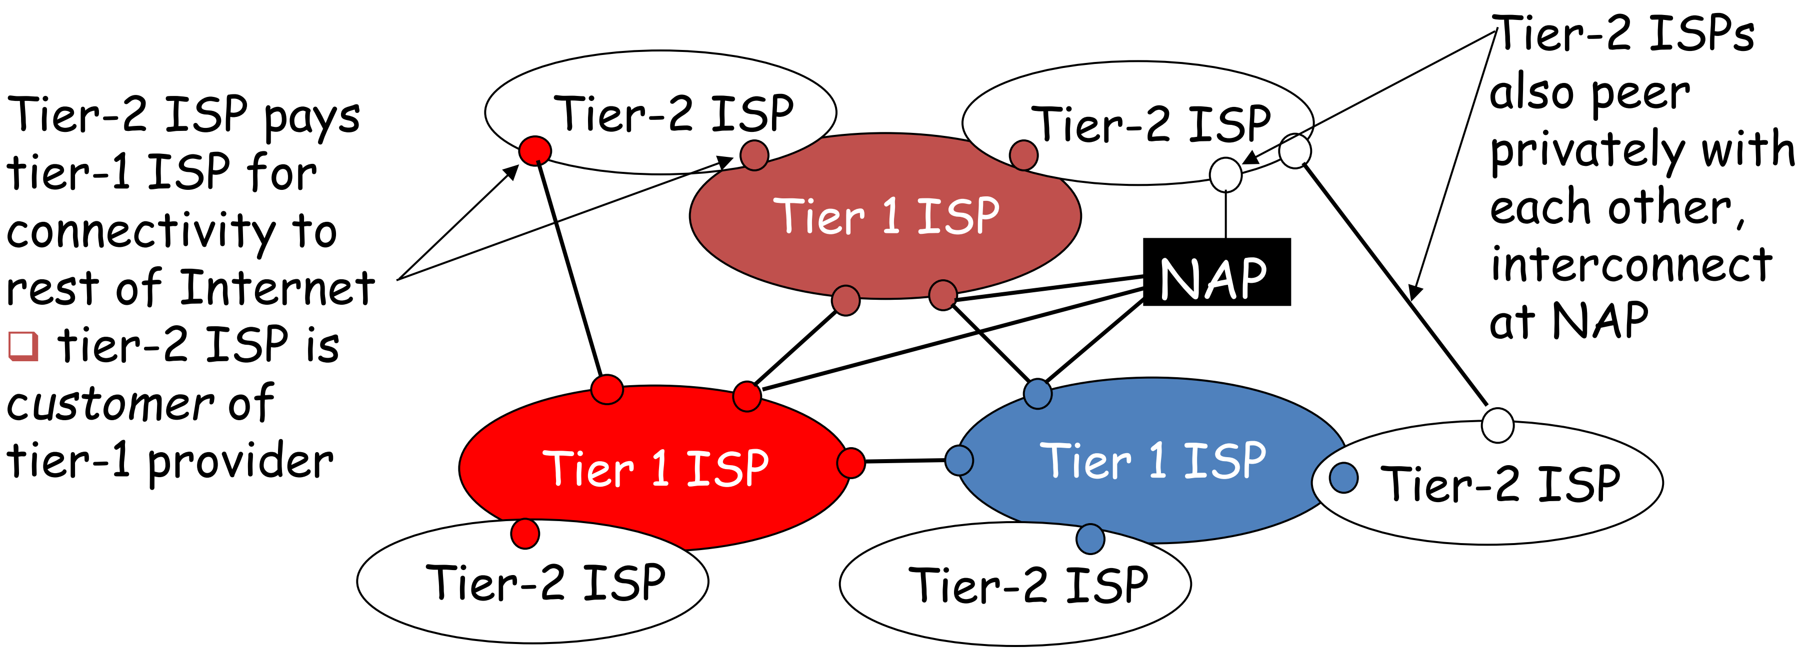

Internet Structure: Network of Networks

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

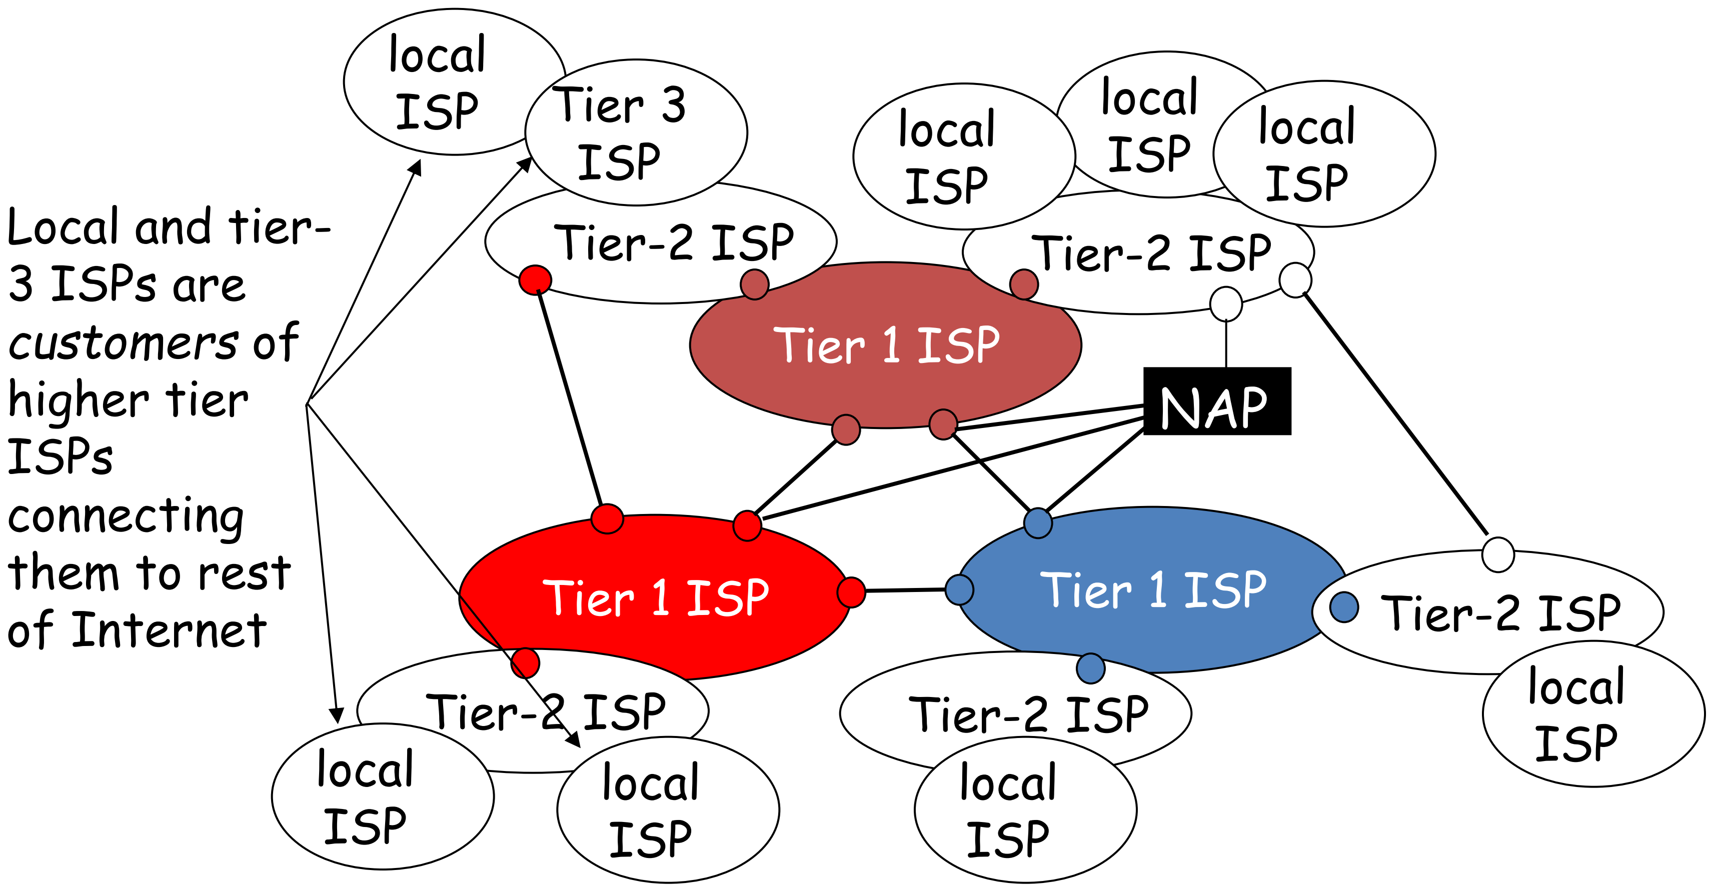

Internet Structure: Network of Networks

- a packet passes through many networks!

  local (taxi) → T1 (bus) → T2 (domestic) → T3 (international)

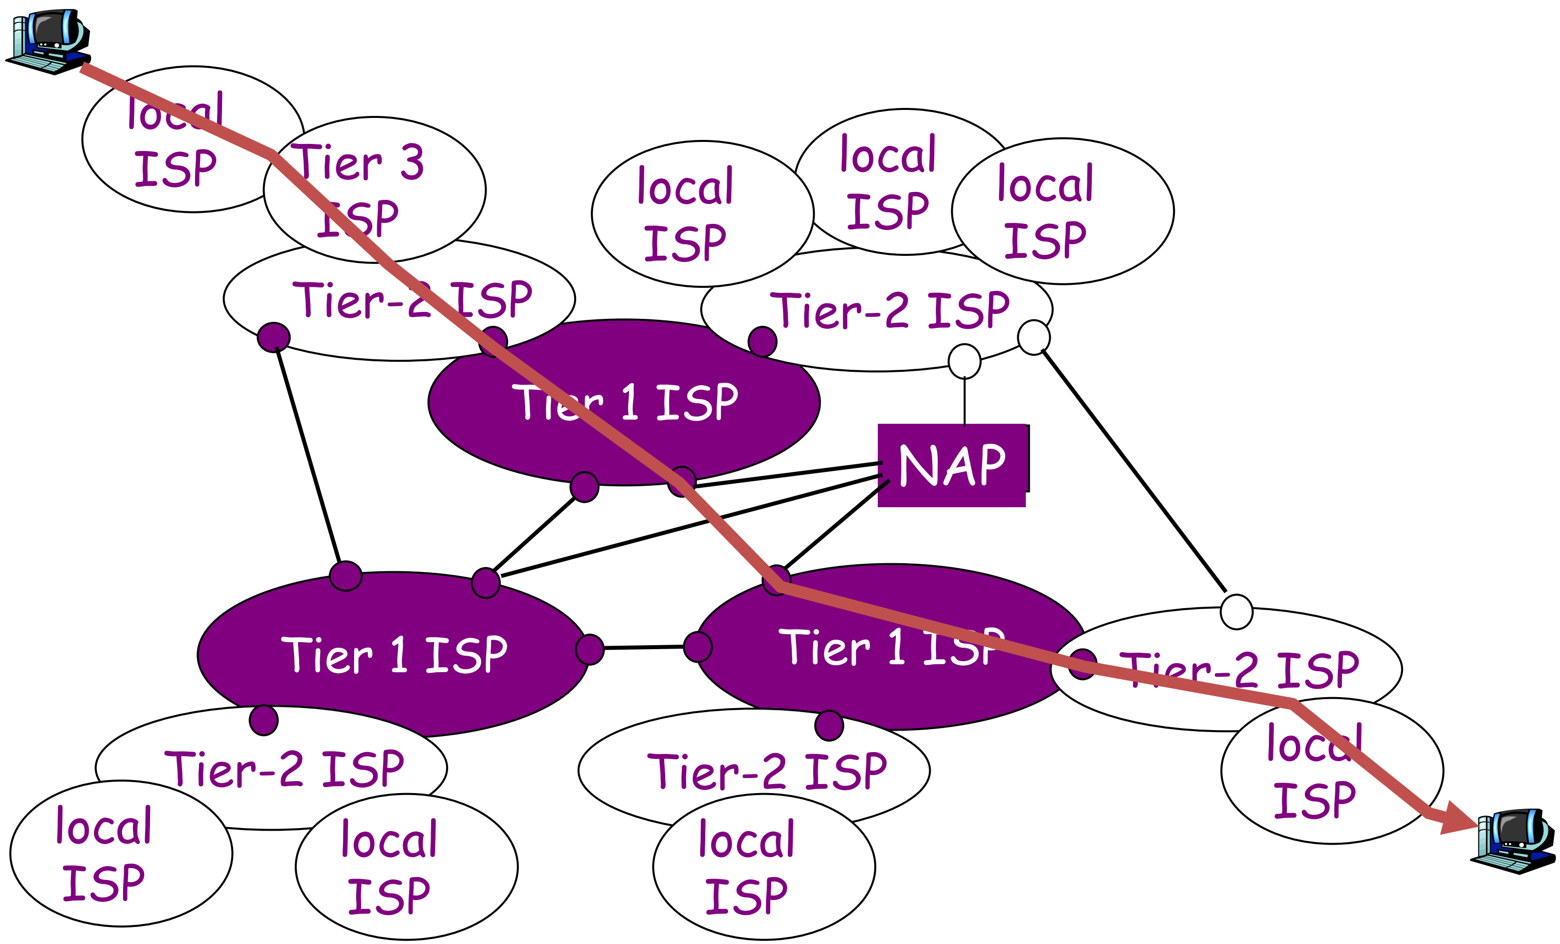

Chapter 1: Roadmap

1.1 What *is* the Internet?

1.2 Network edge

1.3 Network access and physical media

1.4 Network core

1.5 Internet structure and ISPs

1.6 Delay & loss in packet-switched networks

1.7 Protocol layers, service models

1.8 History

How do loss and delay occur?

packets *queue* in router buffers

- packet arrival rate to link exceeds output link capacity
- packets queue, wait for turn

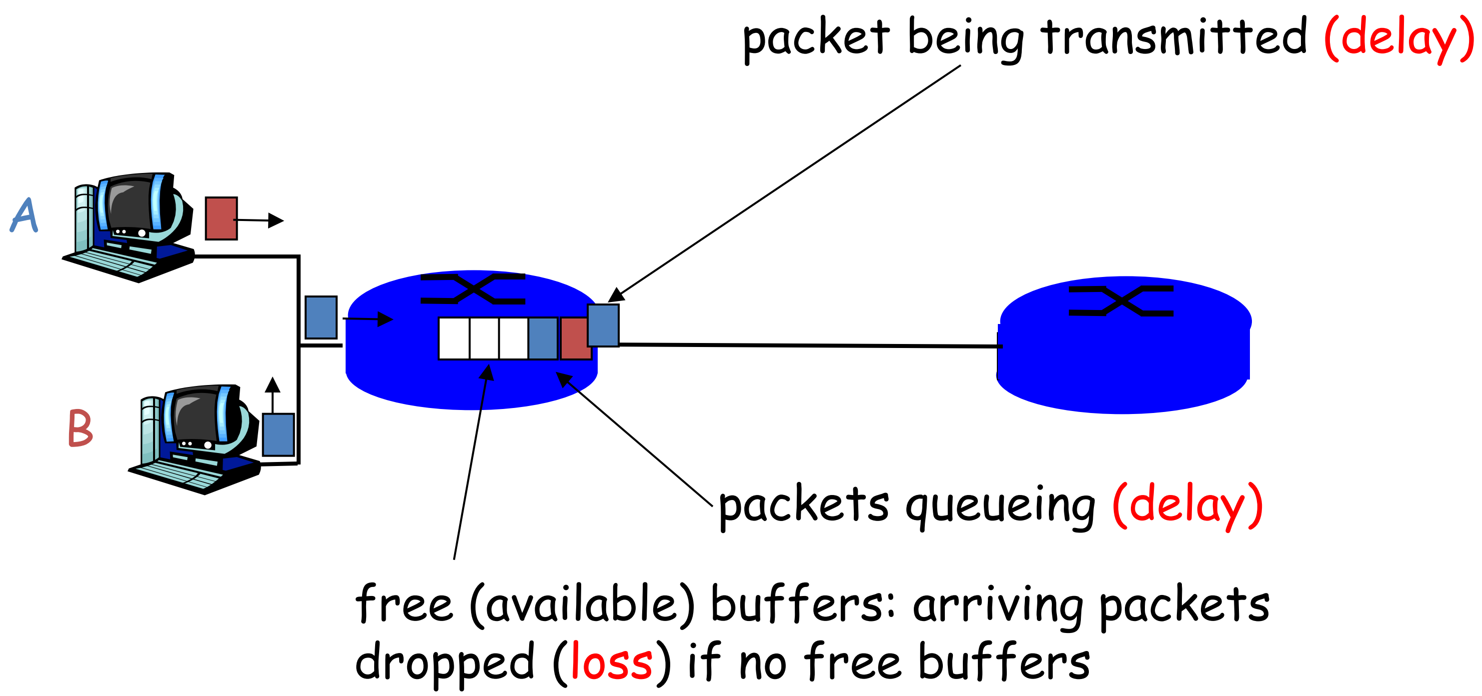

Four Sources of Packet Delay

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

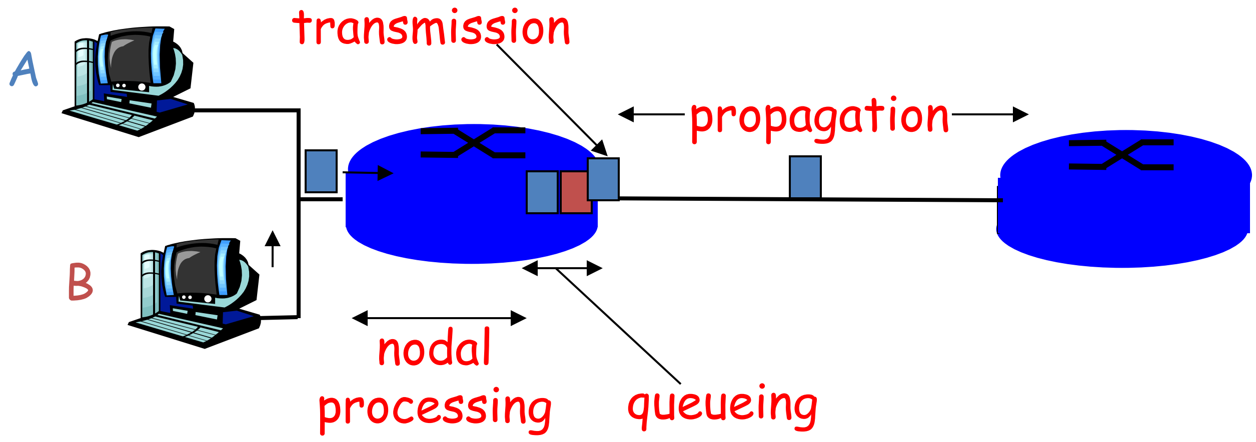

Delay in packet-switched networks

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 (~2x10^8 m/sec)
- propagation delay = d/s

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

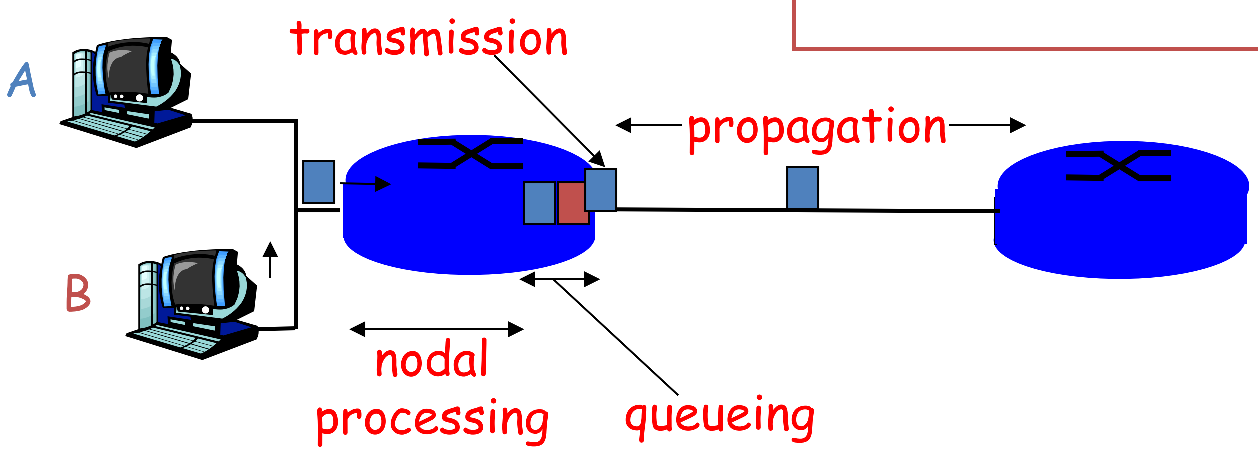

Comparing Transmission & Propagation Delays

- Transmission delay
  - Amount of time required to push out a packet
  - Function of the packet’s length & transmission rate of the link
  - Nothing to do with the distance between the two routers
- Propagation delay
  - Time it takes a bit to propagate from one router to the next
  - Function of the distance between two routers and propagation speed
  - Nothing to do with the packets’ length or transmission rate

Caravan analogy

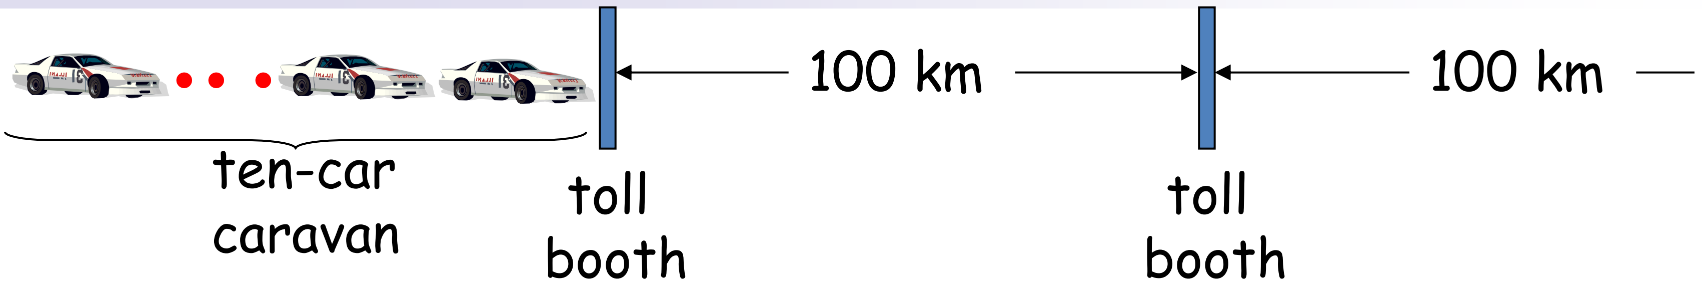

- Cars "propagate" at 100 km/hr
- Toll booth takes 12 sec to service a car (transmission time)
- car~bit; caravan ~ packet
- Q: How long until caravan is lined up before 2nd toll booth?

- Time to "push" entire caravan through toll booth onto highway = 12*10 = 120 sec
- Time for last car to propagate from 1st to 2nd toll both: 100km/(100km/hr)= 1 hr
- A: 62 minutes

Caravan analogy (more)

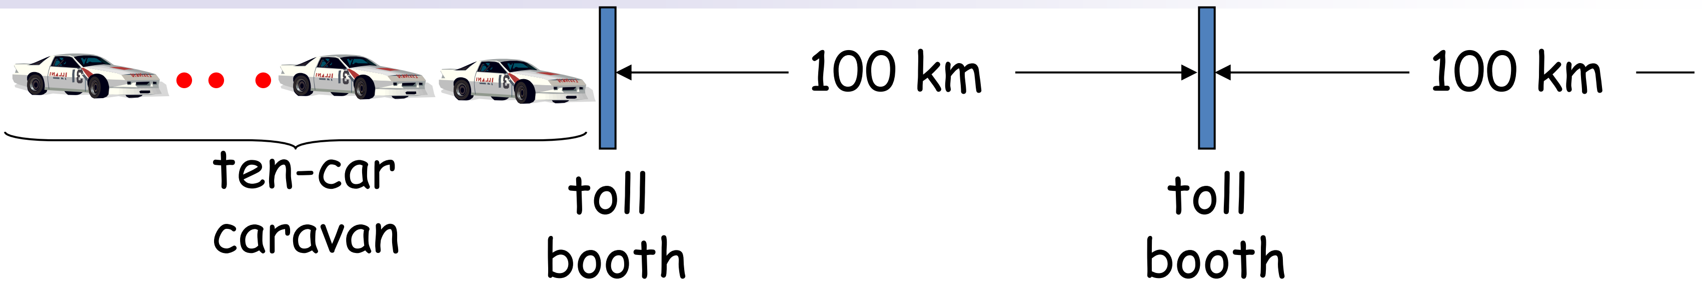

- Cars now "propagate" at 1000 km/hr
- Toll booth now takes 1 min to service a car
- Q: Will cars arrive to 2nd booth before all cars serviced at 1st booth?

- Yes! After 7 min, 1st car at 2nd booth and three cars still at 1st booth.
- 1st bit of packet can arrive at 2nd router before packet is fully transmitted at 1st router!

Nodal delay

$$d_{\text{nodal}} = d_{\text{proc}} + d_{\text{queue}} + d_{\text{trans}} + d_{\text{prop}}$$

- d_{proc} = processing delay
  - typically a few microsecs or less
- d_{queue} = queuing delay
  - depends on congestion
- d_{trans} = transmission delay
  - = L/R, significant for low-speed links
- d_{prop} = propagation delay
  - a few microsecs to hundreds of msecs

Queueing delay (revisited)

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

traffic intensity = La/R

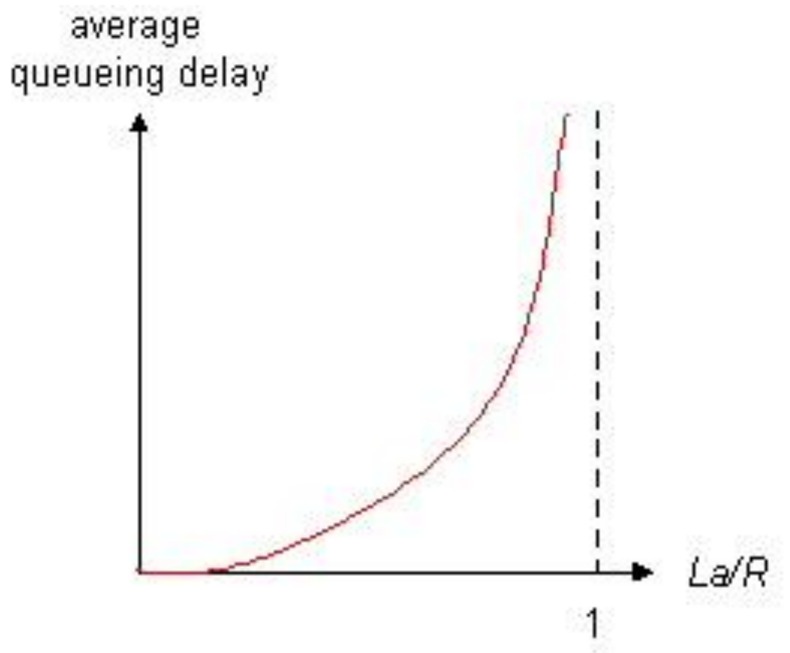

- La/R ~ 0: average queueing delay small
- La/R -> 1: delays become large
- La/R > 1: more "work" arriving than can be serviced, average delay infinite!

Packet loss

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

“Real” Internet delays and routes

- 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

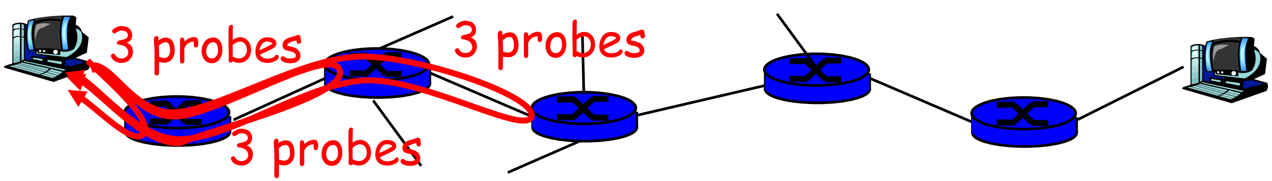

"Real" Internet delays and routes

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 response (probe lost, router not replying)

http://www.traceroute.org

Chapter 1: Roadmap

1.1 What *is* the Internet?

1.2 Network edge

1.3 Network access and physical media

1.4 Network core

1.5 Internet structure and ISPs

1.6 Delay & loss in packet-switched networks

1.7 Protocol layers, service models

1.8 History

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

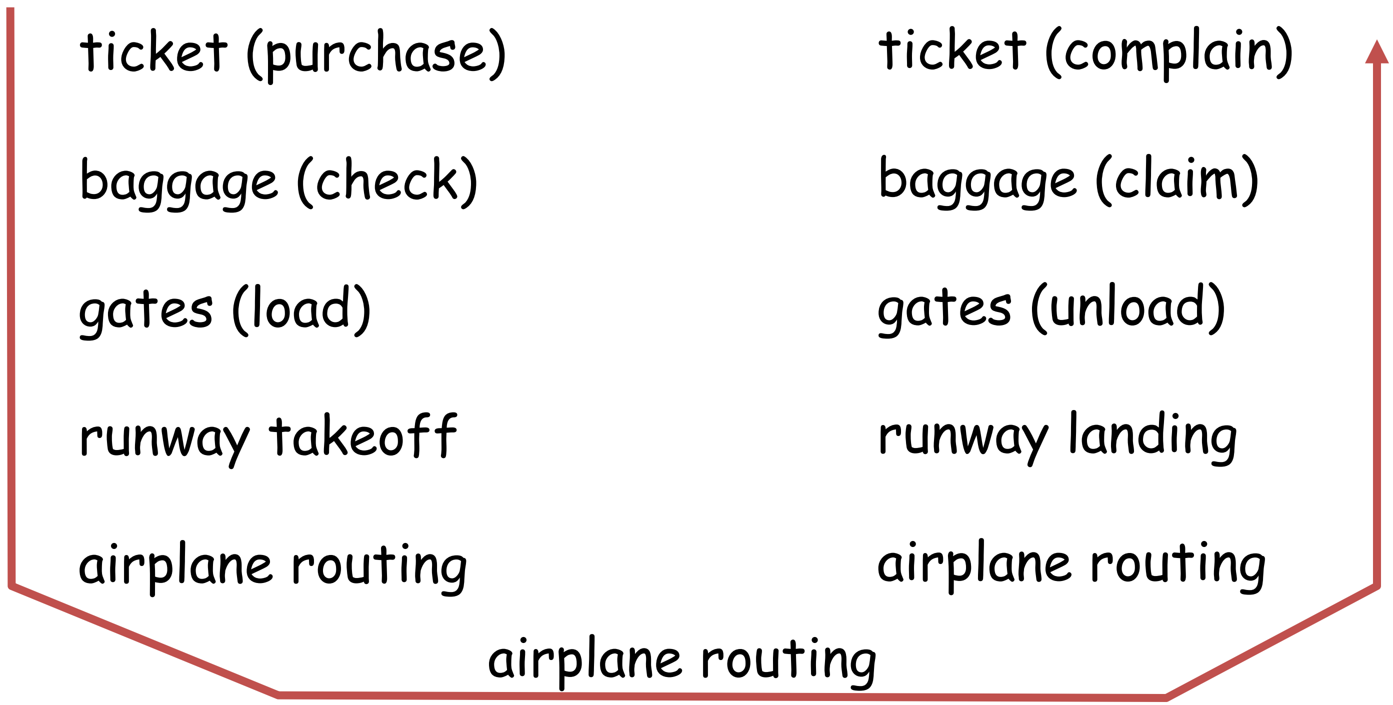

- a series of steps

Layering of airline functionality

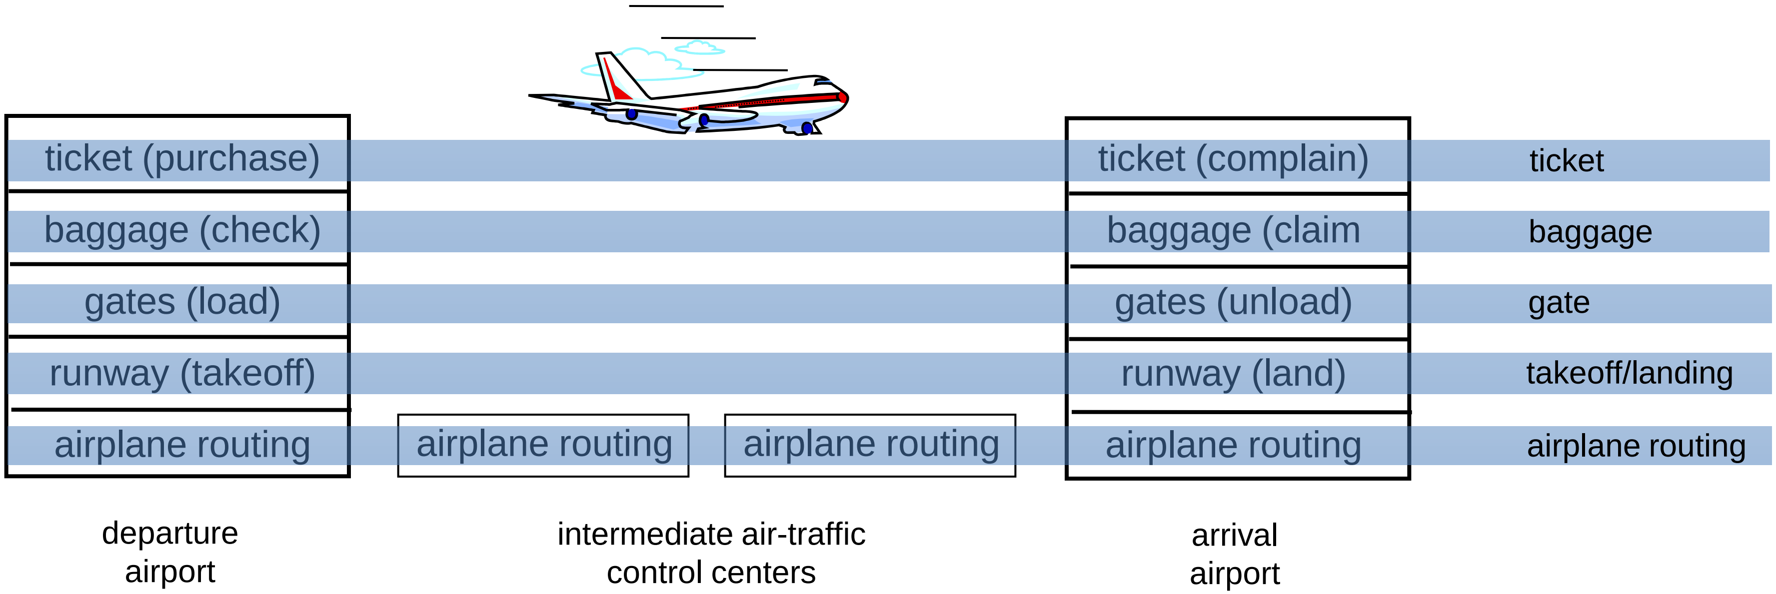

Layers: each layer implements a service

- Same layers communicate
  - Baggage section of BWI only calls baggage section of PHL
- Layers rely on services provided by layer below

Internet protocol stack

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

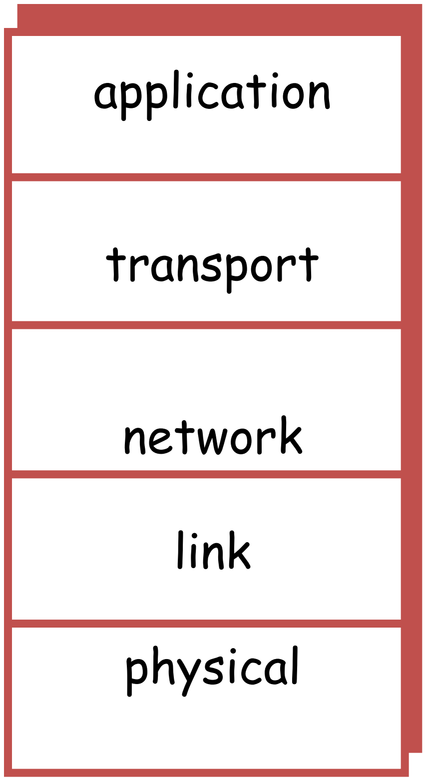

Encapsulation

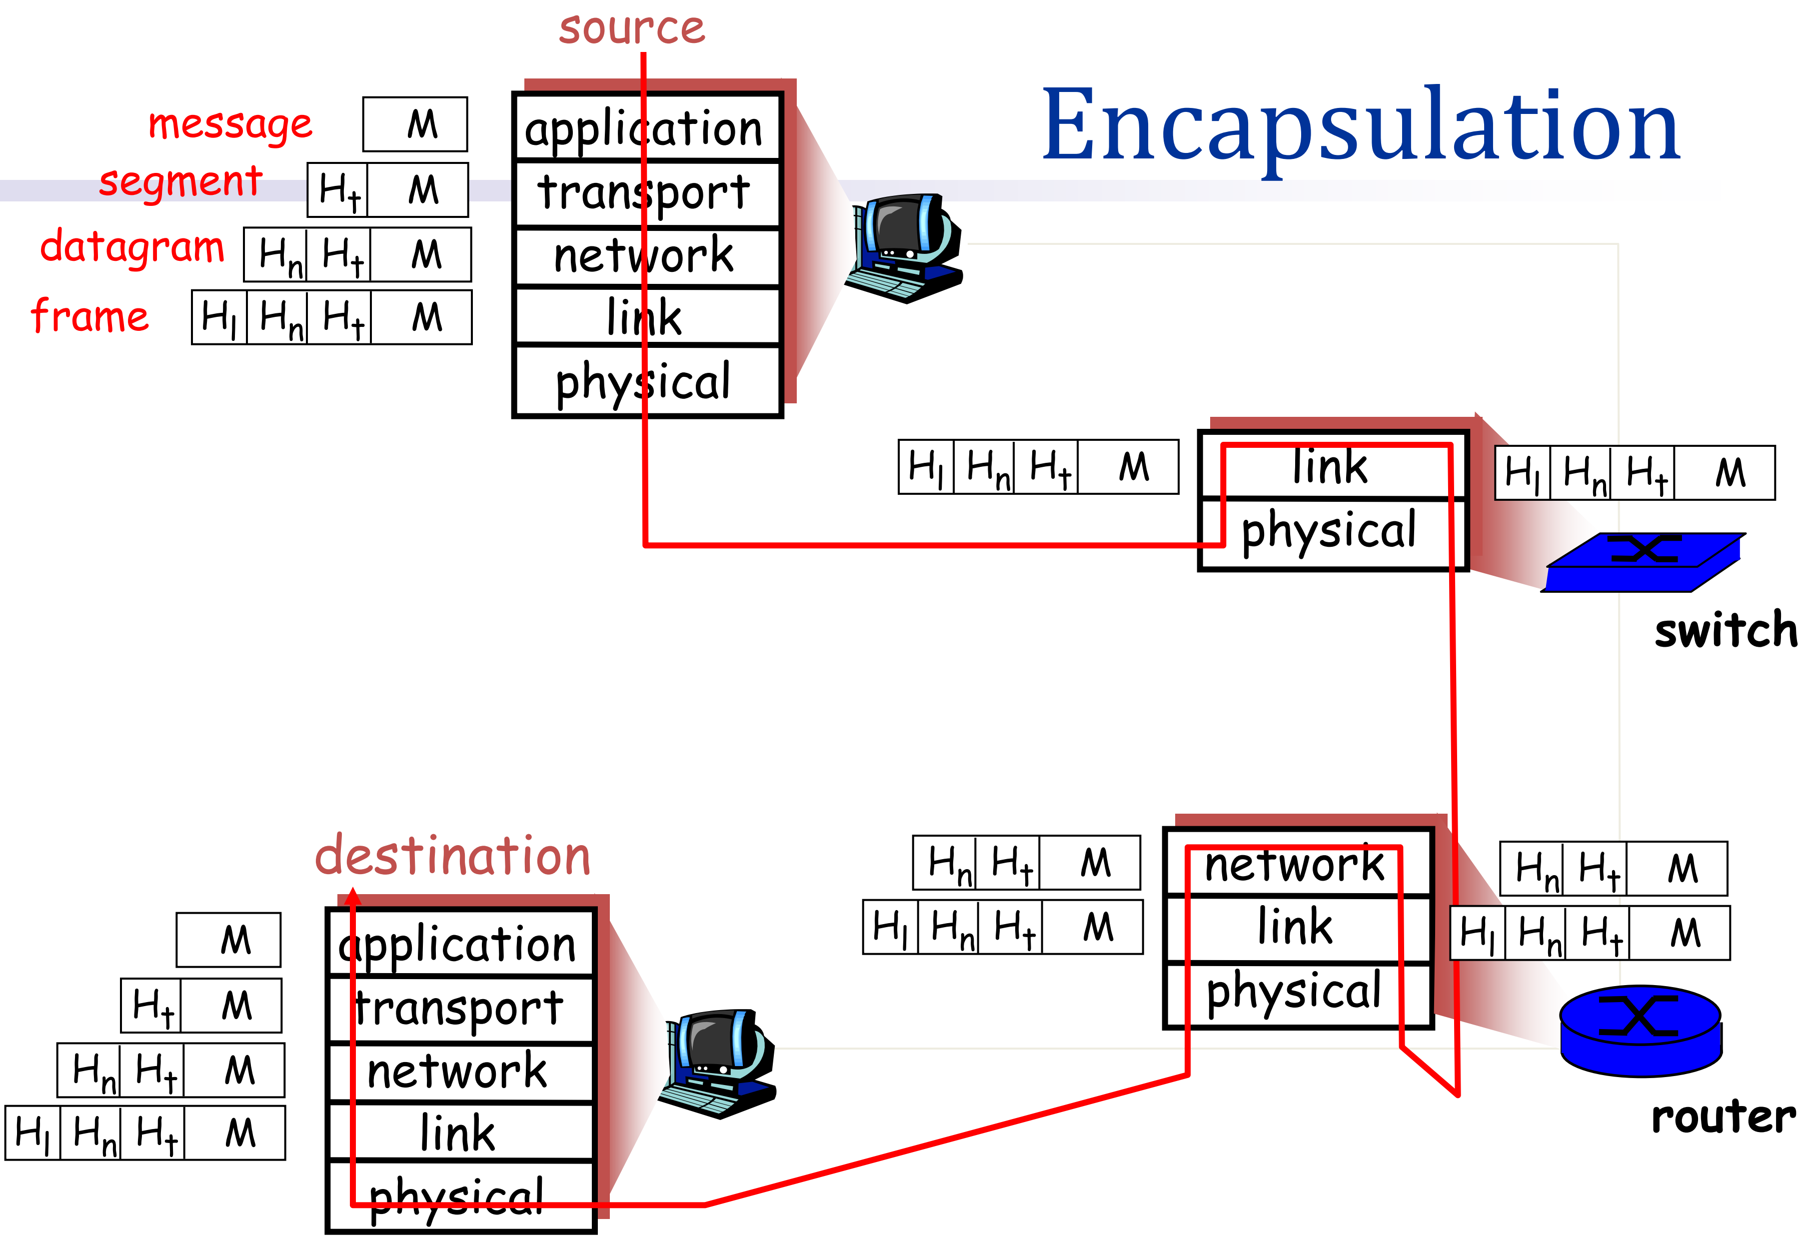

Introduction: Summary

Covered a "ton" of material!

- Internet overview
- what's a protocol?
- network edge, core, access network
  - packet-switching versus circuit-switching
- Internet/ISP structure
- performance: loss, delay
- layering and service models

You now have:

- context, overview, "feel" of networking
- more depth, detail *to follow!*

Questions

?