Multimedia Communications
@CS.NCTU

Lecture 1: Networking Overview

Instructor: Kate Ching-Ju Lin (林靖茹)

Slides modified from
"Computer Networking: A Top-Down Approach" 6th Edition

Outline

- What's the Internet?
- What's a protocol?
- Network edge
  - hosts, access net, physical media
- Network core
  - packet/circuit switching, Internet structure
- Performance
  - loss, delay, throughput
- Protocol layers, service models
- History

Outline

- **What's the Internet?**
- What's a protocol?
- Network edge
    - hosts, access net, physical media
- Network core
    - packet/circuit switching, Internet structure
- Performance
    - loss, delay, throughput
- Protocol layers, service models
- History

What is the Internet?

Two types of description:

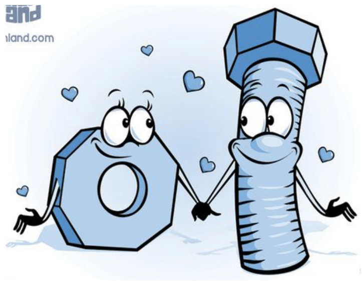

- **Nuts and bolts of the Internet**
  - i.e., hardware and software components
  - from the structure perspective

- **An infrastructure that provides services to applications**
  - including email, Web, games, P2P, VoIP, streaming, social networking, messaging, etc
  - from the functionality perspective

“Nuts and Bolts” View

- billions of connected computing devices:

- hosts = end systems
  - running network apps

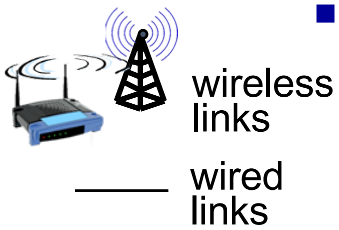

- communication links
  - fiber, copper, radio, satellite
  - transmission rate: bandwidth

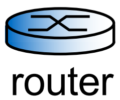

- packet switches: forward packets (chunks of data)
  - routers and switches

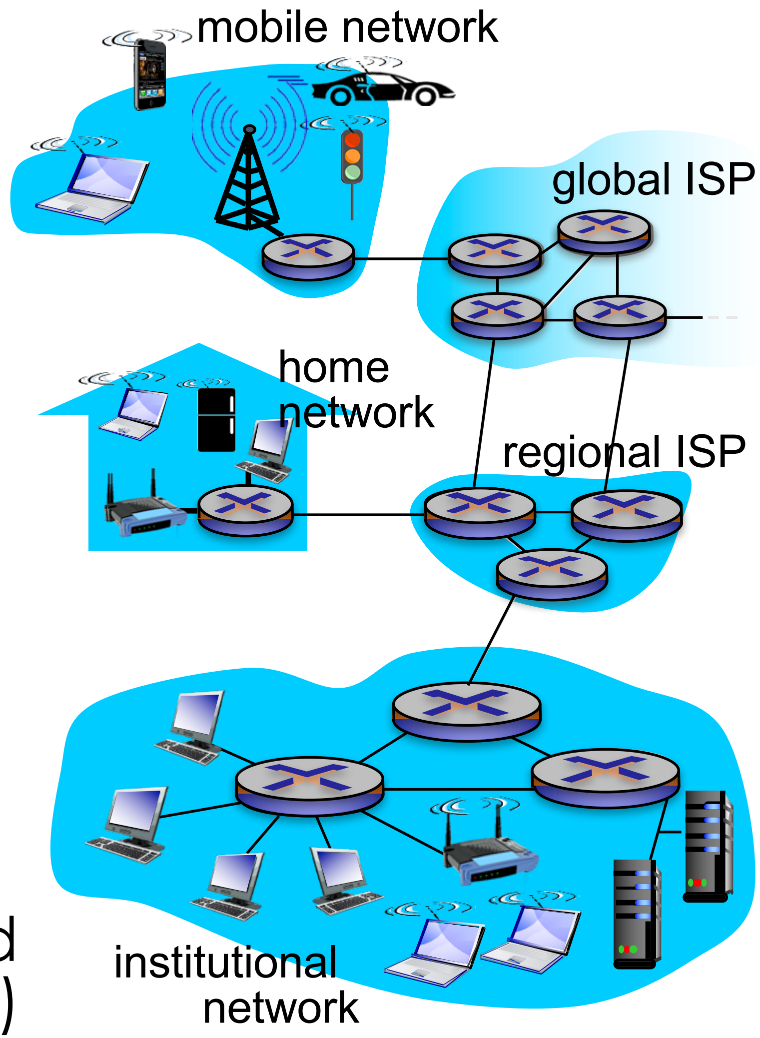

"Fun" Internet-Connected Devices

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

Web-enabled toaster +
weather forecaster

Tweet-a-watt:
monitor energy use

Internet
refrigerator

Slingbox: watch,
control cable TV remotely

sensorized,
bed
mattress

Internet phones

"Nuts and Bolts" View

- Internet: "network of networks"
  - Interconnected ISPs (Internet Service Providers)
- Protocols control sending, receiving of messages
  - e.g., TCP, IP, HTTP, Skype, 802.11
- Internet standards
  - RFC: Request for comments
  - IETF: Internet Engineering Task Force

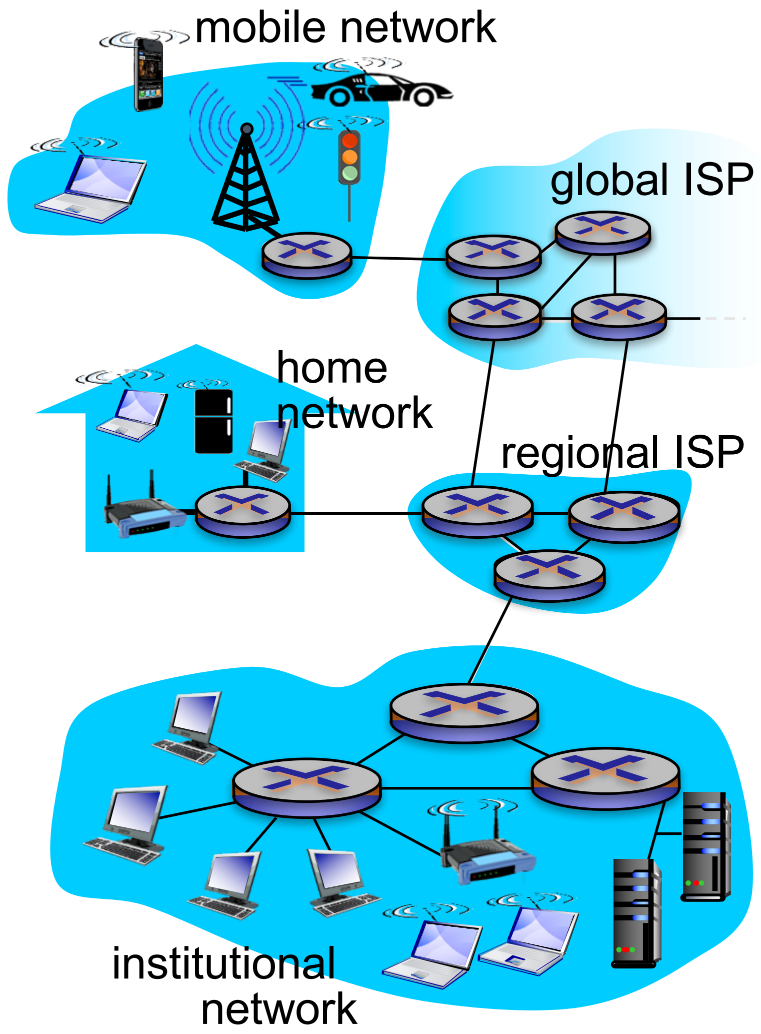

"Service" View

- Infrastructure that provides services to applications:
  - Web, VoIP, email, games, e-commerce, social nets, …
- Provide programming interface to apps
  - hooks that allow sending and receiving app programs to "connect" to Internet
  - provide service options, analogous to postal service

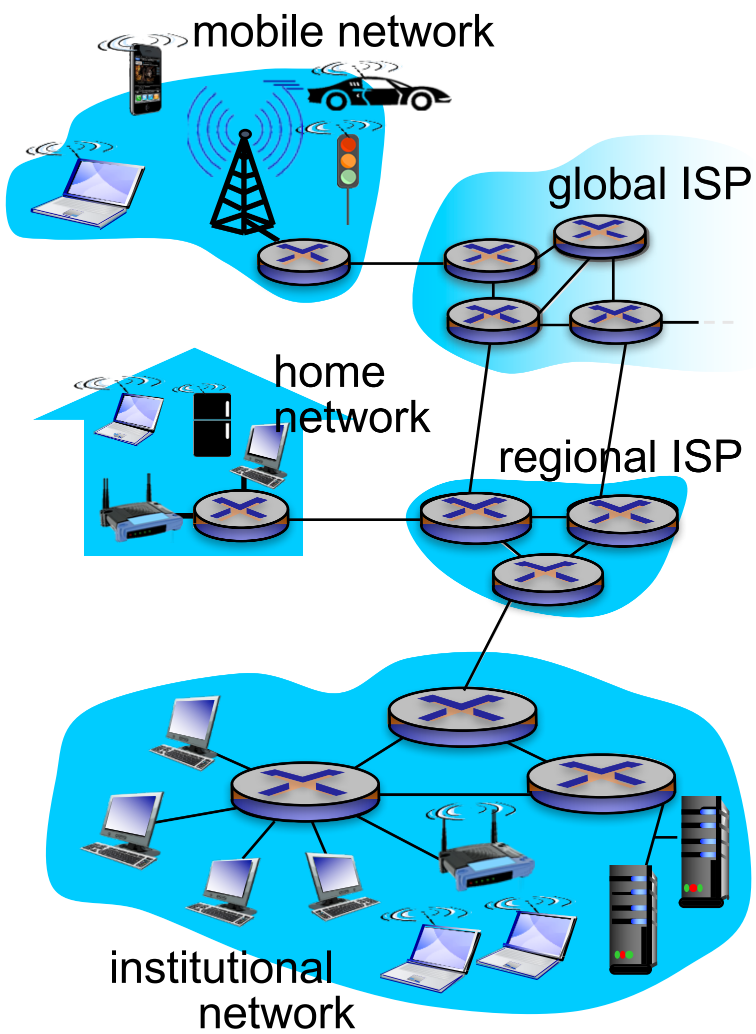

Outline

- What's the Internet?
- **What's a protocol?**
- Network edge
  - hosts, access net, physical media
- Network core
  - packet/circuit switching, Internet structure
- Performance
  - loss, delay, throughput
- Protocol layers, service models
- History

What's a Protocol?

human protocols:

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

... specific messages sent

... specific actions taken when messages received, or other events

network protocols:

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

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

What's a Protocol?

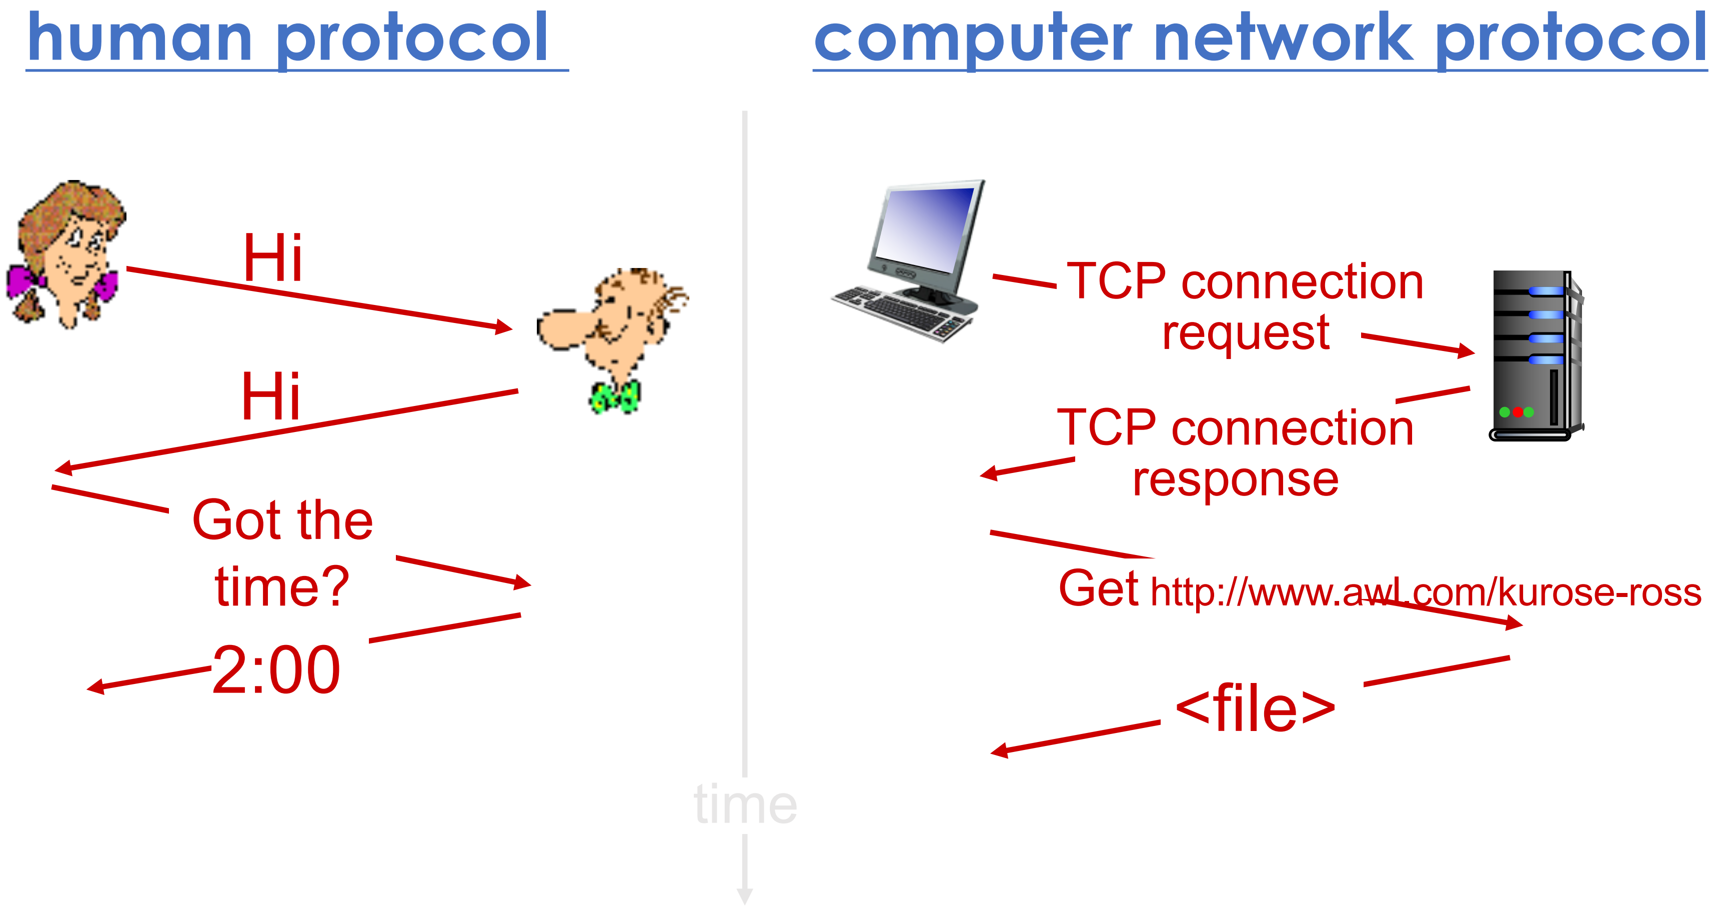

Q: other human protocols?

Outline

- What's the Internet?
- What's a protocol?
- **Network edge:**
  - hosts, access net, physical media
- Network core
  - packet/circuit switching, Internet structure
- Performance
  - loss, delay, throughput
- Protocol layers, service models
- History

Network Structure

- **Network edge:**
  - hosts: clients and servers
  - servers often in data centers
- **Access networks, physical media:**
  - Connect hosts to first routers (edge routers)
- **Network core:**
  - interconnected routers
  - network of networks

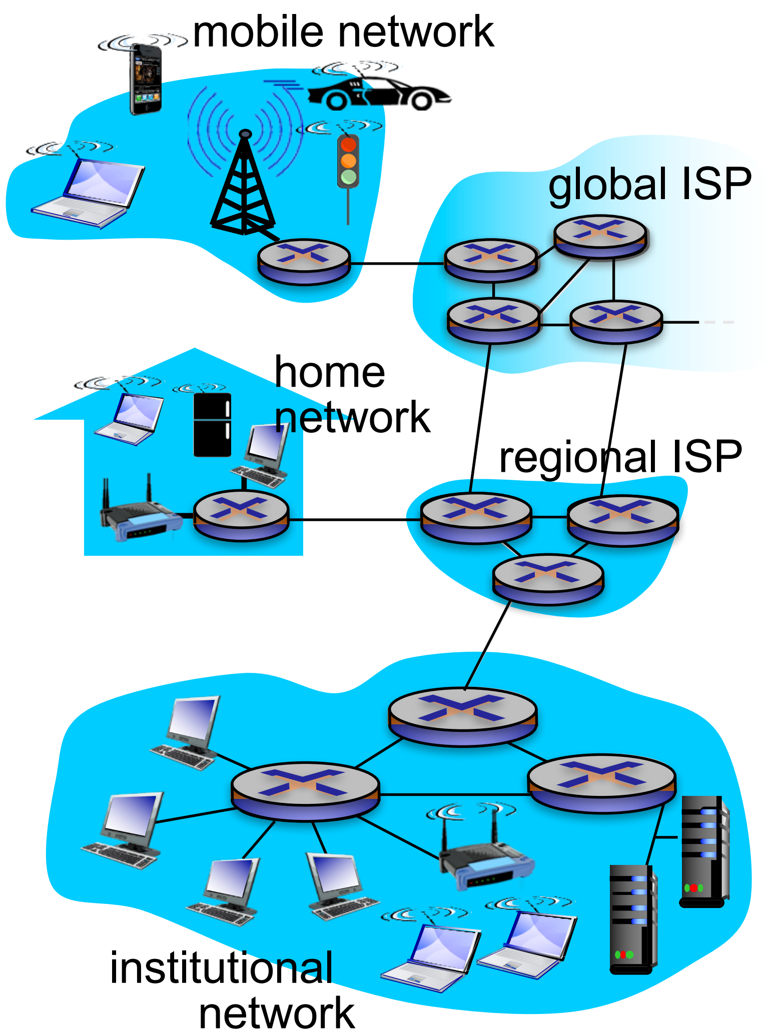

Access Networks and Physical Media

- How to connect end systems to edge router?
  - residential access nets
  - institutional access networks (school, company)
  - mobile access networks

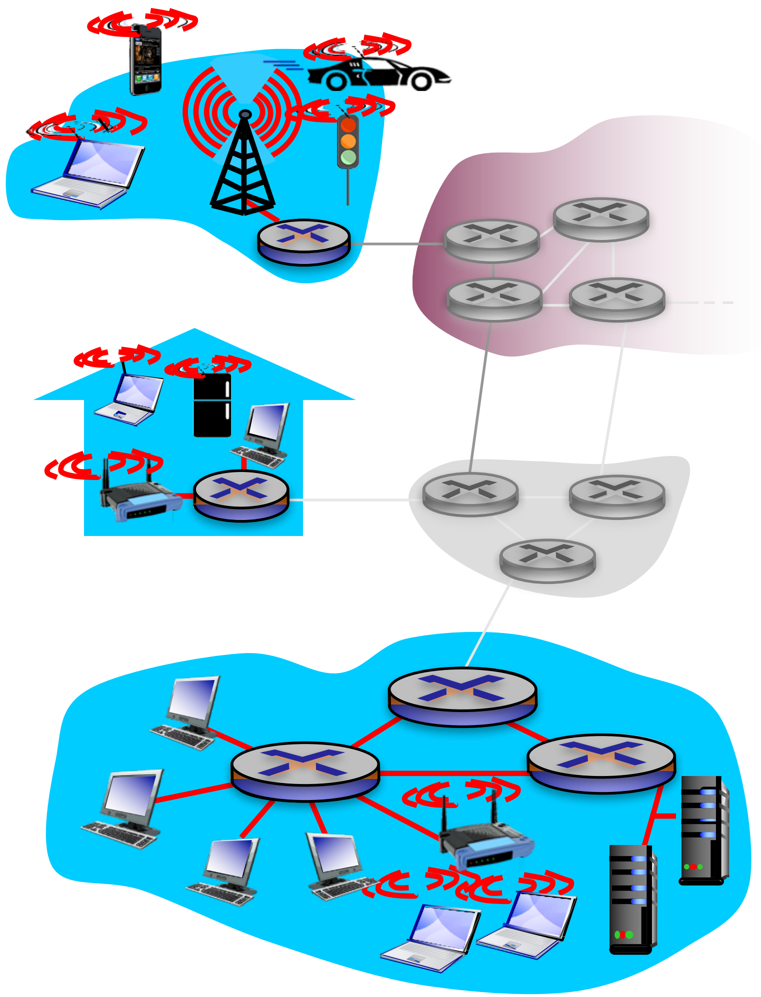

keep in mind

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

Access Network: DSL

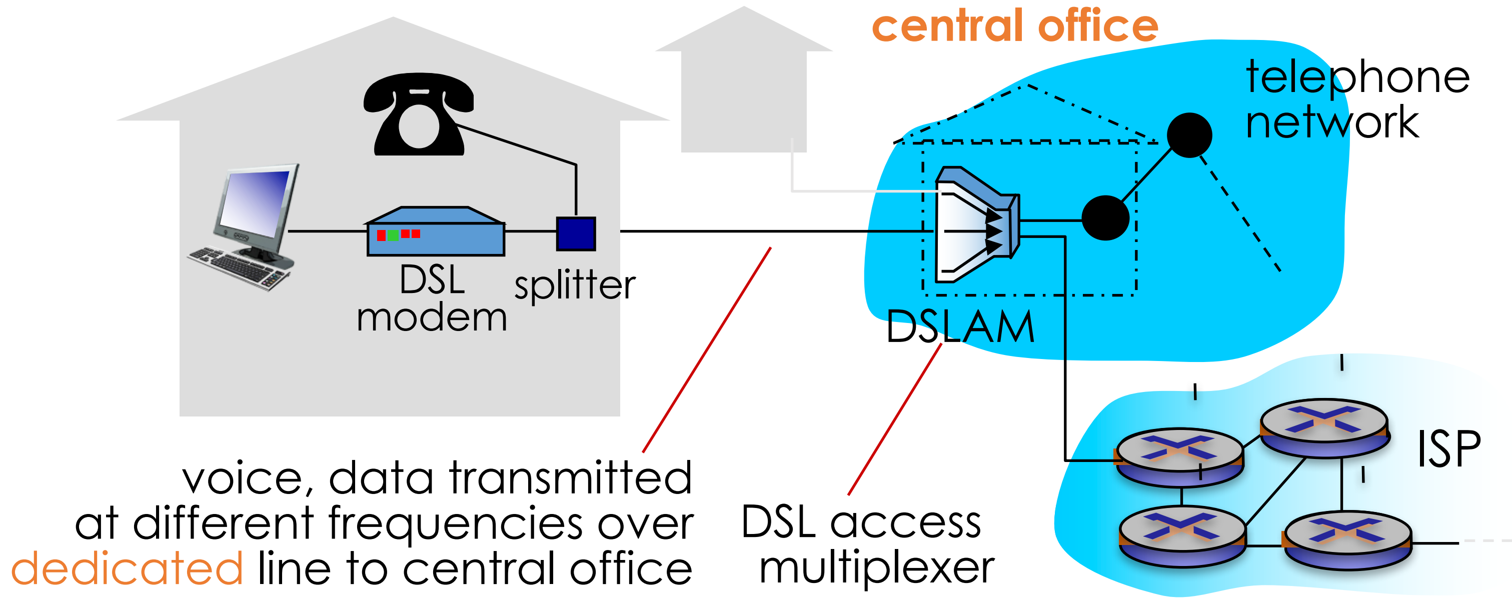

- Use existing telephone line to central office DSLAM (Digital Subscriber Line Access Multiplexer)
  - data over DSL phone line goes to Internet
  - voice over DSL phone line goes to telephone net
- < 2.5 Mbps for upstream (typically < 1 Mbps)
- < 24 Mbps for downstream (typically < 10 Mbps)

→ Asymmetric!

Access Network: Cable Network

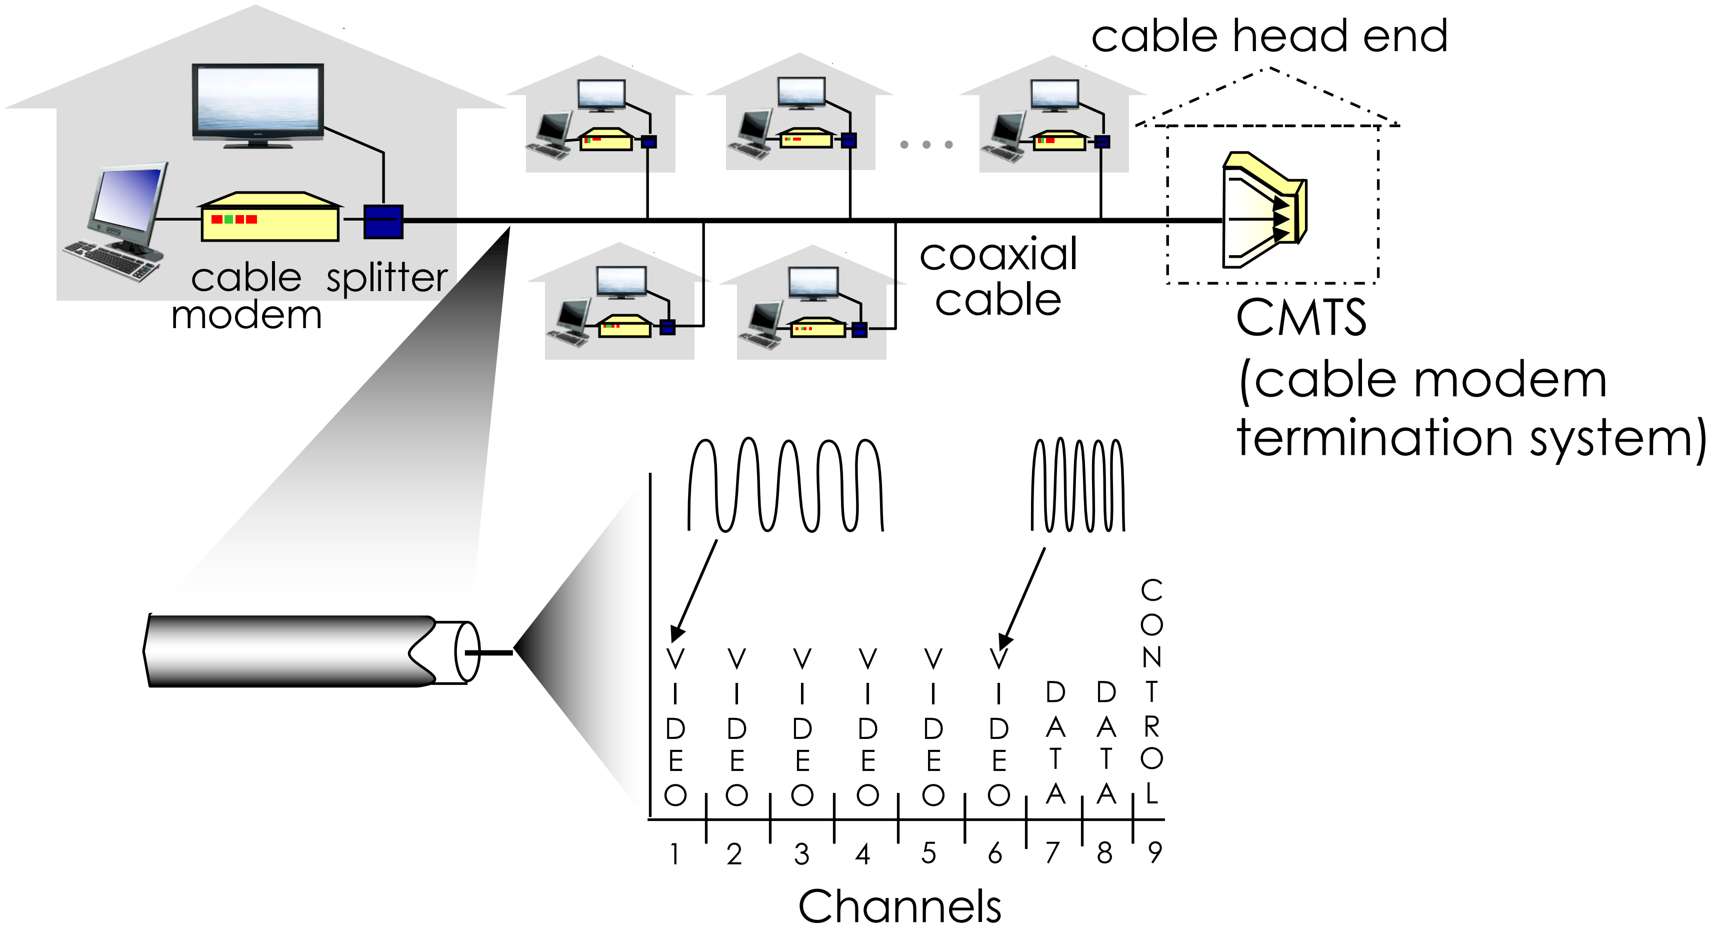

- **frequency division multiplexing**
  - different channels transmitted in different frequency bands

Access Network: Home Network

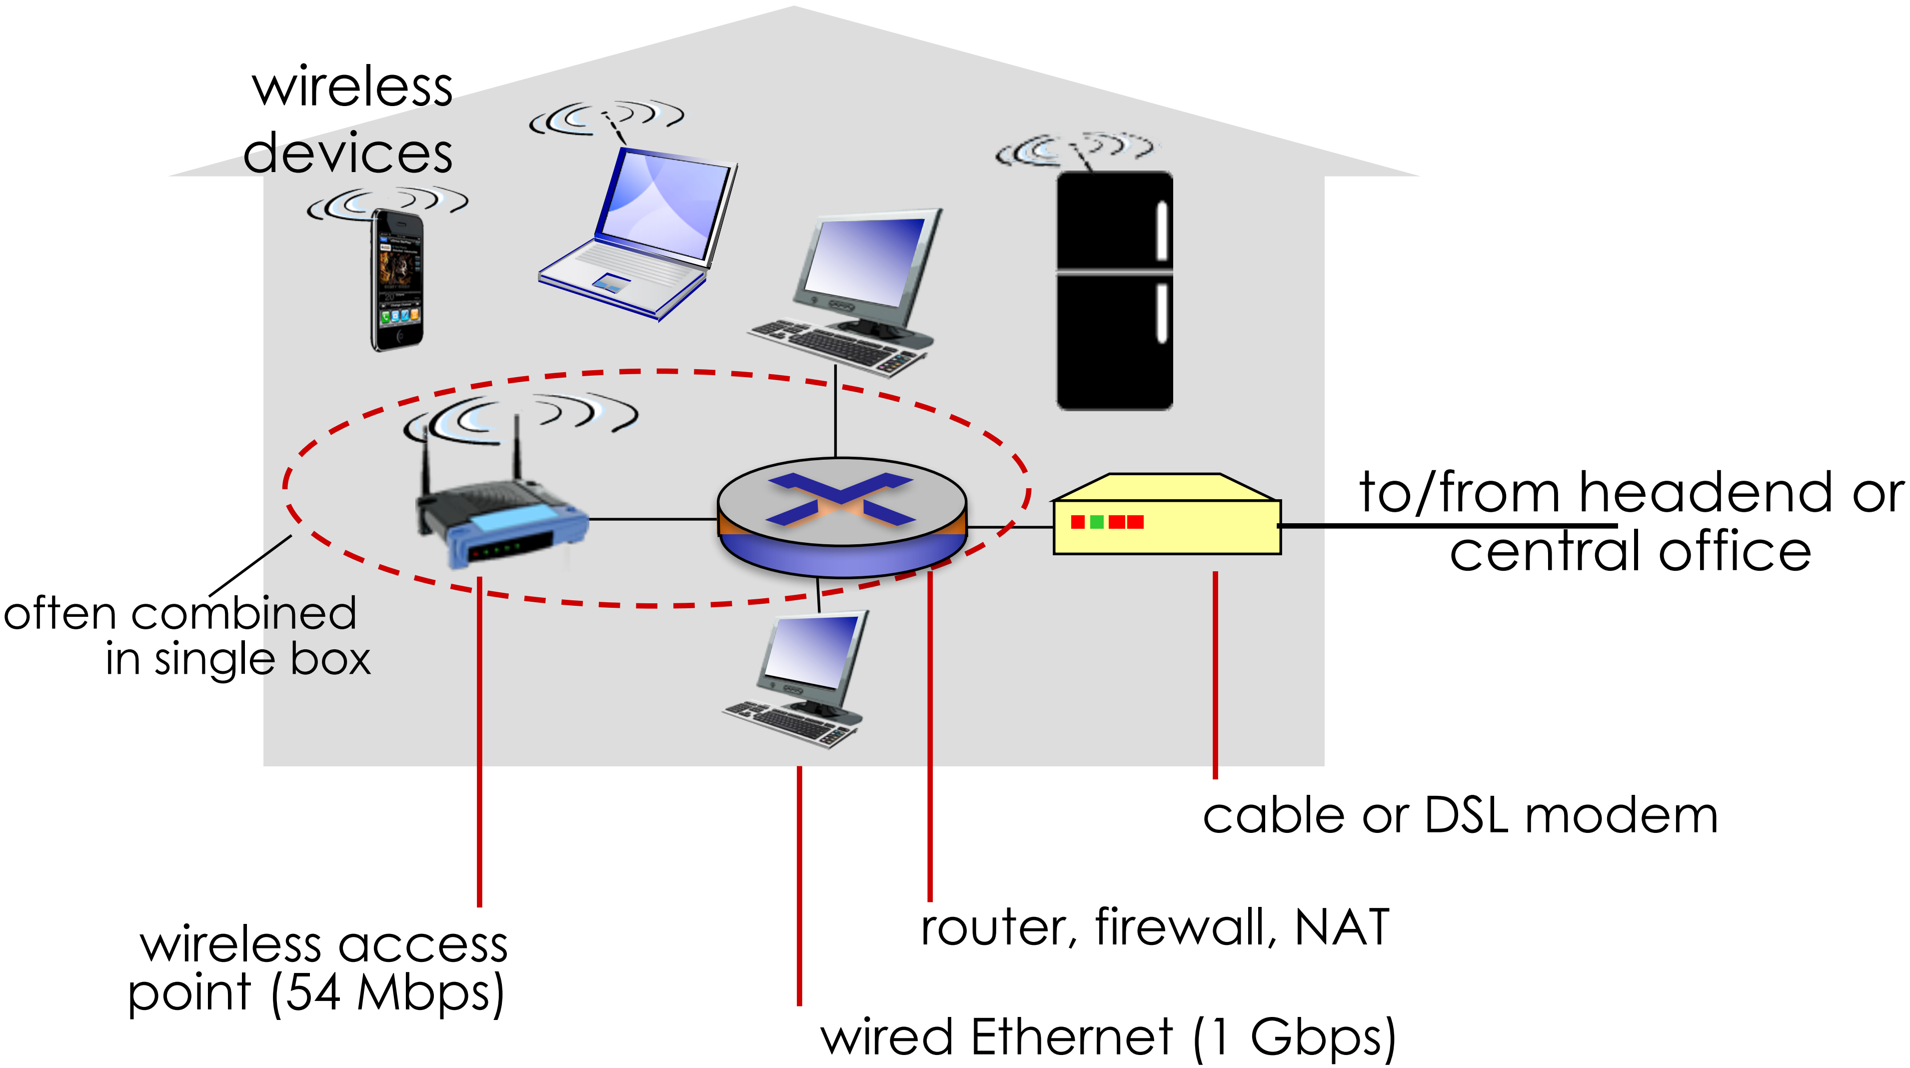

Enterprise Access Networks (Ethernet)

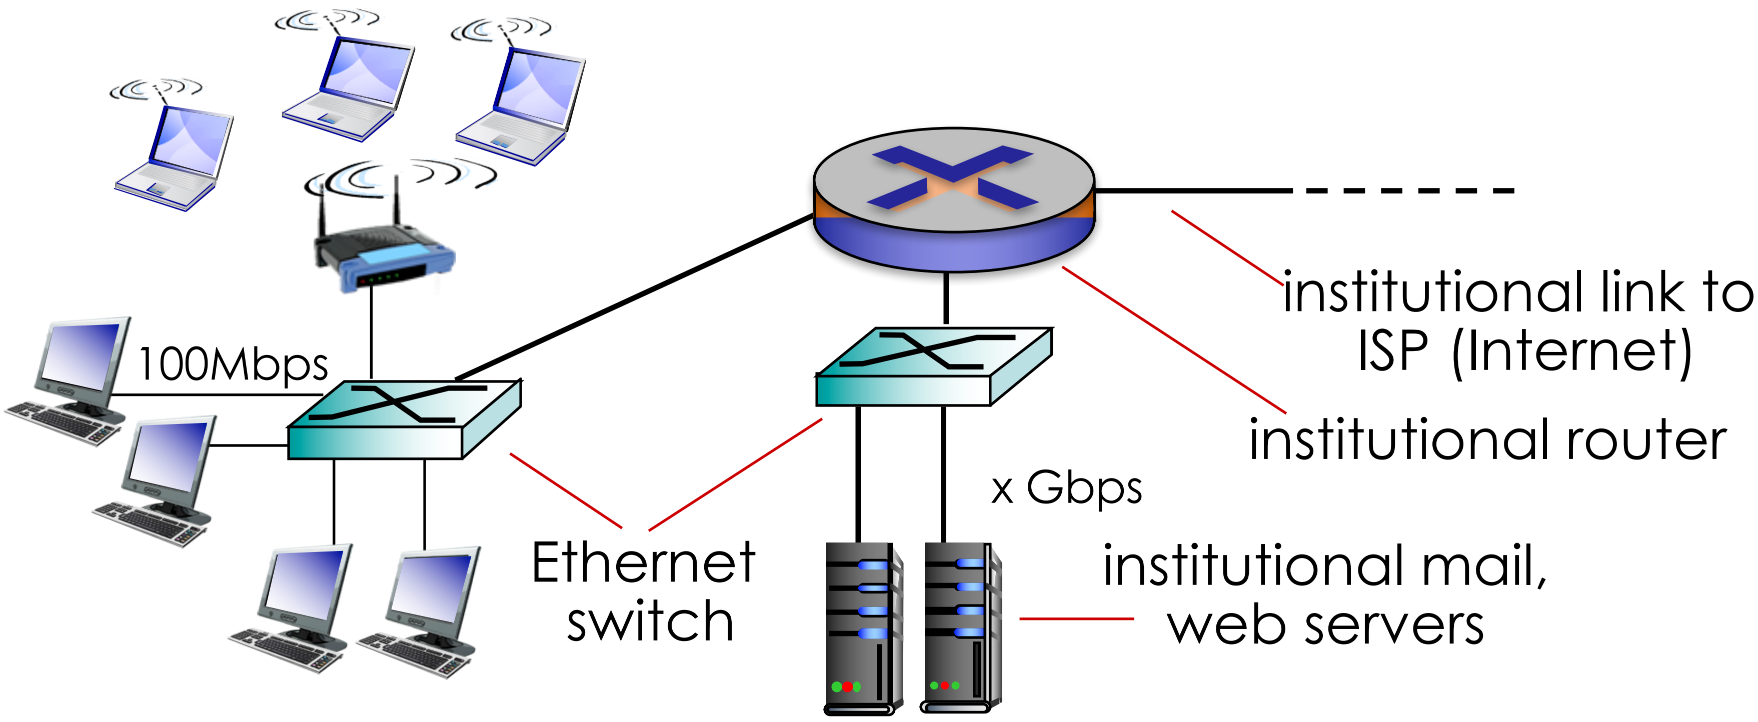

- Typically used in companies, universities, etc.
  - 10 Mbps, 100Mbps, 1Gbps, 10Gbps transmission rates
  - Today, end systems typically connect into Ethernet switch

Wireless Access Networks

- Shared wireless access network connects end system to routers
  - via base station aka “access point” (AP)

wireless LANs

- within building (100 ft.)
- 802.11b/g/n (WiFi): 11, 54, 450 Mbps transmission rate

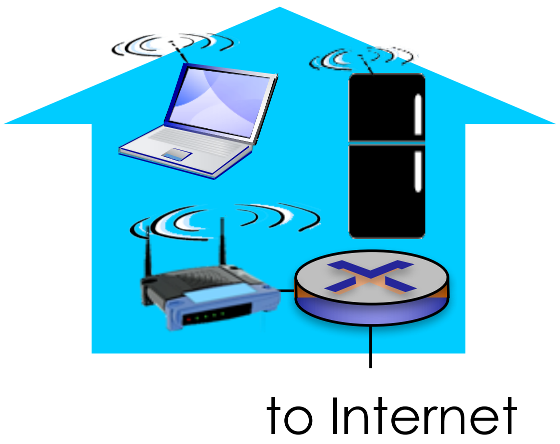

wide-area wireless access

- provided by telco (cellular) operator, 10’s km
- between 1 and 10 Mbps
- 3G, 4G: LTE (Long-Term Evolution)

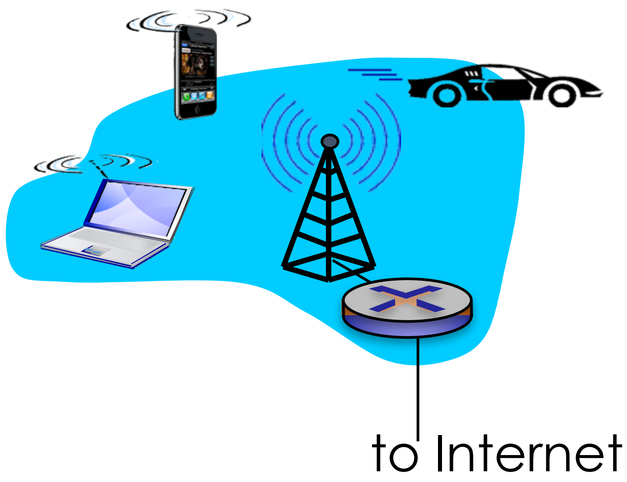

Outline

- What's the Internet?
- What's a protocol?
- Network edge
  - hosts, access net, physical media
- **Network core**
  - packet/circuit switching, Internet structure
- Performance
  - loss, delay, throughput
- Protocol layers, service models
- History

Network Core

- Mesh of interconnected routers
- **Packet-switching**: hosts break application-layer messages into packets
  - Forward packets from one router to the next, across links on path from source to destination
  - Each packet transmitted at full link capacity

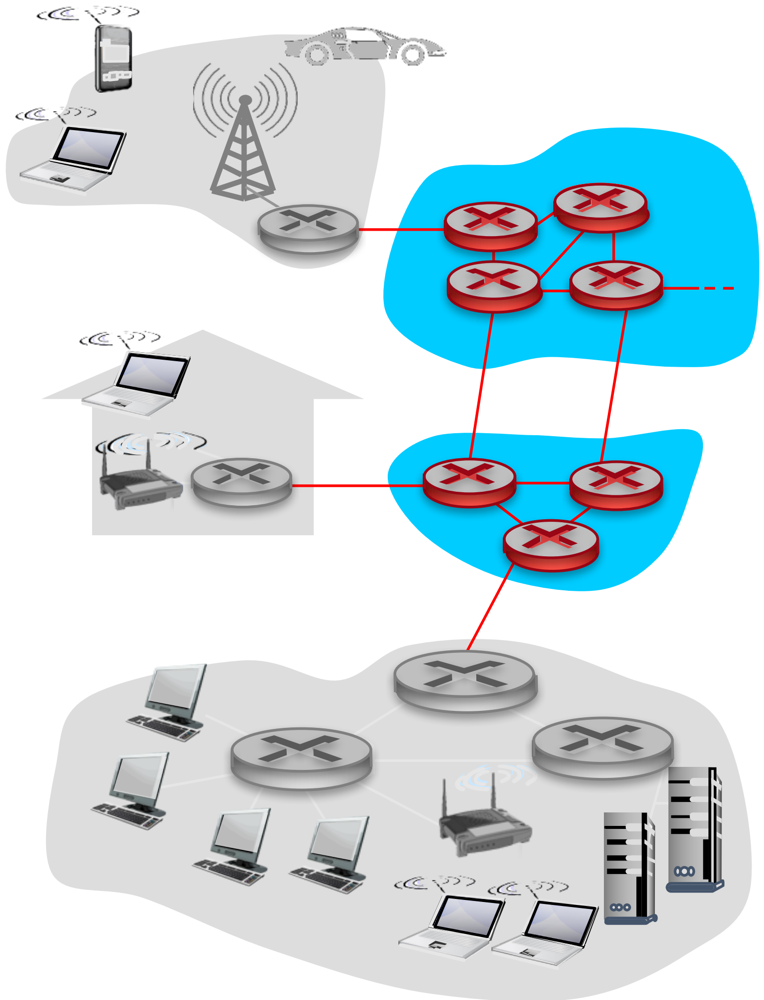

Host: Sends Packets of Data

Host sending function:

- takes application message
- breaks into smaller chunks, known as packets, of length L bits
- transmits packet into access network at transmission rate R
  - link transmission rate, aka link *capacity*, *aka link bandwidth*

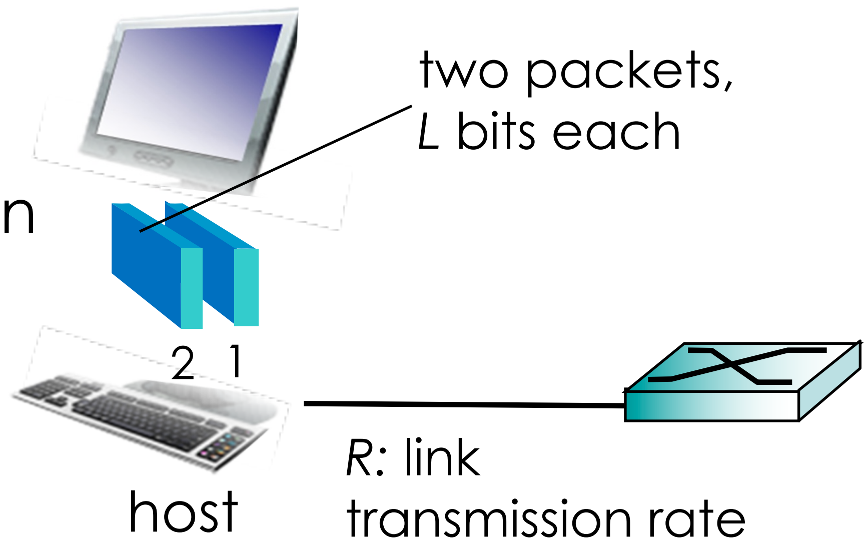

$$\text{packet transmission delay} = \text{time needed to transmit } L\text{-bit packet into link} = \frac{L \text{ (bits)}}{R \text{ (bits/sec)}}$$

Packet Switching

- Take **L/R** seconds to transmit (push out) **L**-bit packet into link at **R** bps
- Store-and-forward transmission
  - Entire packet must arrive at router before it can be transmitted on next link
- N-hop end-end delay = **N*L/R**
  - assuming zero propagation delay

one-hop example:

- L = 7.5 Mbits
- R = 1.5 Mbps
- one-hop transmission delay = 5 sec

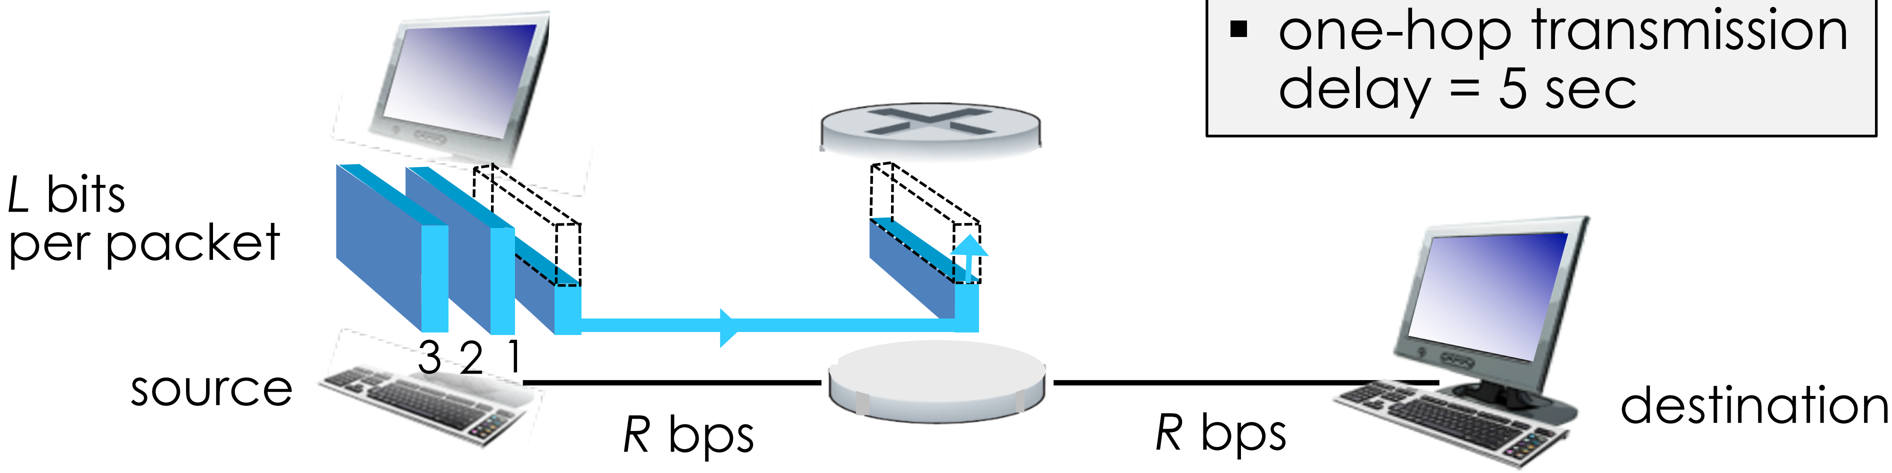

Q: How much time is required to send three packets?

Packet Switching: Queueing Delay, Loss

- **Queueing delay**
  - A packet switch has an **output buffer**
  - Packets buffered in the queue before being forwarded
  - Delay depends on the level of congestion
- **Loss**
  - Given a finite buffer, arriving packets (or some queued packets) must be dropped if the buffer is full

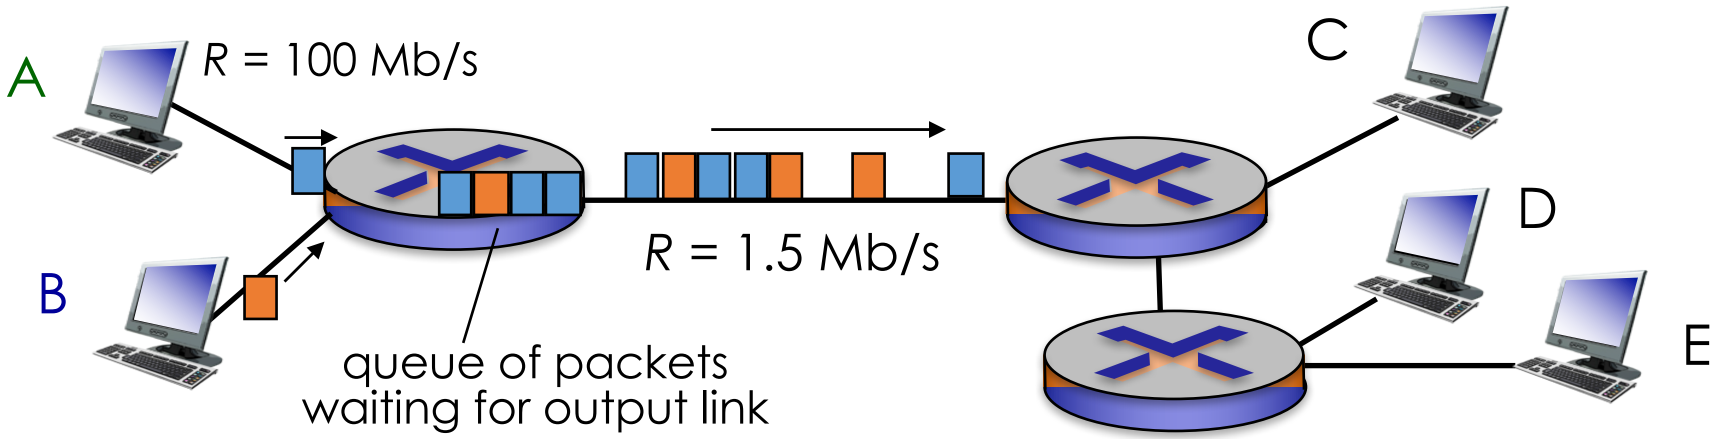

Forwarding and Routing

- **Forwarding**
  - Move packets from router's input to appropriate router output
- **Routing**
  - Determines source-destination route based on **routing algorithms**

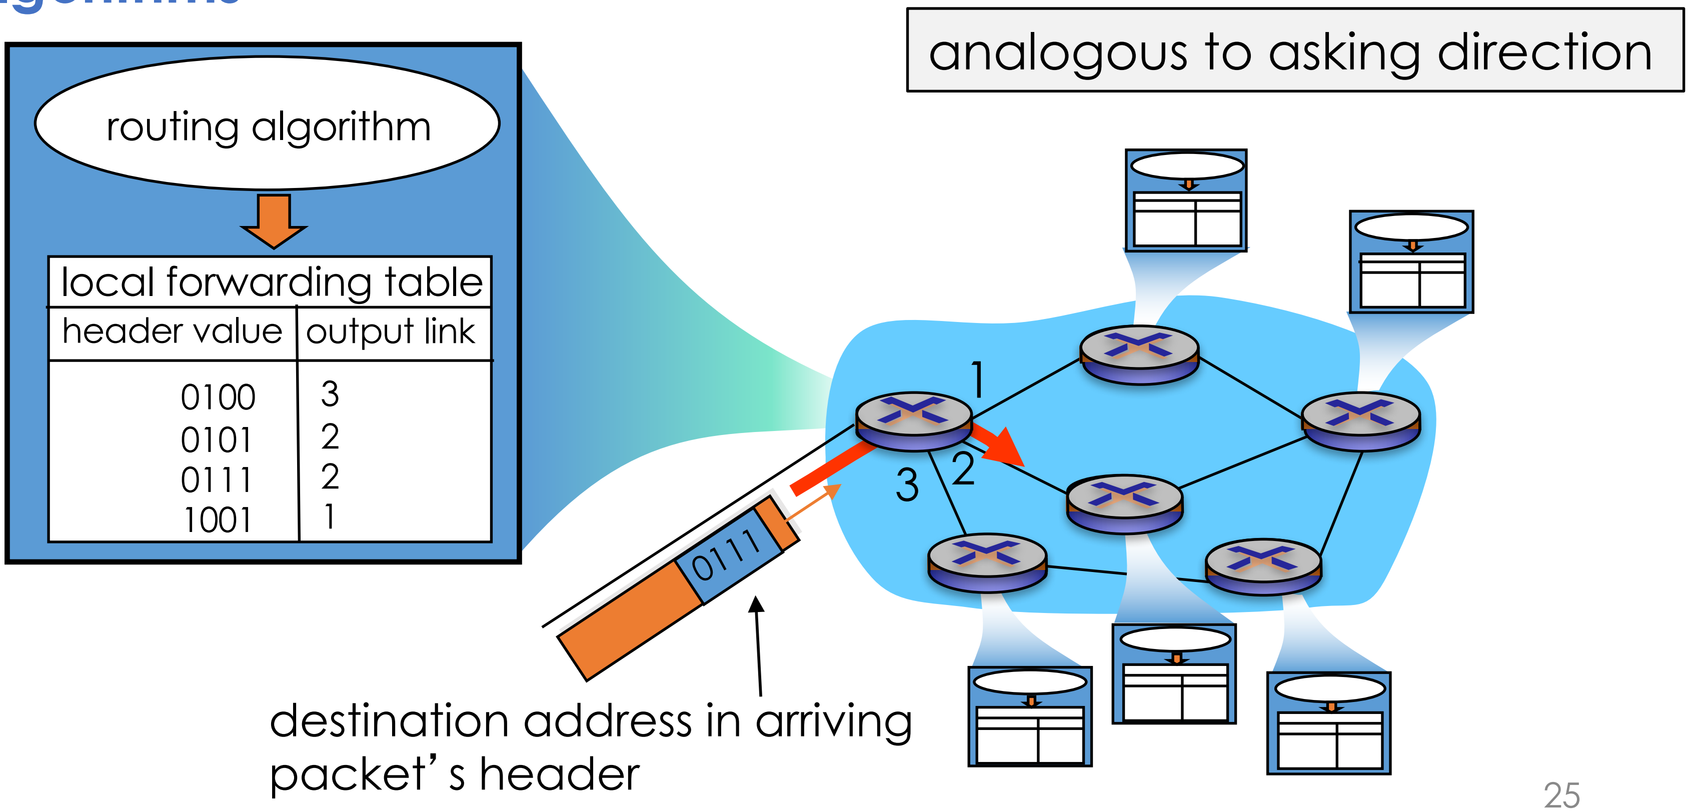

Forwarding and Routing

- **Forwarding**
  - Move packets from router's input to appropriate router output
- **Routing**
  - Determines source-destination route based on **routing algorithms**

Routing protocols

- Used to automatically set the forwarding table
- Possible algorithms
  - Shortest path
  - Fastest path
  - Load balancing

Two Switching Models

- Packet switching
  - Store-ant-forward
  - Link resources are not reserved for any source-destination pairs

- Circuit switching
  - Resources needed along a path are reserved for a duration

Circuit Switching

- End-end resources allocated to, **reserved** for "call" between source & destination
- In diagram, each link has four circuits.
- Dedicated resources
  - **no sharing**
  - **guaranteed performance**
- Commonly used in traditional telephone networks
- Circuit segment idle if not used by call

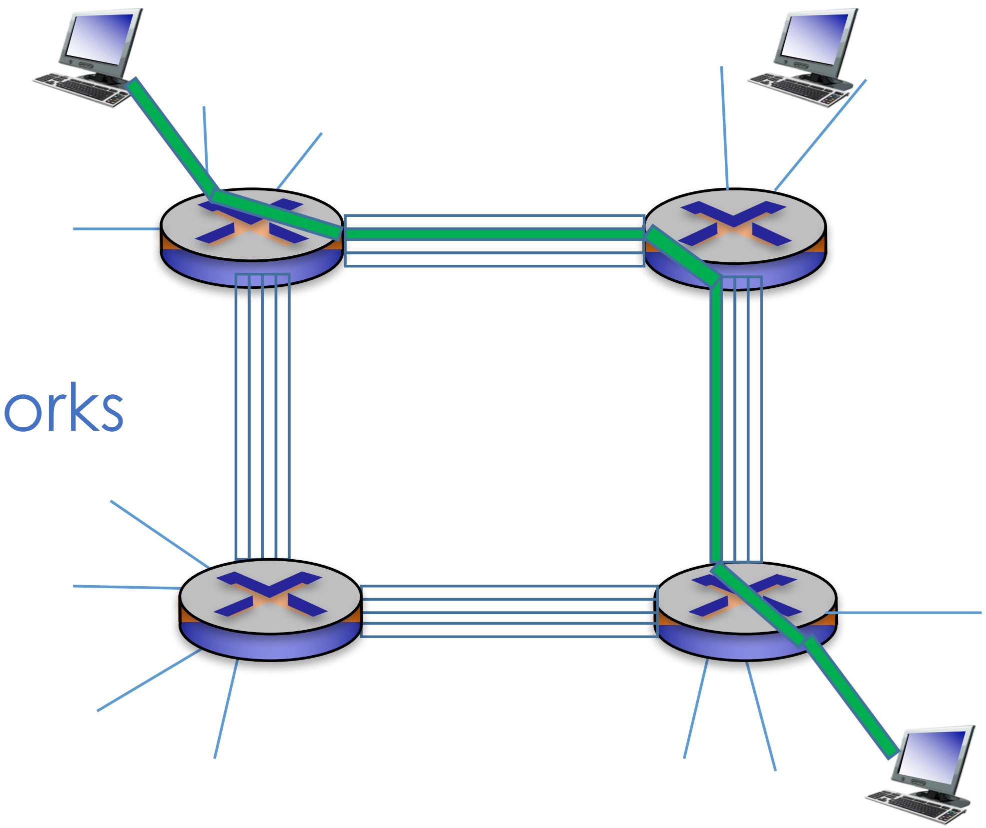

Circuit Switching: FDM vs. TDM

- **Multiplexing**: allocate resources to multiple users

FDM
(Frequency division multiplexing)

Example:
4 users

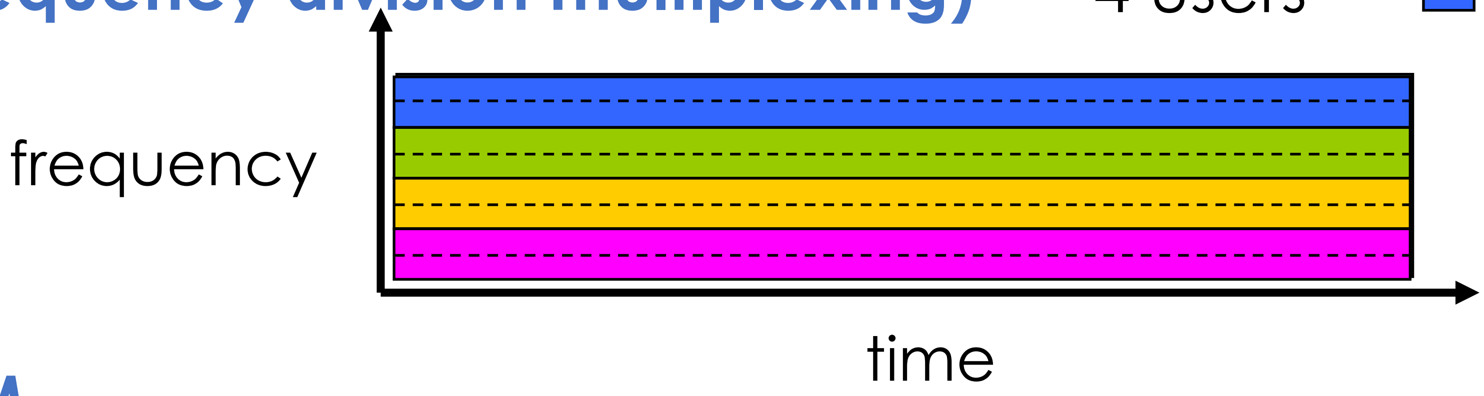

TDM
(Time division multiplexing)

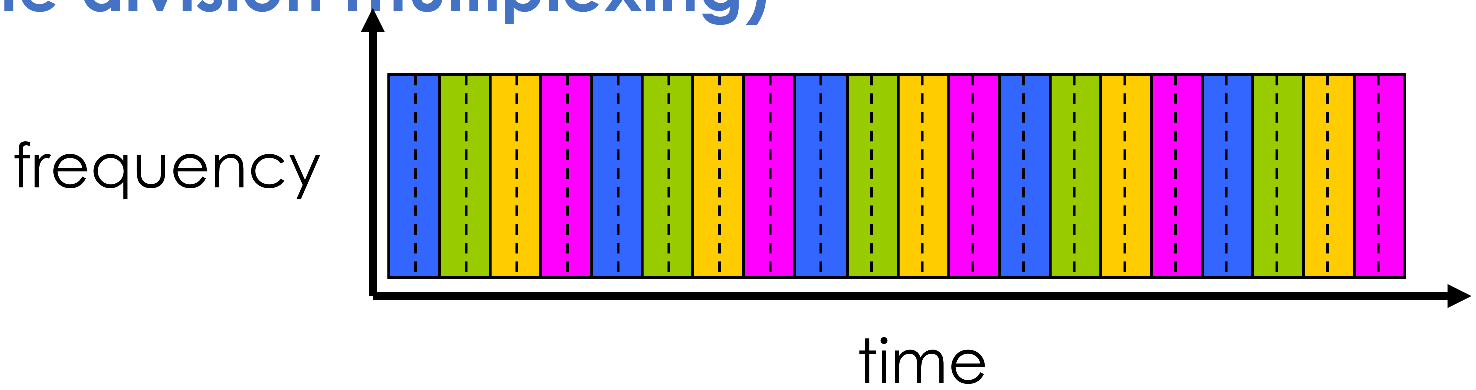

Two Switching Models

- Packet switching
  - Store-ant-forward
  - Link resources are not reserved for any source-destination pairs

☺ better sharing
☺ simpler, more efficient, support more users
☹ loss and delay, might suffer from congestion
☹ less suitable for real-time applications

- Circuit switching
  - Resources needed along a path are reserved for a duration

Example of Sharing

packet switching allows more users to use network!

example:

- 1 Mb/s link
- each user:
  - 100 kb/s when "active"
  - active 10% of time

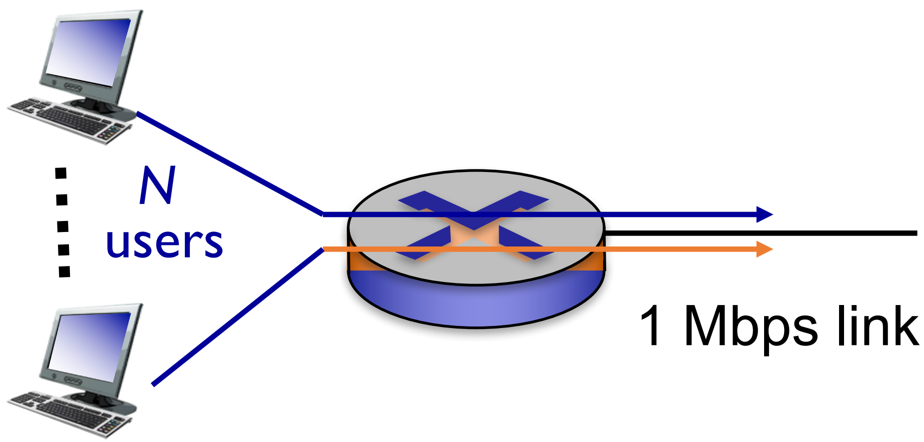

- circuit-switching:
  - 10 users
- packet switching:
  - with 35 users, probability > 10 active at same time is less than .0004

Network of Networks

- End systems connect to Internet via access ISPs (Internet Service Providers)
  - residential, company and university ISPs
- Access ISPs in turn must be interconnected
  - so that any two hosts can send packets to each other
- Resulting network of networks is very complex
  - evolution was driven by economics and national policies (rather than performance consideration)

Internet Structure: Network of Networks

Question: given *millions* of access ISPs, how to connect them together?

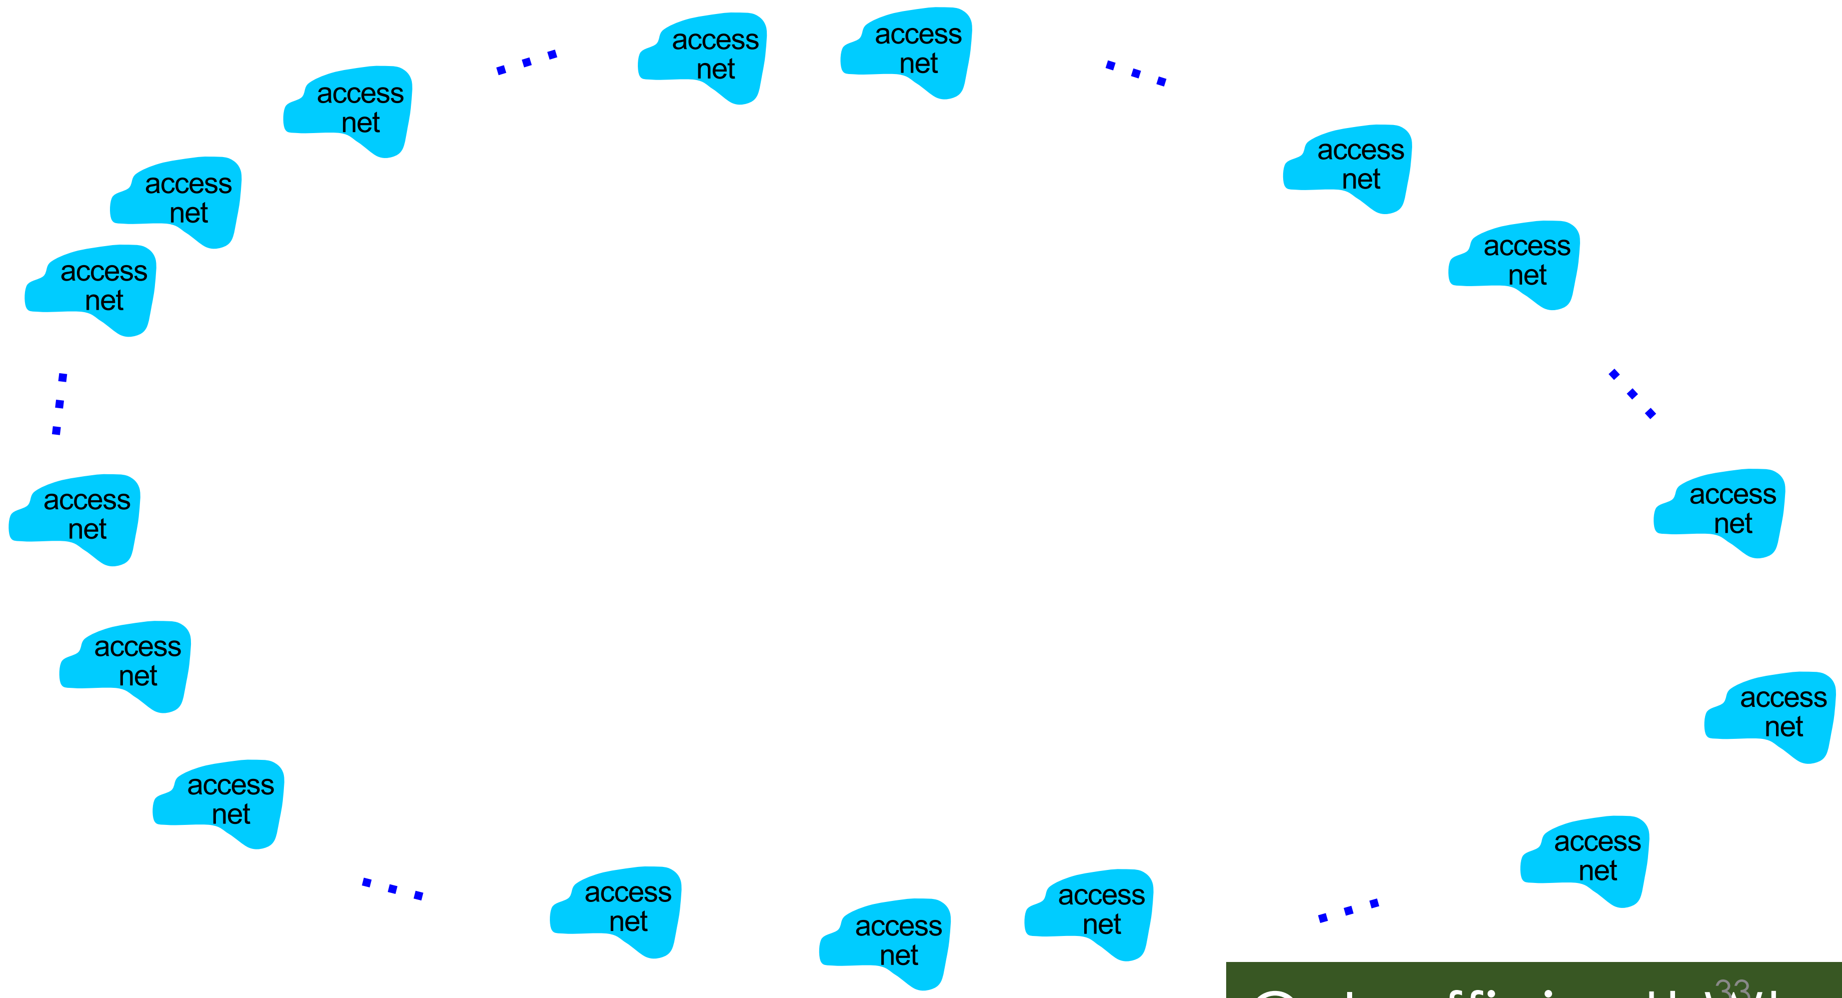

Q: Inefficient! Why?

Internet Structure: Network of Networks

Option: connect each access ISP to every other access ISP?

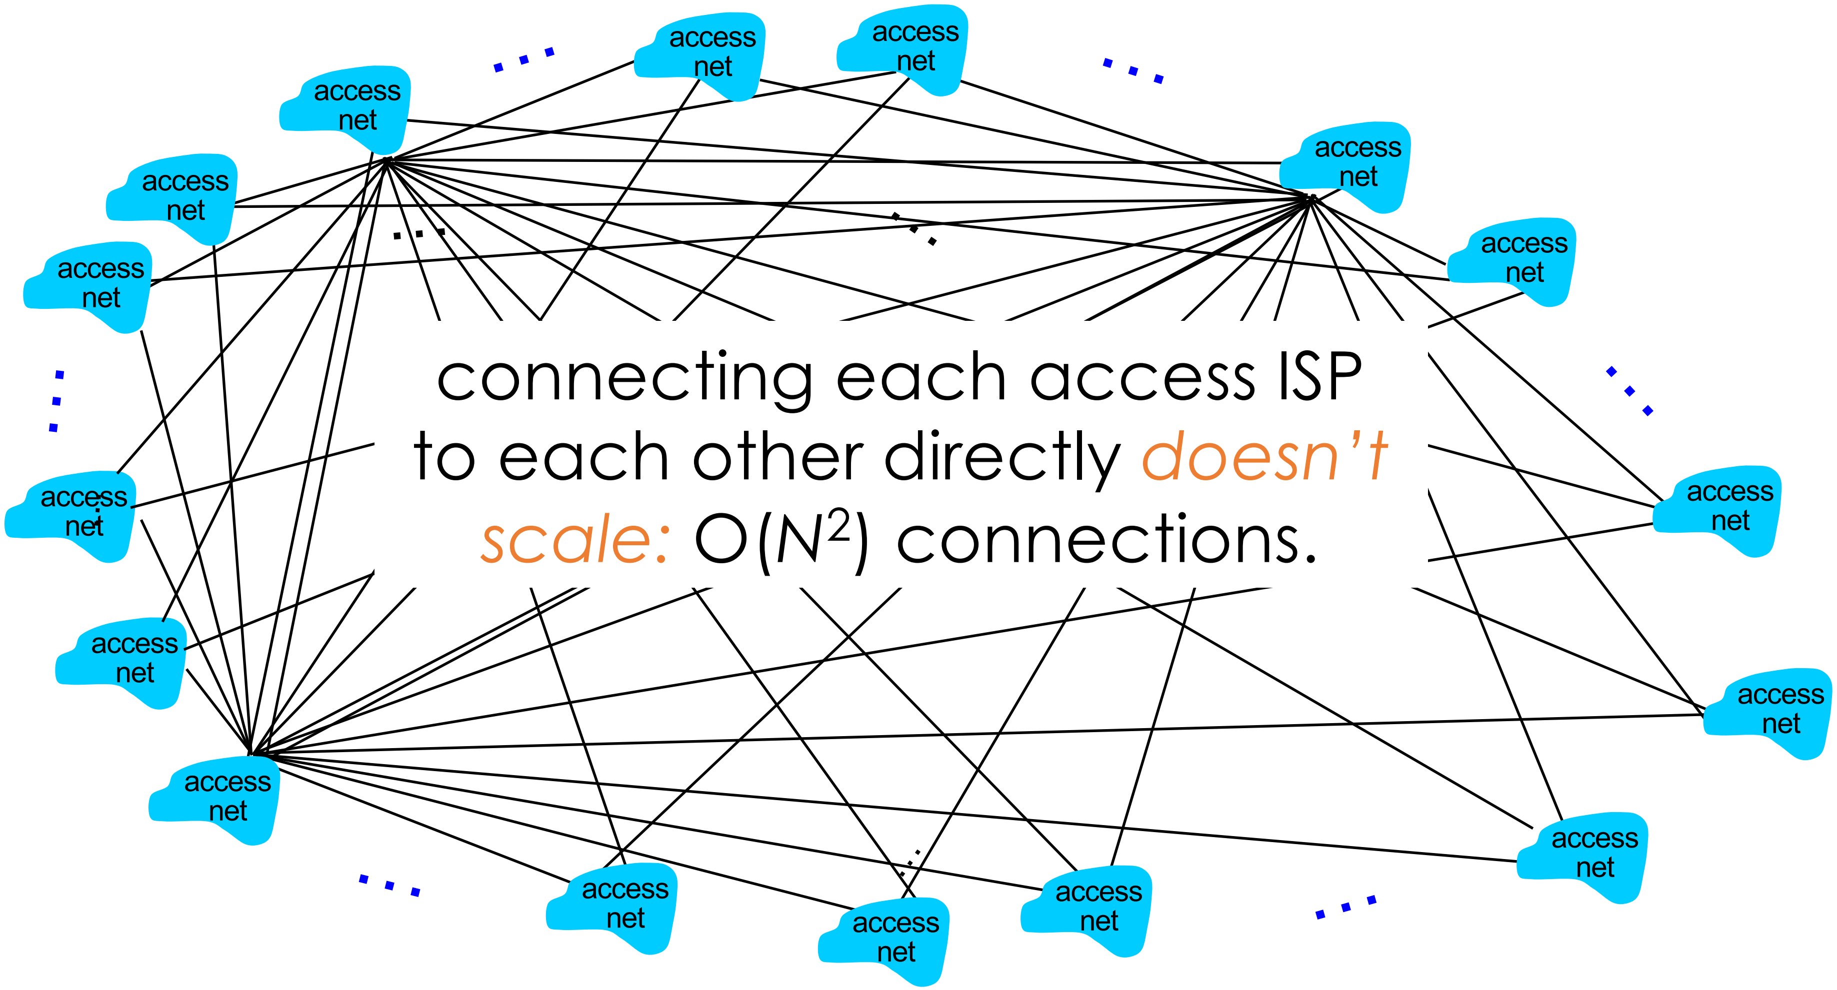

connecting each access ISP to each other directly *doesn't scale:* $O(N^2)$ connections.

Internet Structure: Network of Networks

Option: connect each access ISP to one global transit ISP?
Customer and provider ISPs have economic agreement

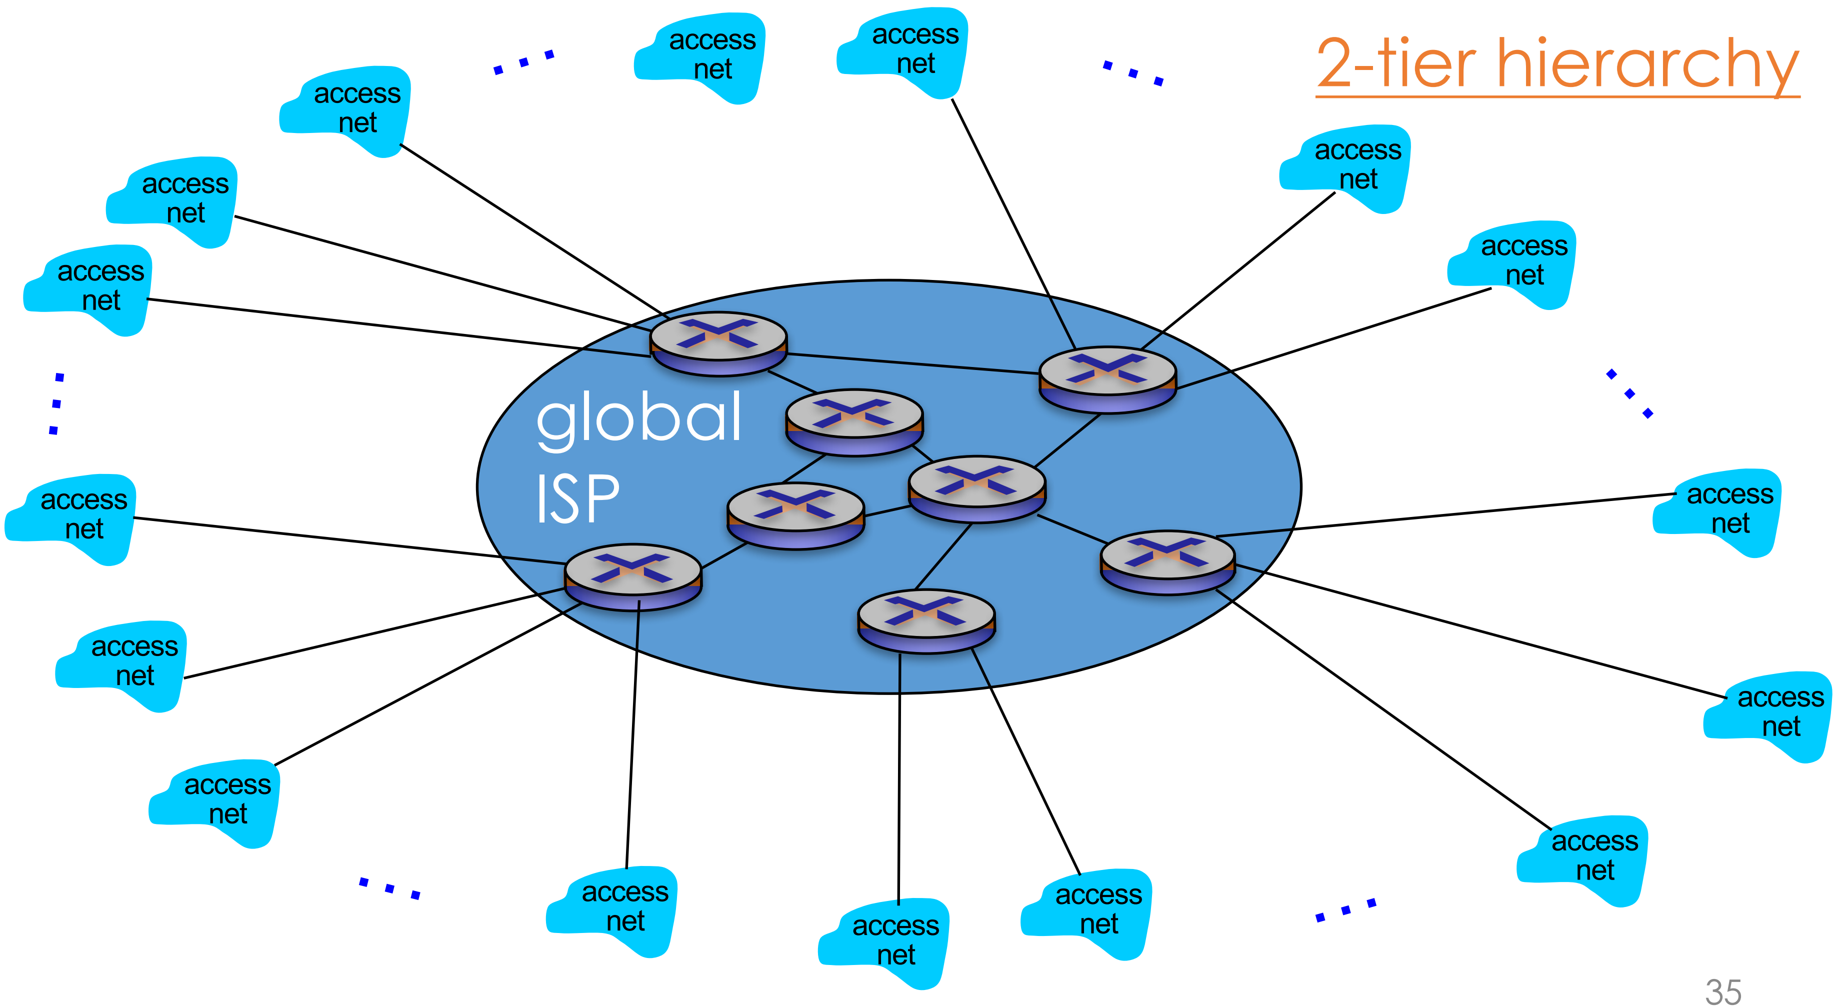

Internet Structure: Network of Networks

But if one global ISP is viable business, there will be competitors

Friend or competitor?

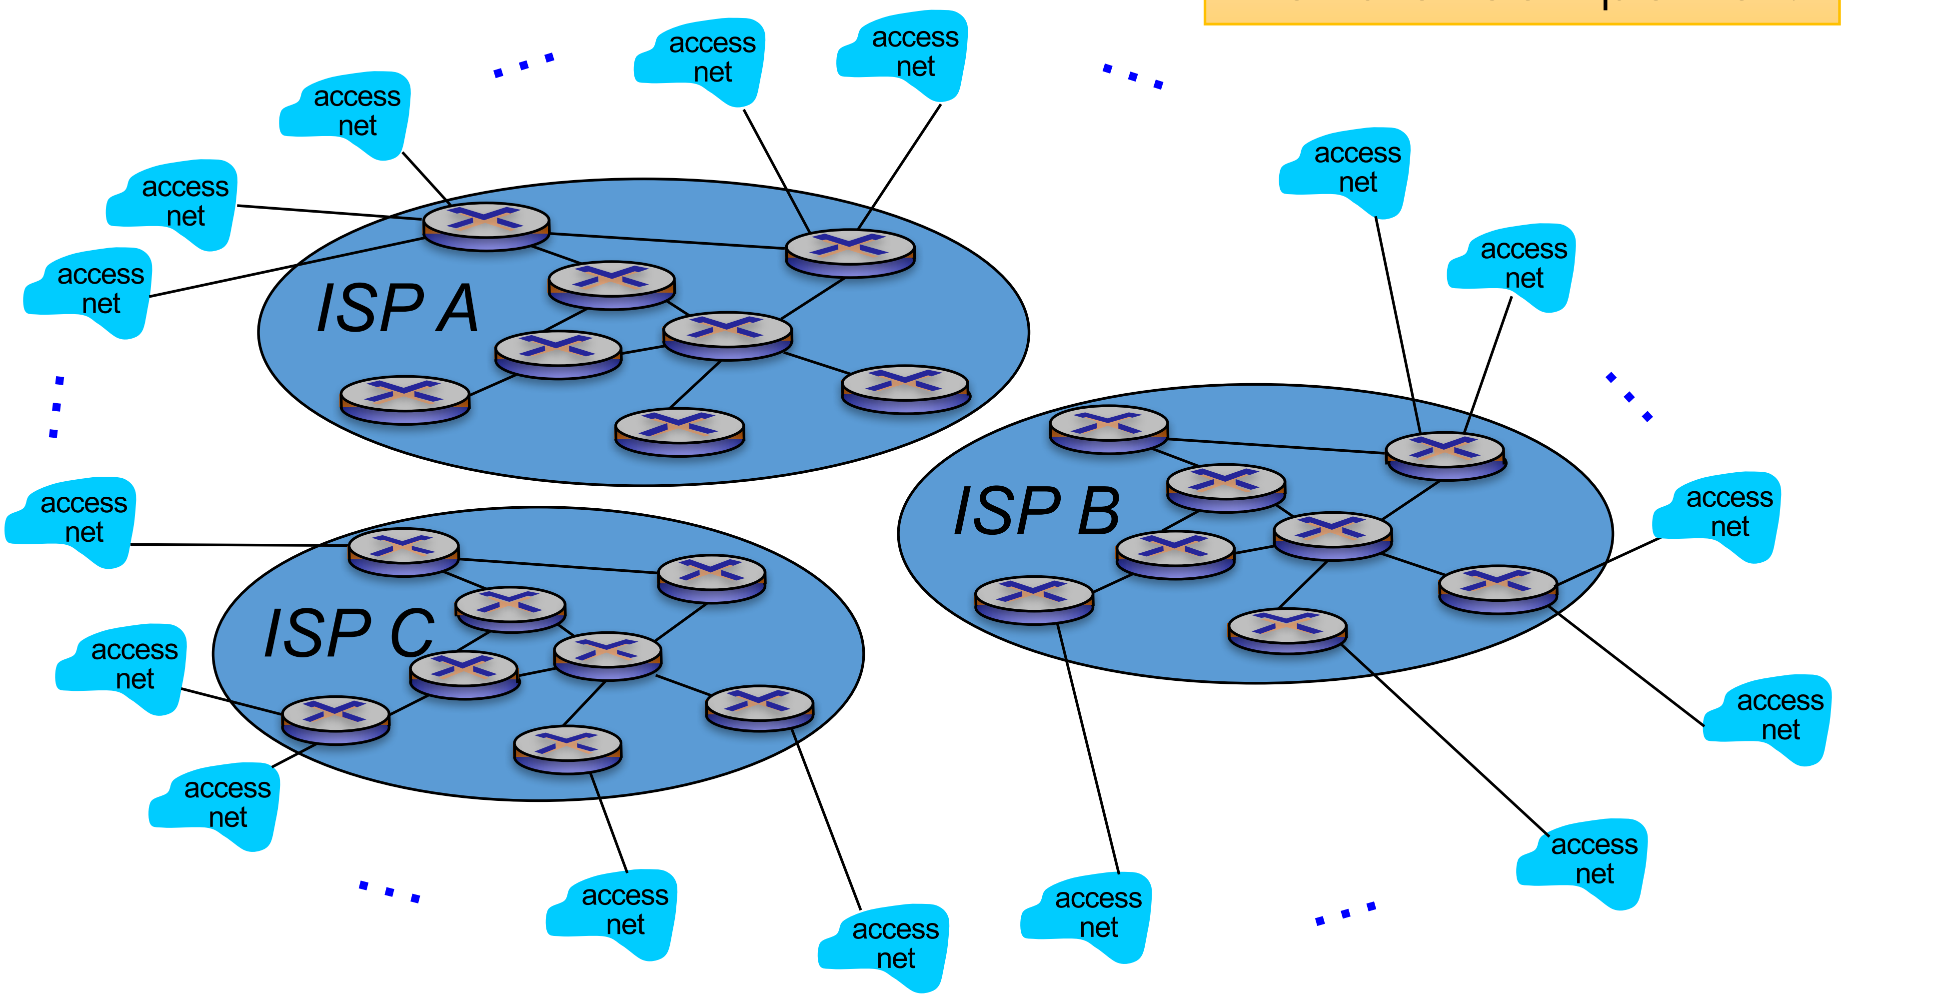

Internet Structure: Network of Networks

But if one global ISP is viable business, there will be competitors
→ Which must be interconnected

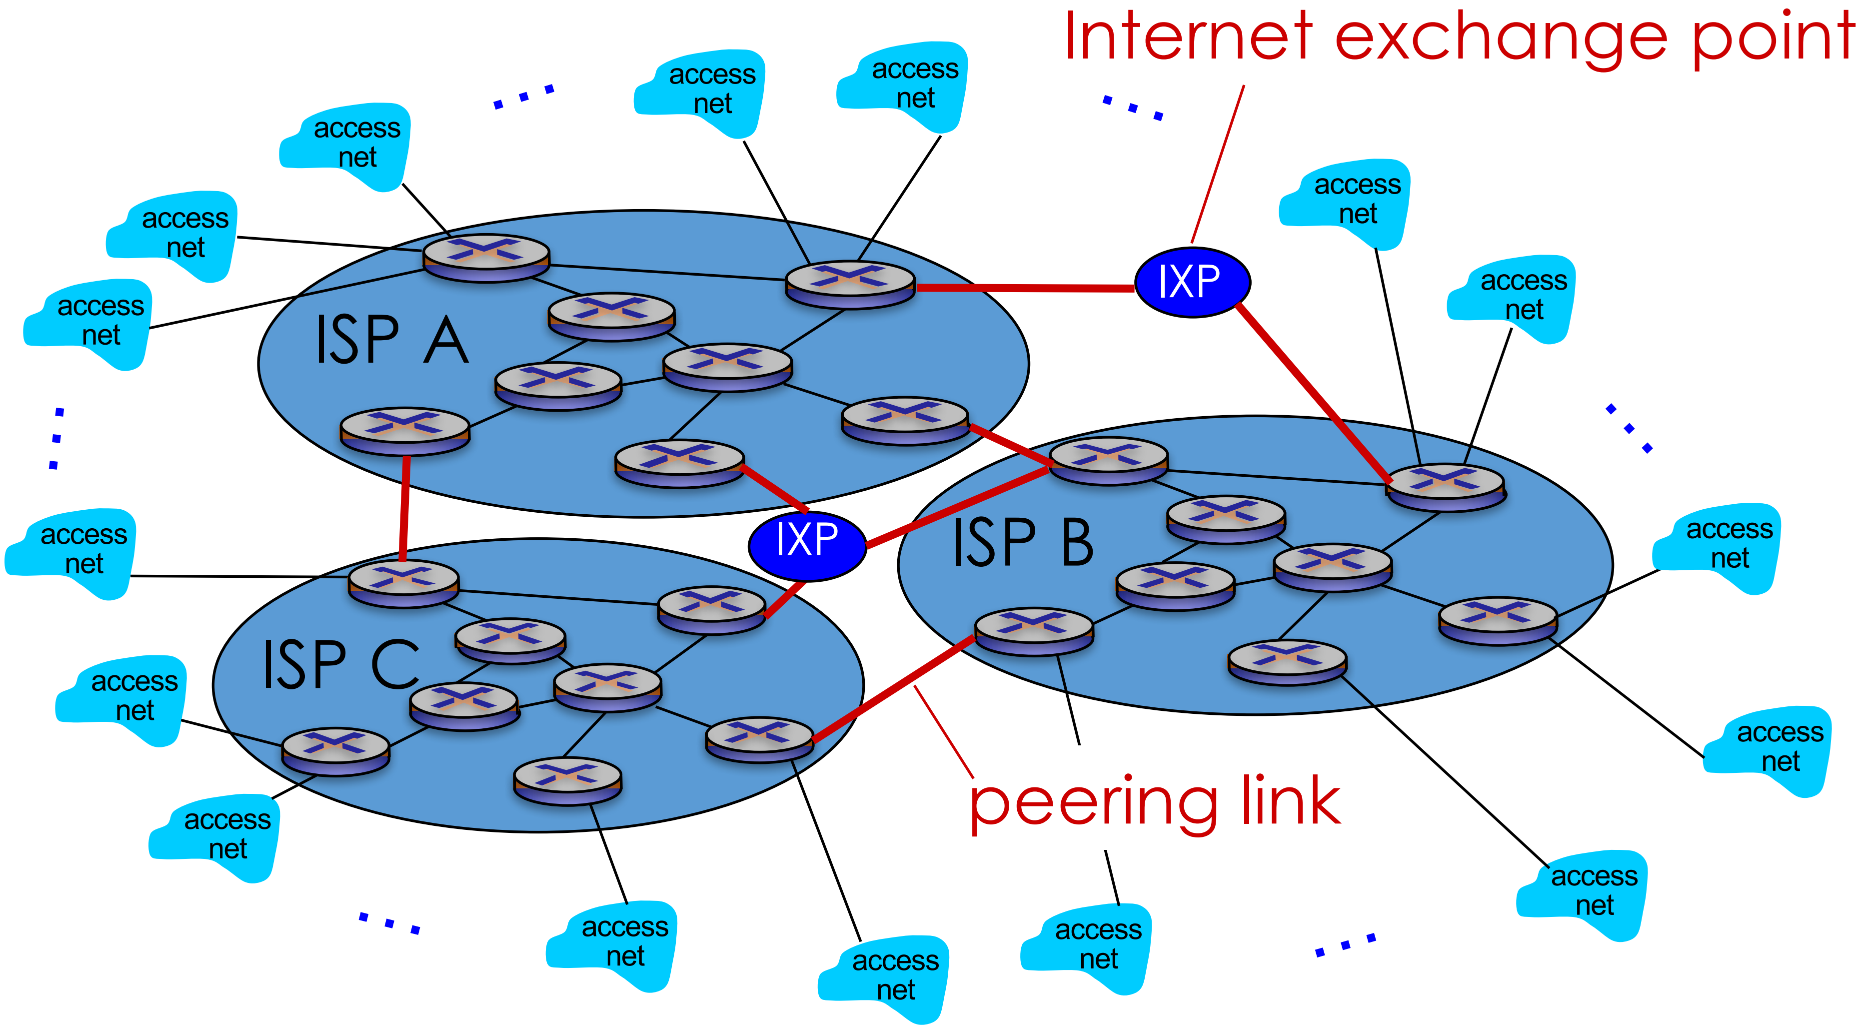

Internet Structure: Network of Networks

Regional networks may arise to connect access nets to ISPs

→ Each access ISP pays the connected regional ISPs

→ Each regional ISP pays tier-1 ISPs

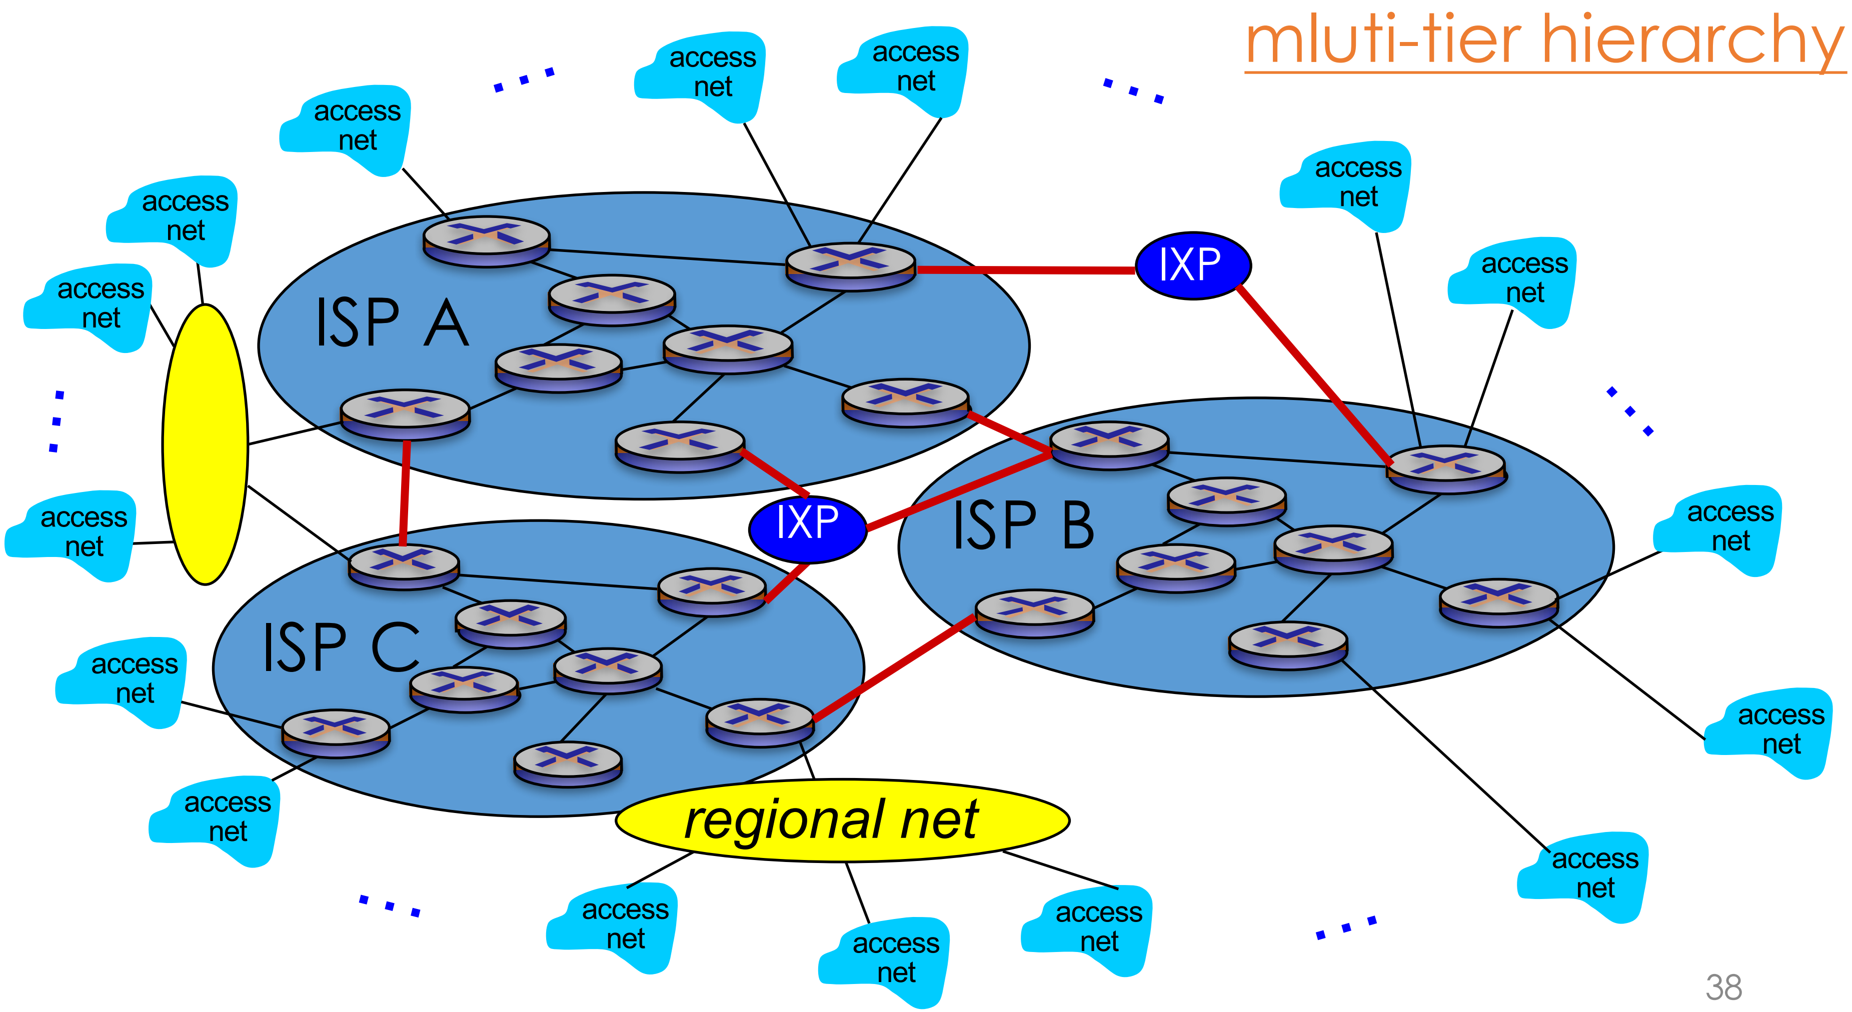

Internet Structure: Network of Networks

Multi-home:
An ISP may connect to several provider ISPs to ensure reliability

mluti-tier hierarchy

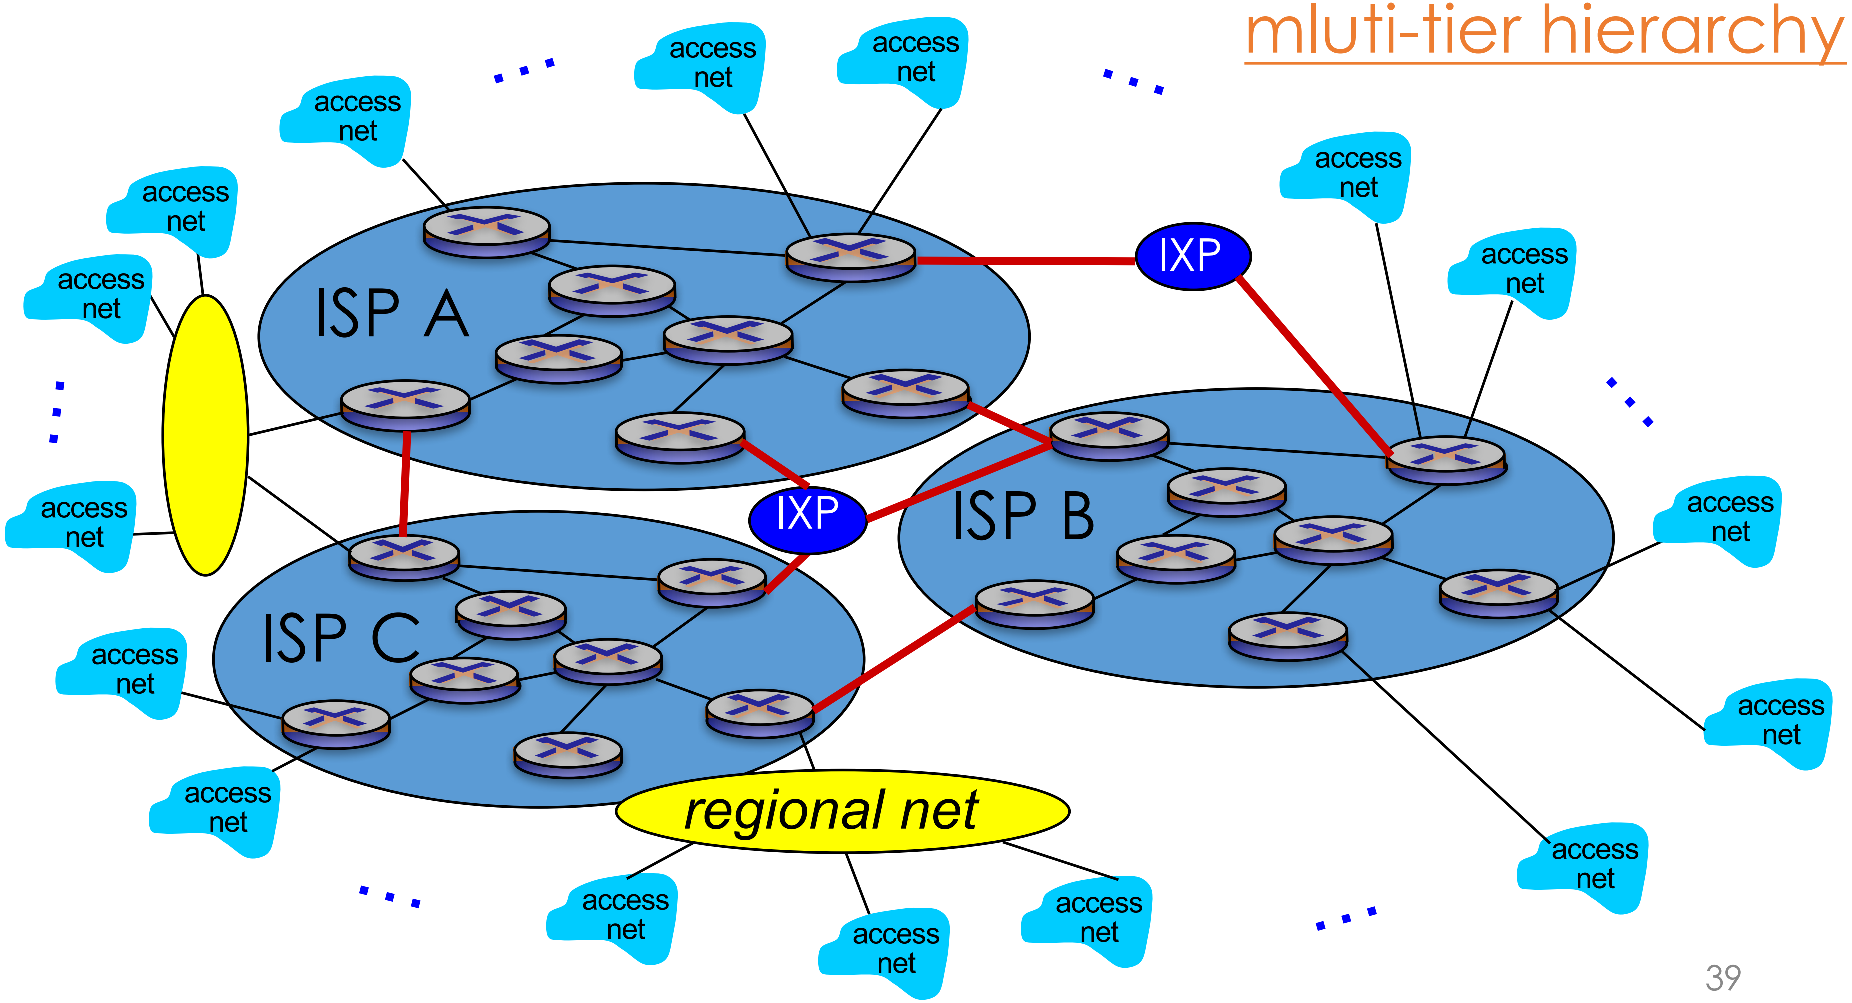

Internet Structure: Network of Networks

Option: content provider networks (e.g., Google, Microsoft, Akamai) may run their own network, to bring services, content close to end users

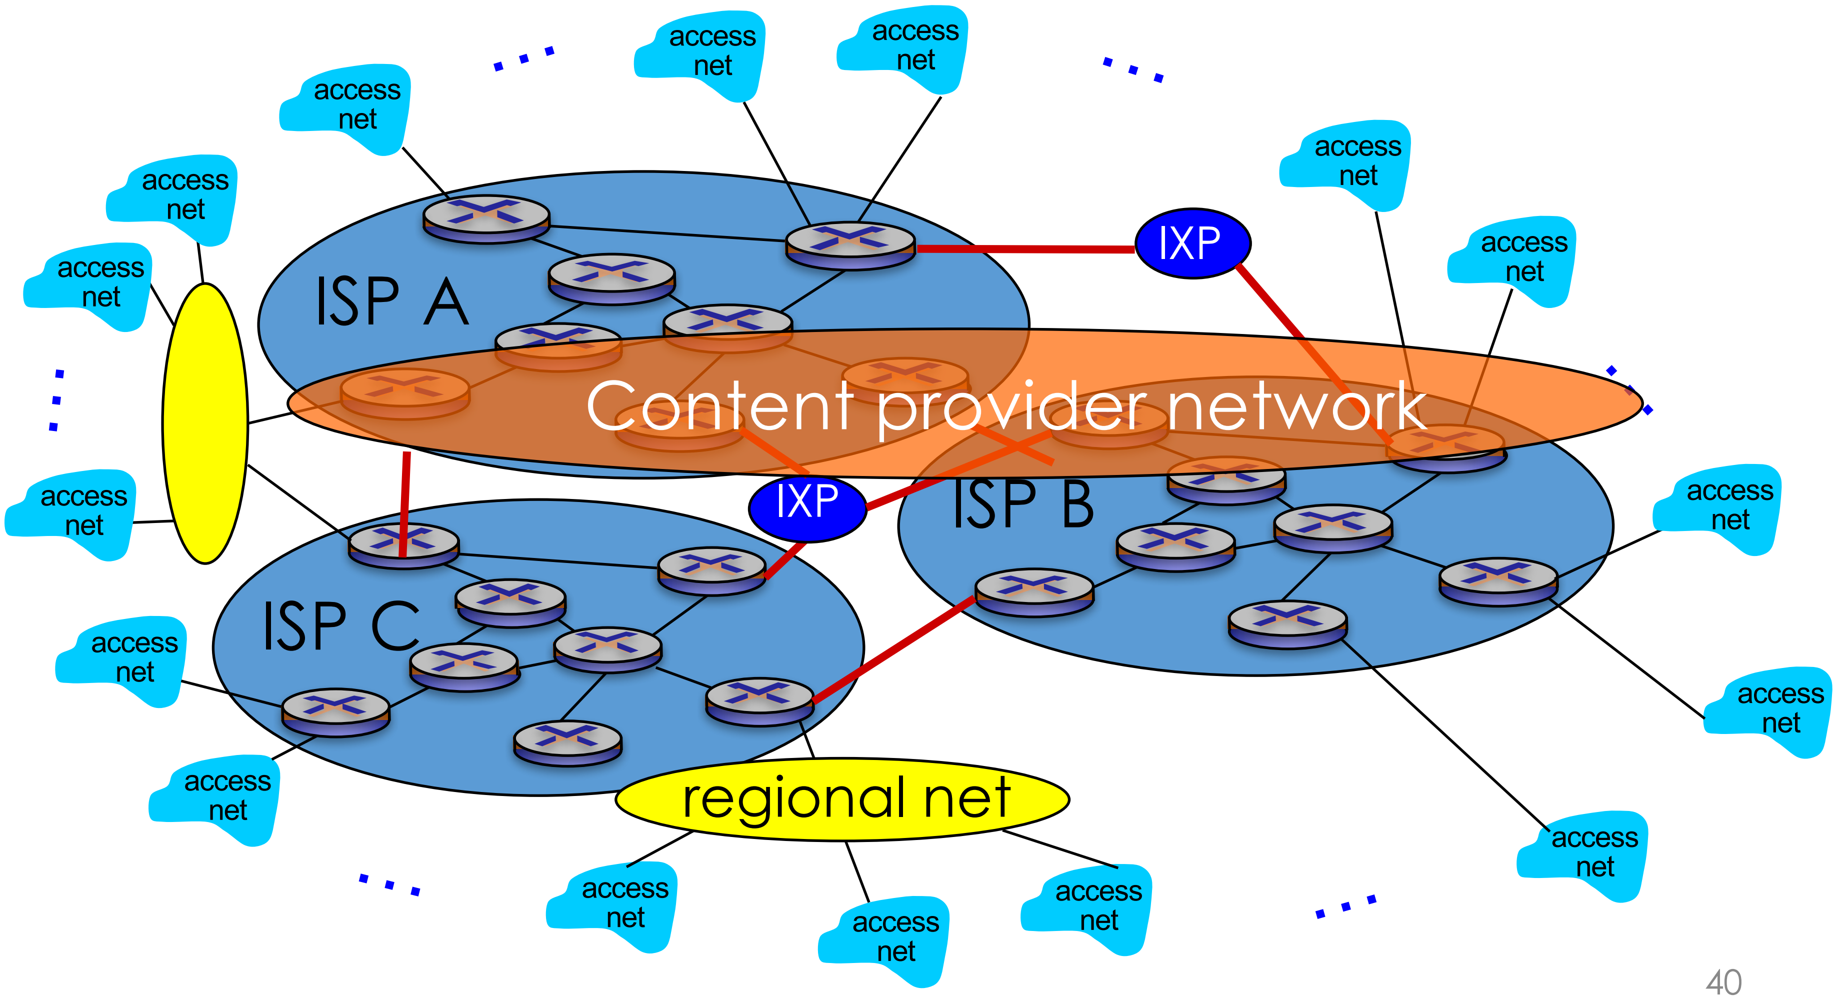

Internet Structure: Network of Networks

- At center: small number of well-connected large networks
  - “tier-1” commercial ISPs (e.g., Level 3, Sprint, AT&T, NTT), national & international coverage
  - content provider network (e.g., Google): private network that connects it data centers to Internet, often bypassing tier-1, regional ISPs

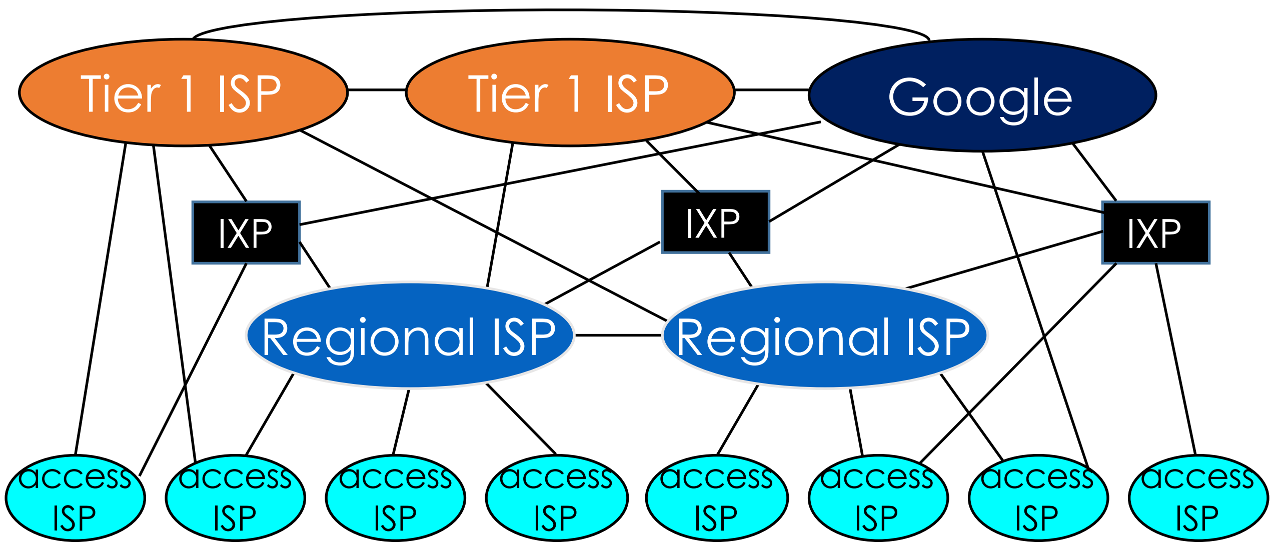

Tier-1 ISP: e.g., Sprint

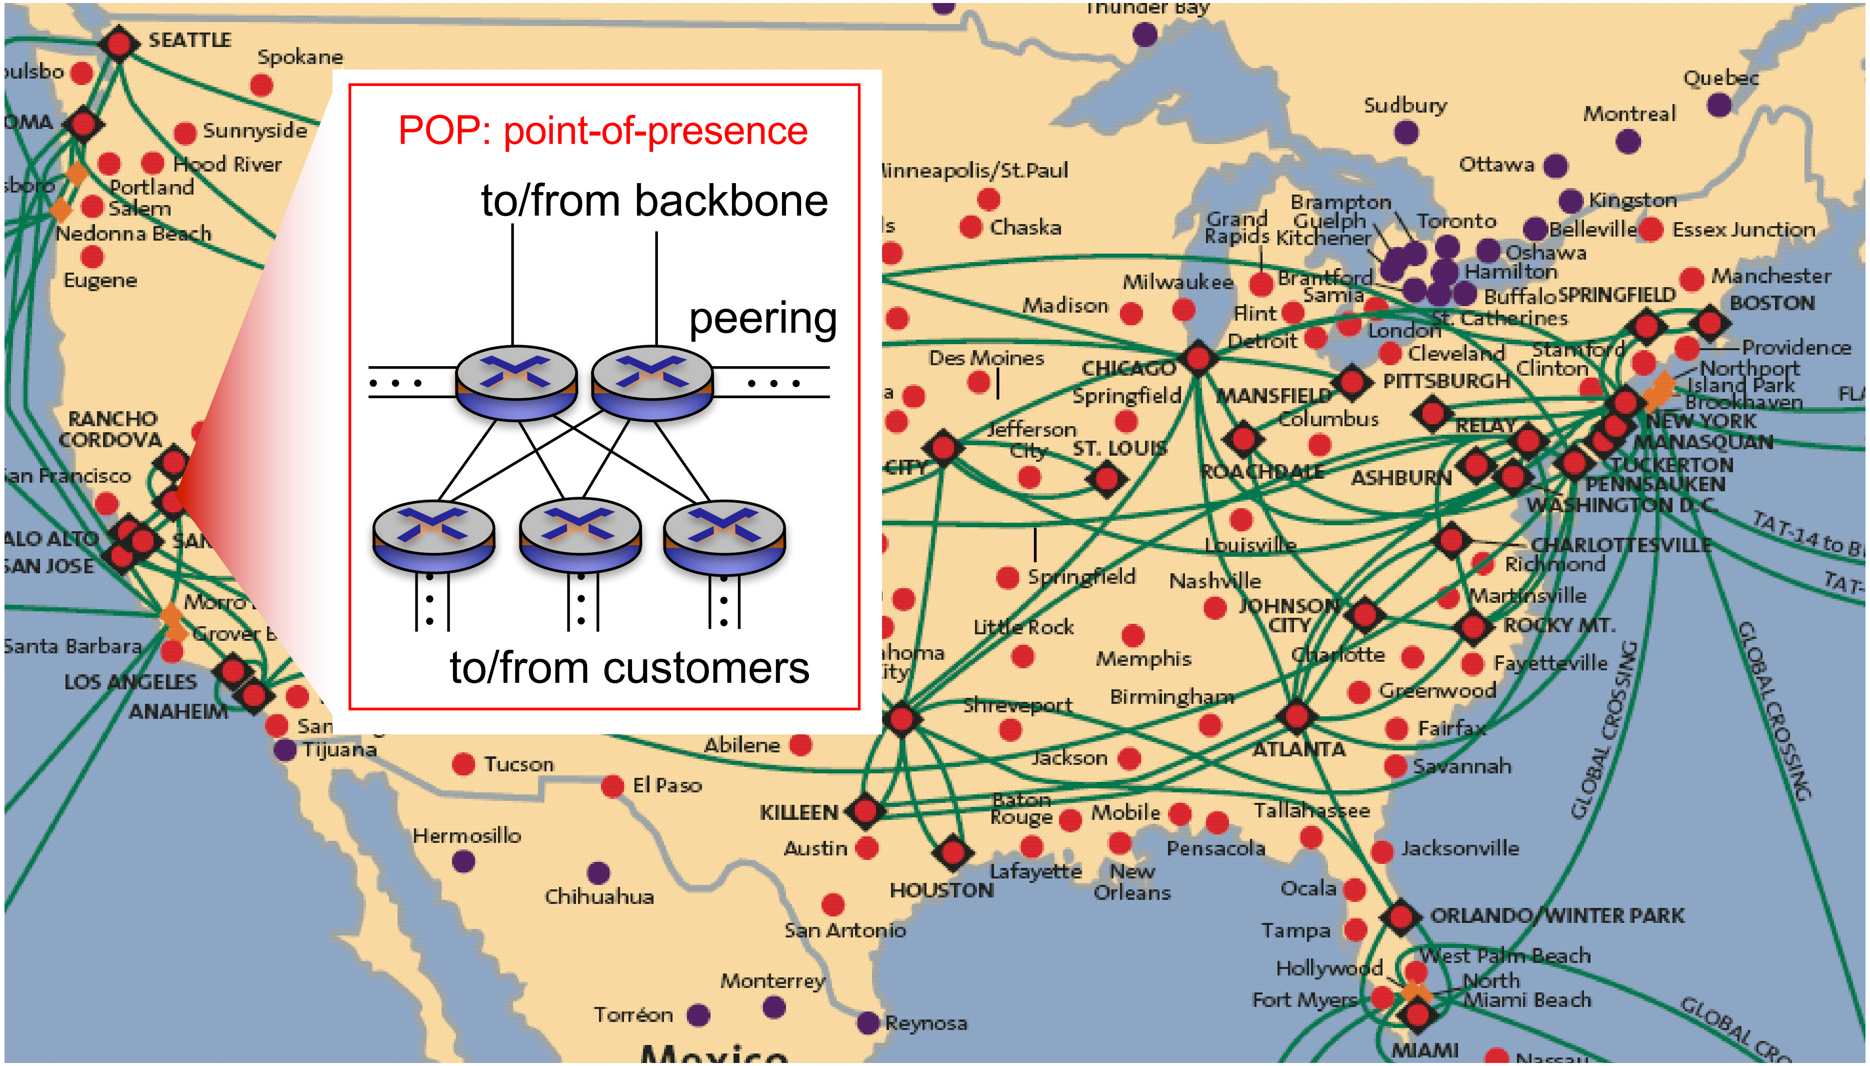

Outline

- What's the Internet?
- What's a protocol?
- Network edge
  - hosts, access net, physical media
- Network core
  - packet/circuit switching, Internet structure
- **Performance**
  - loss, delay, throughput
- Protocol layers, service models
- History

How do loss and delay occur?

packets *queue* in router buffers

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

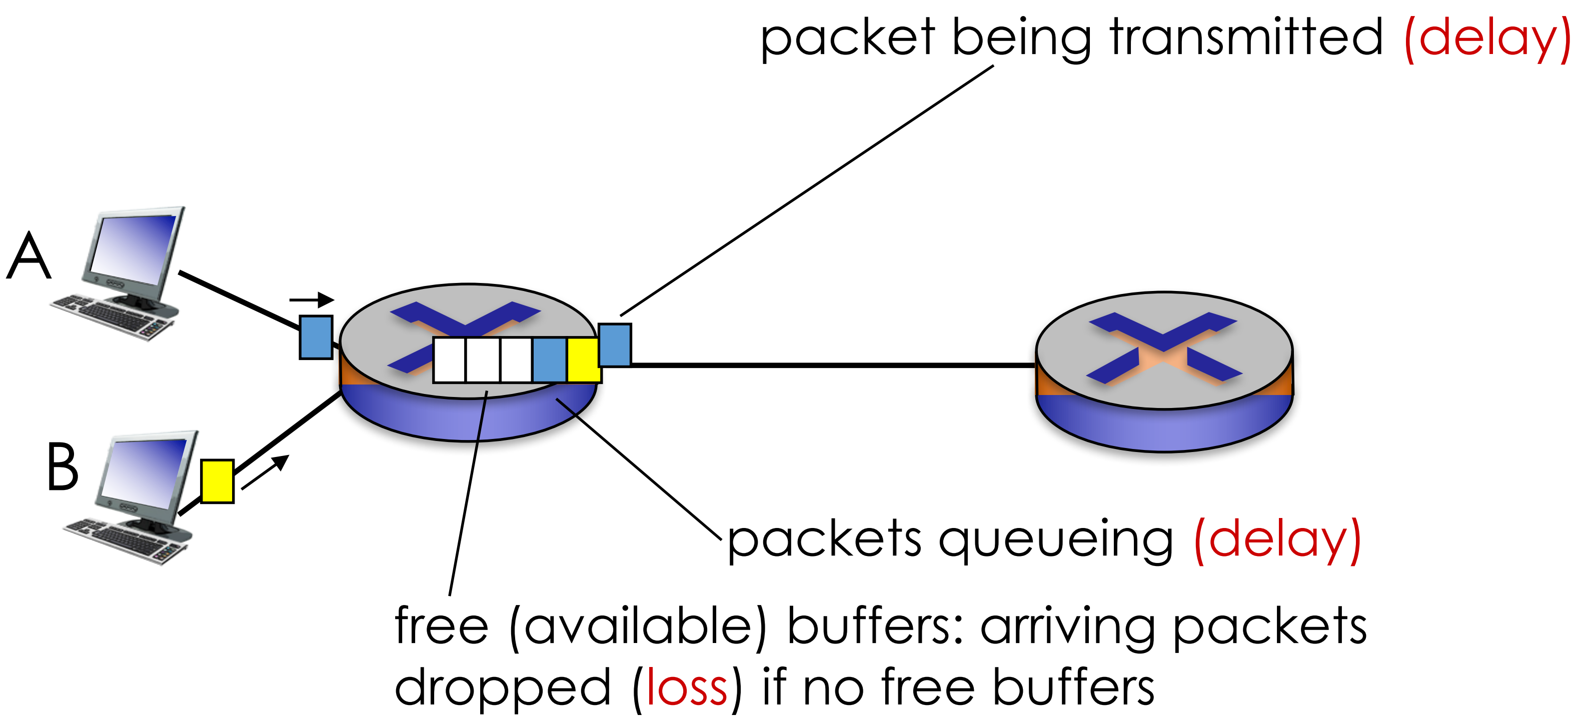

Four Sources of Packet Delay

1. Nodal processing delay
   - Time required to examine the packet's header and determine where to go
2. Queueing delay
   - Wait in the buffer for being transmitted onto the link
3. Transmission delay
   - Time required to push all the packet's bits into the link
4. Propagation delay
   - Time required to propagate from the beginning of the link to another end point

Four Sources of Packet Delay

- d_{proc}: nodal processing
  - check bit errors
  - determine output link
  - typically < msec
- d_{queue}: queueing delay
  - time waiting at output link for transmission
  - depends on congestion level of router

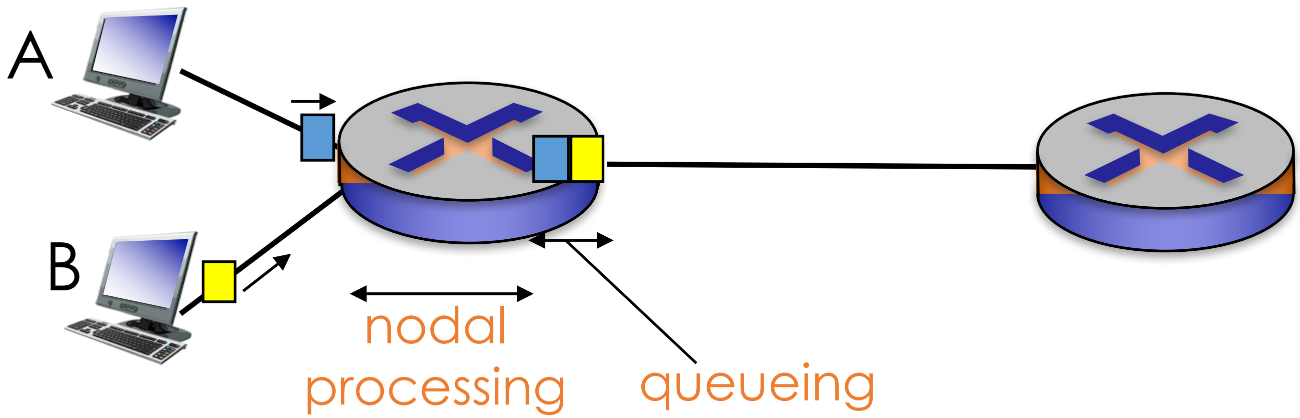

Four Sources of Packet Delay

- d_{trans}: transmission delay
  - L: packet length (bits)
  - R: link bandwidth (bps)
  - d_{trans} = L/R

d_{trans} and d_{prop} very different

- d_{prop}: propagation delay
  - d: length of physical link
  - s: propagation speed (~2x108 m/sec)
  - d_{prop} = d/s

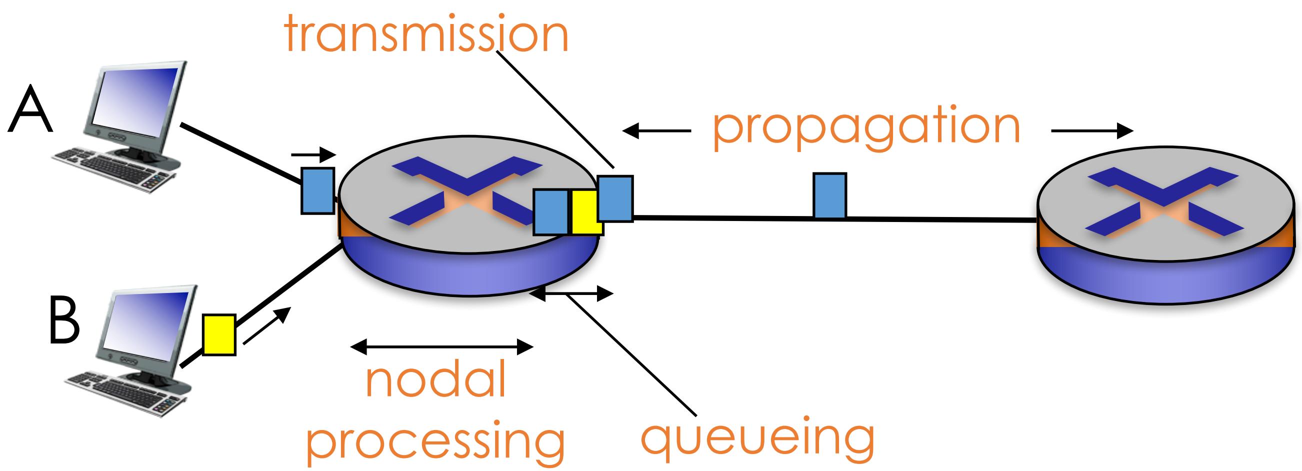

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

Caravan Analogy

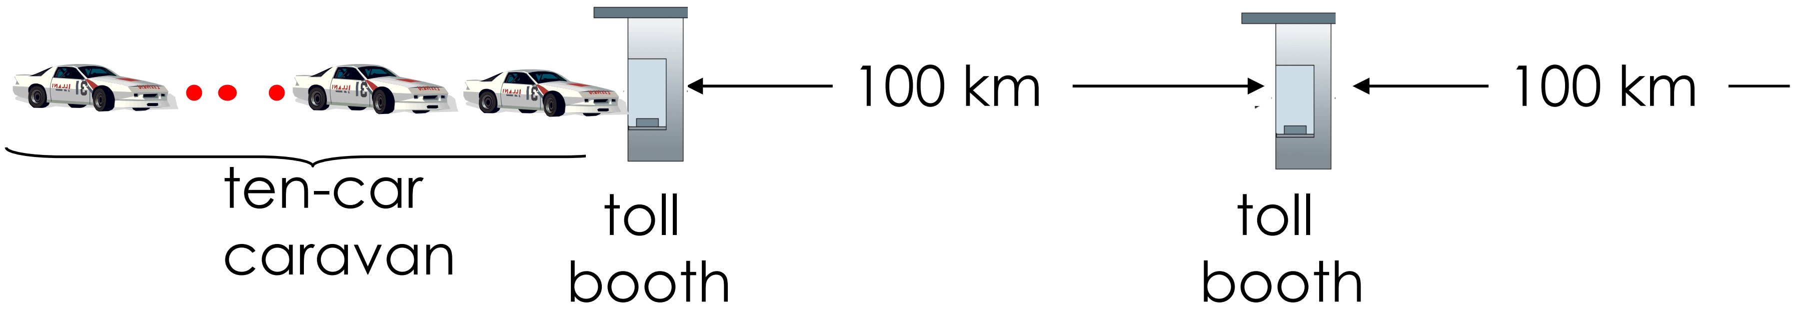

- Cars "propagate" at 100 km/hr
- Toll booth takes 12 sec to service car (bit transmission time)
- car ~ bit; caravan ~ packet
- Q: How long until caravan is lined up before 2nd toll booth?

- 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

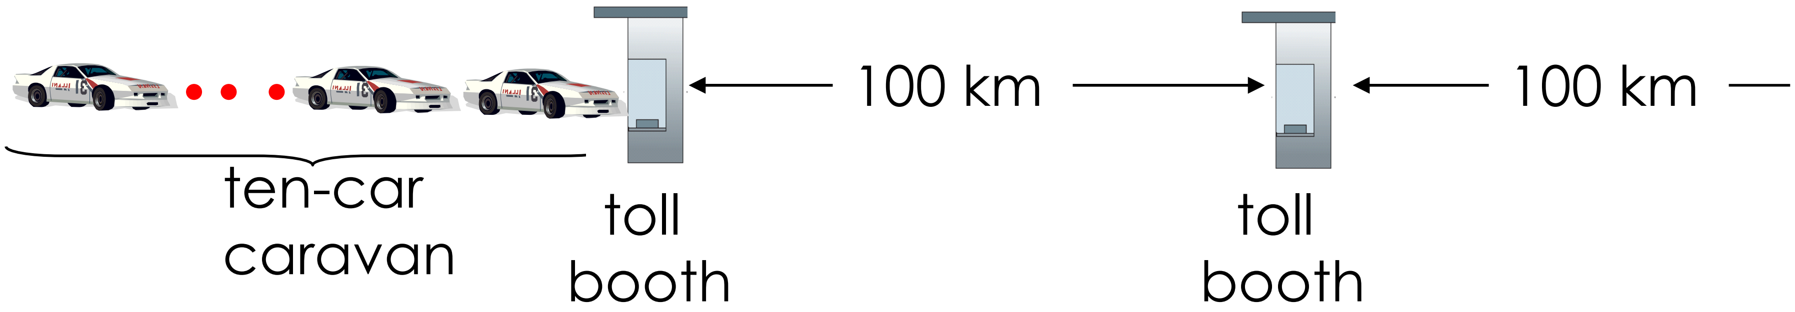

- Suppose cars now "propagate" at 1000 km/hr, and suppose toll booth now takes one min to service a car
- Q: Will cars arrive to 2nd booth before all cars serviced at first booth?

A: Yes! after 7 min, first car arrives at second booth; three cars still at first booth

Queueing Delay (revisited)

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

La/R: traffic intensity

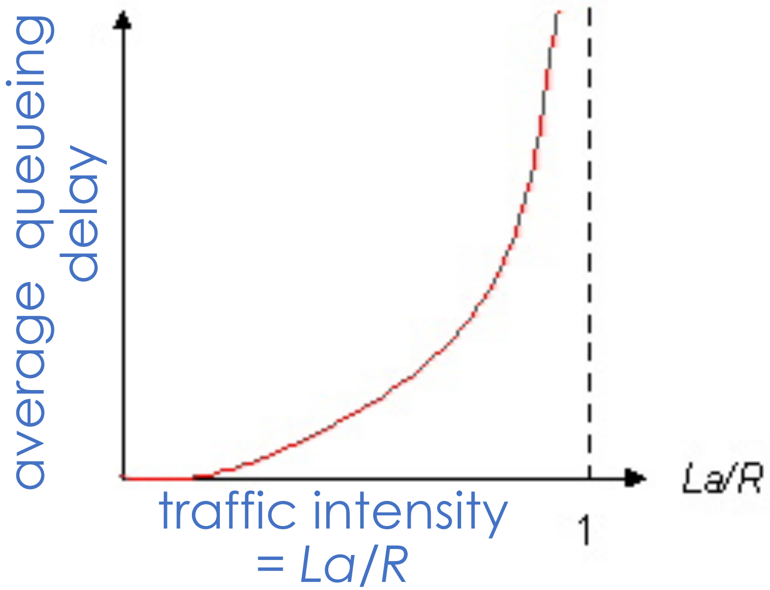

- La/R ~ 0: avg. queueing delay small
- La/R →1: avg. queueing delay large
- La/R > 1: more "work" arriving than can be serviced, average delay **infinite**!

La/R ~ 0

La/R →1

"Real" Internet delays and routes

- What do "real" Internet delay & loss look like?
- `traceroute` program: provides delay measurement from source to router along end-to-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

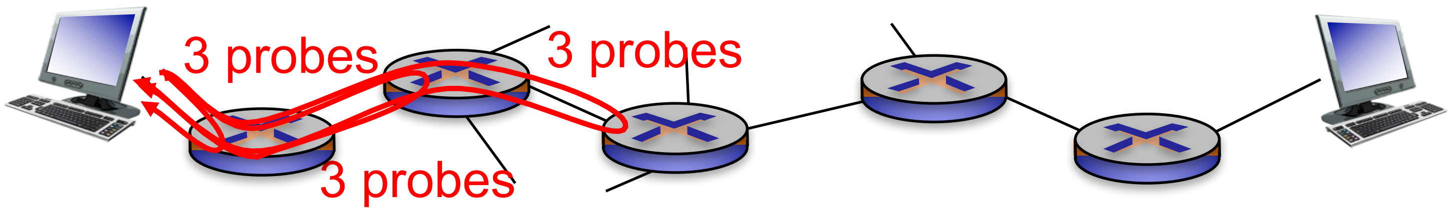

“Real” Internet delays and routes

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

3 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)

Mini-Assignment

- traceroute from linux6.cs.nctu.edu.tw (or your local machine) to www.csail.mit.edu and answer the following questions

1. Copy and paste your results
2. How many hops are there from the sources to the destination?
3. What is the hop with the longest delay?
4. Why sometimes a later router responds faster than earlier routers? (Why sometimes the response latency is decreasing?)

- Save your answers as a pdf file and send to mmcom.nctu@gmail.com

Packet Loss

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

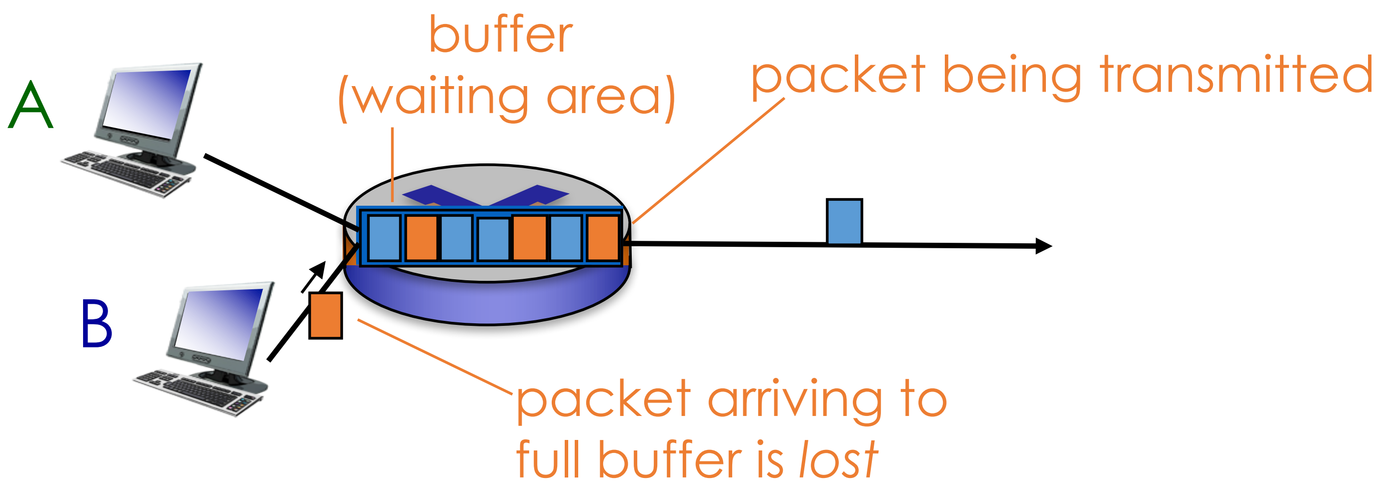

Throughput

- Throughput: rate (bits/time unit) at which bits transferred between sender/receiver
  - instantaneous: rate at a given point in time (how many bits sent in one second)
  - average: rate over longer period of time (how many time required to send a batch of bits)

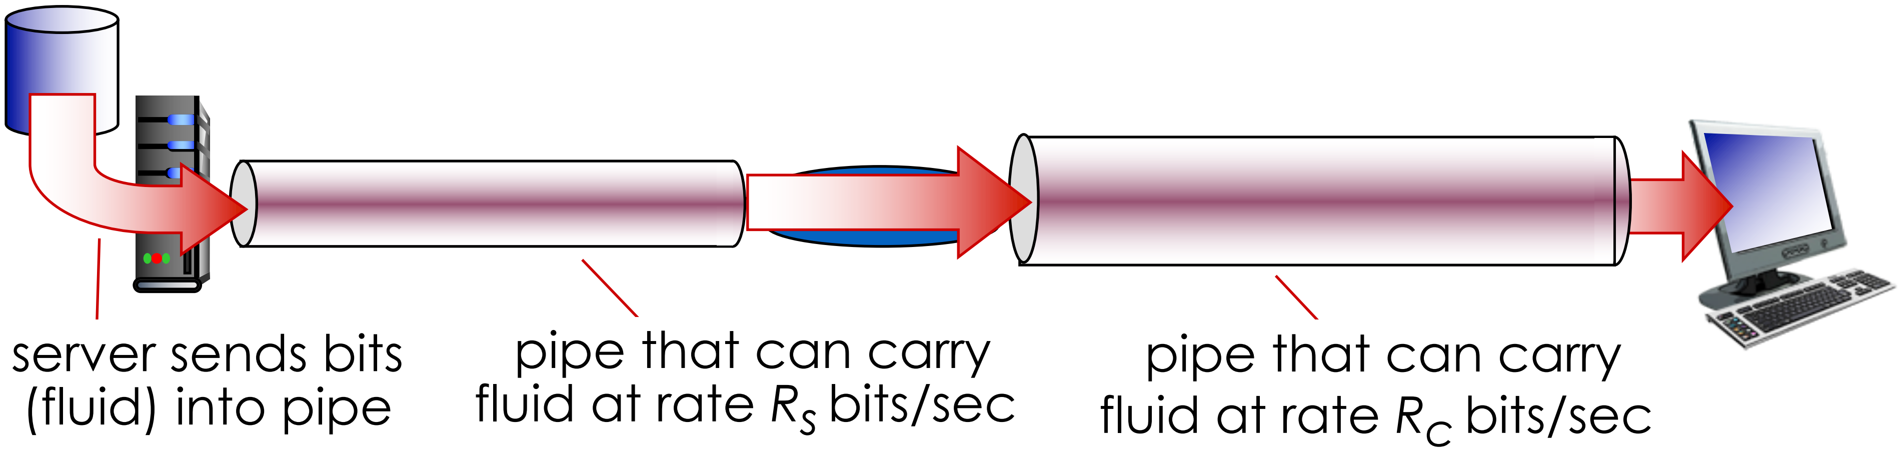

Throughput

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

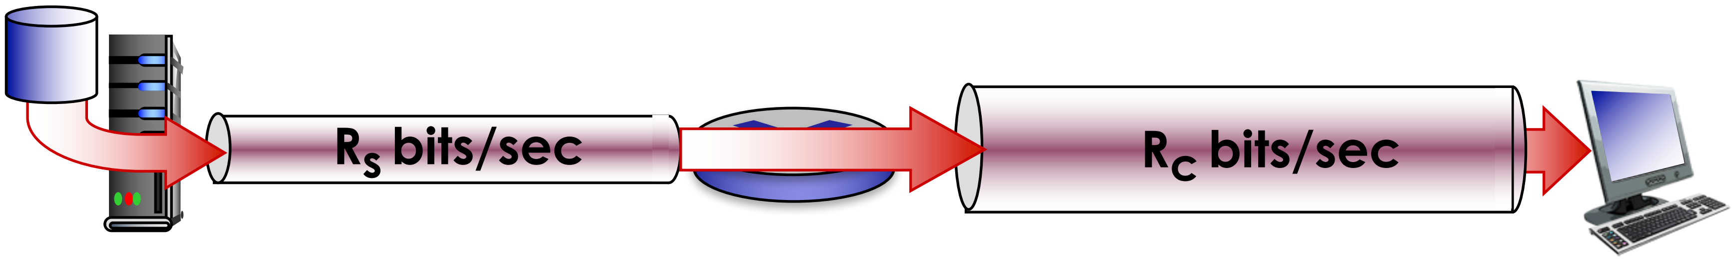

- $R_s > R_c$ What is average end-end throughput?

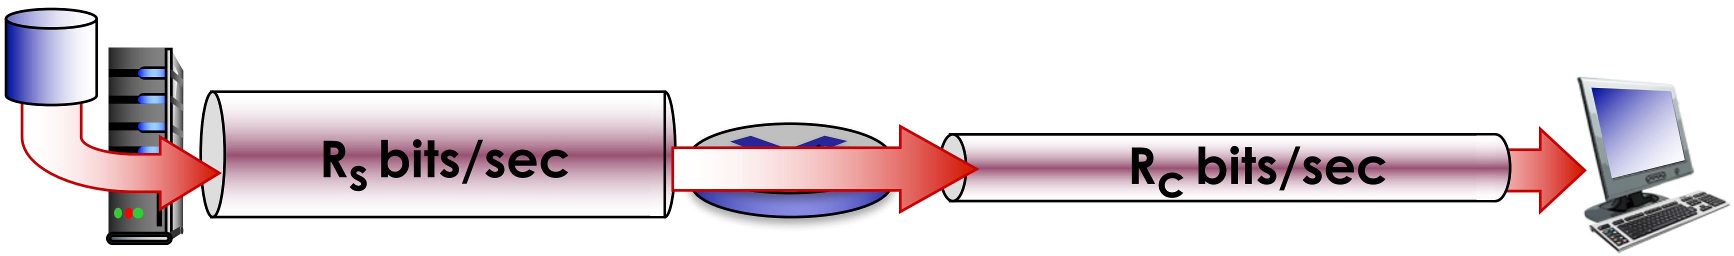

bottleneck link

- The link along a path with the minimum capacity
- The bottleneck link limits the end-end throughput

Throughput: Internet Scenario

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

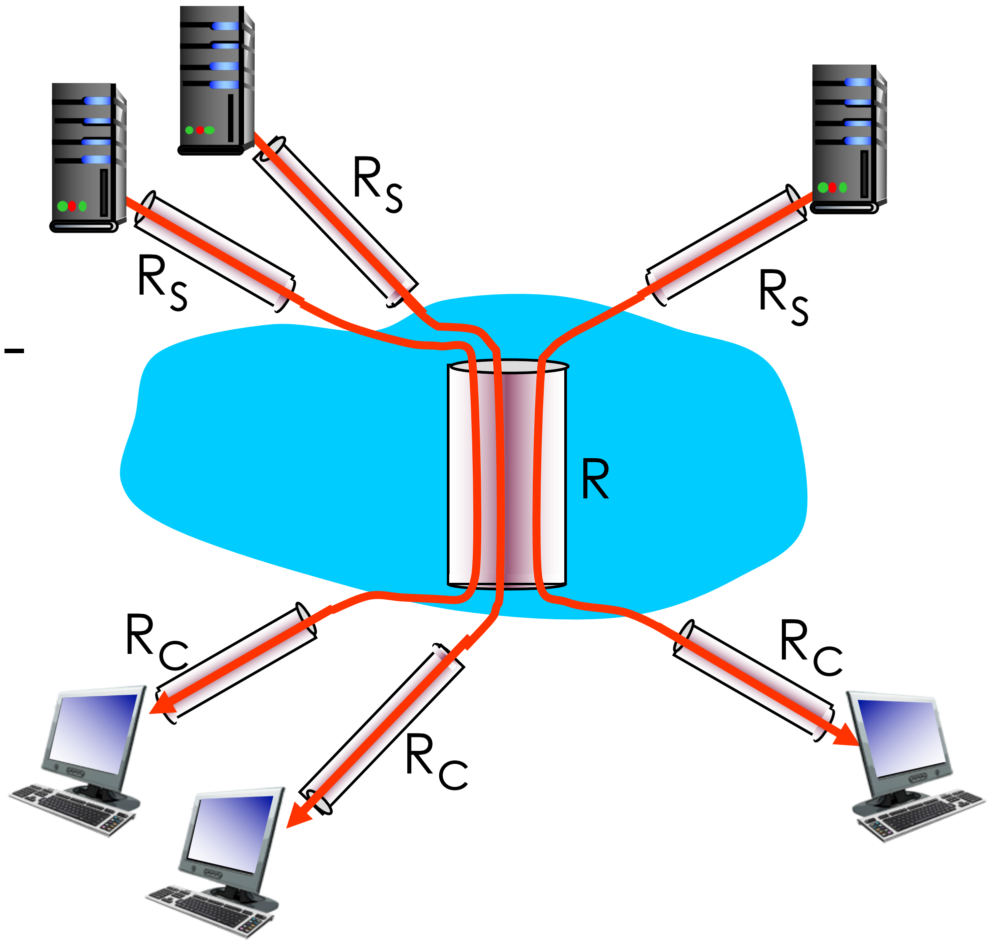

10 connections (fairly) share backbone bottleneck link *R* bits/sec

Outline

- What's the Internet?
- What's a protocol?
- Network edge
  - hosts, access net, physical media
- Network core
  - packet/circuit switching, Internet structure
- Performance
  - loss, delay, throughput
- **Protocol layers, service models**
- History

Protocol "Layers"

- Networks are complex, with many "pieces"
  - hosts
  - routers
  - links of various media
  - applications
  - protocols
  - hardware, software
- How to simplify the organization of a network?
  → **Layering!**
  - Build a structure: divide tasks based on their functionality and assign each task to a proper layer
  - Similar to the airline system

Layering of Airline Functionality

layers: each layer implements a service

- via its own internal-layer actions
- relying on services provided by layer below

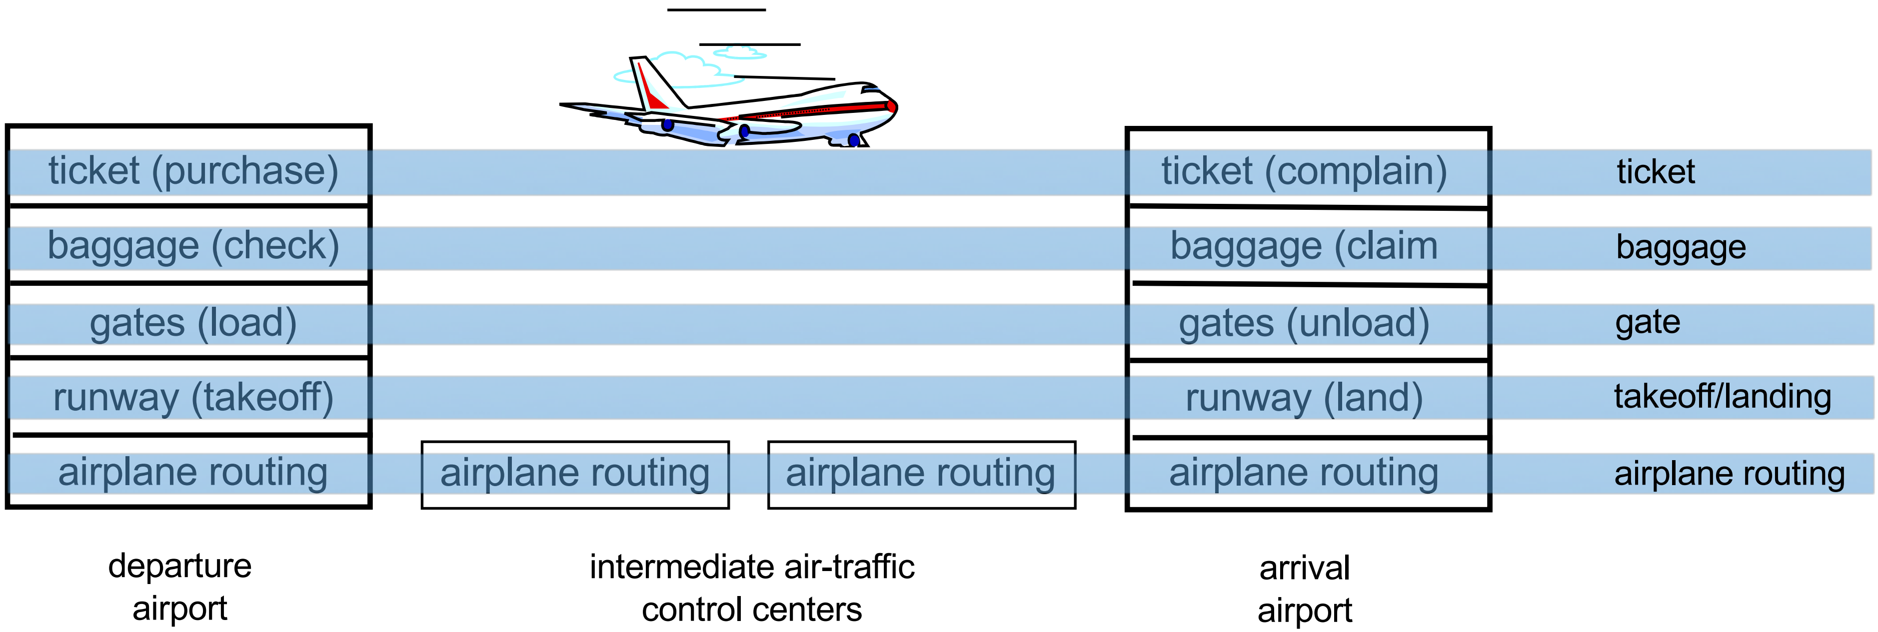

Why Layering?

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 the rest of a system
- Layering considered harmful?
  - May exist dependency between layers
  - If so, cross-layer designs might be preferable

Internet Protocol Stack

- ***Application:***
  - supporting network services
  - FTP, SMTP, HTTP, DNS (message)
- ***Transport:***
  - process-to-process data transfer
  - TCP, UDP (segment)
- ***Network (aka IP):***
  - end-to-end routing from source to destination (along a path)
  - IP, routing protocols (packet)
- ***Link:***
  - data transfer between neighboring network elements (host-to-host)
  - Ethernet, 802.11, PPP (frame)
- ***Physical:***
  - bits on the communication channels, i.e.,"wire" or "air" (symbol)

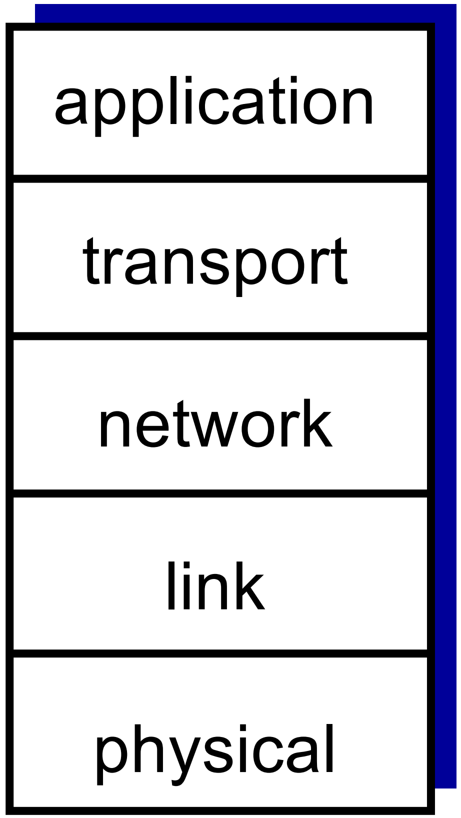

Top-down approach

ISO/OSI Reference model

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

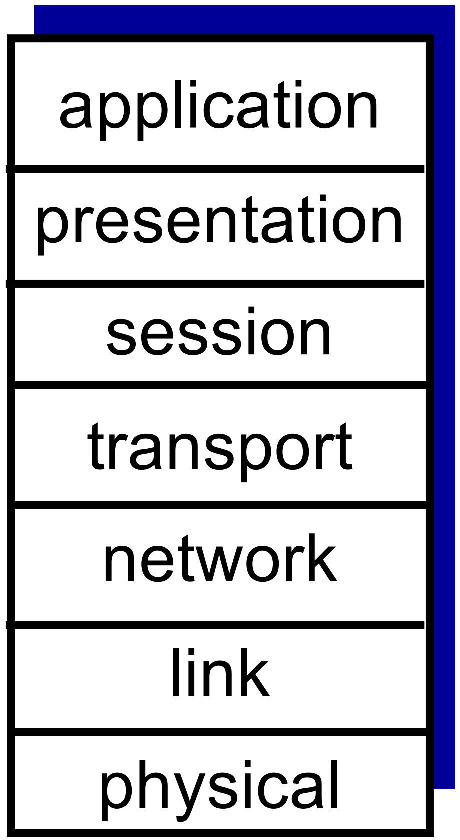

Encapsulation

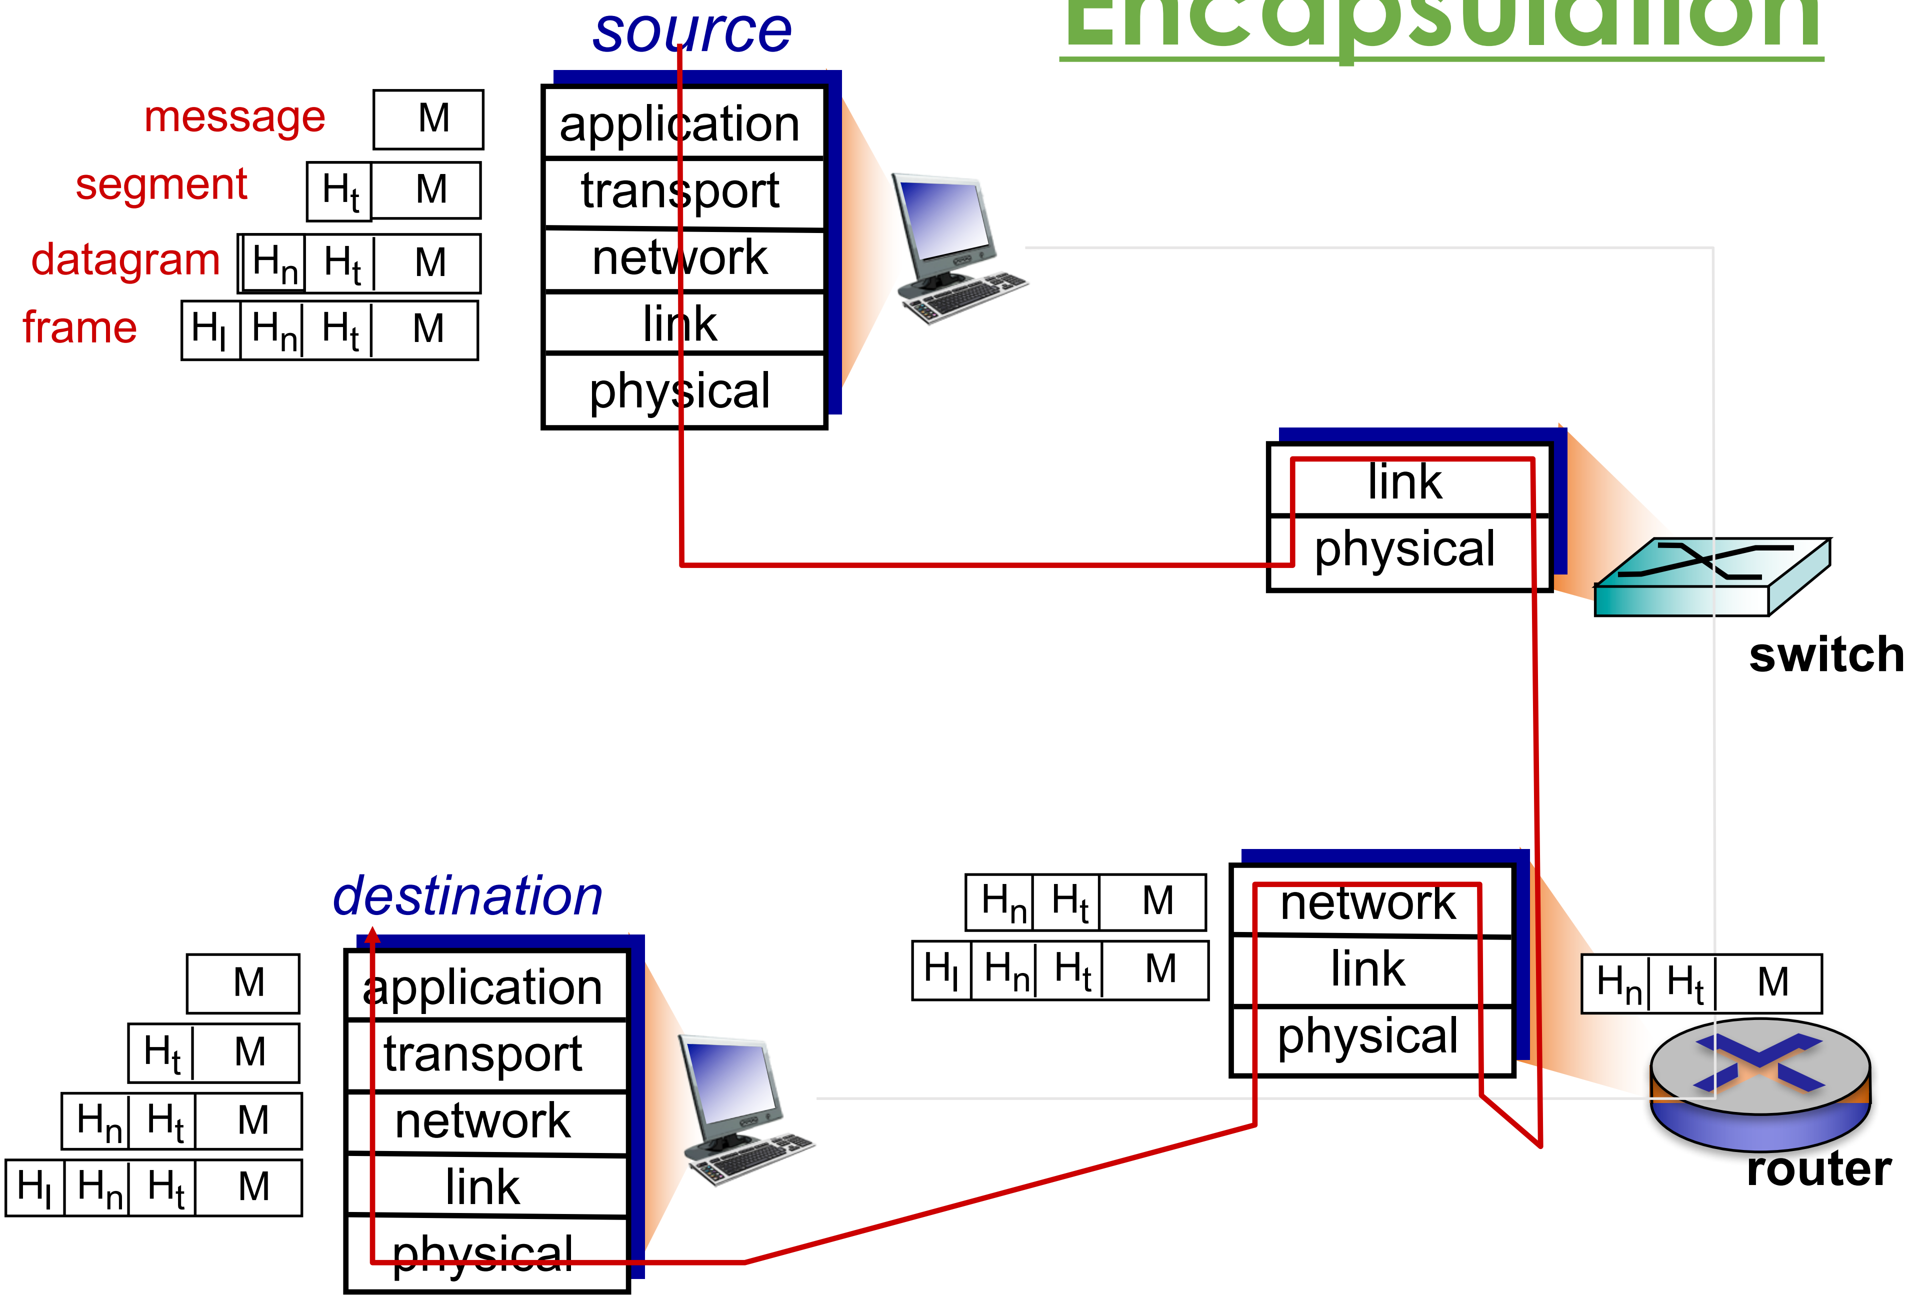

Outline

- What's the Internet?
- What's a protocol?
- Network edge
  - hosts, access net, physical media
- Network core
  - packet/circuit switching, Internet structure
- Performance
  - loss, delay, throughput
- Protocol layers, service models
- **History**

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 demo
  - NCP (Network Control Protocol) first host-host protocol
  - first e-mail program
  - ARPAnet has 15 nodes

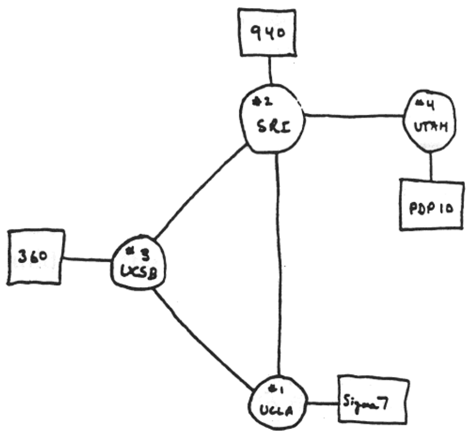

THE ARPA NETWORK

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

Internet History

2005-present

- ~5B devices attached to Internet (2016)
  - smartphones and tablets
- aggressive deployment of broadband access
- increasing ubiquity of high-speed wireless access
- emergence of online social networks:
  - Facebook: ~ one billion users
- service providers (Google, Microsoft) create their own networks
  - bypass Internet, providing “instantaneous” access to search, video content, email, etc.
- e-commerce, universities, enterprises running their services in “cloud” (e.g., Amazon EC2)

Summary

- What's the Internet?
- What's a protocol?
- Network edge
  - hosts, access net, physical media
- Network core
  - packet/circuit switching, Internet structure
- Performance
  - loss, delay, throughput
- Protocol layers, service models
- History