Chapter 1
Introduction

A note on the use of these PowerPoint slides:

We’re making these slides freely available to all (faculty, students, readers). They’re in PowerPoint form so you see the animations; and can add, modify, and delete slides (including this one) and slide content to suit your needs. They obviously represent a *lot* of work on our part. In return for use, we only ask the following:

- If you use these slides (e.g., in a class) that you mention their source (after all, we’d like people to use our book!)
- If you post any slides on a www site, that you note that they are adapted from (or perhaps identical to) our slides, and note our copyright of this material.

For a revision history, see the slide note for this page.

Thanks and enjoy! JFK/KWR

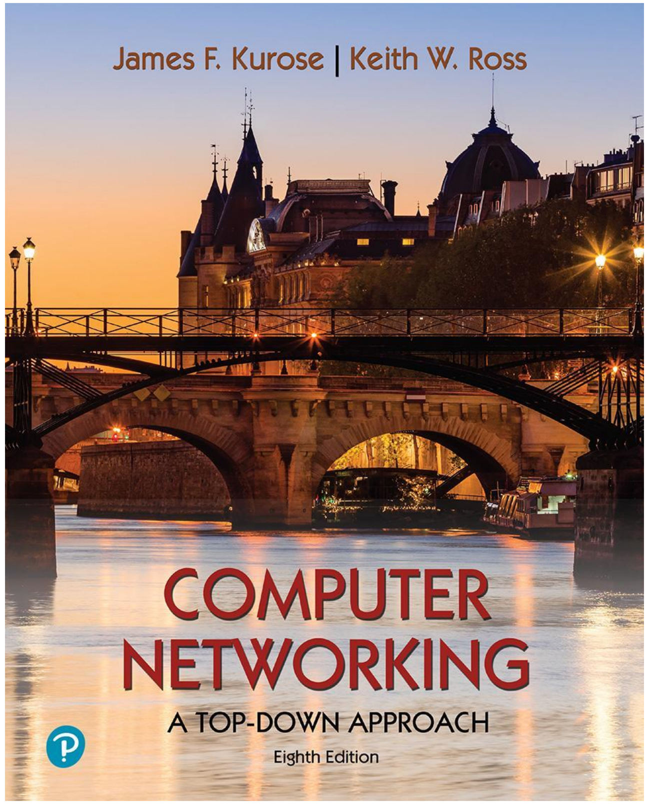

Computer Networking: A Top-Down Approach
8^{th} edition
Jim Kurose, Keith Ross
Pearson, 2020

Part I: The Internet

(with a capital I)

The Internet: a "nuts and bolts" view

Billions of connected computing *devices*:

- *hosts = end systems*
- running *network apps* at Internet's "edge"

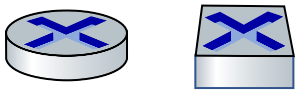

Packet switches: forward packets (chunks of data)

- *routers, switches*

Communication links

- fiber, copper, radio, satellite
- transmission rate: *bandwidth*

Networks

- collection of devices, routers, links: managed by an organization

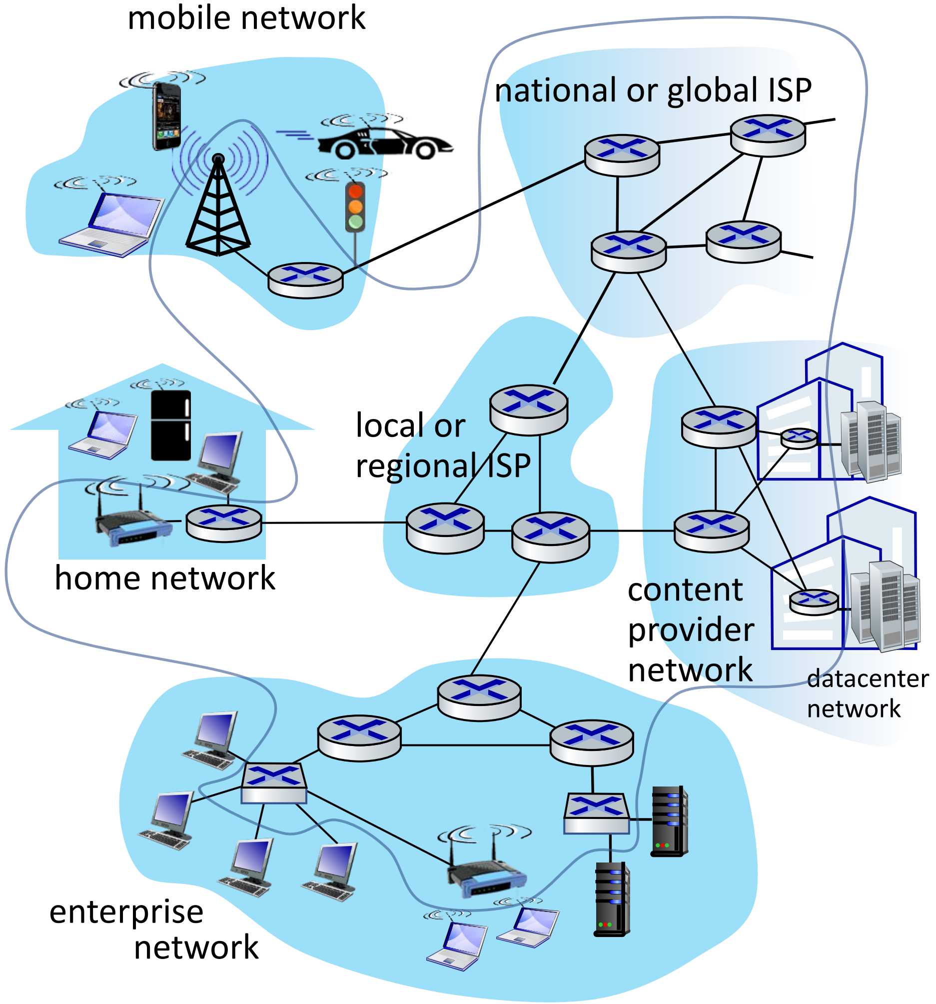

“Fun” Internet-connected devices

Amazon Echo

Internet refrigerator

IP picture frame

Pacemaker & Monitor

Tweet-a-watt: monitor energy use

Web-enabled toaster + weather forecaster

Security Camera

Slingbox: remote control cable TV

AR devices

Internet phones

sensorized, bed mattress

Fitbit

Others?

The Internet: a "nuts and bolts" view

- *Internet: "network of networks"*
  - Interconnected ISPs
- *protocols* are *everywhere*
  - control sending, receiving of messages
  - e.g., HTTP (Web), streaming video, Skype, TCP, IP, WiFi, 4G, Ethernet
- *Internet standards*
  - RFC: Request for Comments
  - IETF: Internet Engineering Task Force

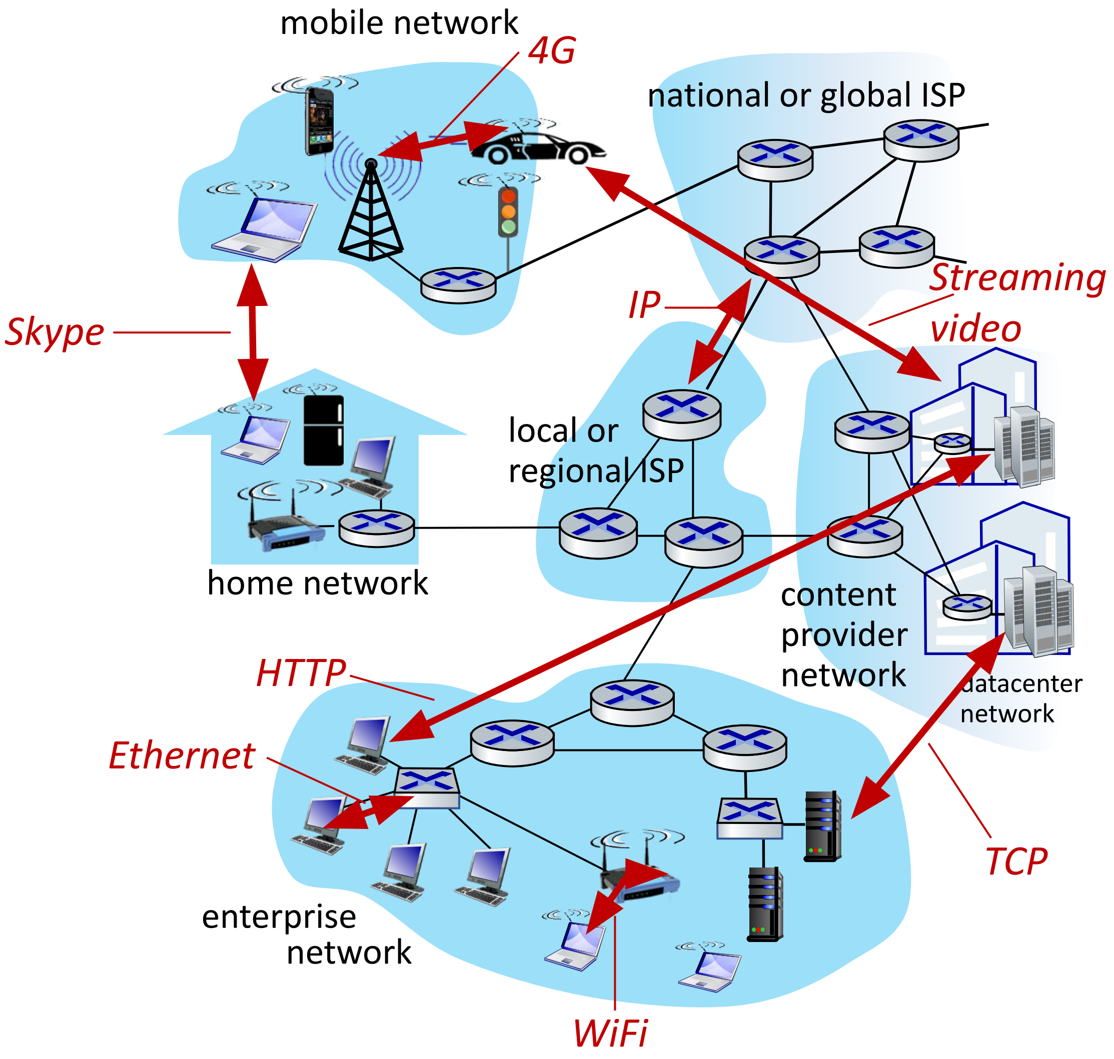

The Internet: a "service" view

- *Infrastructure* that provides services to applications:
  - Web, streaming video, multimedia teleconferencing, email, games, e-commerce, social media, inter-connected appliances, …
- provides *programming interface* to distributed applications:
  - "hooks" allowing sending/receiving apps to "connect" to, use Internet transport service
  - provides service options, analogous to postal service

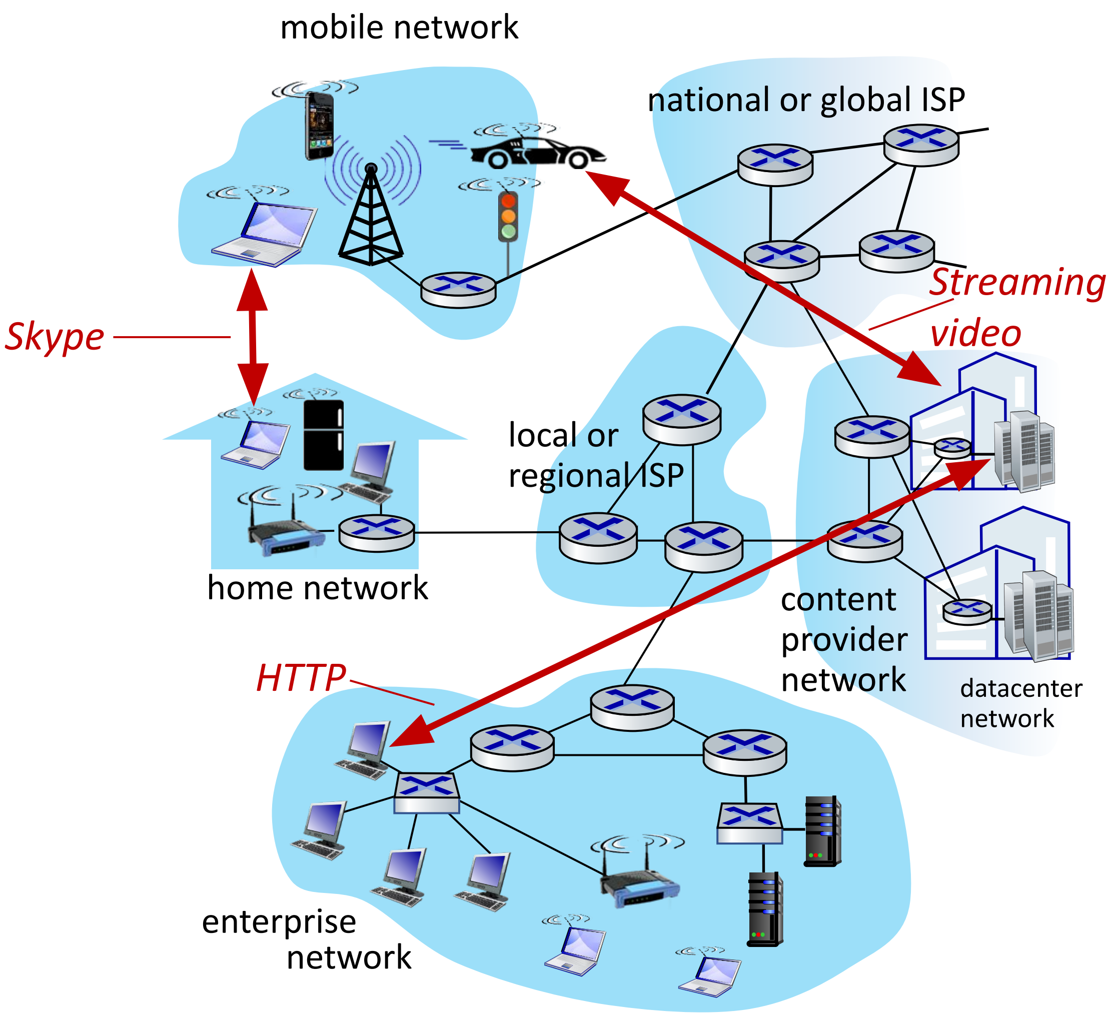

What's a protocol?

Human protocols:

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

… specific messages sent

… specific actions taken when message received, or other events

Network protocols:

- computers (devices) rather than humans
- all communication activity in Internet governed by protocols

Protocols define the format, order of messages 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:

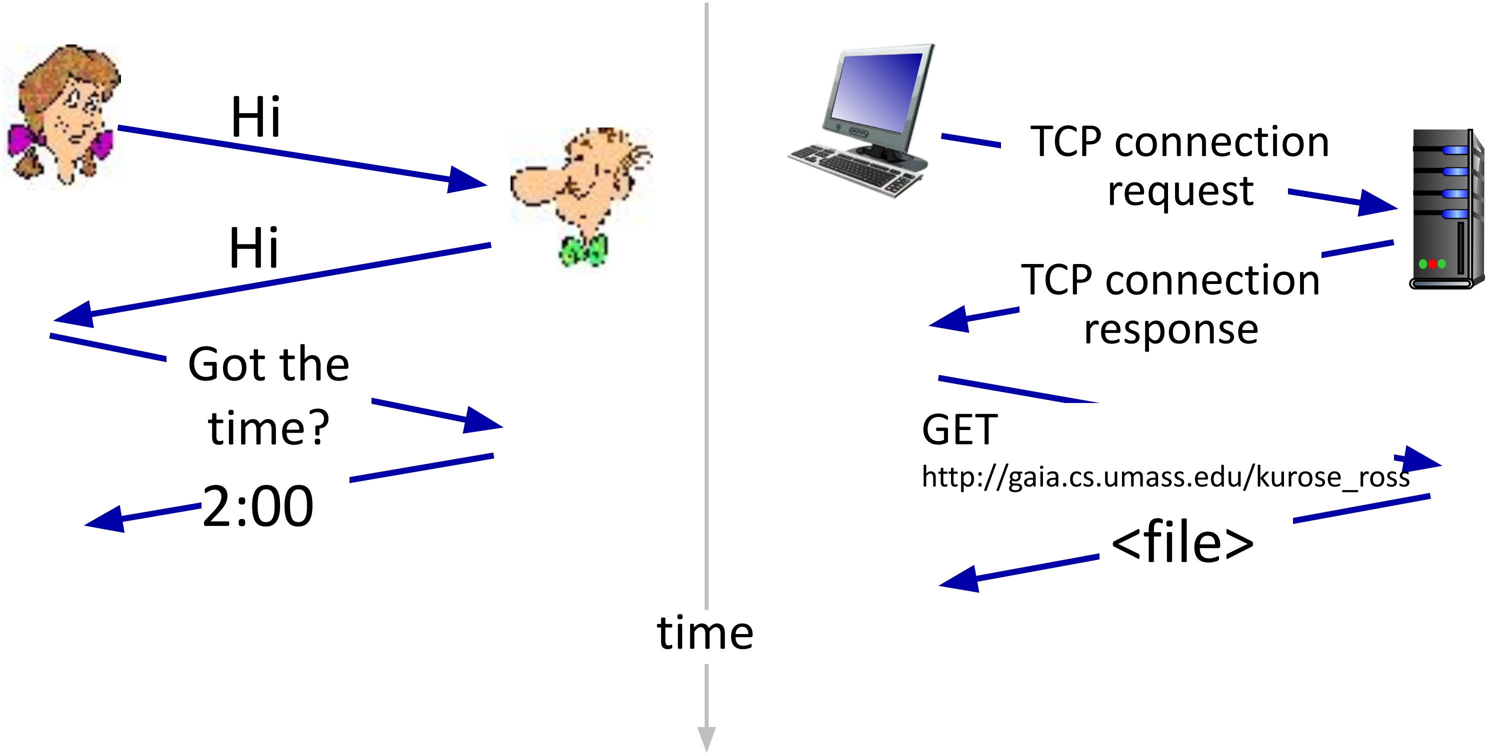

A closer look at Internet structure

Network edge:

- hosts: clients and servers
- servers often in data centers

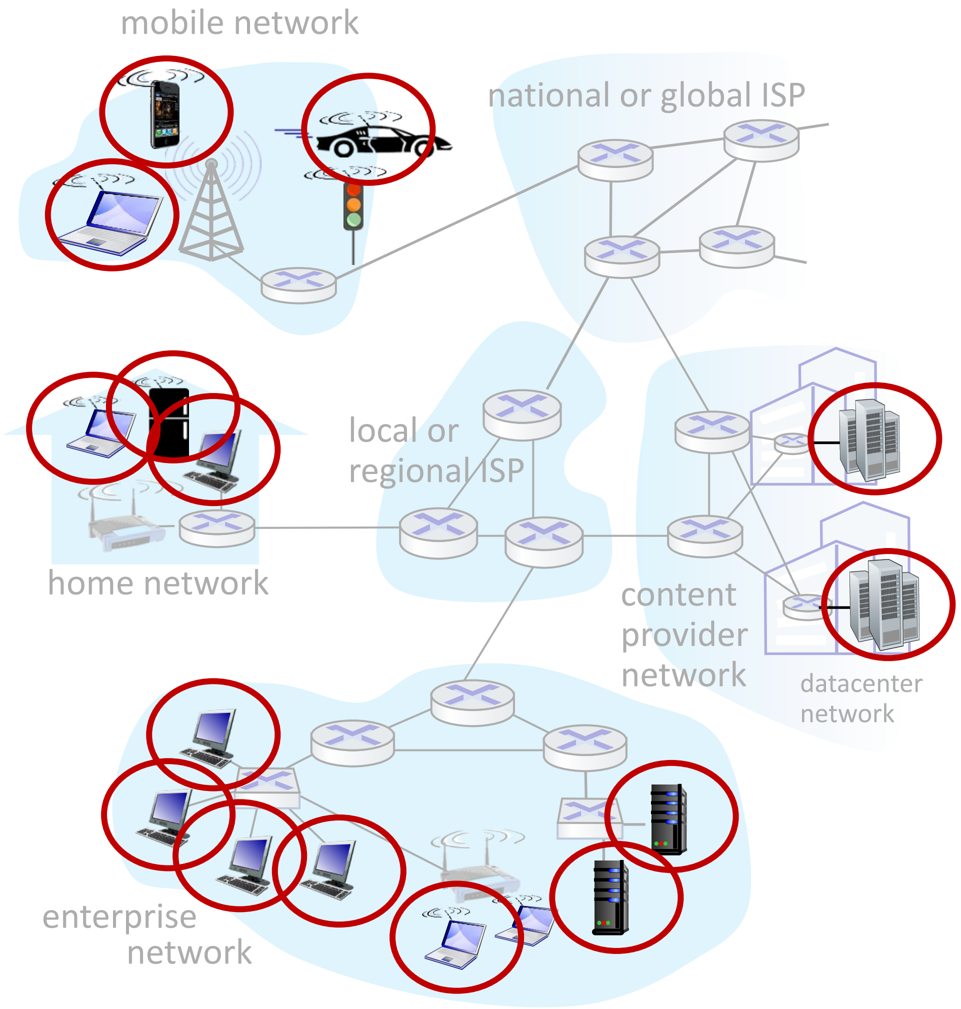

Client vs Server

- When used in connection-oriented protocols:
  - Client: party that creates connection
  - Server: party that awaits connections
- When used in application-level protocols:
  - Client: party that requests something
  - Server: party that provides something
- When used to categorize end hosts:
  - Client: computer that mostly runs programs that act as application-level clients
  - Server: computer that mostly runs programs that act as application-level servers

A closer look at Internet structure

Network edge:

- hosts: clients and servers
- servers often in data centers

Access networks, physical media:

- wired, wireless communication links

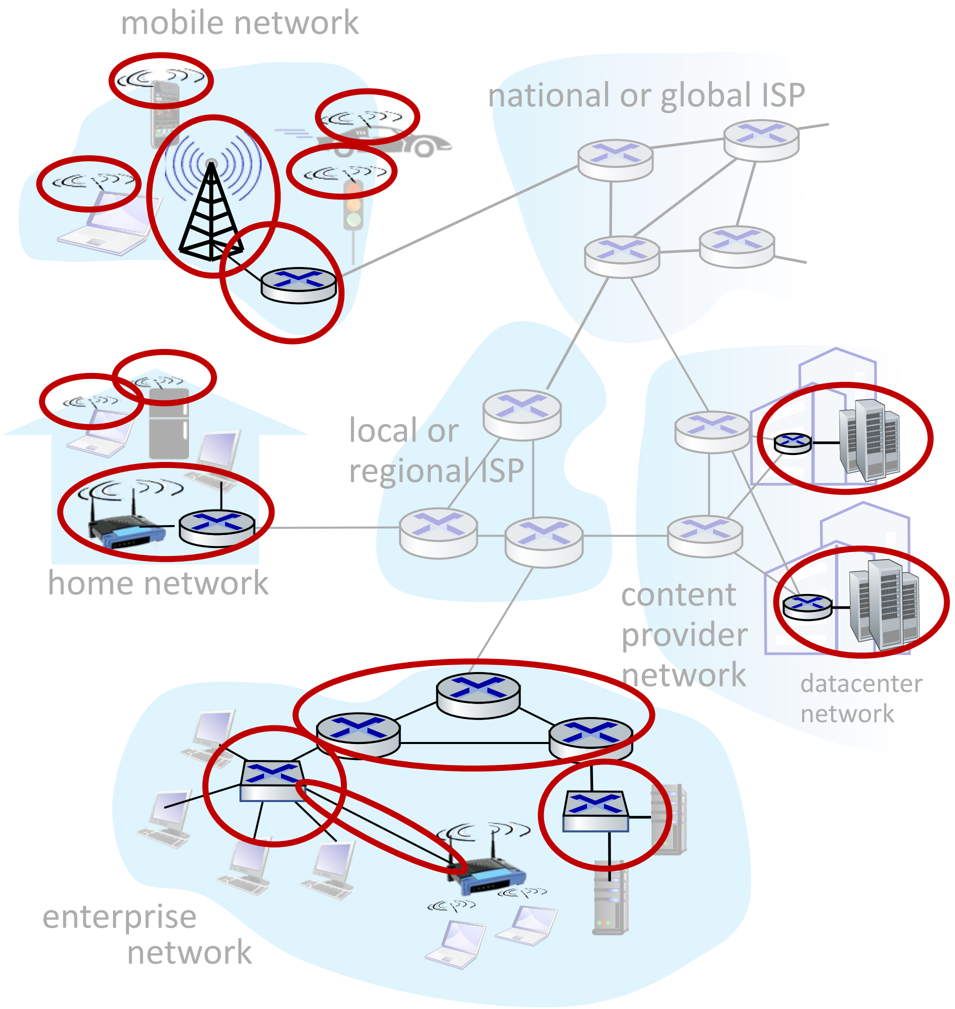

A closer look at Internet structure

Network edge:

- hosts: clients and servers
- servers often in data centers

Access networks, physical media:

- wired, wireless communication links

Network core:

- interconnected routers
- network of networks

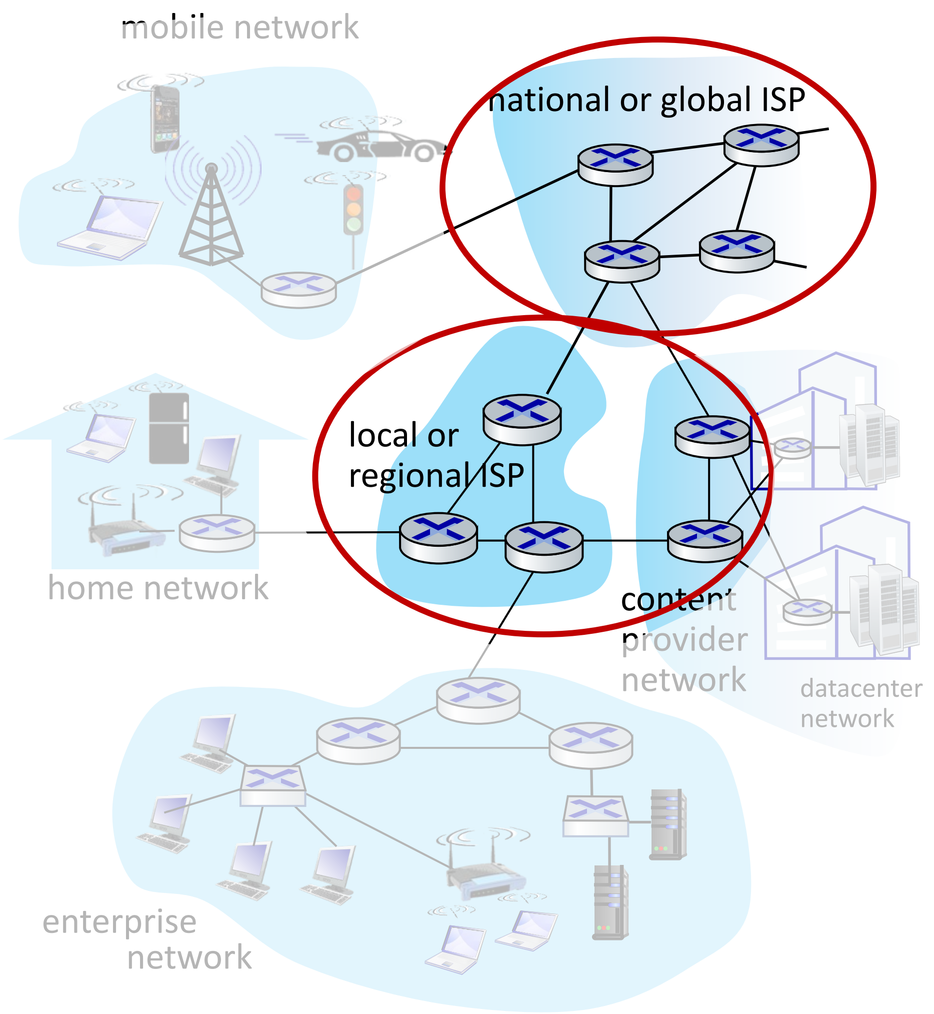

Host: sends *packets* of data

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*

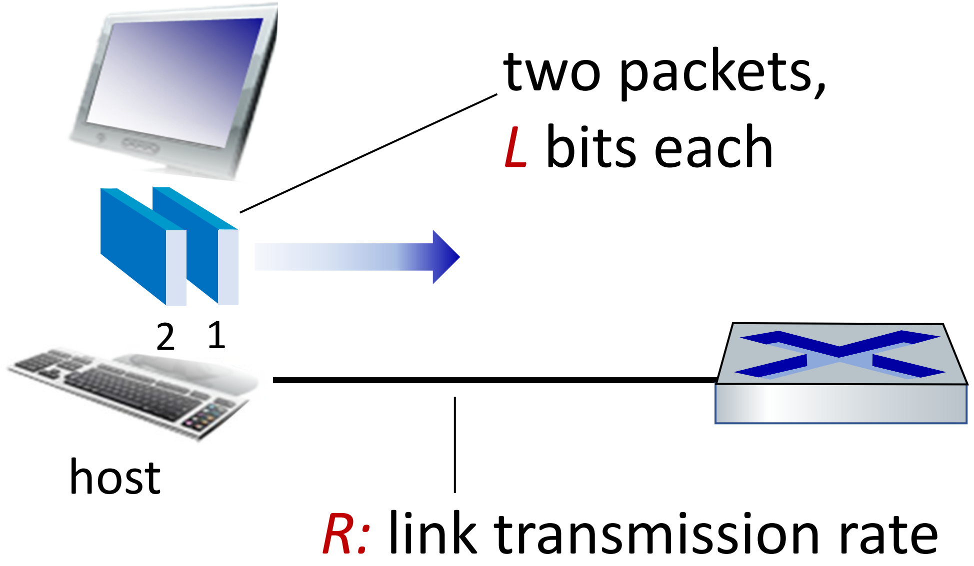

$$\text{packet transmission delay} = \text{time needed to transmit } L\text{-bit packet into link} = \frac{L \text{ (bits)}}{R \text{ (bits/sec)}}$$

Links: physical media

- bit: propagates between transmitter/receiver pairs
- physical link: what lies between transmitter & receiver
- guided media:
  - signals propagate in solid media: copper, fiber, coax
- unguided media:
  - signals propagate freely, e.g., radio

Twisted pair (TP)

- two insulated copper wires
  - Category 5: 100 Mbps, 1 Gbps Ethernet
  - Category 6: 10Gbps Ethernet

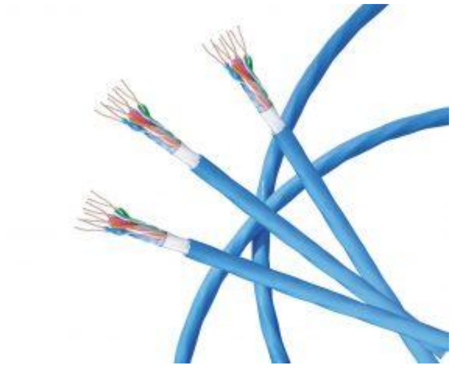

- Other examples:
  Coax, Fiber
  Wireless links; Satellite;

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 (WiFi, 4G/5G)

What to look for:

- transmission rate (bits per second) of access network?
- shared or dedicated access among users?

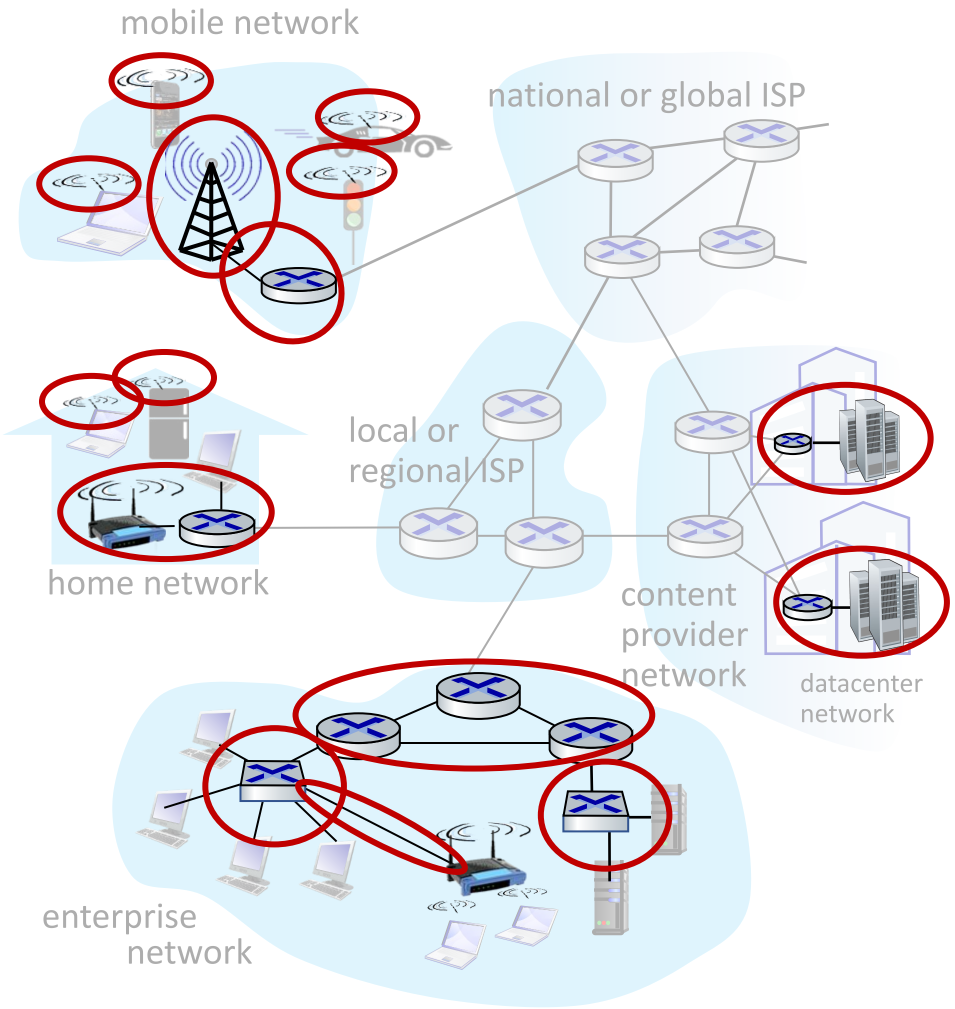

Access networks: cable-based access

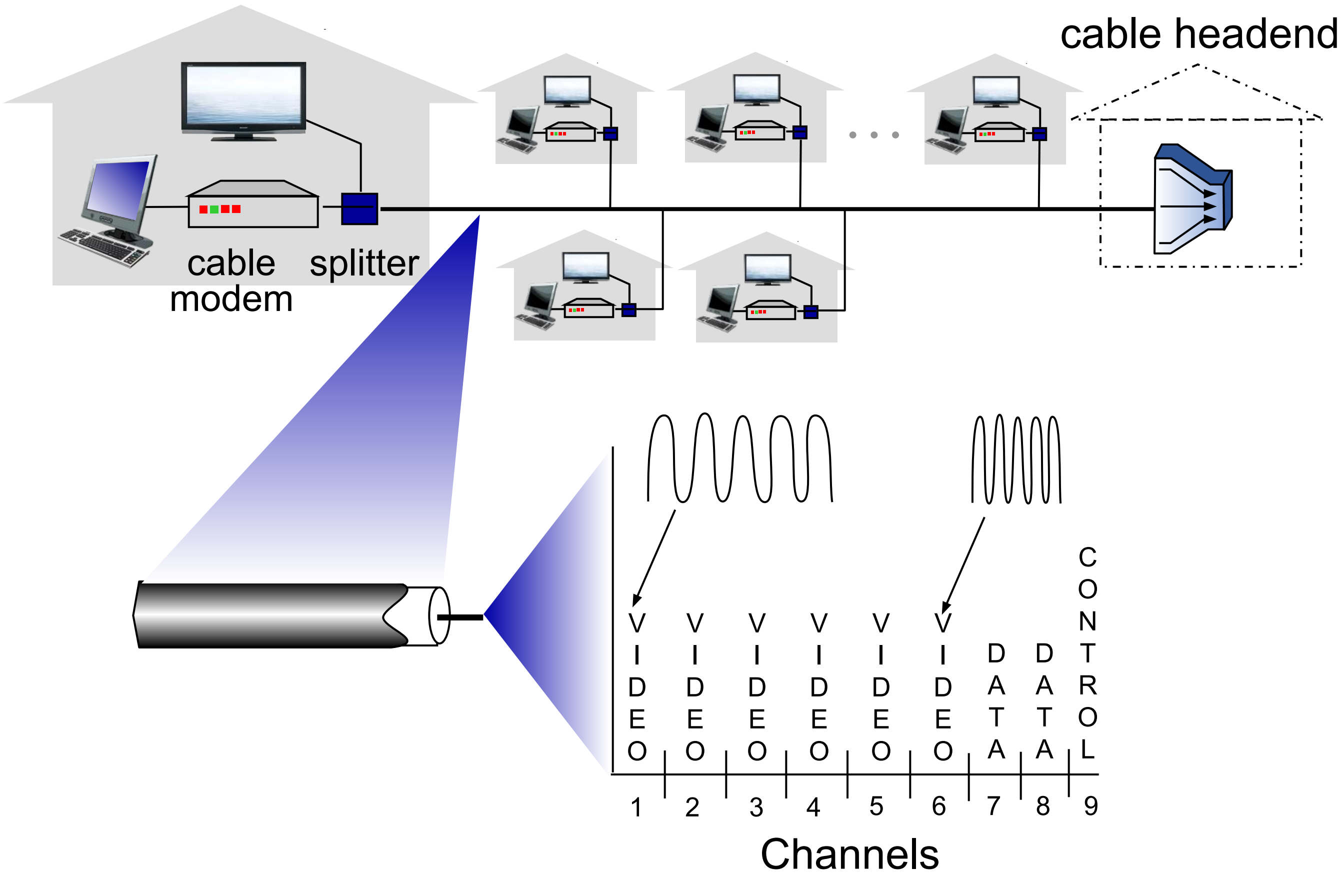

frequency division multiplexing (FDM): different channels transmitted in different frequency bands

Access networks: cable-based access

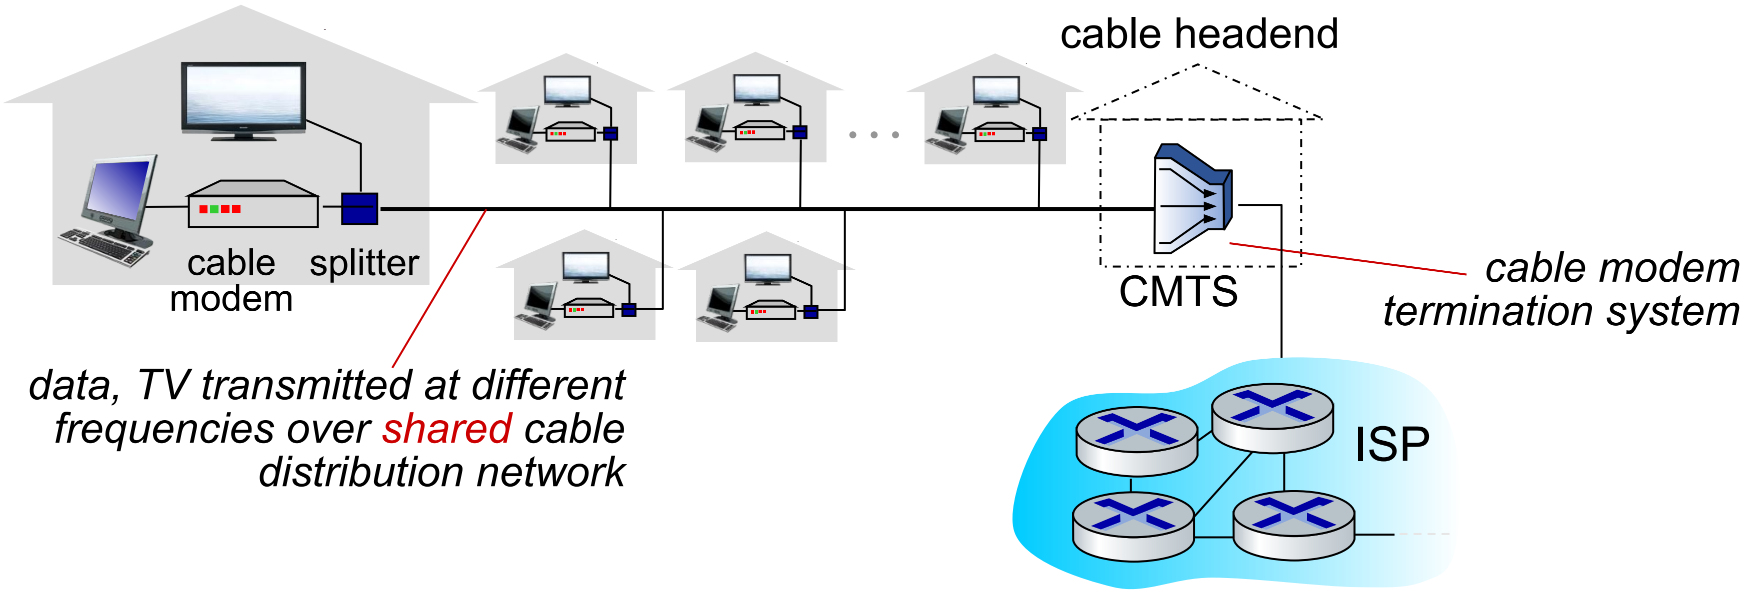

- HFC: hybrid fiber coax
  - asymmetric: up to 40 Mbps – 1.2 Gbs downstream transmission rate, 30-100 Mbps upstream transmission rate
- network of cable, fiber attaches homes to ISP router
  - homes *share access network* to cable headend

192.168.100.1/cmSignal.htm

Not secure | 192.168.100.1/cmSignal.htm

SURFboard

Cable Modem

Status | Signal | Addresses | Configuration | Logs | Open Source | Help

This page provides information about the current upstream and downstream signal status of your Cable Modem.

Downstream	Bonding Channel Value			
Channel ID	34	35	36	37
Frequency	429000000 Hz	435000000 Hz	441000000 Hz	447000000 Hz
Signal to Noise Ratio	37 dB	37 dB	37 dB	37 dB
Downstream Modulation	QAM256	QAM256	QAM256	QAM256
Power Level The Downstream Power Level reading is a snapshot taken at the time this page was requested. Please Reload/Refresh this Page for a new reading	3 dBmV	3 dBmV	3 dBmV	3 dBmV

Example:
Dr. Back's Cable Modem Signal Status page.
4 downstream channels (bonded) 34-37 and
4 upstream channels (25-28)

Upstream	Bonding Channel Value			
Channel ID	26	28	27	25
Frequency	23700000 Hz	36500000 Hz	30100000 Hz	17300000 Hz
Ranging Service ID	13185	13185	13185	13185
Symbol Rate	5.120 Msym/sec	5.120 Msym/sec	5.120 Msym/sec	5.120 Msym/sec
Power Level	38 dBmV	39 dBmV	38 dBmV	39 dBmV
Upstream Modulation	[2] QPSK [1] 32QAM [3] 64QAM	[2] QPSK [1] 32QAM [3] 64QAM	[2] QPSK [1] 32QAM [3] 64QAM	[2] QPSK [1] 32QAM [3] 64QAM
Ranging Status	Success	Success	Success	Success

Access networks: home networks

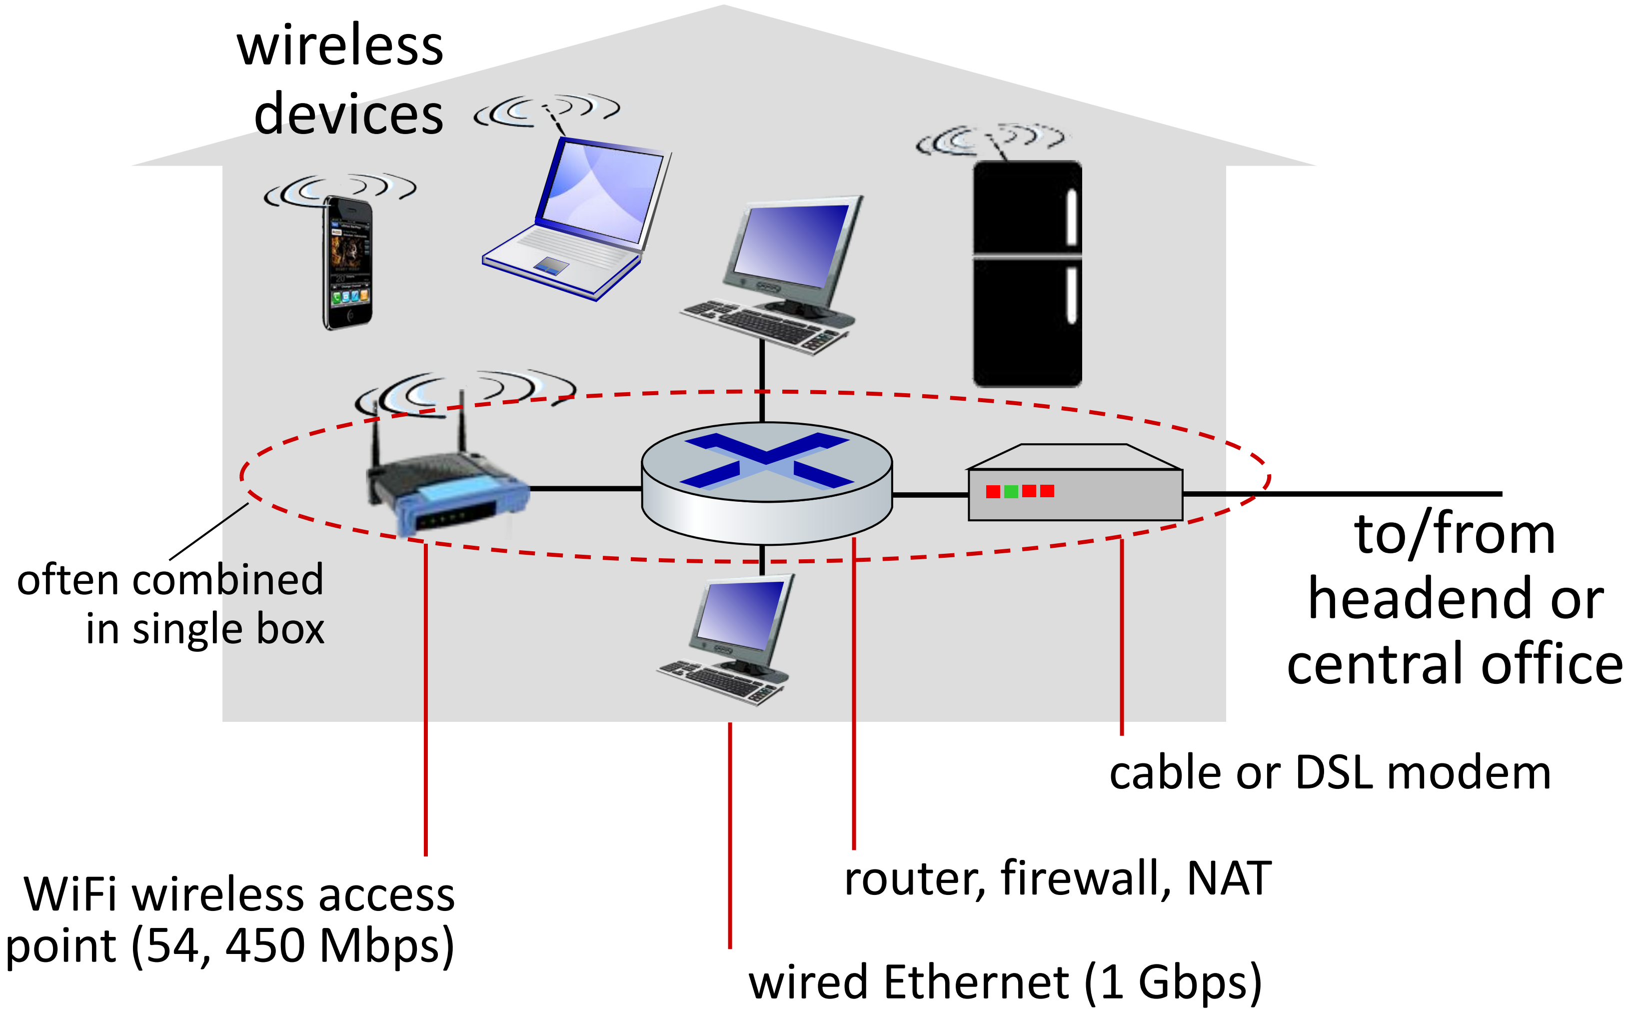

Dr. Back's somewhat improvised structured wiring cabinet with cable modem and Gigabit switch

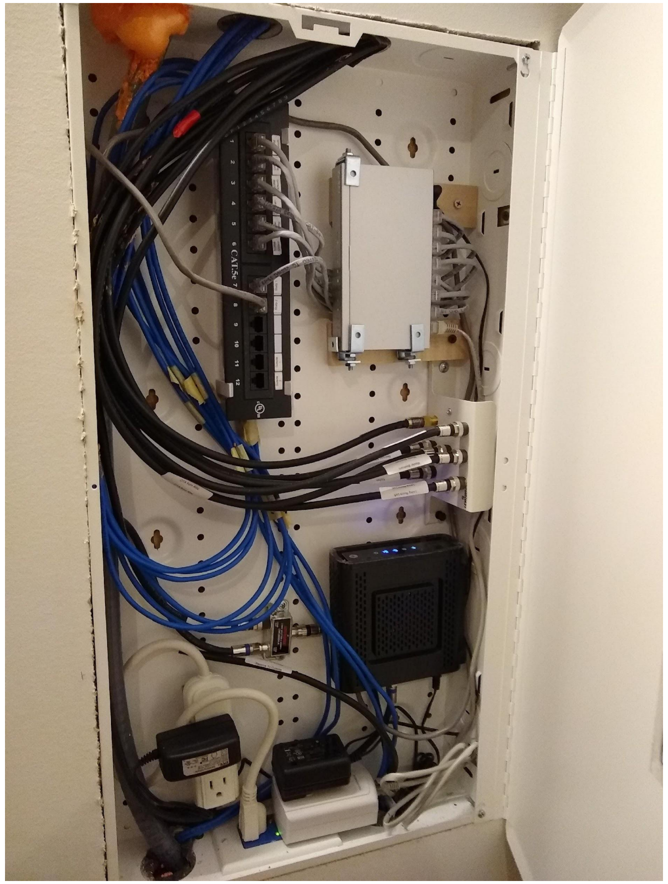

2021 Update - Dr. Back got fiber from WideOpenBlacksburg

```
$ traceroute rlogin.cs.vt.edu
traceroute to rlogin.cs.vt.edu (128.173.236.214), 30 hops max, 60 byte packets
 1  _gateway (192.168.1.1)  0.584 ms  0.548 ms  0.522 ms
 2  216.252.202.1 (216.252.202.1)  1.788 ms  1.758 ms  1.917 ms
 3  1715-rr-transitsw1.bnt.com (216.252.194.90)  1.690 ms  1.662 ms  1.821 ms
 4  isb-border.xe-5-0-3.892.cns.vt.edu (192.70.187.185)  1.359 ms  1.309 ms  1.279 ms
 5  isb-core.lo0.2000.cns.vt.edu (198.82.1.139)  1.677 ms  1.846 ms  1.604 ms
 6  isb-bar.lo0.2000.cns.vt.edu (198.82.1.174)  2.433 ms  5.987 ms  6.129 ms
 7  rlogin.cs.vt.edu (128.173.236.214)  2.012 ms !X  2.112 ms !X  2.056 ms !X
```

(does not leave Blacksburg!!!)

Wireless access networks

Shared *wireless* access network connects end system to router

- via base station aka “access point”

Wireless local area networks (WLANs)

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

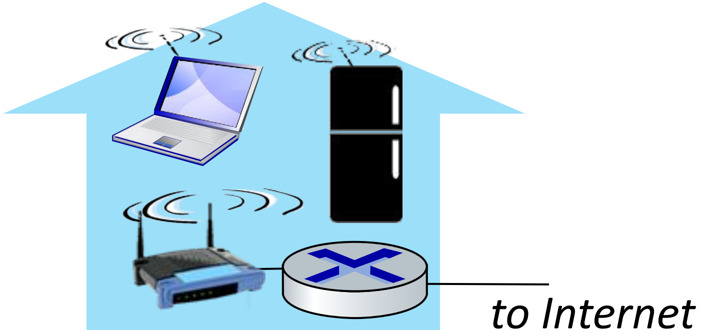

Wide-area cellular access networks

- provided by mobile, cellular network operator (10’s km)
- 10’s Mbps
- 4G cellular networks (5G coming)

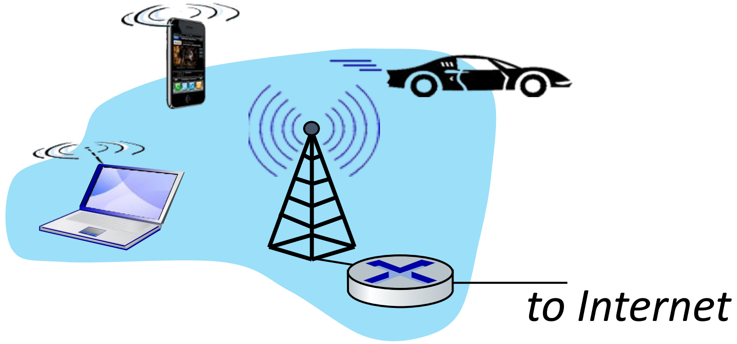

Access networks: enterprise networks

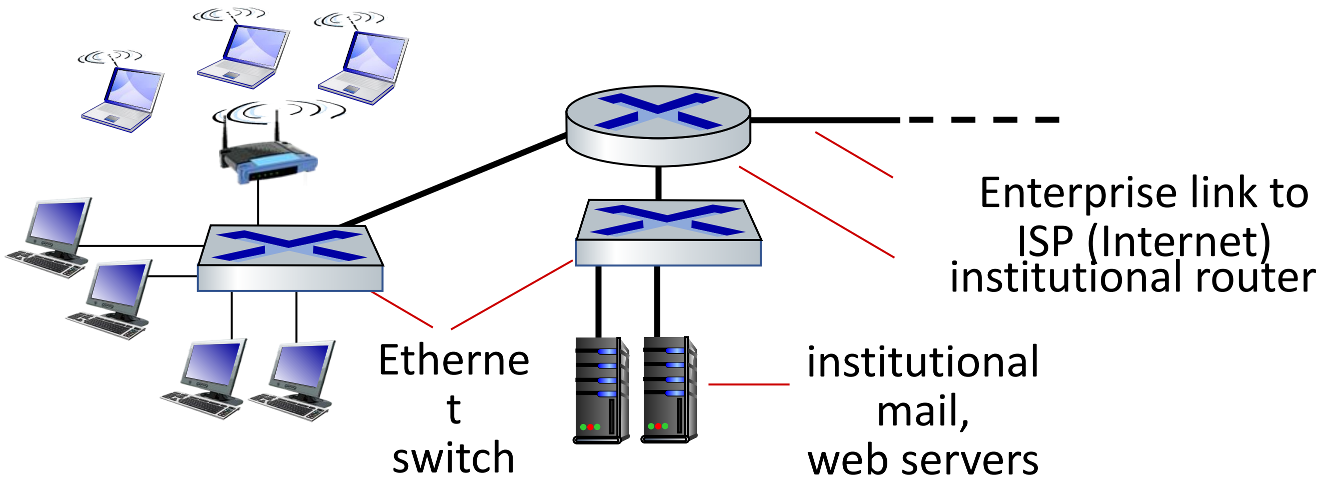

- companies, universities, etc.
- mix of wired, wireless link technologies, connecting a mix of switches and routers (we'll cover differences shortly)
  - Ethernet: wired access at 100Mbps, 1Gbps, 10Gbps
  - WiFi: wireless access points at 11, 54, 450 Mbps

The 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

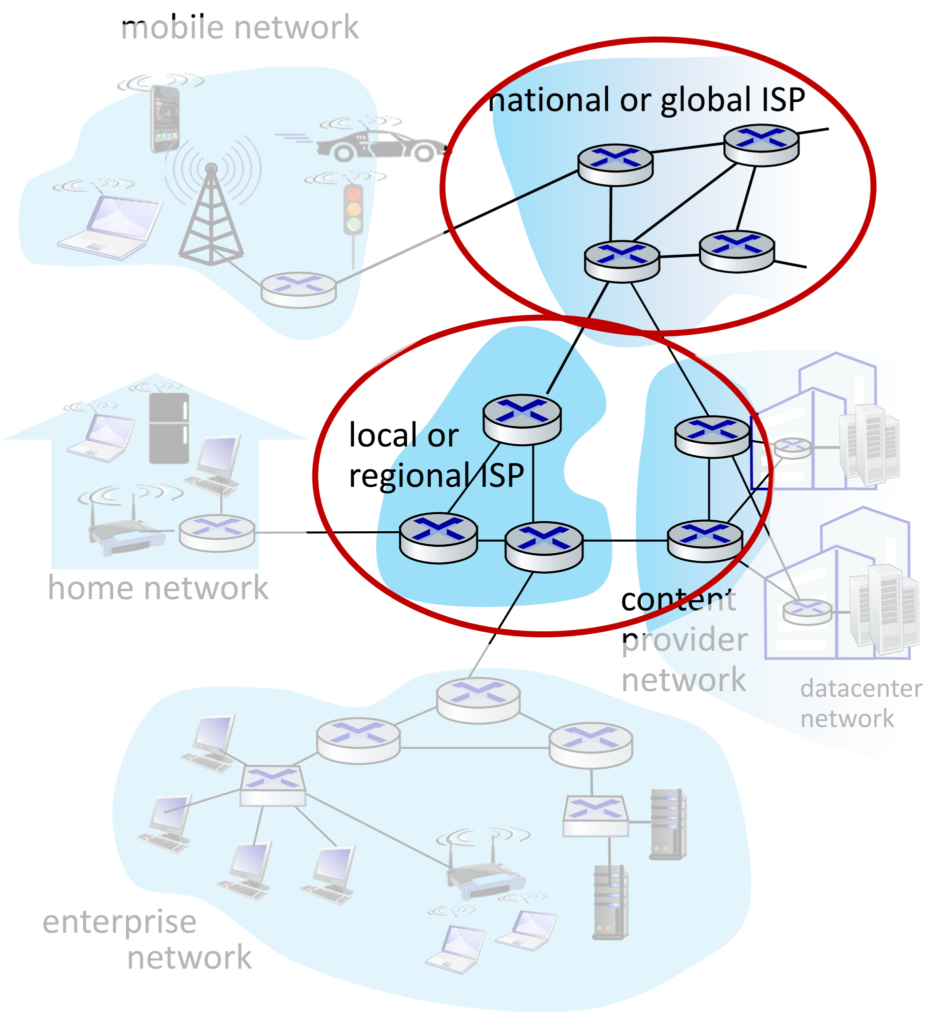

Two key network-core functions

Forwarding:

- *local* action: move arriving packets from router's input link to appropriate router output link

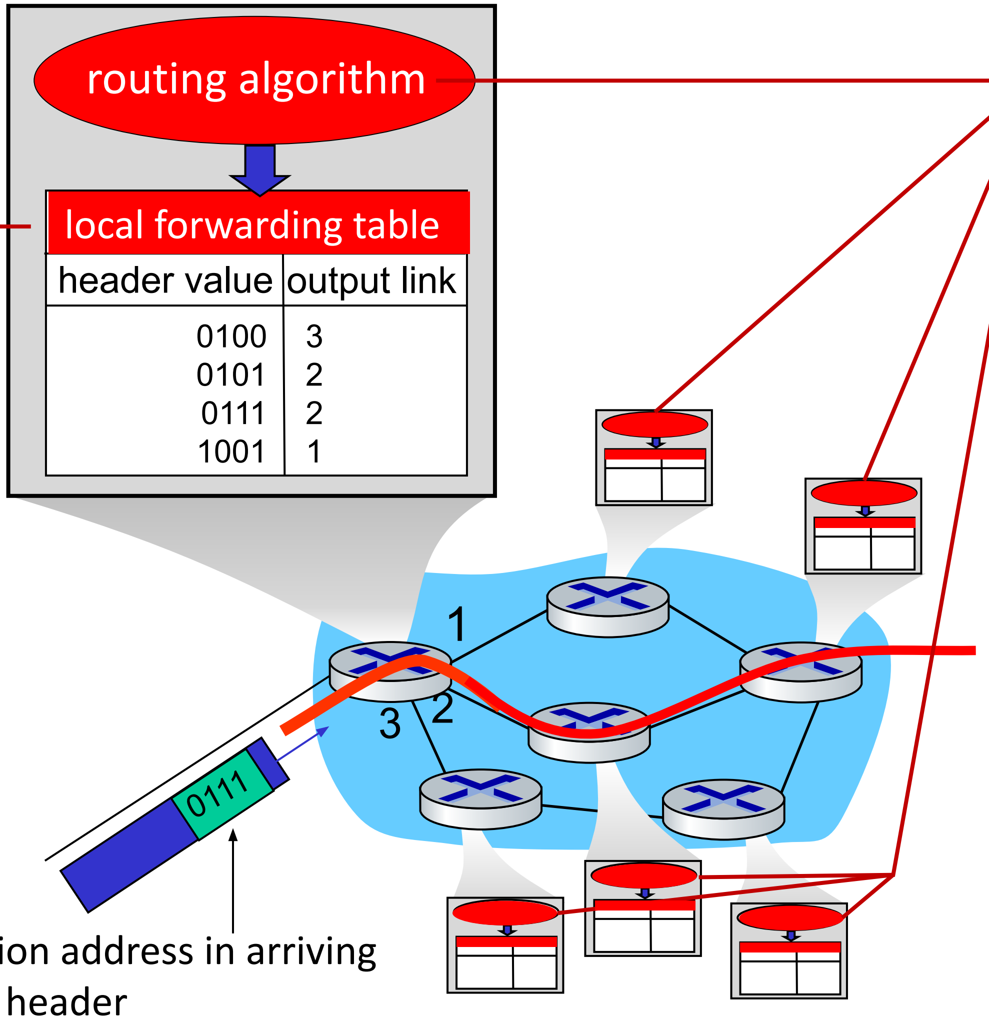

Routing:

- *global* action: determine source-destination paths taken by packets
- routing algorithms

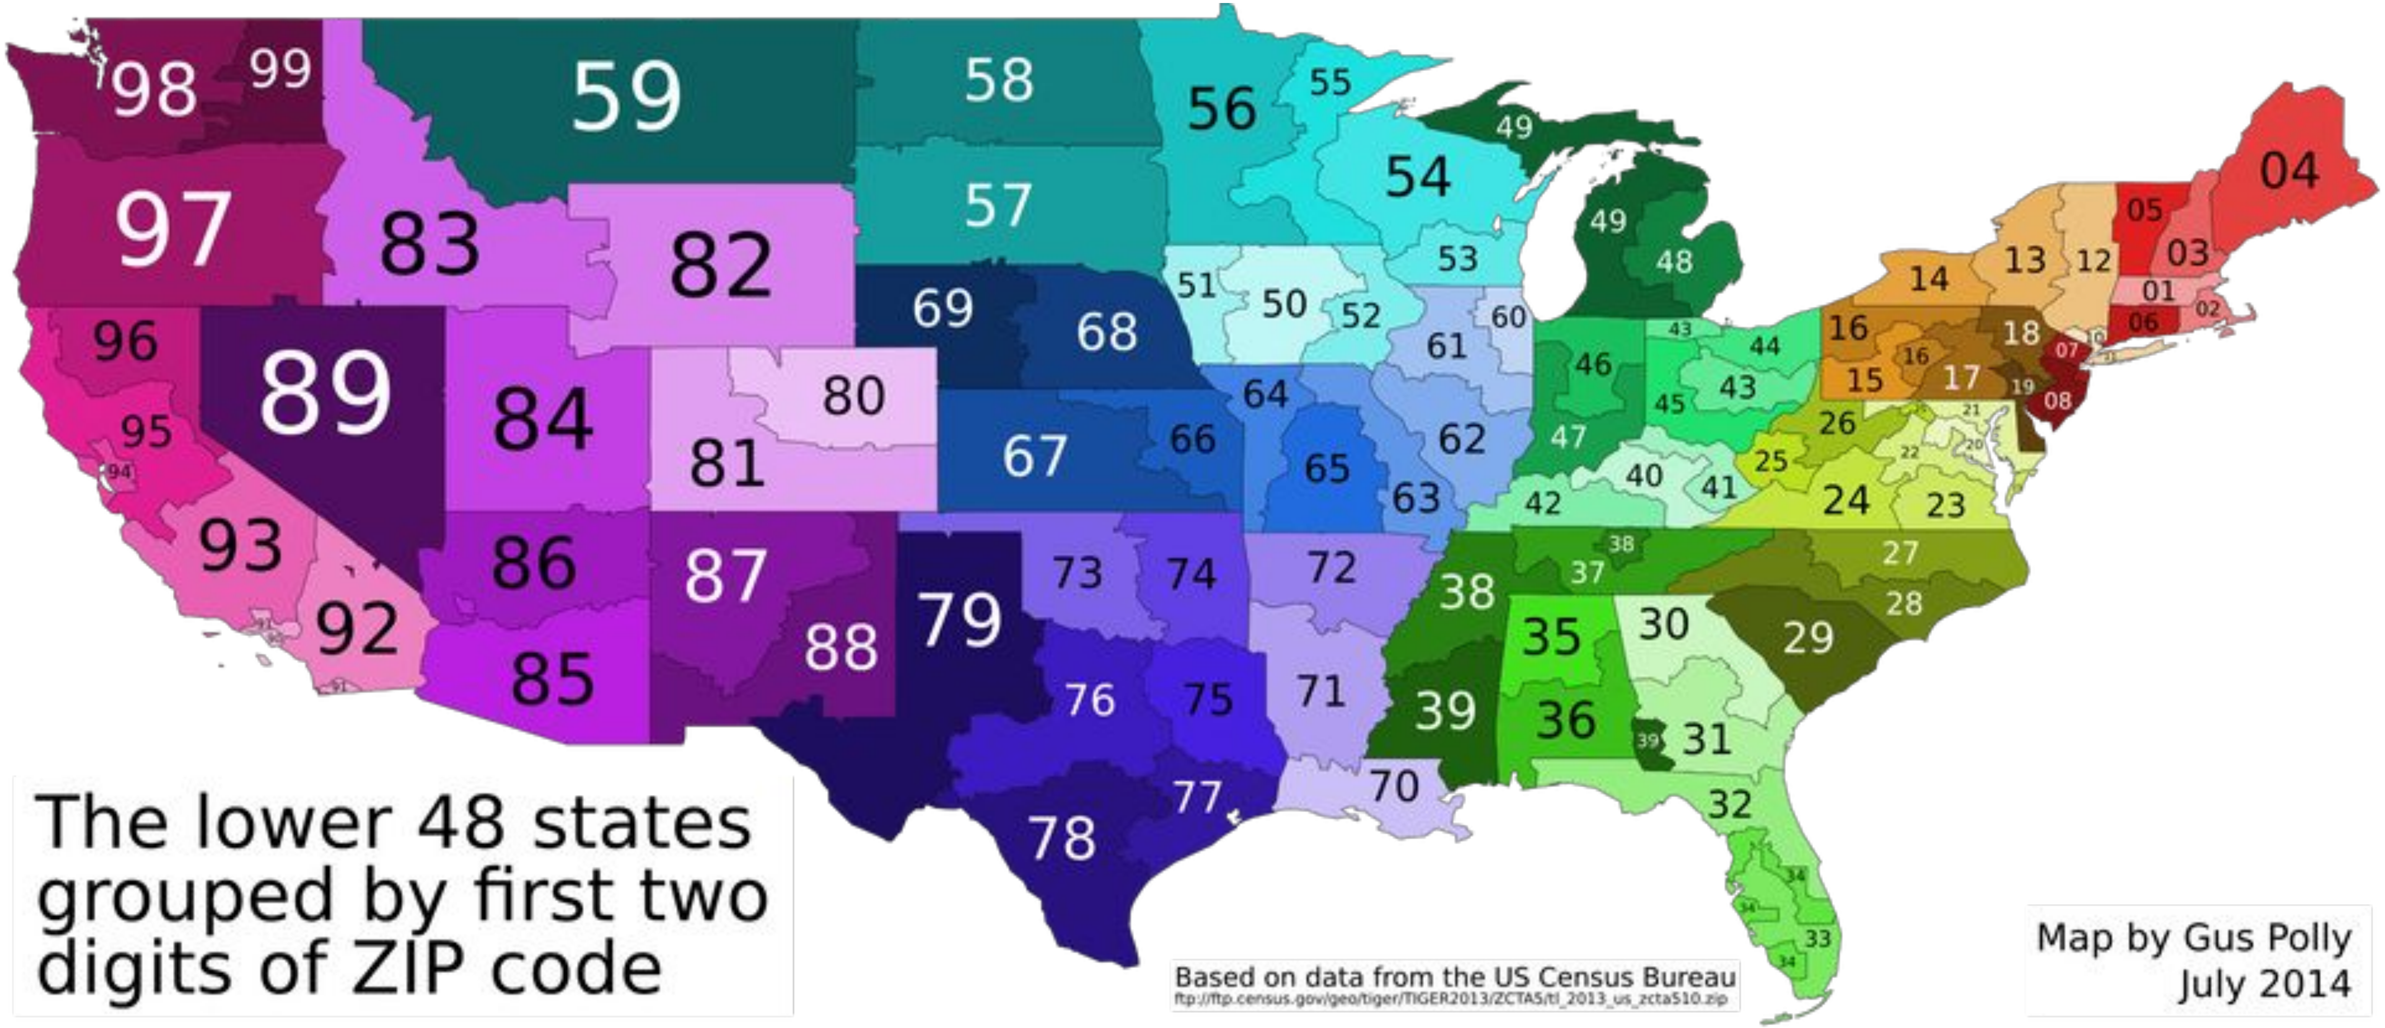

http://i.imgur.com/NJGcg6v.png

Internet structure: a "network of networks"

- Hosts connect to Internet via access Internet Service Providers (ISPs)
  - residential, enterprise (company, university, commercial) 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
- Let's take a stepwise approach to describe current Internet structure

Internet structure: a "network of networks"

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

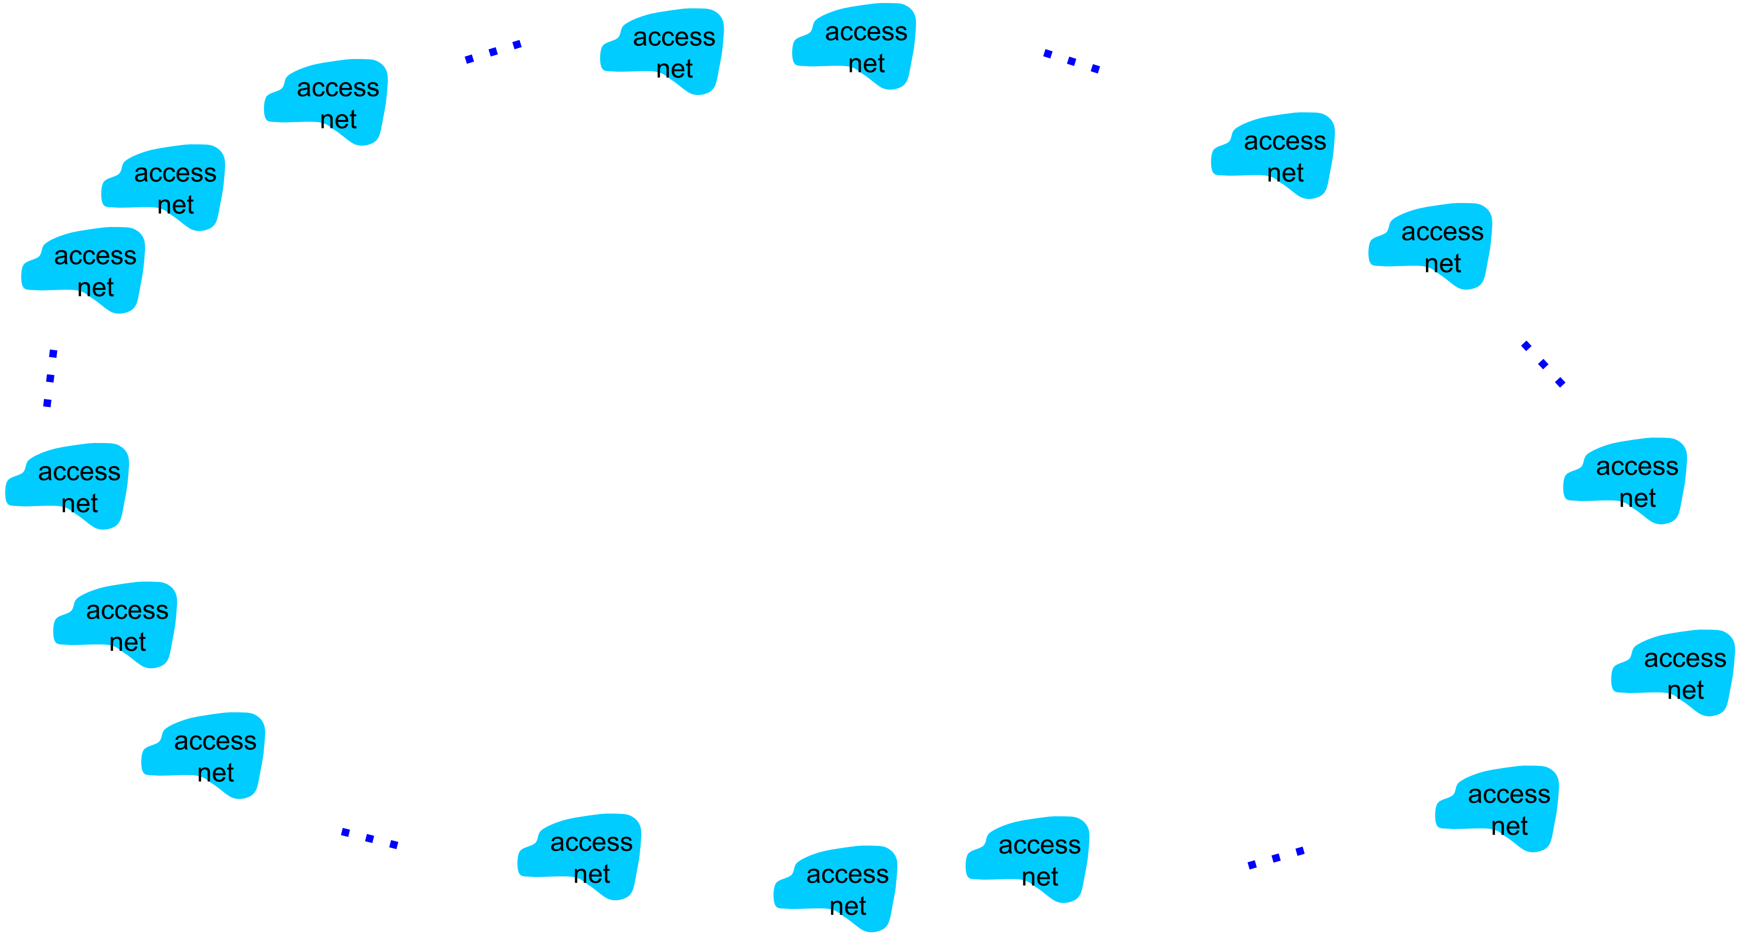

Internet structure: a "network of networks"

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

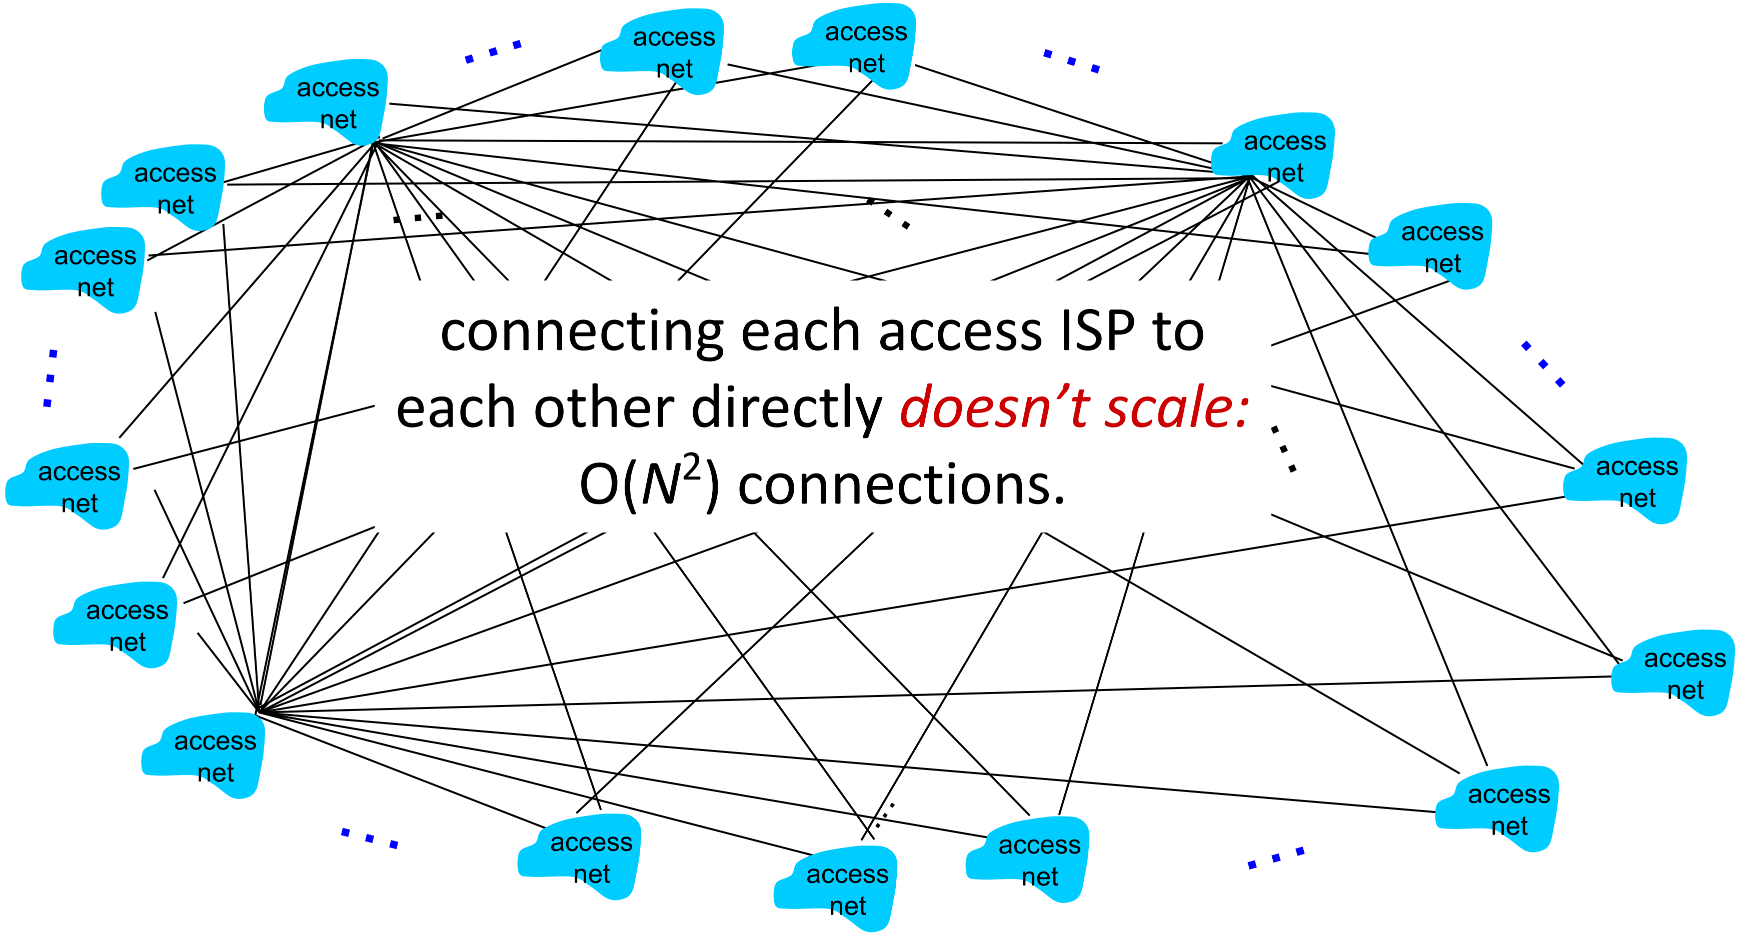

Internet structure: a "network of networks"

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

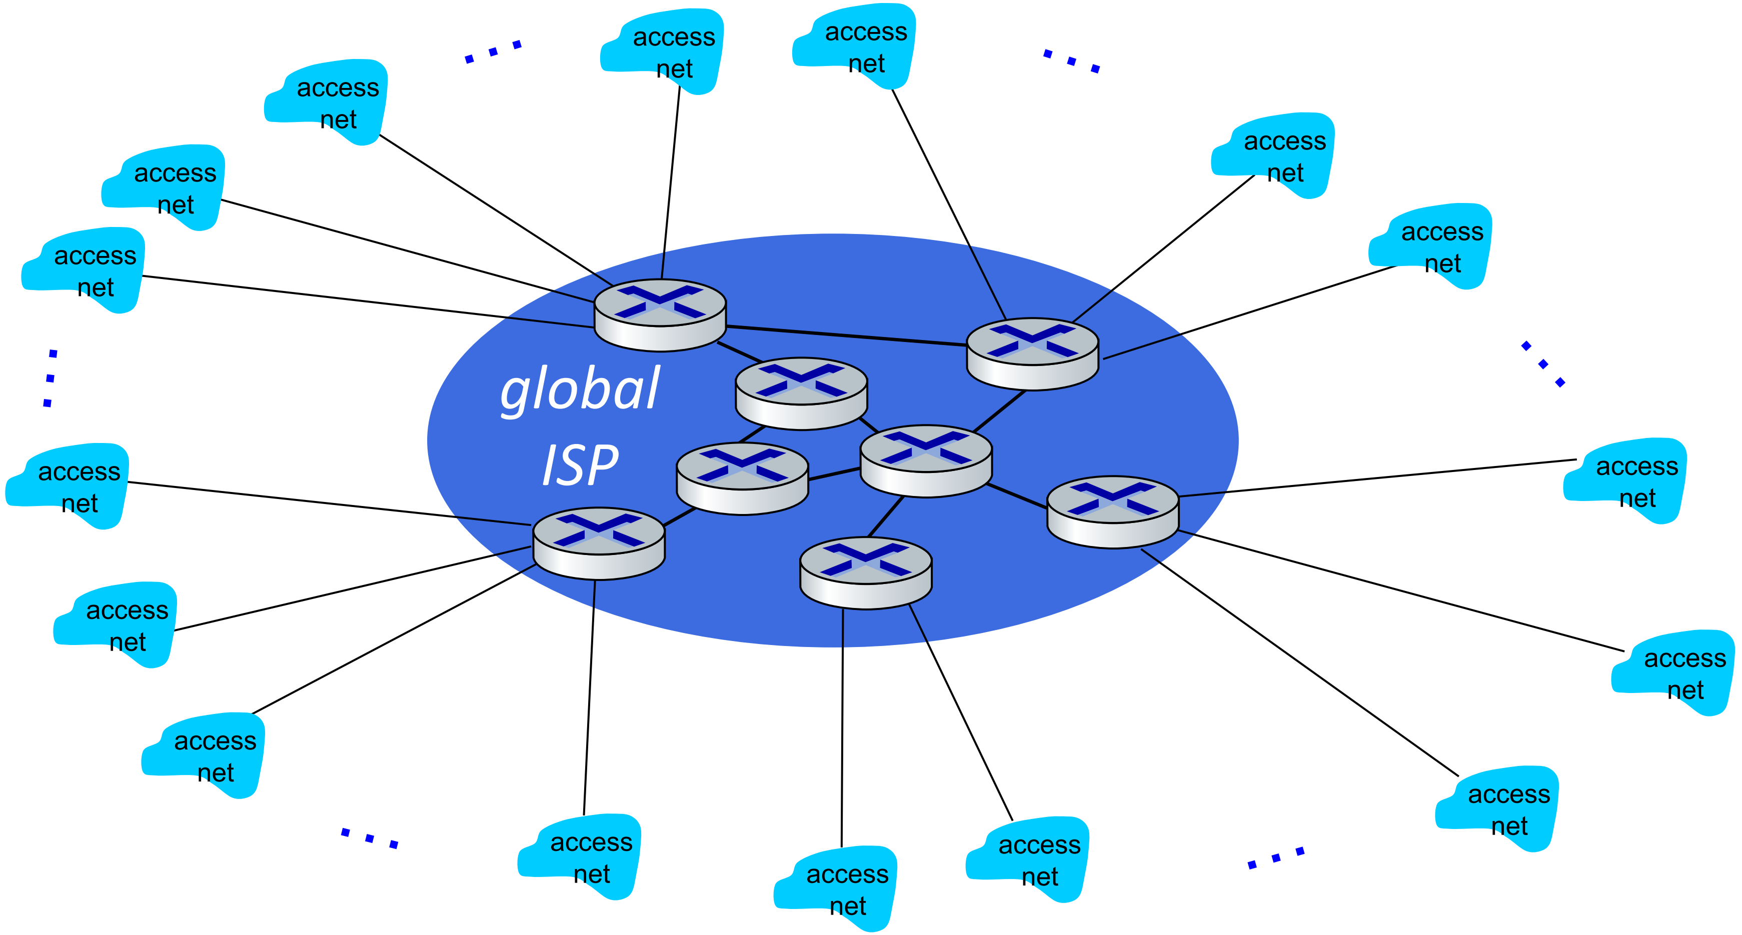

Internet structure: a "network of networks"

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

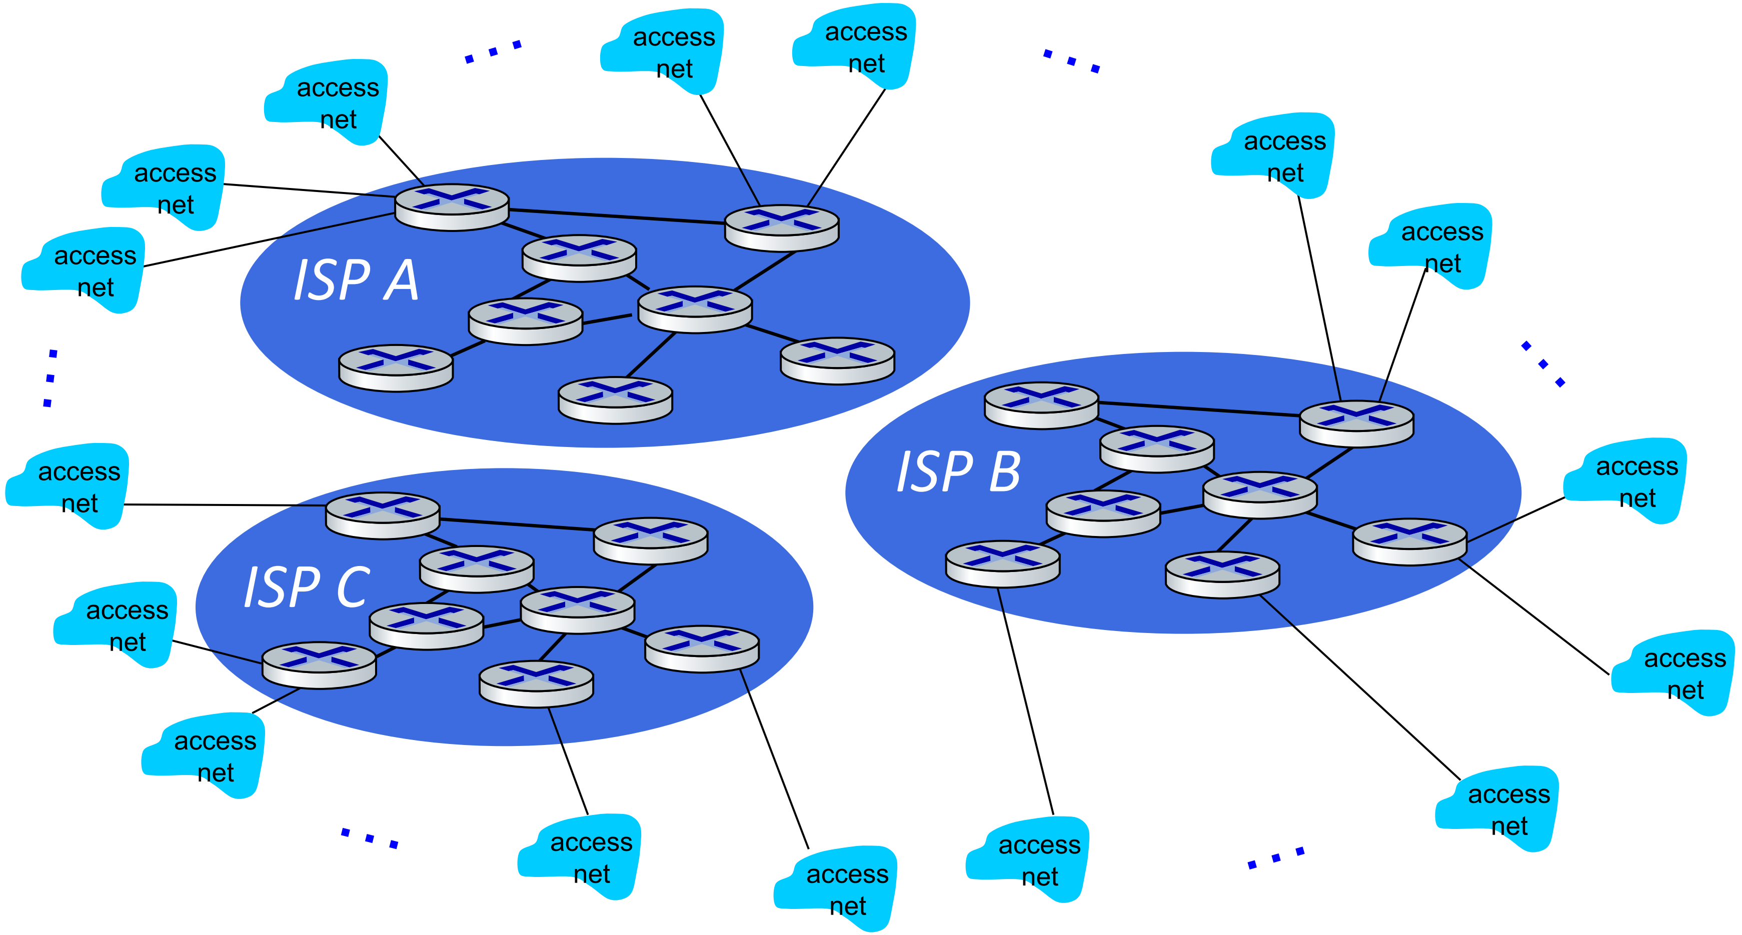

Internet structure: a "network of networks"

But if one global ISP is viable business, there will be competitors …. who will want to be connected

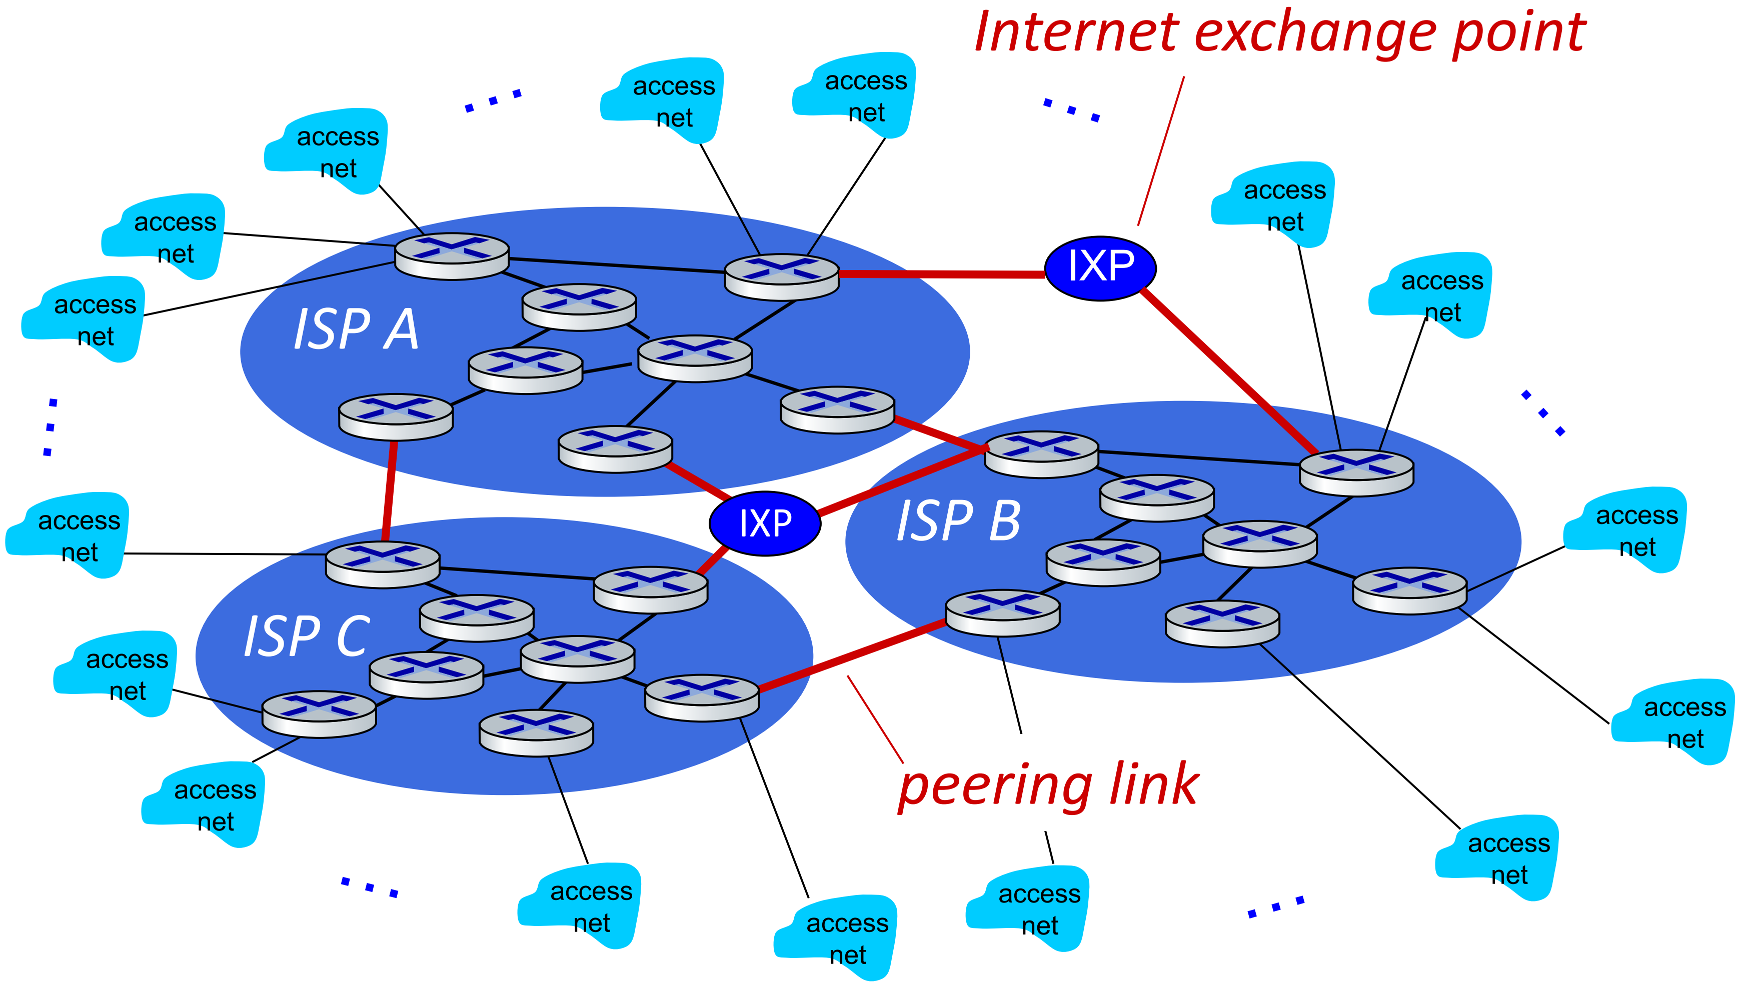

Internet structure: a "network of networks"

… and regional networks may arise to connect access nets to ISPs

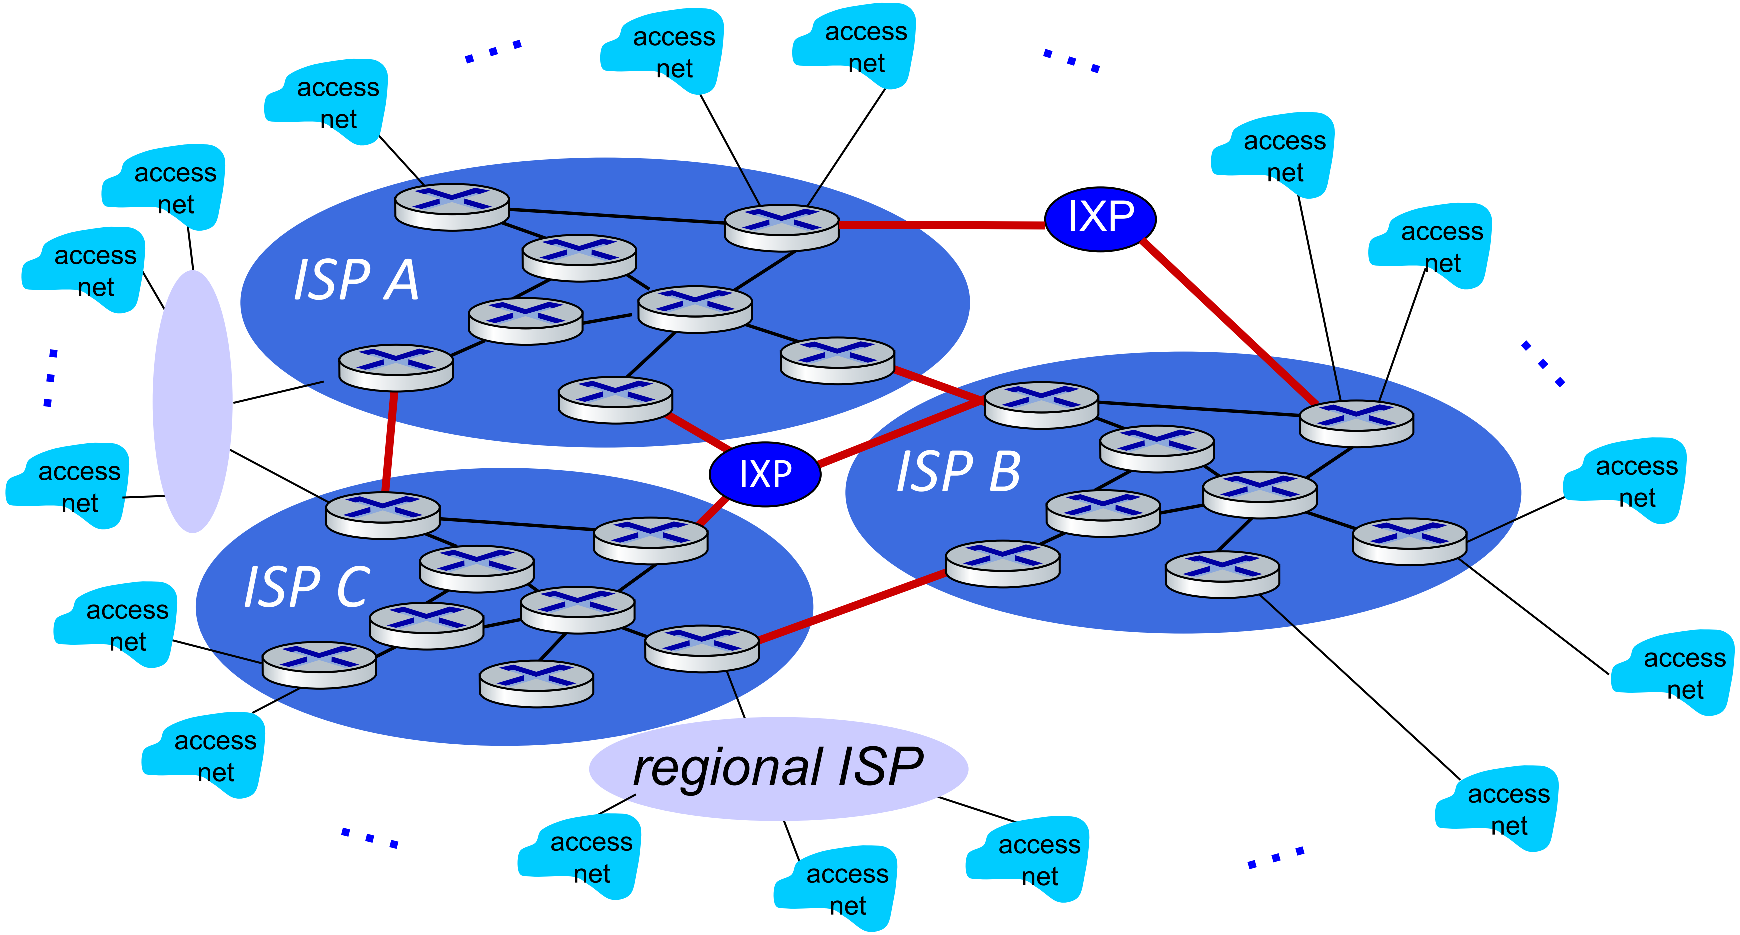

Internet structure: a "network of networks"

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

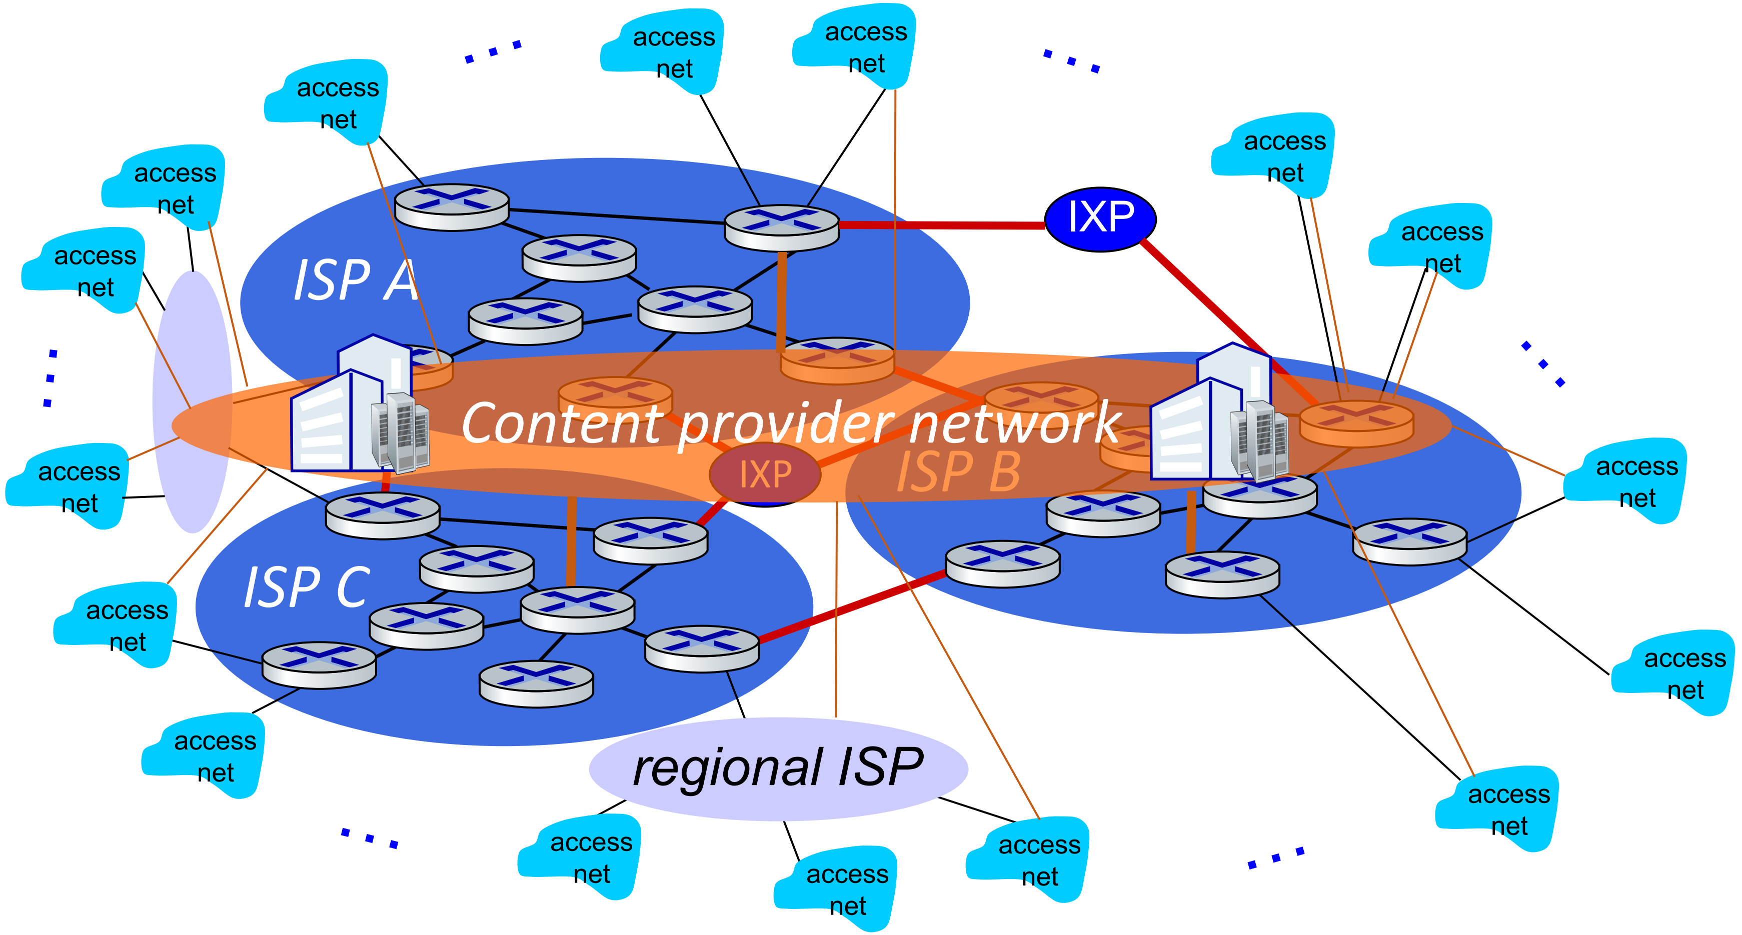

Internet structure: a "network of networks"

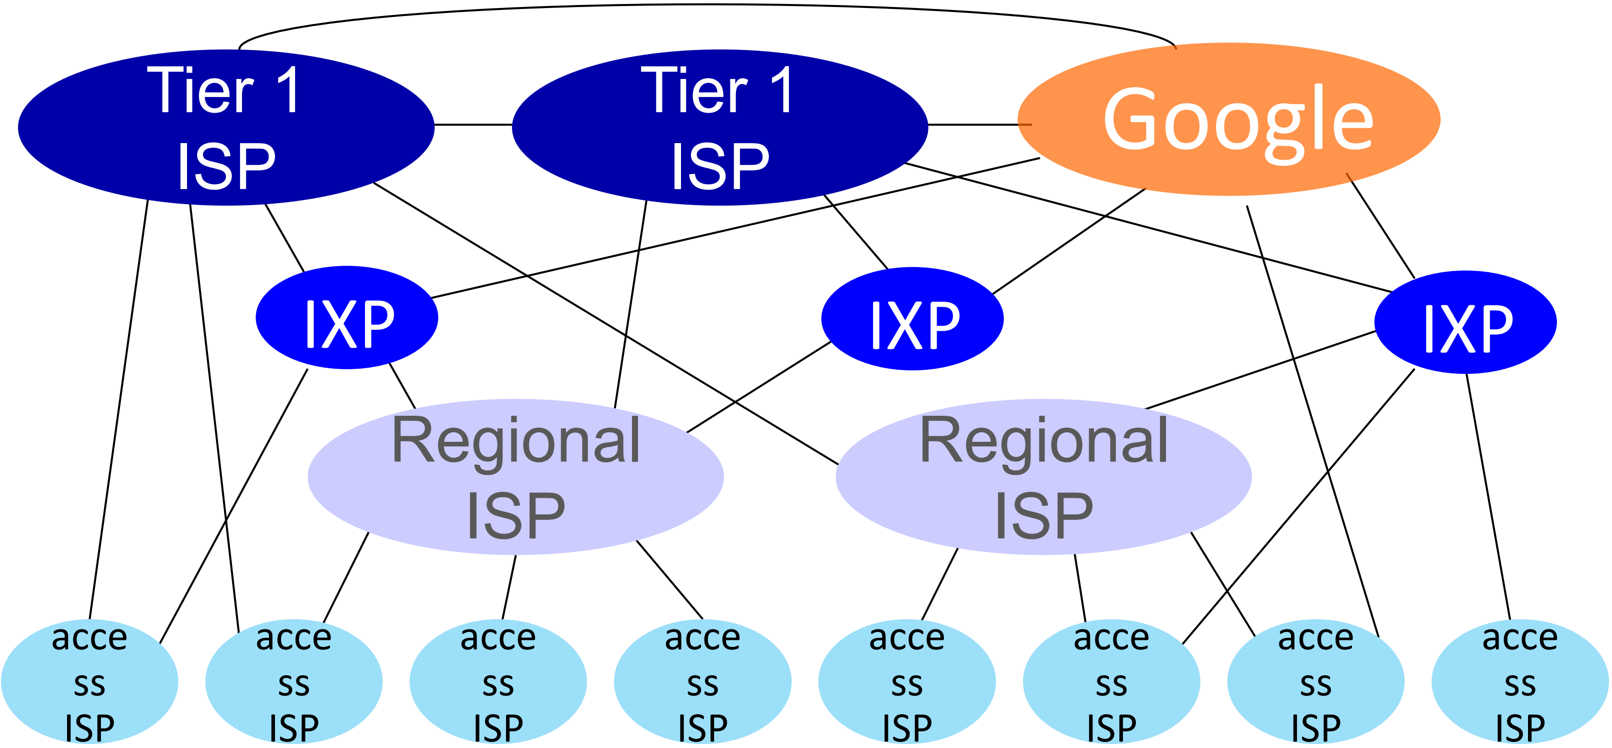

At "center": small # of well-connected large networks

- "tier-1" commercial ISPs (e.g., Level 3, Sprint, AT&T, NTT), national & international coverage
- content provider networks (e.g., Google, Facebook): private network that connects its data centers to Internet, often bypassing tier-1, regional ISPs

Tier-1 ISP Network map: Sprint (2019)

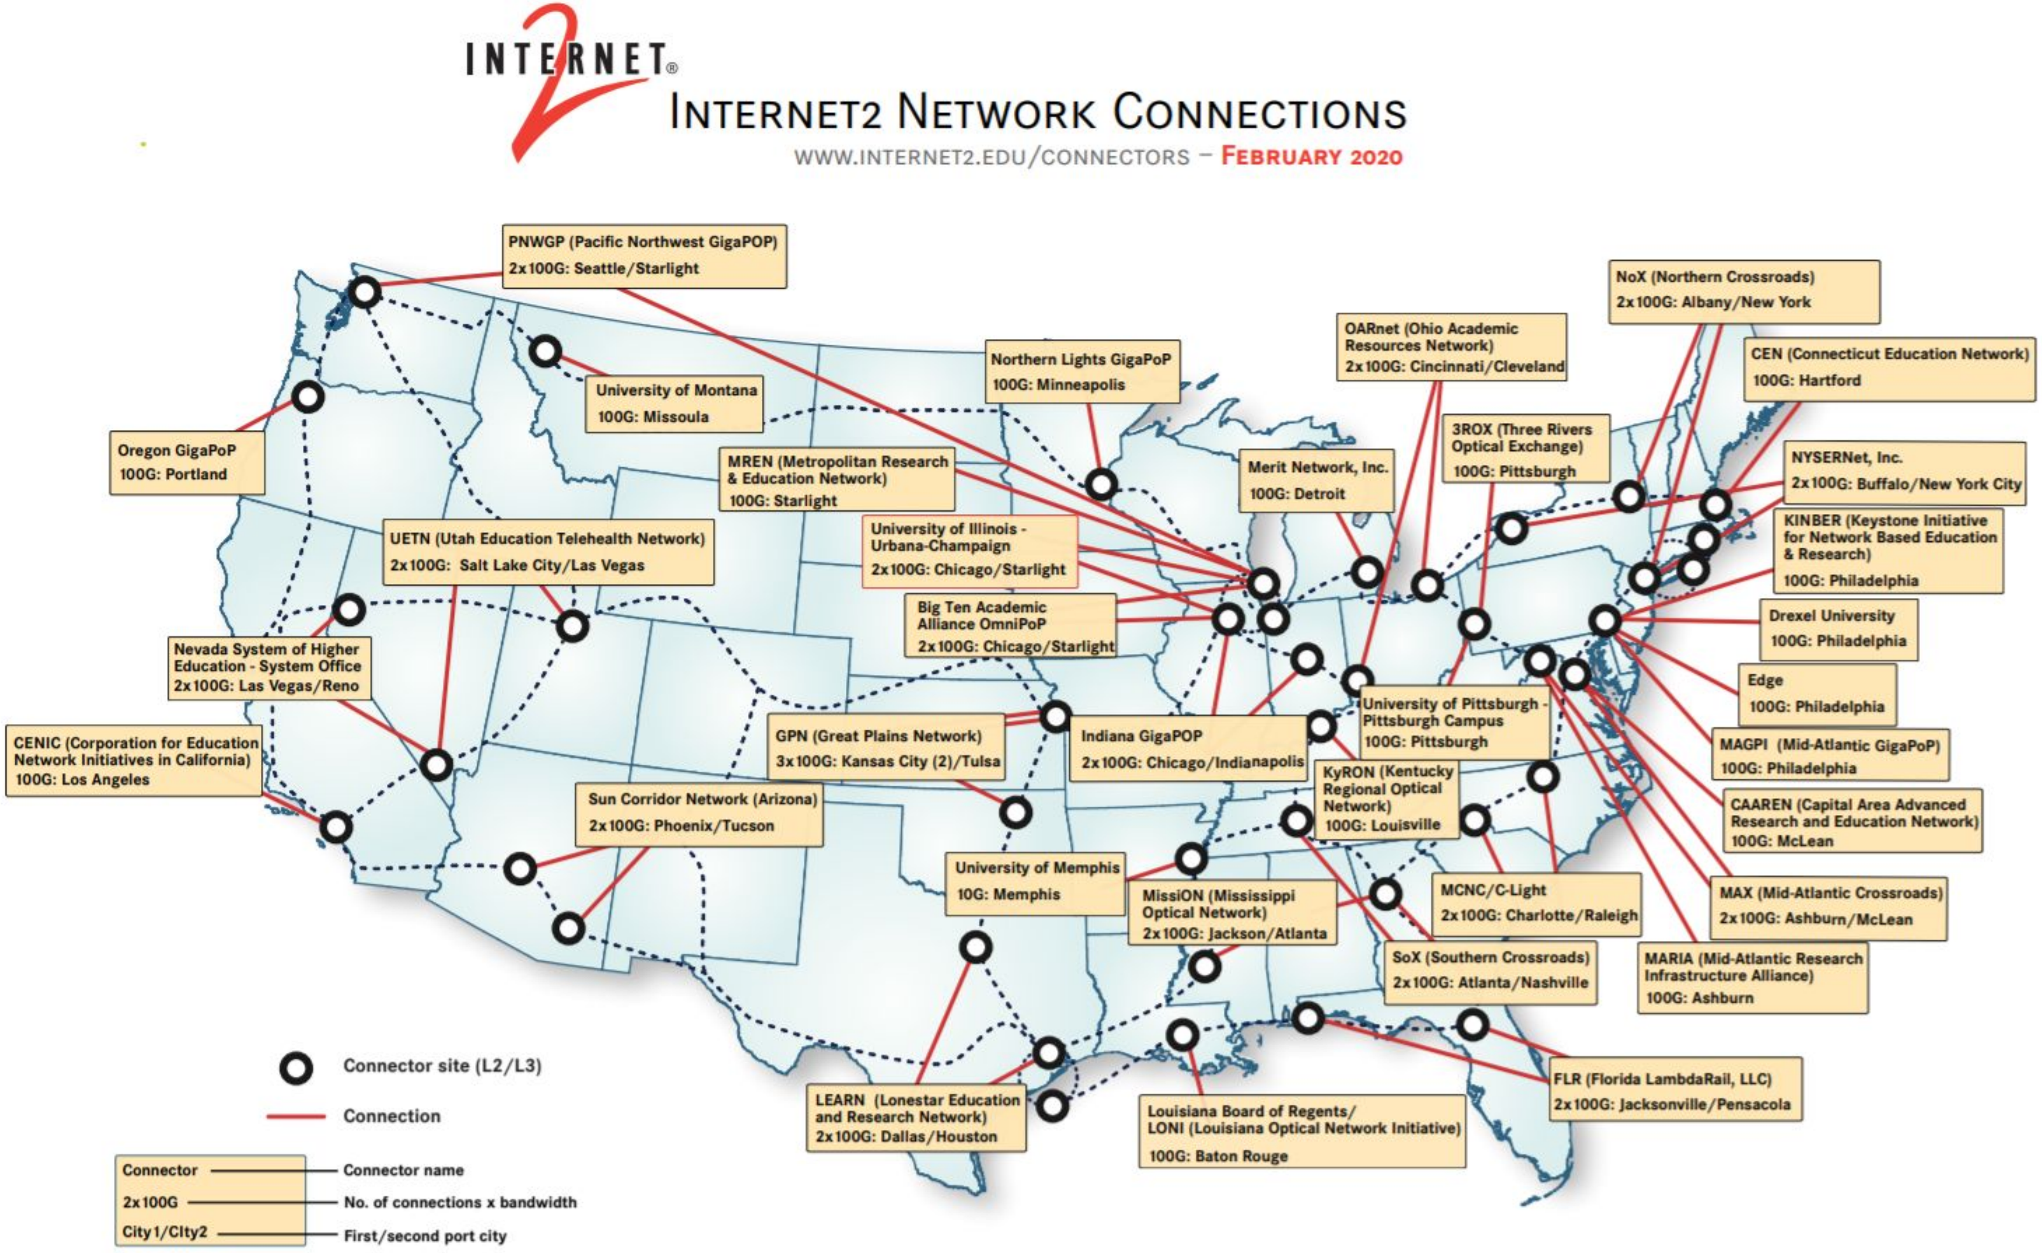

https://www.internet2.edu/media/medialibrary/2020/02/19/ConnectorsMap_202002_XaFyViw.pdf

Part II: On Delay & Throughput

Packet-switching: store-and-forward

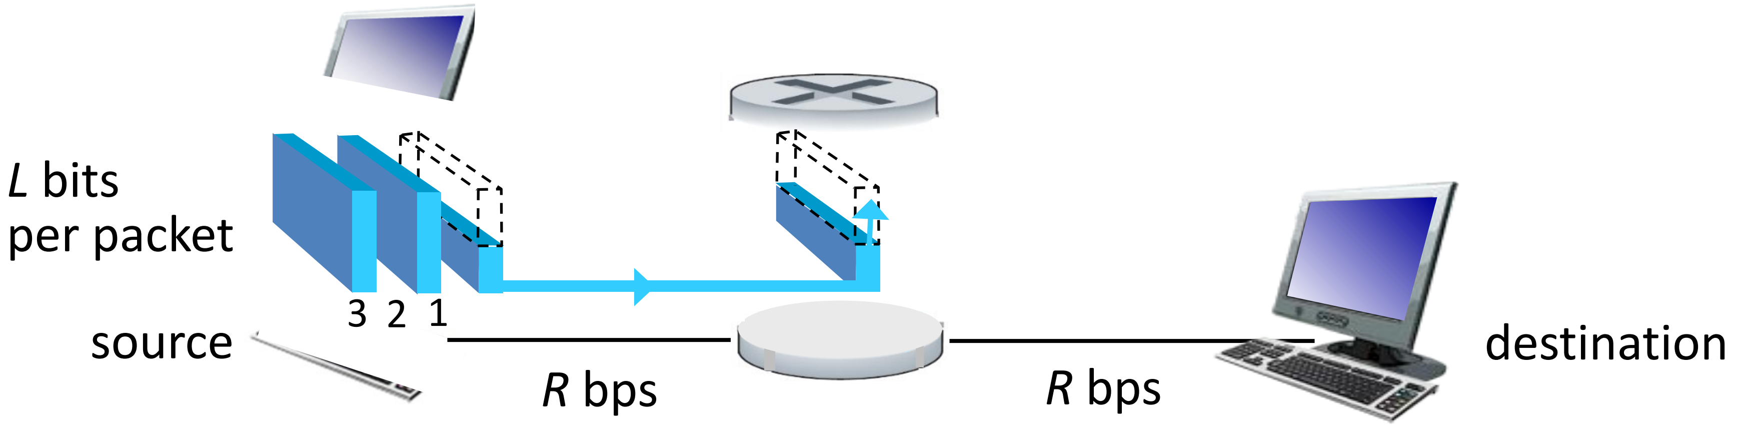

- Transmission delay: takes L/R seconds to transmit (push out) L-bit packet into link at R bps
- *Store and forward: entire* packet must arrive at router before it can be transmitted on next link
- *End-end delay:* $2L/R$ (above), assuming zero propagation delay (more on delay shortly)

One-hop numerical example:

- L = 10 Kbits
- R = 100 Mbps
- one-hop transmission delay = 0.1 msec

Packet-switching: queueing delay, loss

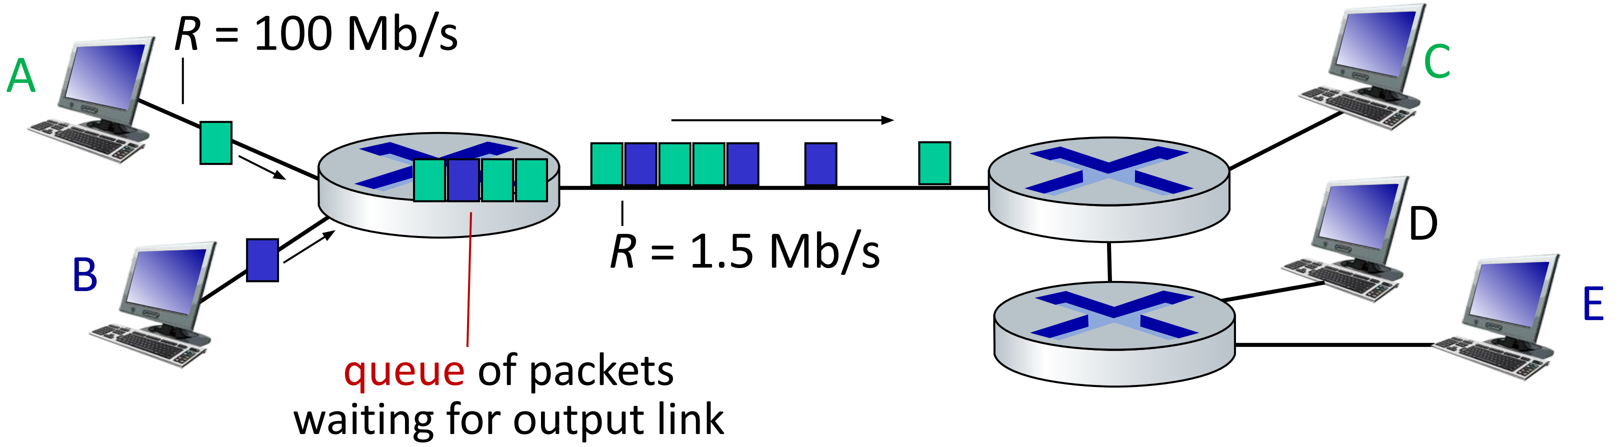

Packet queuing and loss: if arrival rate (in bps) to link exceeds transmission rate (bps) of link for a period of time:

- packets will queue, waiting to be transmitted on output link
- packets can be dropped (lost) if memory (buffer) in router fills up

Packet delay: four sources

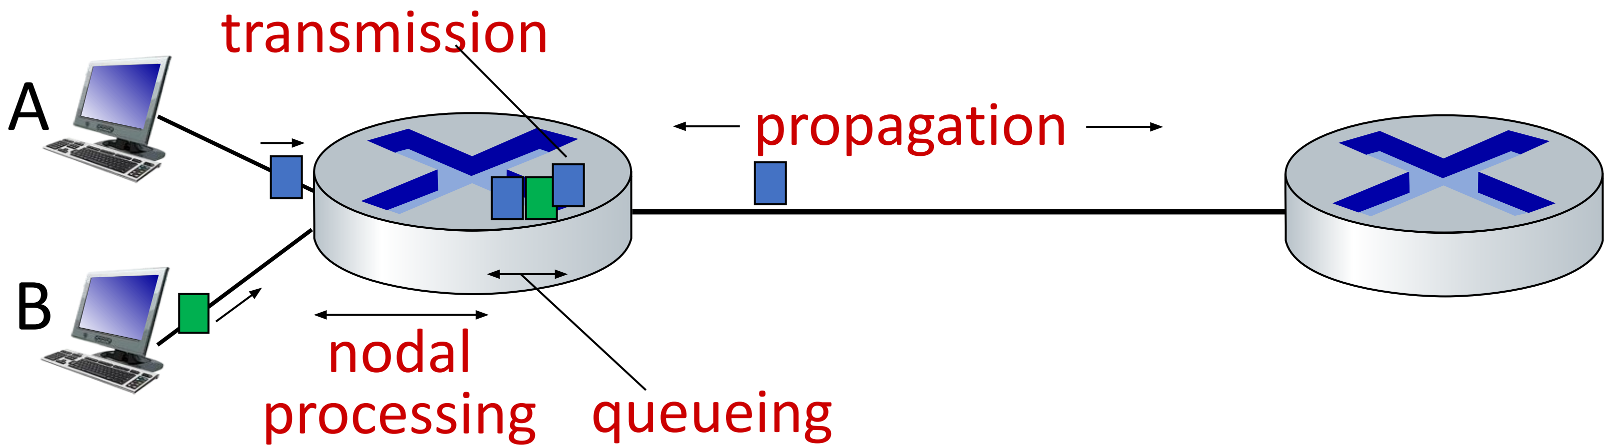

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

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

Packet delay: four sources

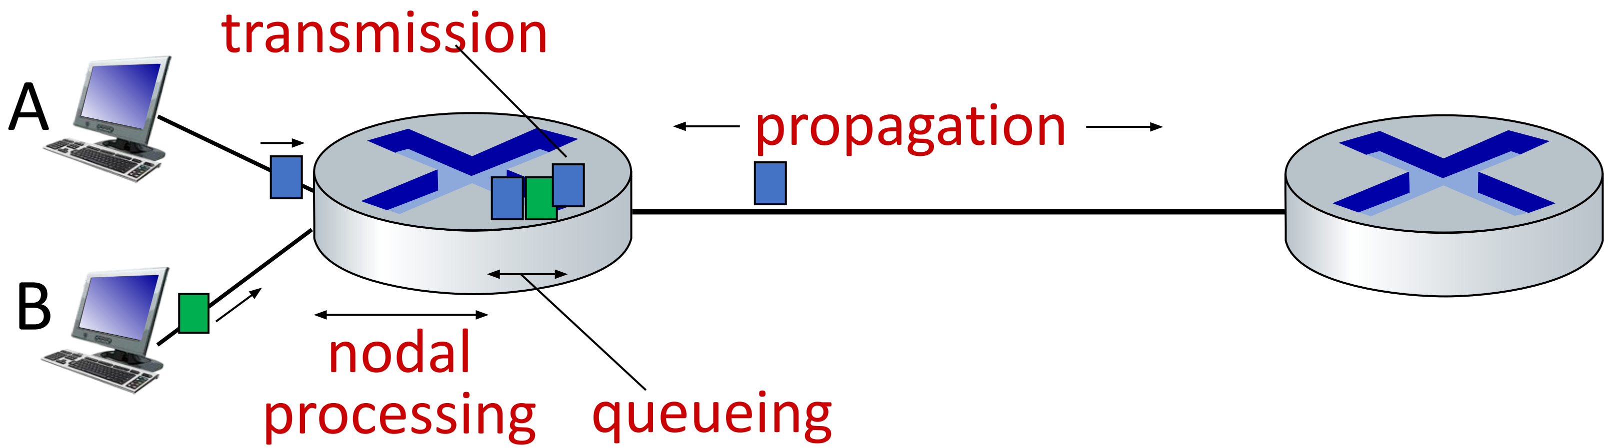

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

d_{trans}: transmission delay:

- L: packet length (bits)
- R: link *transmission rate (bps)*
- $d_{trans} = L/R$

d_{prop}: propagation delay:

- d: length of physical link
- s: propagation speed (~2×10^8 m/sec)
- $d_{prop} = d/s$

d_{trans} and d_{prop} *very* different

* Check out the online interactive exercises: http://gaia.cs.umass.edu/kurose_ross

Packet queueing delay (revisited)

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

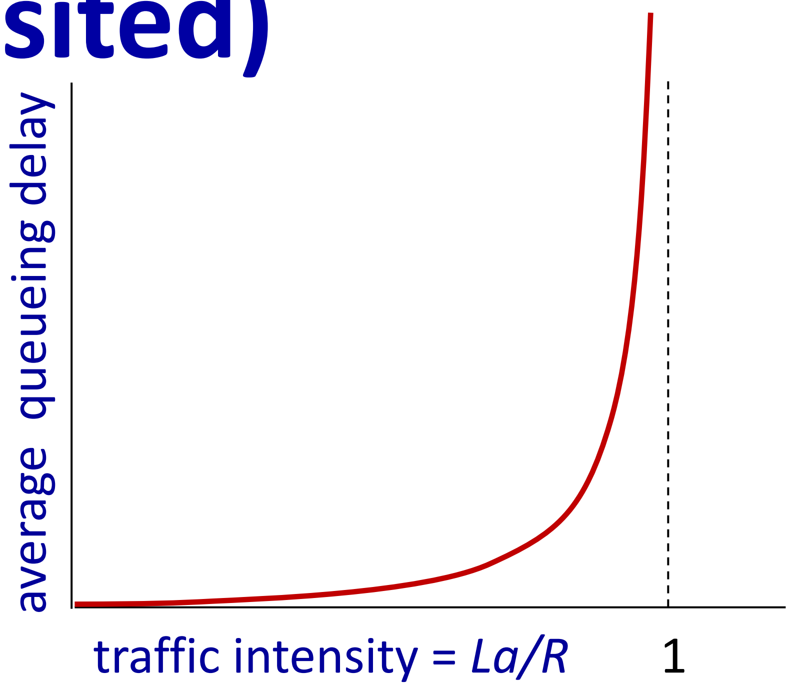

- *La/R* ~ 0: avg. queueing delay small
- *La/R* -> 1: avg. queueing delay large
- *La/R* > 1: more "work" arriving is more than can be serviced - average delay infinite!

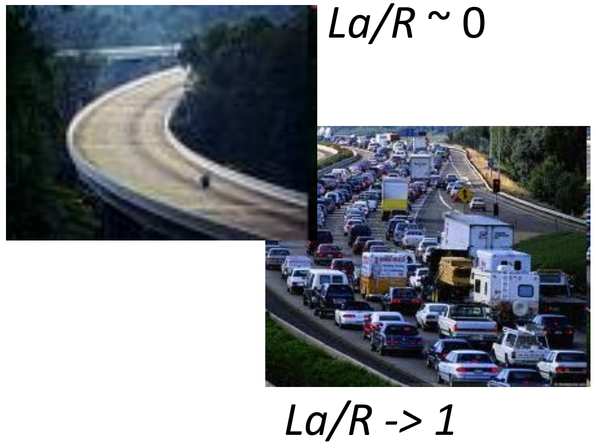

Throughput

- *throughput:* rate (bits/time unit) at which bits are being sent from sender to receiver
  - *instantaneous:* rate at given point in time
  - *average:* rate over longer period of time

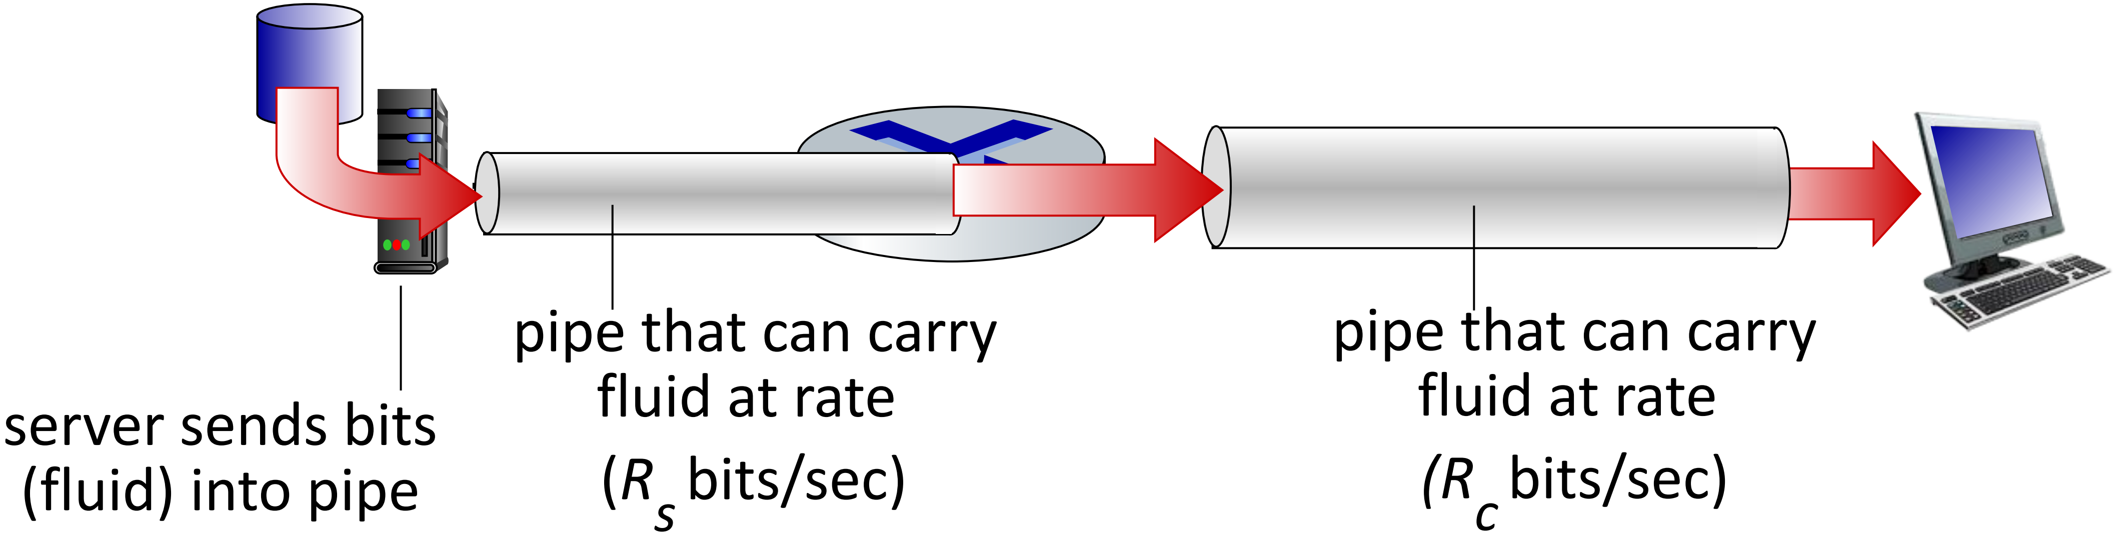

Throughput

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

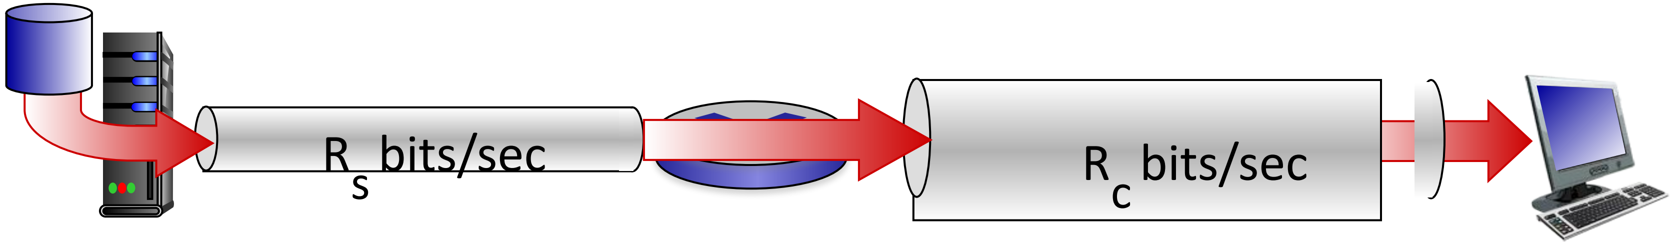

$R_s > R_c$ What is average end-end throughput?

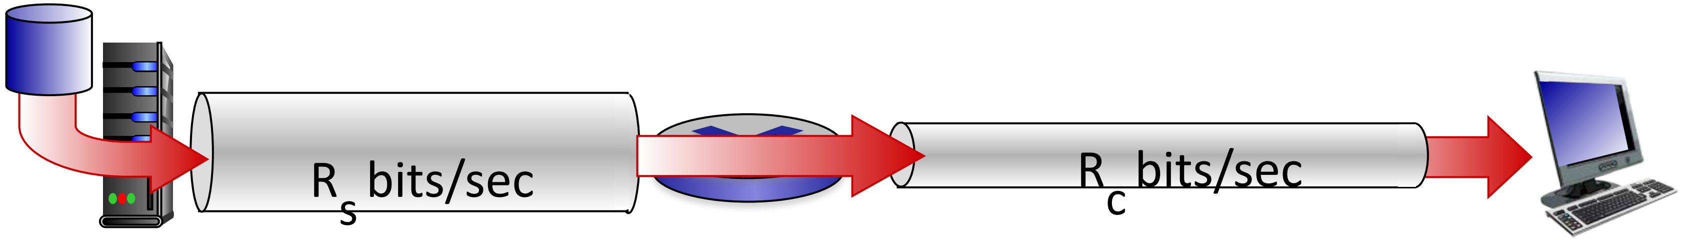

bottleneck link

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

Throughput: network scenario

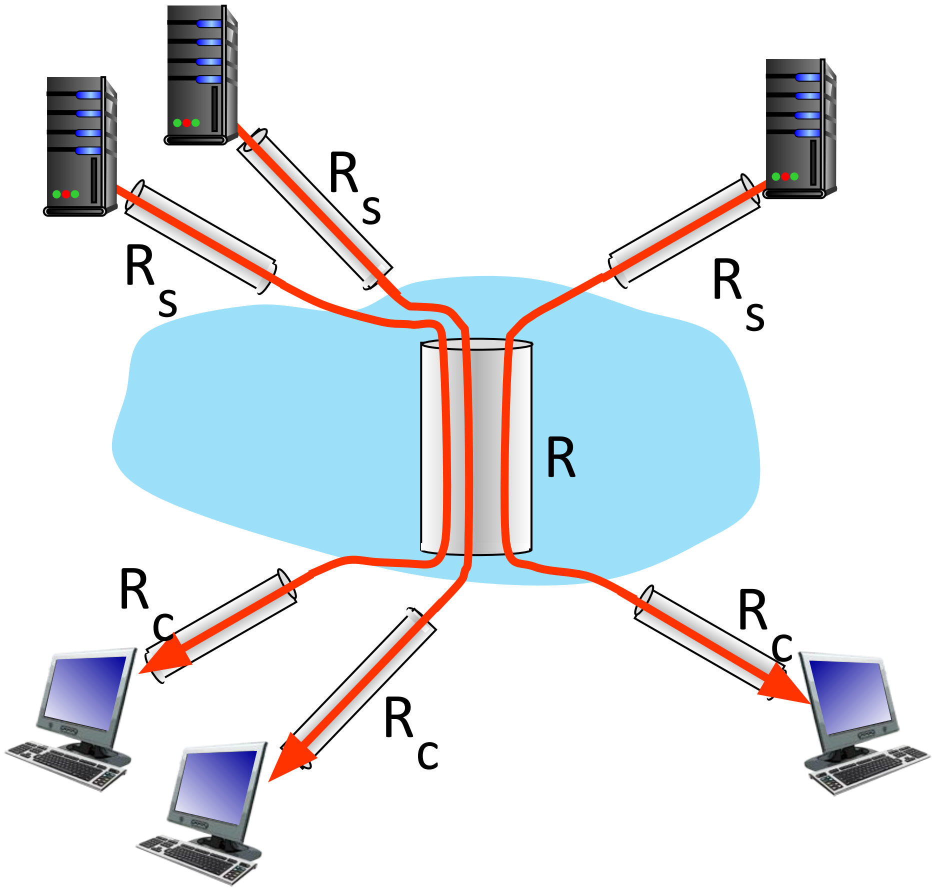

10 connections (fairly) share backbone bottleneck link *R* bits/sec

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

* Check out the online interactive exercises for more examples: http://gaia.cs.umass.edu/kurose_ross/

Part III: On Layering

Protocol “layers” and reference models

Networks are complex, with 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, updating of system
  - change in layer's service *implementation*: transparent to rest of system
- layering considered harmful?
- layering in other complex systems?

Internet protocol stack

- *application:* supporting network applications
  - IMAP, SMTP, HTTP
- *transport:* process-process data transfer
  - TCP, UDP
- *network:* routing of datagrams from source to destination
  - IP, routing protocols
- *link:* data transfer between neighboring network elements
  - Ethernet, 802.11 (WiFi), PPP
- *physical:* bits "on the wire"

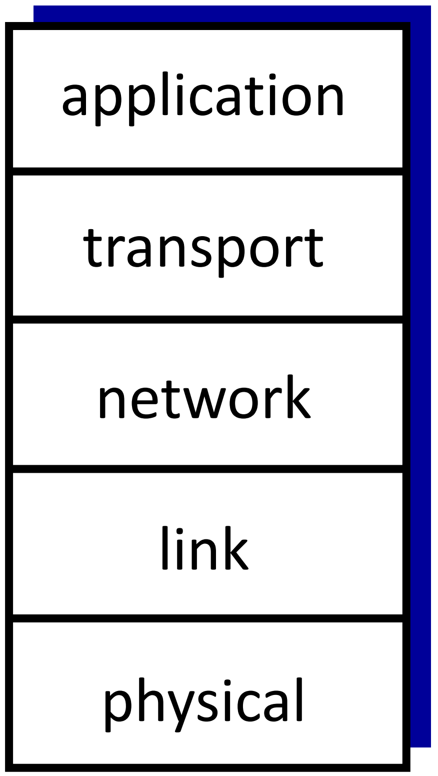

Historical Sidenote

- Before the Internet protocol stack was established as a de facto standard, there were attempts at designing network protocols by committee, the so-called (Open System Interconnect) OSI/ISO reference model
- Never gained traction, but numbering “Layer 3” etc. survived

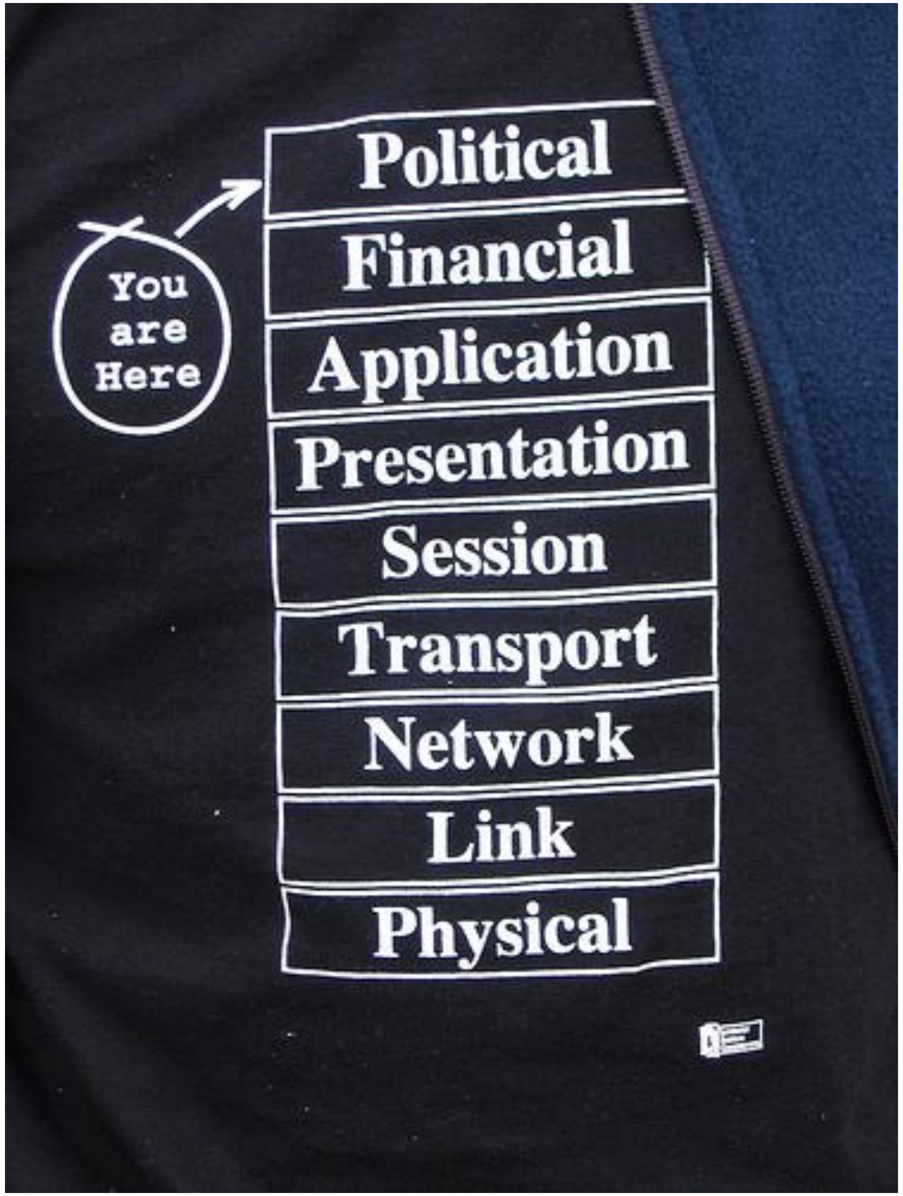

Services vs Protocols

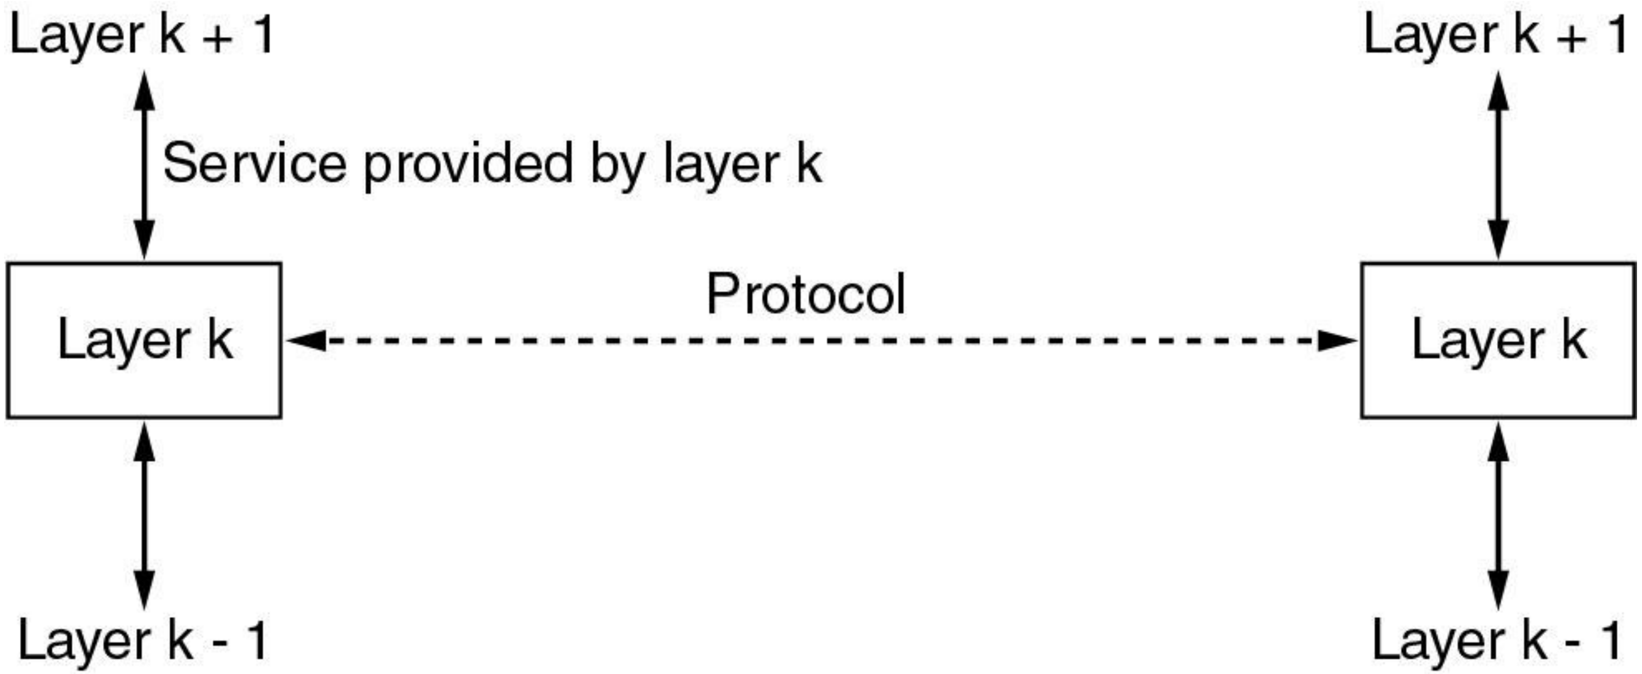

(horizontal component)

- Layer *k* may interact with peer layer *k* only via protocols
  Source: Tanenbaum

Encapsulation

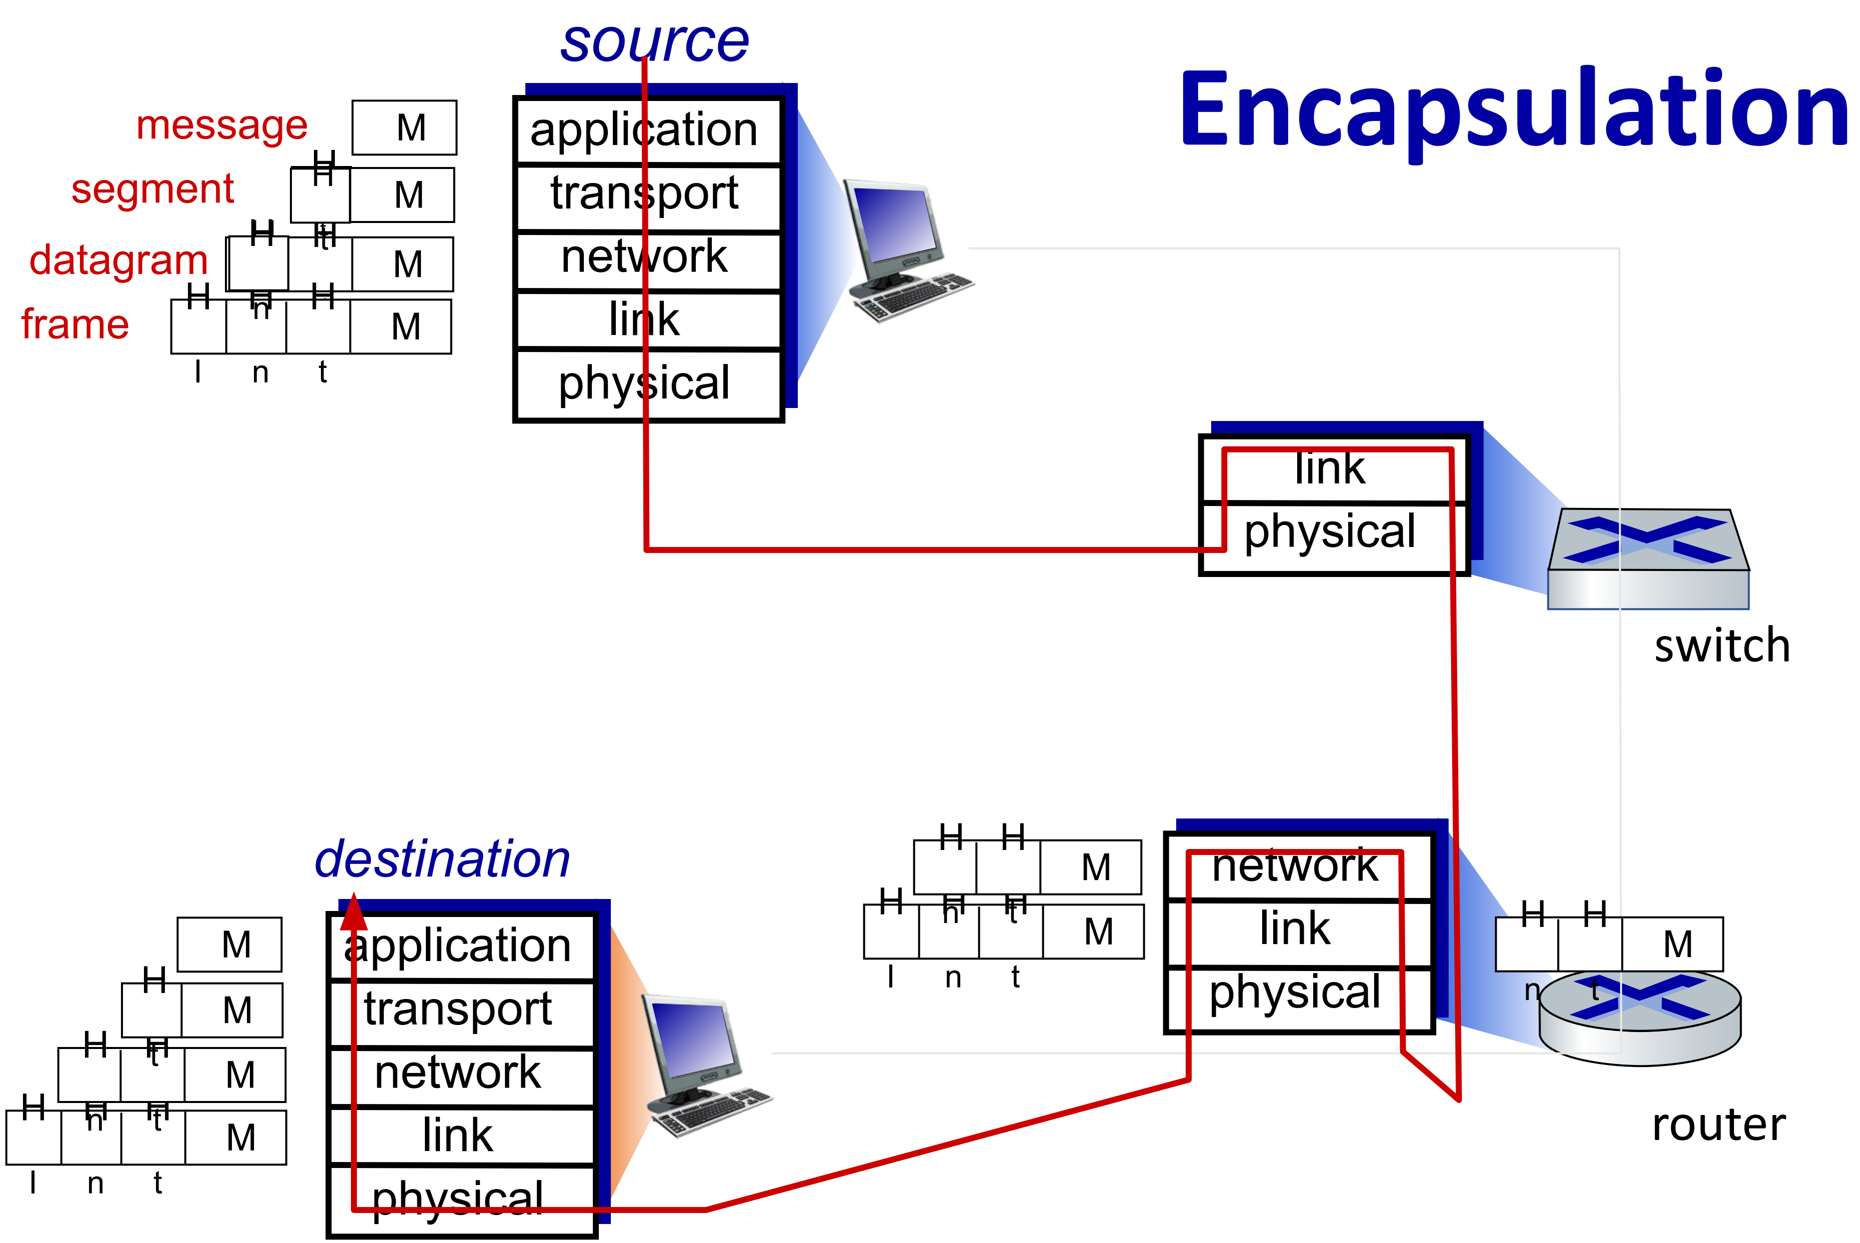

TCP/IP Hourglass View

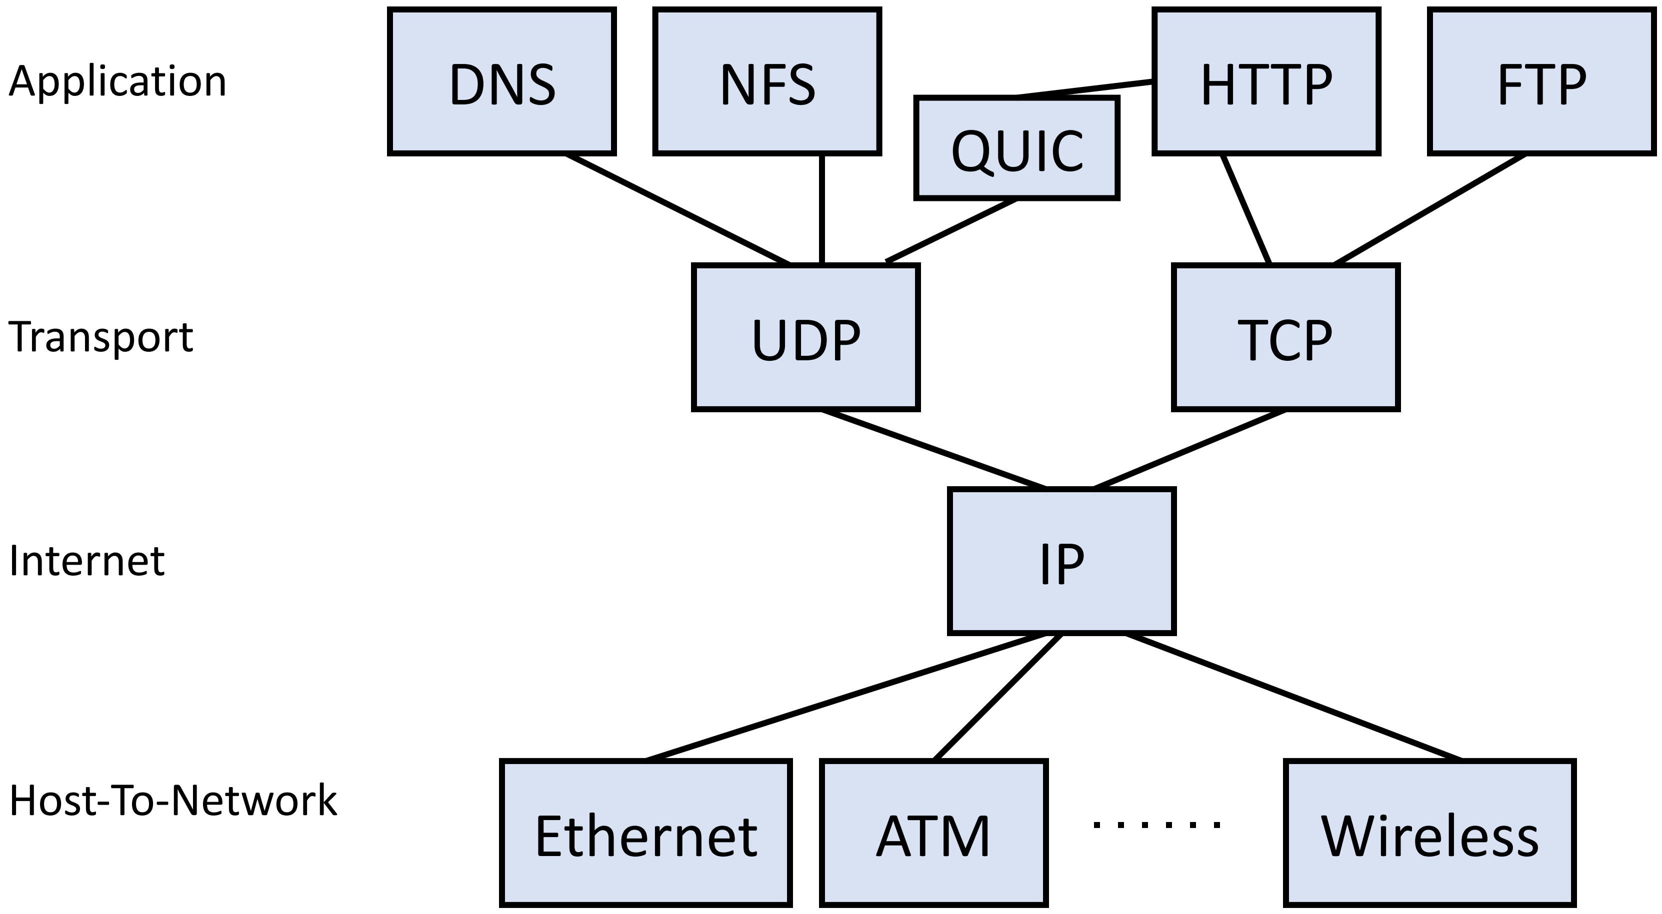

Typical Implementation

may cross multiple boundaries!