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.

Thanks and enjoy! JFK/KWR
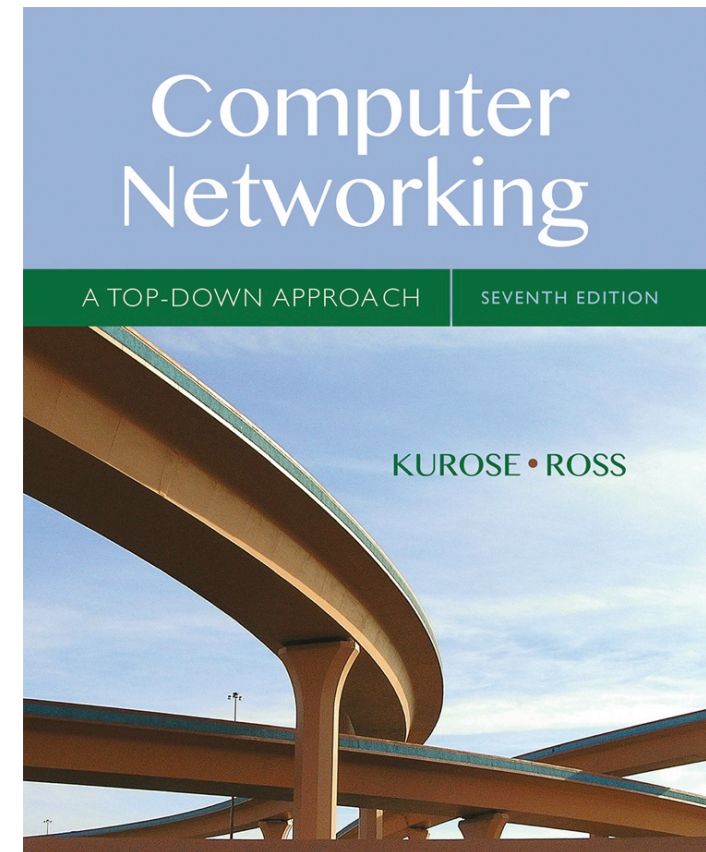
Computer Networking: A Top Down Approach

7th edition
Jim Kurose, Keith Ross
Pearson/Addison Wesley
April 2016

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; hosts, access net, physical media
- network core: packet/circuit switching, Internet structure
- performance: loss, delay, throughput
- security
- protocol layers, service models
- history

Chapter 1: roadmap

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

PC

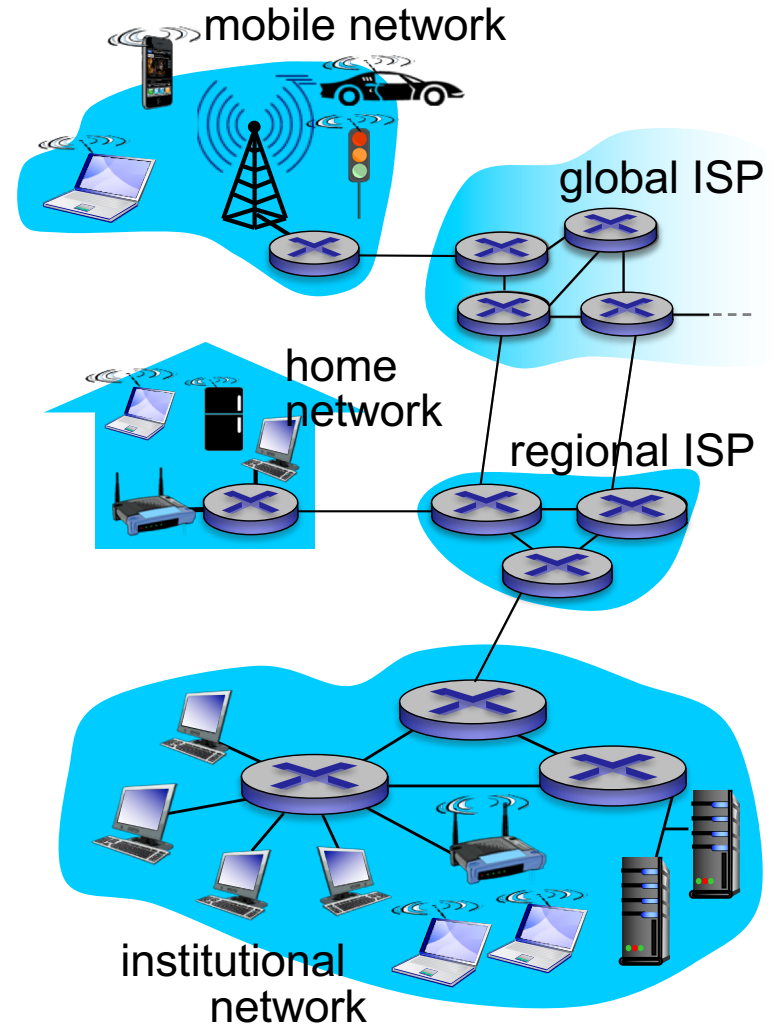server

wireless laptop

smartphone

- billions of connected computing devices:
  - *hosts* = *end systems*
  - running *network apps*
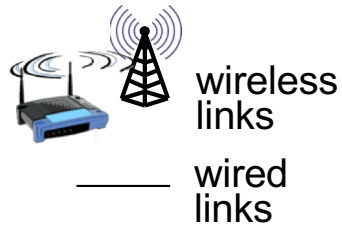
- *communication links*
  - fiber, copper, radio, satellite
 - transmission rate: *bandwidth*

wireless links

wired links

router

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

mobile network

global ISP

home network

regional ISP

institutional network

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

- *Internet:* "network of networks"
  - Interconnected ISPs
- *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

mobile network

global ISP

home network

regional ISP

institutional network
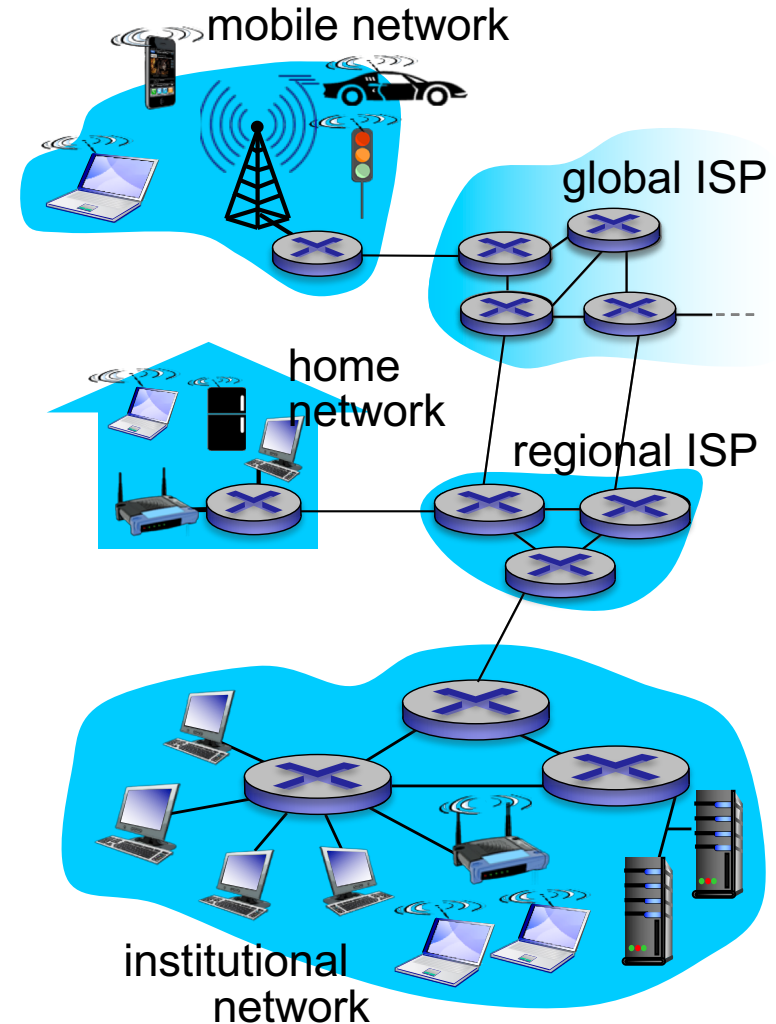
What's the Internet: a service view

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

mobile network

global ISP

home network

regional ISP

institutional network
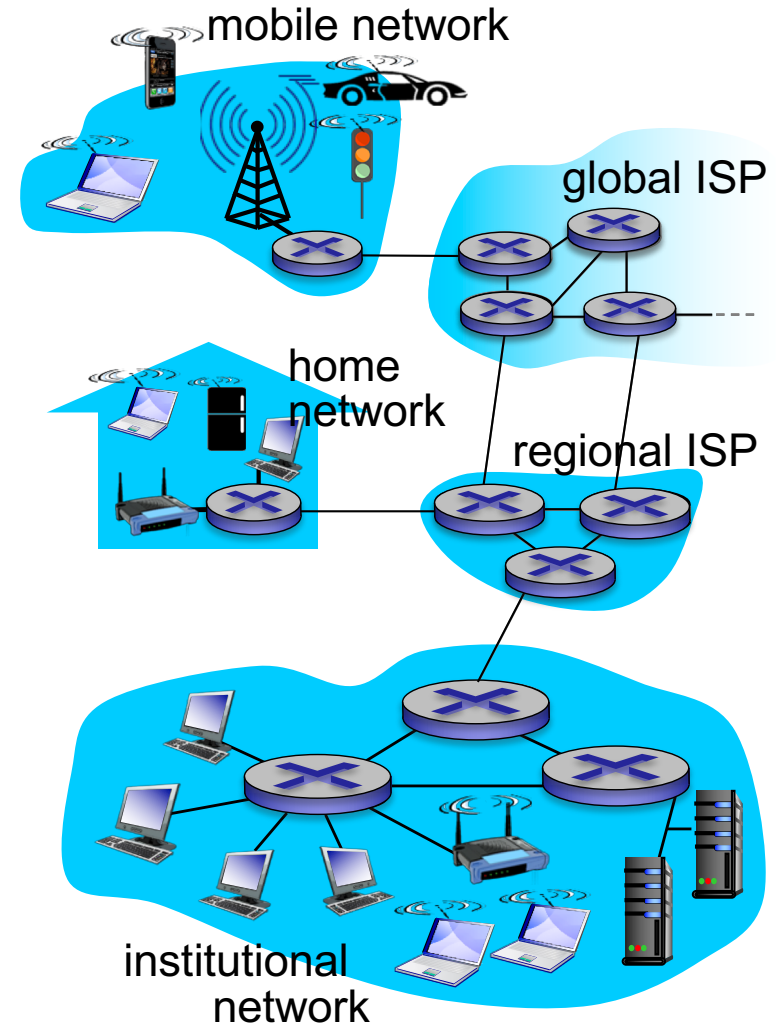
What's a protocol?

human protocols:

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

… specific messages sent

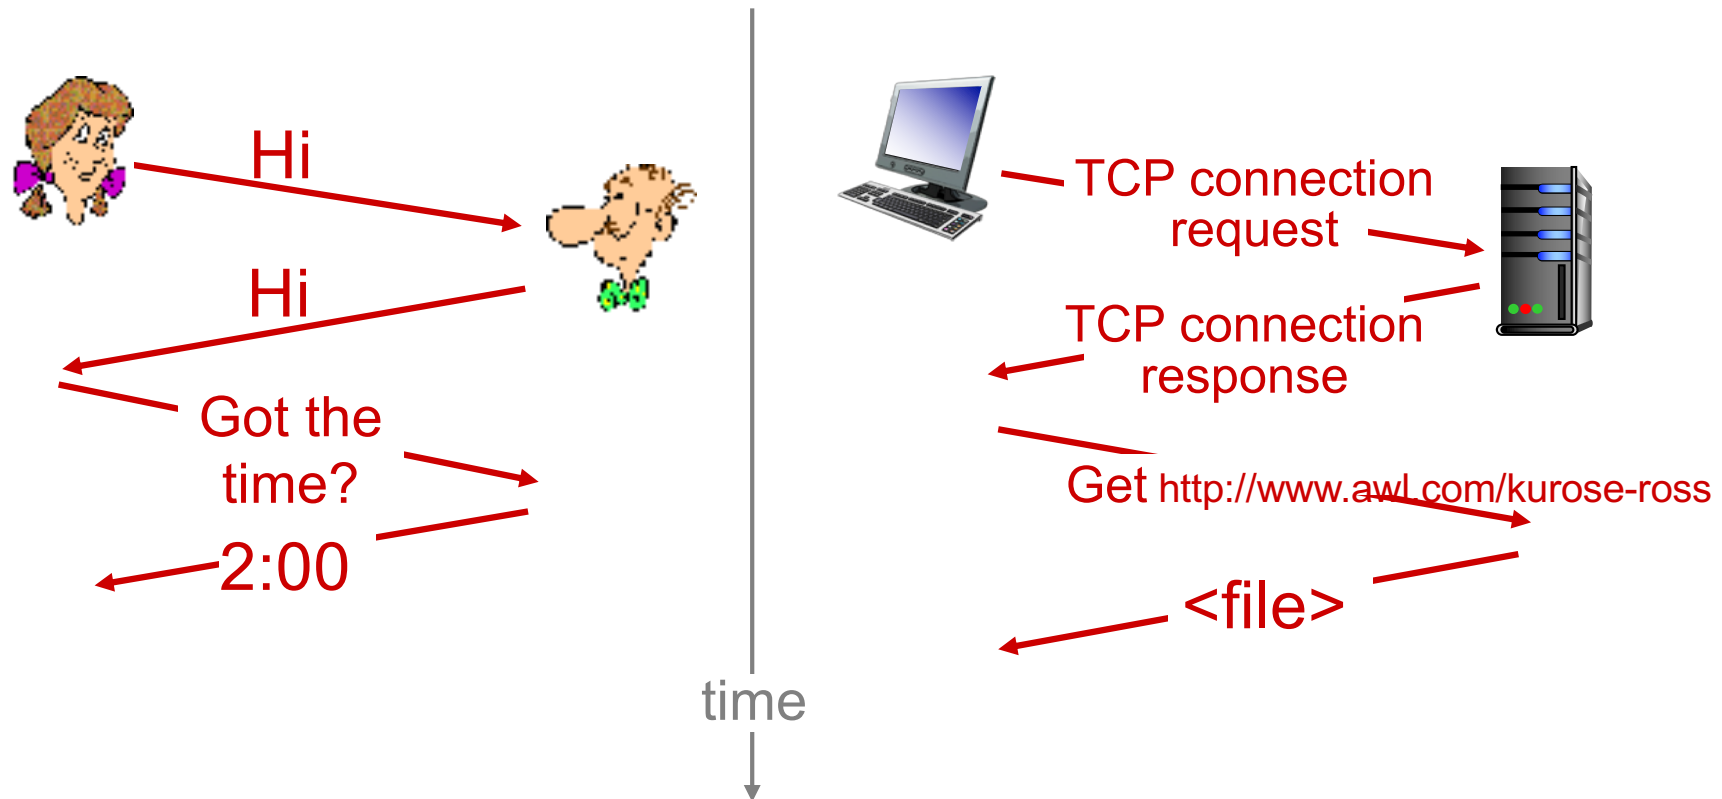… specific actions taken when messages received, or other events

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?

a human protocol and a computer network protocol:

Human Protocol	Computer Network Protocol
Hi →	TCP connection request →
← Hi	← TCP connection response
Got the time? →	Get http://www.awl.com/kurose-ross →
← 2:00	← <file>

time

Q: other human protocols?

Chapter 1: roadmap

A closer look at network structure:

- *network edge:*
  - hosts: clients and servers
  - servers often in data centers

- *access networks, physical media:* wired, wireless communication links

- *network core:*
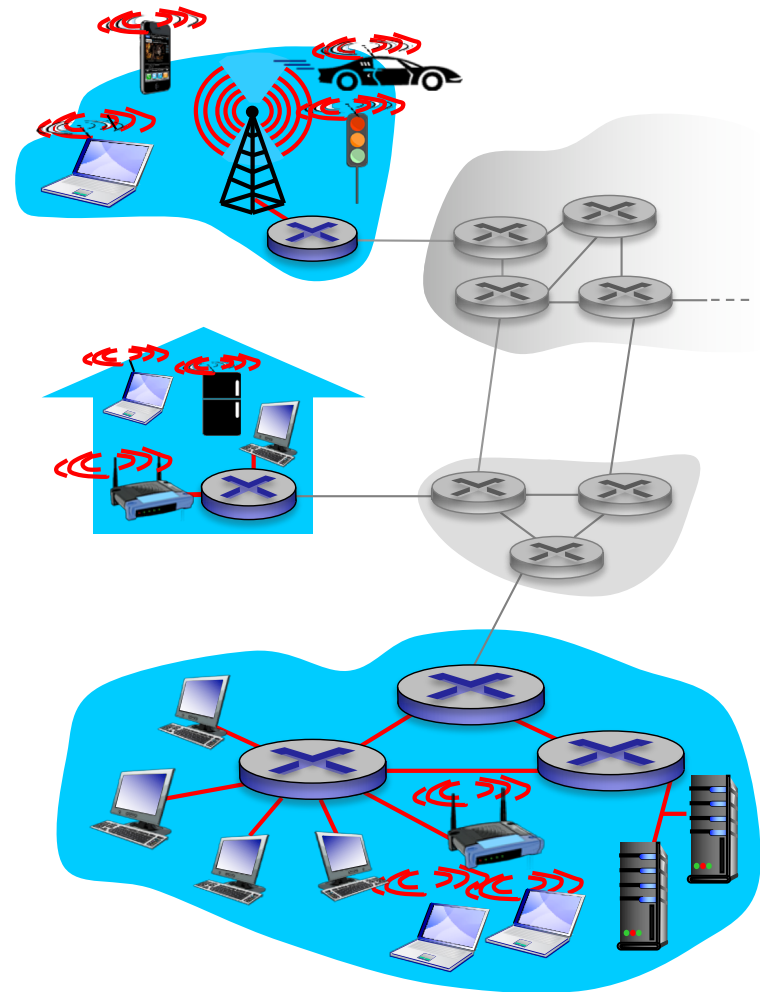 - interconnected routers
  - network of networks
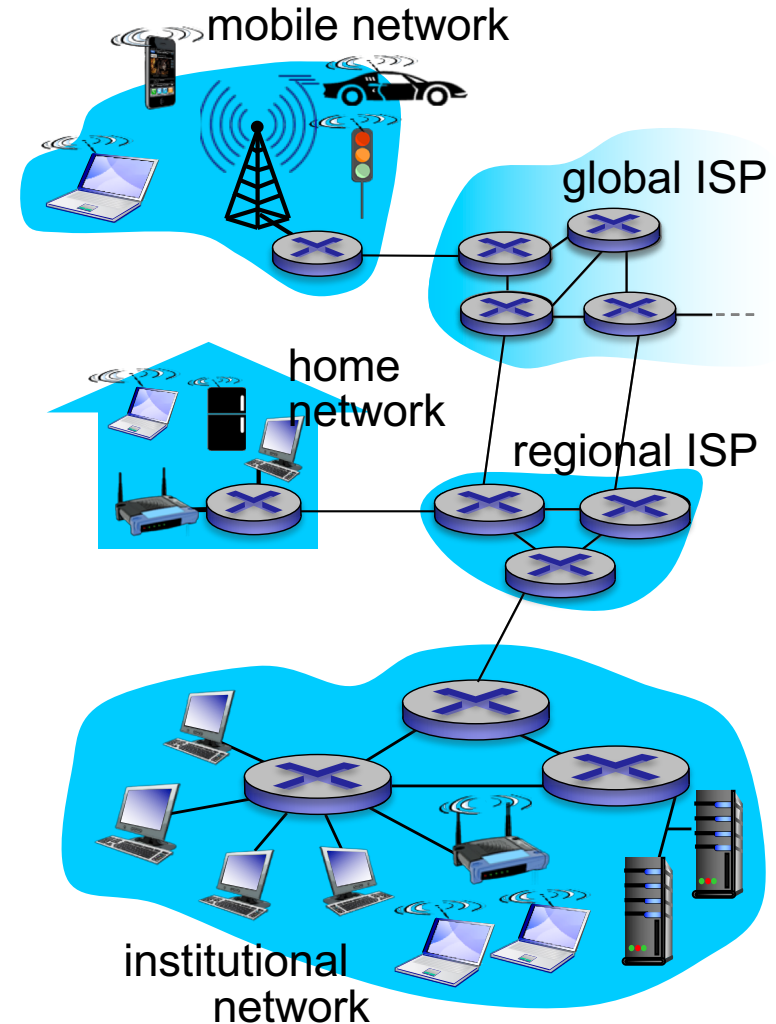
Access networks and physical media

Q: How to connect end systems to edge router?

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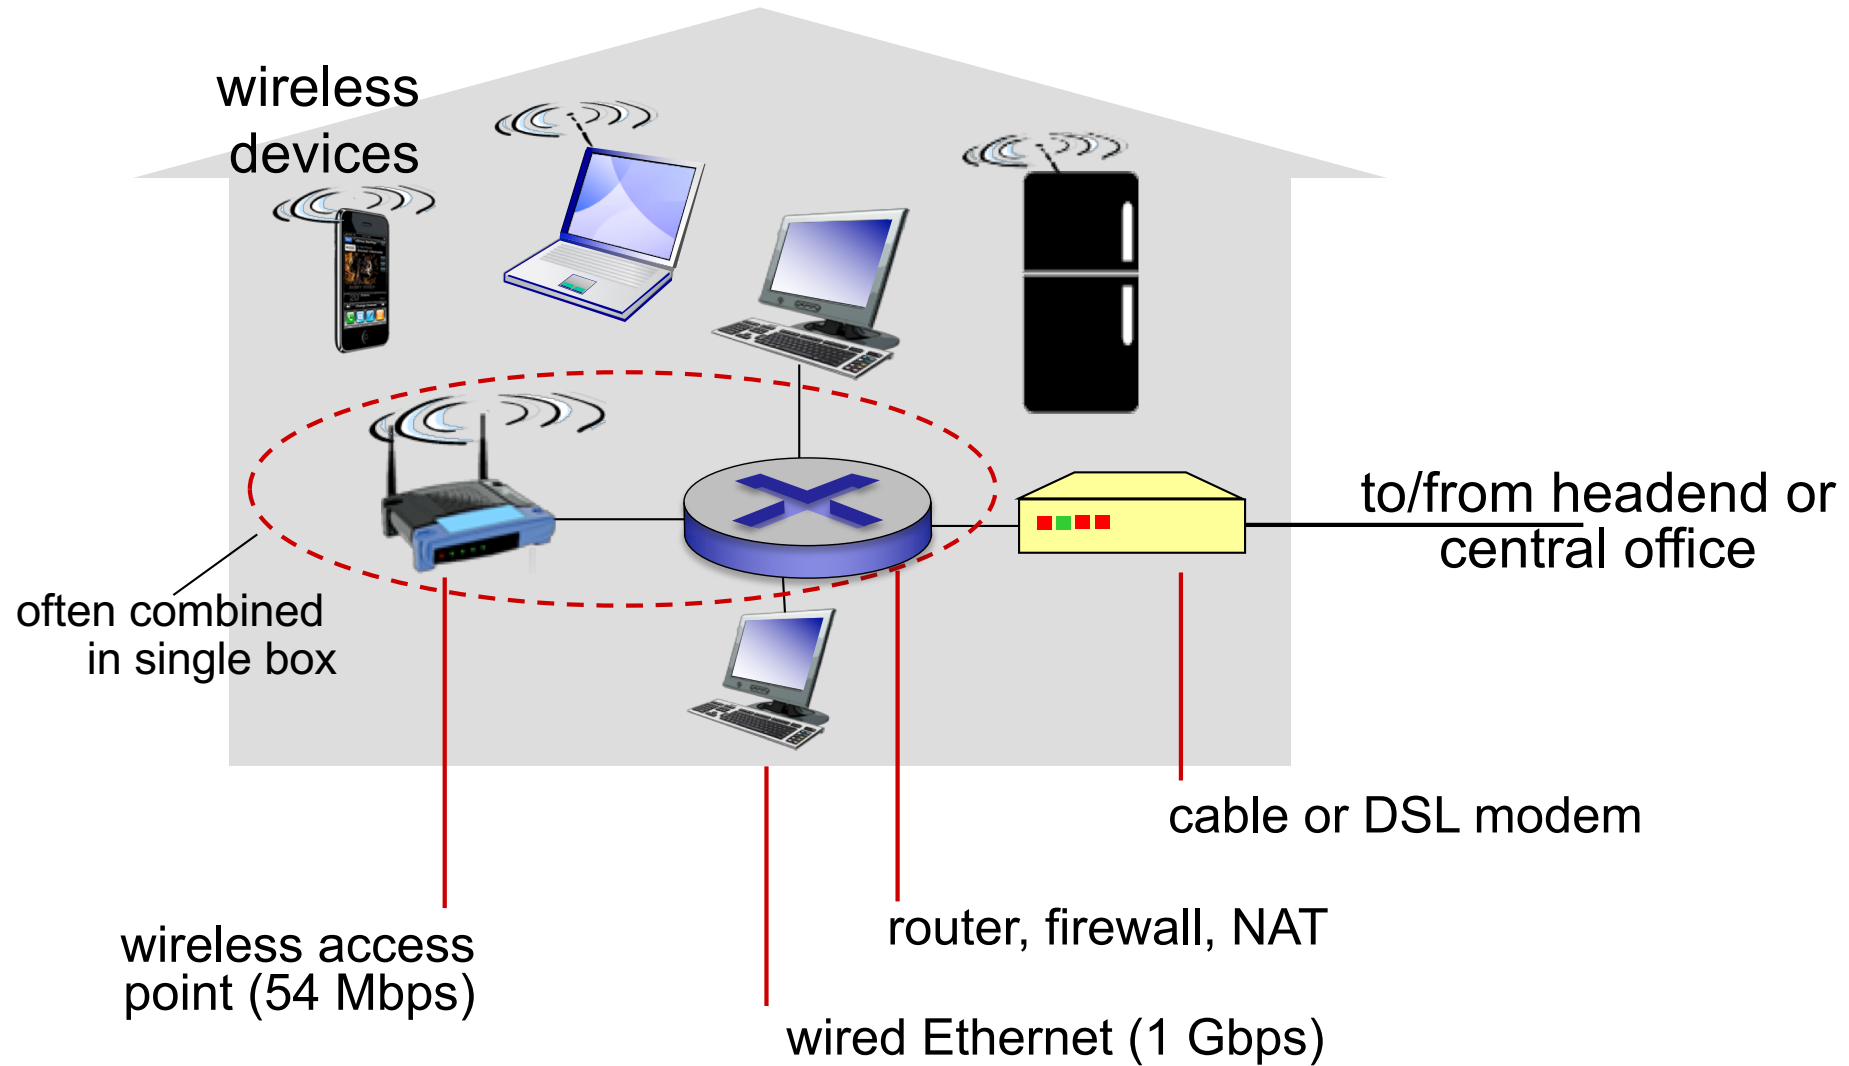- mobile access networks

keep in mind:

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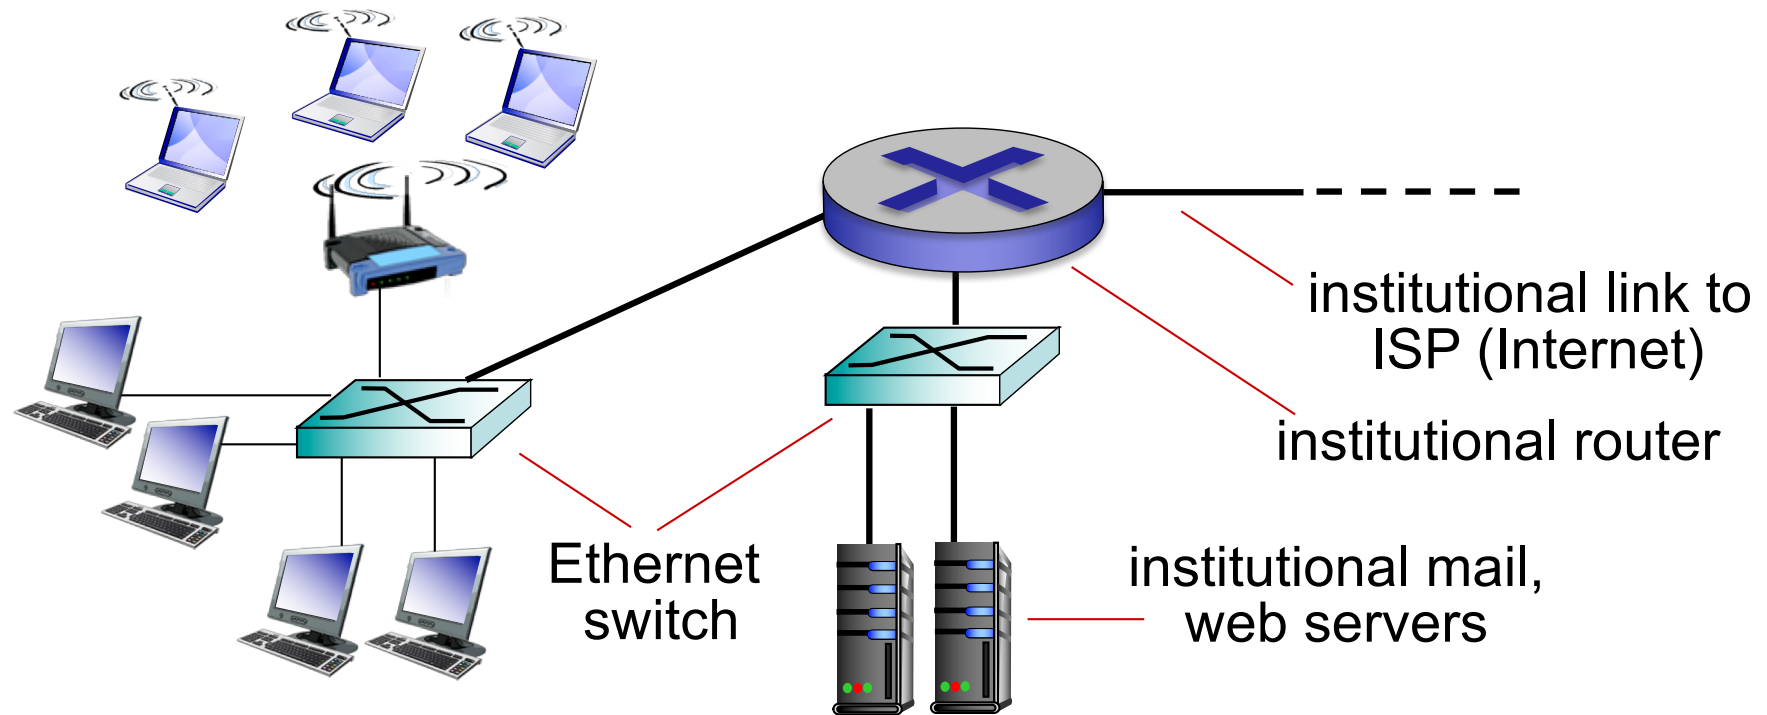- shared or dedicated?

Access network: home network

wireless
devices

often combined
in single box

wireless access
point (54 Mbps)

wired Ethernet (1 Gbps)

router, firewall, NAT

cable or DSL modem

to/from headend or
central office

Enterprise access networks (Ethernet)

institutional link to
ISP (Internet)

institutional router

Ethernet
switch

institutional mail,
web servers

- 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 router
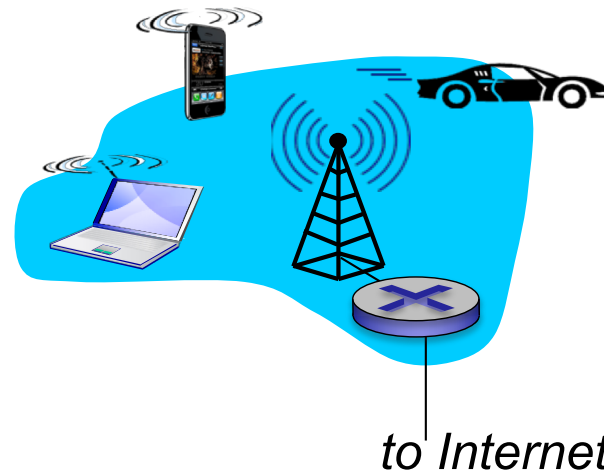 - via base station aka "access point"

wireless LANs:

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

to Internet
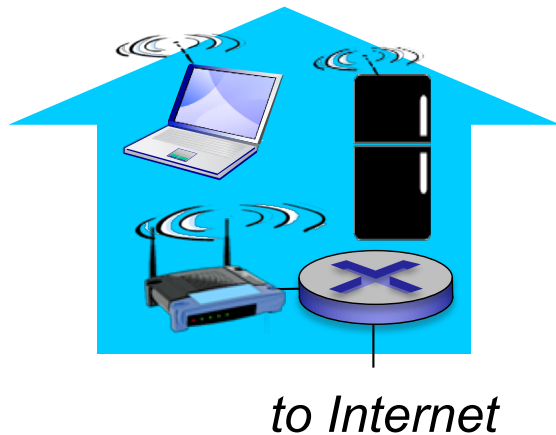
wide-area wireless access

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

to Internet

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*

two packets, *L* bits each

2 1

R: link transmission rate

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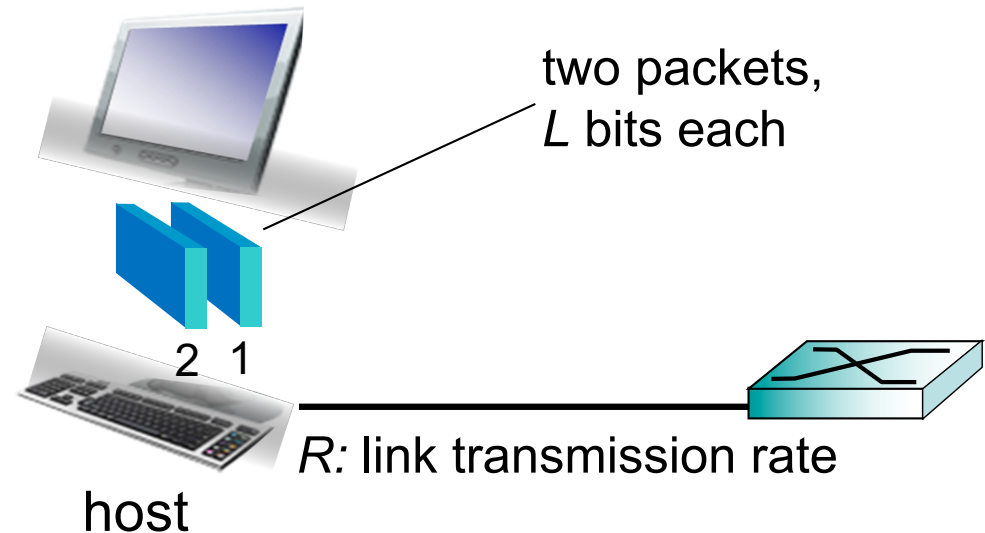
$$\text{packet transmission delay} = \text{time needed to transmit } L\text{-bit packet into link} = \frac{L \text{ (bits)}}{R \text{ (bits/sec)}}$$

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

Chapter 1: roadmap

The network core

- mesh of interconnected routers

- packet-switching: hosts break application-layer messages into *packets*
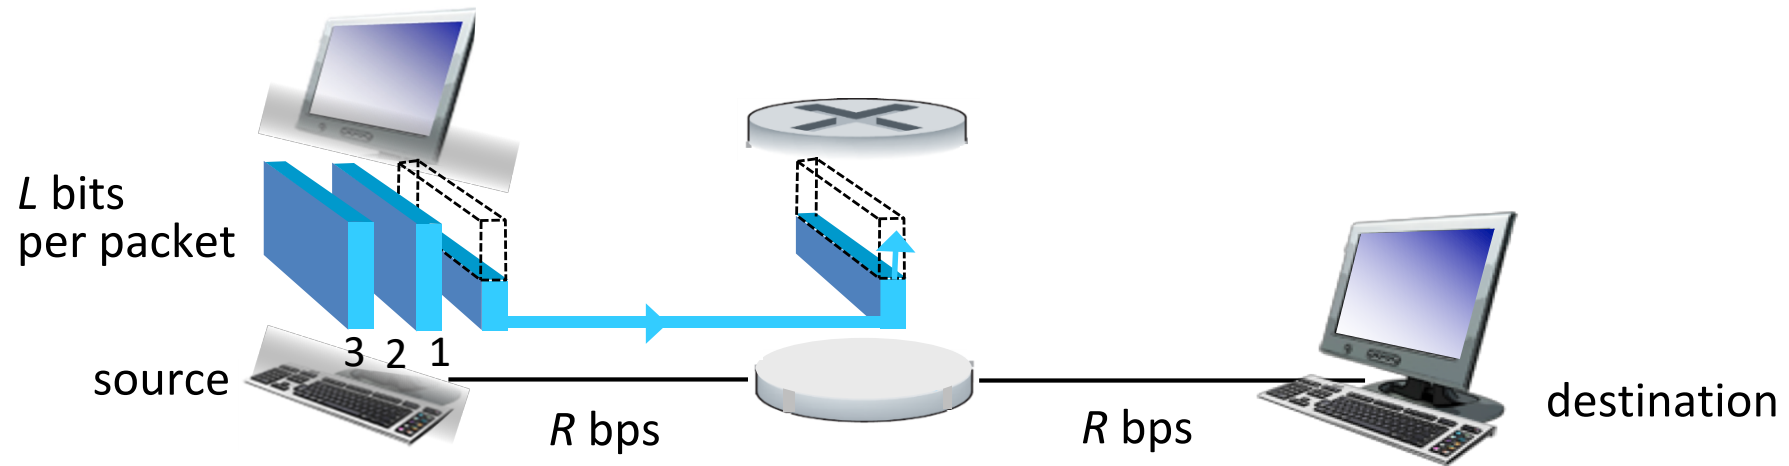
  - forward packets from one router to the next, across links on path from source to destination

  - each packet transmitted at full link capacity
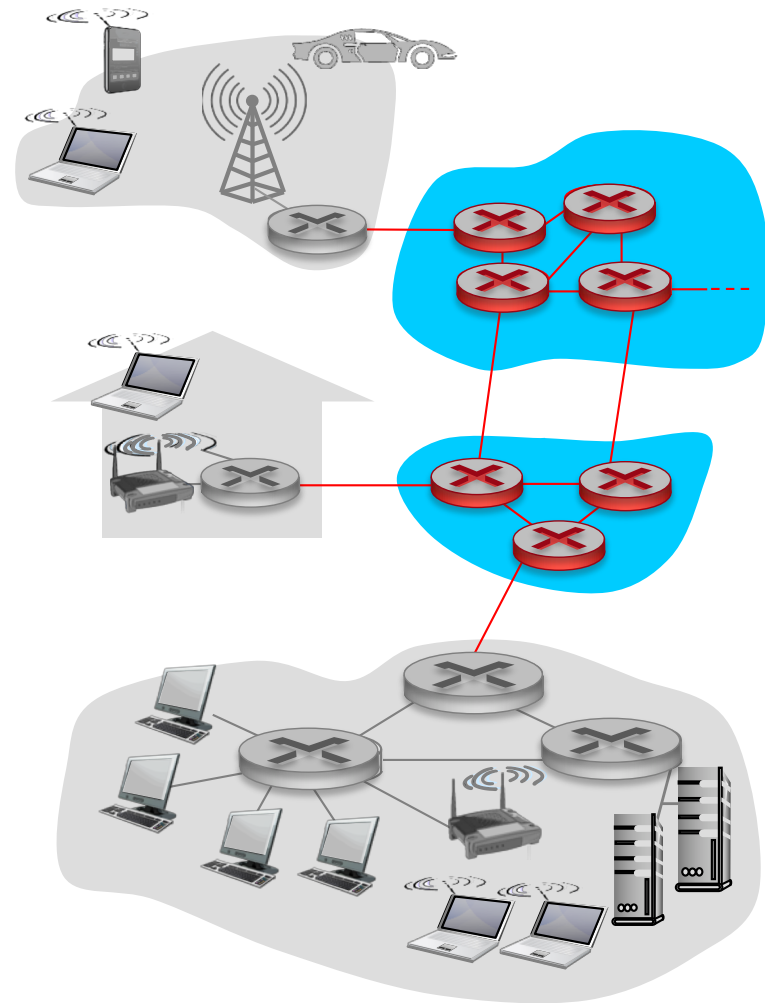
Packet-switching: store-and-forward

- 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 = 2*L/R* (assuming zero propagation delay)

one-hop numerical example:
- *L* = 7.5 Mbits
- *R* = 1.5 Mbps
- one-hop transmission delay = 5 sec

more on delay shortly …

Packet Switching: queueing delay, loss

A   R = 100 Mb/s

B

queue of packets
waiting for output link

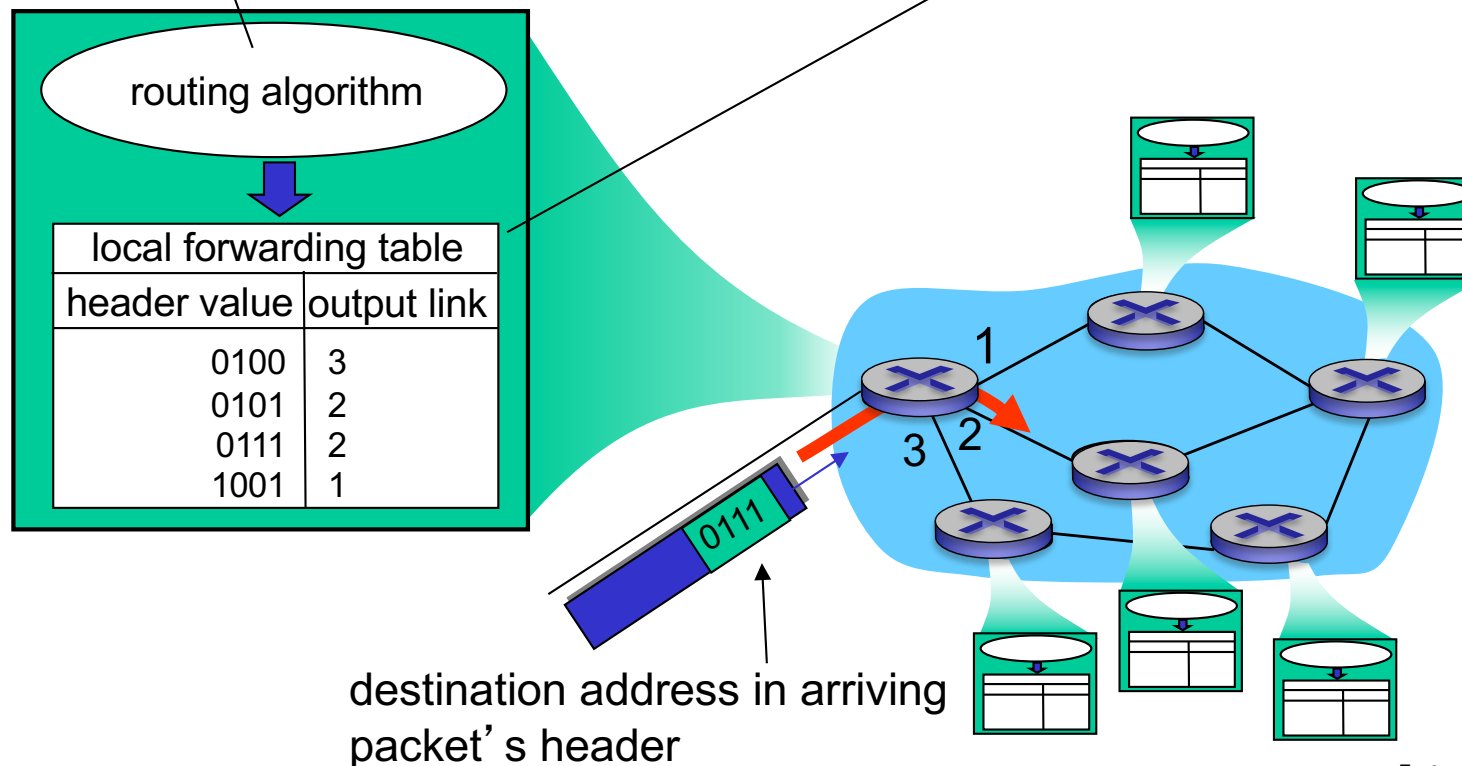R = 1.5 Mb/s

C

D

E
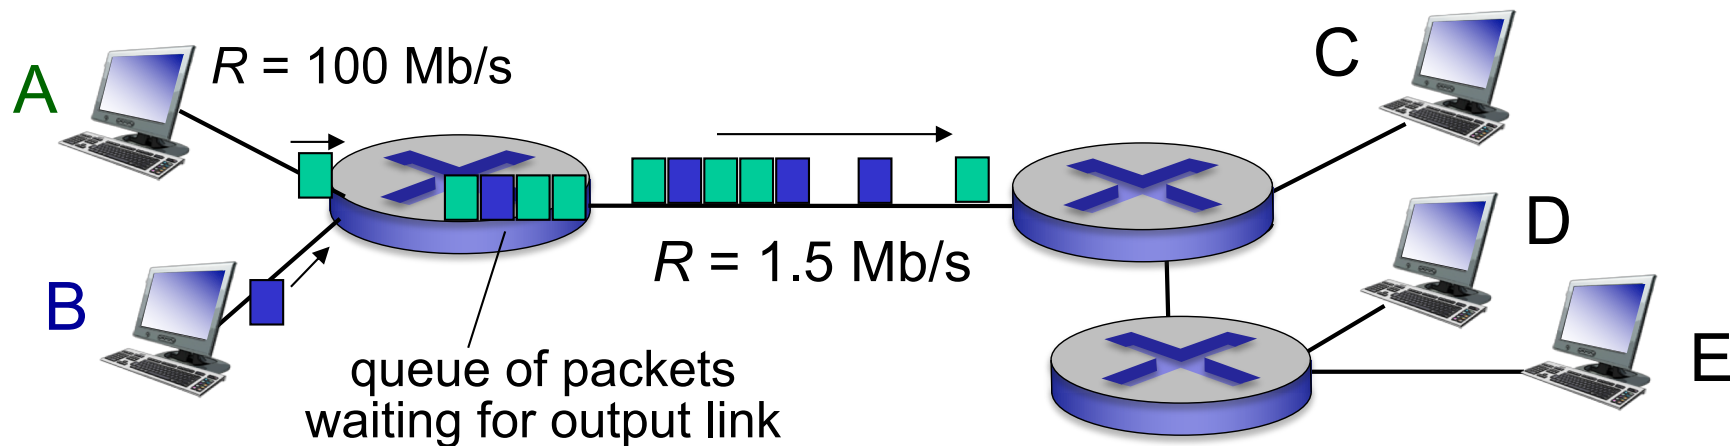
queuing and loss:

- if arrival rate (in bits) to link exceeds transmission rate of link for a period of time:
  - packets will queue, wait to be transmitted on link
  - packets can be dropped (lost) if memory (buffer) fills up

Two key network-core functions

routing: determines source-destination route taken by packets

- *routing algorithms*

forwarding: move packets from router's input to appropriate router output

routing algorithm

↓

local forwarding table	
header value	output link
0100	3
0101	2
0111	2
1001	1

0111

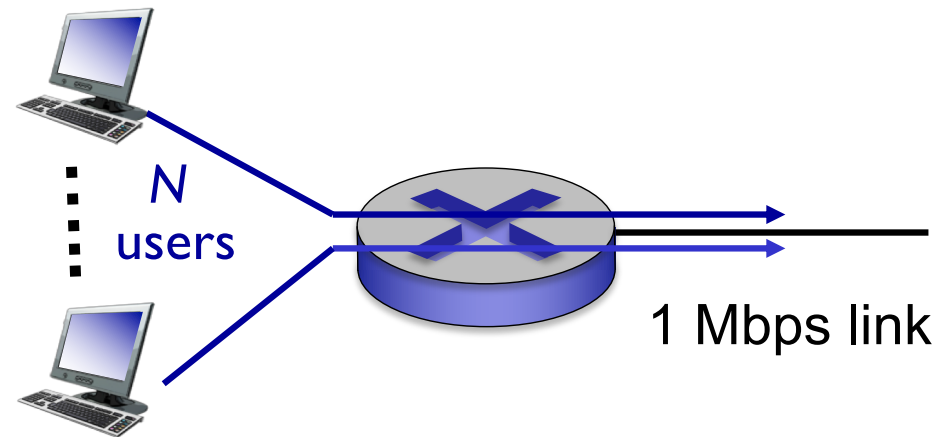destination address in arriving packet's header

1

3 2

Packet switching versus circuit switching

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 *

N users

1 Mbps link

Q: how did we get value 0.0004?

Q: what happens if > 35 users ?

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

Packet switching versus circuit switching

is packet switching a "slam dunk winner?"

- great for bursty data
  - resource sharing
  - simpler, no call setup
- excessive congestion possible: packet delay and loss
  - protocols needed for reliable data transfer, congestion control
- Q: How to provide circuit-like behavior?
  - bandwidth guarantees needed for audio/video apps
 - still an unsolved problem (chapter 7)

Q: human analogies of reserved resources (circuit switching) versus on-demand allocation (packet-switching)?

Internet structure: 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
- Let's take a stepwise approach to describe current Internet structure
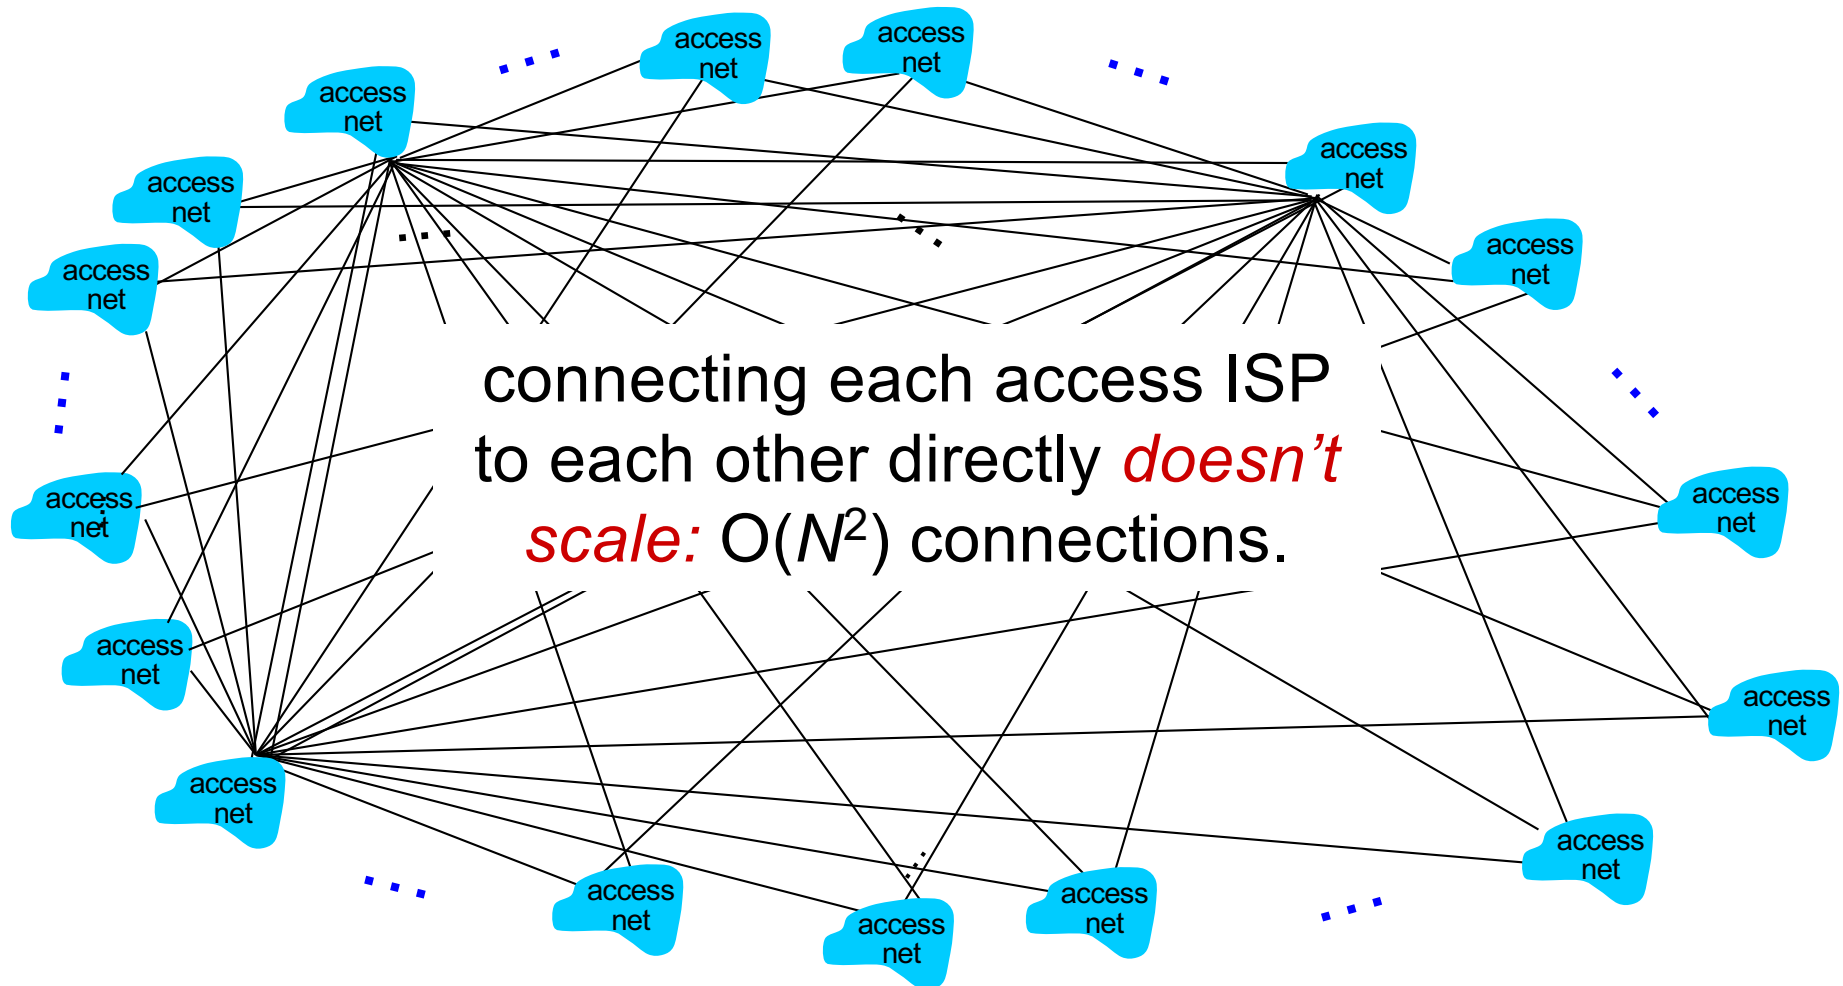
Internet structure: network of networks

Question: given *millions* of access ISPs, how to connect them together?
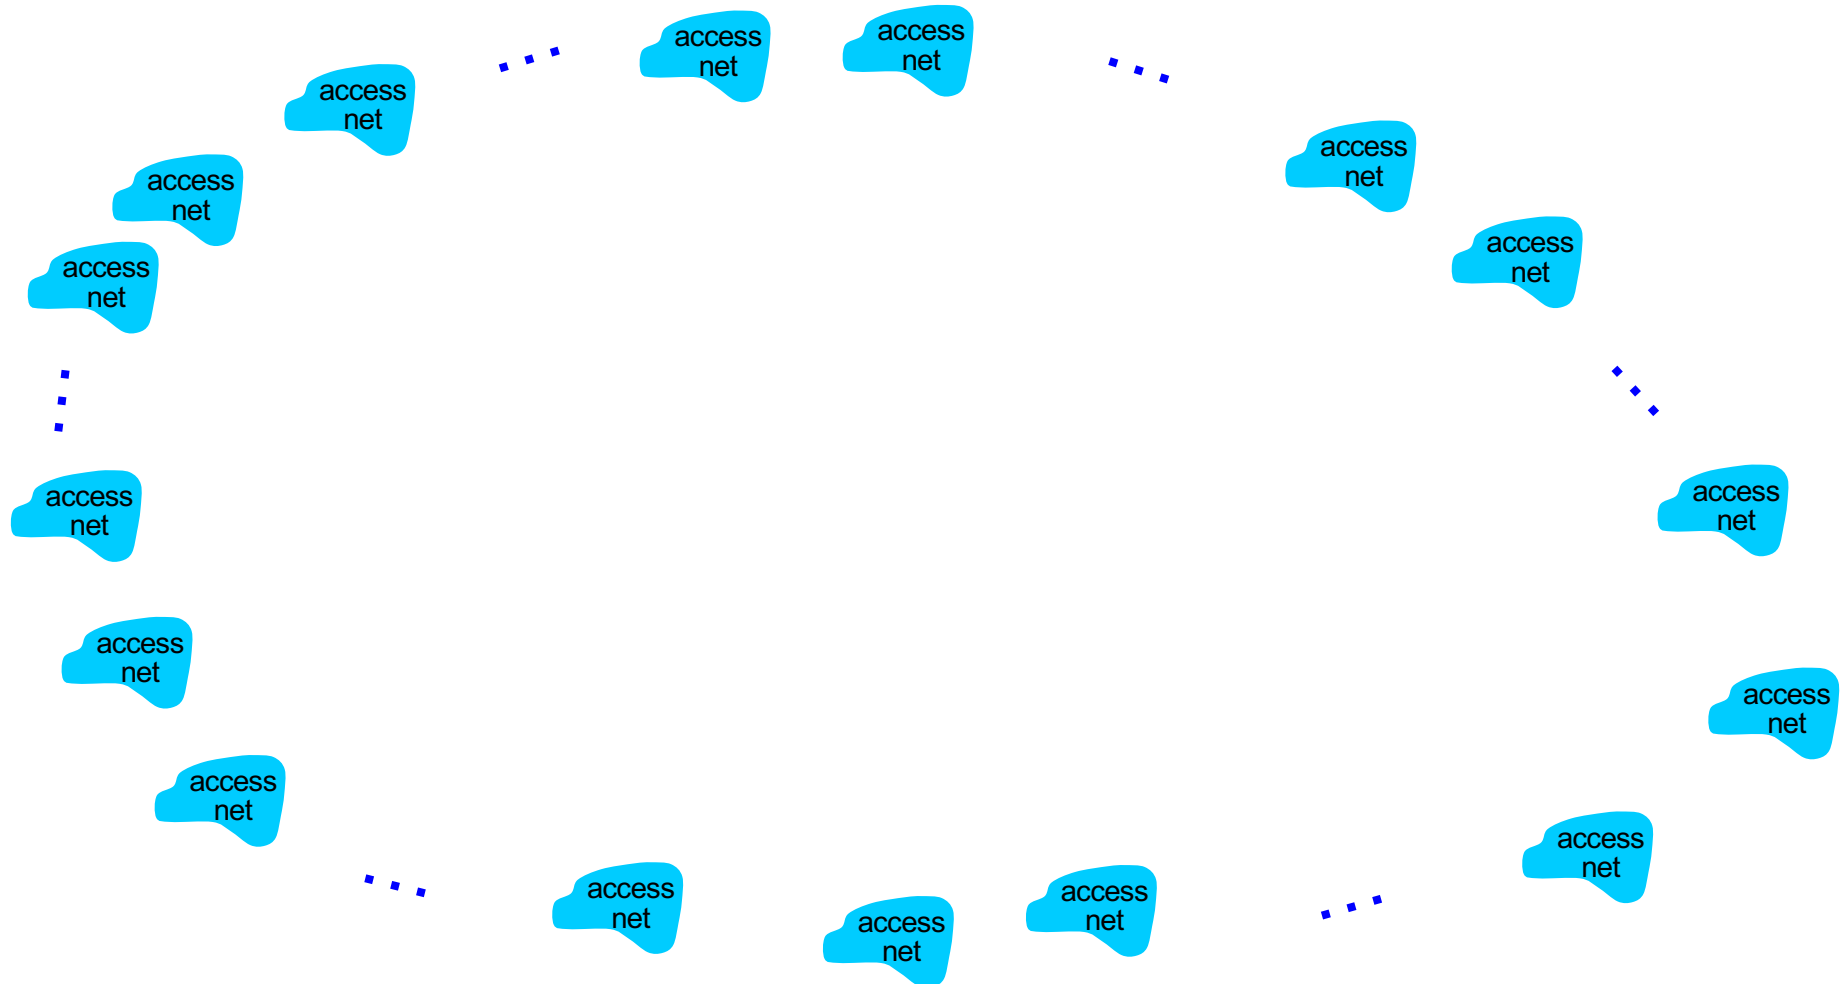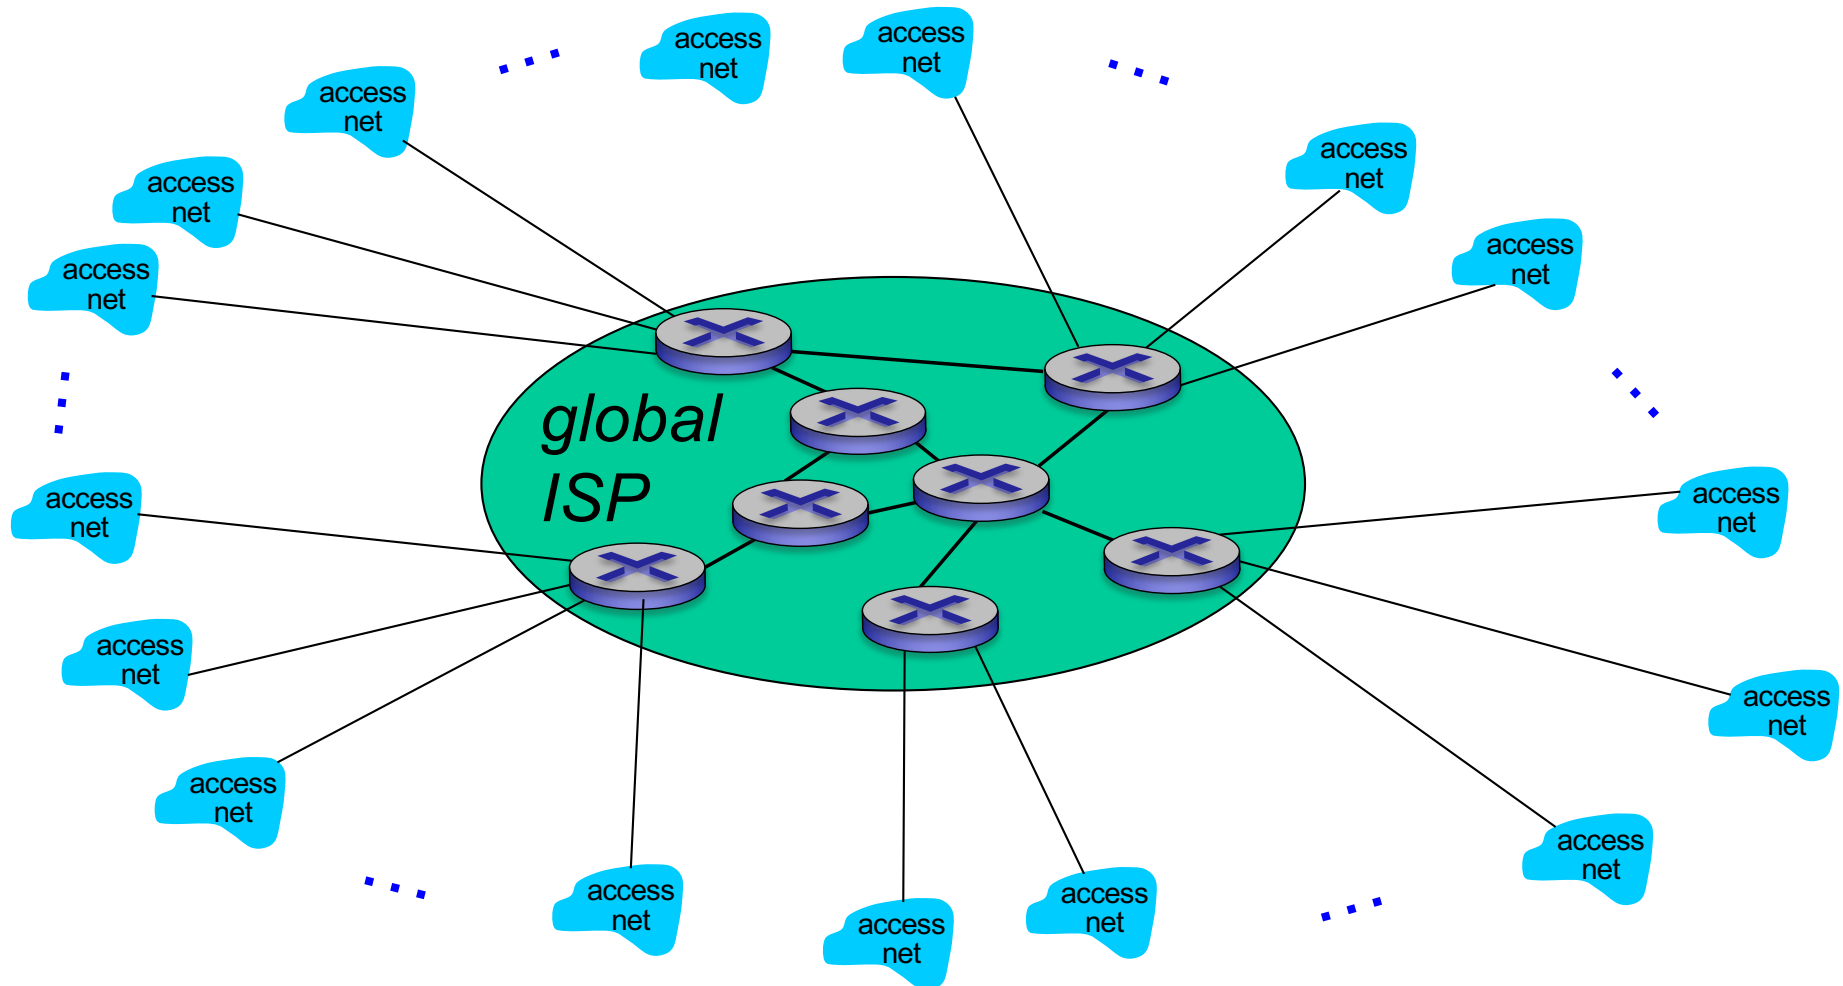
Internet structure: network of networks

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.

access net *(×20)*

Internet structure: network of networks

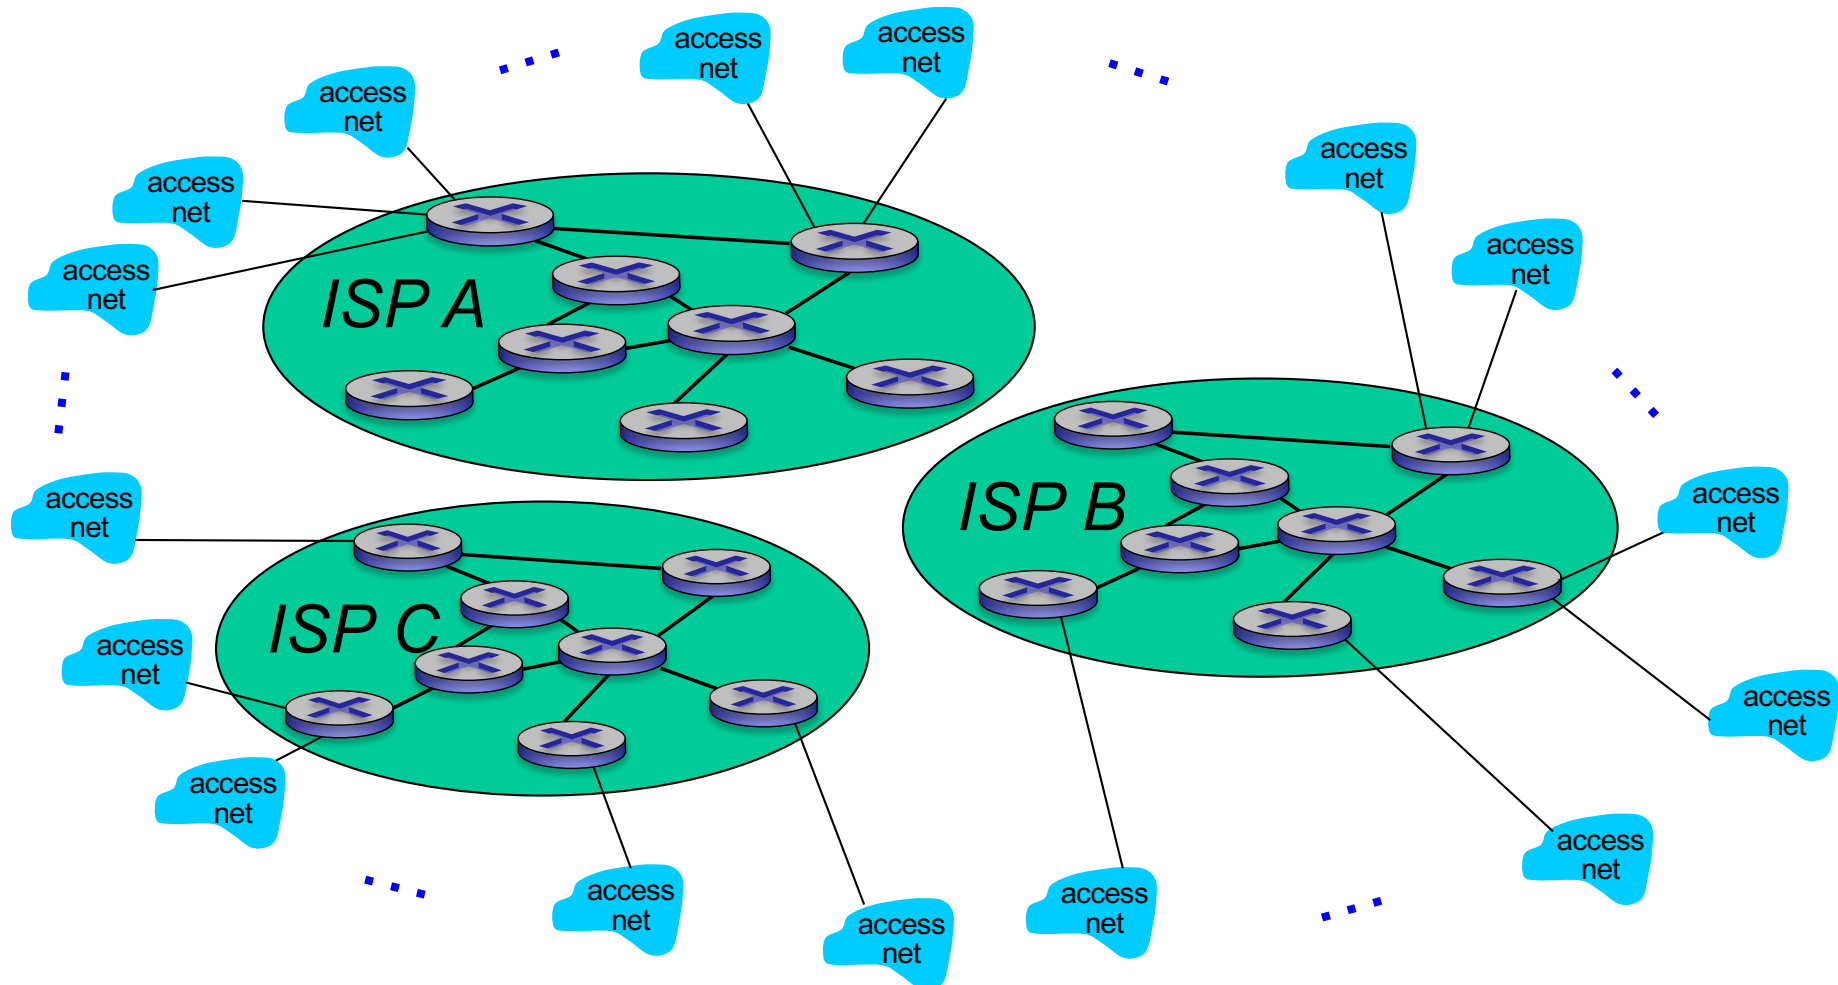*Option:* connect each access ISP to one global transit ISP?
Customer and *provider* ISPs have economic agreement.
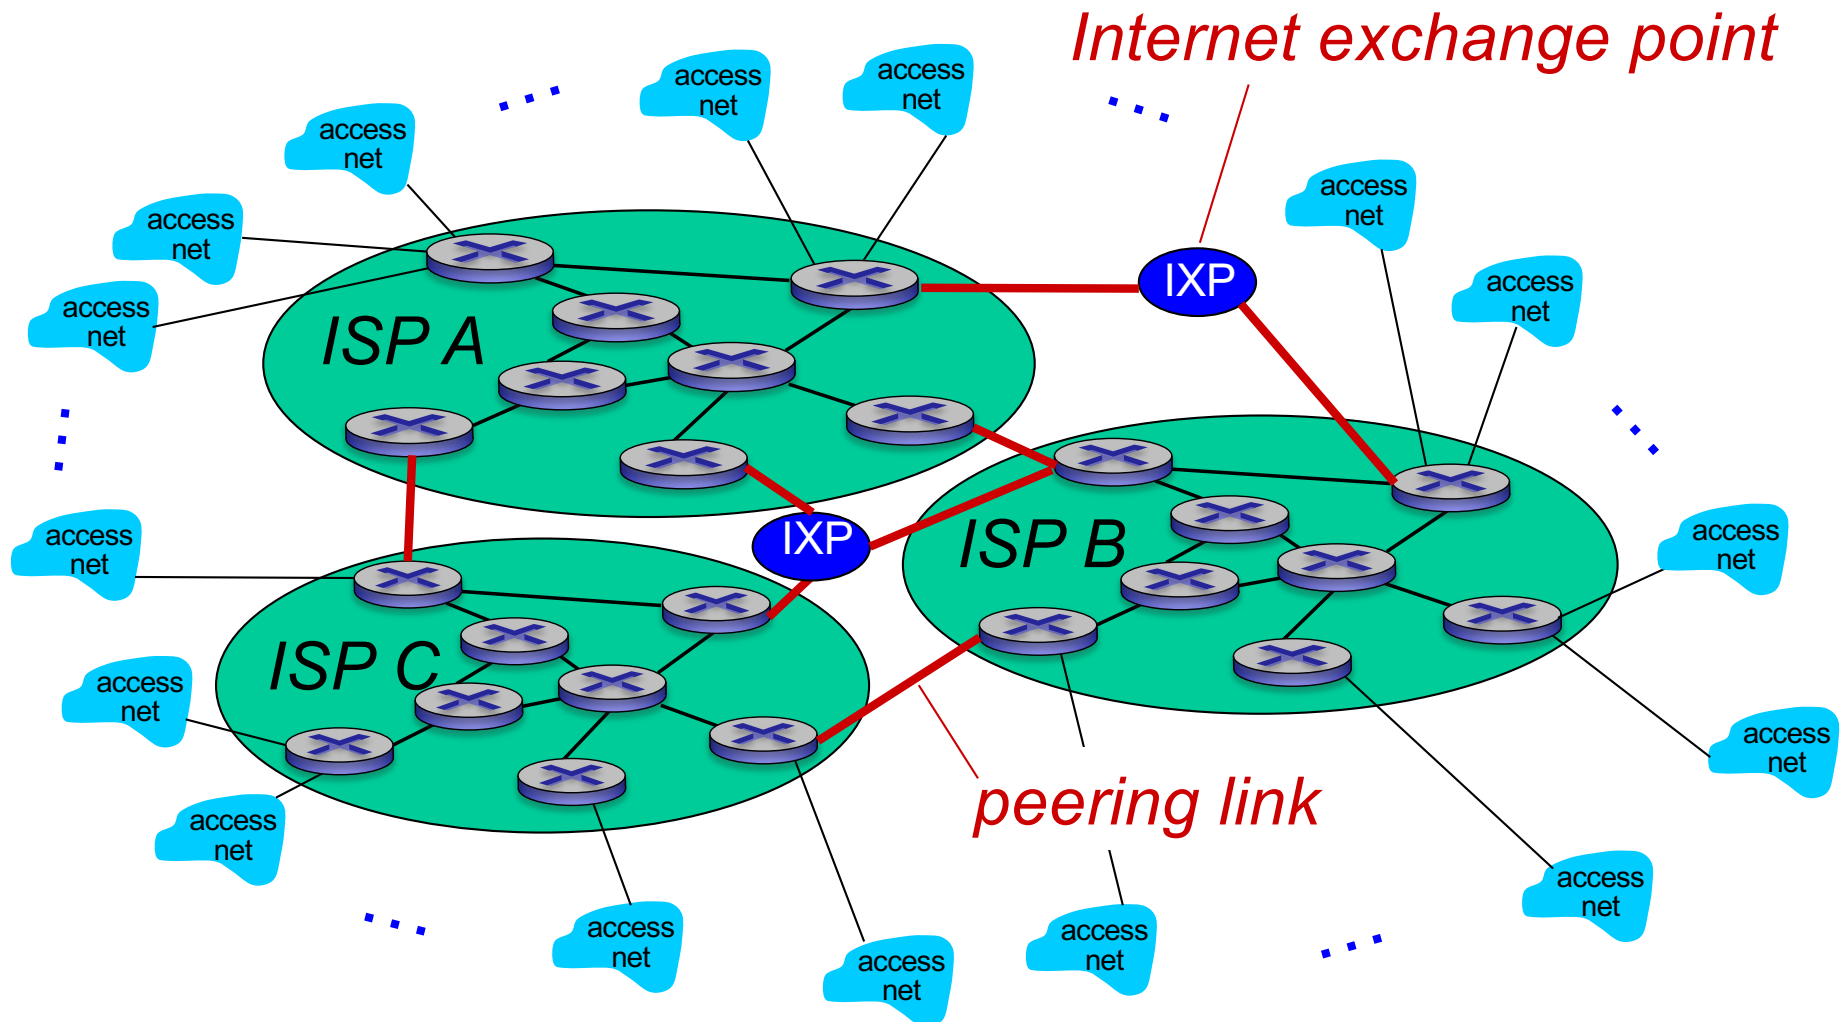
Internet structure: network of networks

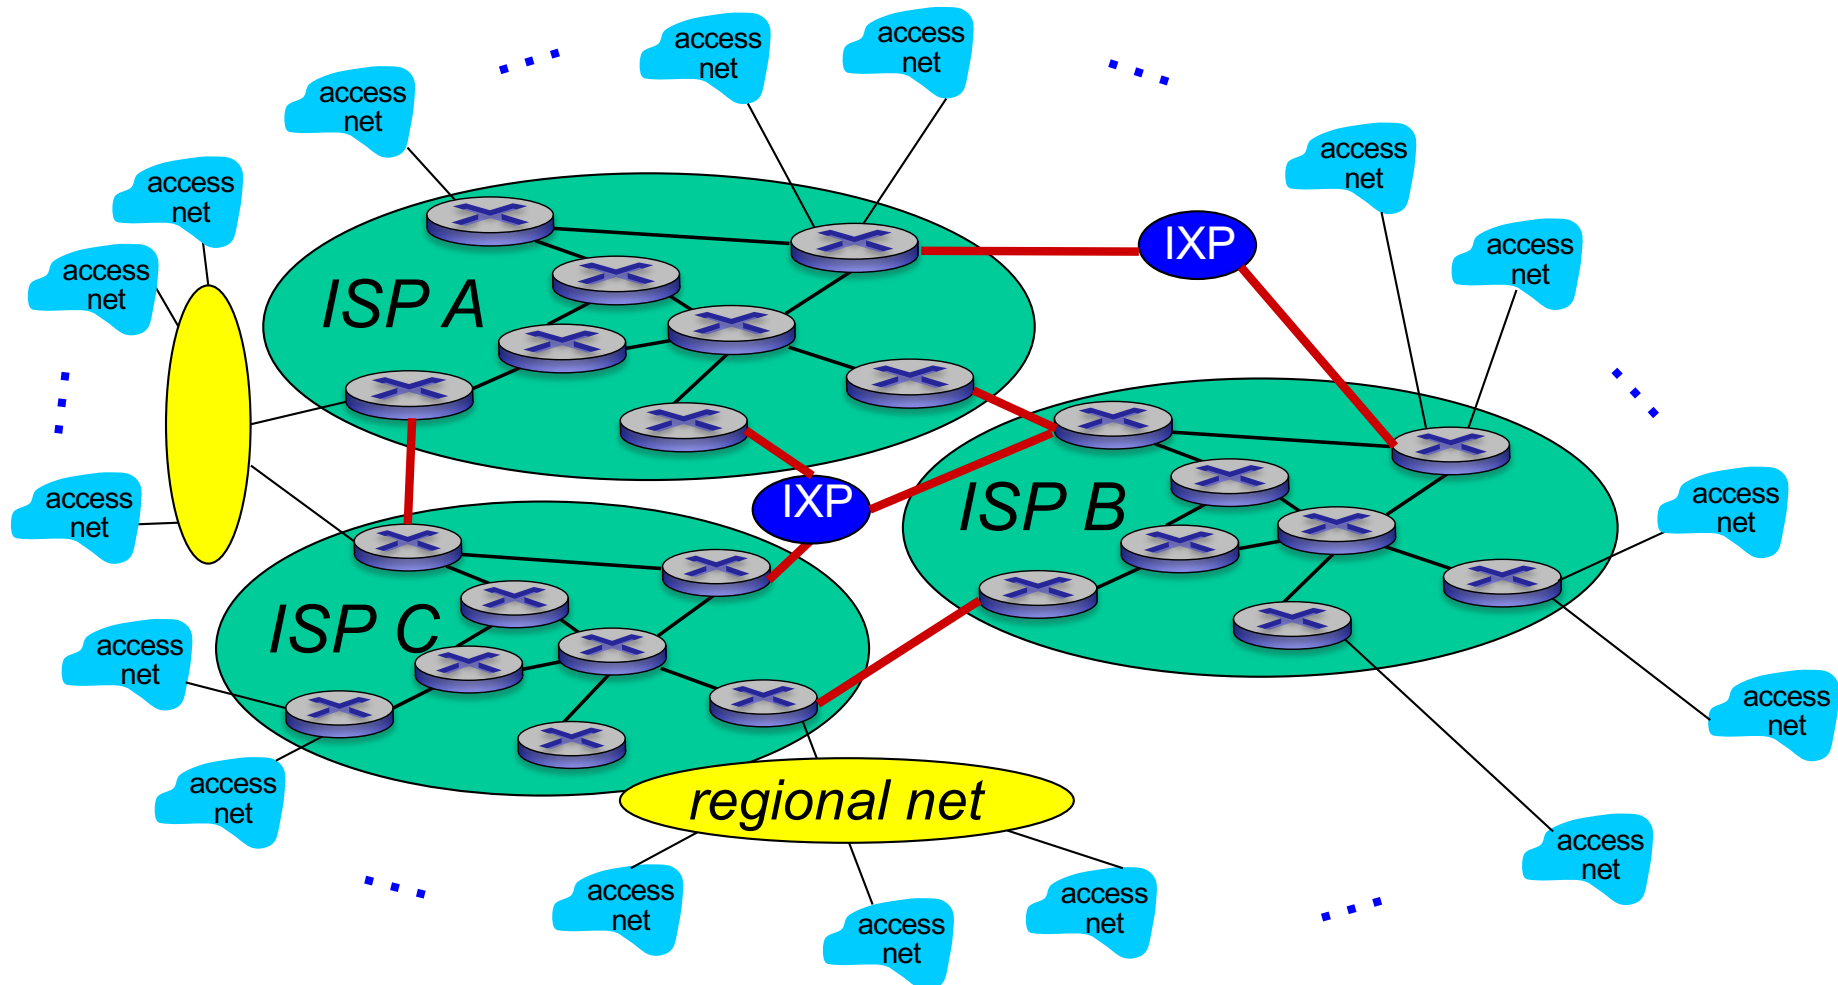But if one global ISP is viable business, there will be competitors
….

Internet structure: network of networks

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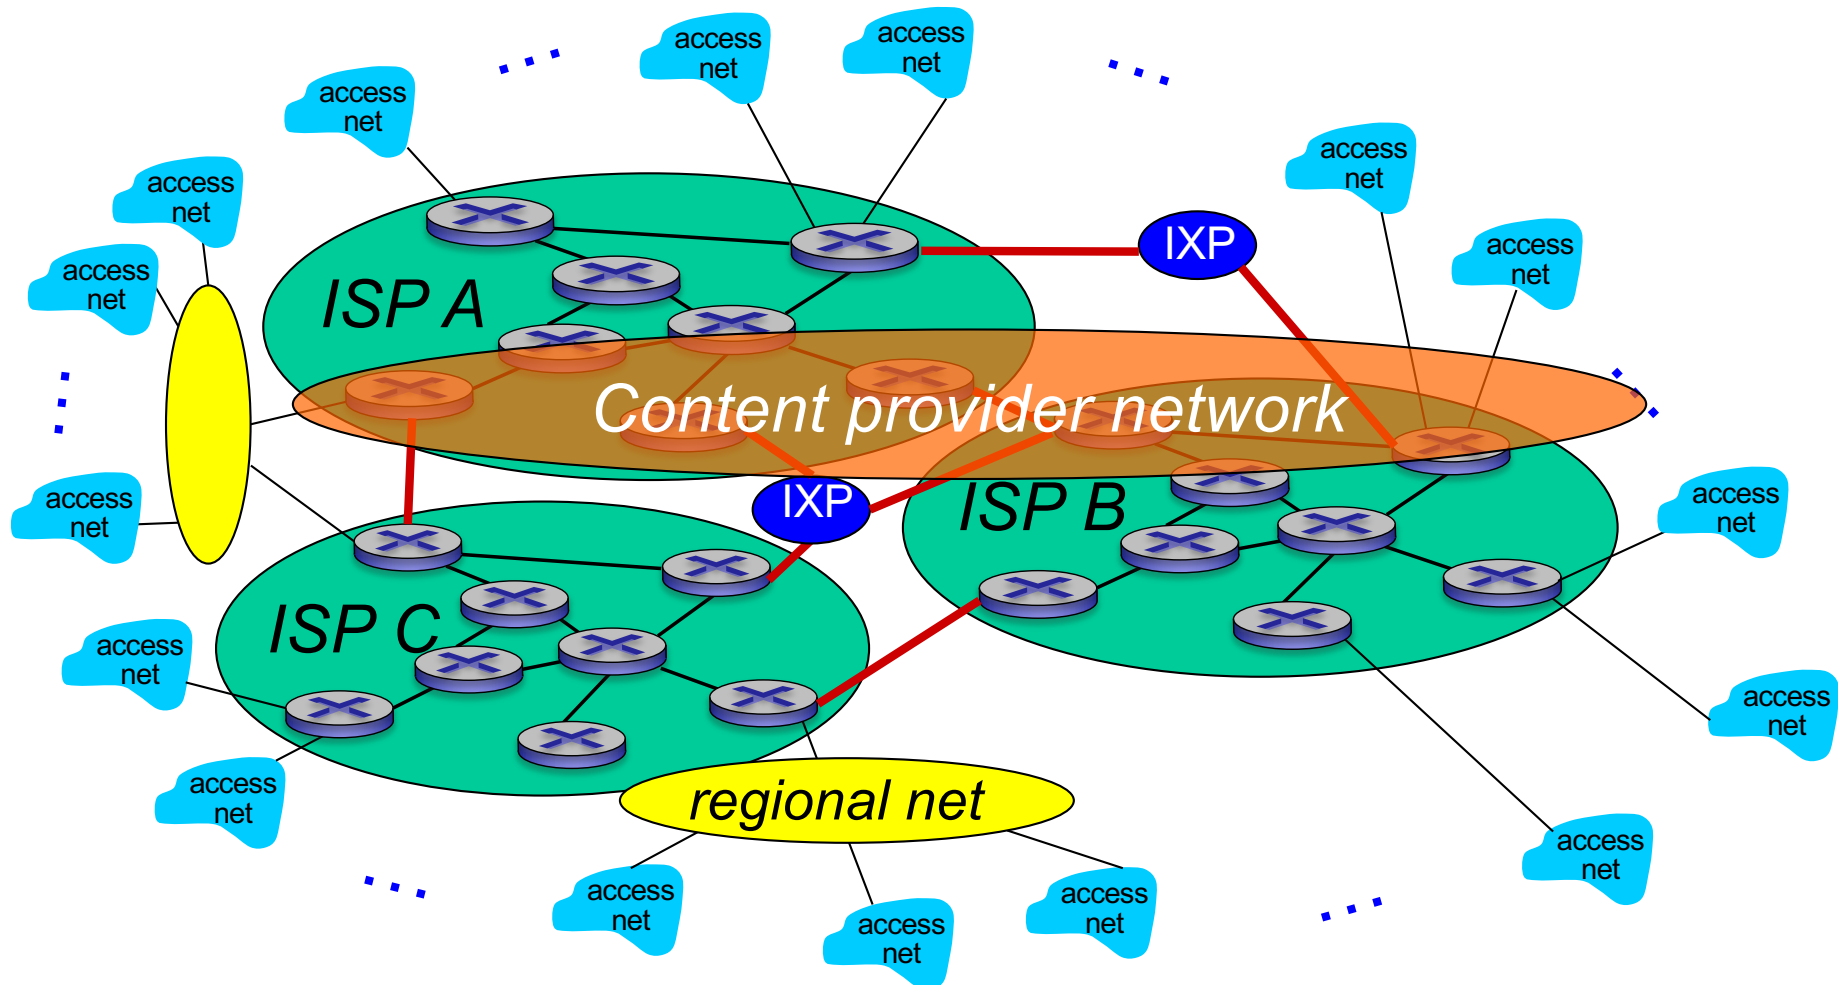…. which must be interconnected

Internet exchange point

peering link

Internet structure: network of networks

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

Internet structure: 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
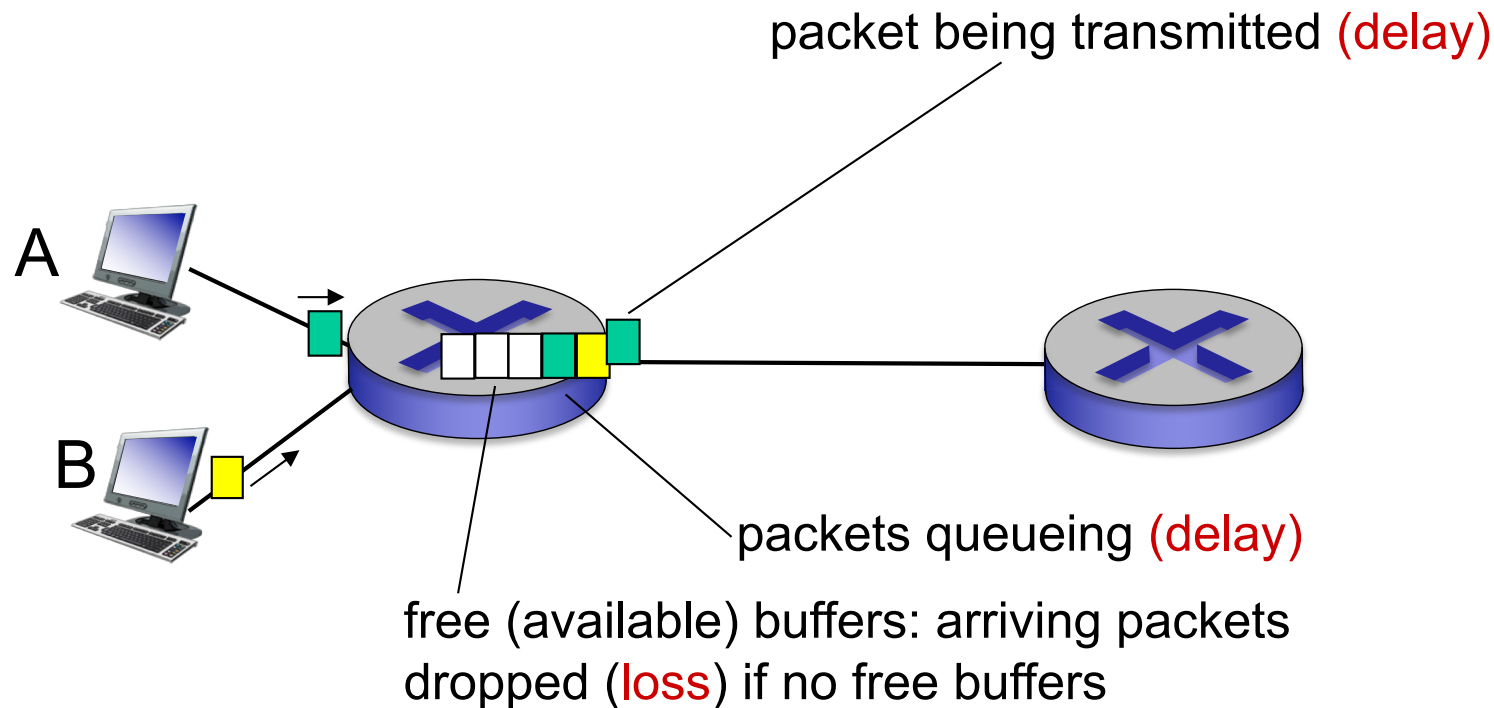
Chapter 1: roadmap

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

packet being transmitted (delay)

A

B

packets queueing (delay)

free (available) buffers: arriving packets
dropped (loss) if no free buffers

Four sources of packet delay
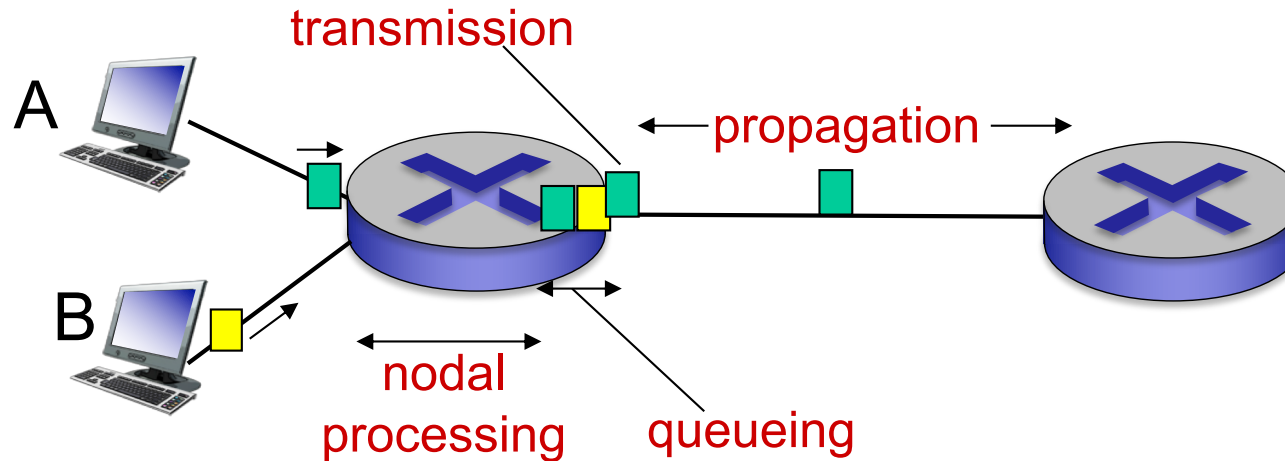
$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

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

Four sources of packet delay

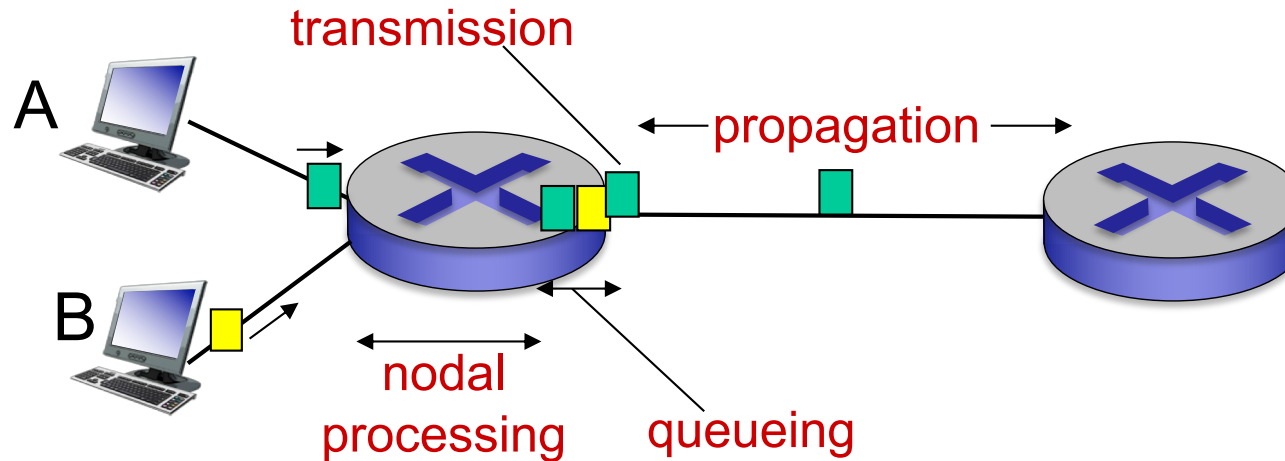$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

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 (~2x10^8 m/sec)
- d_{prop} = d/s

* Check out the online interactive exercises for more examples: http://gaia.cs.umass.edu/kurose_ross/interactive/
* Check out the Java applet for an interactive animation on trans vs. prop delay

Packet loss

- queue (aka buffer) preceding link in buffer 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

buffer
(waiting area)

packet being transmitted

A

B

packet arriving to
full buffer is *lost*
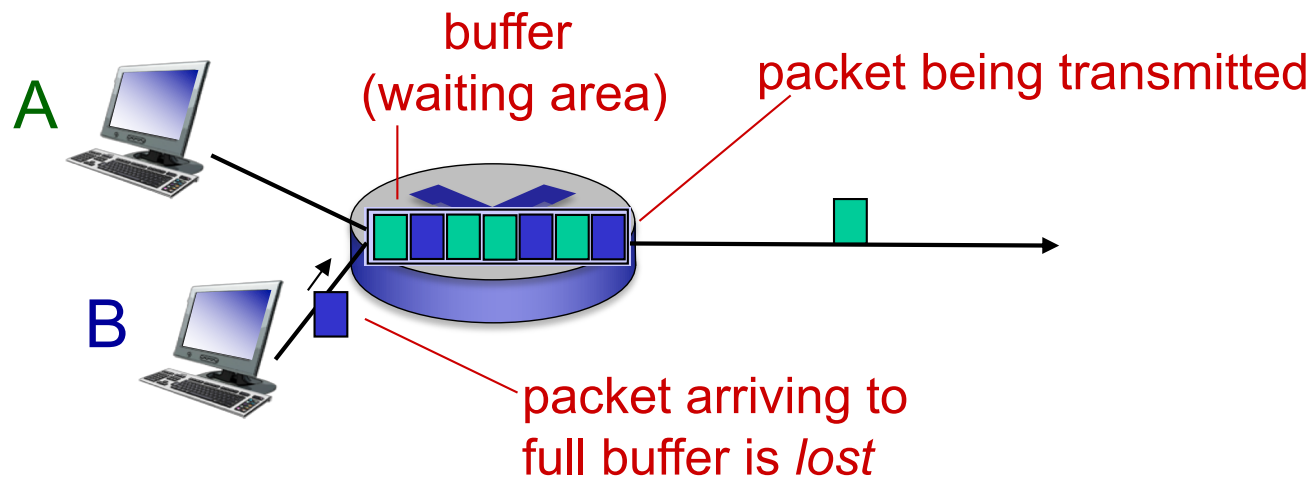
* Check out the Java applet for an interactive animation on queuing and loss

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

server sends bits
(fluid) into pipe

pipe that can carry
fluid at rate
R_s bits/sec)

pipe that can carry
fluid at rate
R_c bits/sec)
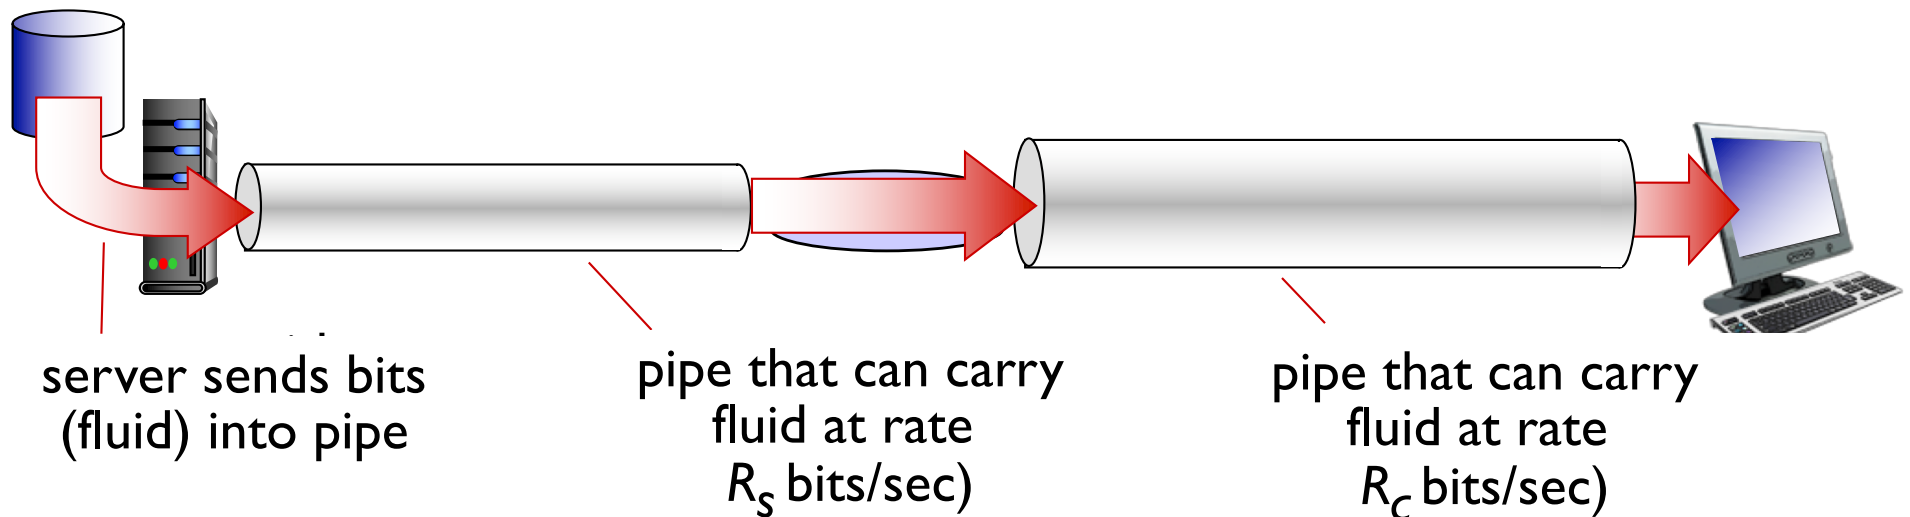
Throughput (more)

- $R_s < R_c$ What is average end-end throughput?
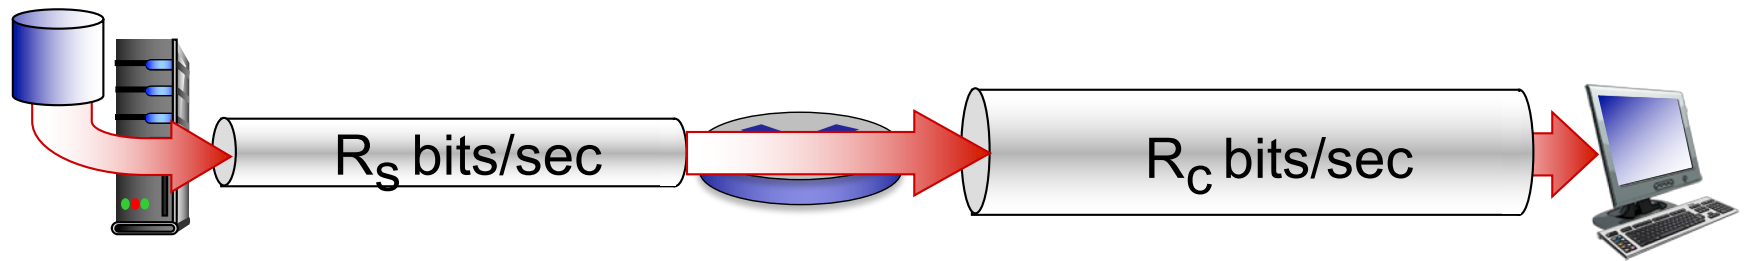
- $R_s > R_c$ What is average end-end throughput?
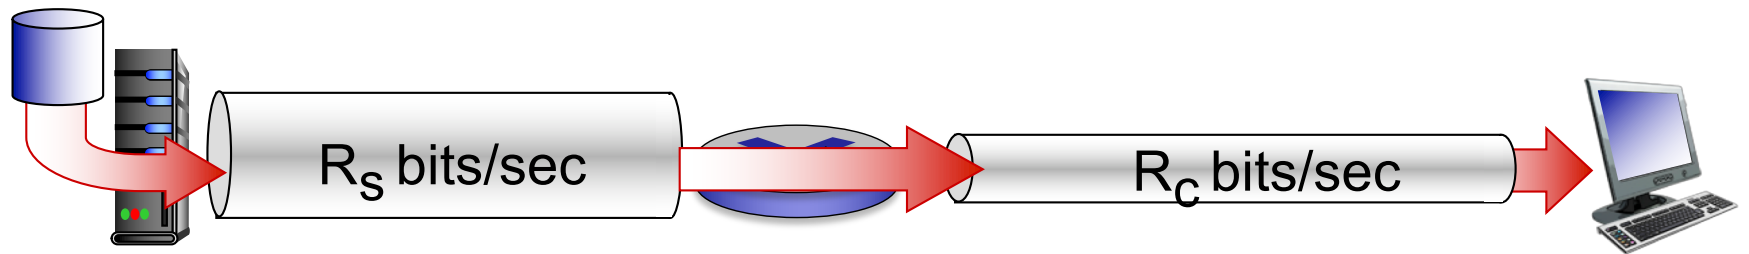
bottleneck link

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

Chapter 1: roadmap

Protocol "layers"

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 of implementation of layer's service transparent to rest of system
  - e.g., change in gate procedure doesn't affect rest of system
- layering considered harmful?

Internet protocol stack

- *application:* supporting network applications
  - FTP, 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, 802.111 (WiFi), PPP

- *physical:* bits "on the wire"

application
transport
network
link
physical
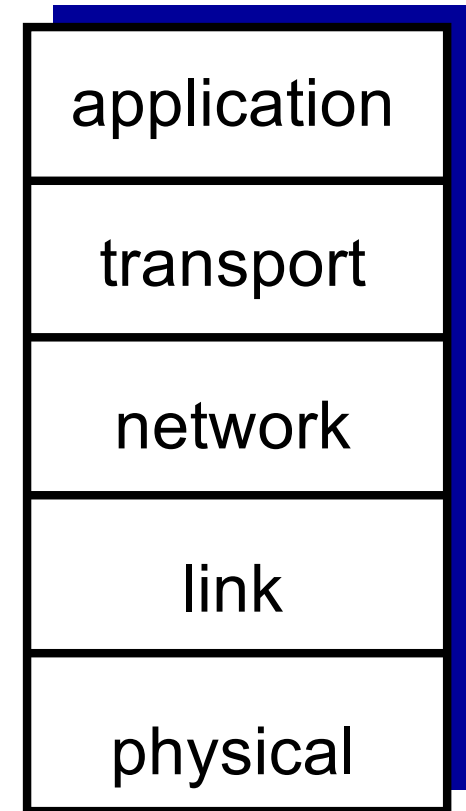
ISO/OSI reference model

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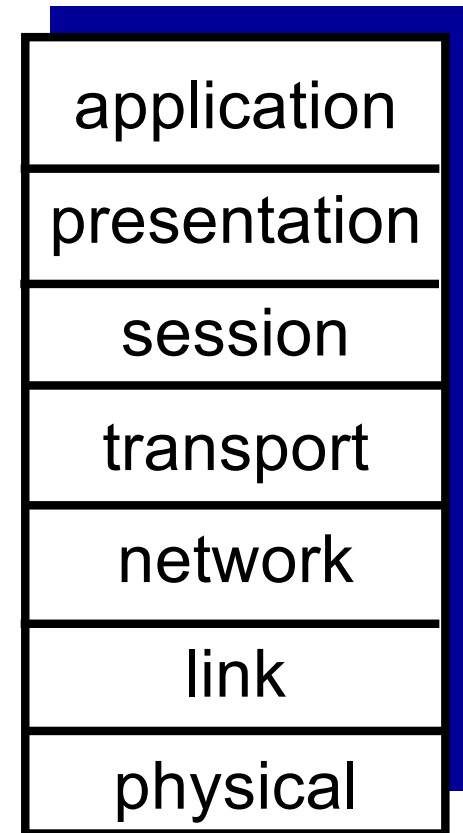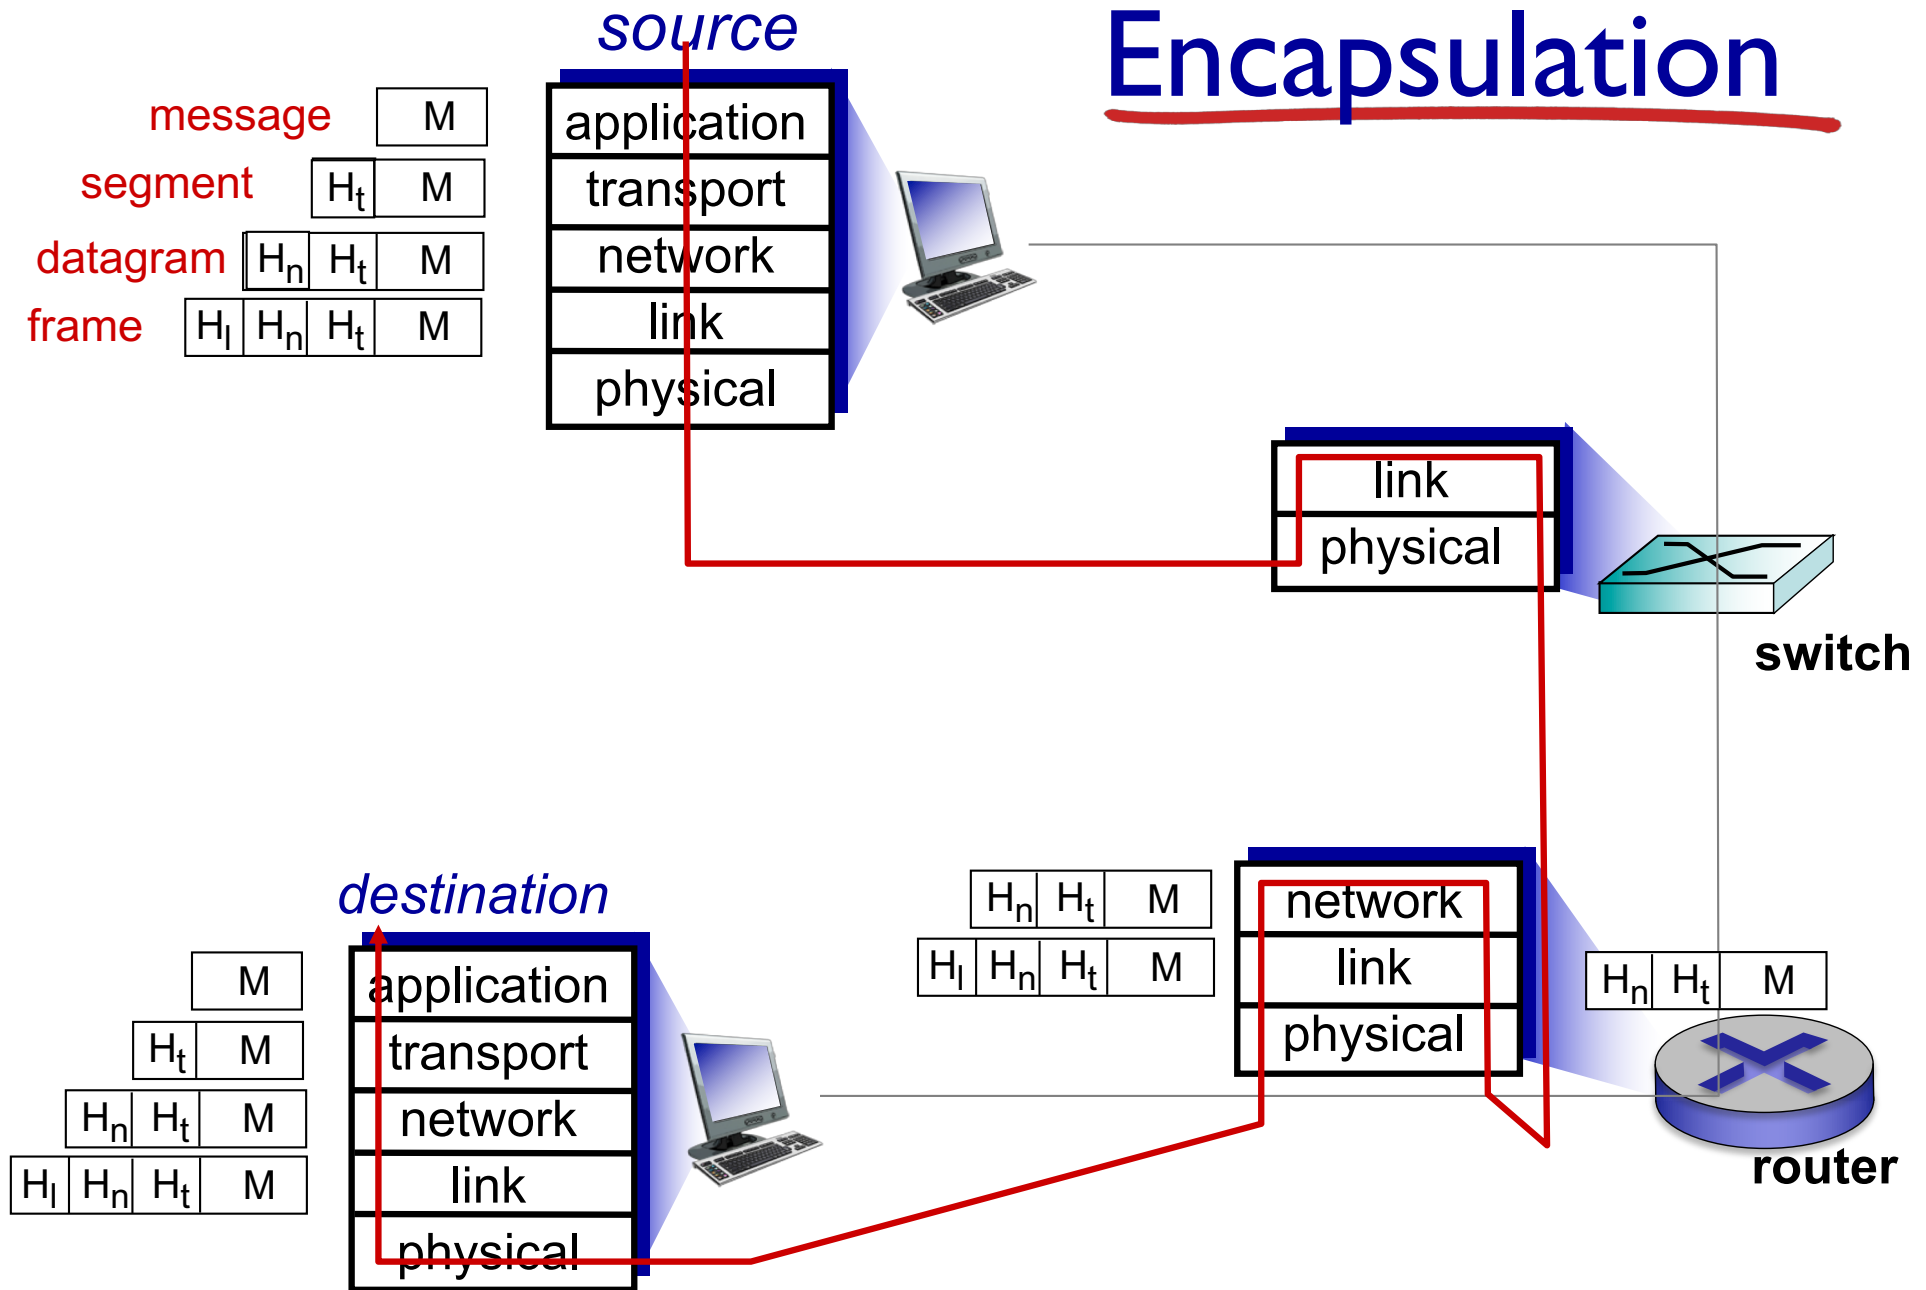
application
presentation
session
transport
network
link
physical

Encapsulation

source

message	M
segment	H_t	M
datagram	H_n	H_t	M
frame	H_l	H_n	H_t	M

application
transport
network
link
physical

link
physical

switch

destination

| M |
H_t	M
H_n	H_t	M
H_l	H_n	H_t	M

application
transport
network
link
physical

| H_n | H_t | M |
| H_l | H_n | H_t | M |

network
link
physical

| H_n | H_t | M |

router

Chapter 1: roadmap

Network security

- **field of network security:**
  - 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: put malware into hosts via Internet

- **malware can get in host from:**

  - *virus:* self-replicating infection by receiving/executing object (e.g., e-mail attachment)

  - *worm:* self-replicating infection by passively receiving object that gets itself executed

- **spyware malware** can record keystrokes, web sites visited, upload info to collection site

- infected host can be enrolled in **botnet,** used for spam. DDoS attacks

Bad guys: attack server, 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)
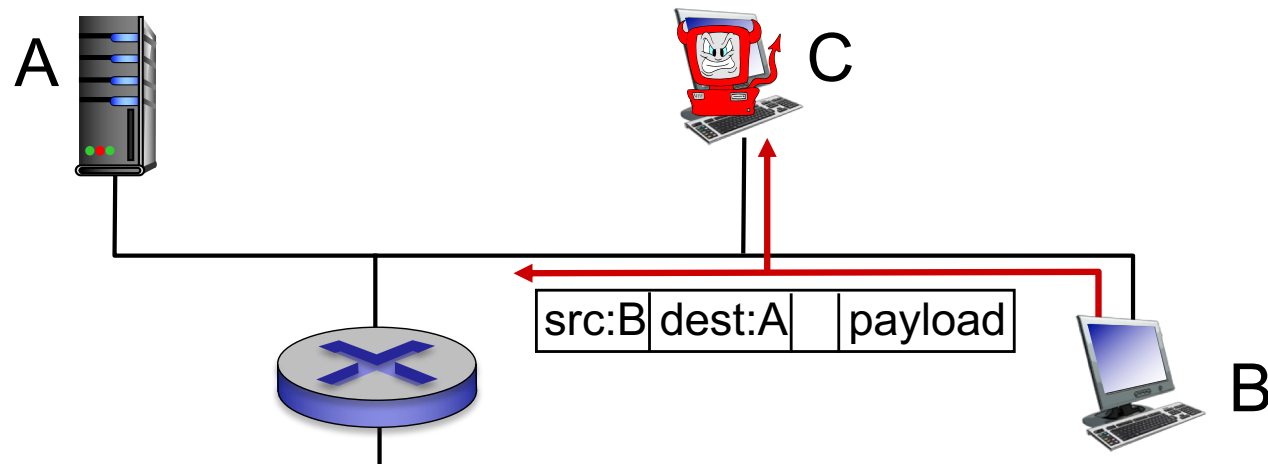
3. send packets to target from compromised hosts

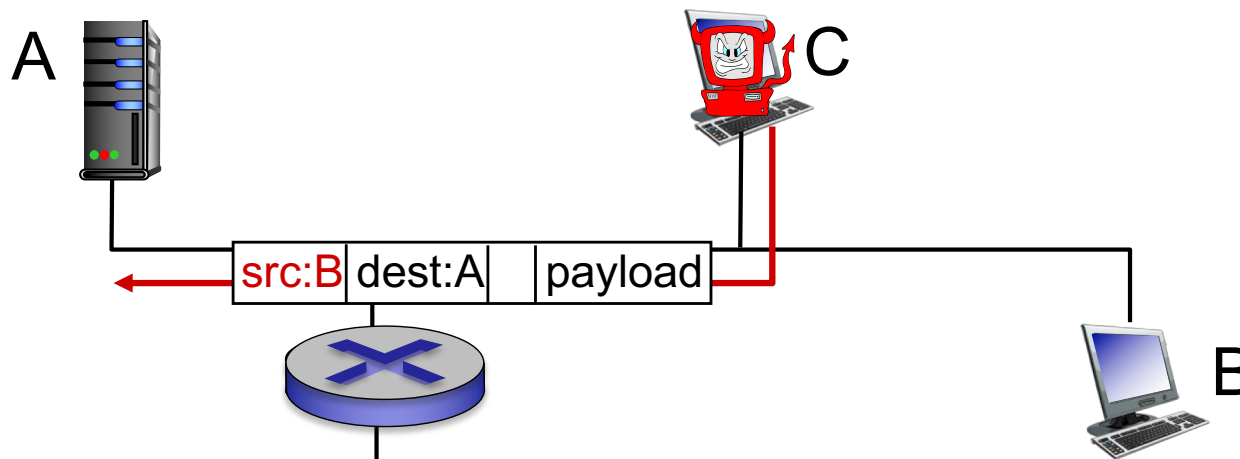target
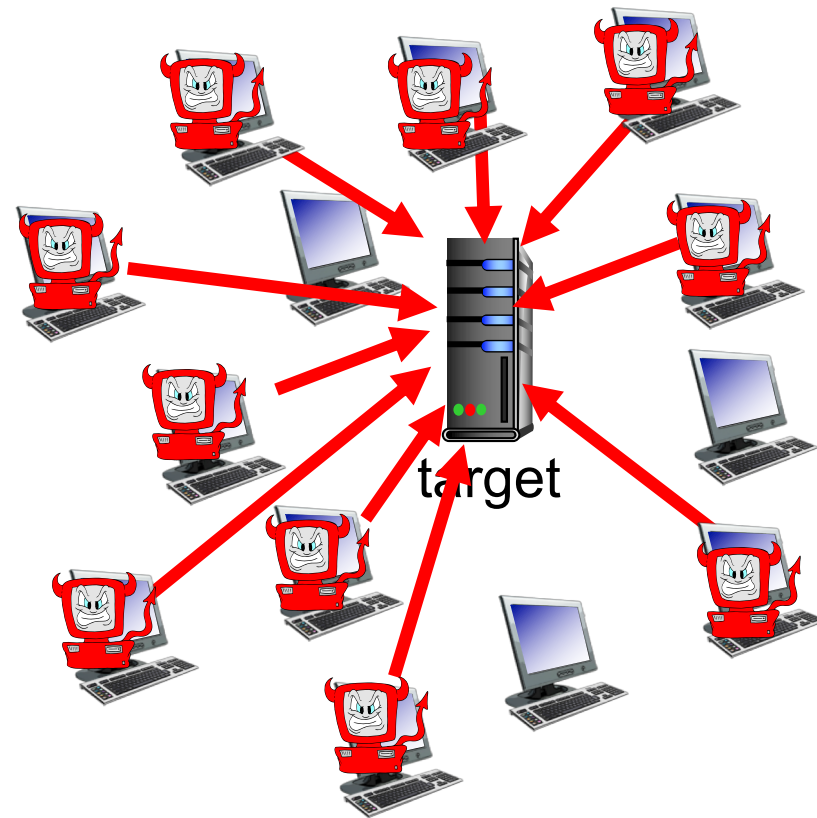
Bad guys can sniff packets

packet "sniffing":

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

| src:B | dest:A | | payload |

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

Bad guys can use fake addresses

IP spoofing: send packet with false source address

… lots more on security (throughout, Chapter 8)

Chapter 1: roadmap

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
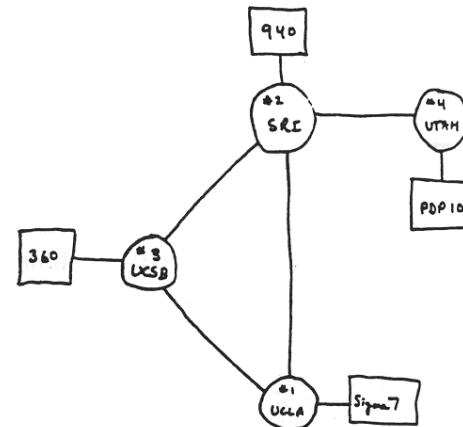
THE ARPA NETWORK

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)

Introduction: summary

covered a "ton" of material!

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

you now have:

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