UNIVERSITEIT VAN AMSTERDAM

Networks and network security

General information and course material to be found on Canvas

Who–when–what?

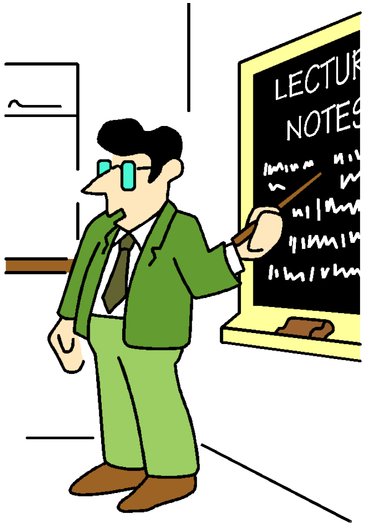

Lecturer:
Dr. Paola Grosso – p.grosso@uva.nl
https://staff.fnwi.uva.nl/p.grosso/
Room: C3.155
Phone: 020. 525.7533

Assistants:

- Johannes Blasier (1st year Master student)
- Frederick Kreuk (3rd year bachelor student)
- Kyrian Maat (1st year Master student)
- Niek van Noort (3rd year bachelor student)
- Cees Portiegies (1st year Master student)
- Lu Zhang (PhD candidate)
- Zi Long Zhu (3rd year bachelor student)

To contact the assistants email: nns18-tas@list.uva.nl

Communication:

- Send email with [NNS18] in the subject!

Who-when-what?

Lectures of 2 hours on Mondays and Thursdays

Labs of 2 hours

- Check the schedule and the group division.
- Check carefully the room assignments.

Note: labs are not mandatory, but highly recommended. We register presence for organizational reasons (joining or splitting groups).

Who-when-what?

An introductory course to computer networks and their security.

Top-down approach, as we dive in more and more in the network architecture as we go.

The 7th edition contains some new material. (Beware if you are using an older edition)

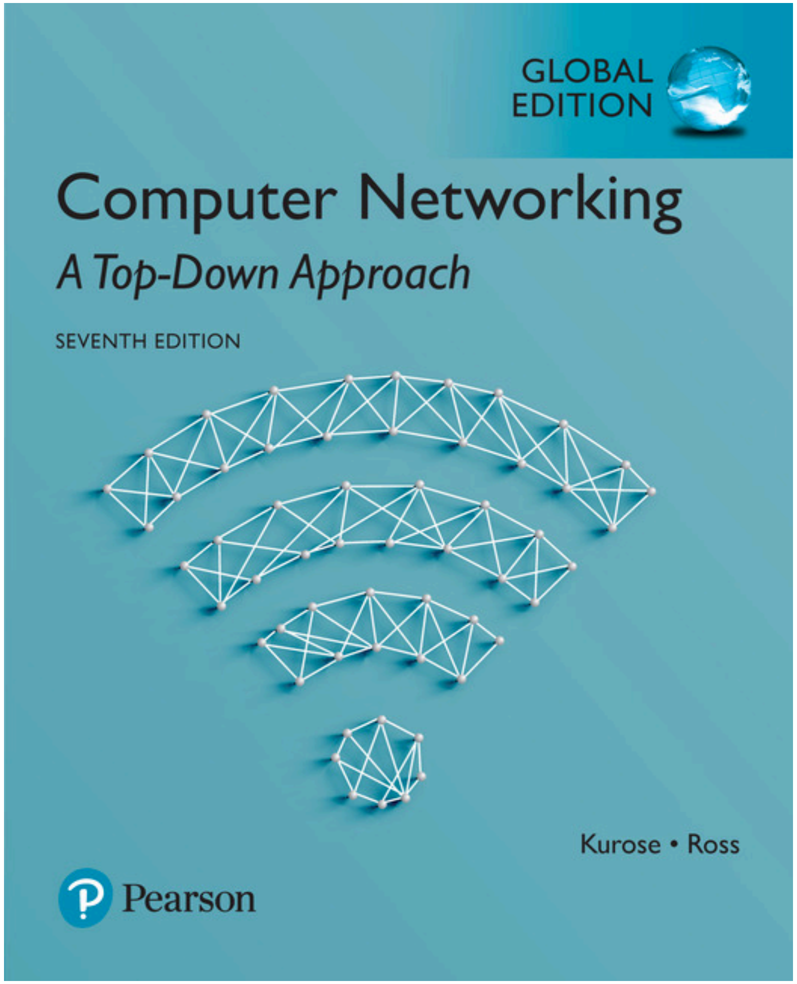

Grades

Final grade composed of:

- 10% from the home assignment
- 40% from labs
- 50% from final exam

Home assignment

In groups of 4/5 you will prepare a short essay (max 2 A4 pages) and a short presentation on a chosen topic.

Topics will be chosen during fourth week of the course. At this time you will also form the groups.

Grade will be the average of the essay (graded by the lecturer and assistants) and of the presentation (graded by the lecturers, assistants and colleagues).

Essay due before the start of week 7.

Presentations given during the lab sessions of week 7.

Labs

Labs count for 40% of the final grade.
Labs give points.

Plus some bonus points to be added at the end.

Some labs are done individually, others in group.

You choose your partner(s).

Topics related to the lecture content:

Programming
Networking tools

Online after the lecture (you can start preparing at home)
Deadlines 11:59pm Mondays and Thursdays

Labs have hard deadlines
A missed lab is 0 (zero) points.
You don't get to repeat it!

Not INF students? Beware….

- You need to be able to program in Python already to do the labs.

This is not the course where you learn to program.

- You need to have labs working on Linux.

See:

https://student.uva.nl/inc/content/az/laptop-minimumeisen/byod/minimumeisen-laptop-informatica.html

Written exam

Written exam count for 50% of the final grade.

Closed book exam with multiple choice, open questions and exercises.

Previous exams to exercise are available online on Canvas.

Exam date: Wednesday October 24 9am-11am

Retake date: Thursday December 20 6pm-9pm (time might change)

Passing

Two conditions to pass (it is an AND!):

Lab grade >= 5.0

Otherwise:
- Will repeat next year

Exam grade >= 5.5

Otherwise:
- ‘herkansing’

No condition set for the home assignment.

If lab grade >= 5.0 and exam >=5.5 but the total grade is not >= 5.5 we will evaluate each case individually and decide repeat course or provide extra ‘herkansing’

Warning!

- I will not tolerate any kind of plagiarism/cheating/copying.
  - Unless specified otherwise the submissions are INDIVIDUAL.

- All submitted labs (text and code) are checked against plagiarism.
- Any improper conduct will be reported to the Exam Commission.

Why plagiarism is bad?

… it is t mostly about your *integrity*!

… it disrespects your peers that have done the work alone

… it frauds the faculty checking your work!

UNIVERSITEIT VAN AMSTERDAM

Lecture 01: Introduction (I)

Internet history and architecture
Network performance

Why this course?

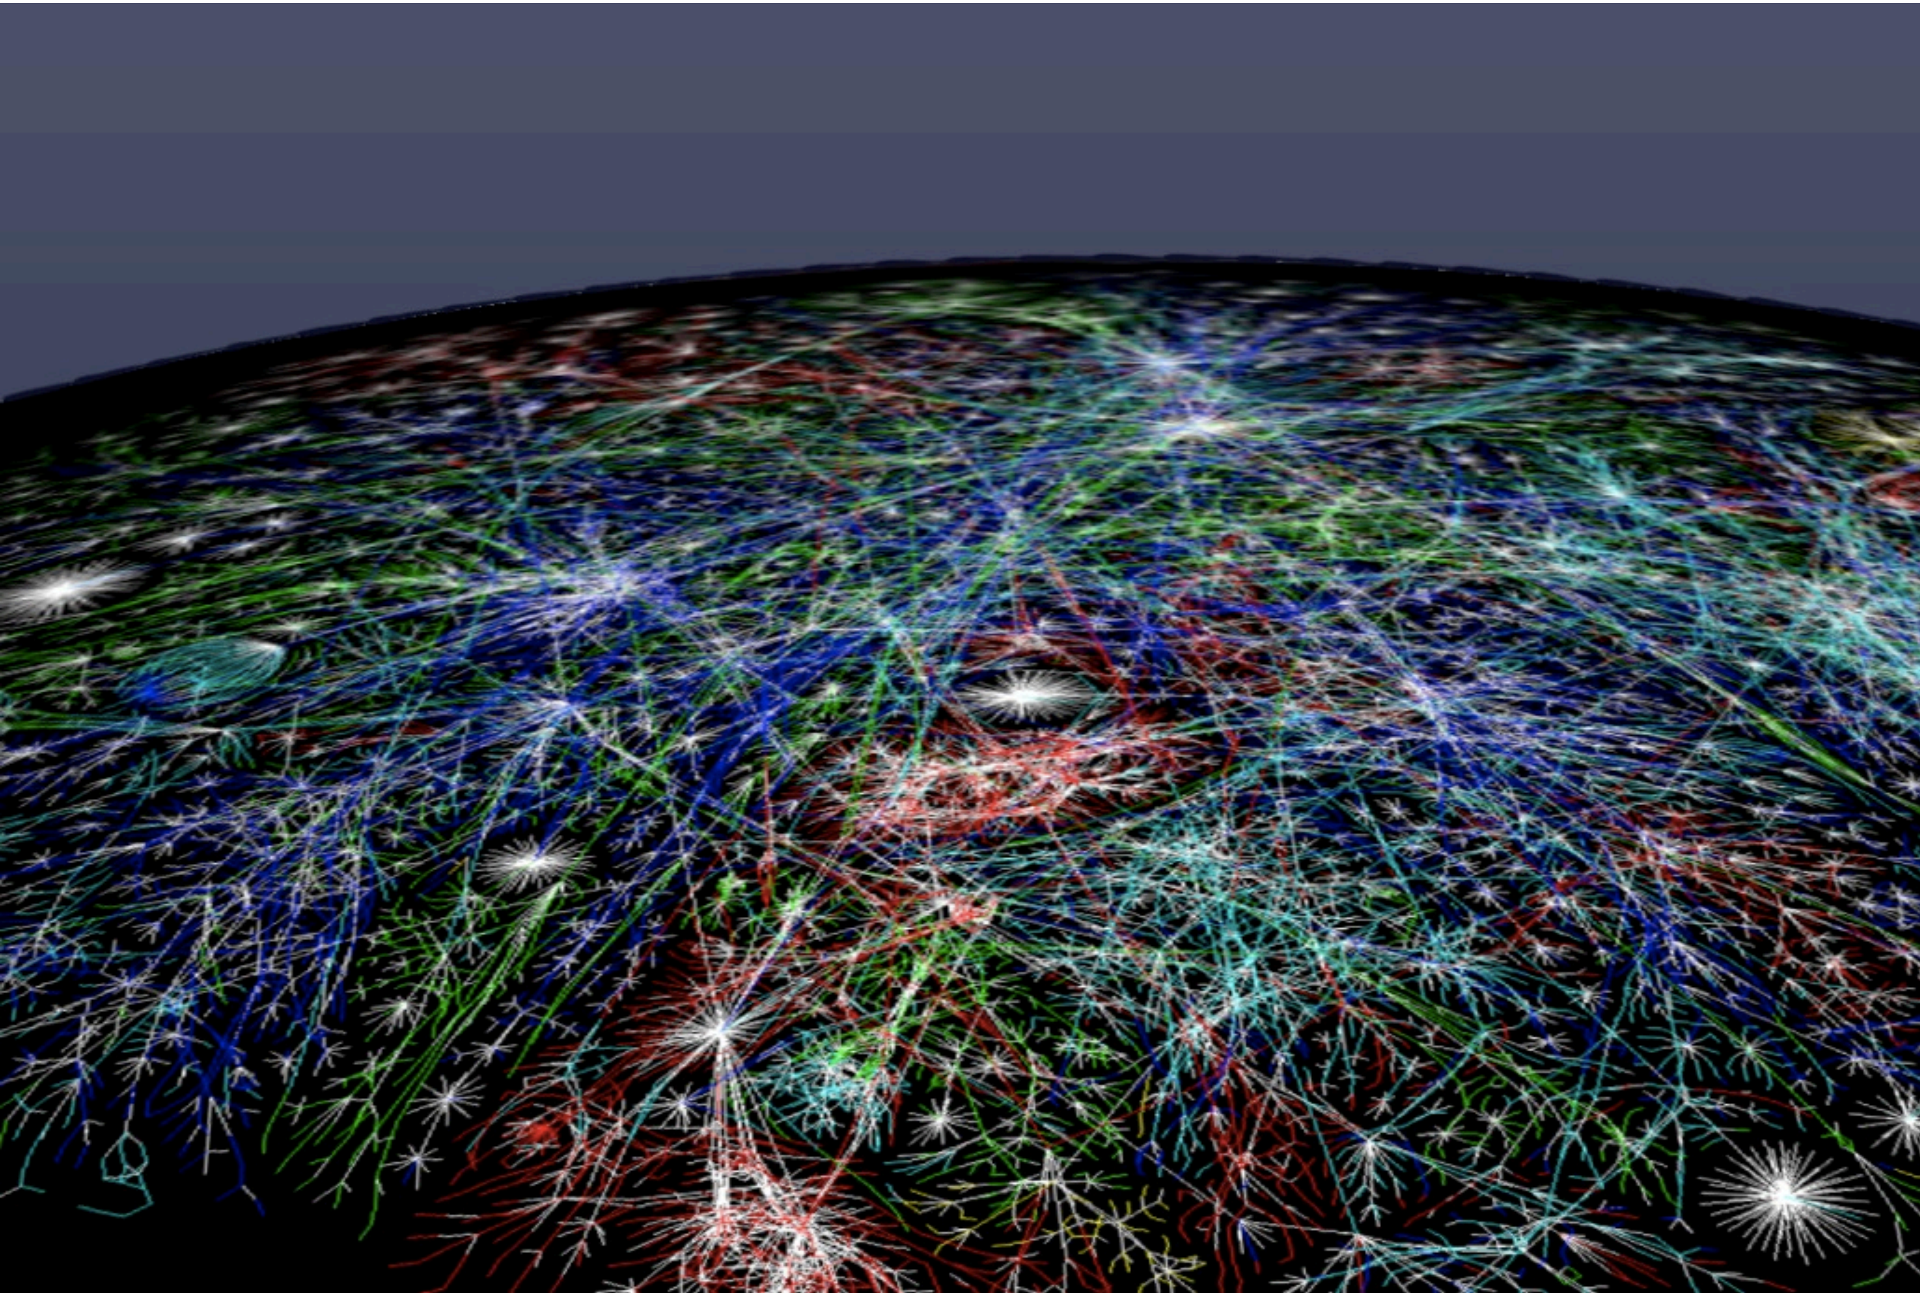

One second on the Internet

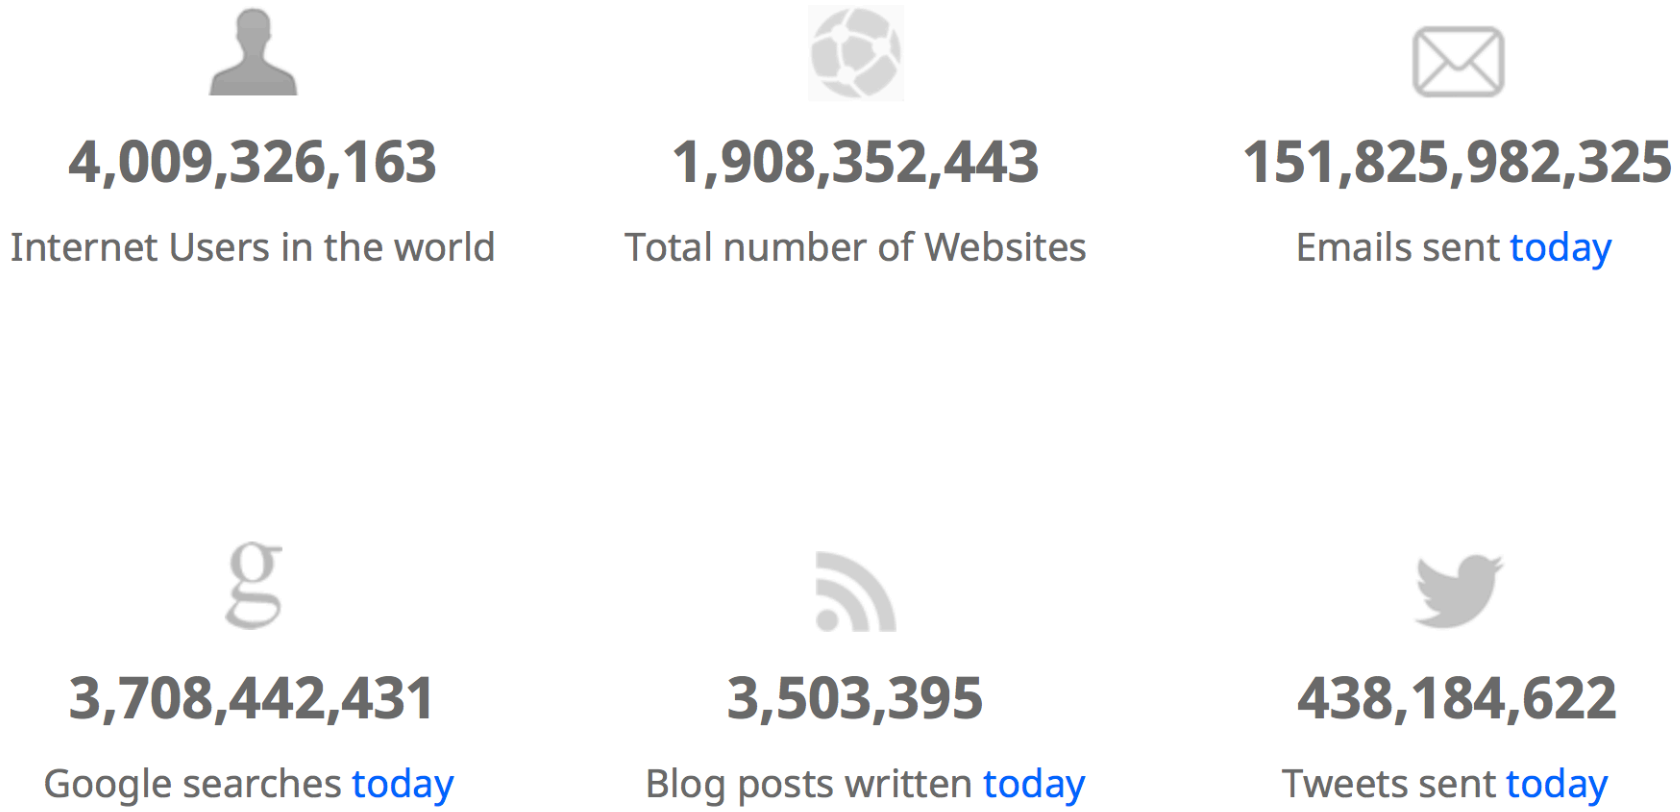

4,009,326,163

Internet Users in the world

1,908,352,443

Total number of Websites

151,825,982,325

Emails sent today

3,708,442,431

Google searches today

3,503,395

Blog posts written today

438,184,622

Tweets sent today

The Internet

The Internet is not the WWW (web).

The Internet connects computers, the web connects information (an information sharing model running above the Internet).

The Internet evolves.

The type of prevalent traffic, and its requirement on the network performance, are changing the way we build and operate the Internet.

The Internet grows.

More and more devices are connected to the Internet, not just any longer ‘computers’.

Consumer internet traffic

Internet video traffic will rise from 60 to 75 percent of total consumer internet traffic by 2018, according to estimates by Cisco.

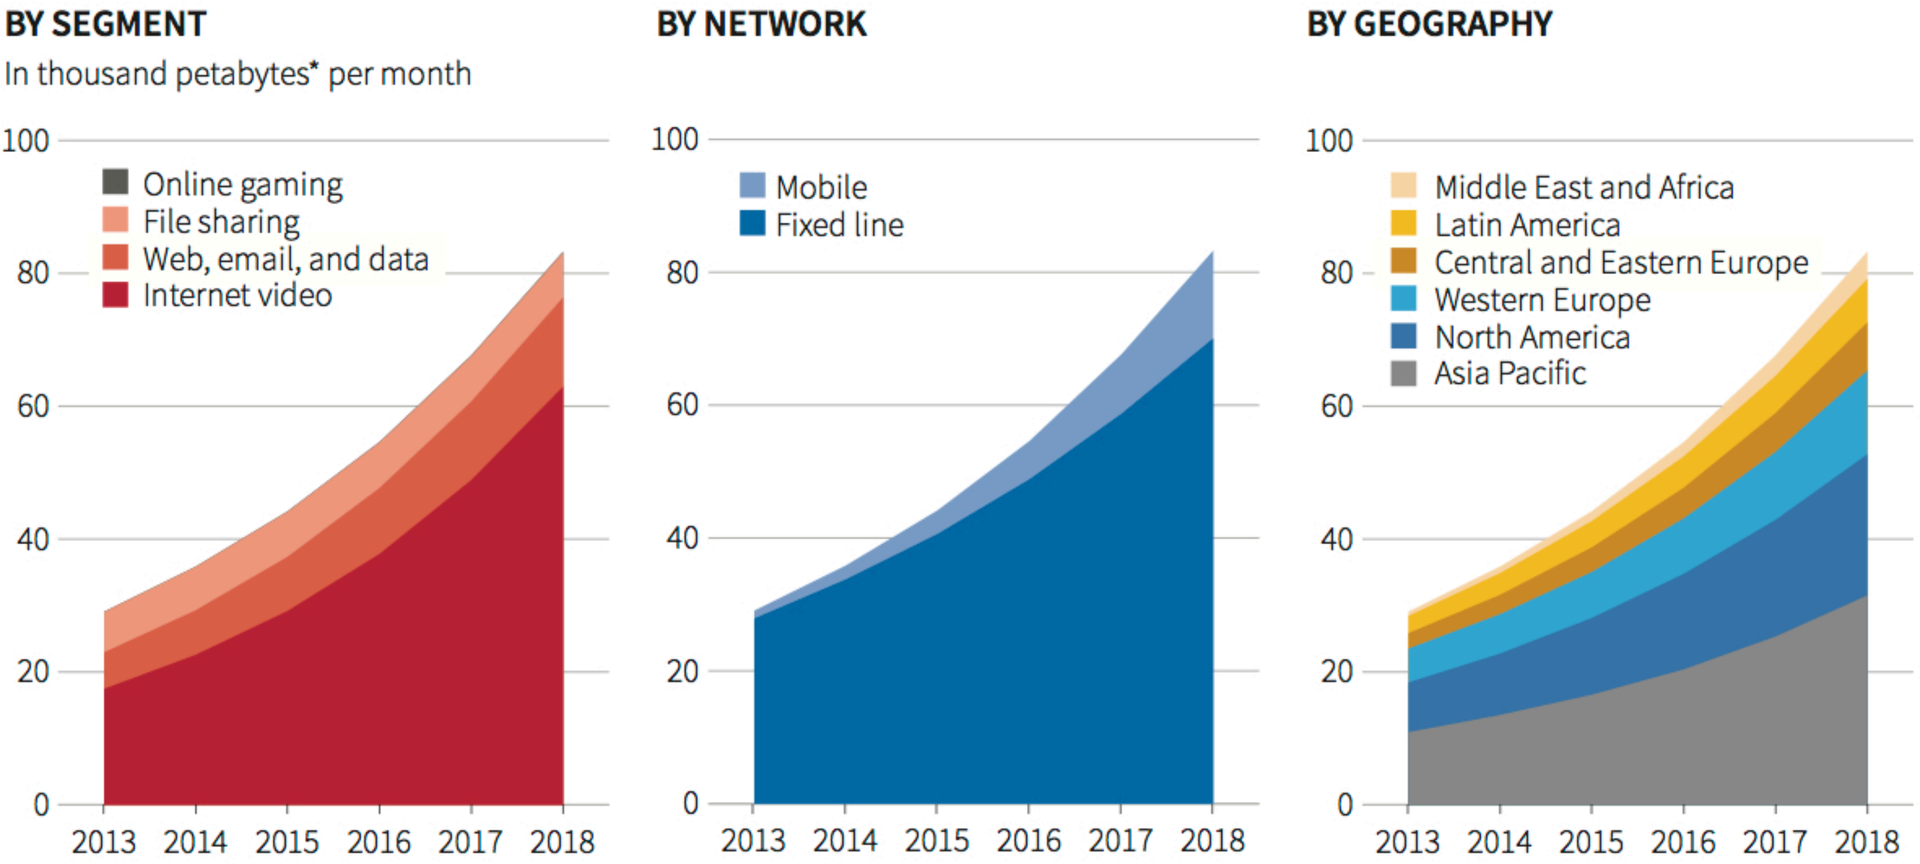

BY SEGMENT

In thousand petabytes* per month

BY NETWORK

BY GEOGRAPHY

Source: Cisco. *Petabyte is equivalent to 1,000 terabytes.

C. Inton, 04/02/2015

REUTERS

The Internet of Things

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

Internet phones

Network security

Main concerns:

- 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"

(D)DoS

Distributed 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

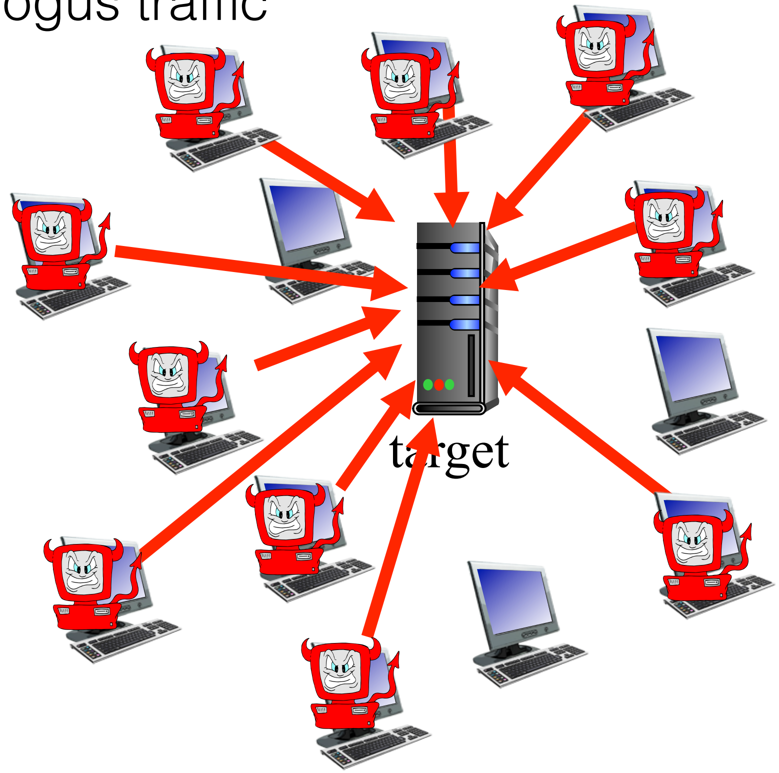

Packets sniffing and IP spoofing

A promiscuous network interface reads/records all packets passing by

Send packets with 'false' source address

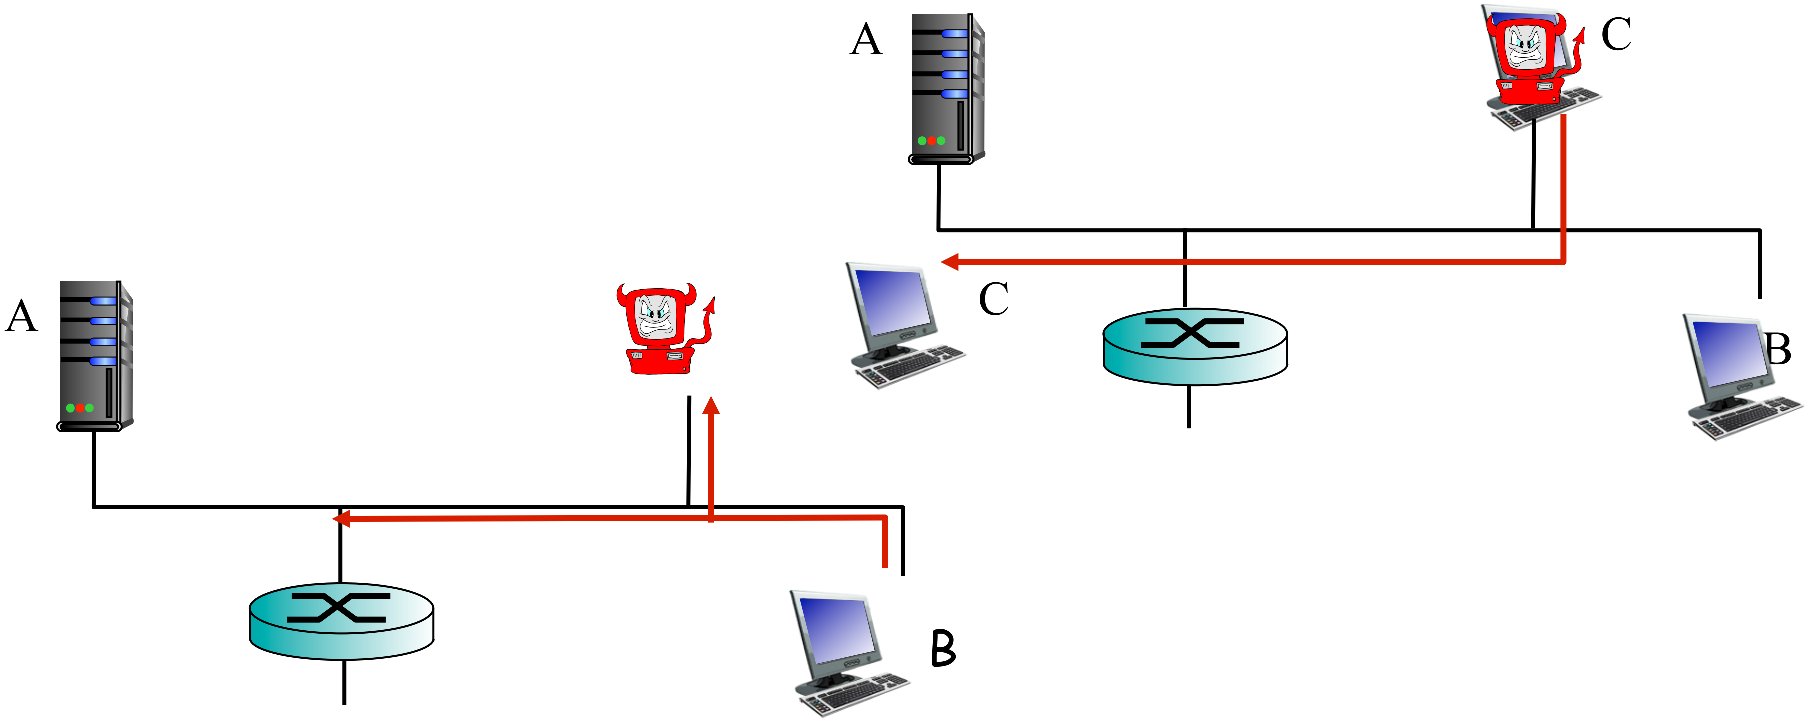

The birth of the Internet

More information?
Watch:
https://vimeo.com/2696386?pg=embed&sec=2696386

Three phases

- Phase 1: 1960s--original development of **packet switching** principles and the ARPANet (Taylor and Roberts);

- Phase 2: 1970s--Development of **internetworking** and extended the packet switching concept to radio and satellite networks (Kahn);

- Phase 3: 1980s and 1990s--**shift**
  - From defense-based sponsorship of the network,
  - To the National Science Foundation and its use to interconnect university-based researchers;
  - To private sector management of the Internet.

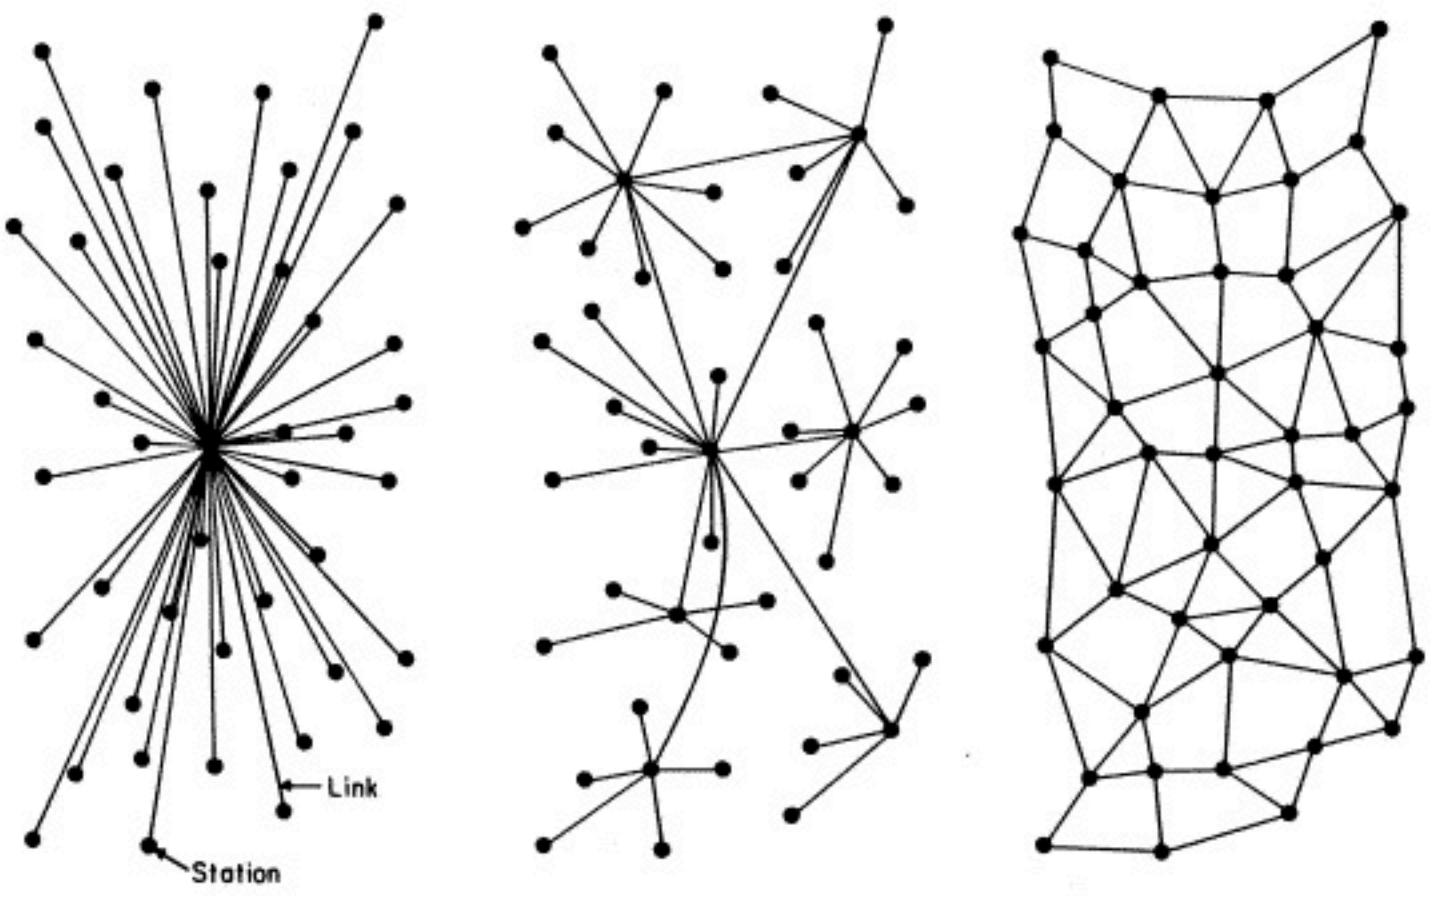

Source: Paul Baran

Early packet-switching principles

- 1964: Paul Baran and Donald Davies packet-switching in military nets

- 1969: the four node ARPANET with IMP (Interface Message Processors), i.e. first generation routers)

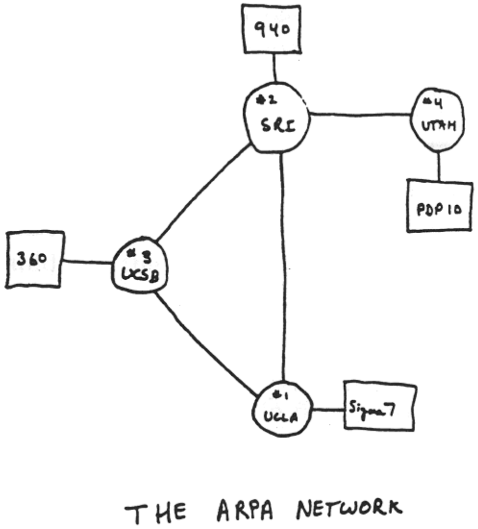

New nets and internetworking

- 1970: ALOHAnet satellite network in Hawaii
- 1974: Cerf and Kahn - architecture for interconnecting networks:
  - minimalism, autonomy - no internal changes required to interconnect networks
  - best effort service model
  - stateless routers
  - decentralized control
- 1979: ARPAnet has 200 nodes

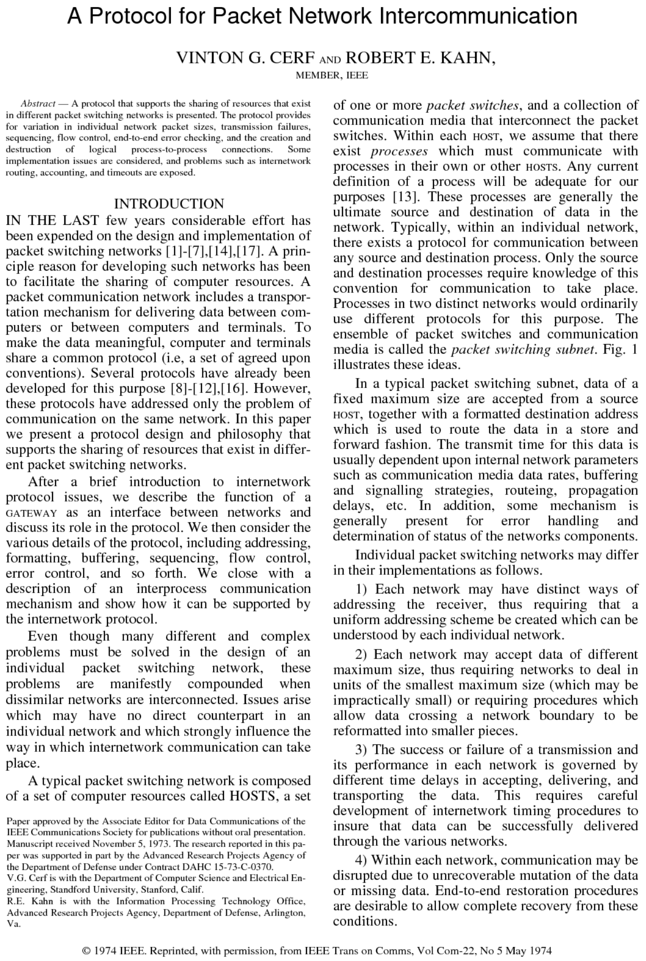

A Protocol for Packet Network Intercommunication

VINTON G. CERF AND ROBERT E. KAHN,
MEMBER, IEEE

Abstract — A protocol that supports the sharing of resources that exist in different packet switching networks is presented. The protocol provides for variation in individual network packet sizes, transmission failures, sequencing, flow control, end-to-end error checking, and the creation and destruction of logical process-to-process connections. Some implementation issues are considered, and problems such as internetwork routing, accounting, and timeouts are exposed.

INTRODUCTION

IN THE LAST few years considerable effort has been expended on the design and implementation of packet switching networks [1]-[7],[14],[17]. A principle reason for developing such networks has been to facilitate the sharing of computer resources. A packet communication network includes a transportation mechanism for delivering data between computers or between computers and terminals. To make the data meaningful, computer and terminals share a common protocol (i.e, a set of agreed upon conventions). Several protocols have already been developed for this purpose [8]-[12],[16]. However, these protocols have addressed only the problem of communication on the same network. In this paper we present a protocol design and philosophy that supports the sharing of resources that exist in different packet switching networks.

After a brief introduction to internetwork protocol issues, we describe the function of a GATEWAY as an interface between networks and discuss its role in the protocol. We then consider the various details of the protocol, including addressing, formatting, buffering, sequencing, flow control, error control, and so forth. We close with a description of an interprocess communication mechanism and show how it can be supported by the internetwork protocol.

Even though many different and complex problems must be solved in the design of an individual packet switching network, these problems are manifestly compounded when dissimilar networks are interconnected. Issues arise which may have no direct counterpart in an individual network and which strongly influence the way in which internetwork communication can take place.

A typical packet switching network is composed of a set of computer resources called HOSTS, a set of one or more *packet switches*, and a collection of communication media that interconnect the packet switches. Within each HOST, we assume that there exist *processes* which must communicate with processes in their own or other HOSTS. Any current definition of a process will be adequate for our purposes [13]. These processes are generally the ultimate source and destination of data in the network. Typically, within an individual network, there exists a protocol for communication between any source and destination process. Only the source and destination processes require knowledge of this convention for communication to take place. Processes in two distinct networks would ordinarily use different protocols for this purpose. The ensemble of packet switches and communication media is called the *packet switching subnet*. Fig. 1 illustrates these ideas.

In a typical packet switching subnet, data of a fixed maximum size are accepted from a source HOST, together with a formatted destination address which is used to route the data in a store and forward fashion. The transmit time for this data is usually dependent upon internal network parameters such as communication media data rates, buffering and signalling strategies, routeing, propagation delays, etc. In addition, some mechanism is generally present for error handling and determination of status of the networks components.

Individual packet switching networks may differ in their implementations as follows.

1) Each network may have distinct ways of addressing the receiver, thus requiring that a uniform addressing scheme be created which can be understood by each individual network.

2) Each network may accept data of different maximum size, thus requiring networks to deal in units of the smallest maximum size (which may be impractically small) or requiring procedures which allow data crossing a network boundary to be reformatted into smaller pieces.

3) The success or failure of a transmission and its performance in each network is governed by different time delays in accepting, delivering, and transporting the data. This requires careful development of internetwork timing procedures to insure that data can be successfully delivered through the various networks.

4) Within each network, communication may be disrupted due to unrecoverable mutation of the data or missing data. End-to-end restoration procedures are desirable to allow complete recovery from these conditions.

Paper approved by the Associate Editor for Data Communications of the IEEE Communications Society for publications without oral presentation. Manuscript received November 5, 1973. The research reported in this paper was supported in part by the Advanced Research Projects Agency of the Department of Defense under Contract DAHC 15-73-C-0370.

V.G. Cerf is with the Department of Computer Science and Electrical Engineering, Standford University, Stanford, Calif.
R.E. Kahn is with the Information Processing Technology Office, Advanced Research Projects Agency, Department of Defense, Arlington, Va.

Proliferation time

- 1983: deployment of TCP/IP

- 1990: 100,000 hosts connected to confederation of networks

- early 1990s: Web
  - hypertext [Bush 1945, Nelson 1960’s]
  - HTML, HTTP: Berners-Lee
  - 1994: Mosaic, later Netscape
  - late 1990’s: commercialization of the Web

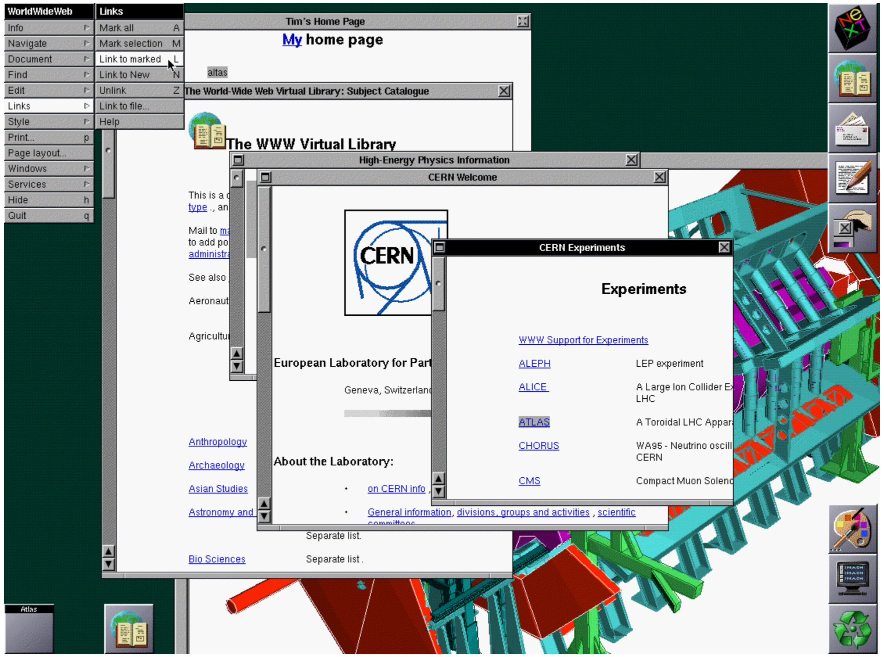

The world as we know it….

- 1991: NSF lifts restrictions on commercial use of NSFnet (decommissioned, 1995)

Current Internet structure

Tiered structure

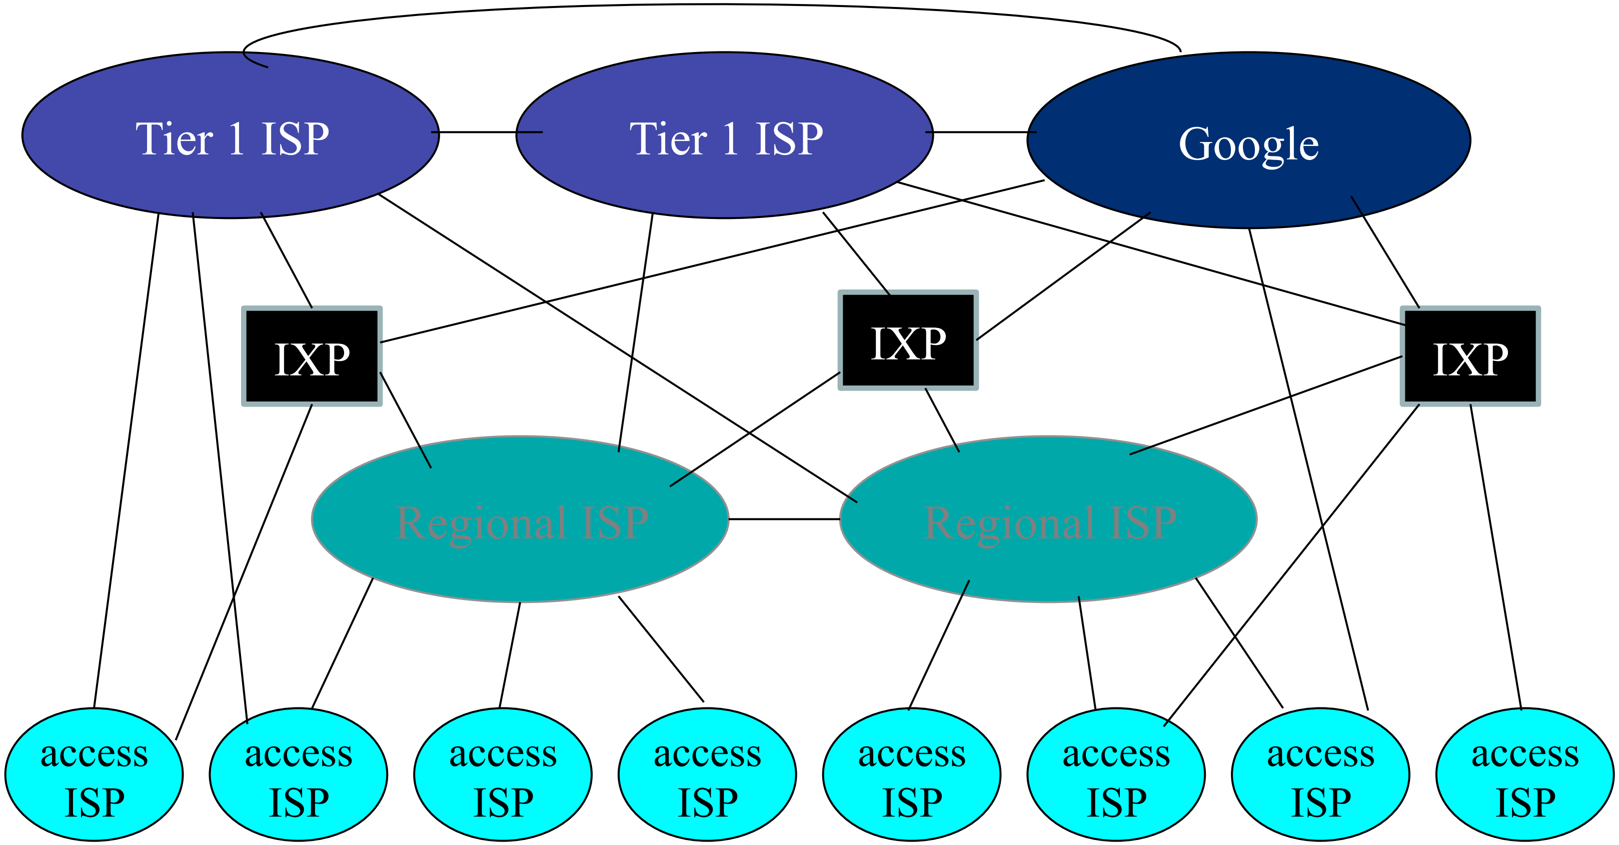

Internet terminology

- millions of connected computing devices:
  - *hosts = end systems*
  - running *network apps*

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

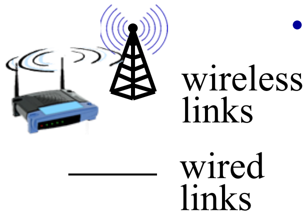

- Packet switches: forward packets (chunks of data)
  - routers and switches

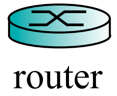

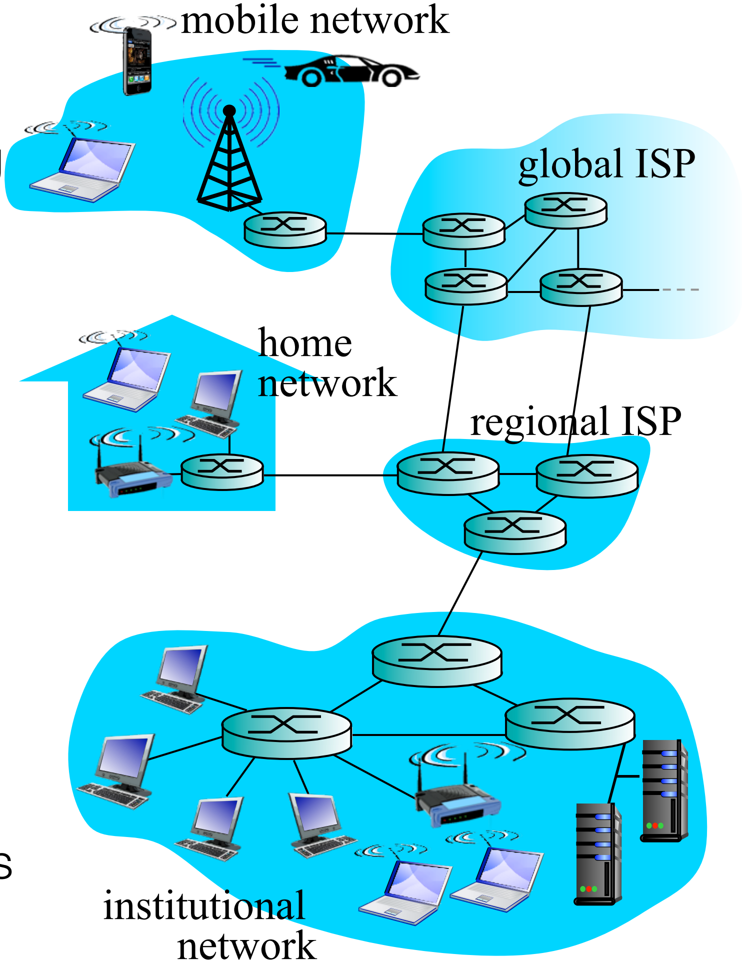

Network of networks

- *Internet:* "network of networks"
  - Interconnected ISPs
- *protocols* control sending, receiving of msgs
  - e.g., TCP, IP, HTTP, Skype, 802.11
- *Internet standards*
  - RFC: Request for comments
  - IETF: Internet Engineering Task Force

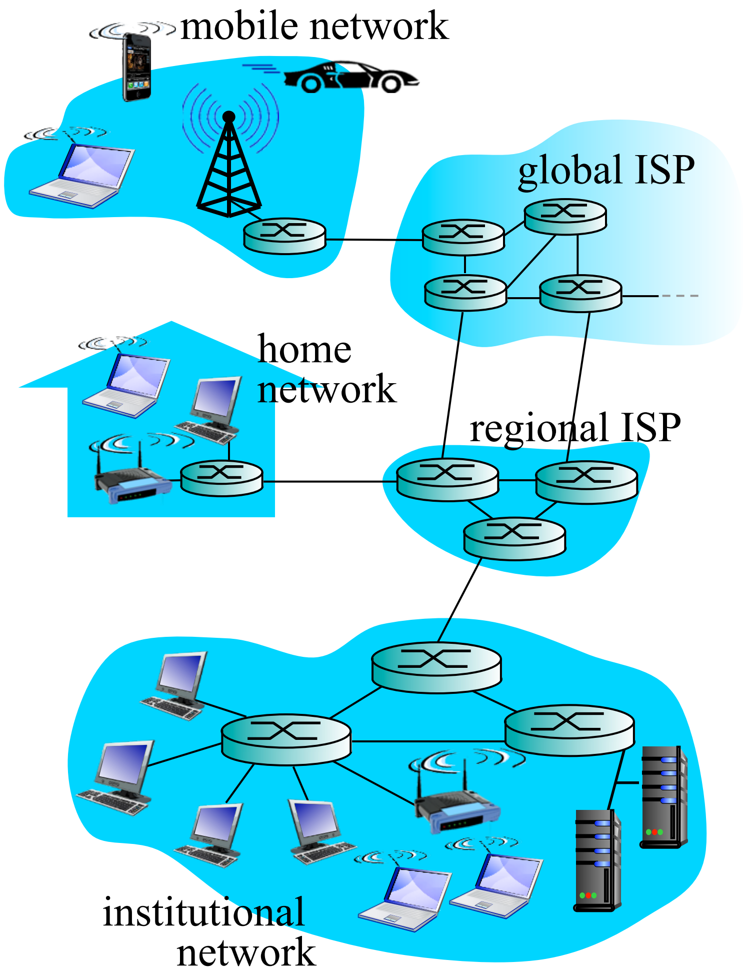

Pause

A closer look

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

2. access networks:
   - Using different physical medias
     - wired, wireless communication links

3. network core:
   - interconnected routers

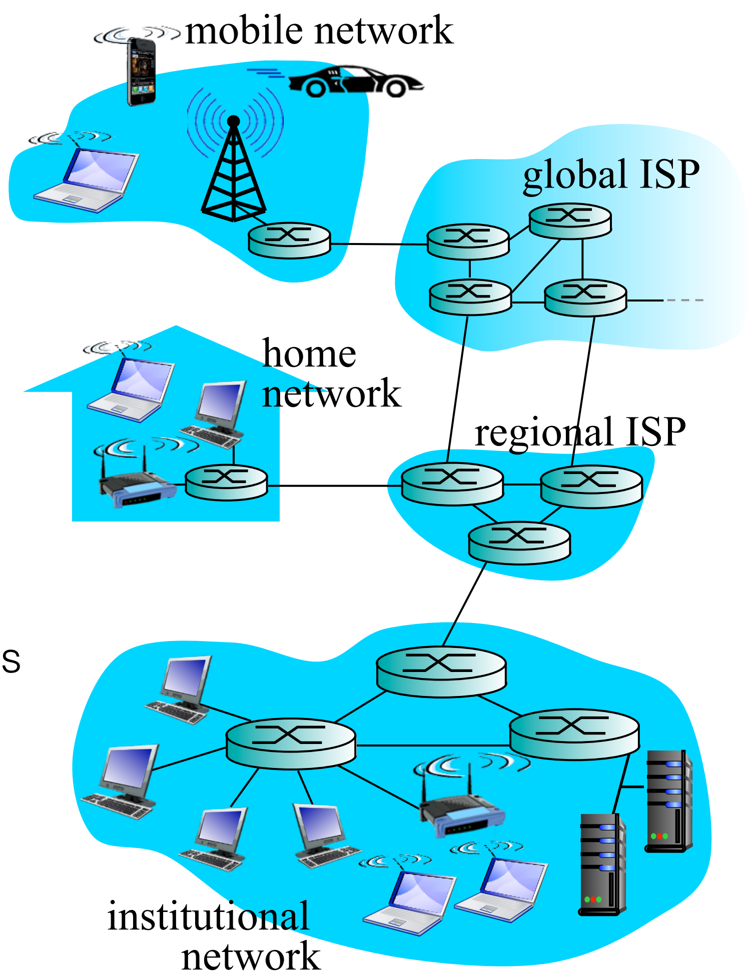

Network edge

End hosts: sends packets into the network

- takes application message
- breaks into smaller chunks, known as packets,
- transmits packet into access network

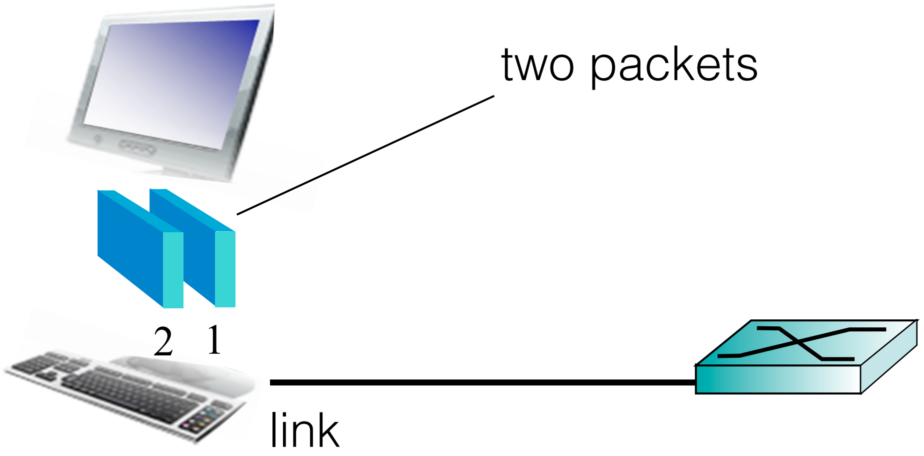

Host/server

Transmission medium

The transmission medium is the physical path between transmitter/ sender and receiver.

Two types of media:

- Guided media
  - Copper twisted pair
  - Copper coaxial cables
  - Optical fibers
- Unguided media
  - Air
  - Water

Physical media: twisted pair, coax

Data is transmitted as electrical pulses

twisted pair (TP)

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

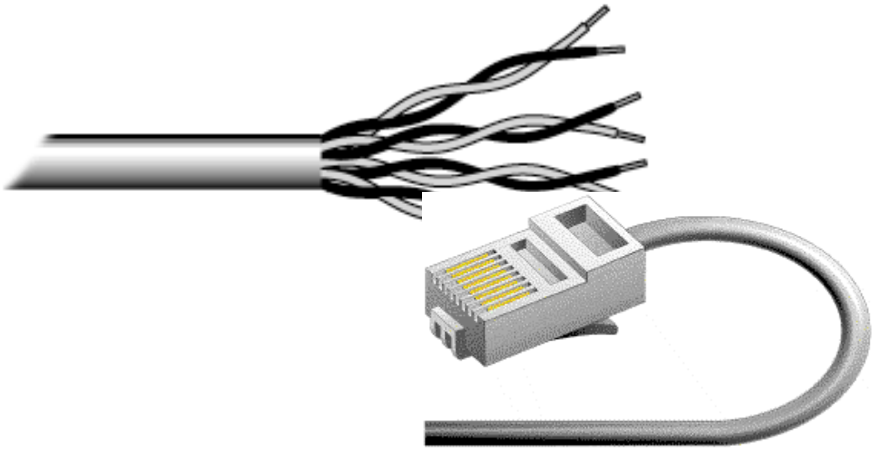

coaxial cable:

- two concentric copper conductors
- bidirectional
- broadband:
  - multiple channels on cable
  - HFC

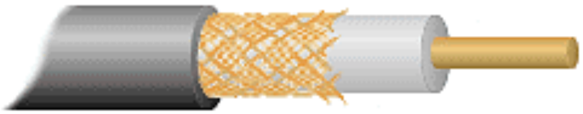

Physical media: fiber

fiber optic cable:

- glass fiber carrying light pulses, each pulse a bit
- high-speed operation:
- high-speed point-to-point transmission (e.g., 10’s-100’s Gbps transmission rate)
- low error rate:

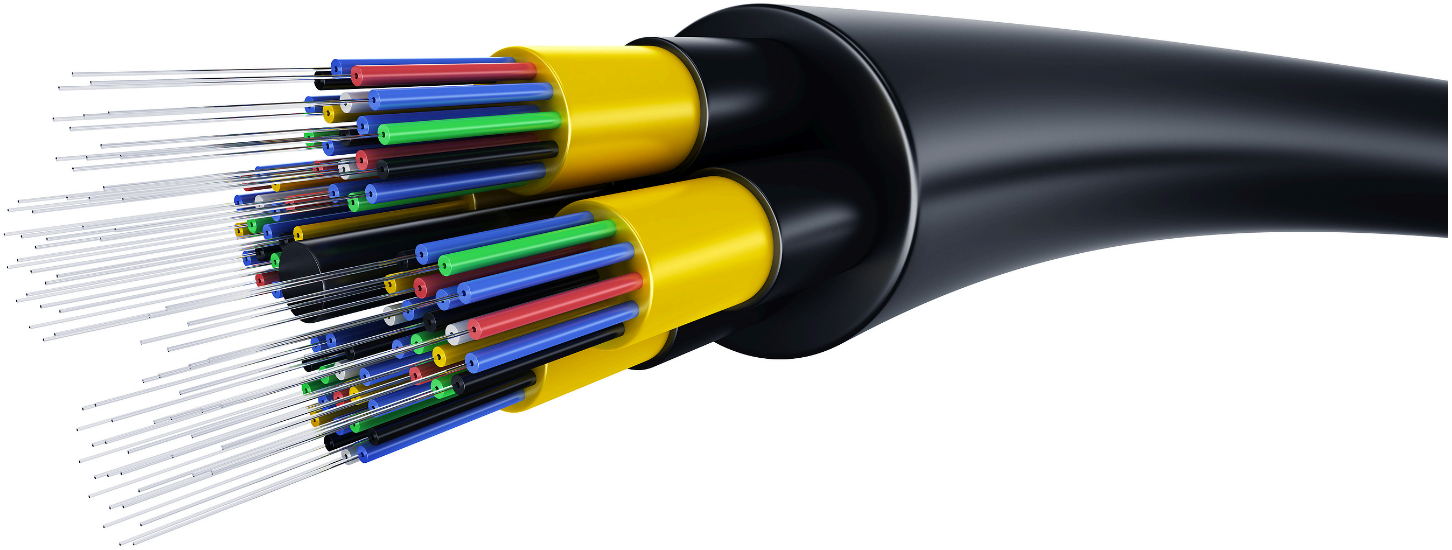

Physical media: radio

- signal carried in electromagnetic spectrum
- no physical “wire”
- bidirectional
- propagation environment effects:
  - reflection
  - obstruction by objects
  - interference

radio link types:

- terrestrial microwave
  - e.g. up to 45 Mbps channels
- LAN (e.g., WiFi)
  - 54 Mbps
- wide-area (e.g., cellular)
  - 4G cellular: ~ 10 Mbps
- satellite
  - Kbps to 45Mbps channel (or multiple smaller channels)

Access networks and physical media

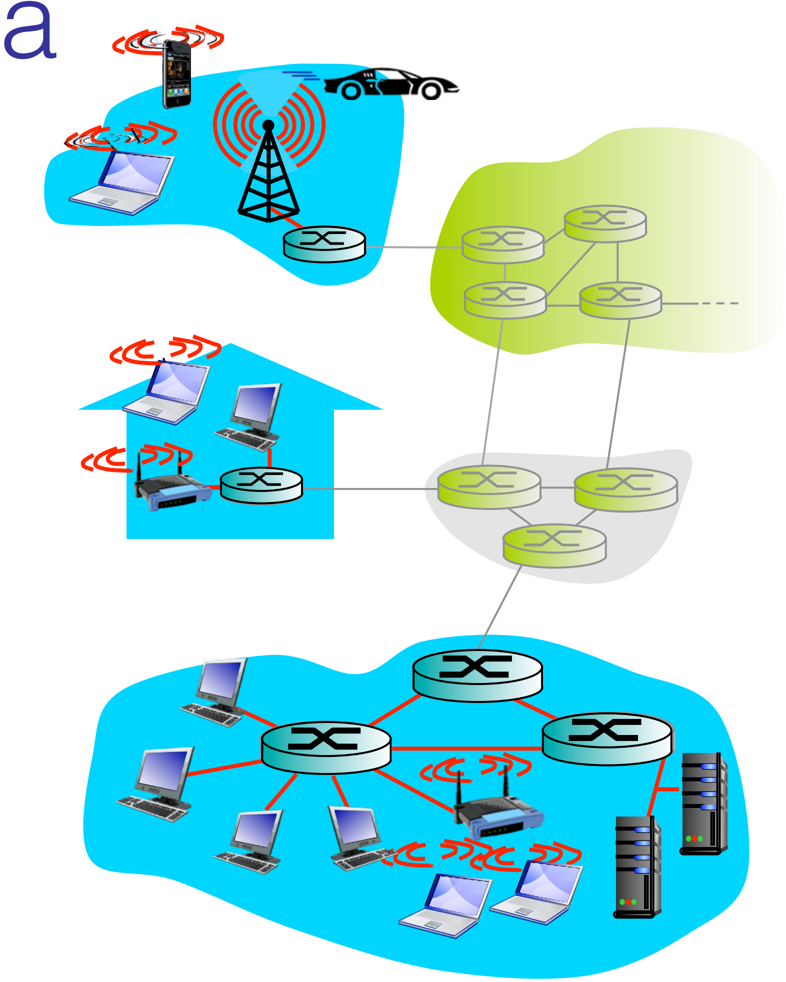

How to connect end systems to edge router?

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

Need to keep in mind:

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

Home network

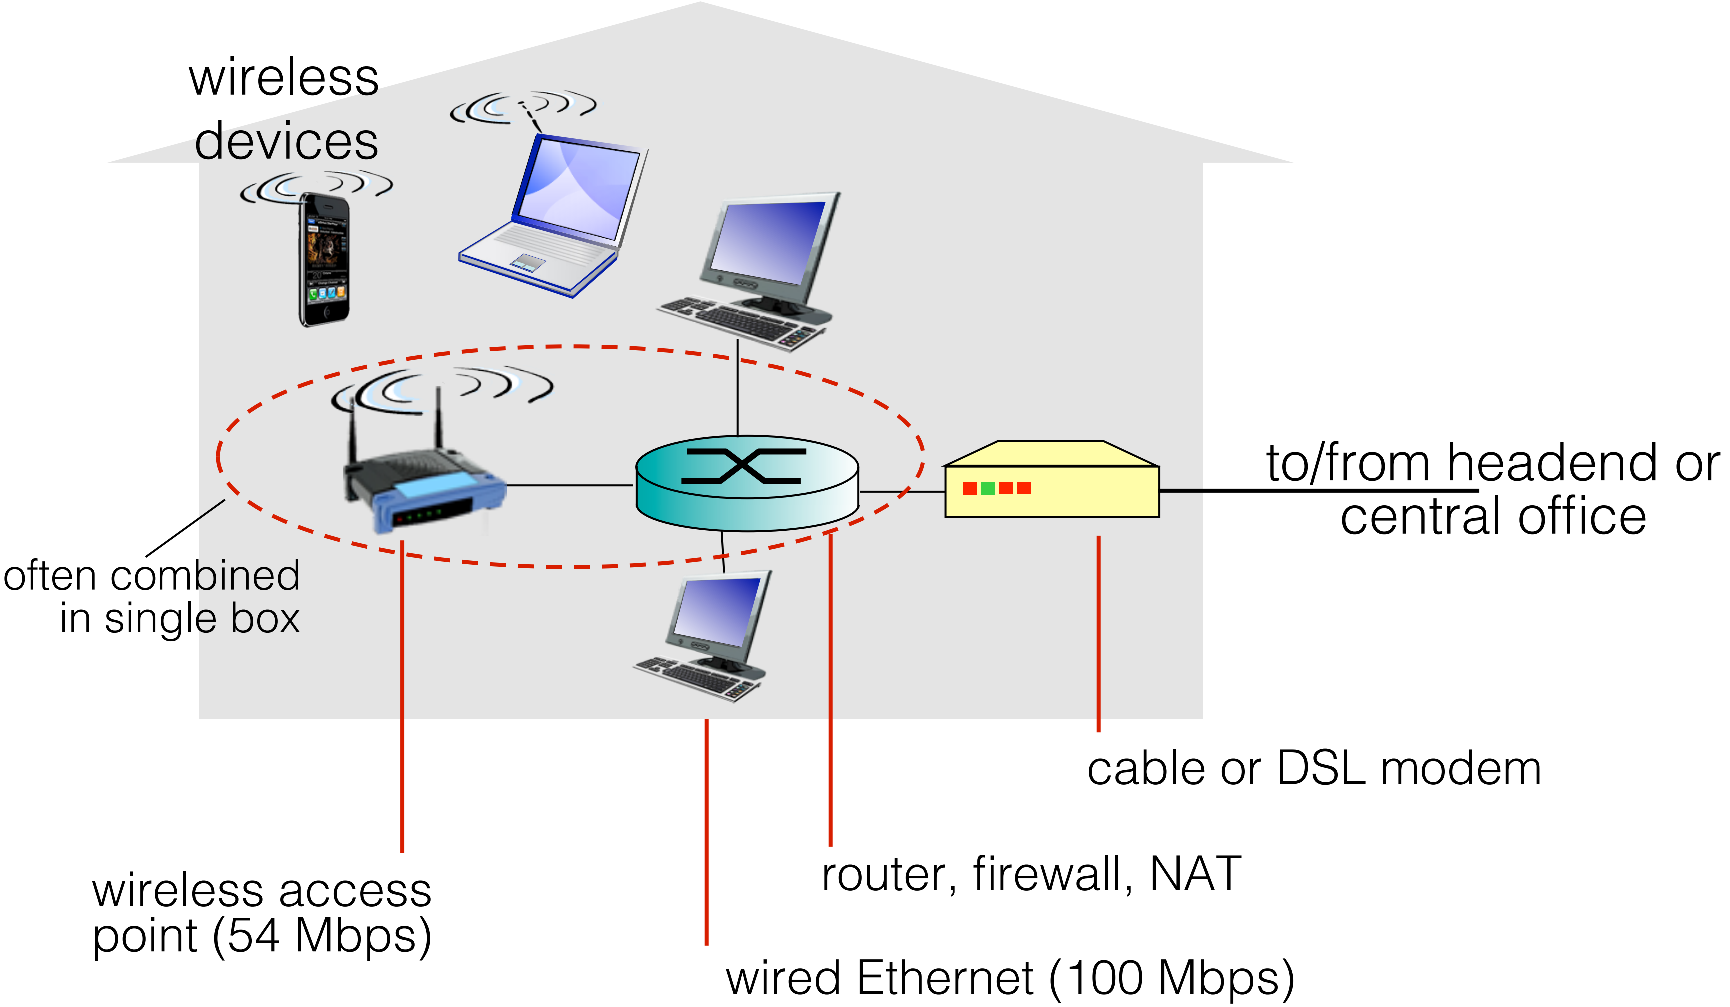
wireless devices
often combined in single box
to/from headend or central office
cable or DSL modem
router, firewall, NAT
wireless access point (54 Mbps)
wired Ethernet (100 Mbps)

Enterprise access networks (Ethernet)

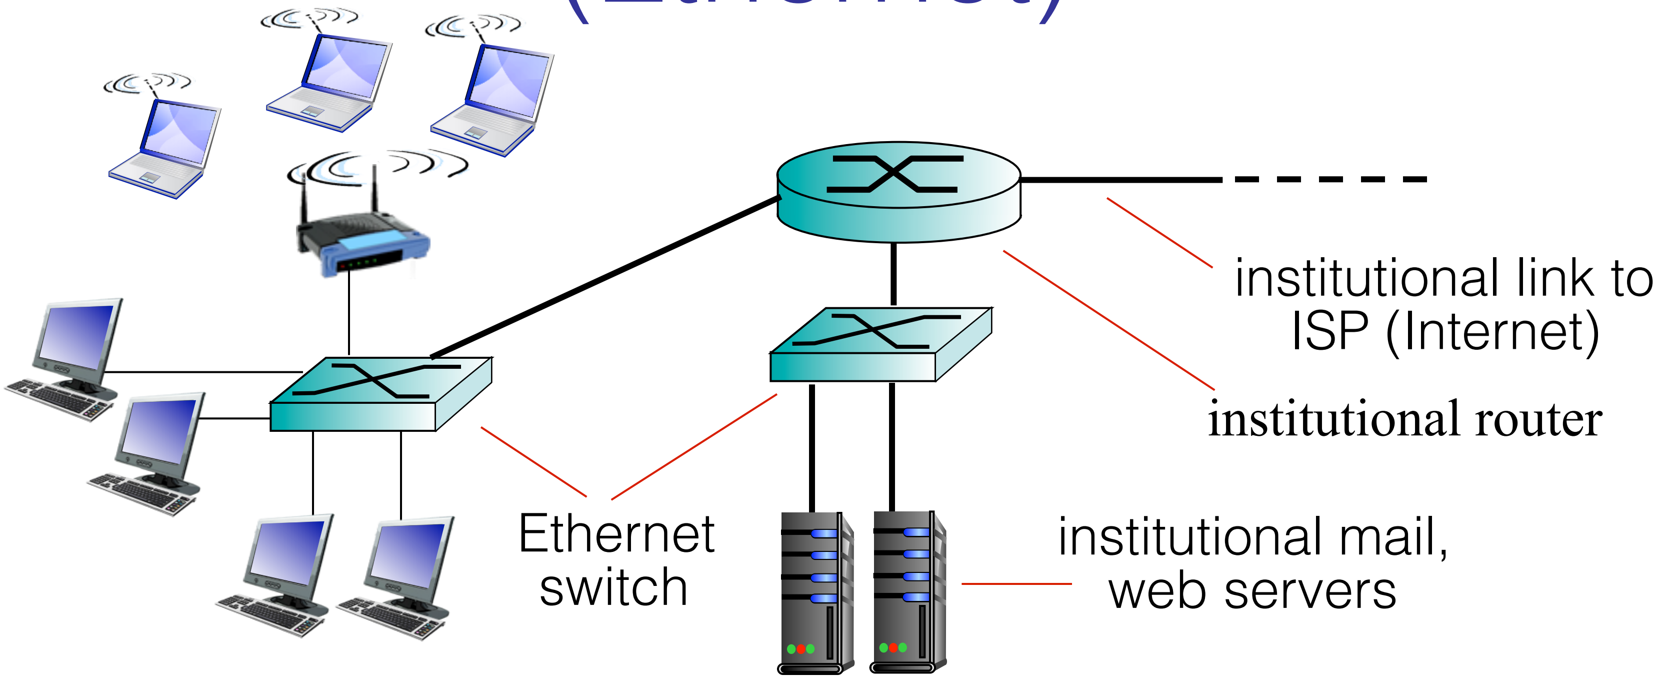
institutional link to ISP (Internet)
institutional router
Ethernet switch
institutional mail, web servers

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

Wireless access networks

wireless LANs:

- within building (100 meters)
- 802.11b/g/n (WiFi)

wide-area wireless access

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

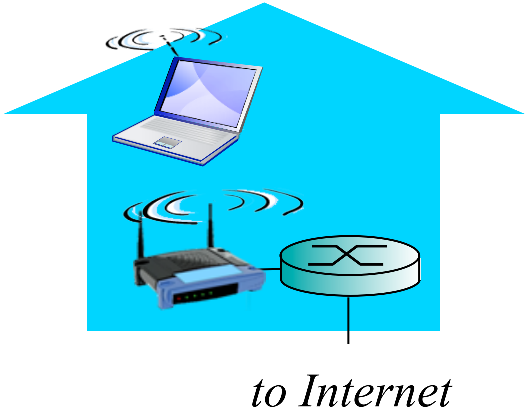

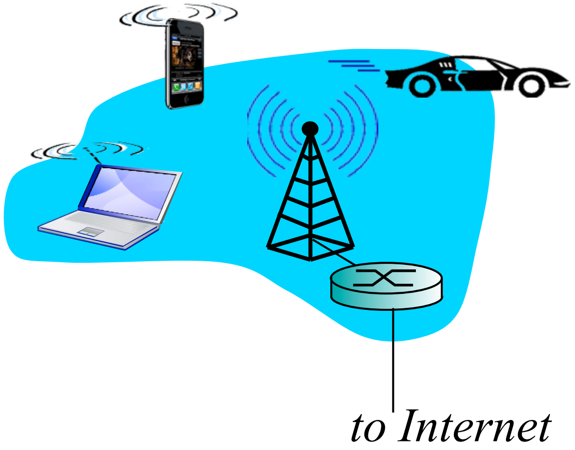

- shared wireless access network connects end system to router
  - via base station aka "access point"

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

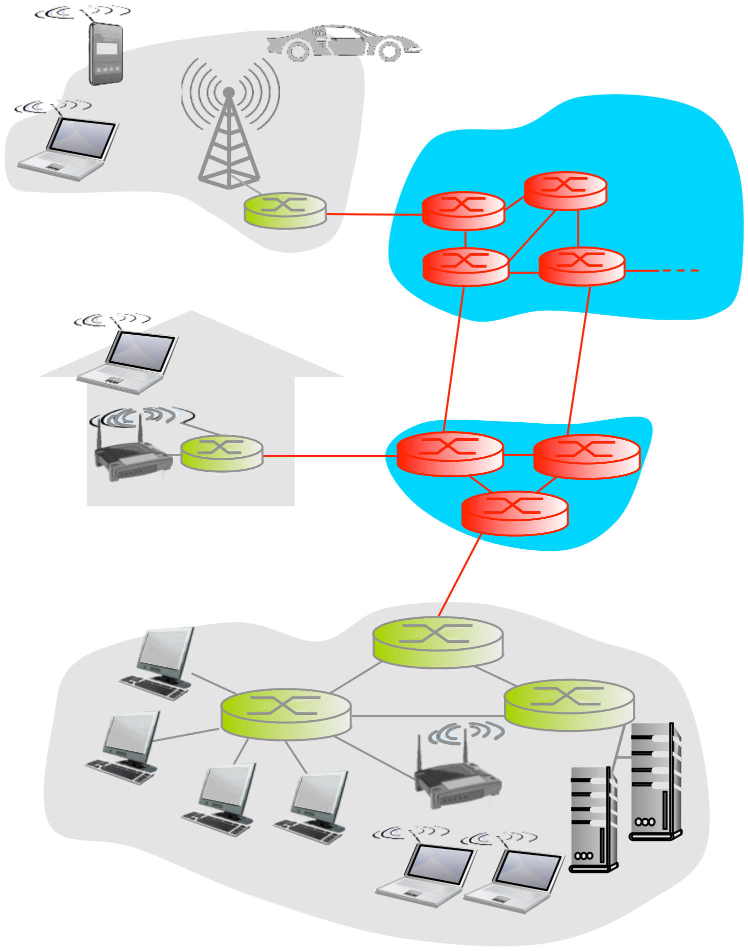

Two key network-core functions

routing:

determines source-destination route taken by packets

- routing algorithms

forwarding:

move packets from router’ s input to appropriate router output

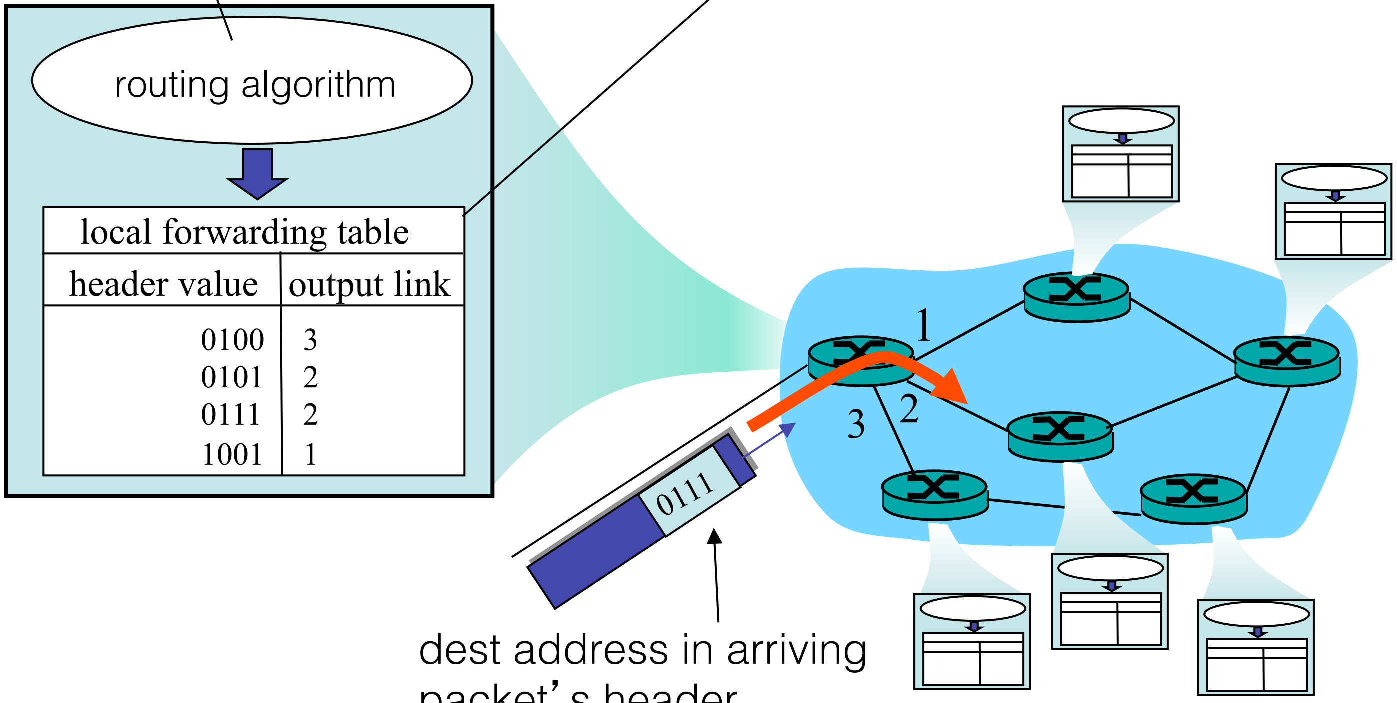

local forwarding table	
header value	output link
0100	3
0101	2
0111	2
1001	1

dest address in arriving packet’ s header

Packet-switching: store-and-forward

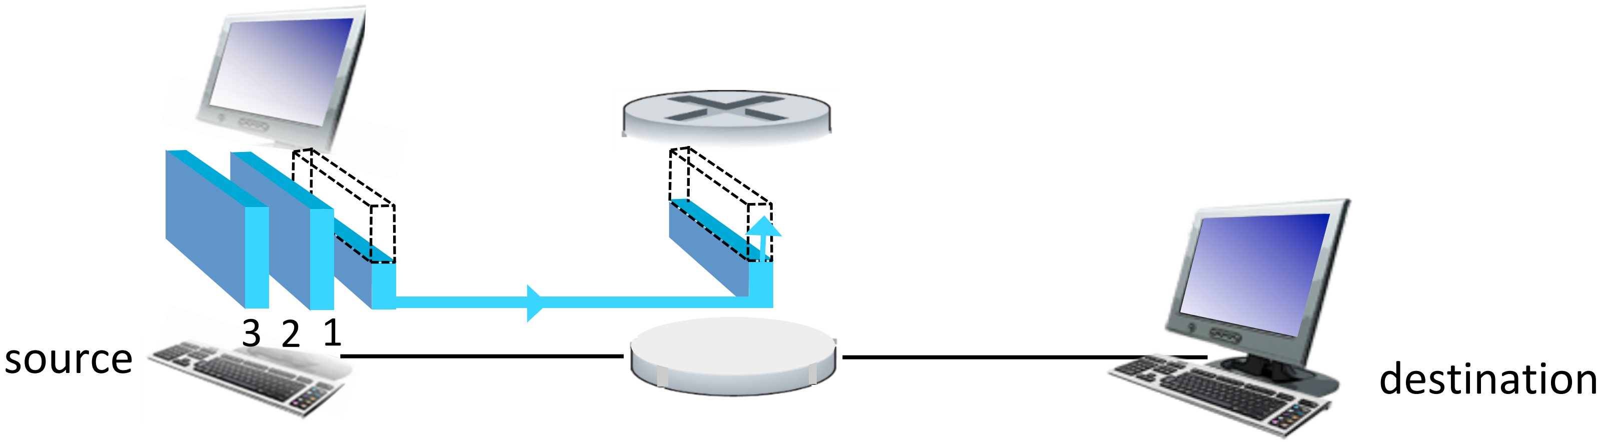

- *Store and forward:* entire packet must arrive at router before it can be transmitted on next link

- *Cut-through:* used to lower the total latency (used expecially in data centers or high frequency trading)

Statistical multiplexing gain

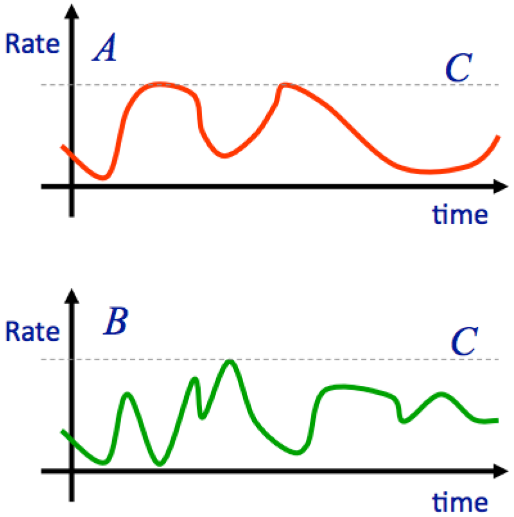

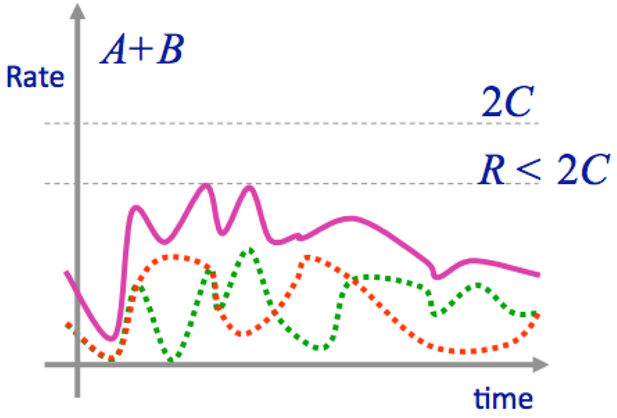

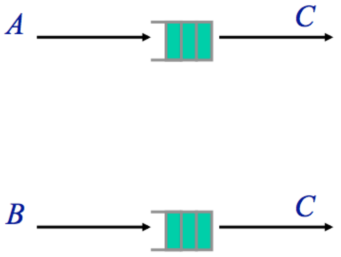

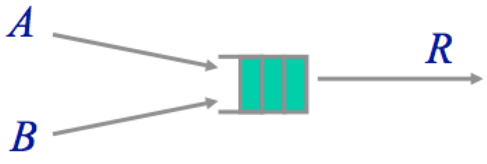

$$Gain = {}^{2C}\!/_{R}$$

Alternative core: circuit switching

End-end resources reserved for "call" between source & dest:

- dedicated resources: no sharing
  - circuit-like (guaranteed) performance
- circuit segment idle if not used by call (no sharing)

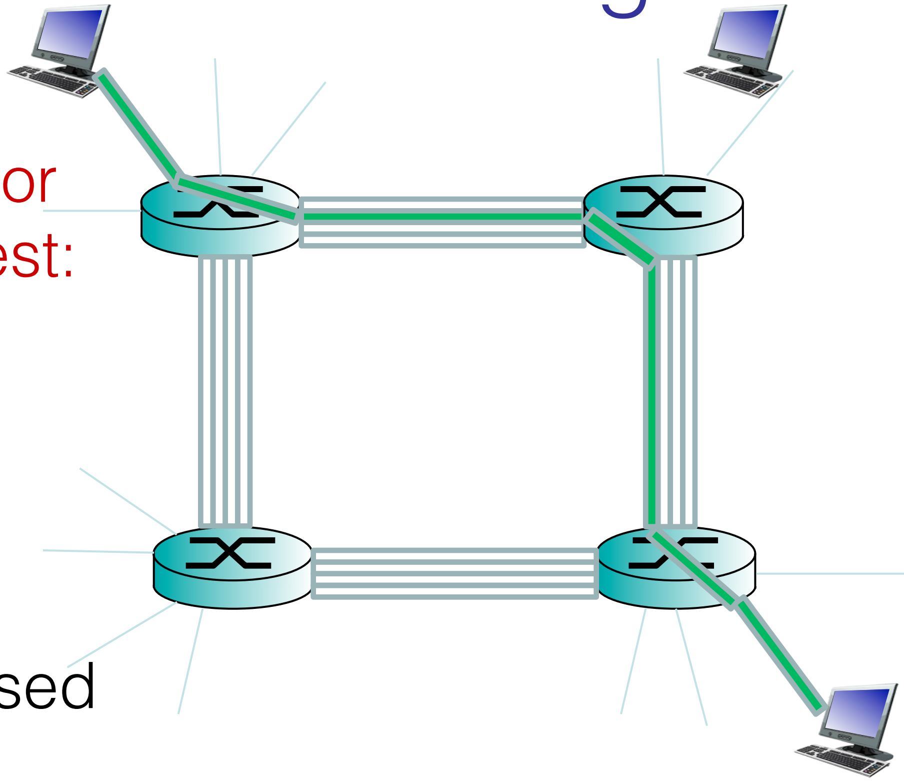

Circuit switching: FDM versus TDM

FDM – Frequency Division Multiplexing

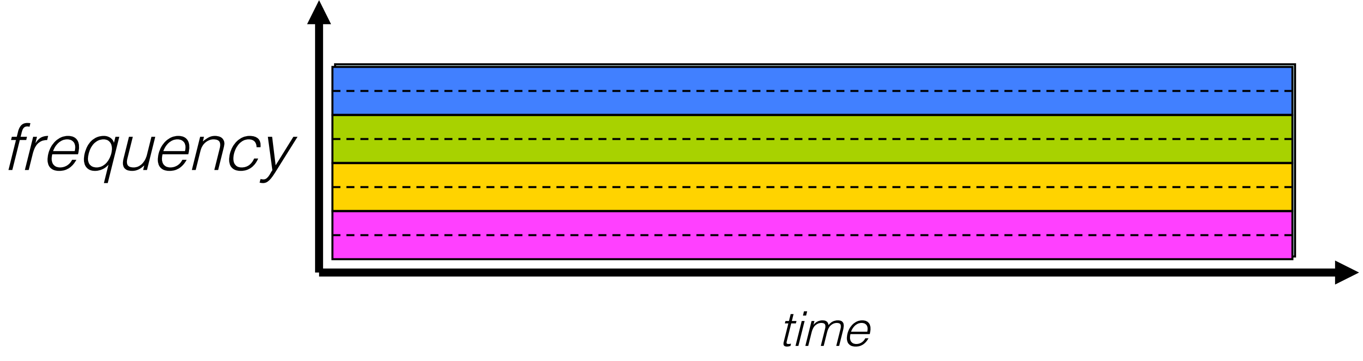

TDM – Time Division Multiplexing

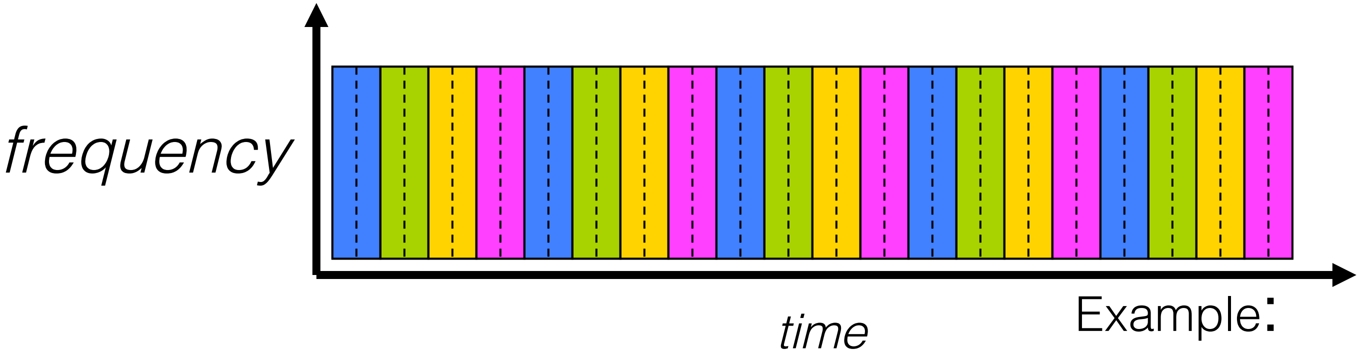

Example:

4 users 🟦 🟩 🟨 🟪

Network performance

What are the functionalities you care about?

You use a computer network, or the Internet, to communicate.

Performance

- Two factors related:
  - Time spent waiting for the data to arrive
  - Throughput

What are the ‘problems’ you can encounter?

- Delays determine/increase time needed to receive the data
- Losses of data reduce the throughput

Delay

How does delay occur?

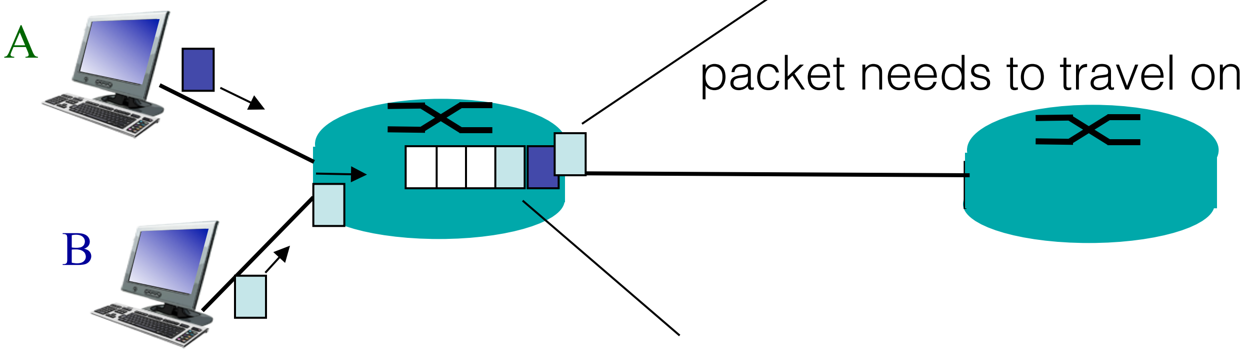

packet needs to be transmitted

packet needs to be processed

packet needs to travel on link

packets needs to wait its turn queueing

Four sources of packet delays

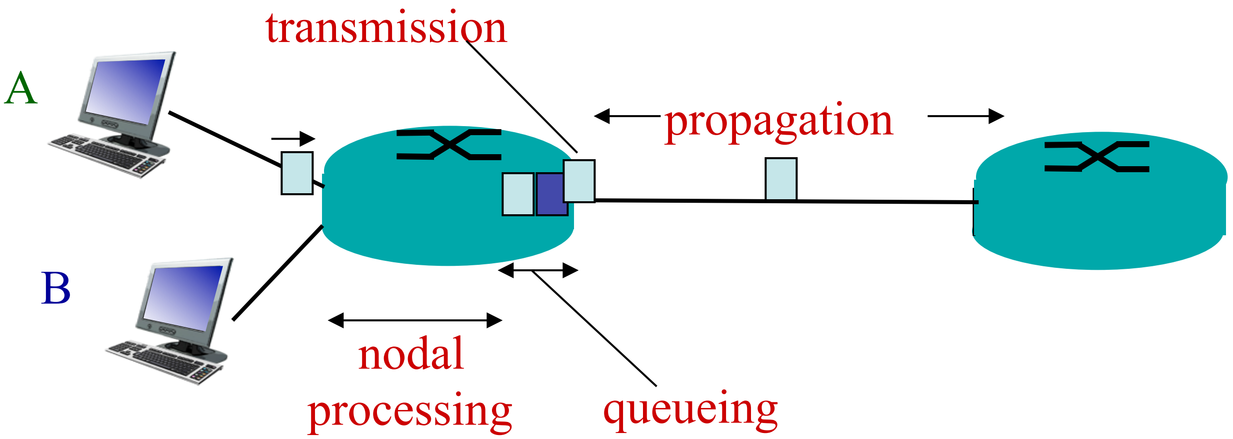

transmission

A

propagation

B

nodal processing

queueing

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

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 in medium (~2×10^8 m/sec in optical fiber)
- $d_{prop} = d/s$

End-to-end delay

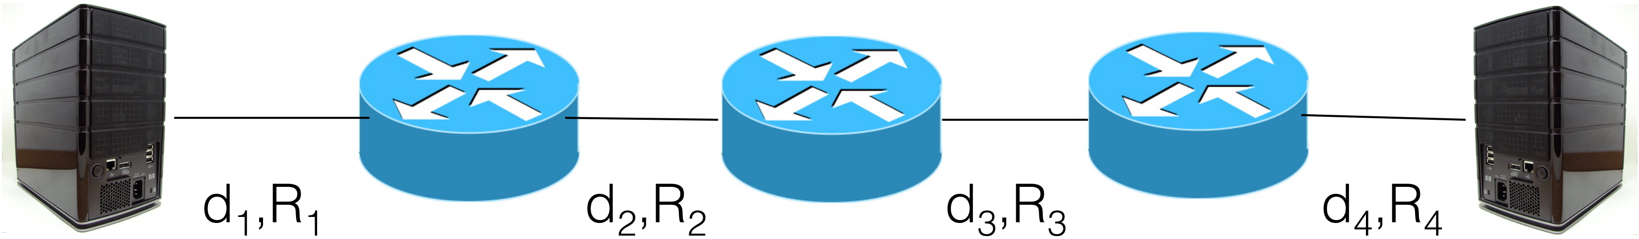

$$d = \sum_i \left(\frac{L}{R_i} + \frac{d_i}{s}\right)$$

What’s happening?

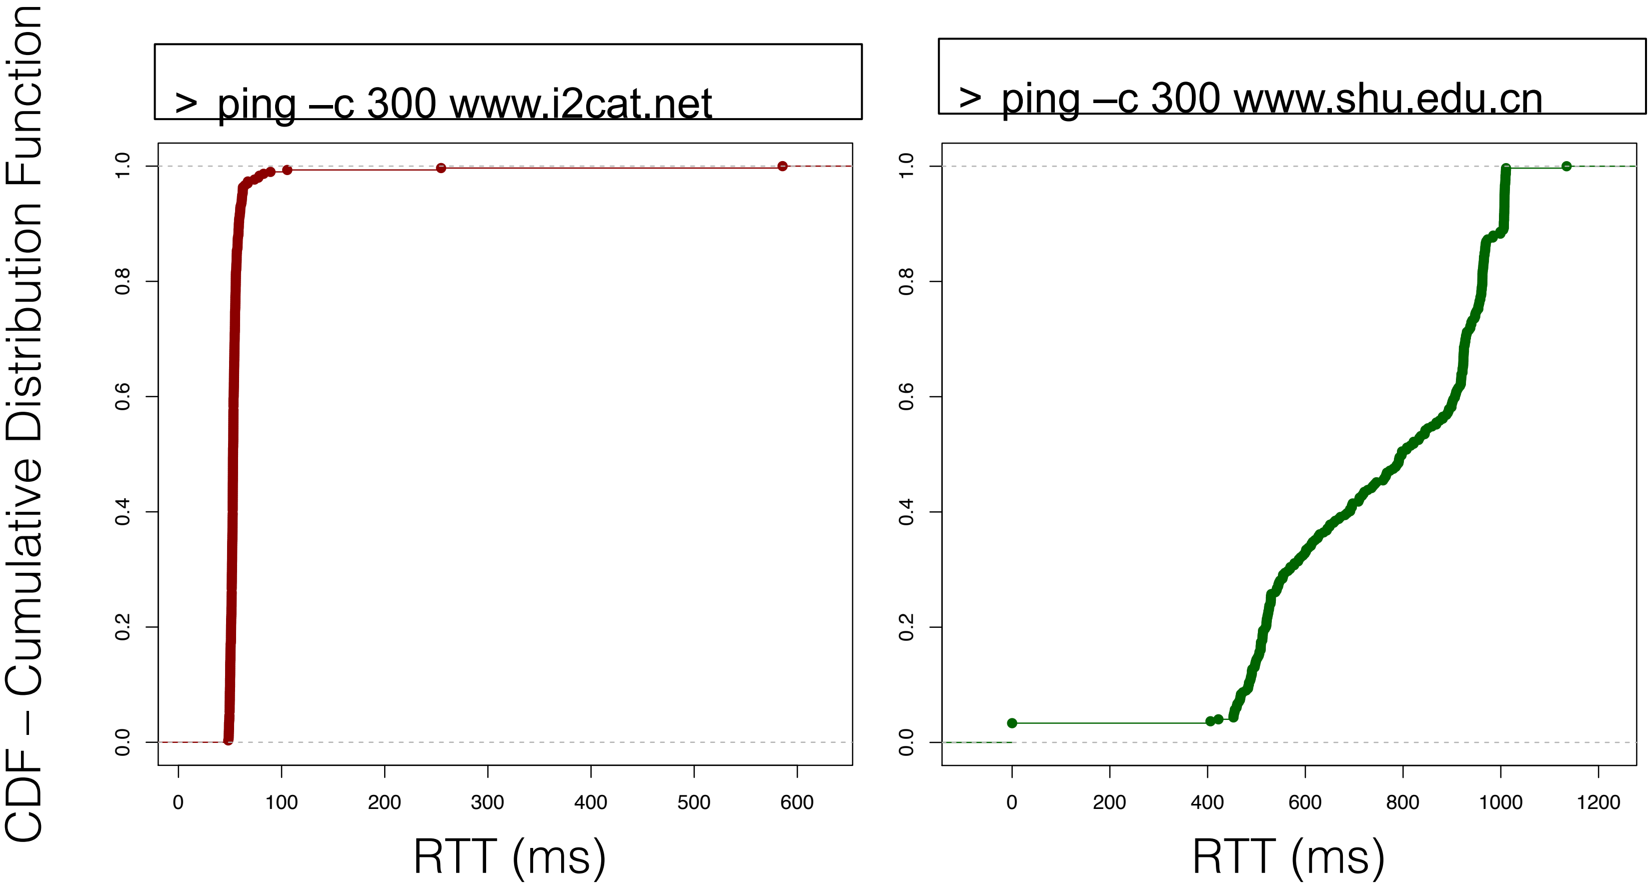

```
> ping –c 300 www.i2cat.net
```

```
> ping –c 300 www.shu.edu.cn
```

Four sources of packet delay

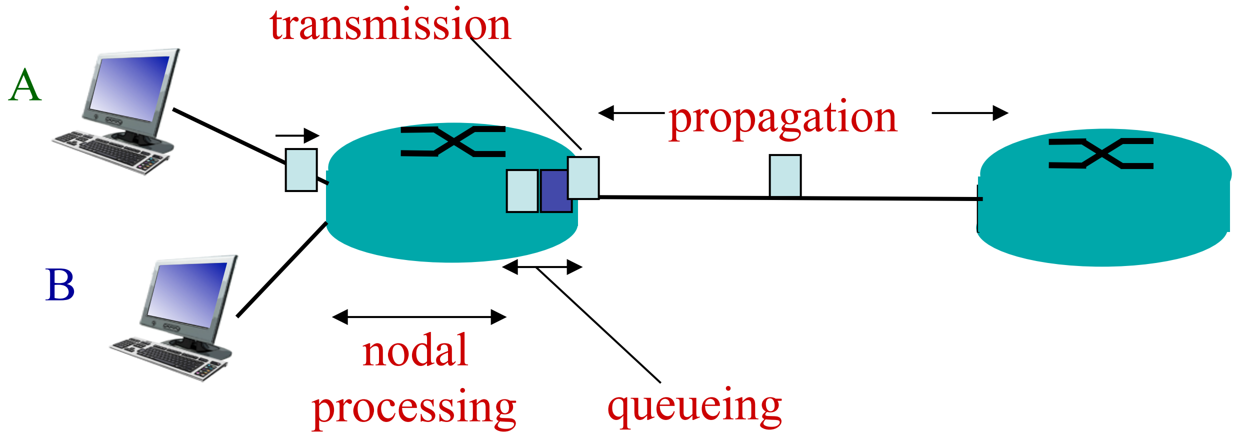

$$d = d_{\text{proc}} + d_{\text{queue}} + d_{\text{trans}} + d_{\text{prop}}$$

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

Queuing delay and loss

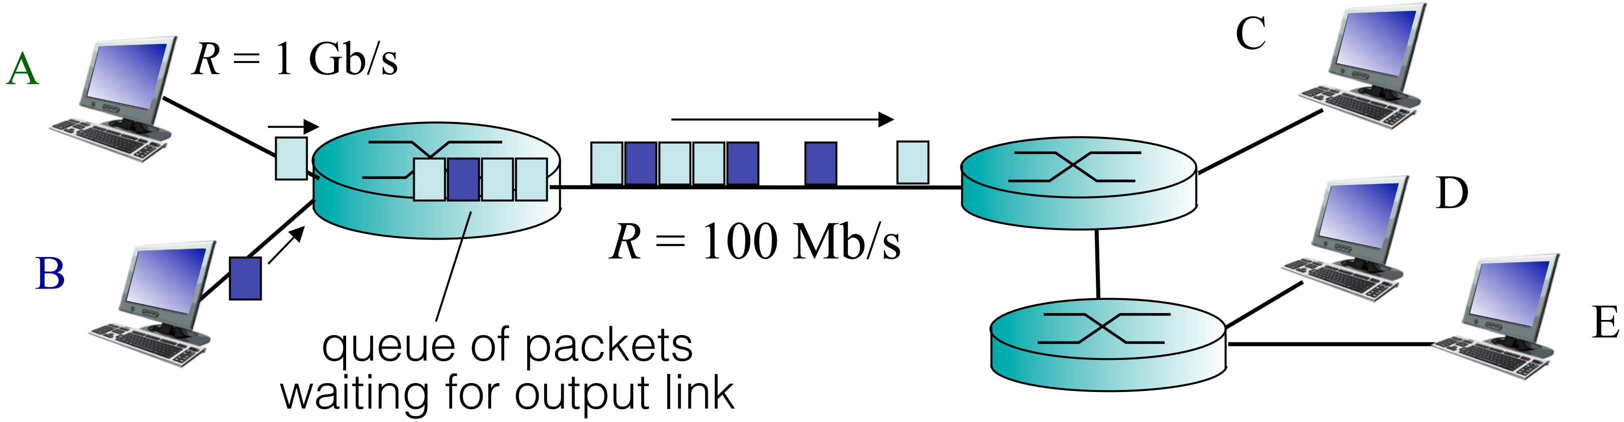

Queuing and loss:

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

Queuing delay

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

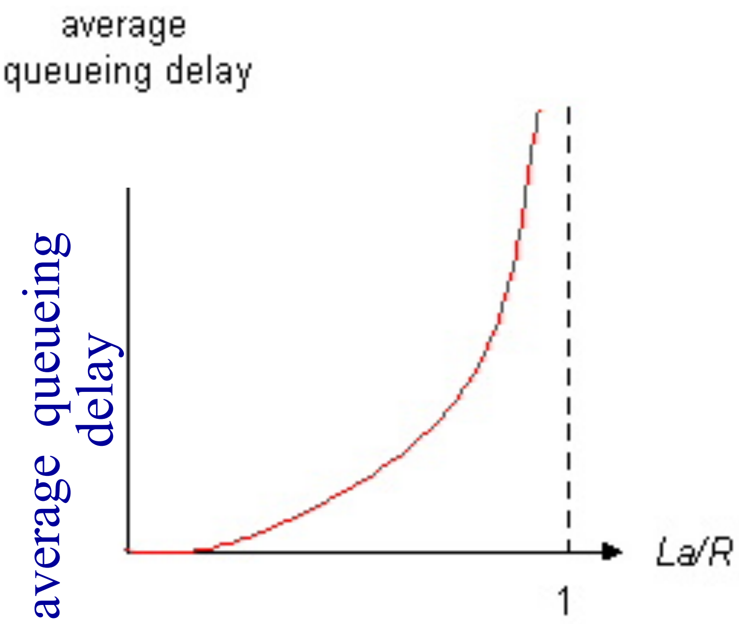

traffic intensity = La/R

- 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

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

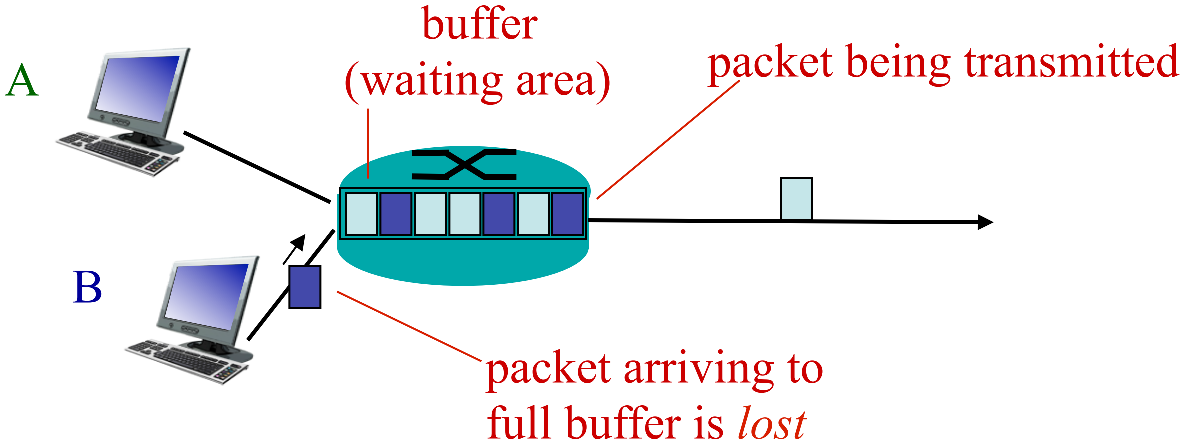

End-to-end delay revised

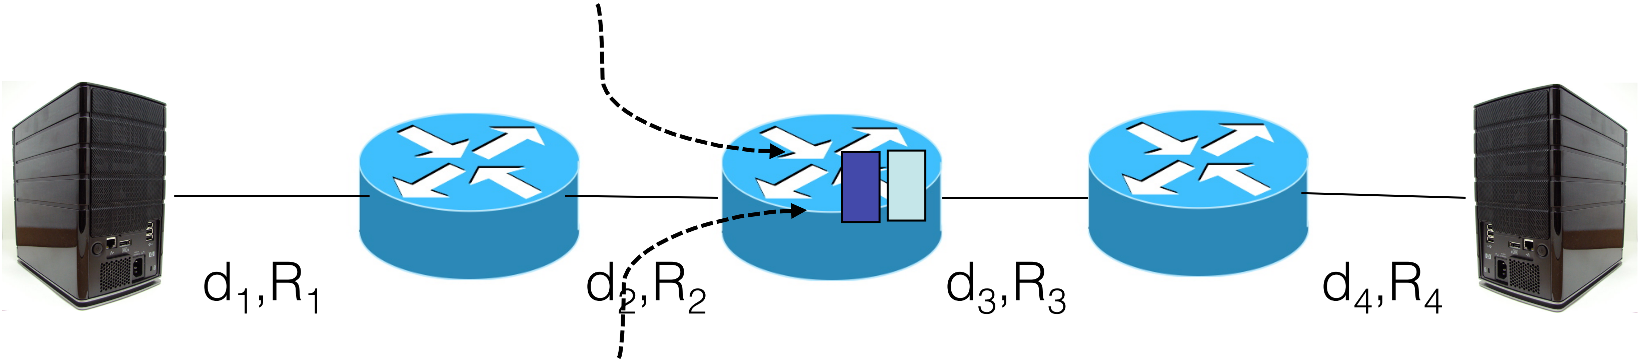

$$d = \sum_i \left(\frac{L}{R_i} + \frac{d_i}{s} + Q_i(t)\right)$$

Throughput

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

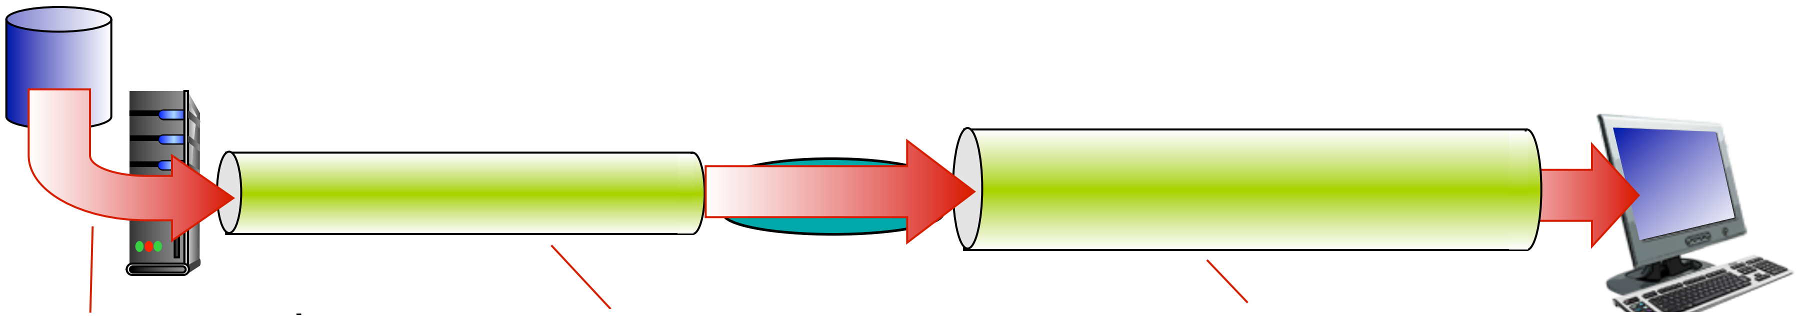

server sends bits (fluid) into pipe

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?

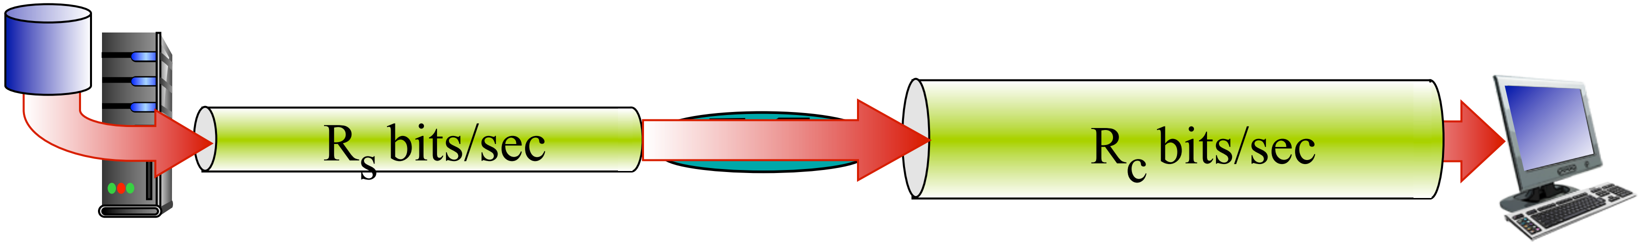

- $R_s > R_c$ What is average end-end throughput?

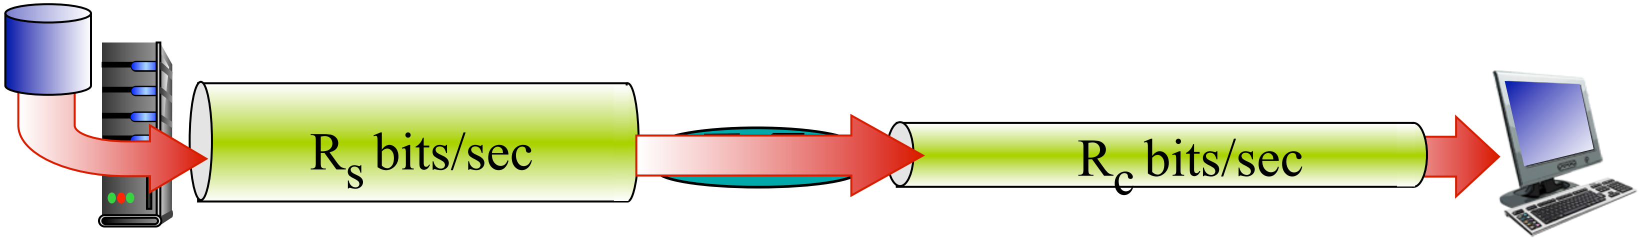

bottleneck link

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

Throughput: Internet scenario

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

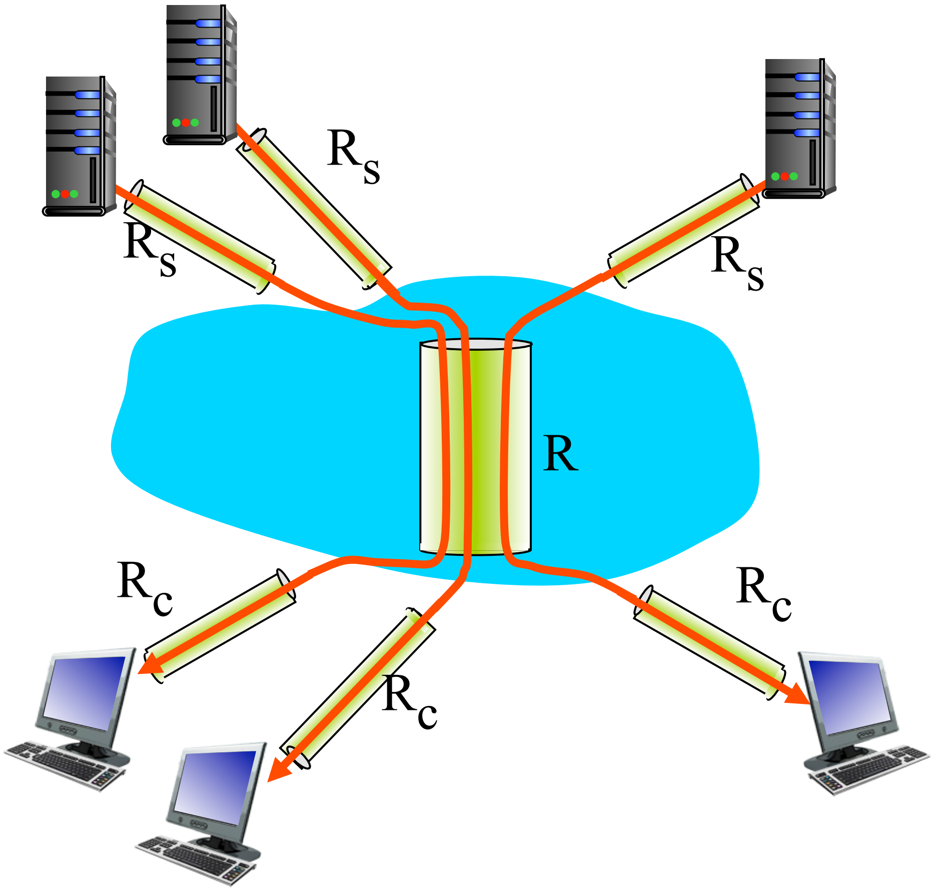

- per-connection end-end throughput: $\min(R_c, R_s, R/6)$
- in practice: R_c or R_s is often bottleneck

What have we learned today?

- How we are going to work in the coming seven weeks.
- Why this course is given.

- How was the Internet born and what are its main features.
- How has the Internet evolved, and what are the challenges ahead.
- What is packet switching
- What are the main performance issues in the Internet.
  - With focus on delays

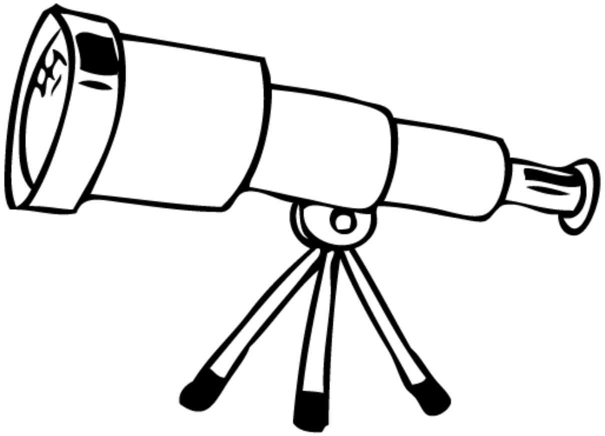

What will we learn next time?

- What is layering and encapsulation?
- What are the basic functions of the application layer?
- How are data passed from the application to the transport layer?
- Basic application layer protocols

Textbook

- Chapter 1:
  - Section 1.1 – What is the Internet?
  - Section 1.2 – Access Networks
  - Section 1.3 – The Network Core
  - Section 1.4 – Delay, Loss and Throughput
  - Section 1.6 – Networks under attack
  - Section 1.7 – History of Computer Networking

Home reading

END-TO-END ARGUMENTS IN SYSTEM DESIGN

By Saltzer et al.

ACM Transactions in Computer Systems 2, 4, November, 1984, pages 277-288

Available at:

http://web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf

and on Canvas in the Home readings folder