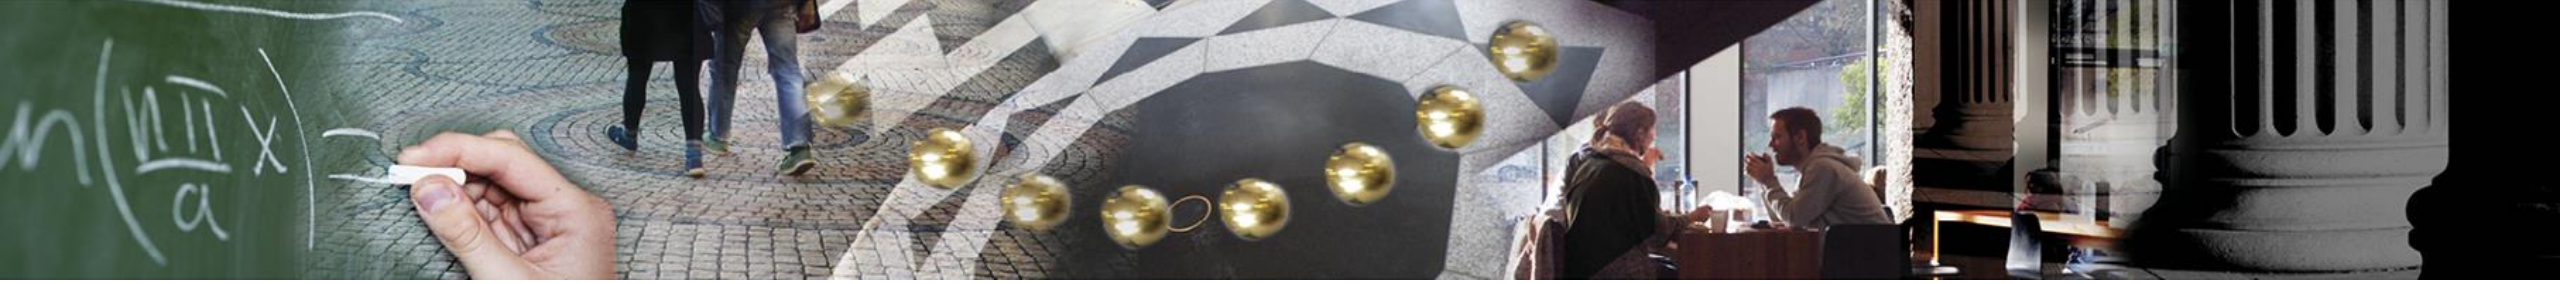

UiO : University of Oslo

Information technology in the health sector (DIGHEL4360)

Internet and the World Wide Web

Introduction

- Nils Gruschka
  - University Kiel (Diploma in Computer Science)
  - T-Systems, Hamburg
  - University Kiel (PhD in Computer Science)
  - NEC Laboratories Europe, Bonn + Heidelberg
  - University of Applied Science Kiel
  - University of Oslo
- Contact:
  - nilsgrus@ifi.uio.no
- Areas of interest:
  - Security: Network, Web, Cloud Computing, Industrial Networks
  - Privacy, Data Protection

Acknowledgement

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

All material copyright 1996-2020
J.F Kurose and K.W. Ross, All Rights Reserved

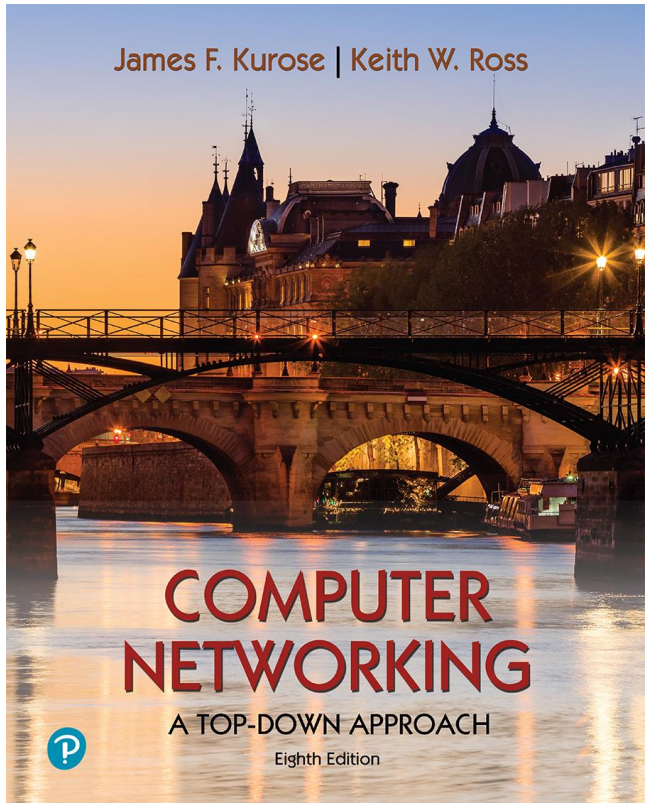

Computer Networking: A Top-Down Approach

8th edition
Jim Kurose, Keith Ross
Pearson, 2020

Online-Tutorial

- https://www.youtube.com/playlist?list=PLzdnOPI1iJNfMRZm5DDxco3UdsFegvuB7

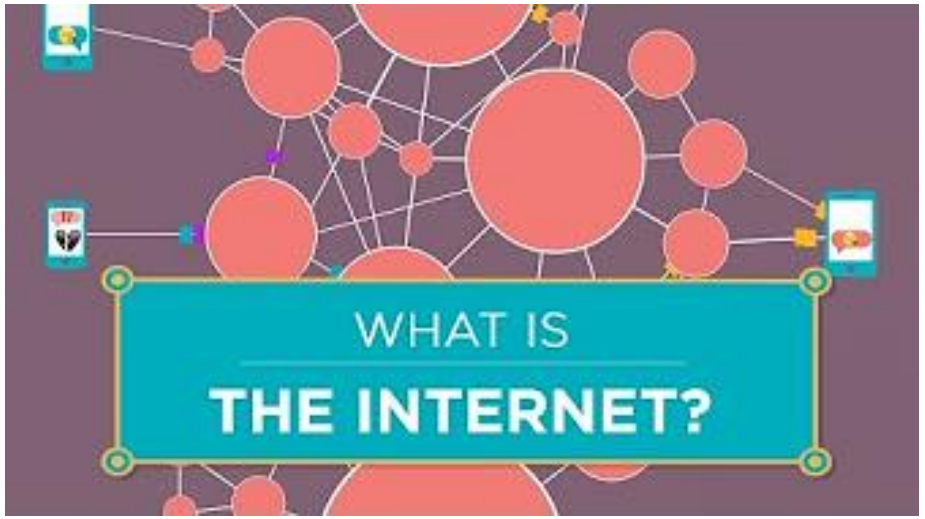

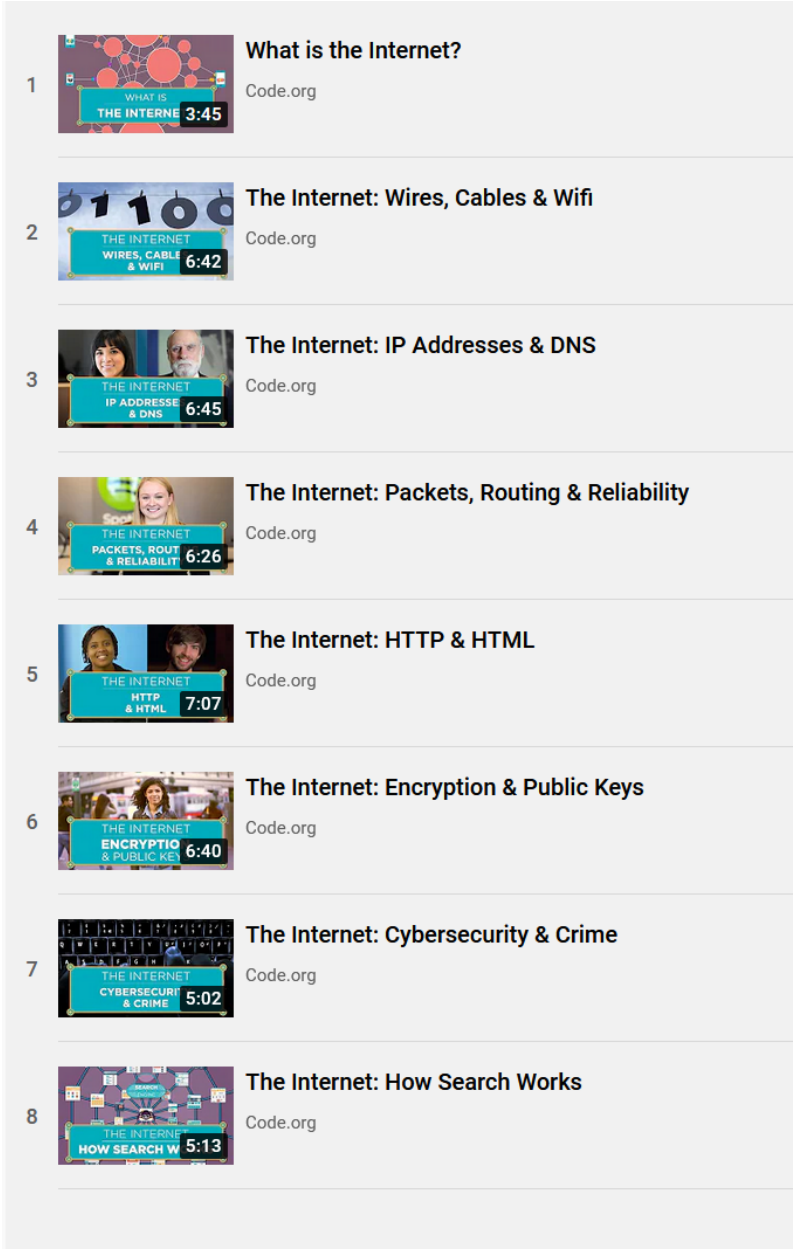

The Internet

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

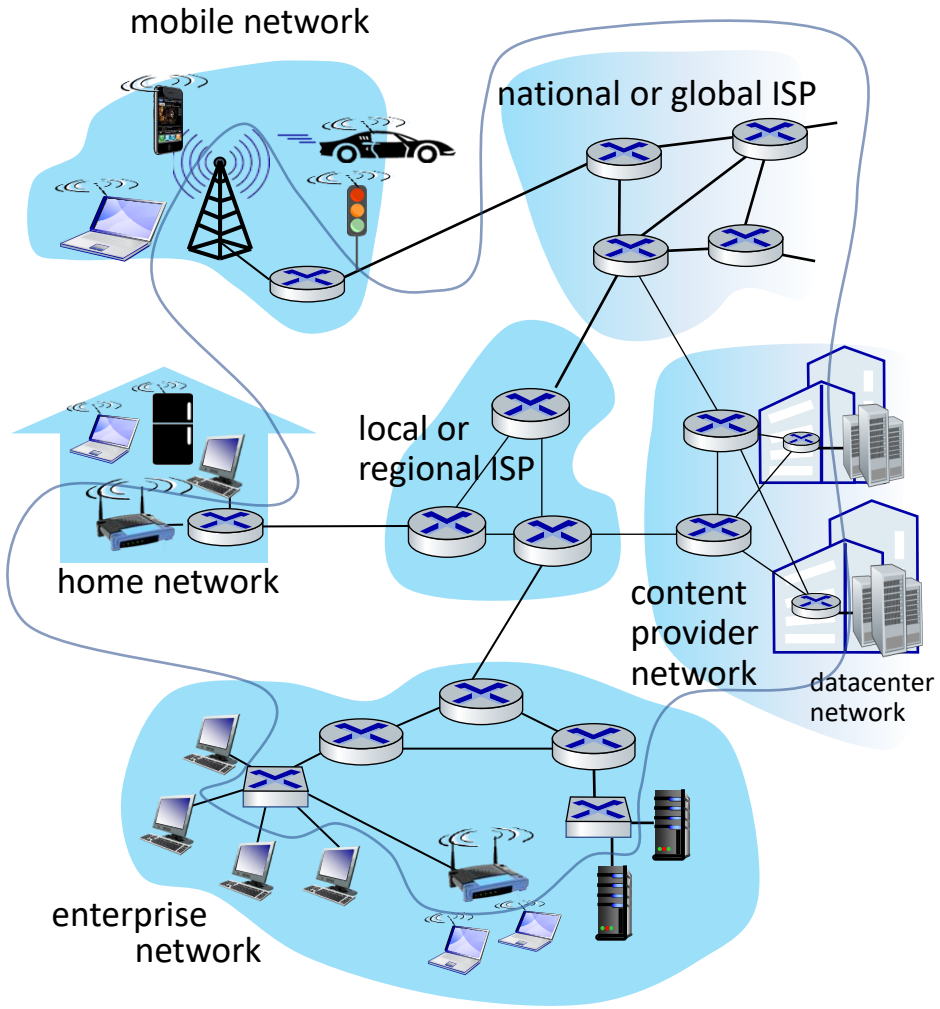

“Fun” Internet-connected devices

Amazon Echo

Internet refrigerator

IP picture frame

Pacemaker & Monitor

Tweet-a-watt: monitor energy use

bikes

cars

scooters

Security Camera

Slingbox: remote control cable TV

Web-enabled toaster + weather forecaster

AR devices

Internet phones

Gaming devices

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

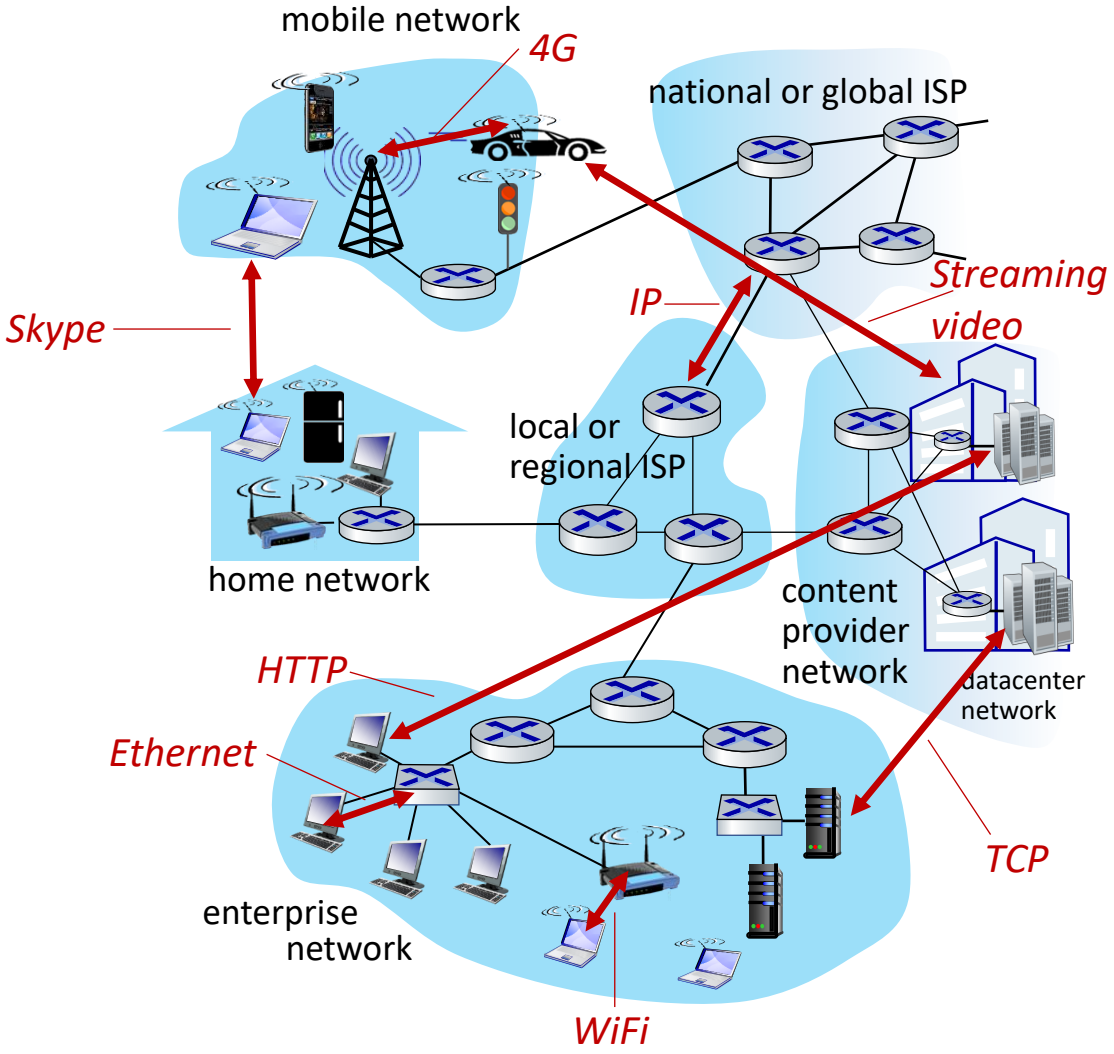

What's a protocol?

A human protocol and a computer network protocol:

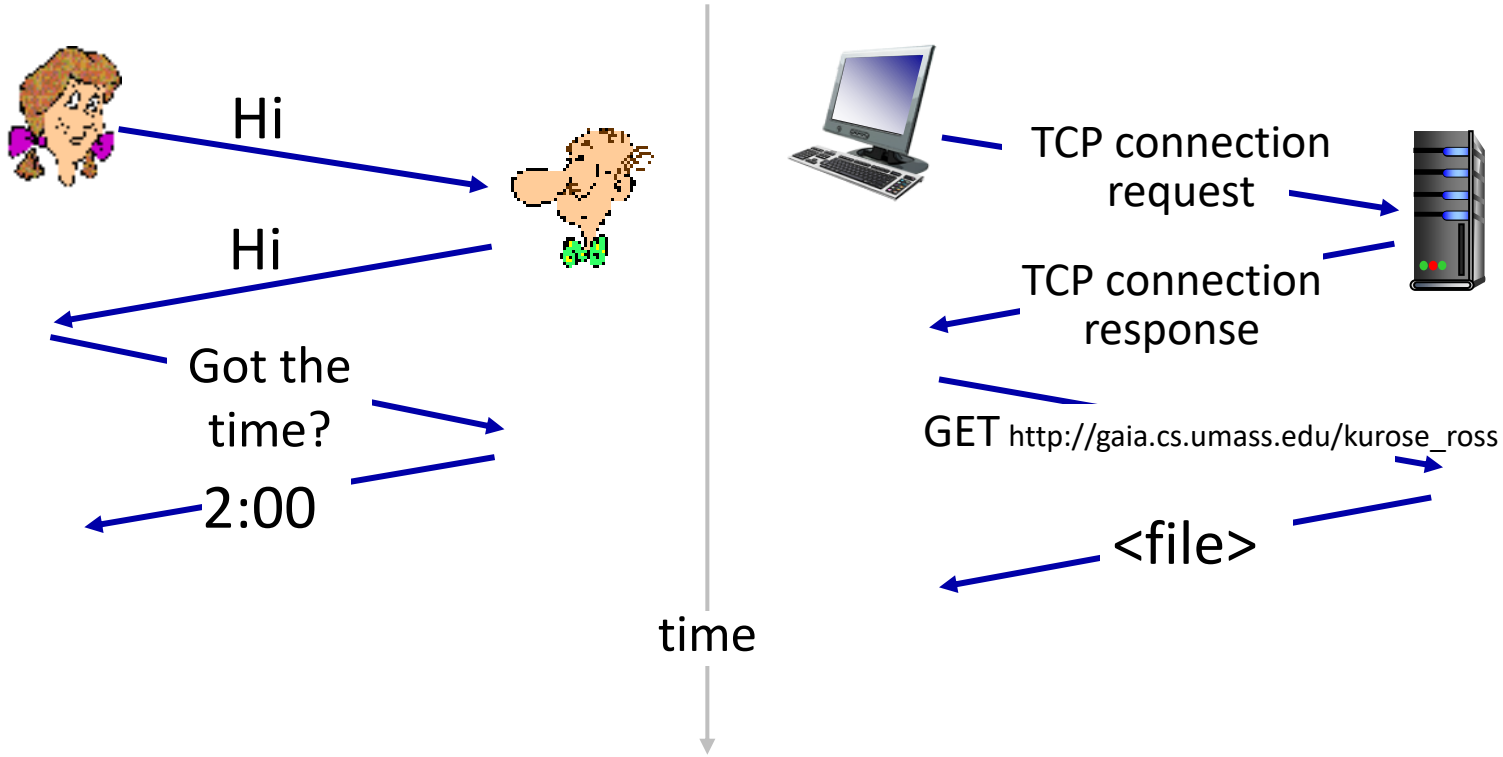

Q: other human protocols?

What's a protocol?

Human protocols:

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

Rules for:

… 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 message transmission, receipt

A closer look at Internet structure

Network edge:

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

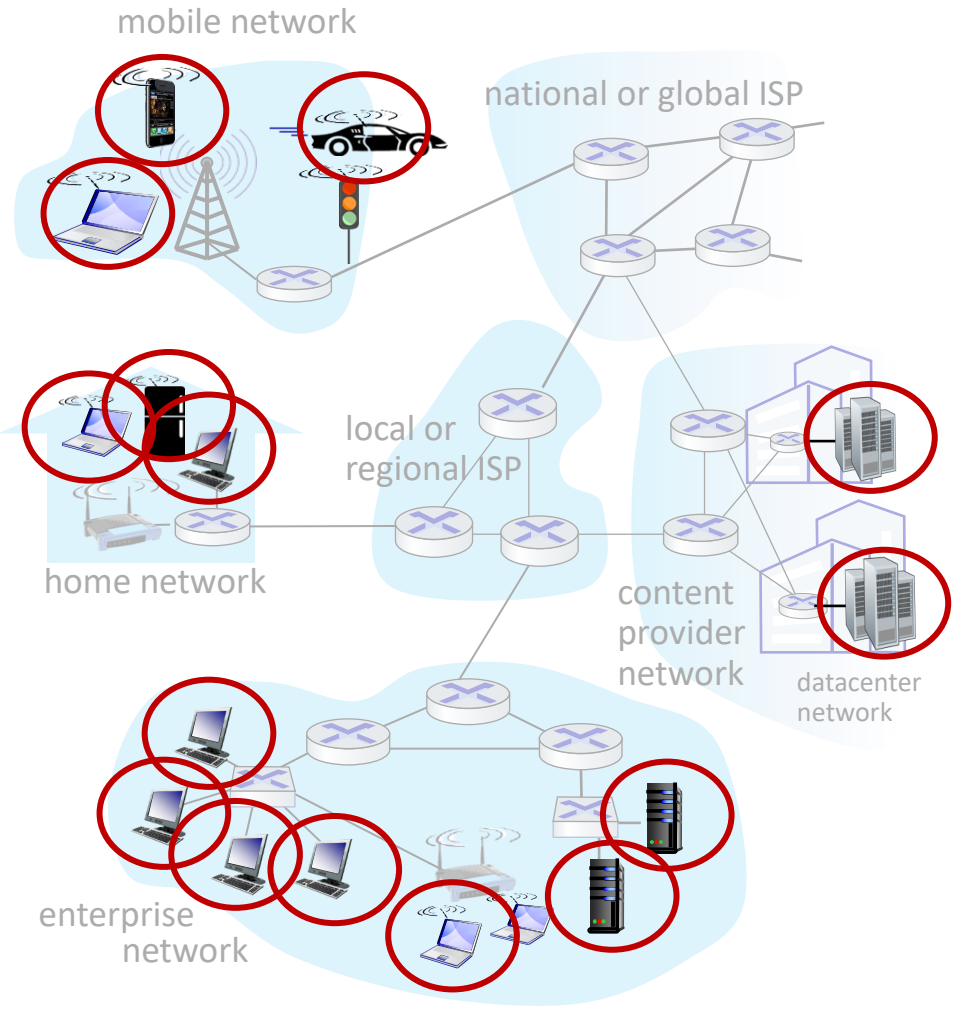

A closer look at Internet structure

Network edge:

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

Access networks, physical media:

- wired, wireless communication links

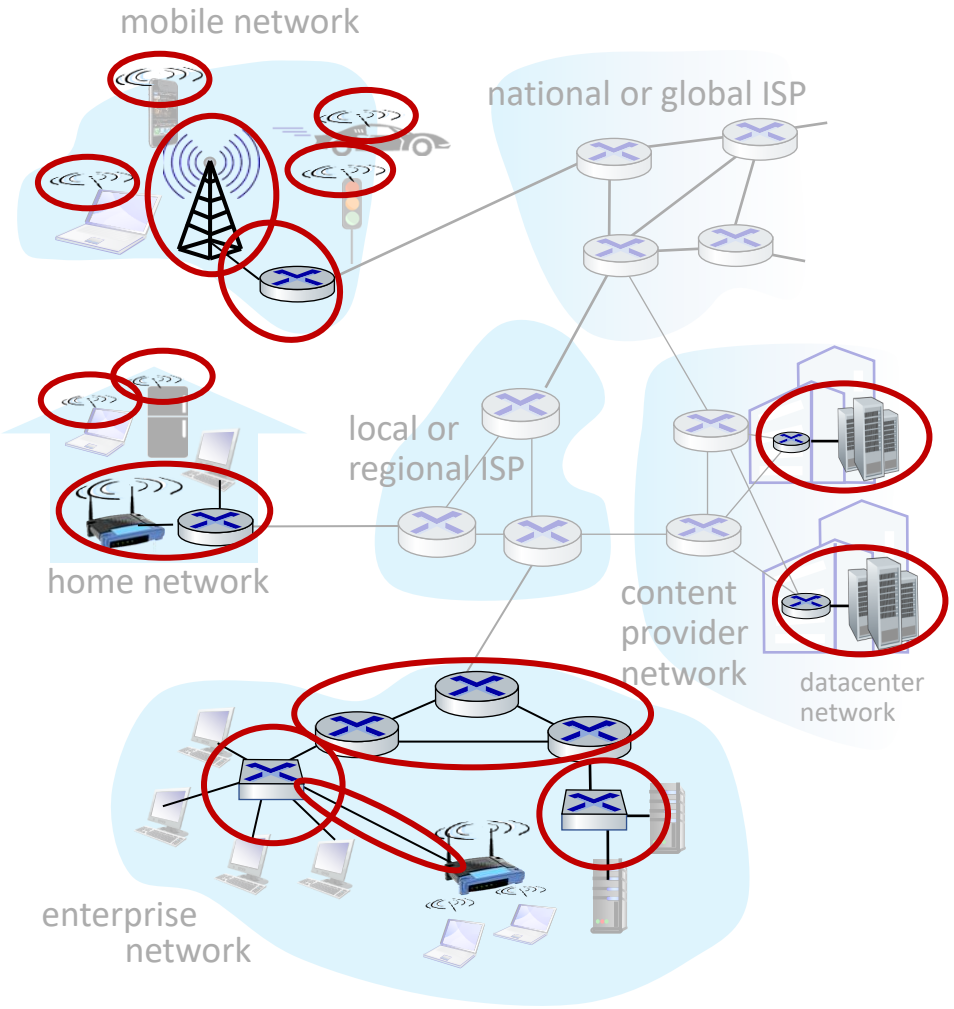

Wireless access networks

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

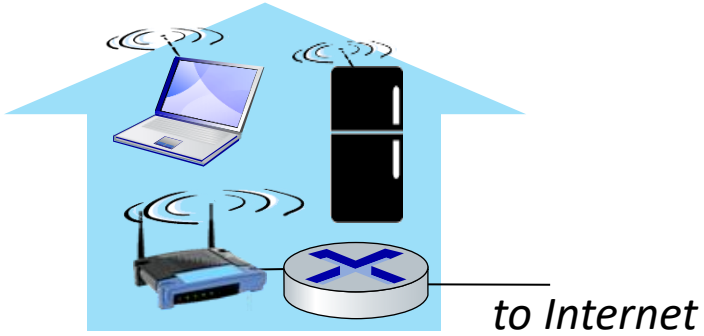

Wide-area cellular access networks

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

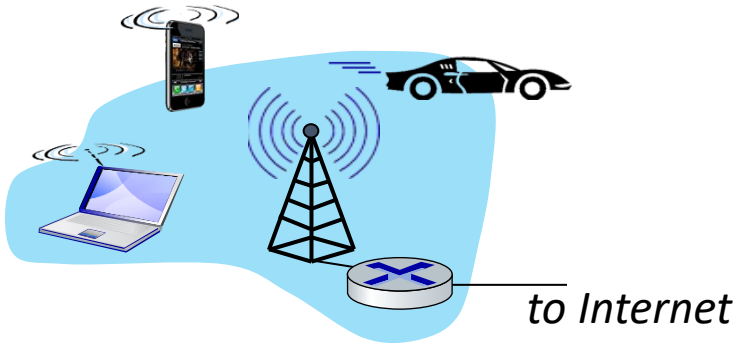

A closer look at Internet structure

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

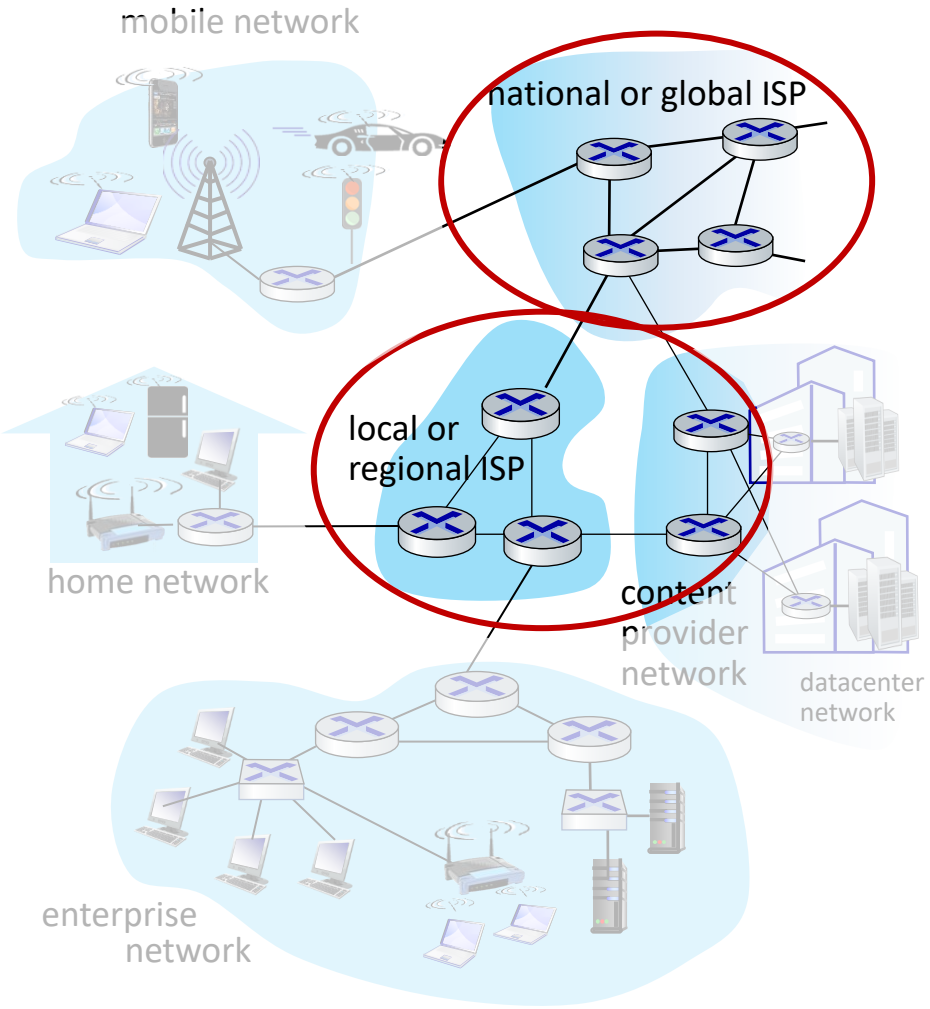

Routers

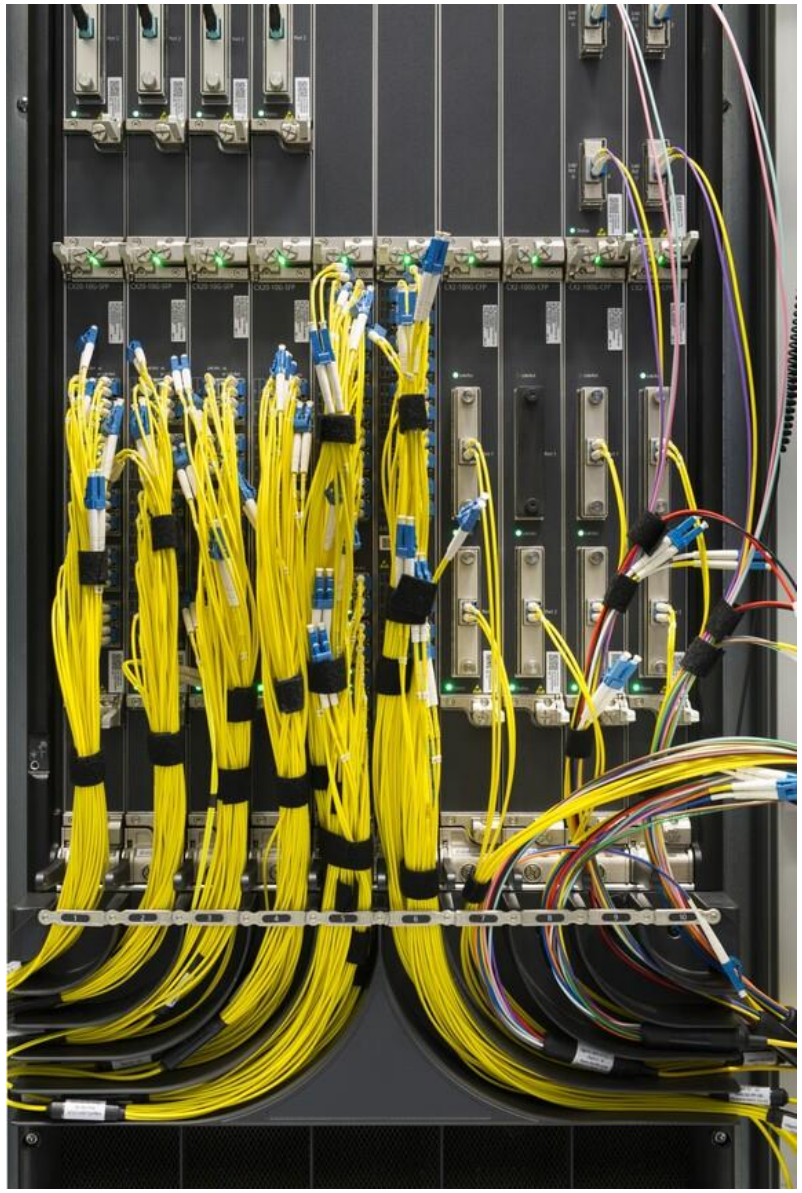

Image sources: ASUS, DE-CIX

The network core

- mesh of interconnected routers
- packet-switching: hosts break application-layer messages into *packets*
  - network forwards packets from one router to the next, across links on path from source to destination

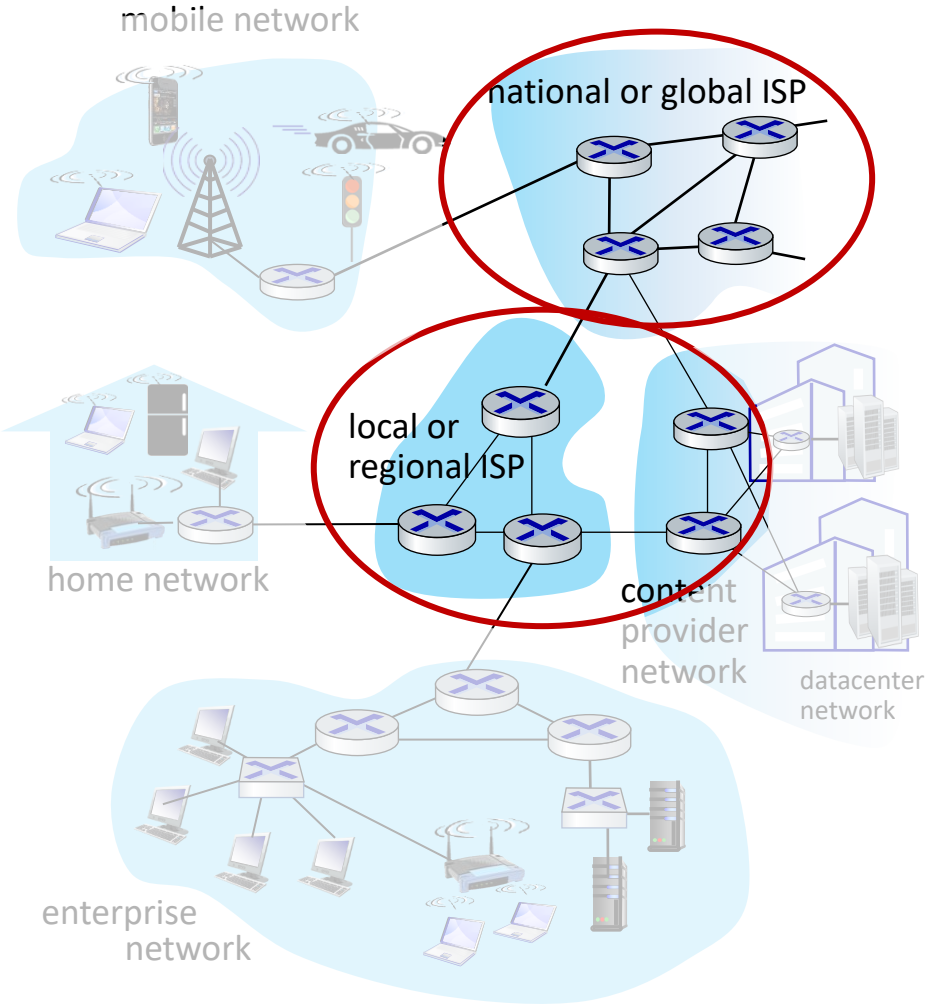

Two key network-core functions

Forwarding:

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

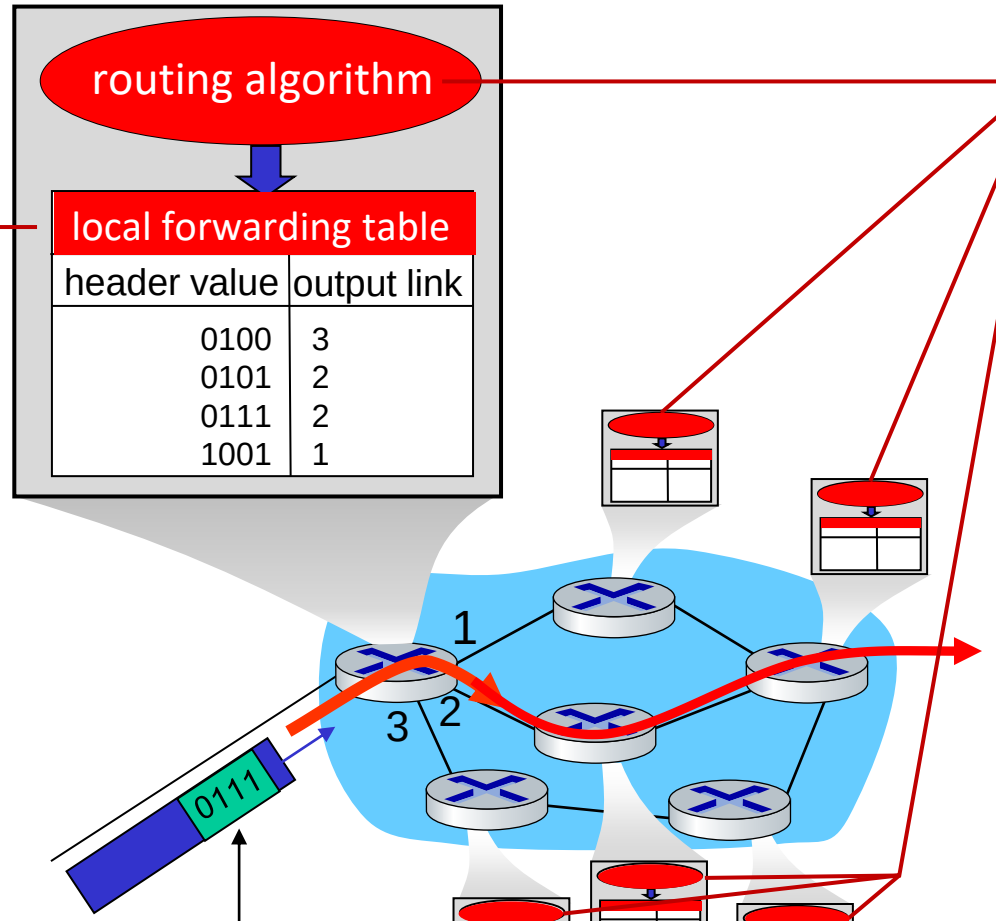

Routing:

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

Internet structure: a "network of networks"

- hosts connect to Internet via access Internet Service Providers (ISPs)
- access ISPs in turn must be interconnected
  - so that *any* two hosts *(anywhere!)* can send packets to each other
- resulting network of networks is very complex
  - evolution driven by economics, national policies

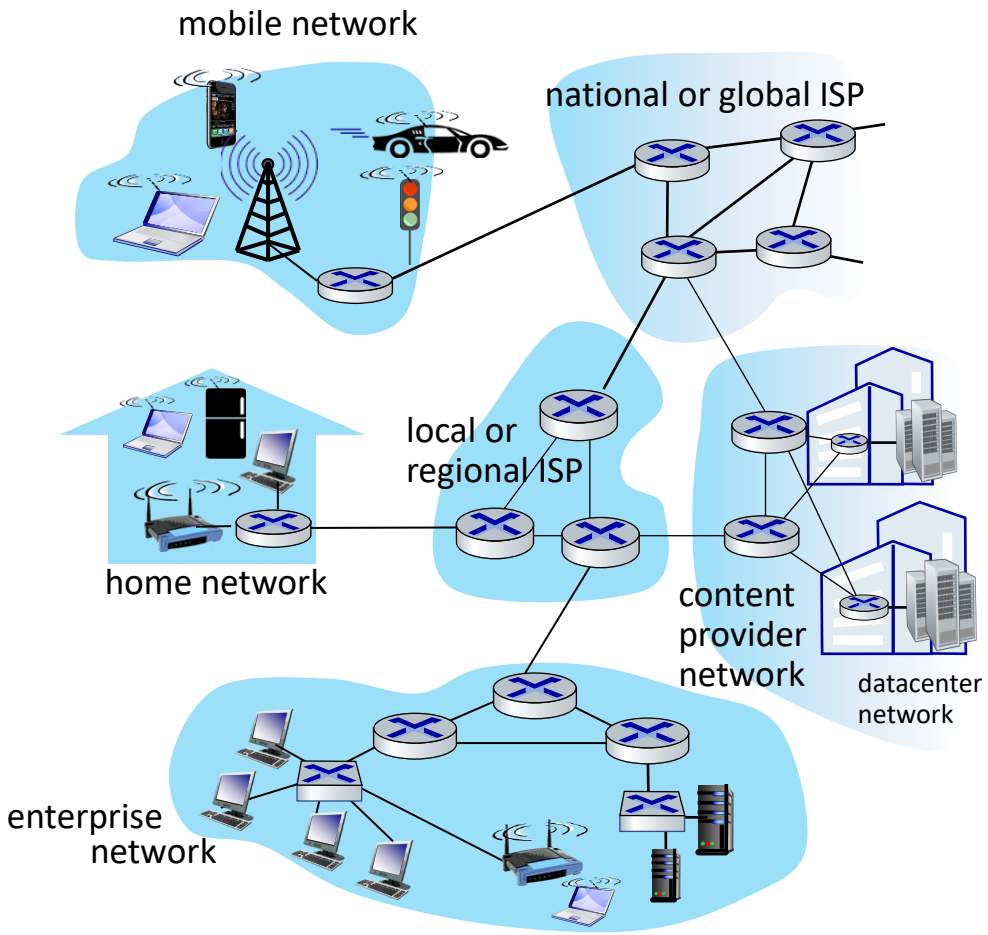

Let's take a stepwise approach to describe current Internet structure

Internet structure: a "network of networks"

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

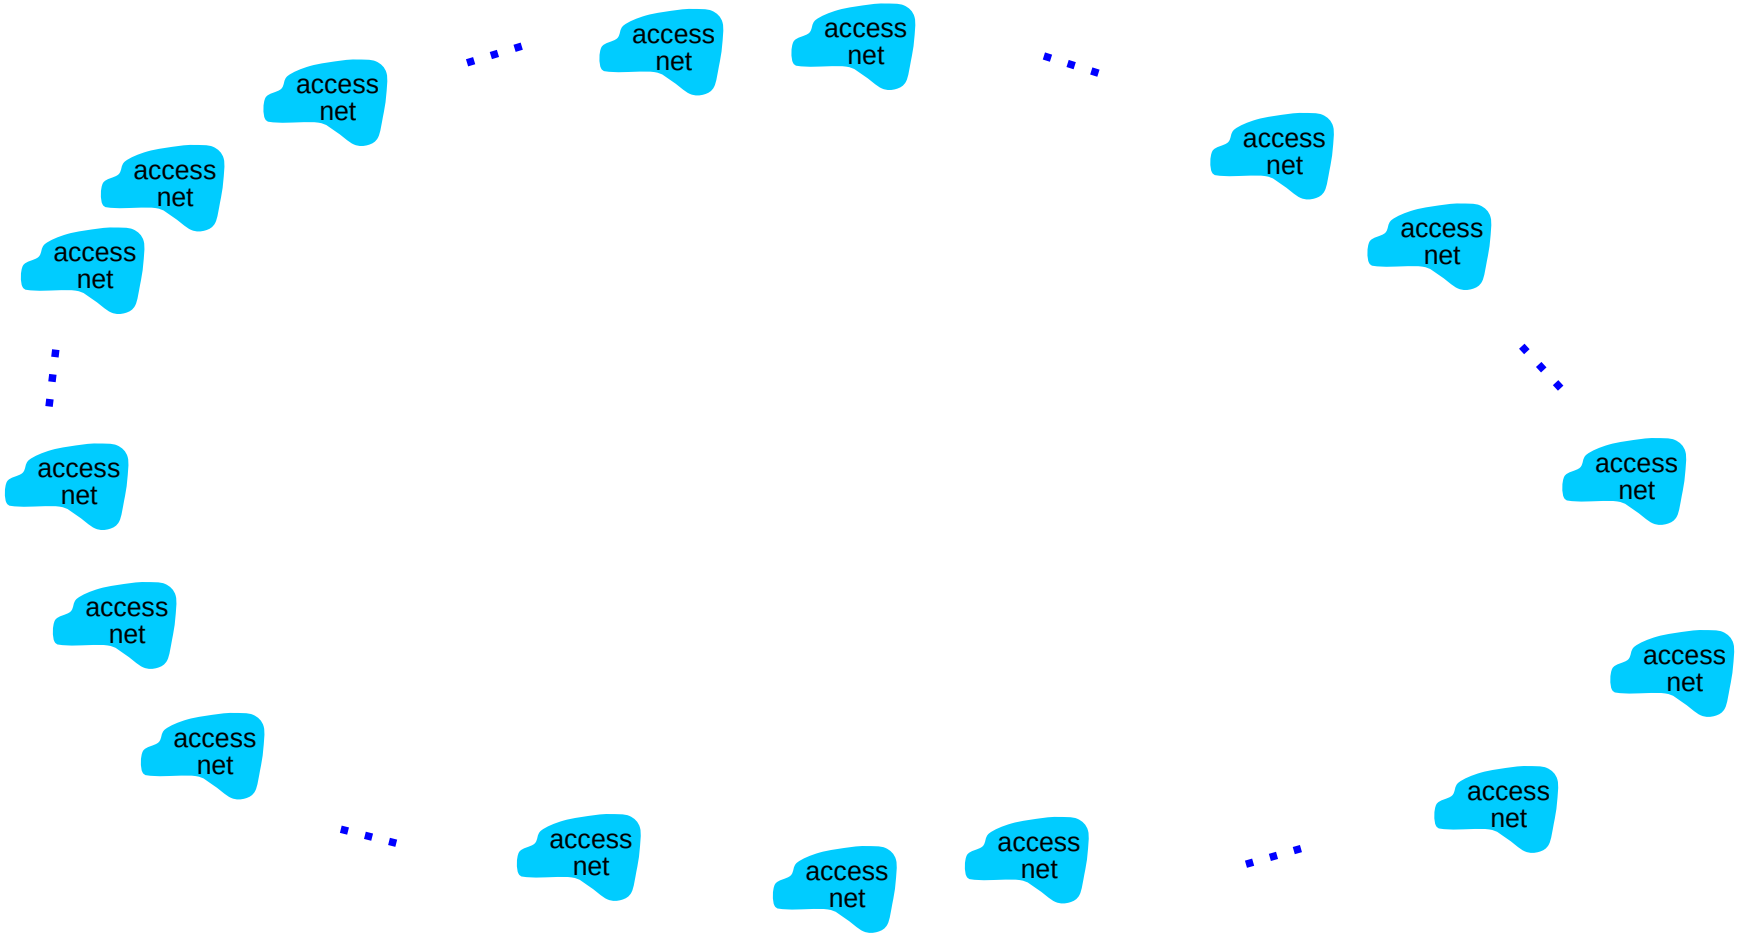

Internet structure: a "network of networks"

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

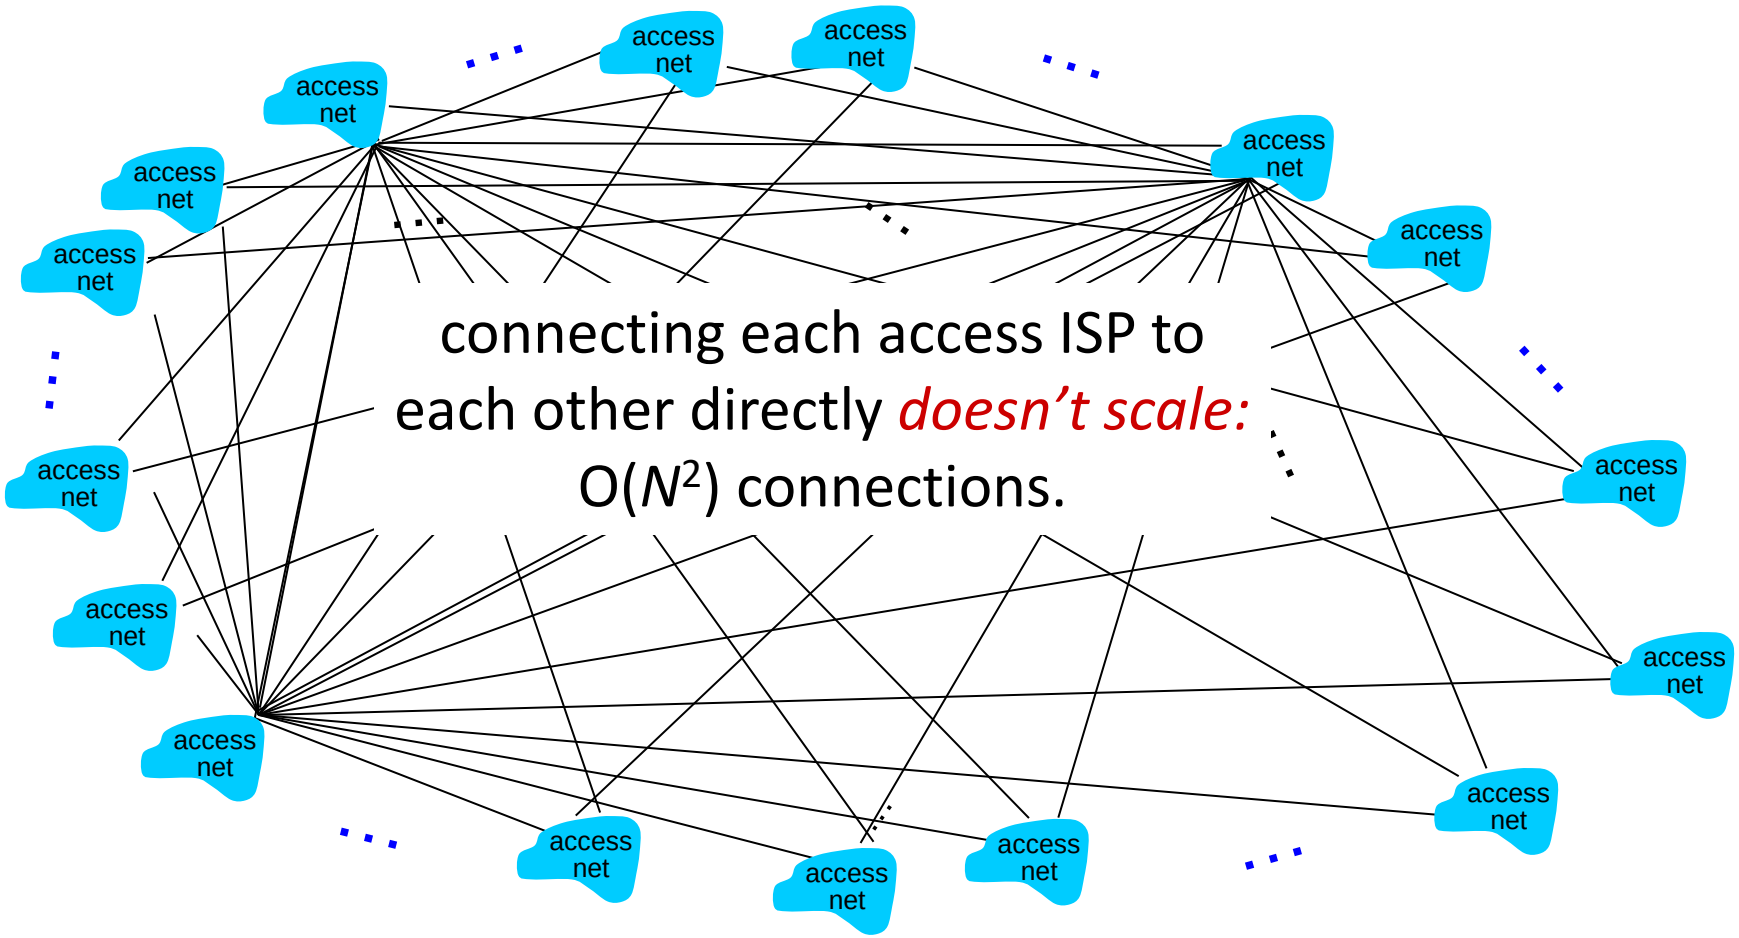

Internet structure: a "network of networks"

Option: connect each access ISP to one global transit ISP?

Customer and provider ISPs have economic agreement.

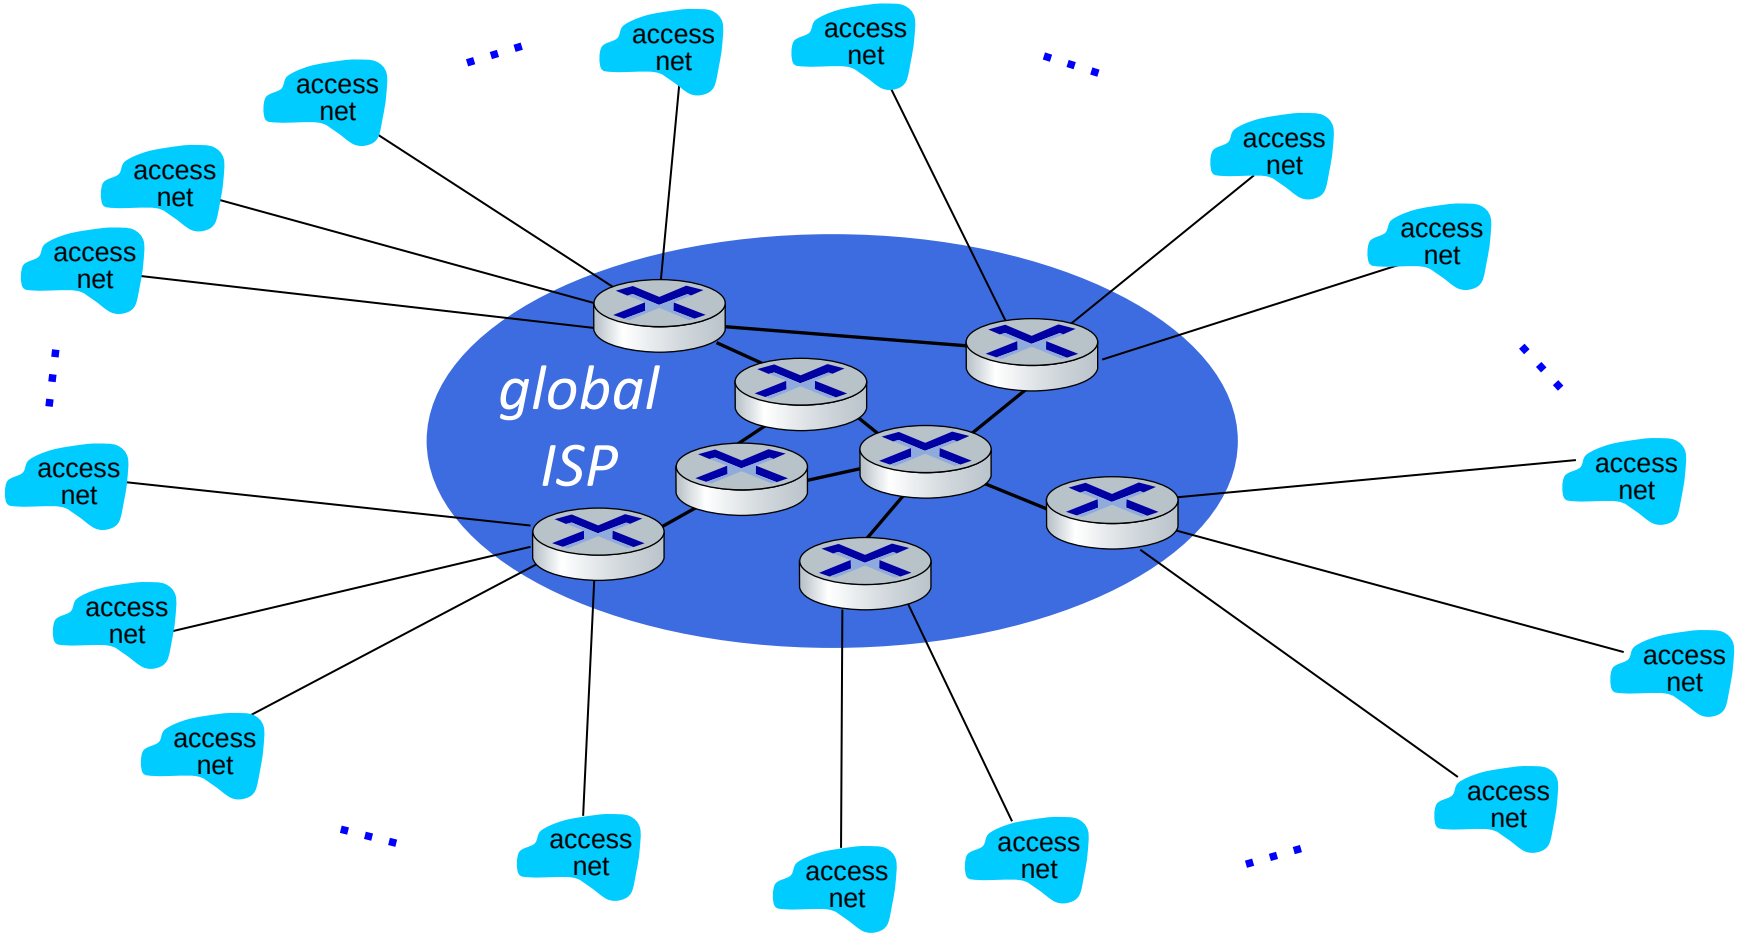

Internet structure: a "network of networks"

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

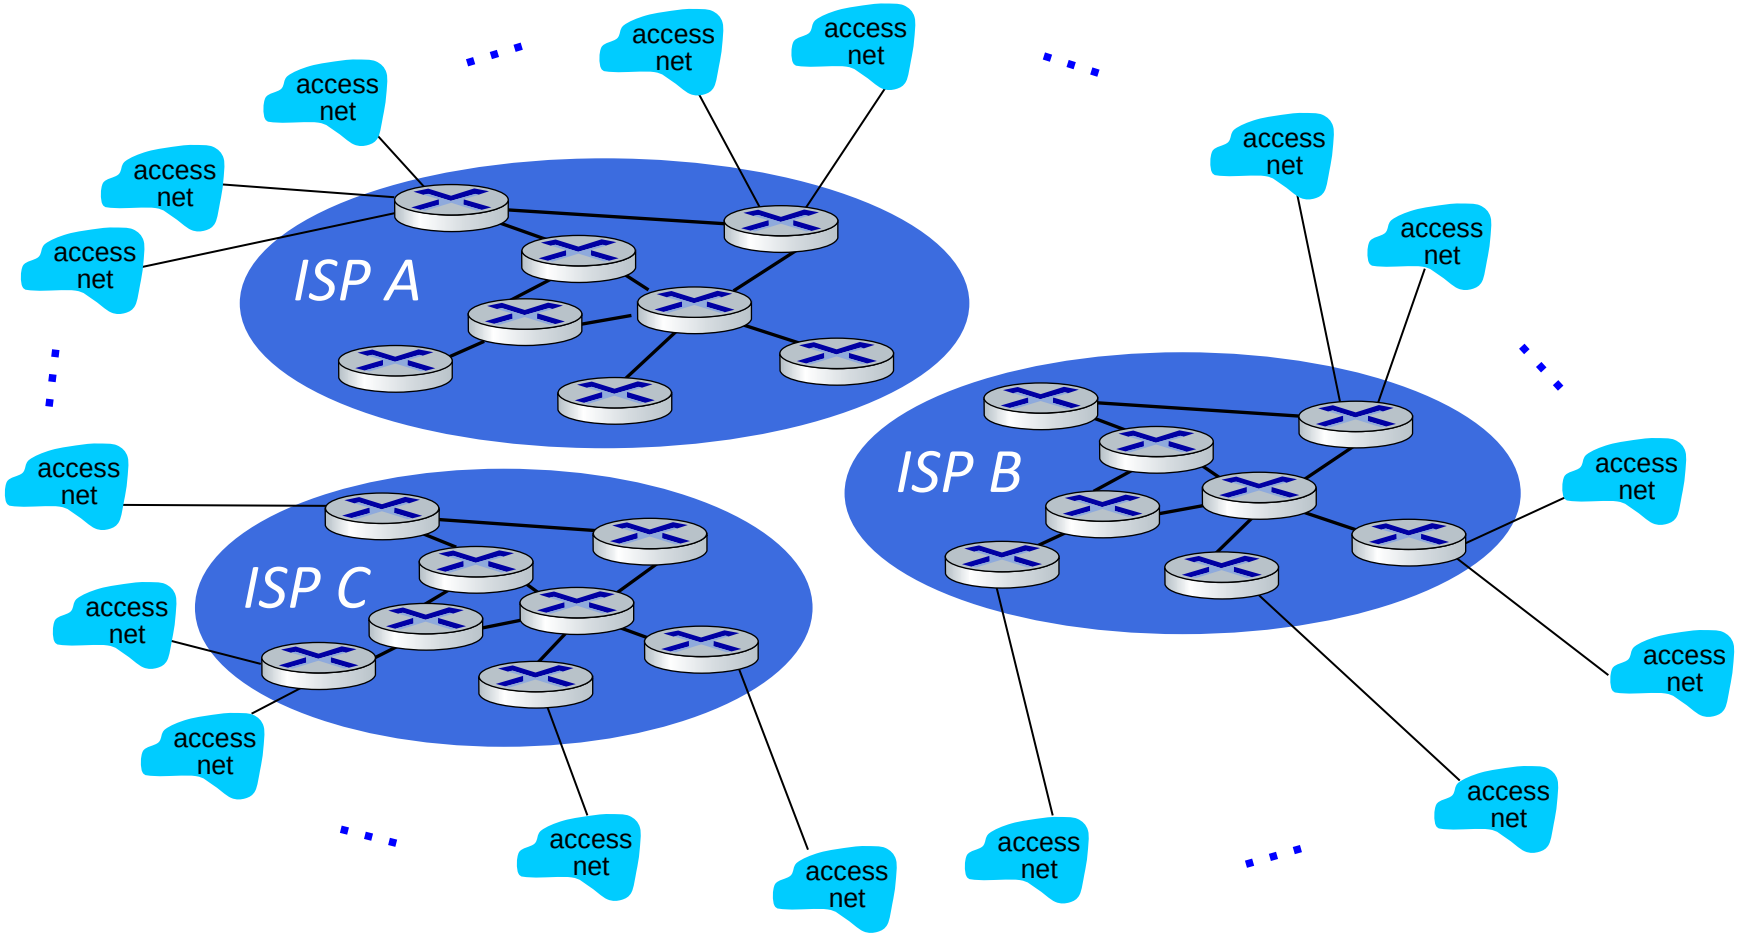

Internet structure: a "network of networks"

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

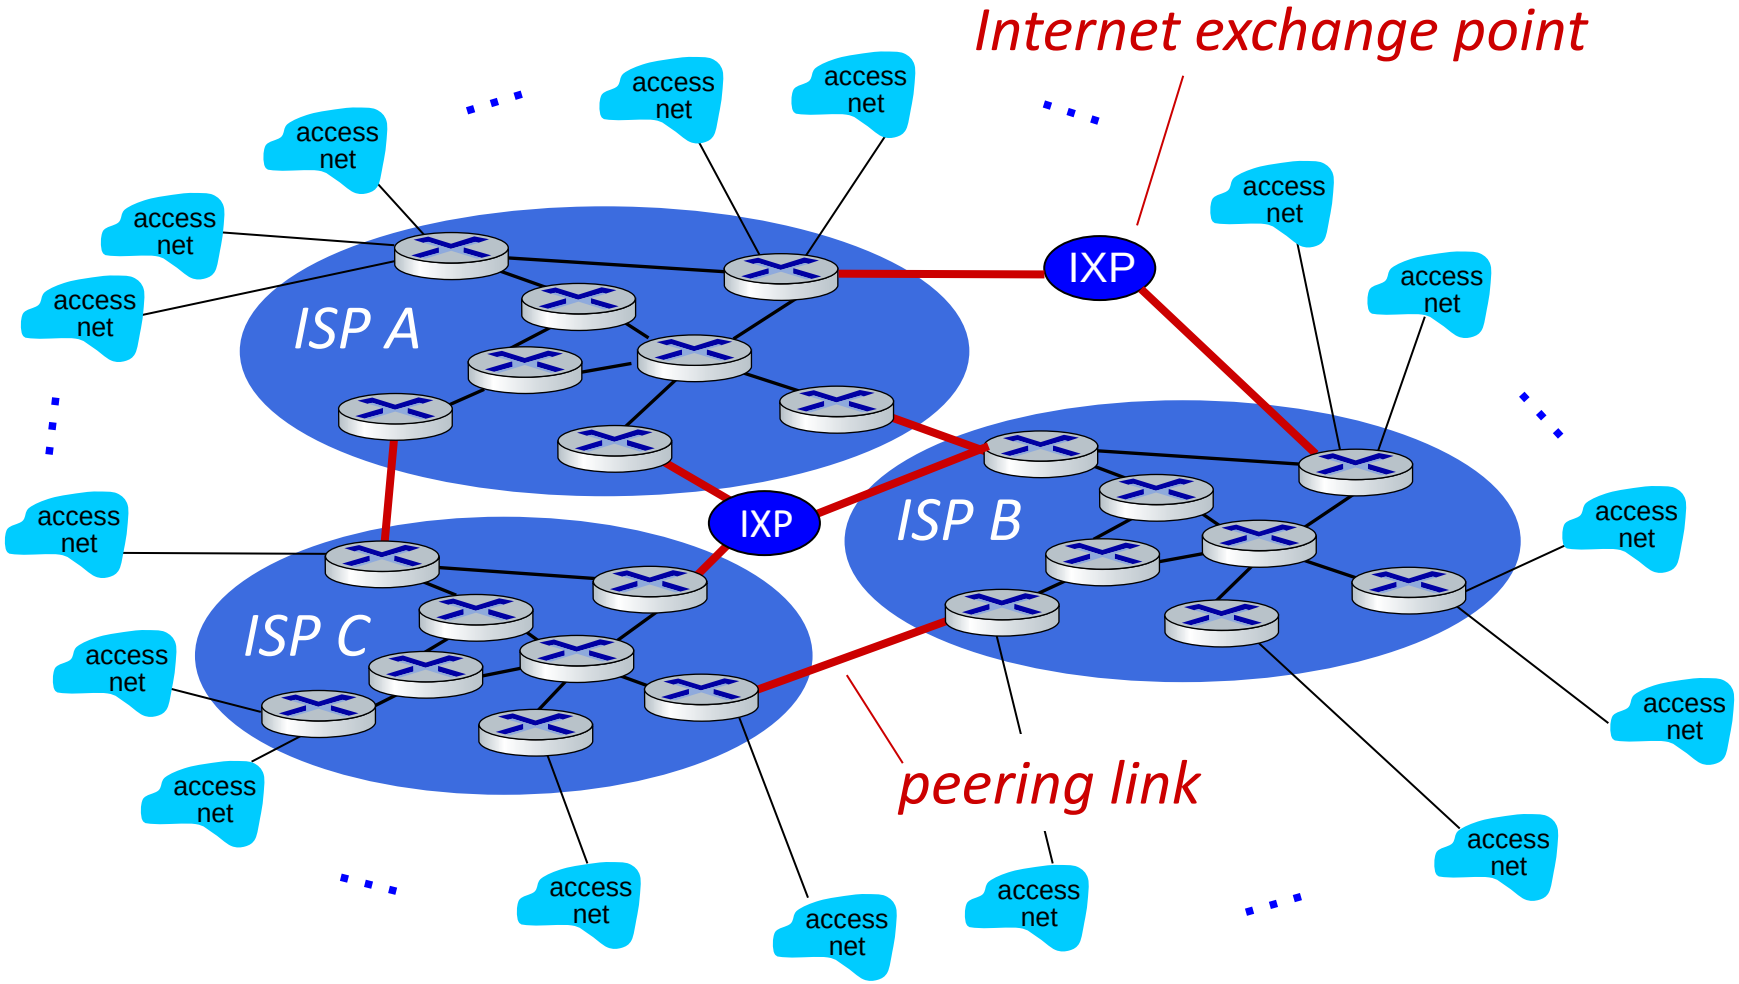

Internet structure: a "network of networks"

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

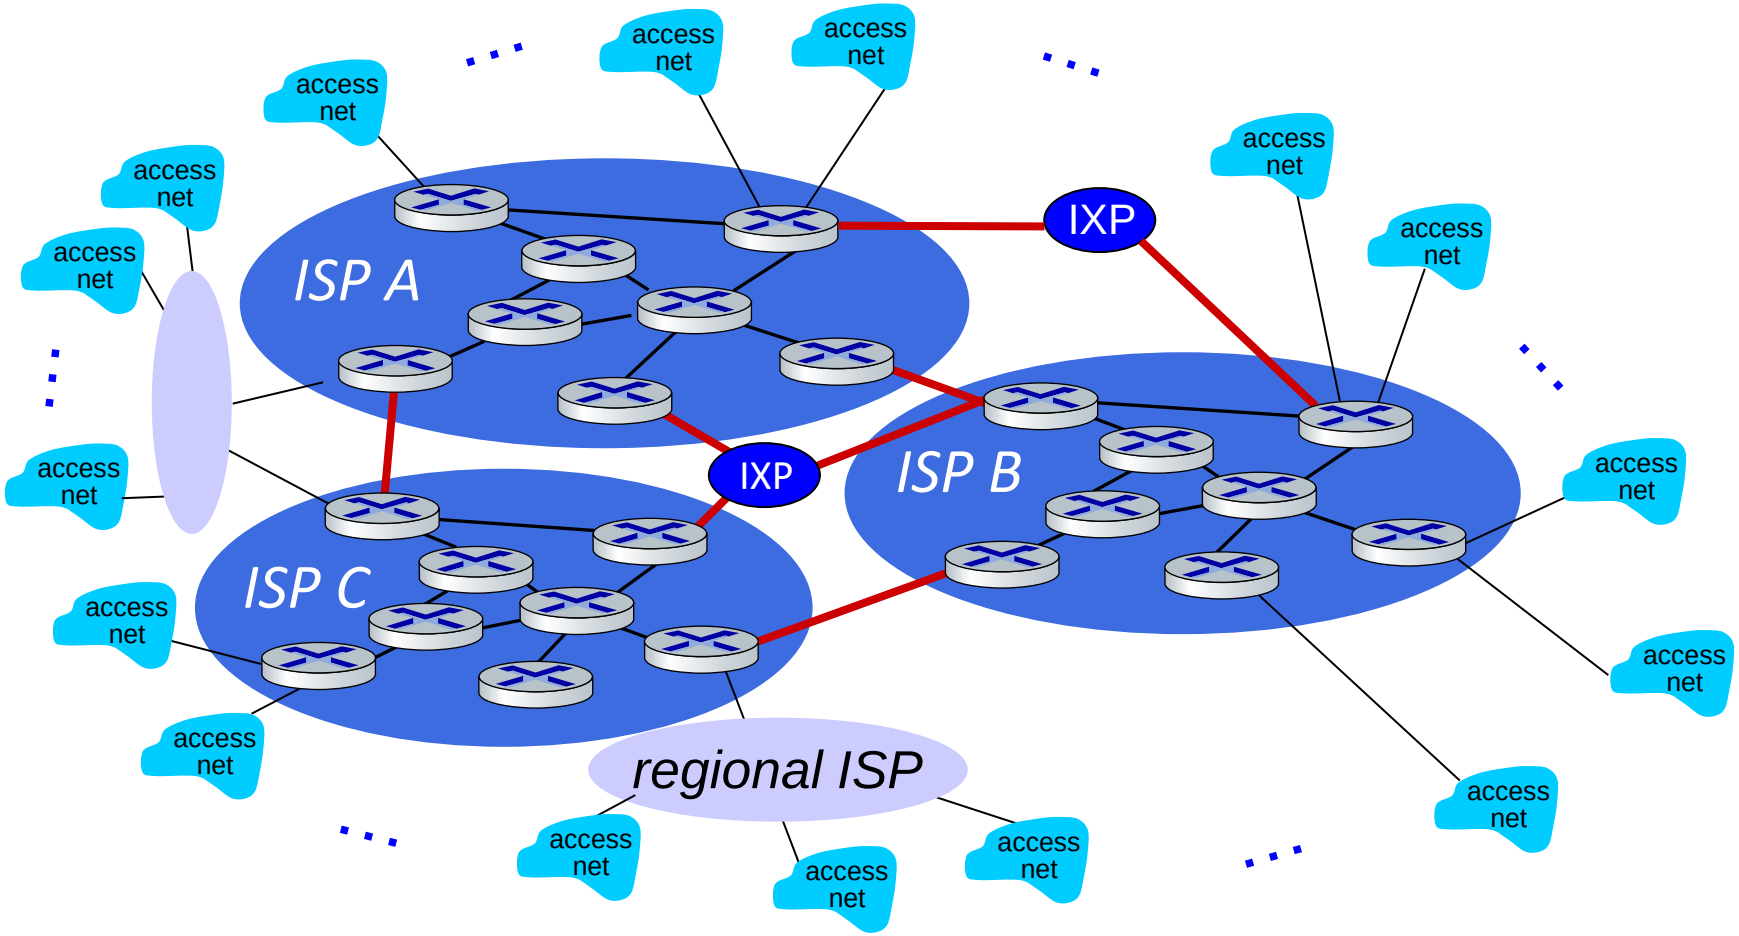

Internet structure: a "network of networks"

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

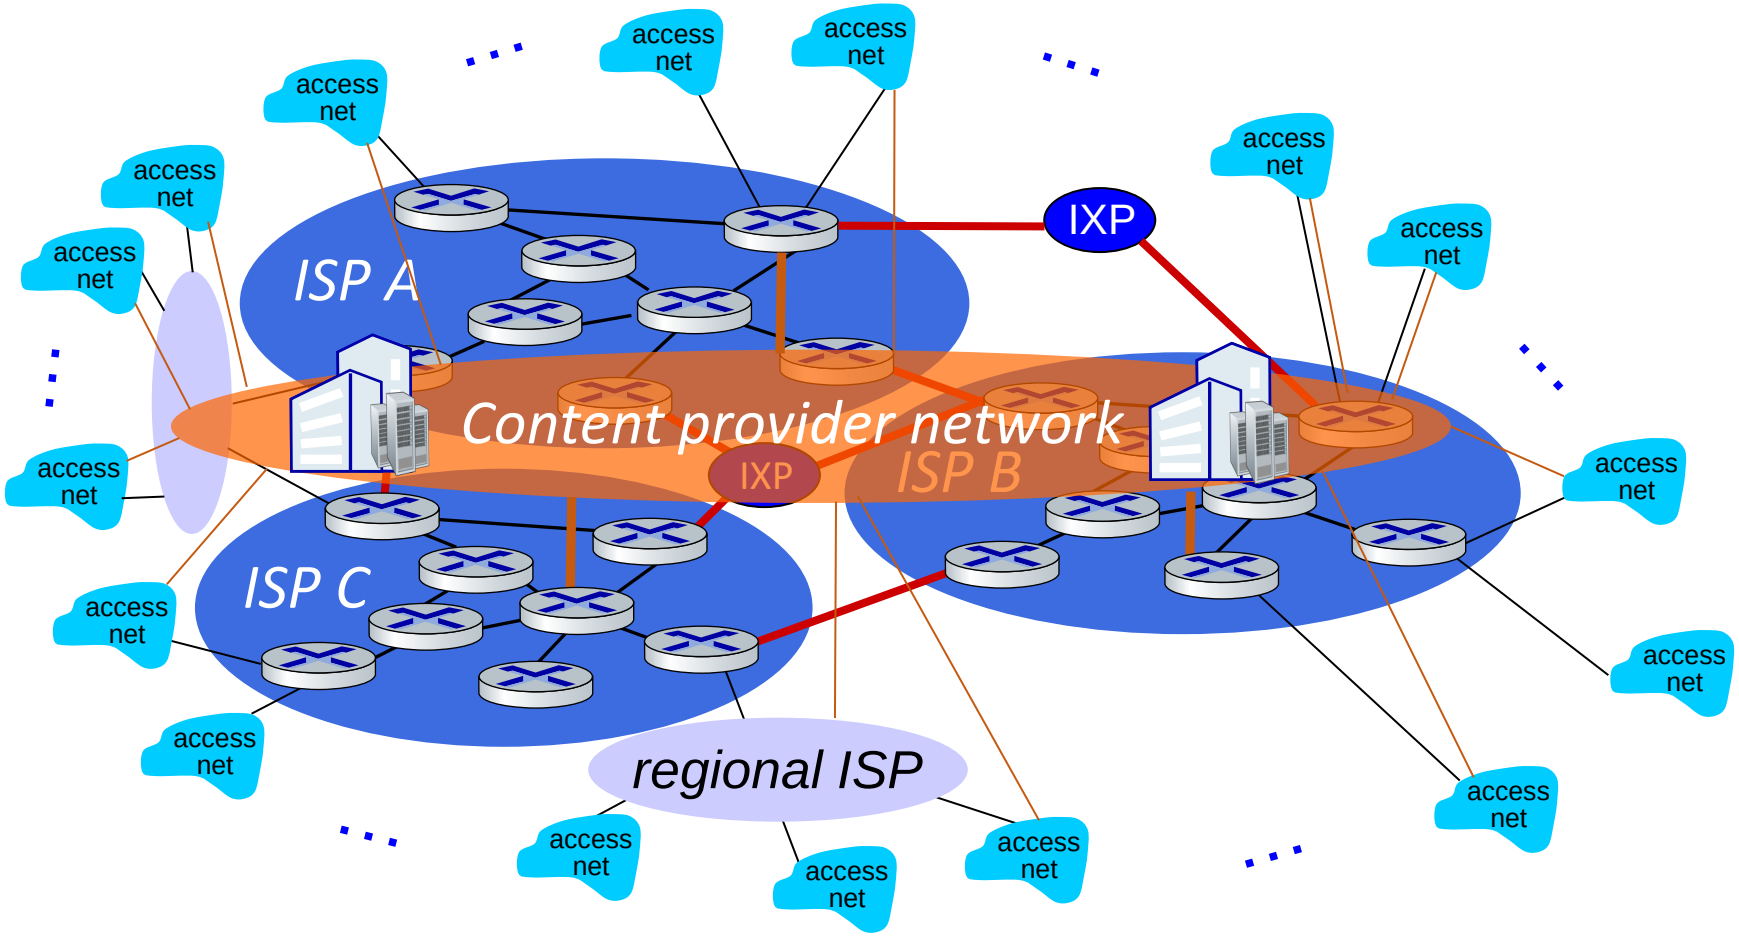

Internet structure: a "network of networks"

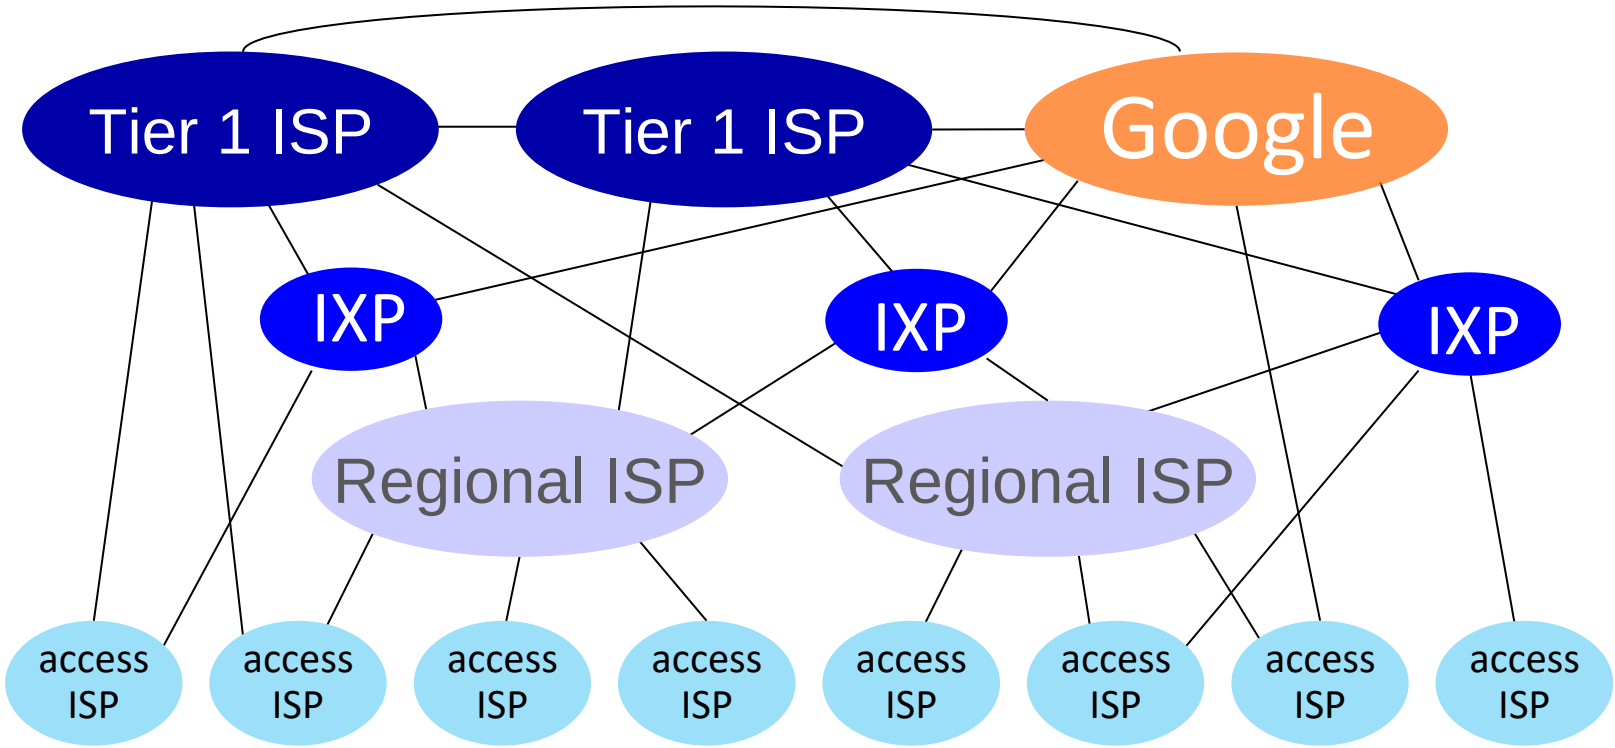

At "center": small # of well-connected large networks

- "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

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?

- and/or our *discussion* of networks?

Example: „Researchers on Boats“

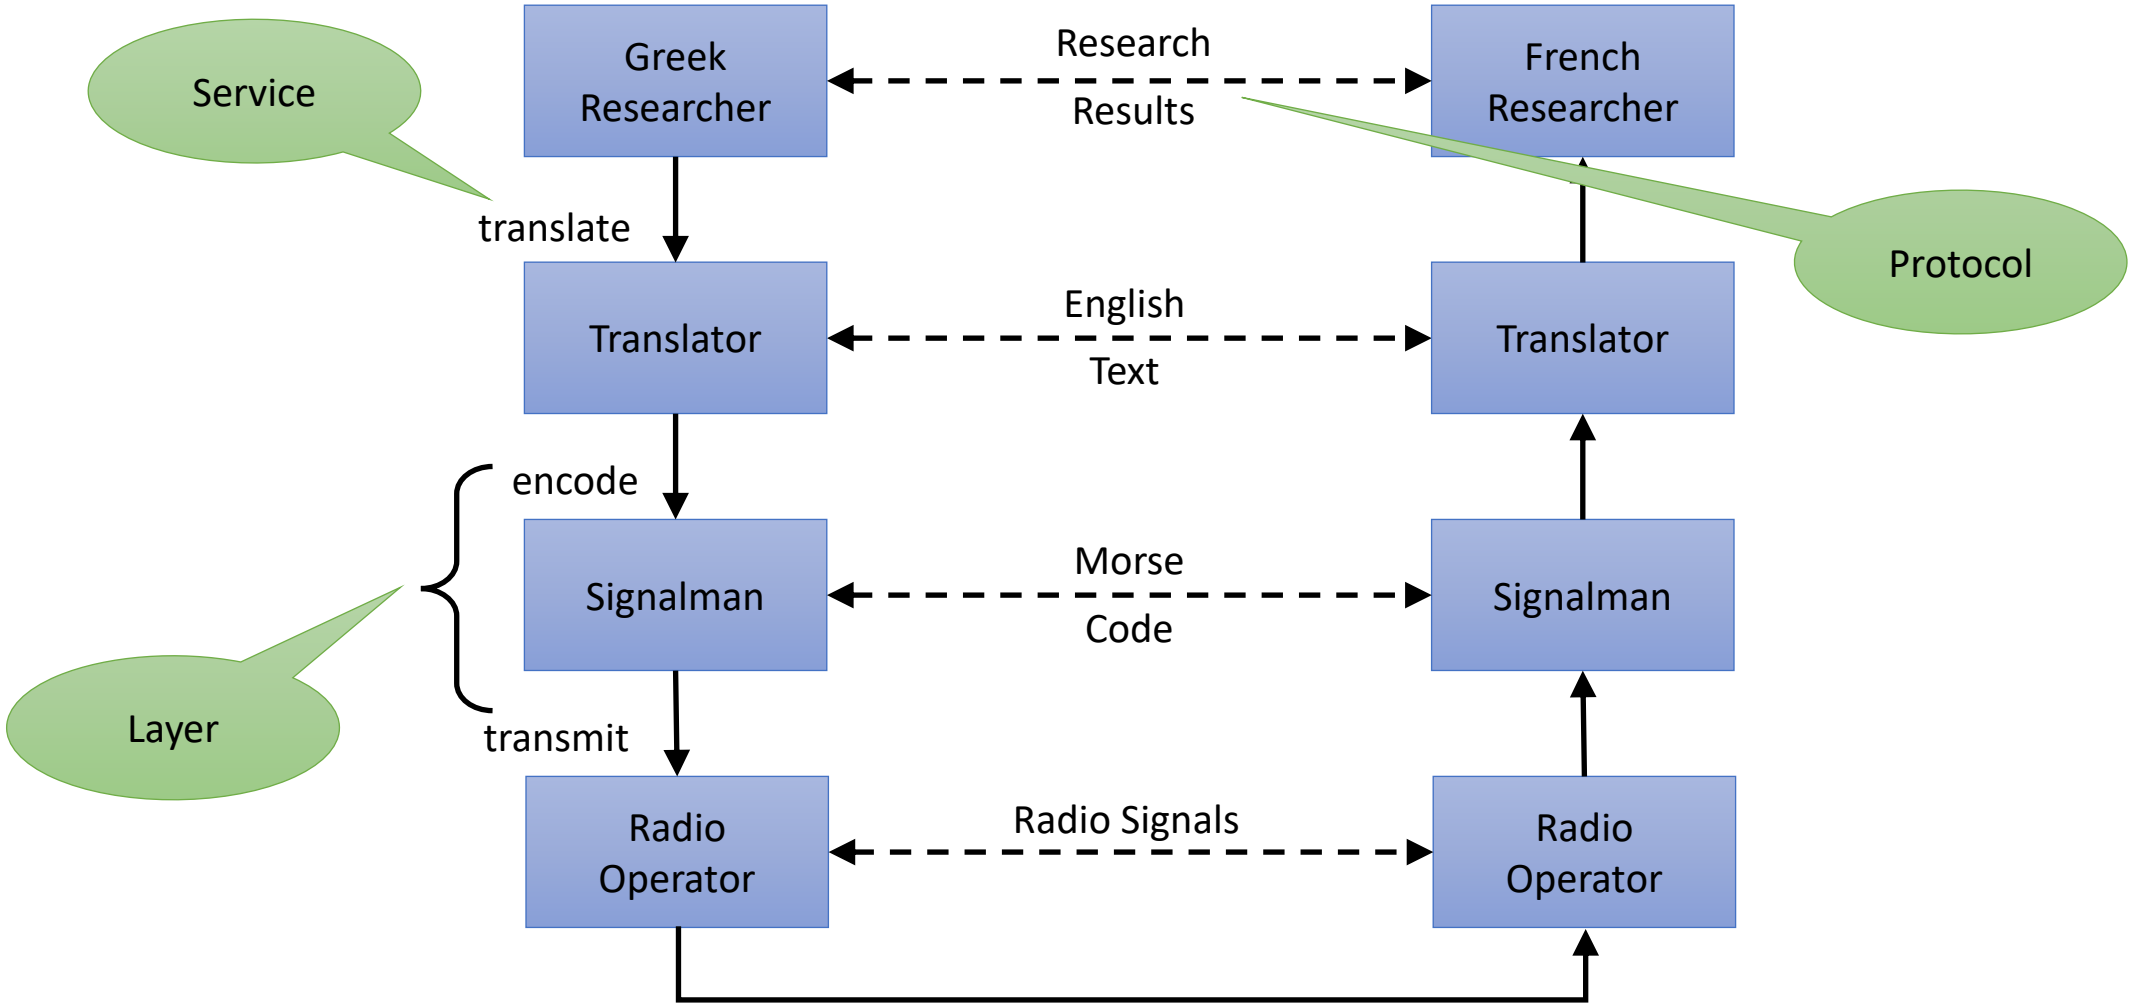

Why layering?

Approach to designing/discussing complex systems:

- explicit structure allows identification, relationship of 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, e.g.:
    - radio transmission can be easily exchanged by light signals
    - all lower layers work the same way for "German and Italian cook exchange soup recipes"

Layered Internet protocol stack

- *application:* supporting network applications
  - HTTP, IMAP, SMTP, DNS
- *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”

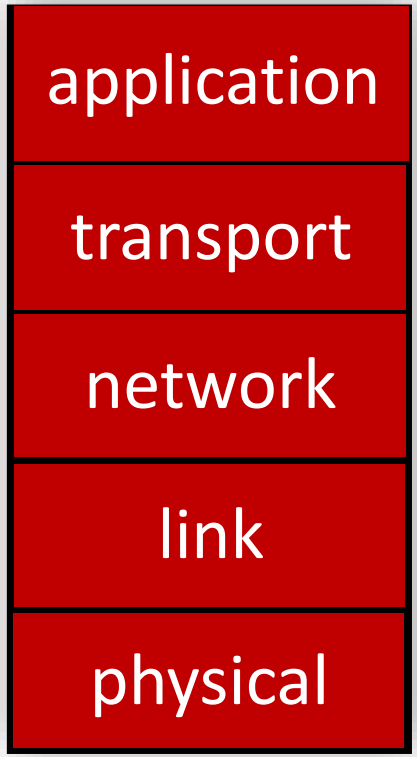

Encapsulation: an end-end view

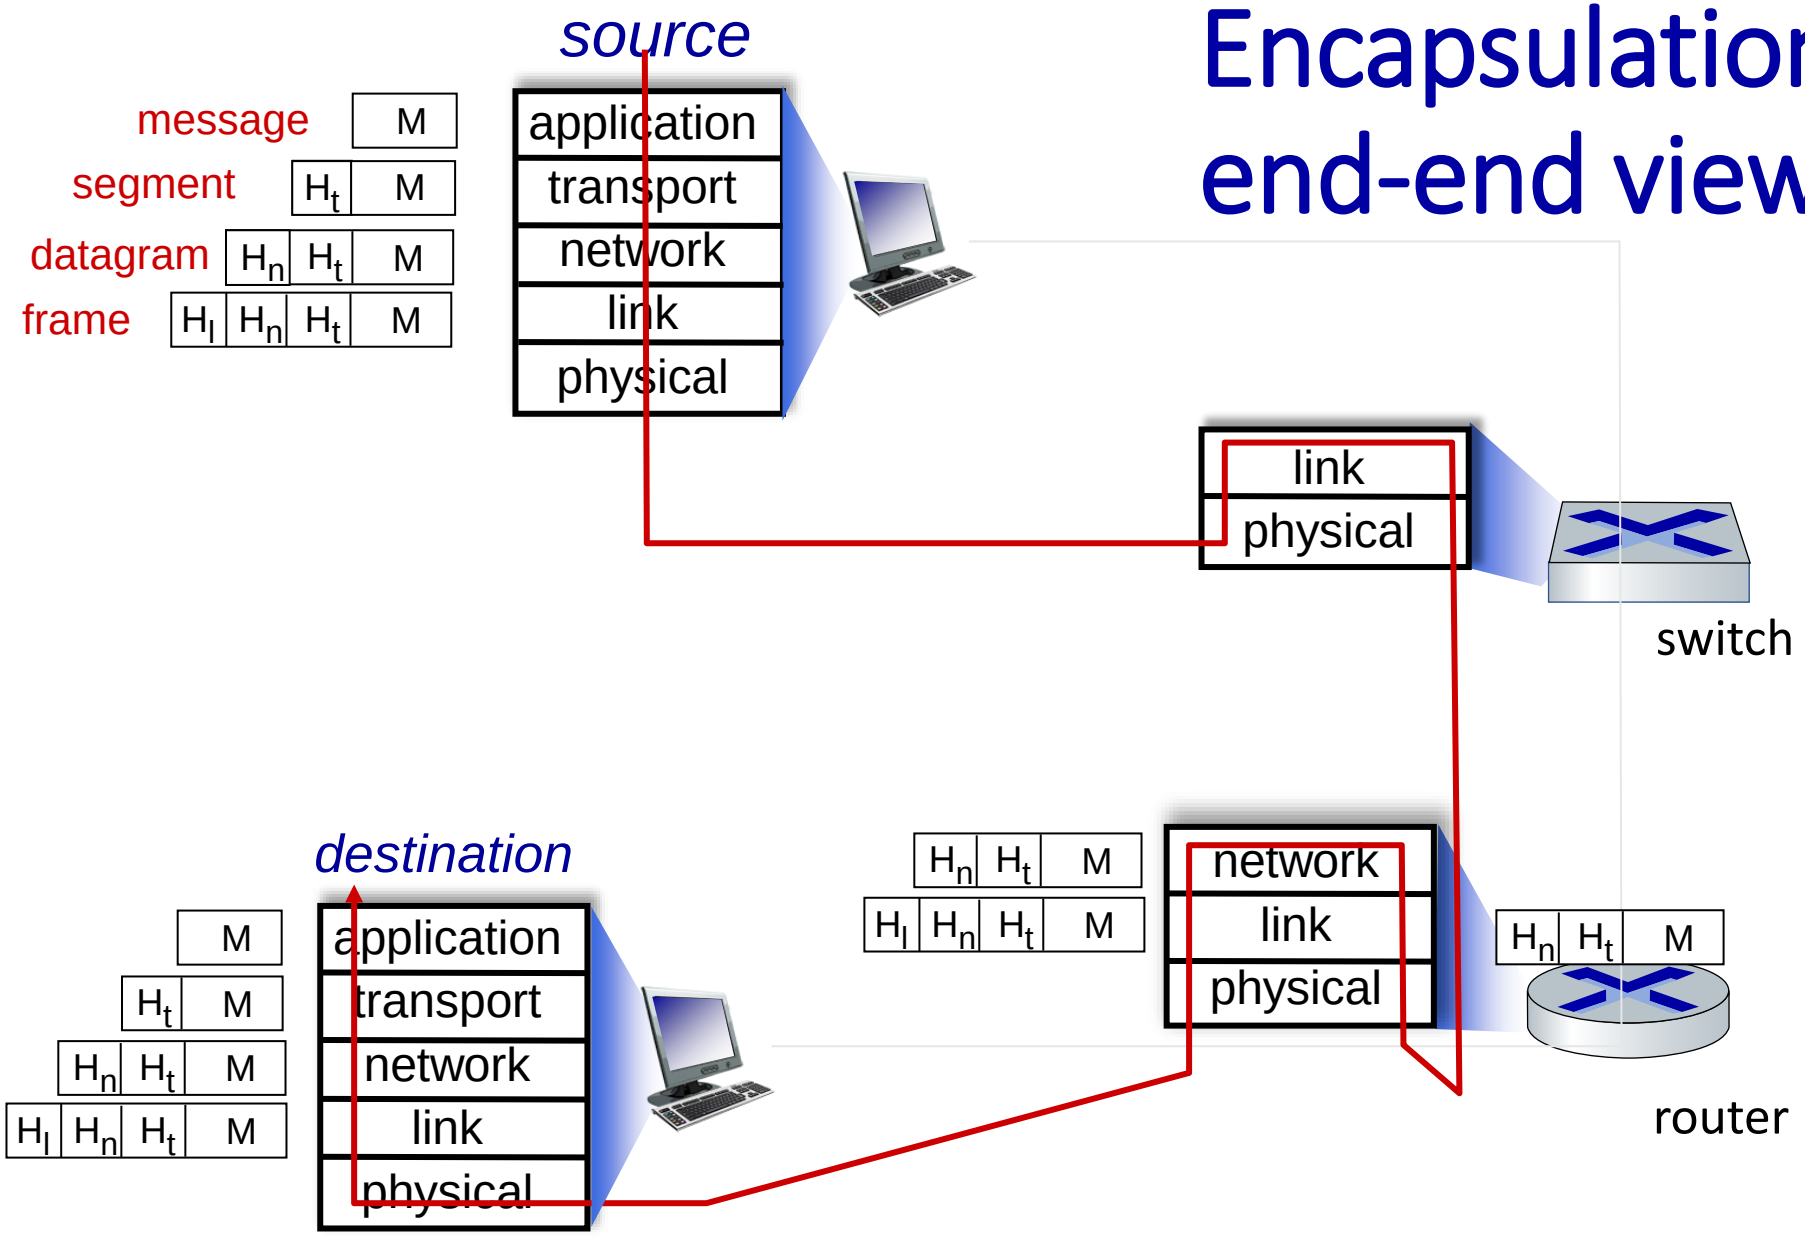

Internet history

- 1957:
  - USSR launches first satellite “Sputnik”
  - In response, the US establishes the Advanced Research Projects Agency (ARPA)
  - Goal: “*to establish US lead in science and technology applicable to the military*”
- 1962:

The Computer as a Communication Device

J.C.R. Licklider and Robert W. Taylor

Reprinted from *Science and Technology*, April 1968.

©Science and Technology 1968

Internet history

- 1965:
  - E-mail allows different users of a time-share mainframe computer to communicate with each other
- 1969:
  - The Internet predecessor ARPANET starts
  - One of the first experiments:
  - "Leonard Kleinrock, a pioneering computer scientist at UCLA, and his small group of graduate students hoped to log onto the Stanford computer and try to send it some data. They would start by typing 'login', and seeing if the letters appeared on the far off monitor."

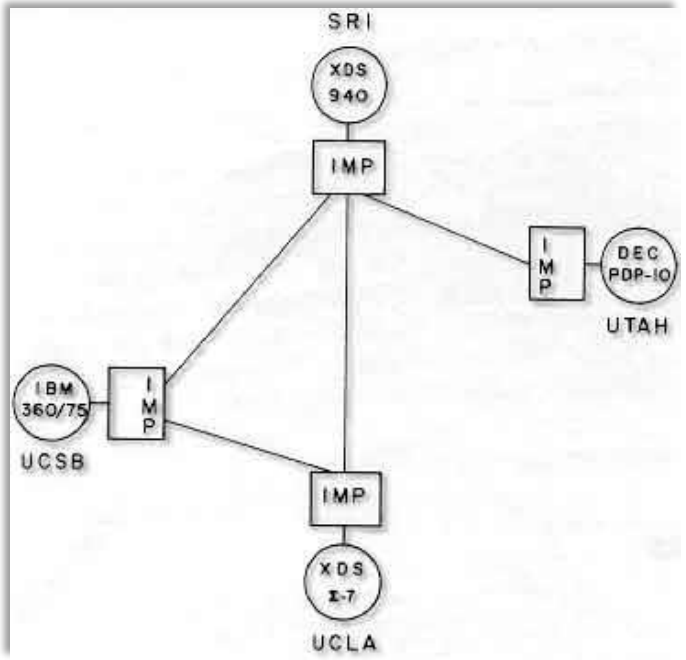

Quelle: http://www.netvalley.com/intval1.html

Internet history

- 1969 (contd.):
  - Kleinrock in an interview:
  - "We set up a telephone connection between us and the guys at SRI … We typed the L and we asked on the phone:
  - "Do you see the L?"
    "Yes, we see the L," came the response.
    We typed the O, and we asked,
    "Do you see the O."
    "Yes, we see the O.“
  - Then we typed the G,
    and the system crashed …
  - Yet a revolution had begun."

Leonard Kleinrock

Internet history

- The revolutionary idea for the Internet: **packet switching**
  - invented by Paul Baran, Donald Davies and Leonard Kleinrock
  - originally: better utilization of line capacity
  - in ARPANET: "survivability of communications systems in the event of a nuclear attack"
- “Proof” of this property:
  - An important Internet node was below the World Trade Center
  - 9/11 attacks have not disrupted the Internet
  - Packets were automatically sent via alternate routes

Internet history

- 1969:
  - Telnet: Login to a remote computer
- 1971:
  - FTP: File Transfer Protocol
- 1973:
  - RFC-561: First proposal for a standardized email format
- 1978:
  - Multi User Dungeon: First online multiplayer game

Source: Wikipedia

Internet history

- 1981
  - RFC-791: Internet Protocol (Version 4)
  - RFC-793: Transmission Control Protocol
- 1983:
  - Start of the Internet

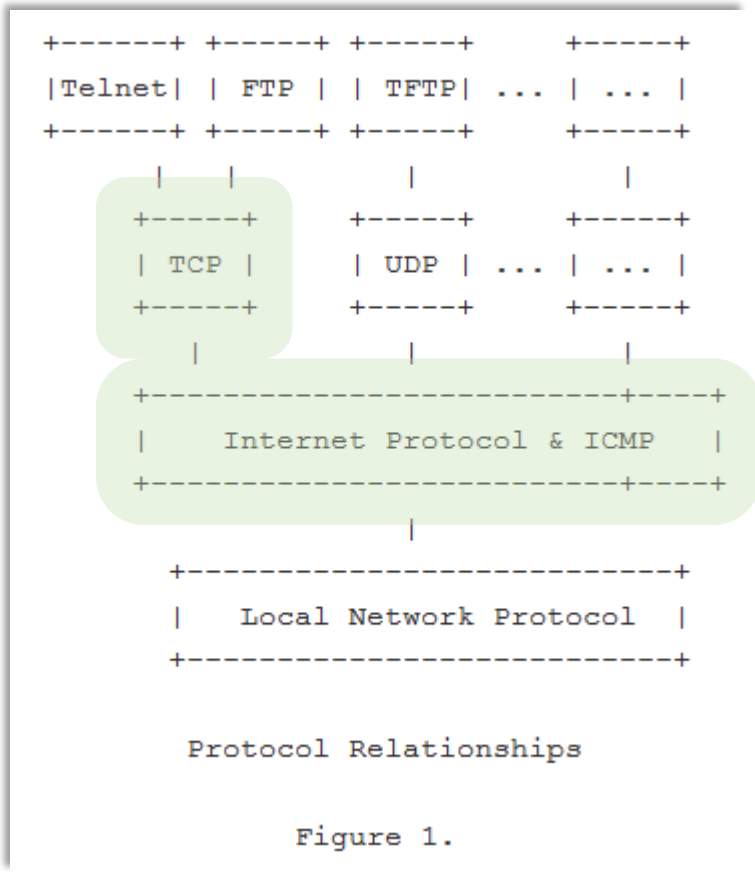

from RFC-791

Internet history

- 1992: World Wide Web:
  - developed by Tim Berners-Lee and Robert Cailliau

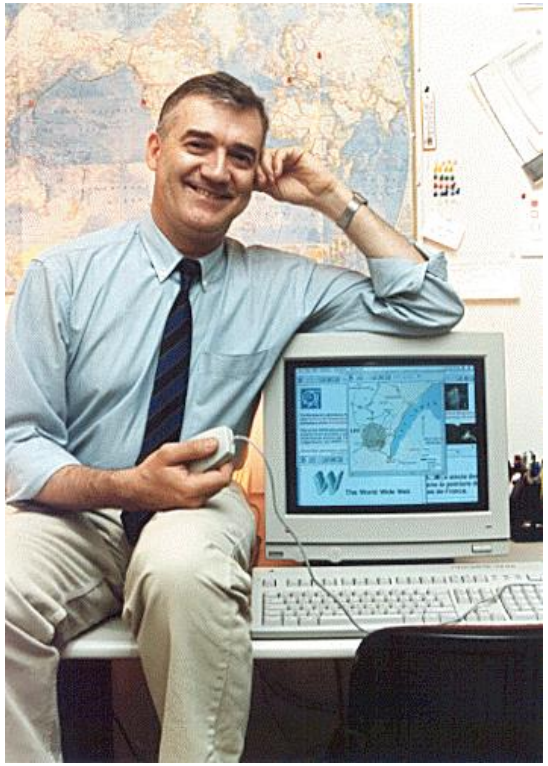

Internet history

- World Wide Web:
  - Original purpose: Access to information within CERN

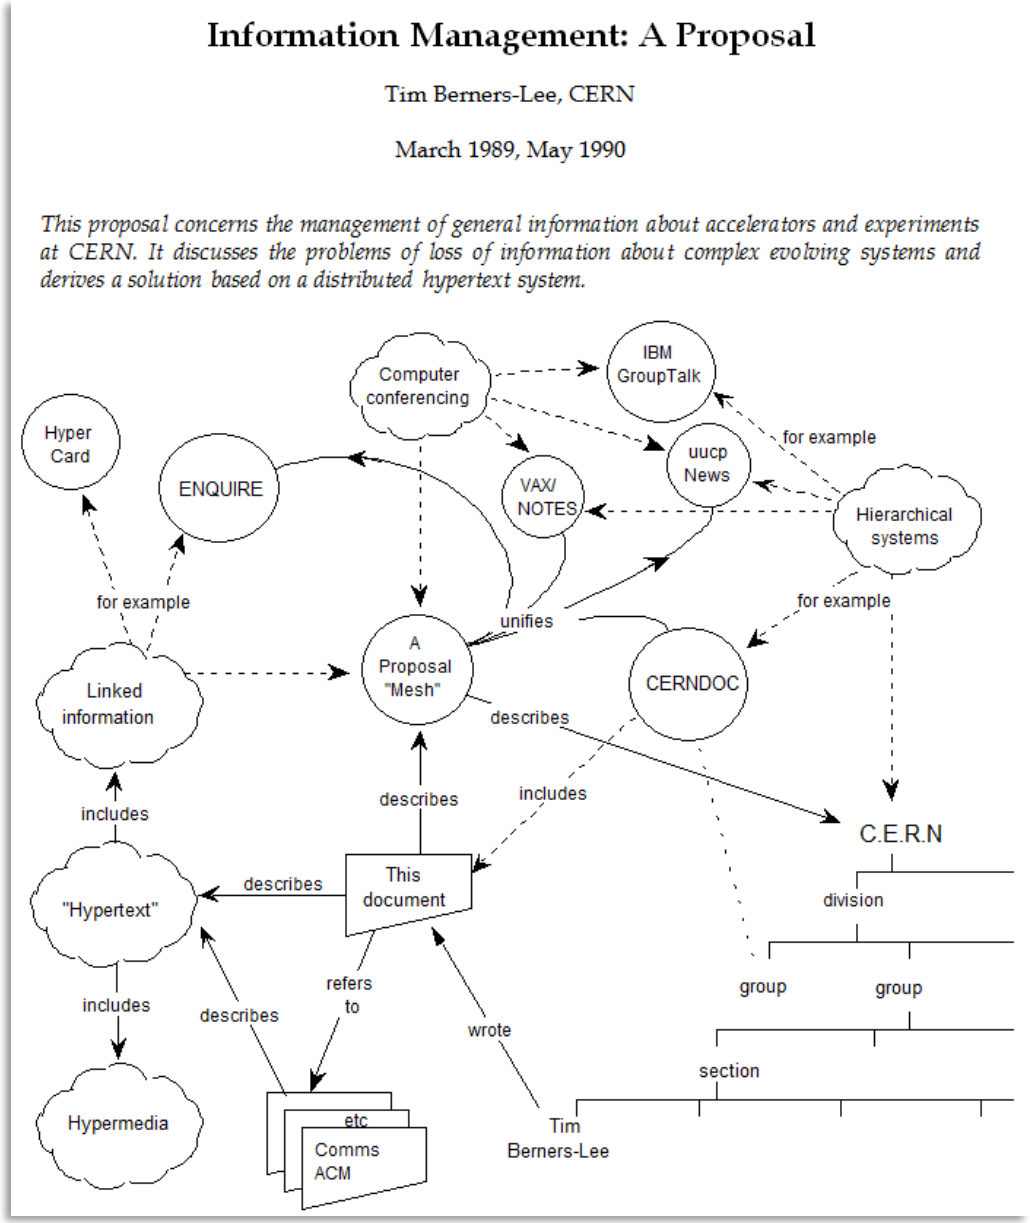

Information Management: A Proposal

Tim Berners-Lee, CERN

March 1989, May 1990

This proposal concerns the management of general information about accelerators and experiments at CERN. It discusses the problems of loss of information about complex evolving systems and derives a solution based on a distributed hypertext system.

Internet history

- 1994:
  - Amazon.com
- 1995:
  - Ebay
- 1998:
  - Google
- 2001:
  - Wikipedia
- 2004:
  - Facebook

World Wide Web (WWW)

Web and HTTP

First, a quick review…

- web page consists of *objects,* each of which can be stored on different Web servers
- object can be HTML file, JPEG image, Java applet, audio file,…
- web page consists of *base HTML-file* which includes *several referenced objects, each* addressable by a *URL,* e.g.,

`www.someschool.edu/someDept/pic.gif`

host name | path name

HTTP overview

HTTP: hypertext transfer protocol

- Web's application-layer protocol
- client/server model:
  - *client:* browser that requests, receives, (using HTTP protocol) and "displays" Web objects
  - *server:* Web server sends (using HTTP protocol) objects in response to requests

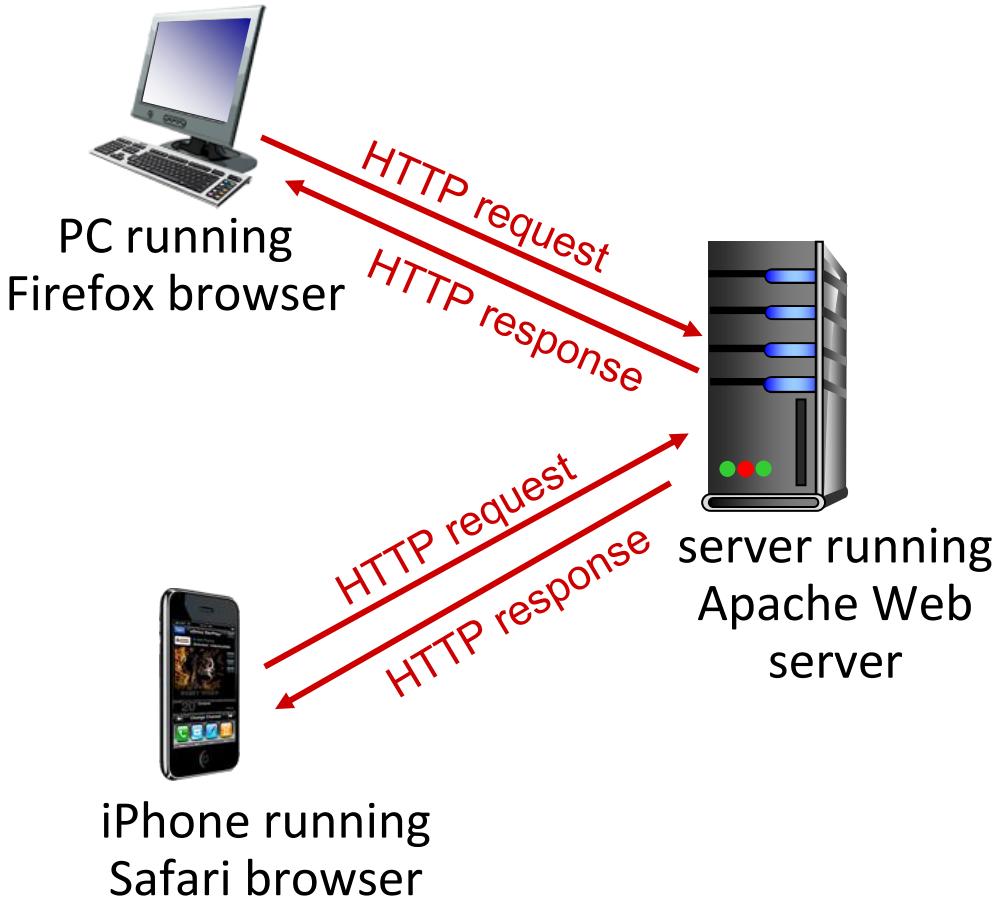

HTTP request message

- two types of HTTP messages: *request, response*
- HTTP request message:
  - ASCII (human-readable format)

carriage return character
line-feed character

request line (GET, POST,
HEAD commands)

carriage return, line feed
at start of line indicates
end of header lines

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

HTTP response message

status line (protocol status code status phrase) → `HTTP/1.1 200 OK`

HTTP response status codes

- status code appears in 1st line in server-to-client response message.
- some sample codes:

200 OK

- request succeeded, requested object later in this message

301 Moved Permanently

- requested object moved, new location specified later in this message (in Location: field)

400 Bad Request

- request msg not understood by server

404 Not Found

- requested document not found on this server

505 HTTP Version Not Supported

Maintaining user/server state: cookies

HTTP GET/response interaction is *stateless*

- server maintains *no* information about past client requests
- no notion of multi-step exchanges of HTTP messages to complete a Web "transaction"
  - no need for client/server to track "state" of multi-step exchange
  - all HTTP requests are independent of each other
  - no need for client/server to "recover" from a partially-completed-but-never-completely-completed transaction
- However …
  - Some applications require a "state", e.g.
    - Shopping: Which items are in the shopping cart?
    - Banking: Is the user already logged in?

Maintaining user/server state: cookies

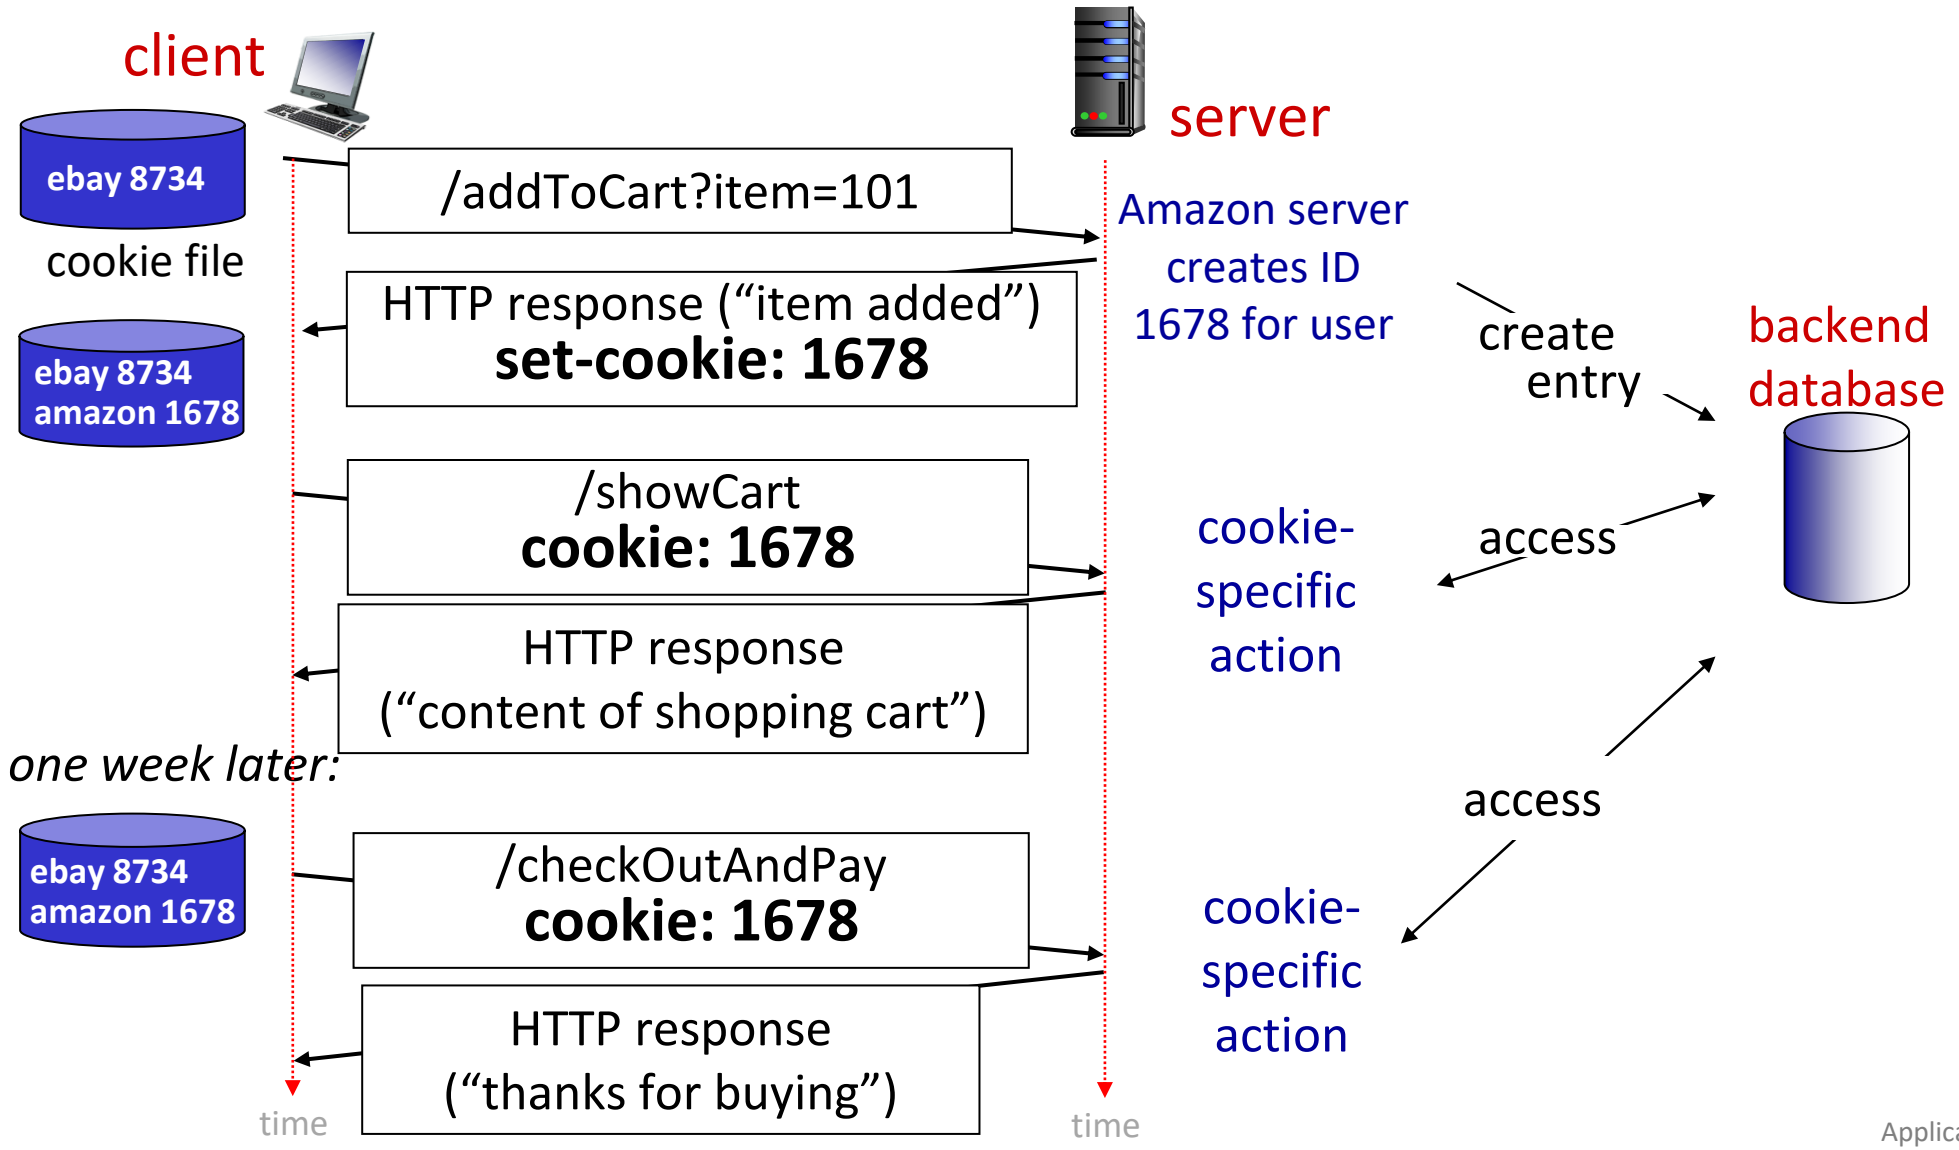

HTTP cookies: comments

What cookies can be used for:

- authorization
- shopping carts
- recommendations
- user session state (Web e-mail)

Challenge: How to keep state?

- *at protocol endpoints:* maintain state at sender/receiver over multiple transactions
- *in messages:* cookies inHTTP messages carry state

aside

cookies and privacy:

- cookies permit sites to *learn* a lot about you on their site.
- third party persistent cookies (tracking cookies) allow common identity (cookie value) to be tracked across multiple web sites

Network security

- Internet + WWW 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!
- We now need to think about:
  - how bad guys can attack computer networks
  - how we can defend networks against attacks
  - how to design architectures that are immune to attacks