LECTURE 1 - INTRODUCTION

Jacob Aae Mikkelsen

GOALS TODAY

- History of networks
- Get "feel" and terminology
  - More depth and details *later* in the course
- Approach: Use Internet as example

ORIGIN OF THE INTERNET

ARPANET 1969

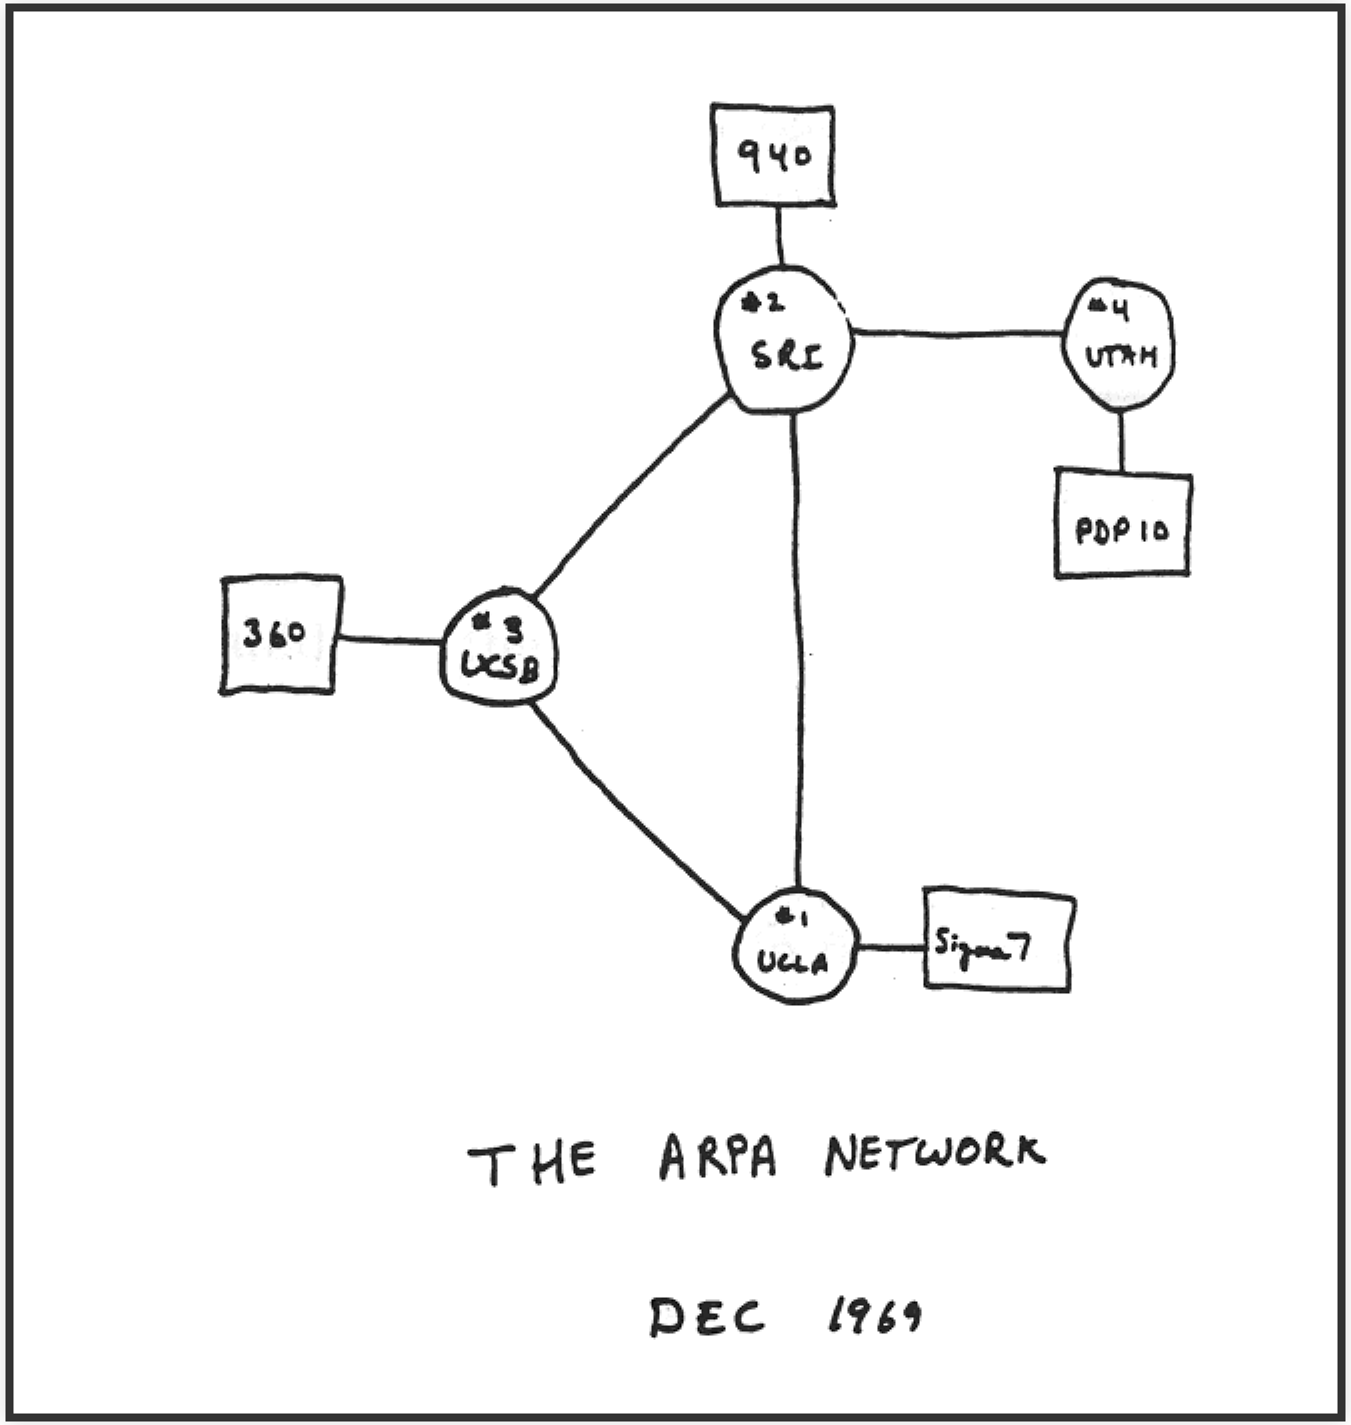

EARLY PACKET-SWITCHING PRINCIPLES

- [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
- [1970] ALOHAnet satellite network in Hawaii

ARPANET 1969

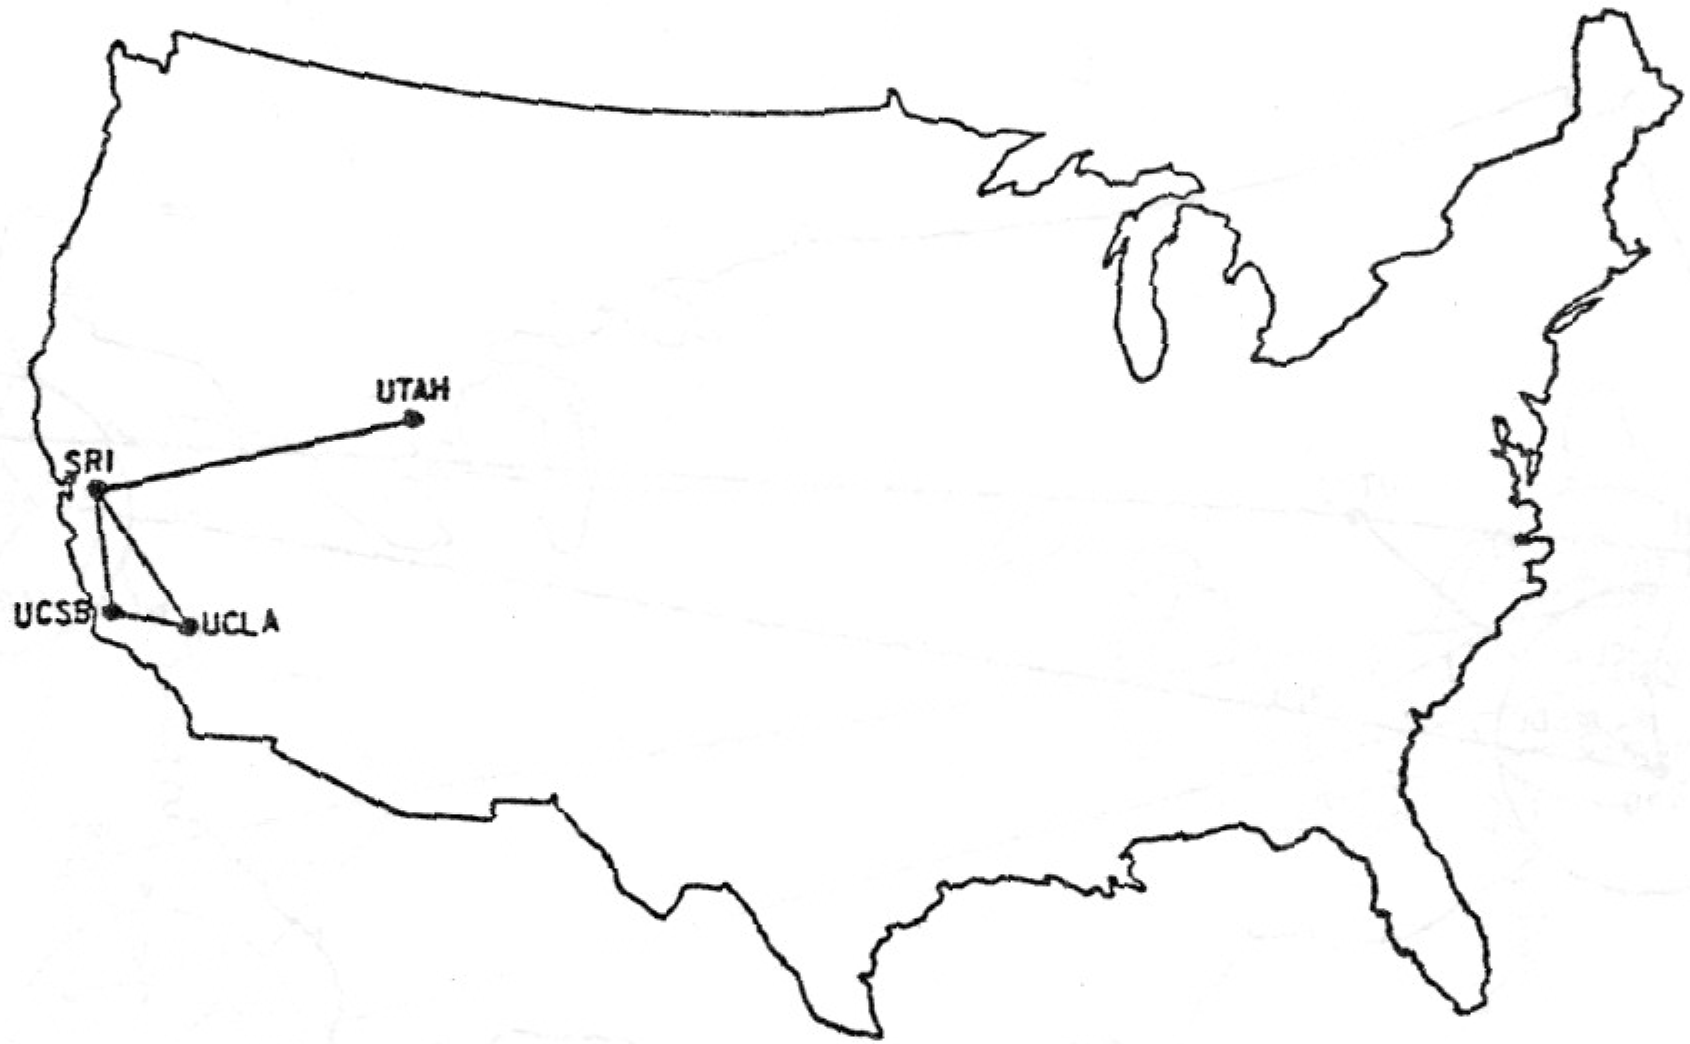

The ARPANET in December 1969

ARPANET 1970

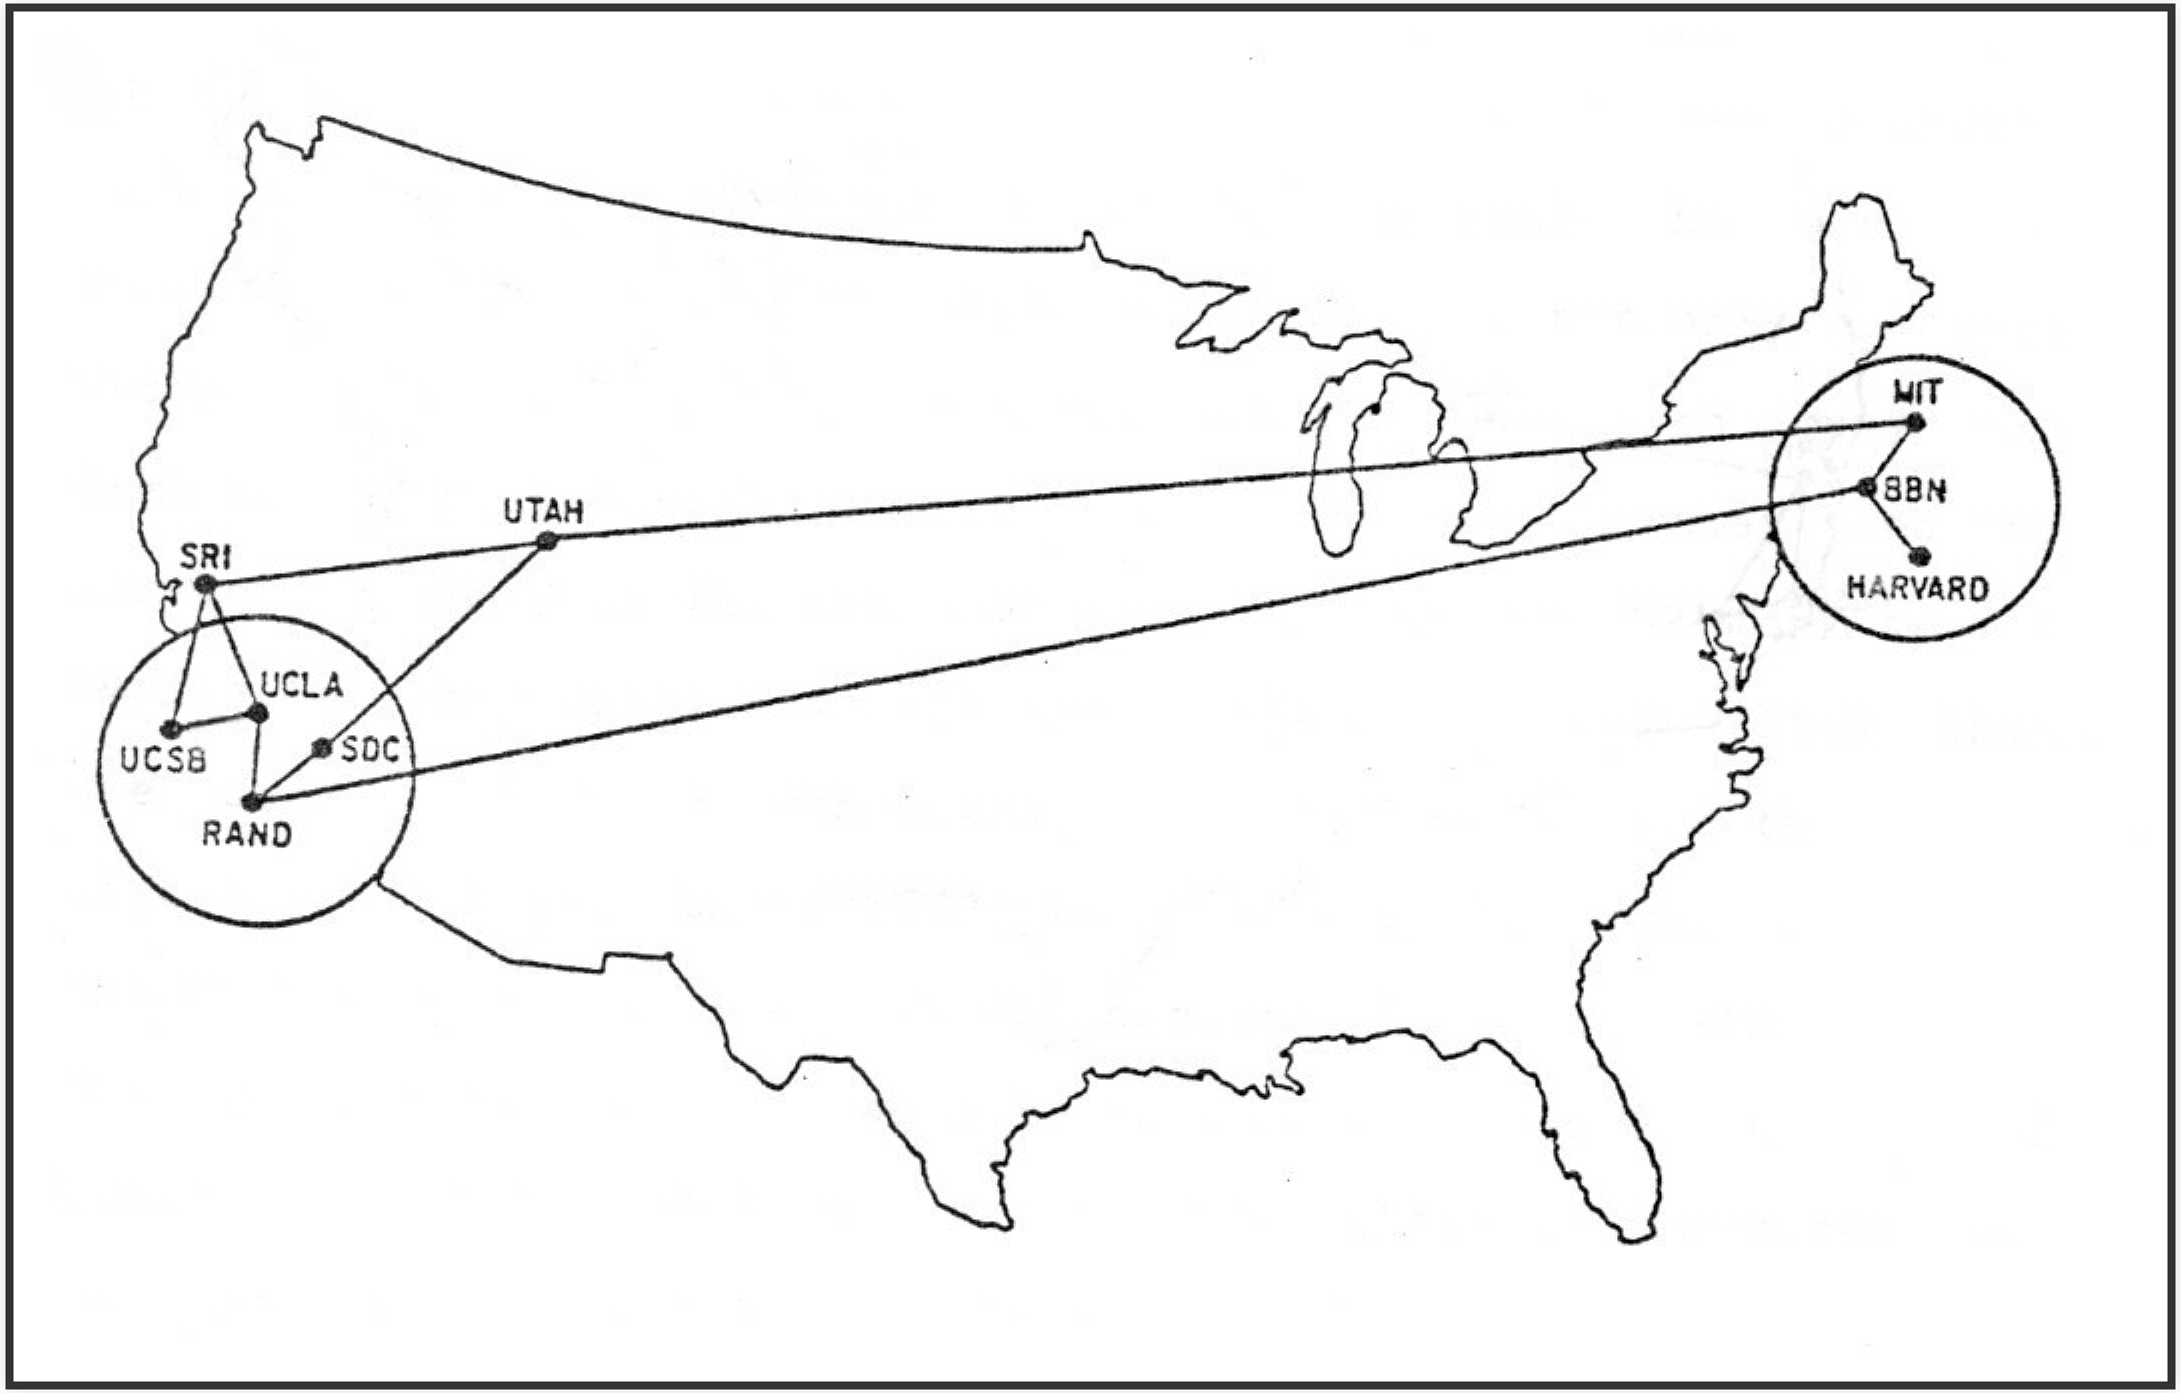

ARPANET 1972

- [1972] ARPAnet public demo
  - NCP (Network Control Protocol) first host-host protocol
  - First e-mail program
  - ARPAnet has 15 nodes

ARPANET SEP 1973

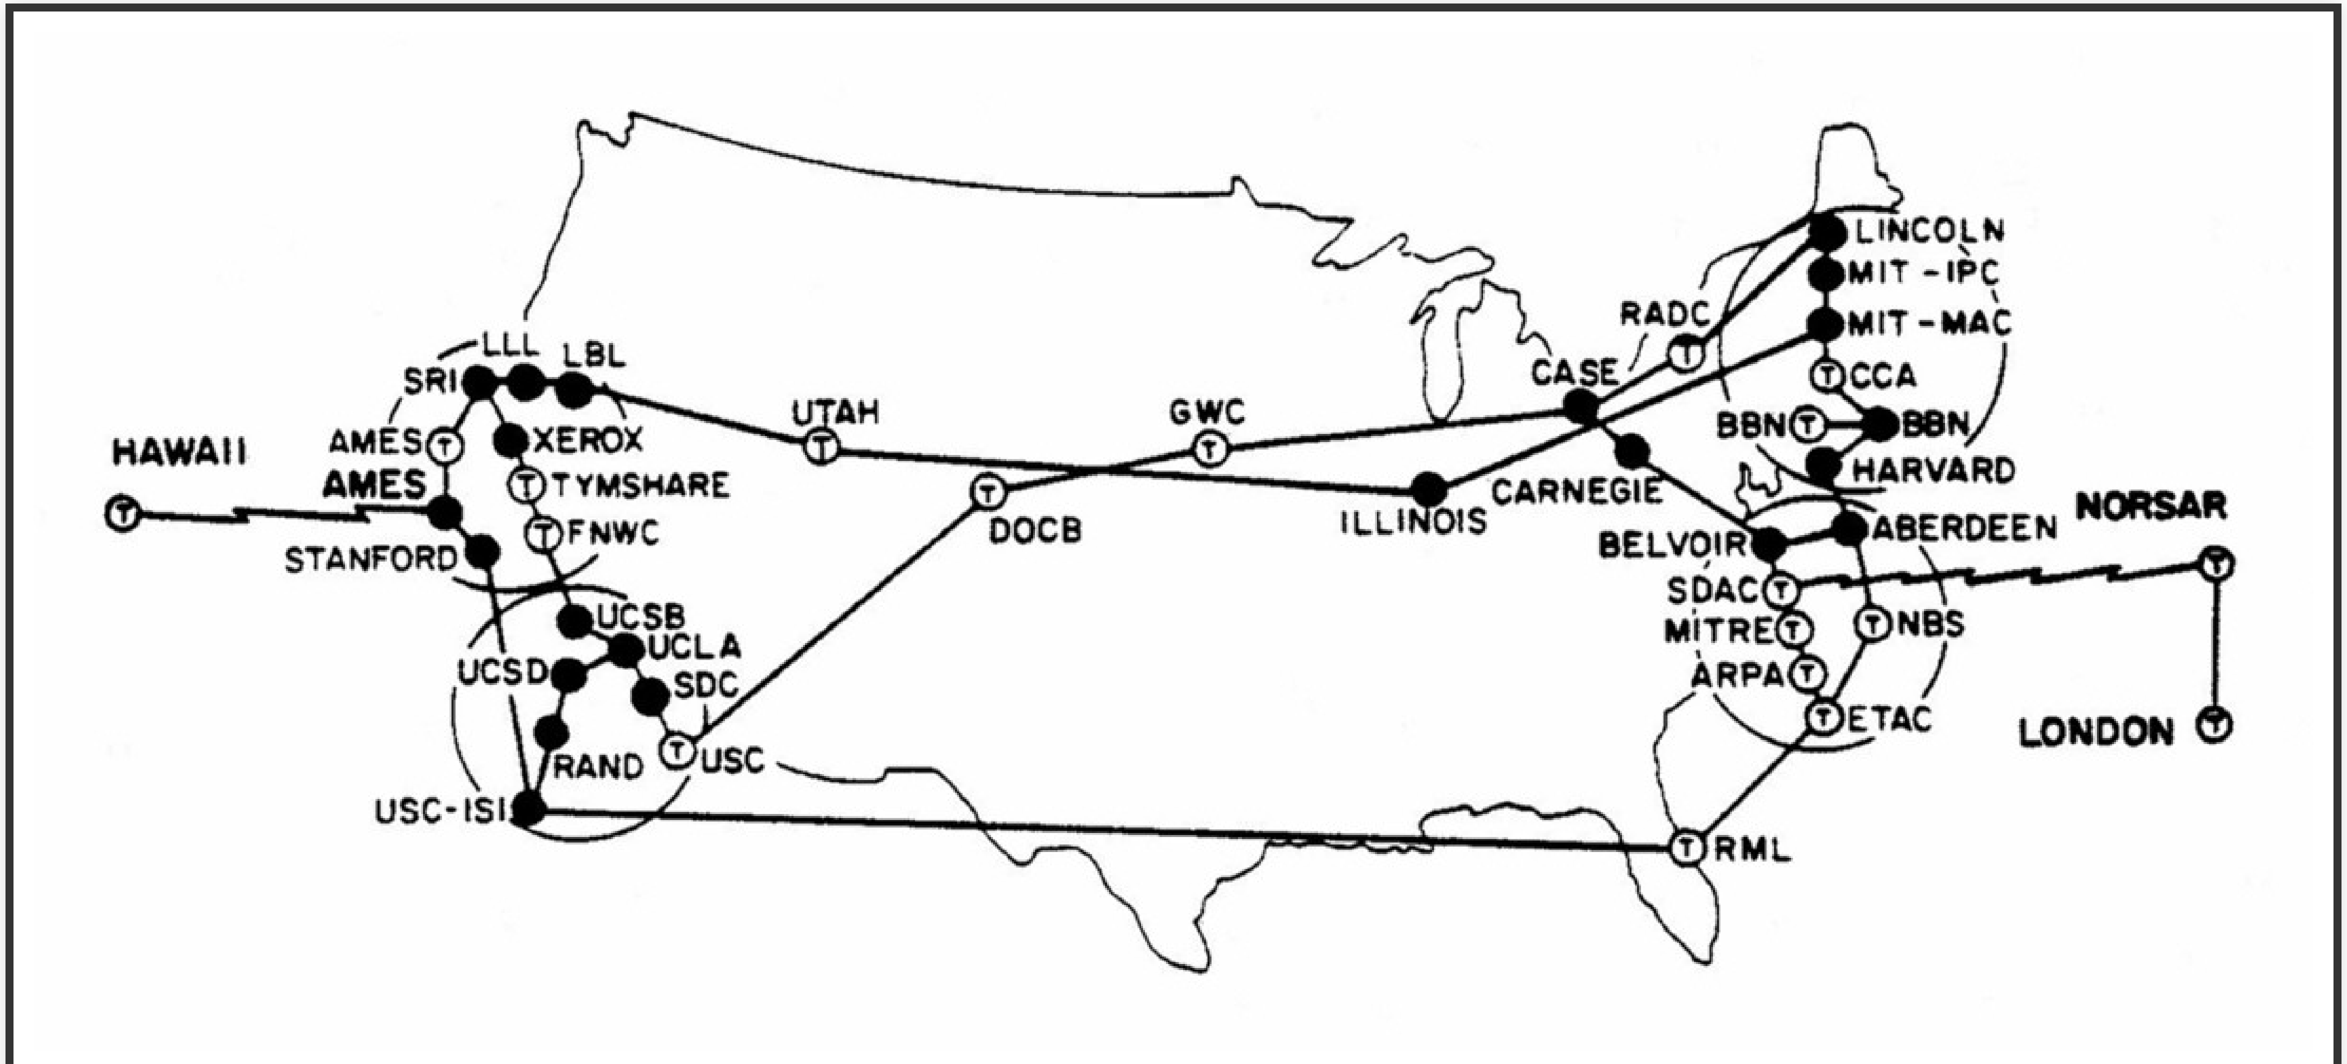

INTERNETWORKING

- [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

ARPANET FEB. 1982

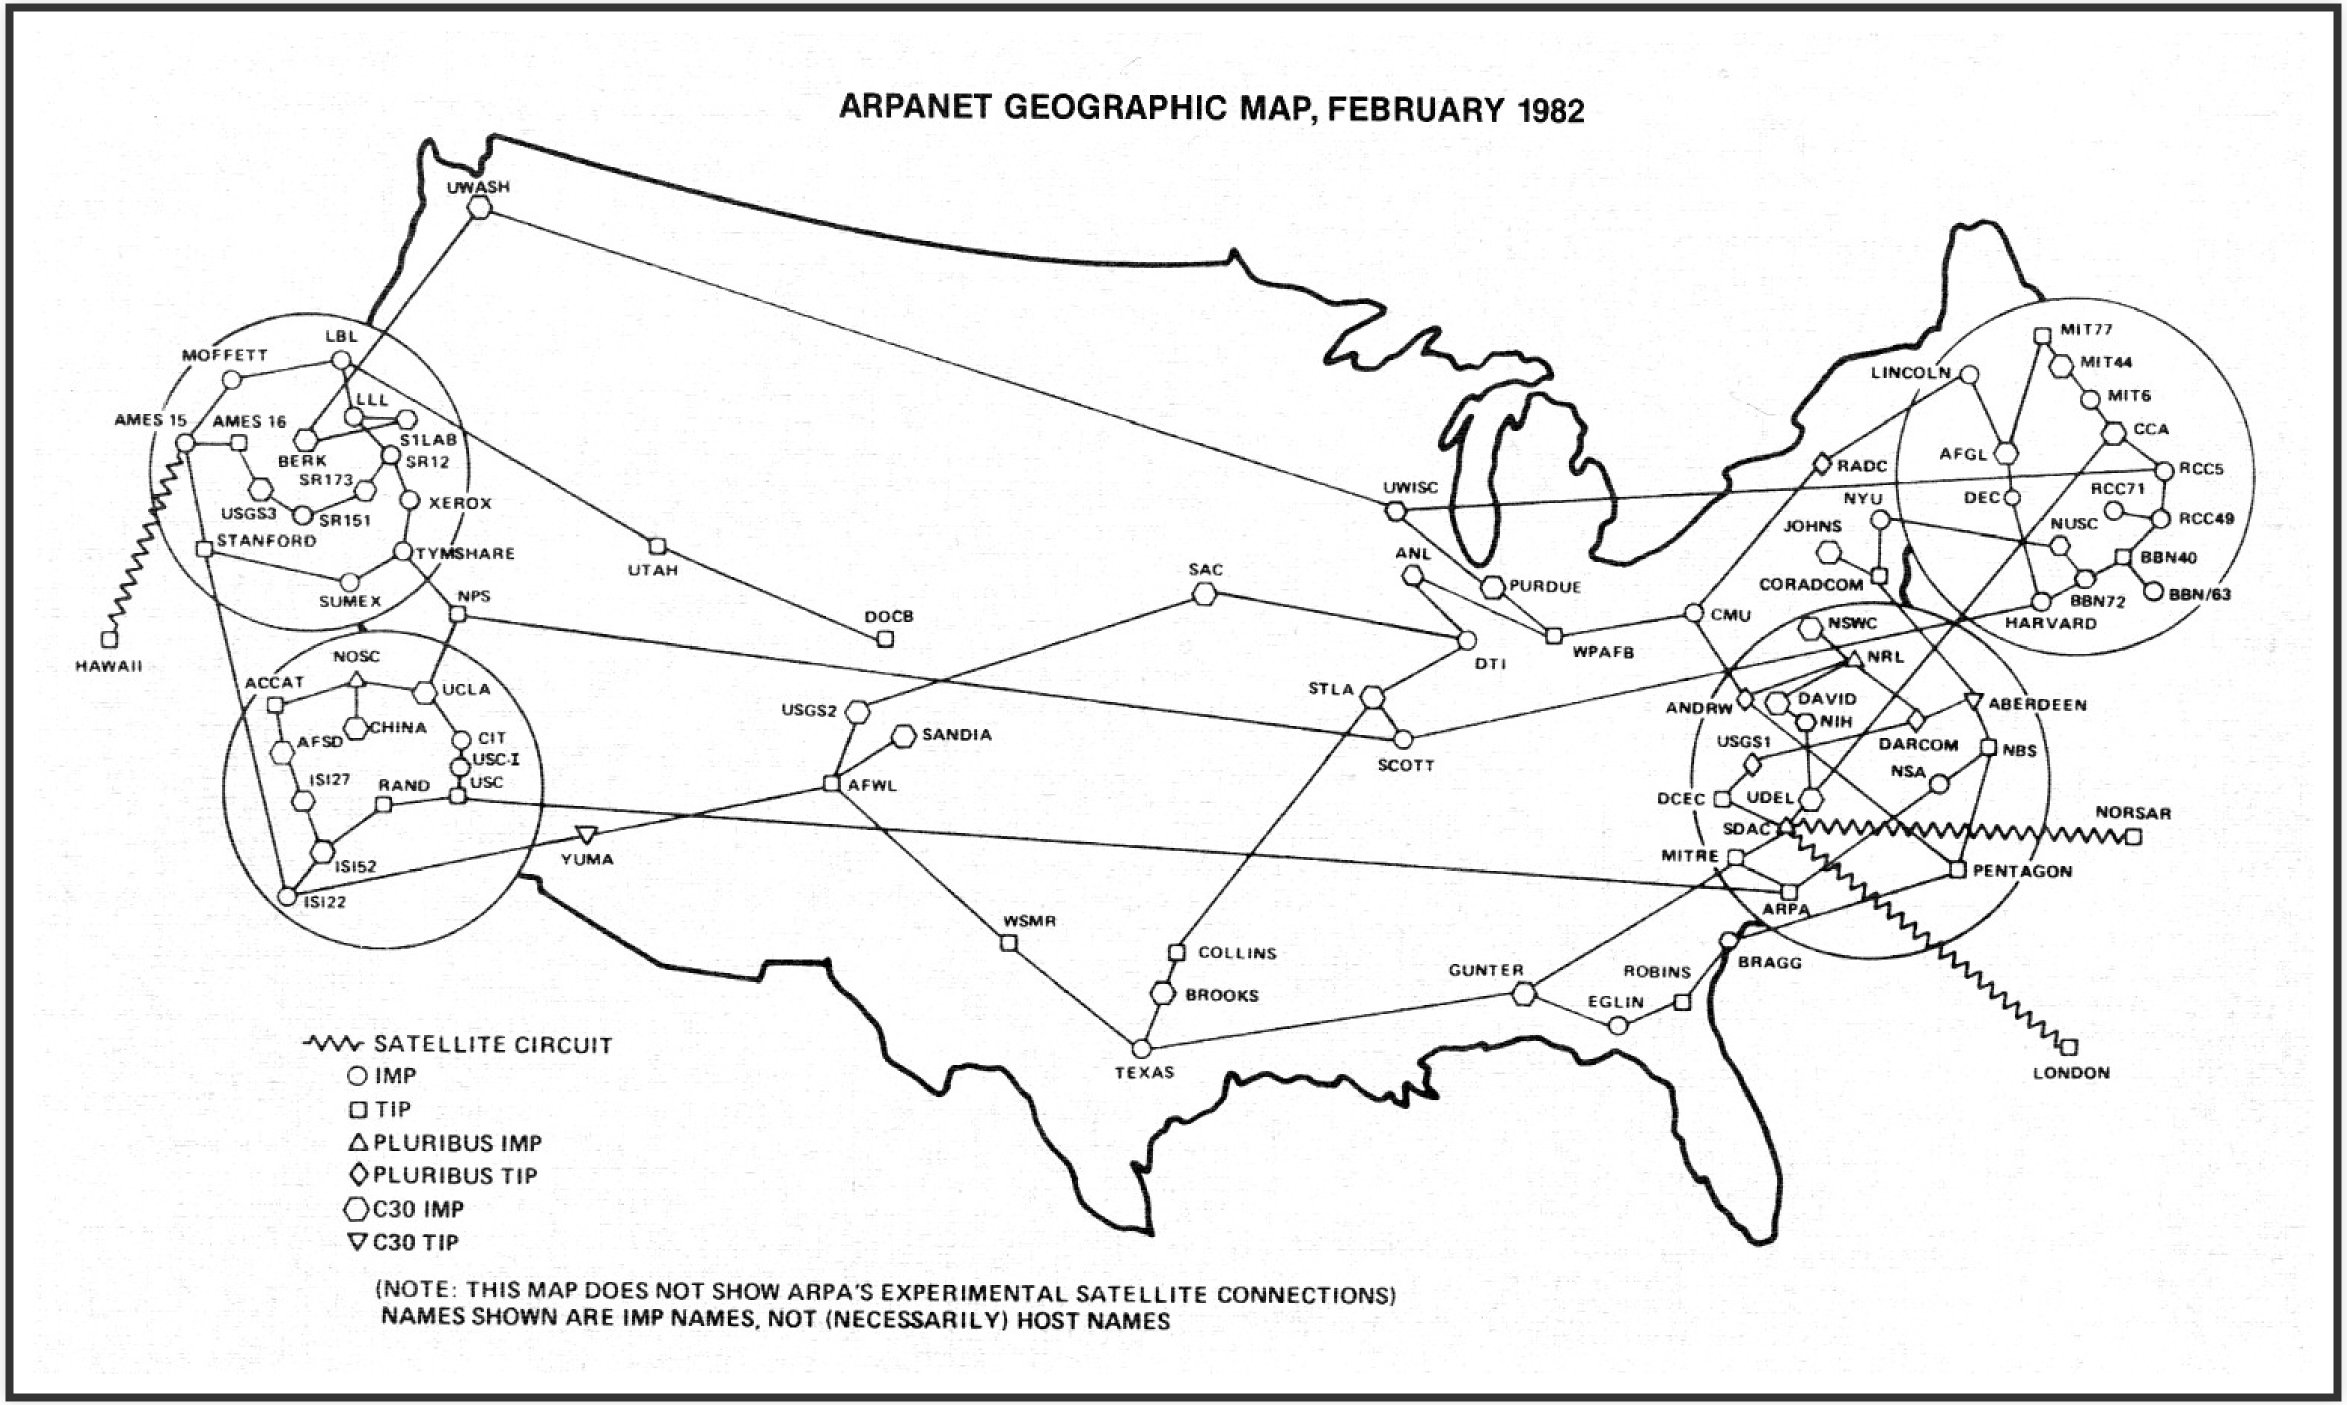
ARPANET GEOGRAPHIC MAP, FEBRUARY 1982
UWASH
MOFFETT
LBL
LLL
AMES 15
AMES 16
S1LAB
BERK
SR12
SR173
XEROX
USGS3
SR151
STANFORD
TYMSHARE
SUMEX
NPS
UTAH
DOCB
HAWAII
NOSC
ACCAT
UCLA
CHINA
AFSD
CIT
USC-I
ISI27
RAND
USC
ISI52
ISI22
YUMA
USGS2
SANDIA
AFWL
WSMR
SAC
UWISC
ANL
PURDUE
DTI
WPAFB
STLA
SCOTT
COLLINS
BROOKS
TEXAS
GUNTER
ROBINS
EGLIN
BRAGG
CMU
RADC
LINCOLN
MIT77
MIT44
MIT6
CCA
AFGL
RCC5
DEC
RCC71
NYU
JOHNS
NUSC
RCC49
BBN40
CORADCOM
BBN72
BBN/63
HARVARD
NSWC
NRL
ANDRW
DAVID
NIH
ABERDEEN
USGS1
DARCOM
NBS
NSA
DCEC
UDEL
SDAC
NORSAR
MITRE
PENTAGON
ARPA
LONDON
SATELLITE CIRCUIT
IMP
TIP
PLURIBUS IMP
PLURIBUS TIP
C30 IMP
C30 TIP
(NOTE: THIS MAP DOES NOT SHOW ARPA'S EXPERIMENTAL SATELLITE CONNECTIONS)
NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES

INTERNETWORKING, PRINCIPLES

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

ARPANET JUN. 1984

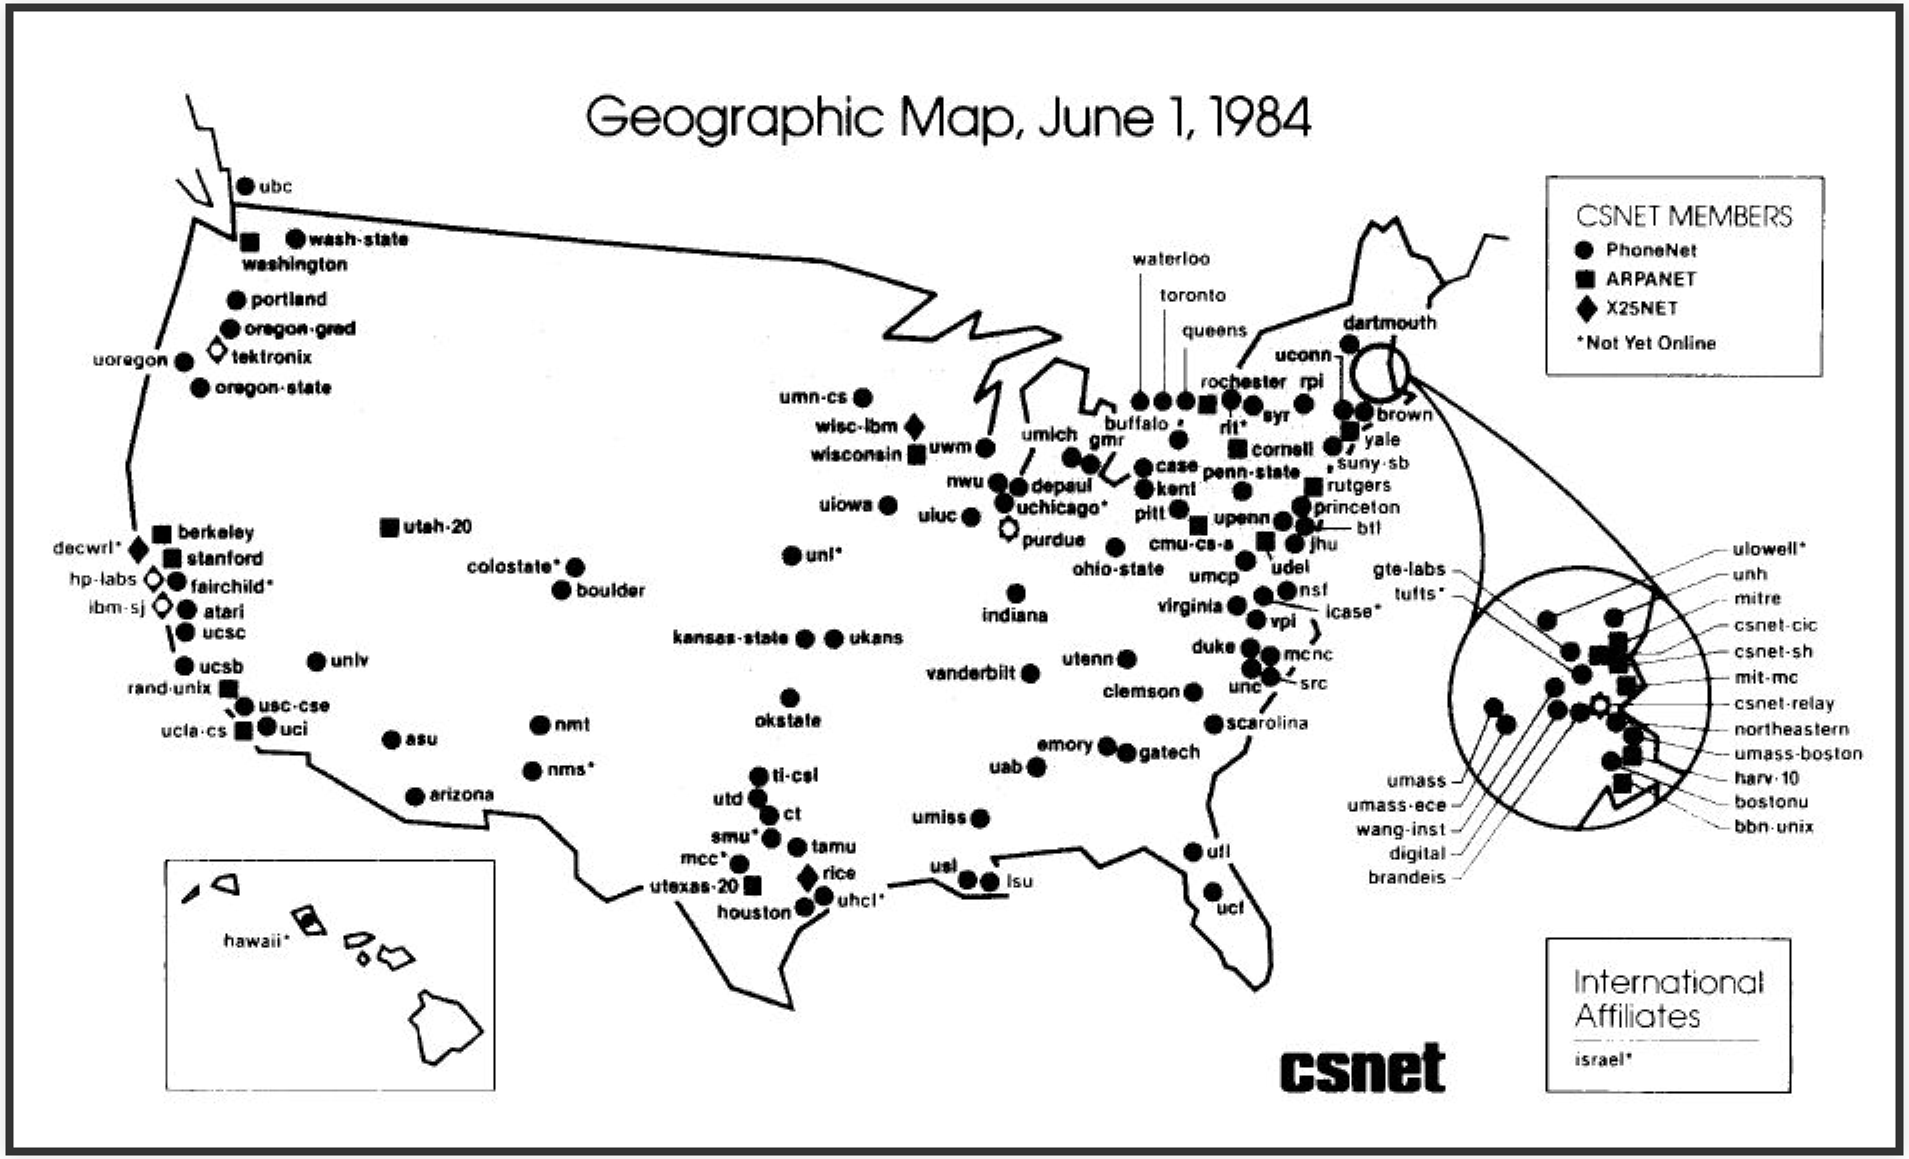

NEW PROTOCOLS

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

NSFNET

NSFNET T3 Network 1992

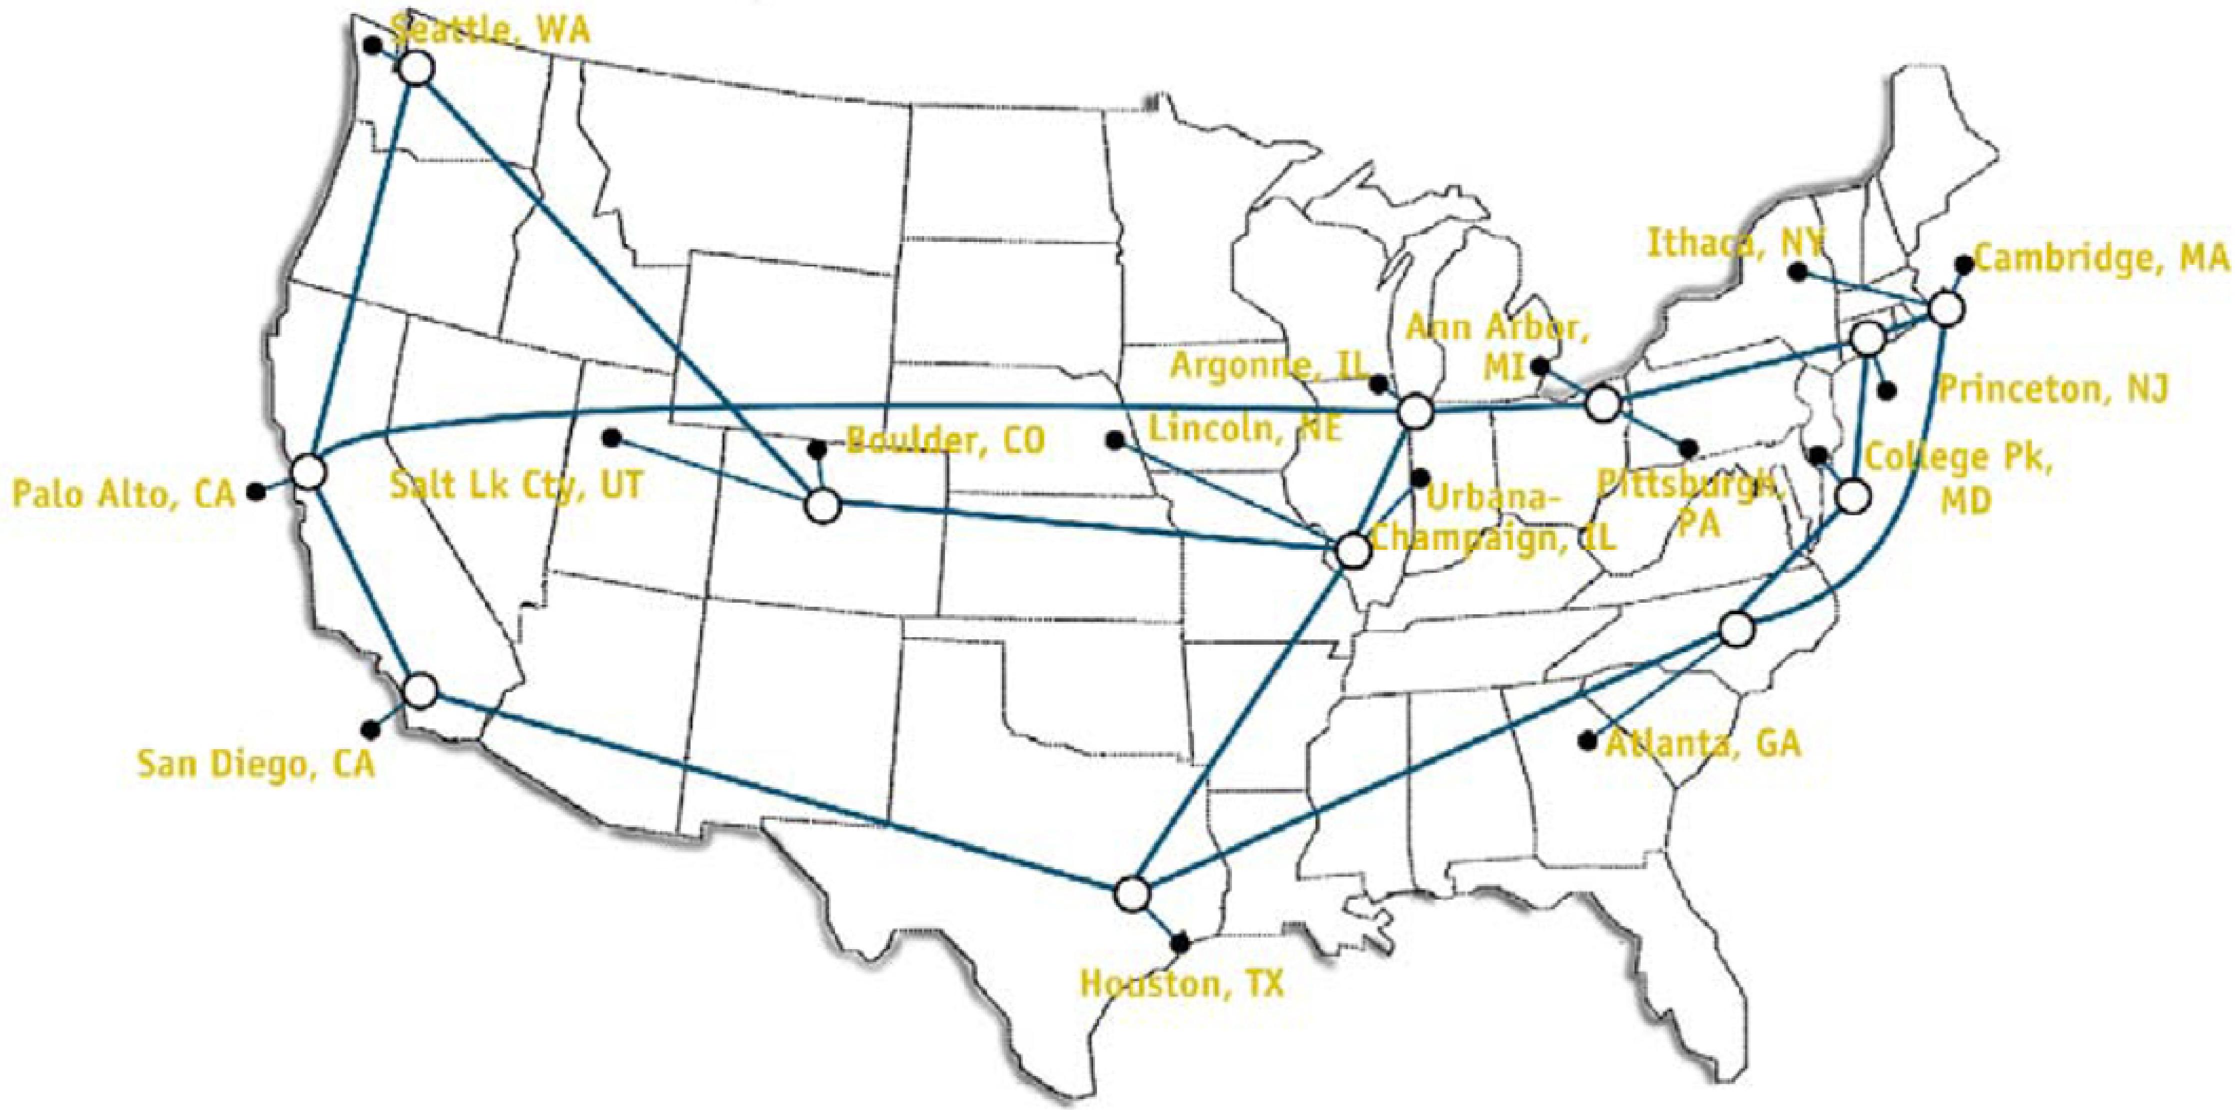

TIM BERNER-LEE

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

GLOBAL NETWORK - USENET

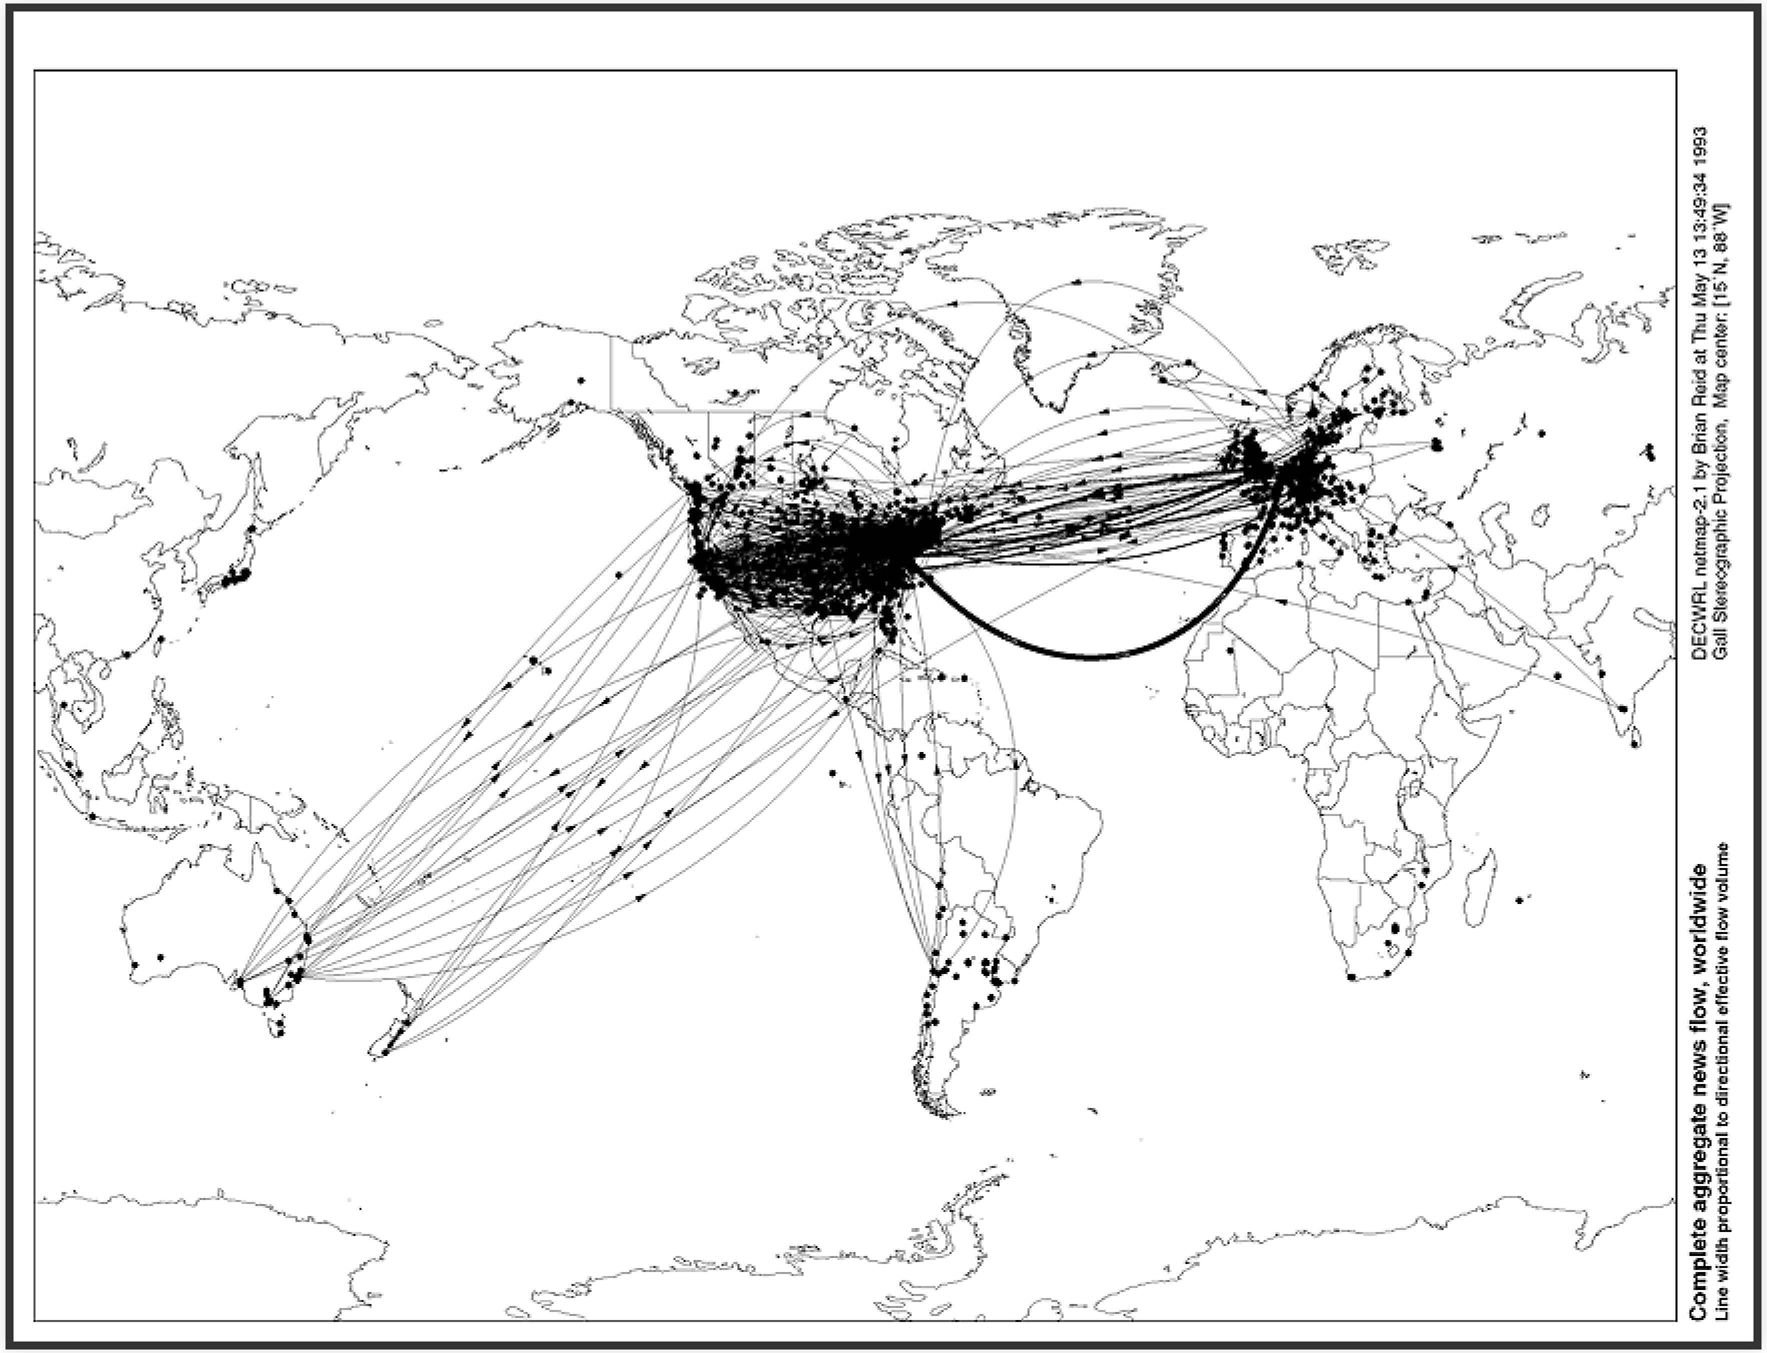

PRIVATIZATION

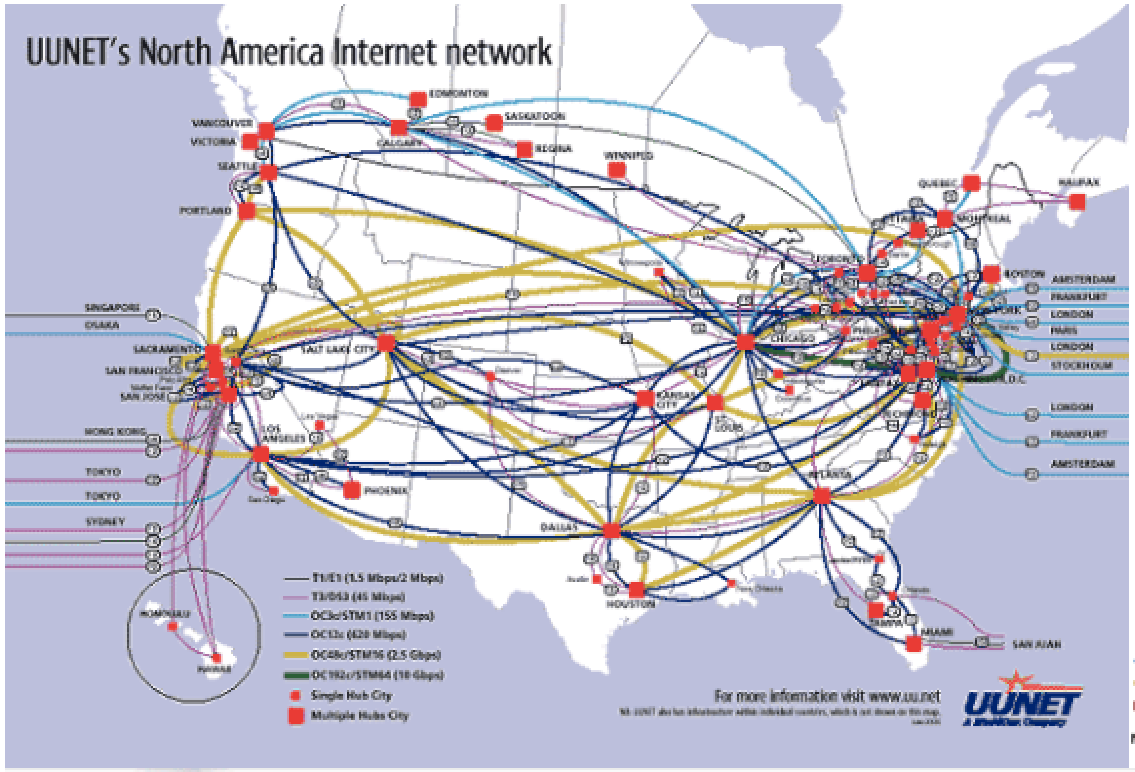
UUNET's North America Internet network
T1/E1 (1.5 Mbps/2 Mbps)
T3/DS3 (45 Mbps)
OC3c/STM1 (155 Mbps)
OC12c (620 Mbps)
OC48c/STM16 (2.5 Gbps)
OC192c/STM64 (10 Gbps)
Single Hub City
Multiple Hubs City
For more information visit www.uu.net
UUNET
SEATTLE
PORTLAND
SACRAMENTO
SAN FRANCISCO
SAN JOSE
SALT LAKE CITY
CHICAGO
KANSAS CITY
DALLAS
HOUSTON
ATLANTA
MIAMI
PHOENIX
SAN JUAN
VANCOUVER
VICTORIA
EDMONTON
SASKATOON
CALGARY
REGINA
WINNIPEG
QUEBEC
HALIFAX
MONTREAL
TORONTO
BOSTON
AMSTERDAM
FRANKFURT
LONDON
PARIS
STOCKHOLM
SINGAPORE
OSAKA
HONG KONG
TOKYO
SYDNEY

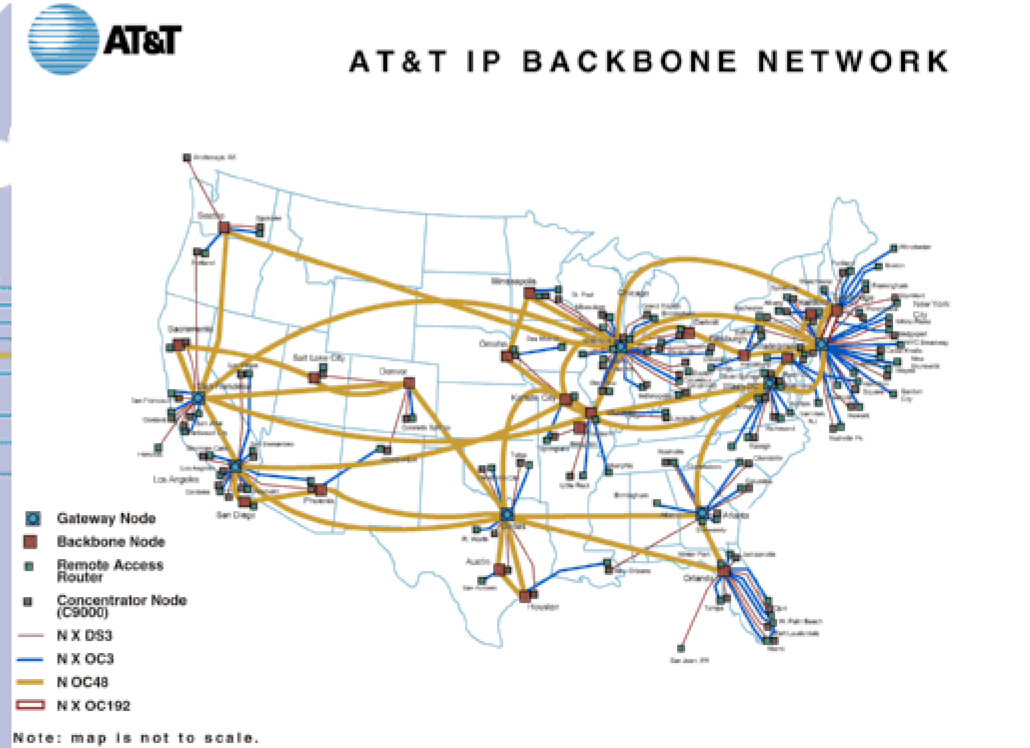
AT&T
AT&T IP BACKBONE NETWORK
Gateway Node
Backbone Node
Remote Access Router
Concentrator Node (C9000)
N X DS3
N X OC3
N OC48
N X OC192
Note: map is not to scale.

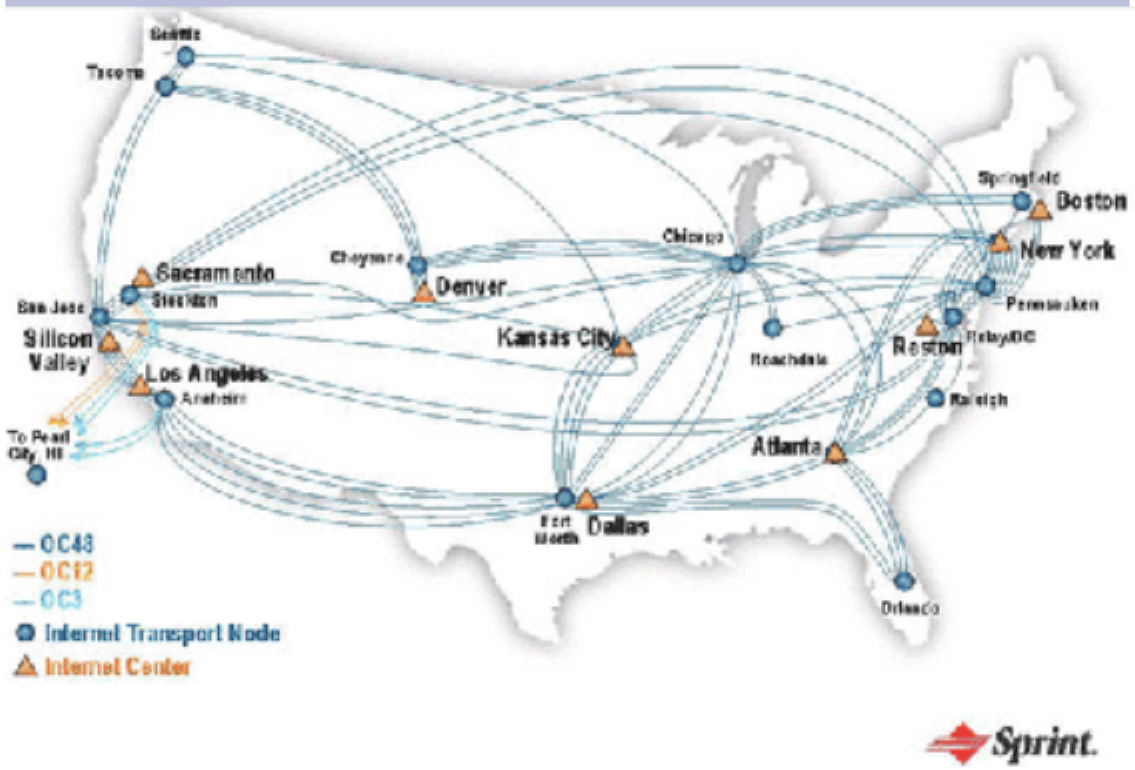
Seattle
Tacoma
Sacramento
Stockton
San Jose
Silicon Valley
Los Angeles
Anaheim
To Pearl City, HI
Cheyenne
Denver
Kansas City
Chicago
Roachdale
Fort Worth
Dallas
Atlanta
Orlando
Raleigh
Reston
Relay, MD
Pennsauken
New York
Springfield
Boston
OC48
OC12
OC3
Internet Transport Node
Internet Center
Sprint.


Level(3)
Seattle
Portland
Sacramento
San Francisco
Oakland
San Jose
San Luis Obispo
Los Angeles
Orange
San Diego
Las Vegas
Phoenix
Salt Lake City
Denver
El Paso
Omaha
Kansas City
Chicago
St. Louis
Detroit
Cleveland
Indianapolis
Louisville
Cincinnati
Nashville
Memphis
Dallas
Fort Worth
Austin
San Antonio
Houston
New Orleans
Atlanta
Jacksonville
Orlando
Tampa
Miami
Charlotte
Raleigh
Richmond
Pittsburgh
Buffalo
Albany
Manchester
Boston
Providence
Hartford
Stamford
Long Island
New York City
Jersey City
Newark
Princeton
Philadelphia
Baltimore
Wilmington
Washington, D.C.

PRESENT

- ~750 million hosts
- Smartphones and tablets
- Aggressive deployment of broadband access
- Increasing ubiquity of high-speed wireless access (3G, 4G)

PRESENT

- Emergence of online social networks:
- Facebook: ~ one billion users
- Service providers (Google, Microsoft) create their own networks
- Bypass Internet, providing “instantaneous” access to search, email, etc.
- E-commerce, universities, enterprises running their services in “cloud” (eg, Amazon EC2)

WHAT IS THE INTERNET?

Nuts-and-bolts description

WHAT IS THE INTERNET?

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

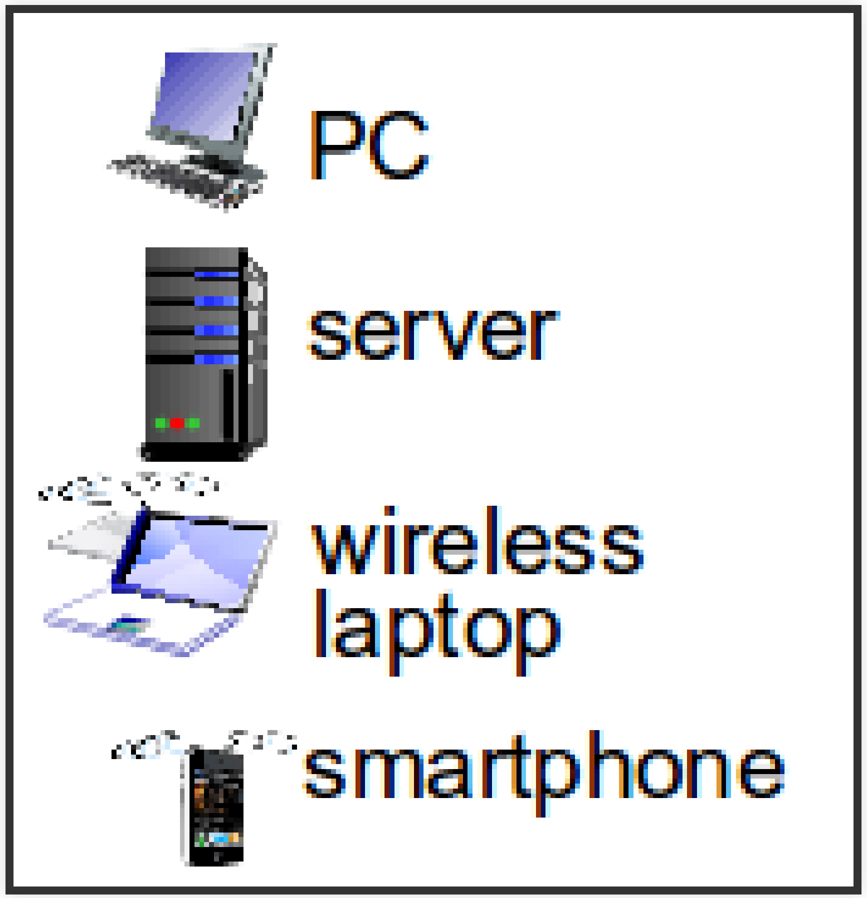

WHAT IS THE INTERNET?

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

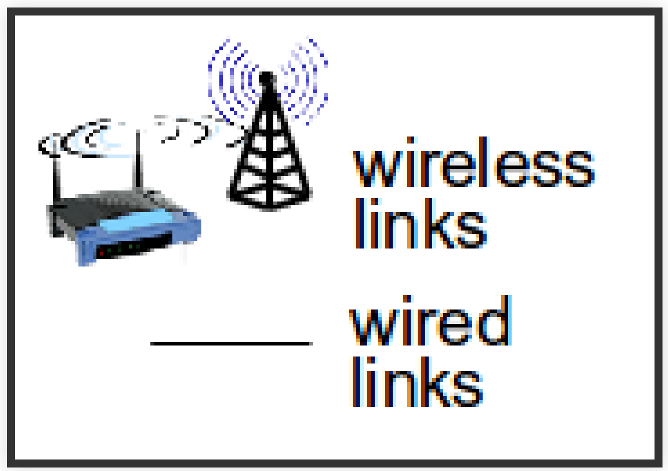

WHAT IS THE INTERNET?

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

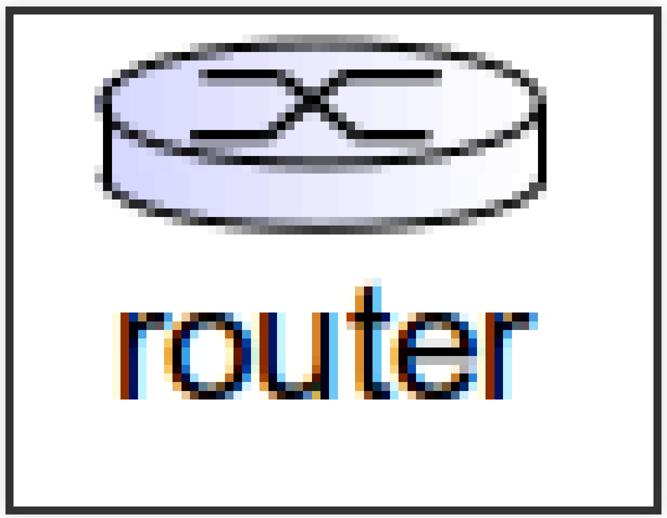

WHAT IS THE INTERNET?

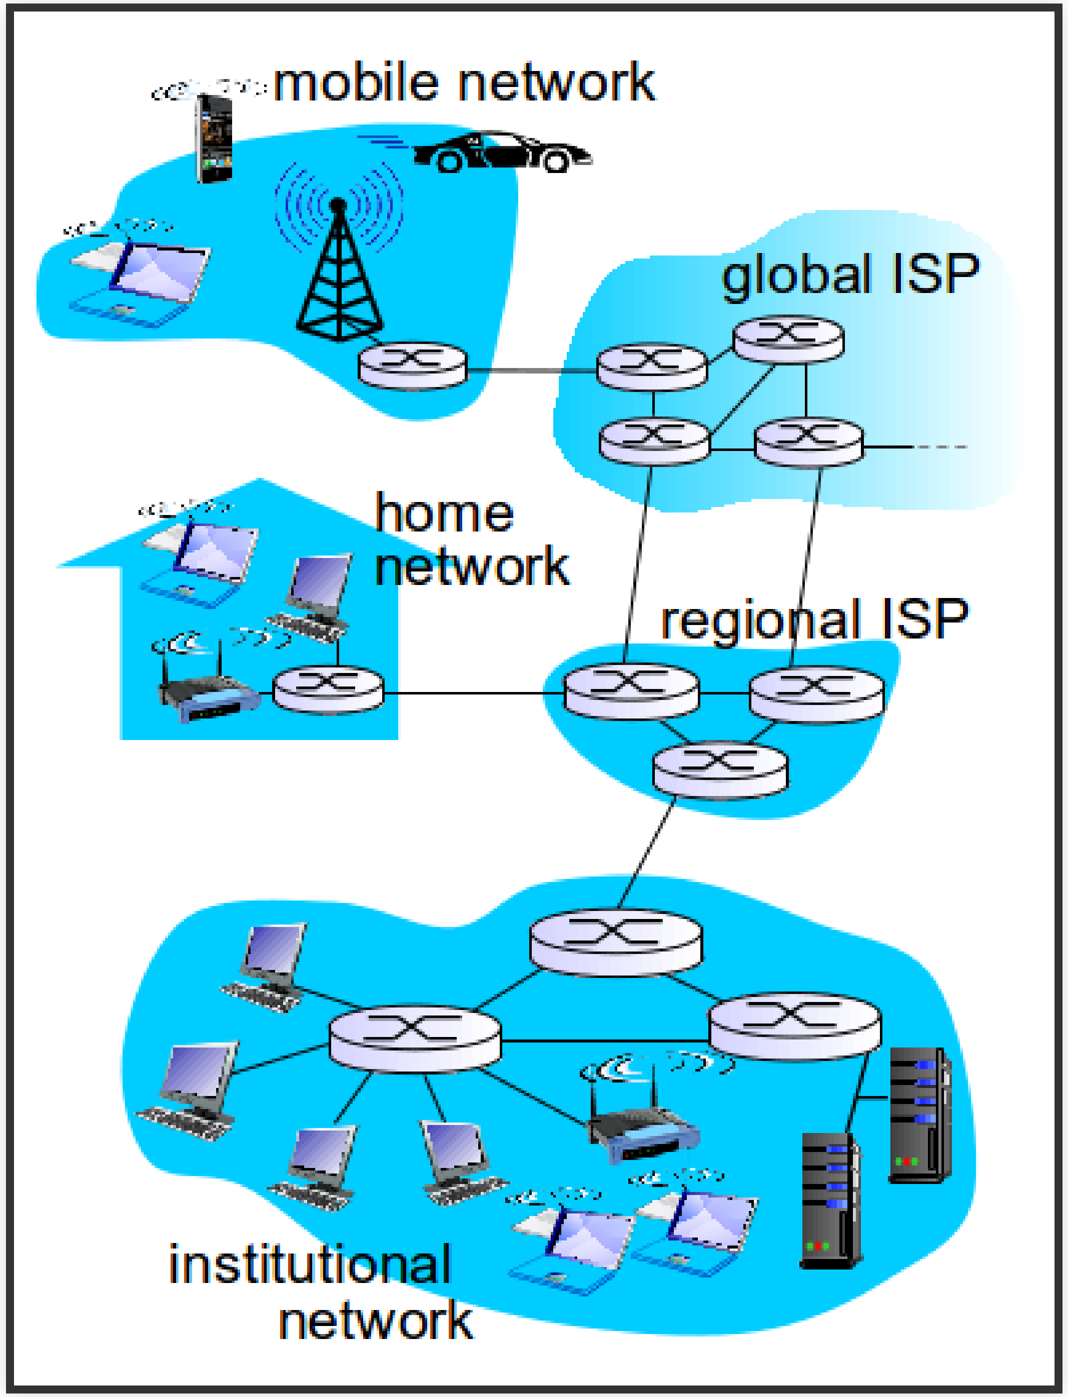

WHAT MAKES UP THE INTERNET?

- **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

WHAT MAKES UP THE INTERNET?

- **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

WHAT'S A PROTOCOL?

A **protocol** define format, order of messages sent and received among network entities, and actions taken on message transmission and/or receipt

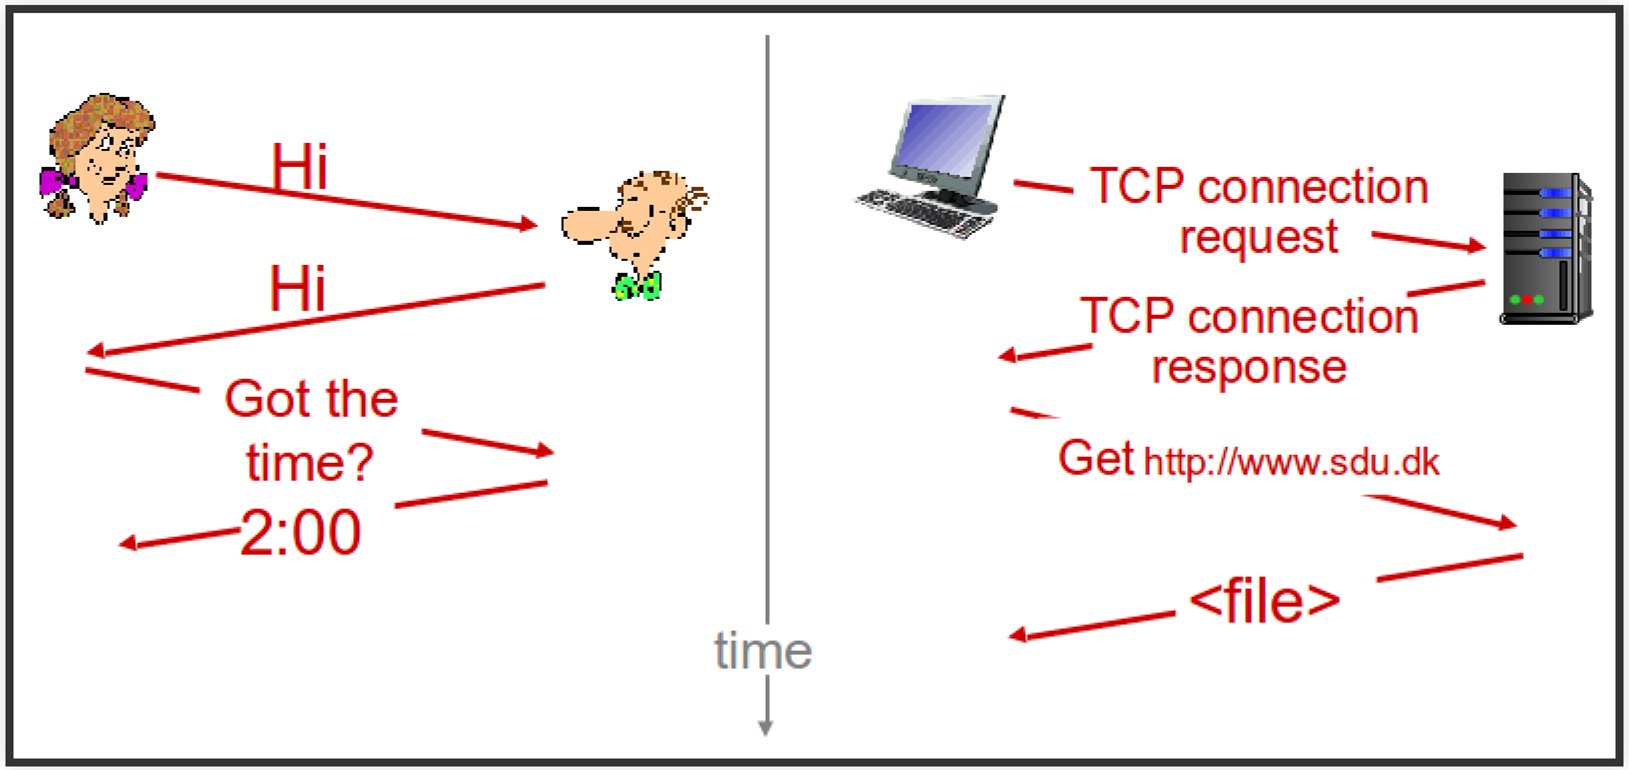

THE NETWORK EDGE

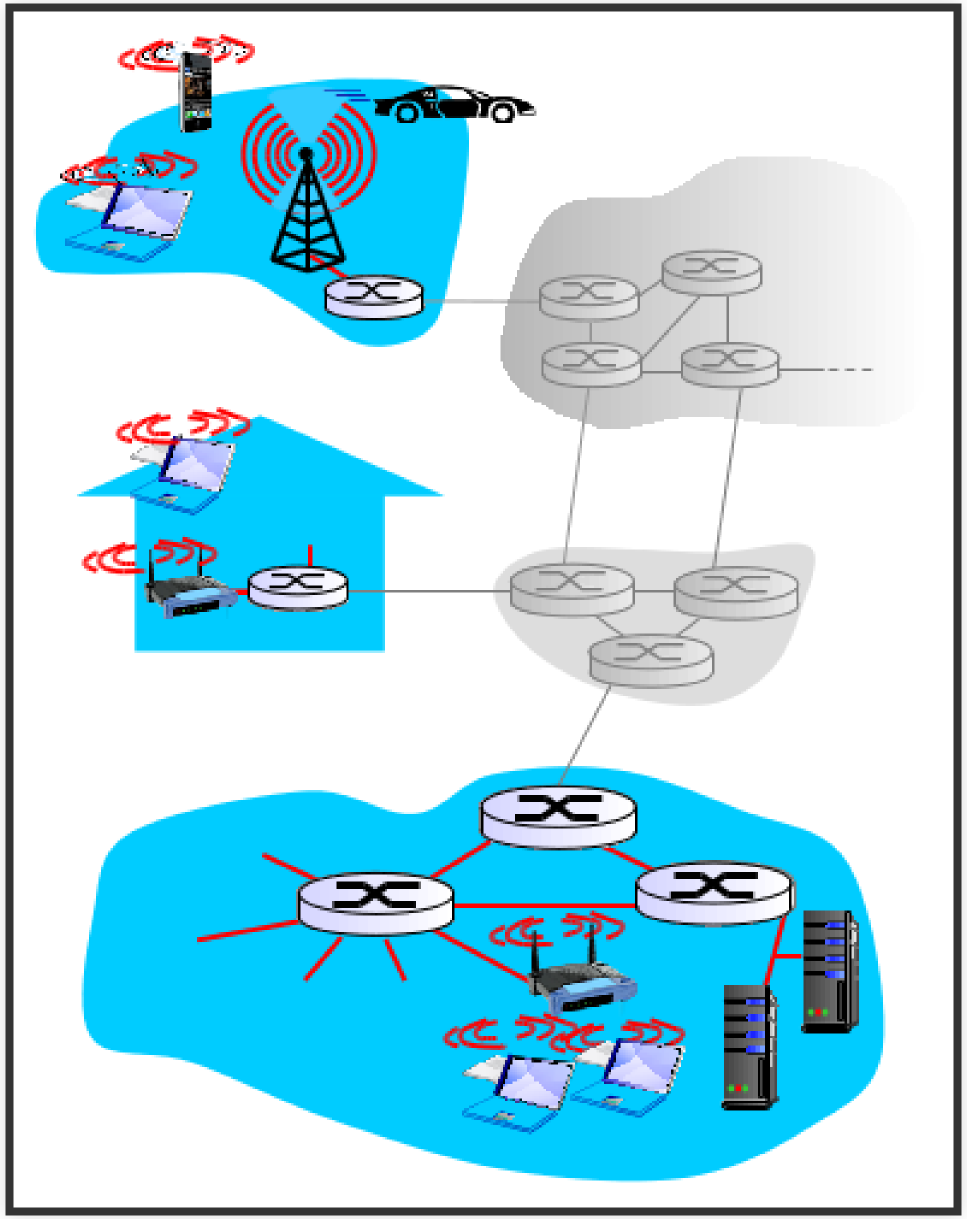

THE NETWORK EDGE

The edge of the network contains the components we are most familiar with

- Computers
- Smartphones
- Other devices used on a day to day basis

Typically called hosts, but also divided into

- Clients
- Servers (often in data centers)

THE NETWORK EDGE

How do we connect end systems to edge routers

- 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 NET: DSL

Digital Subscriber Line

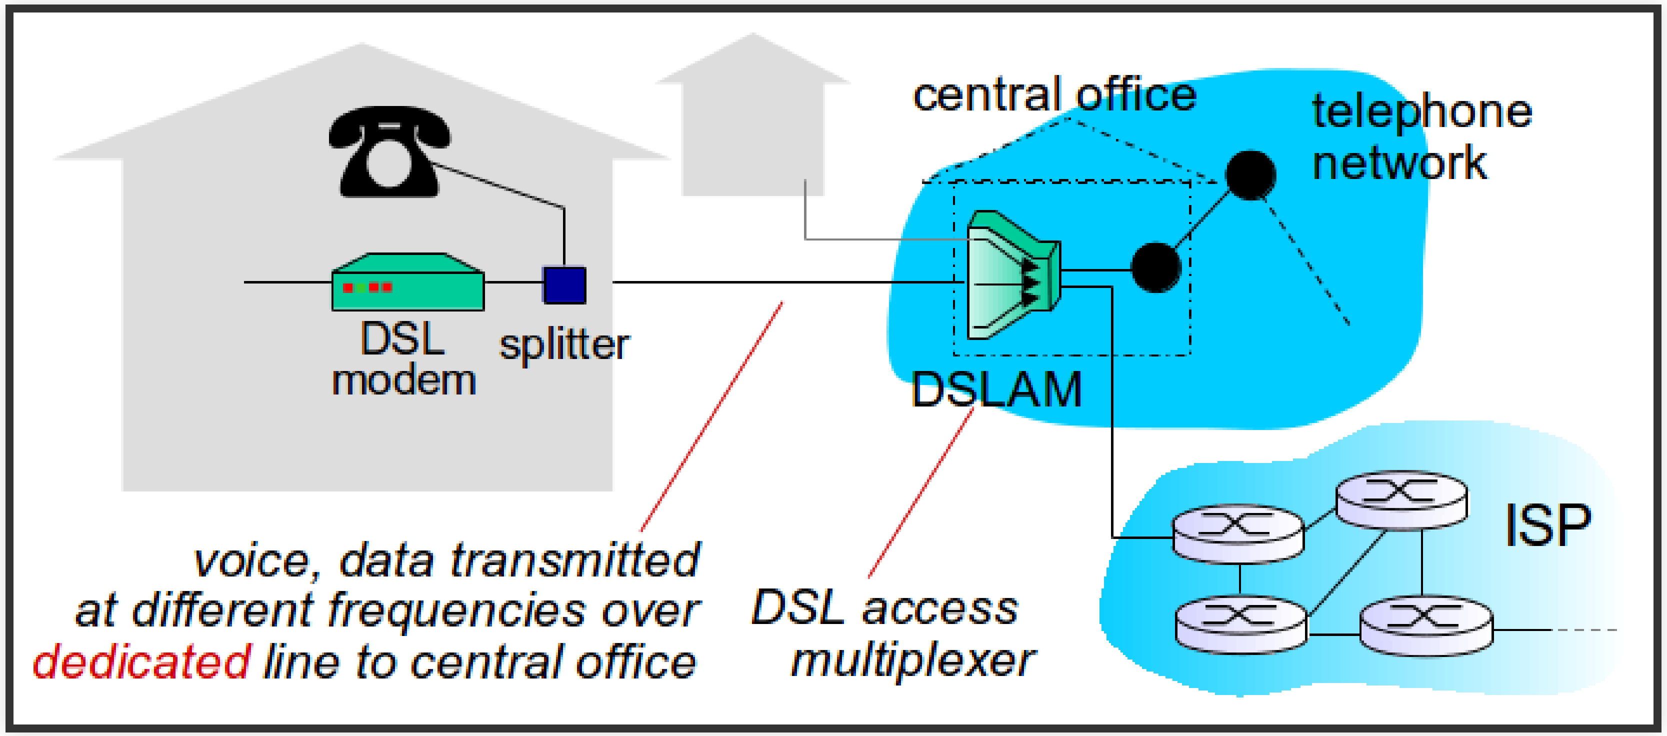

ACCESS NET: CABLE NETWORK

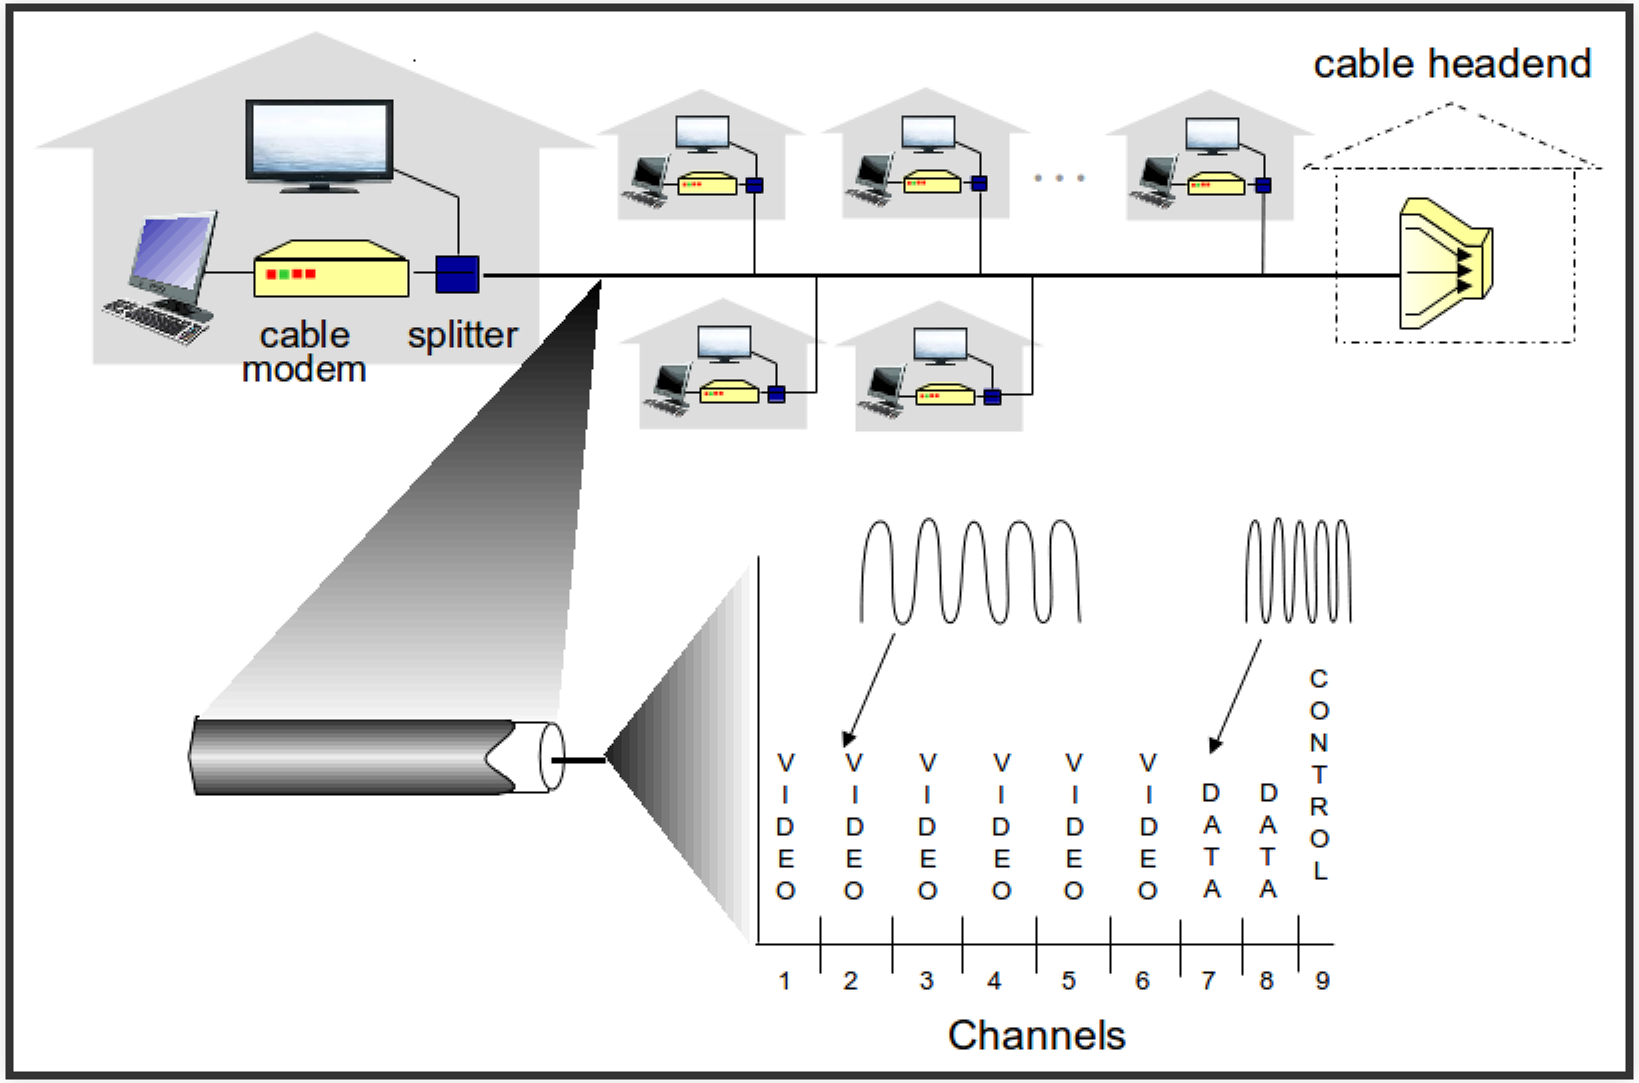

ACCESS NET: HOME NETWORK

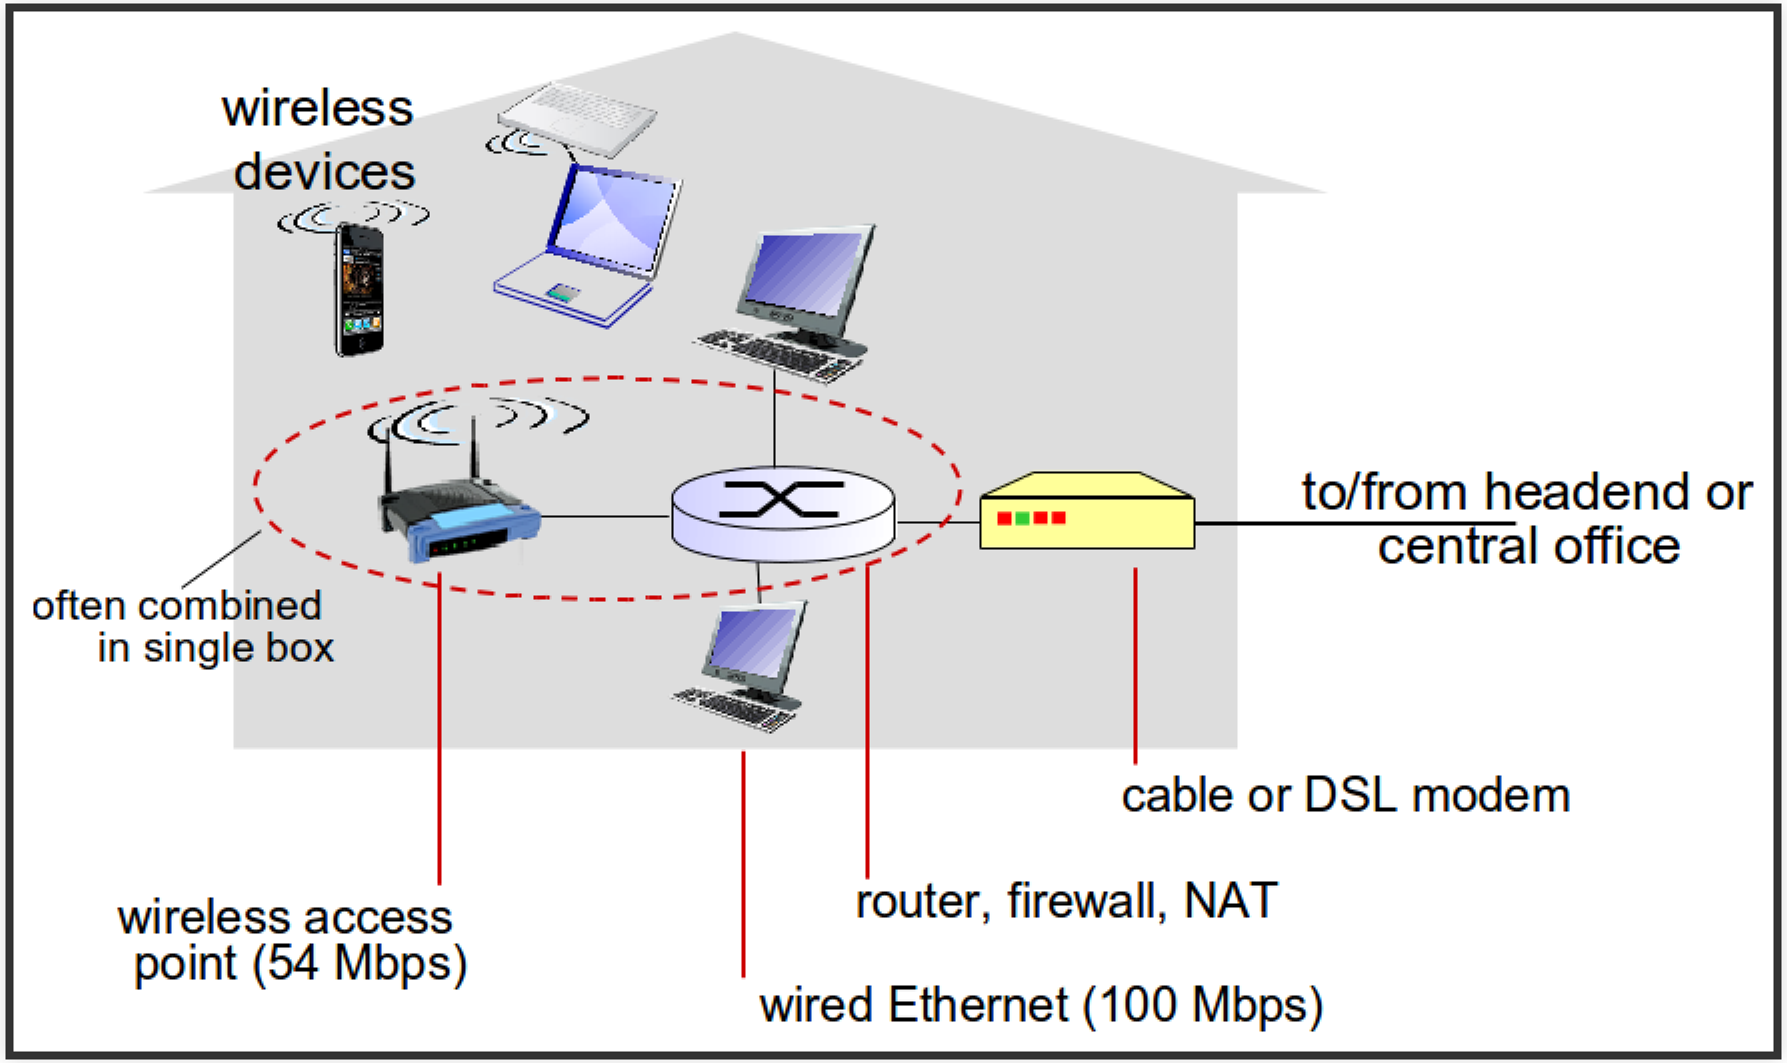

ENTERPRICE ACCESS NETWORKS (ETHERNET)

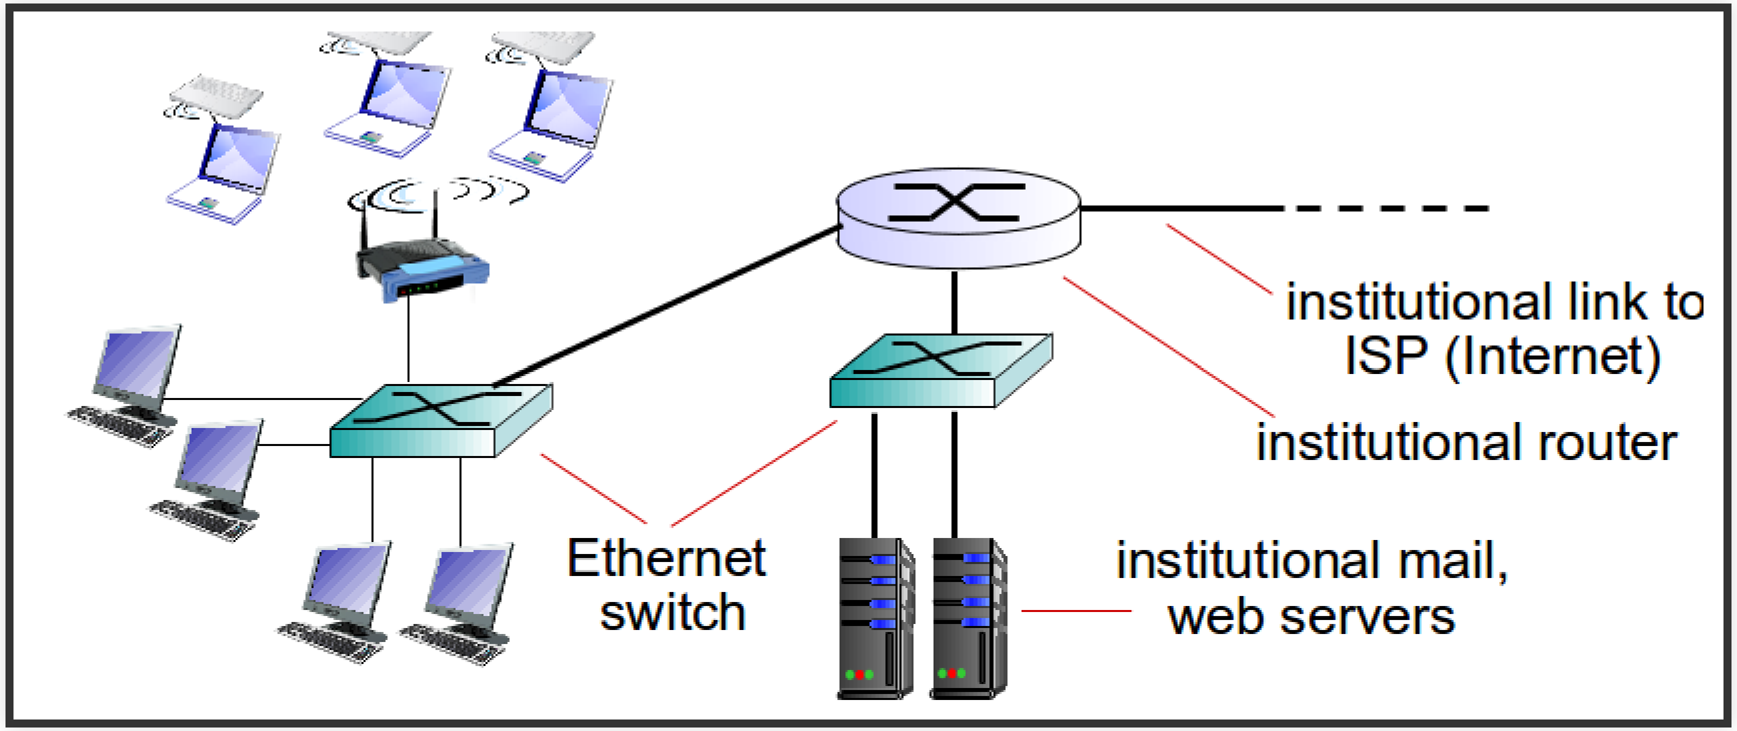

ENTERPRICE ACCESS NETWORKS (ETHERNET)

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

WIRELESS ACCESS NETWORKS

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

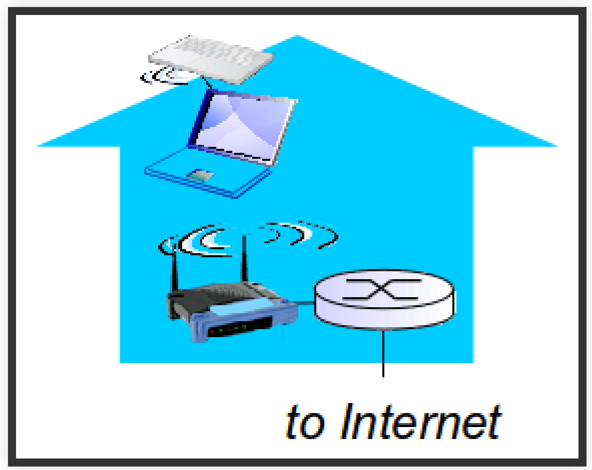

HOST: SENDS PACKETS OF DATA

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

```
packet transmission delay =
time needed to transmit L-bit packet into link =
L (bits) /  R (bits/sec)
```

PHYSICAL MEDIA

- **bit**: propagates between transmitter/receiver pairs
- **physical link**: what lies between transmitter and 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 Gpbs Ethernet
  - Category 6: 10Gbps

TWISTED PAIRS

PHYSICAL MEDIA: COAX

Coaxial cable

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

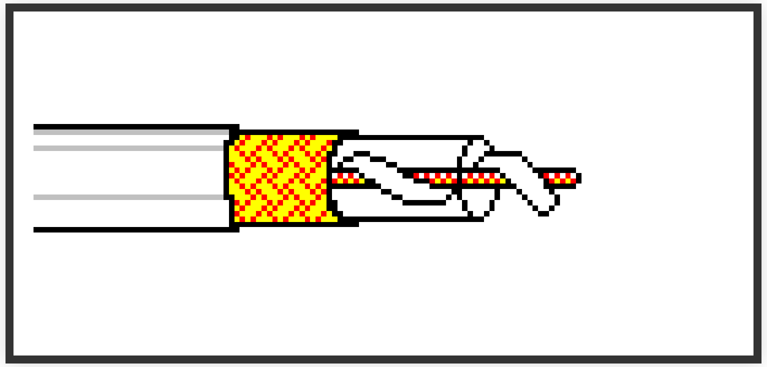

PHYSICAL MEDIA: FIBER

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 Gpbs transmission rate)
- low error rate
  - Repeaters spaced far apart
  - Immune to electromagnetic noise

PHYSICAL MEDIA: FIBER

PHYSICAL MEDIA: RADIO

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

PHYSICAL MEDIA: RADIO

Radio link types:

- **terrestrial microwave** e.g. up to 45 Mbps channels
- **LAN** (e.g., WiFi) 11Mbps, 54 Mbps
- **wide-area** (e.g., cellular) 3G cellular: ~ few Mbps
- **satellite**
  - Kbps to 45Mbps channel (or multiple smaller channels)
  - 270 msec end-end delay
  - geosynchronous versus low altitude

NETWORK CORE

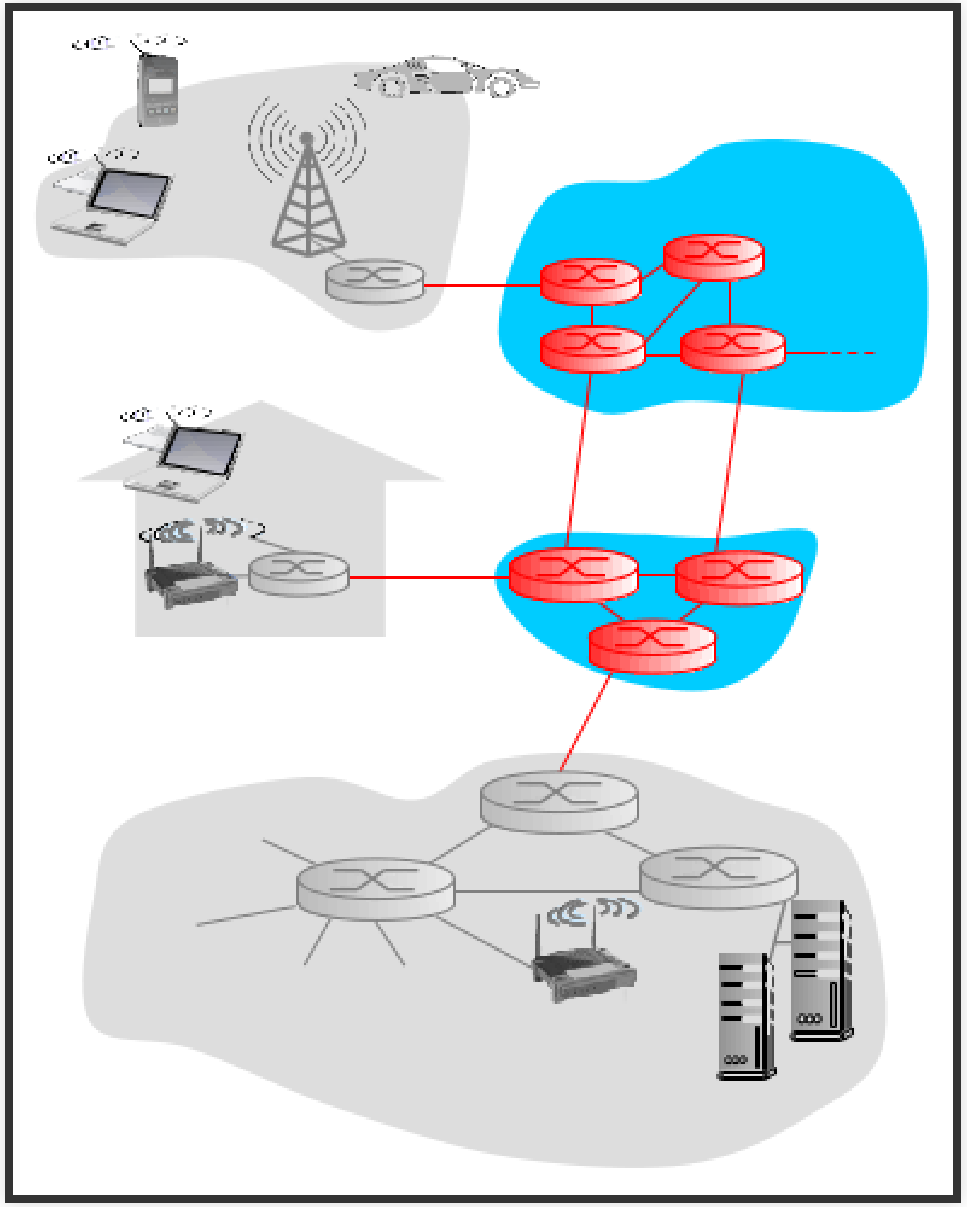

NETWORK CORE

- 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

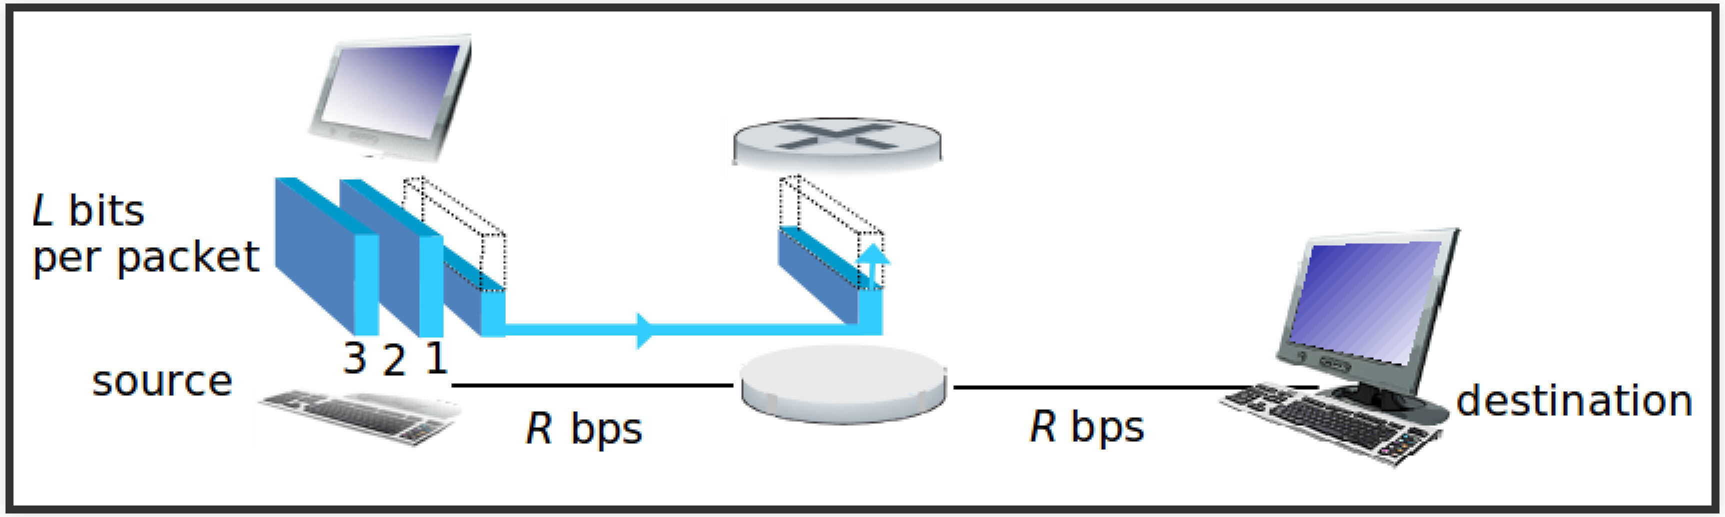

Store-and-forward

```
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 = 2L/R (assuming zero propagation delay)
```

PACKET-SWITCHING

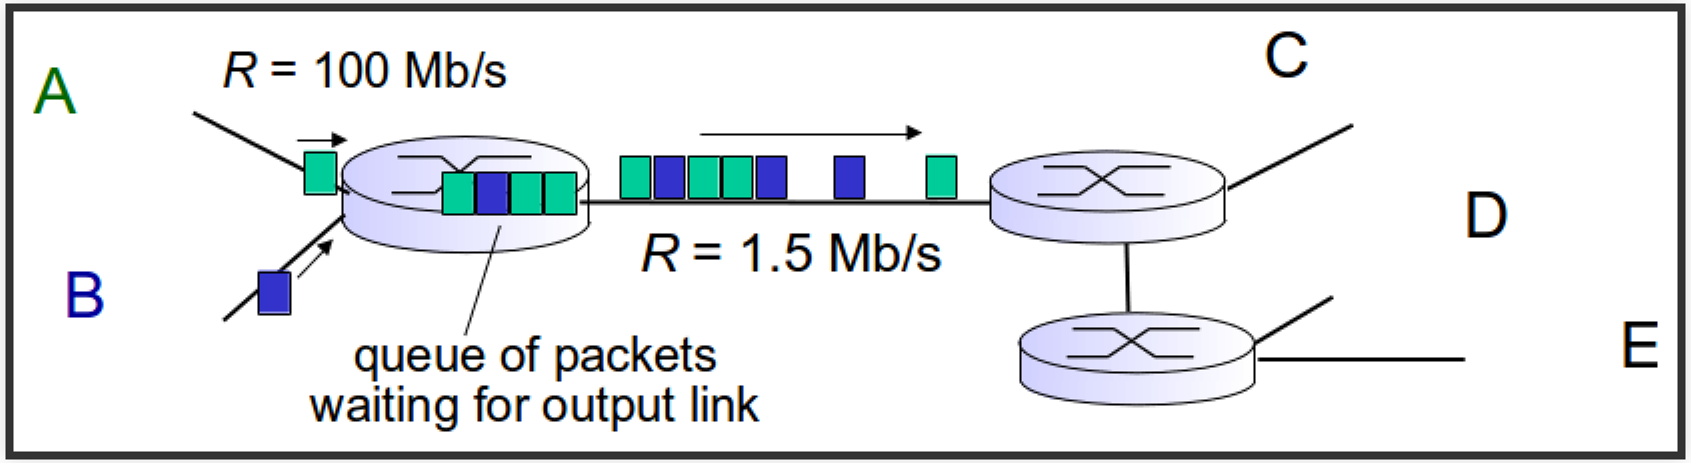

Queueing delay and loss

- 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

CIRCUIT SWITCHING

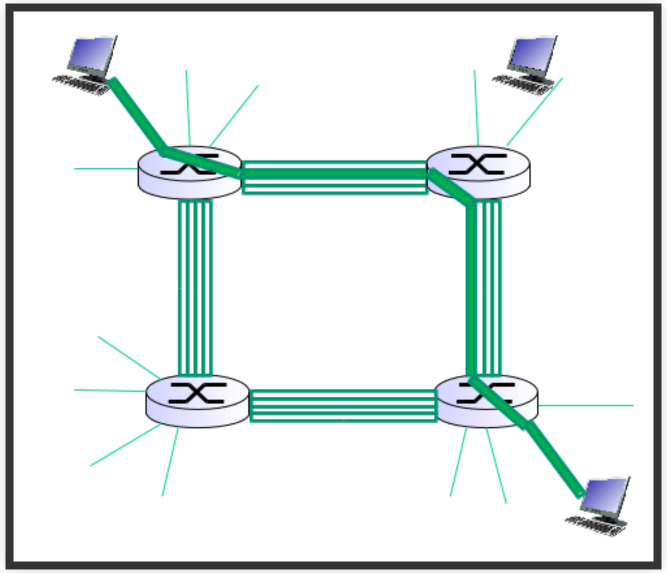

- In diagram, each link has four circuits.
- Dedicated resources: no sharing
  - circuit-like (guaranteed) performance
- circuit segment idle if not used by call (no sharing)
- Commonly used in traditional telephone networks

PACKET- VS. CIRCUIT-SWITCHING

Packet switching allows more users to use network!

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

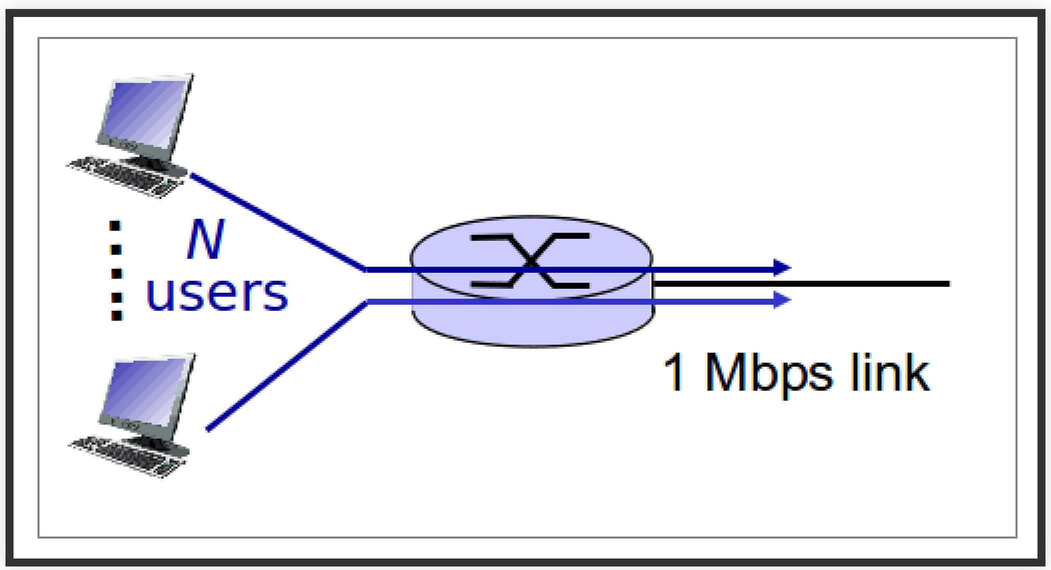

PACKET- VS. CIRCUIT-SWITCHING

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

Q: what happens if > 35 users ?

PACKET- VS. 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

PACKET- VS. CIRCUIT-SWITCHING

Q: How to provide circuit-like behavior?

- Bandwidth guarantees needed for audio/video apps
- still an unsolved problem (chapter 9, not covered)

ROUTING AND FORWARDING

Routing

- Determines source-destination route taken by packets.
- Uses routing algorithms

Forwarding

- Move packets from router's input to appropriate router output

ROUTING AND FORWARDING

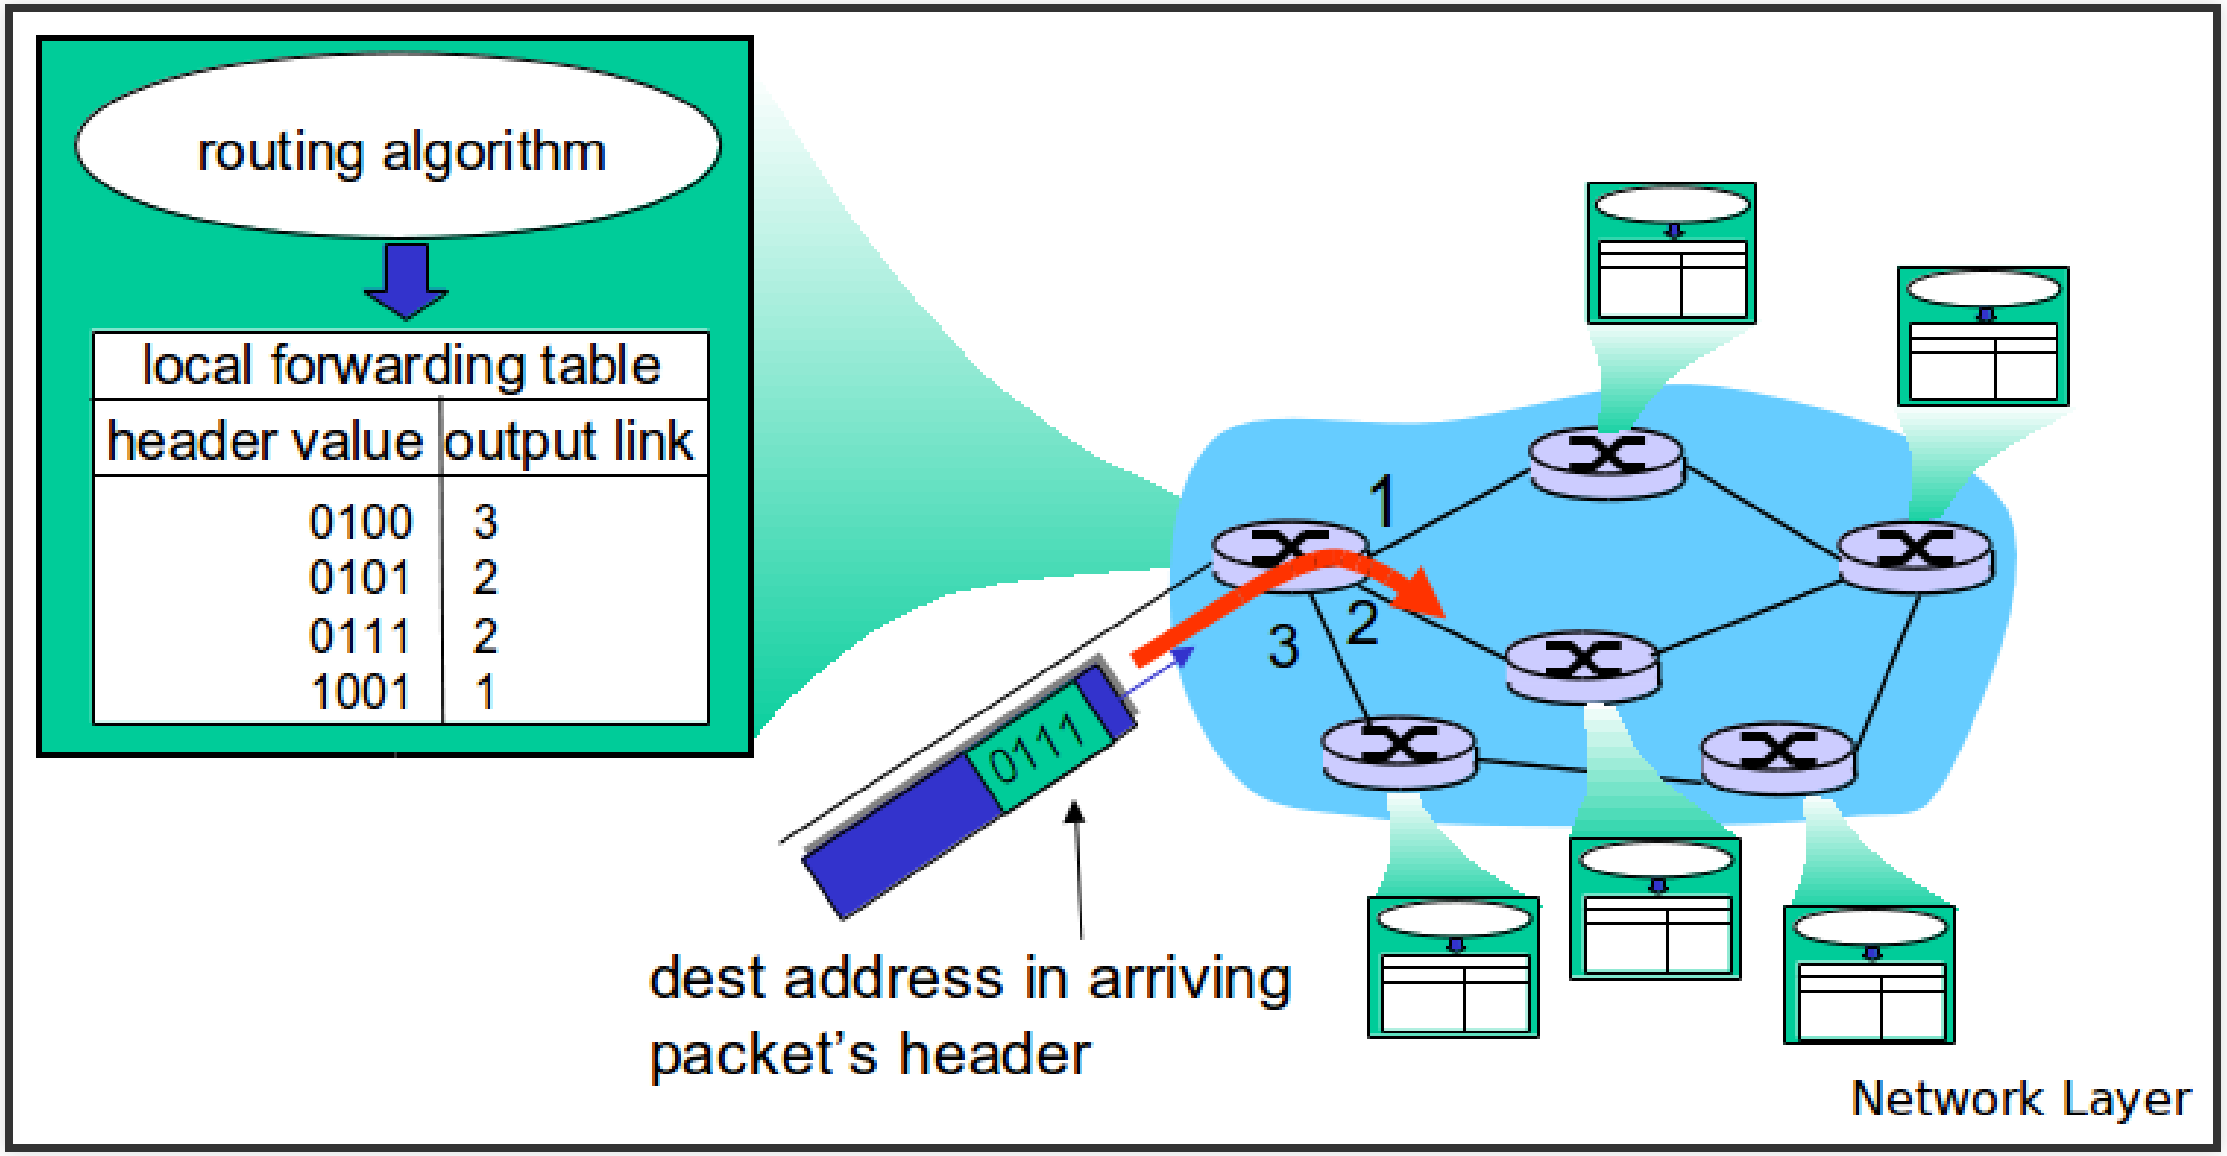

INTERNET STRUCTURE: NETWORK OF NETWORKS

- 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

Question Given millions of access ISPs, how to connect them together?

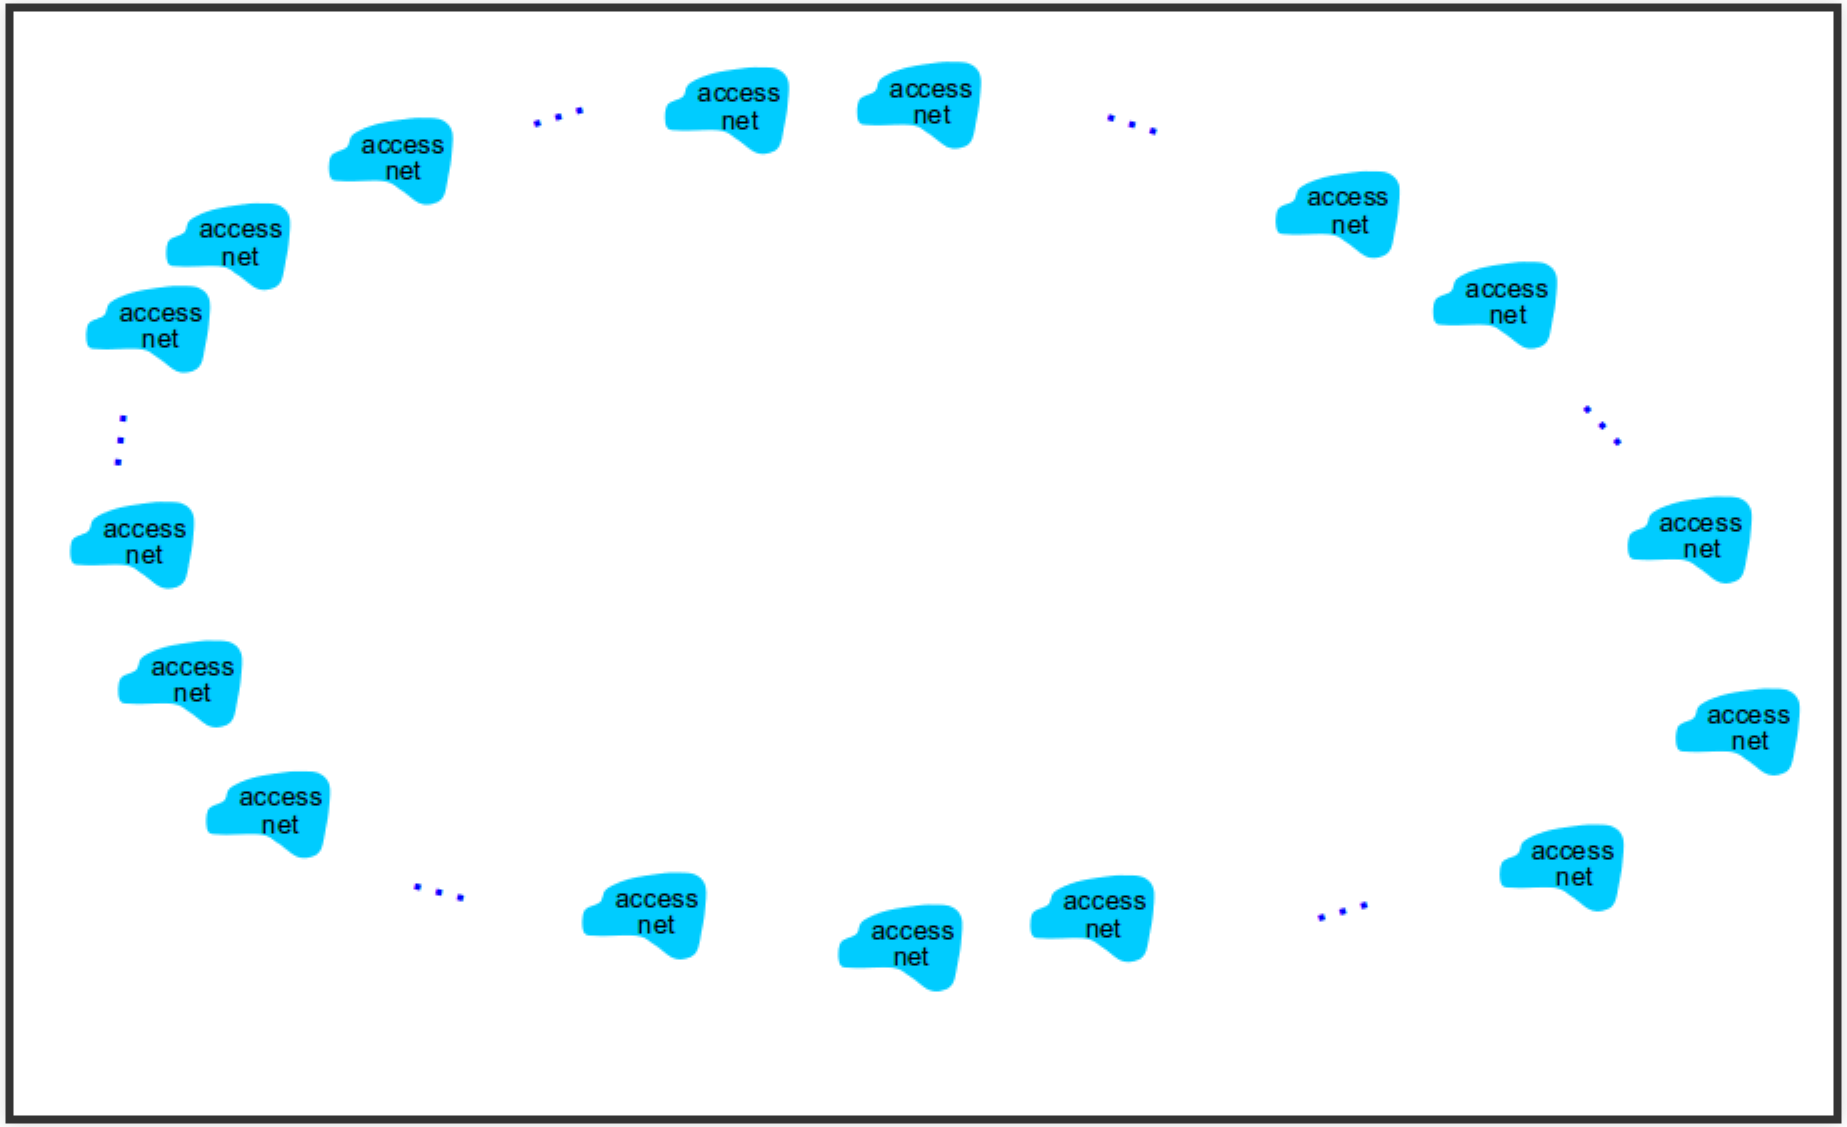

INTERNET STRUCTURE

Option Connect each access ISP to every other access ISP?

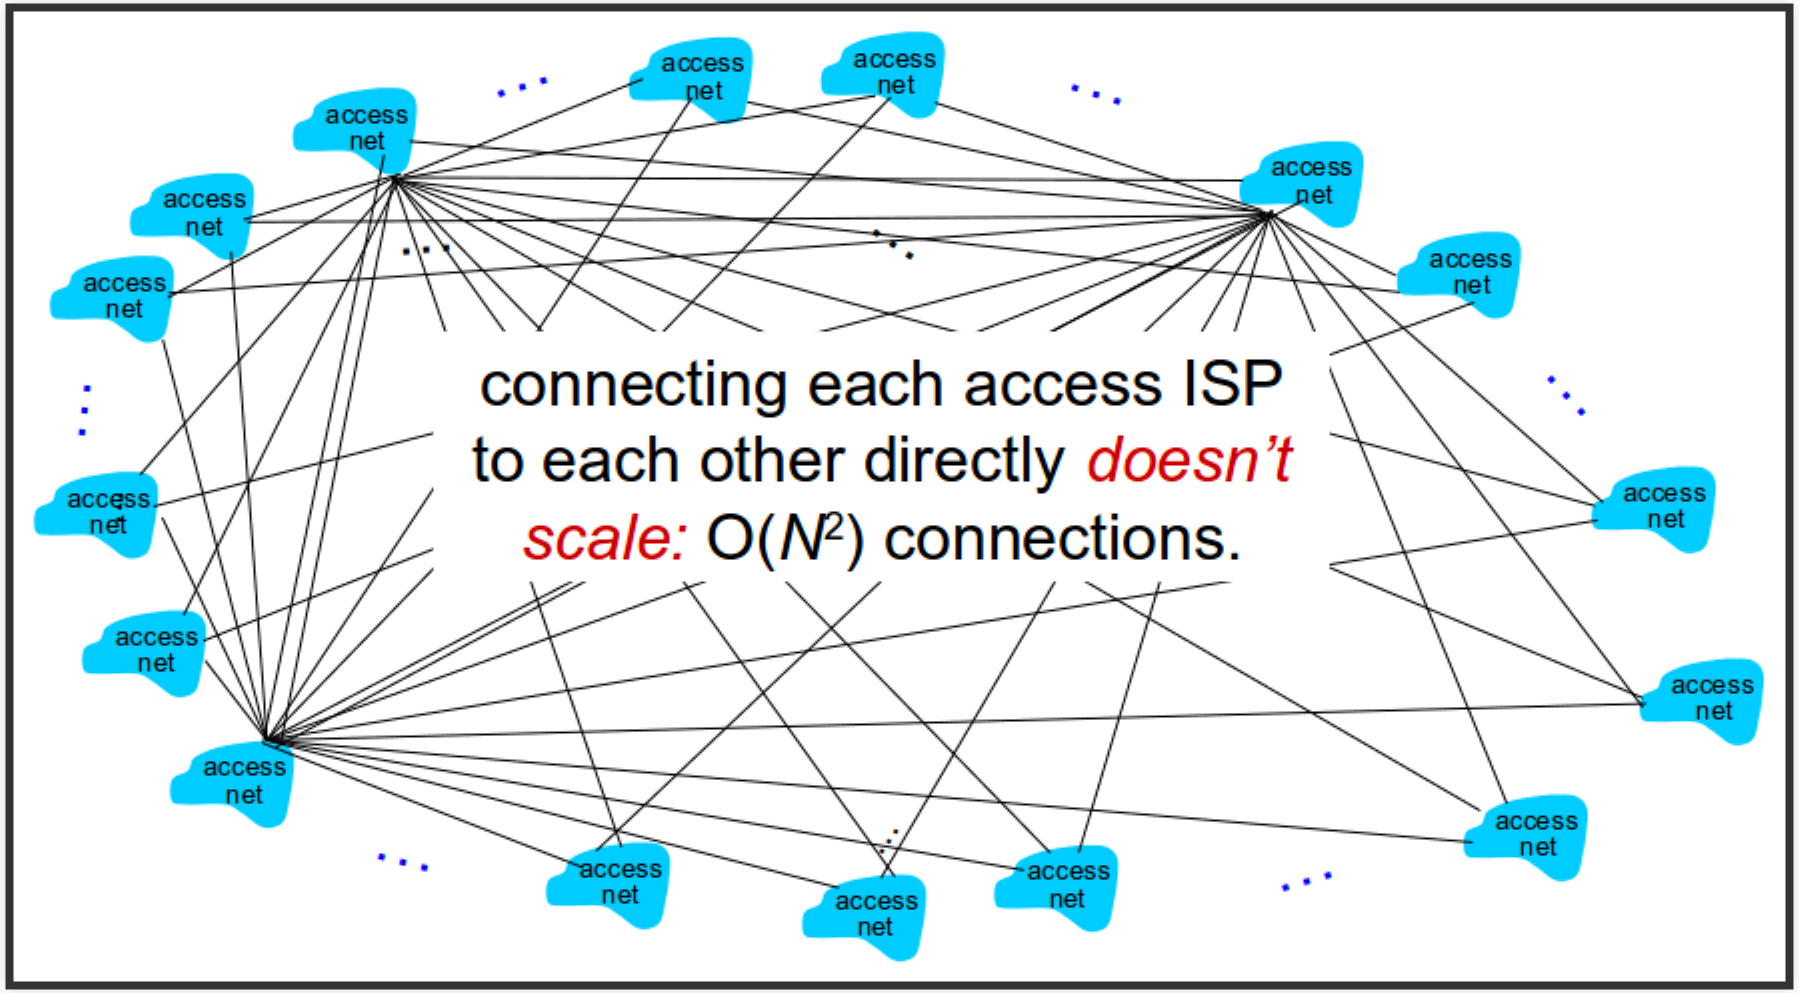

INTERNET STRUCTURE

Option Connect each access ISP to a global transit ISP?
Customer and provider ISPs have economic agreement.

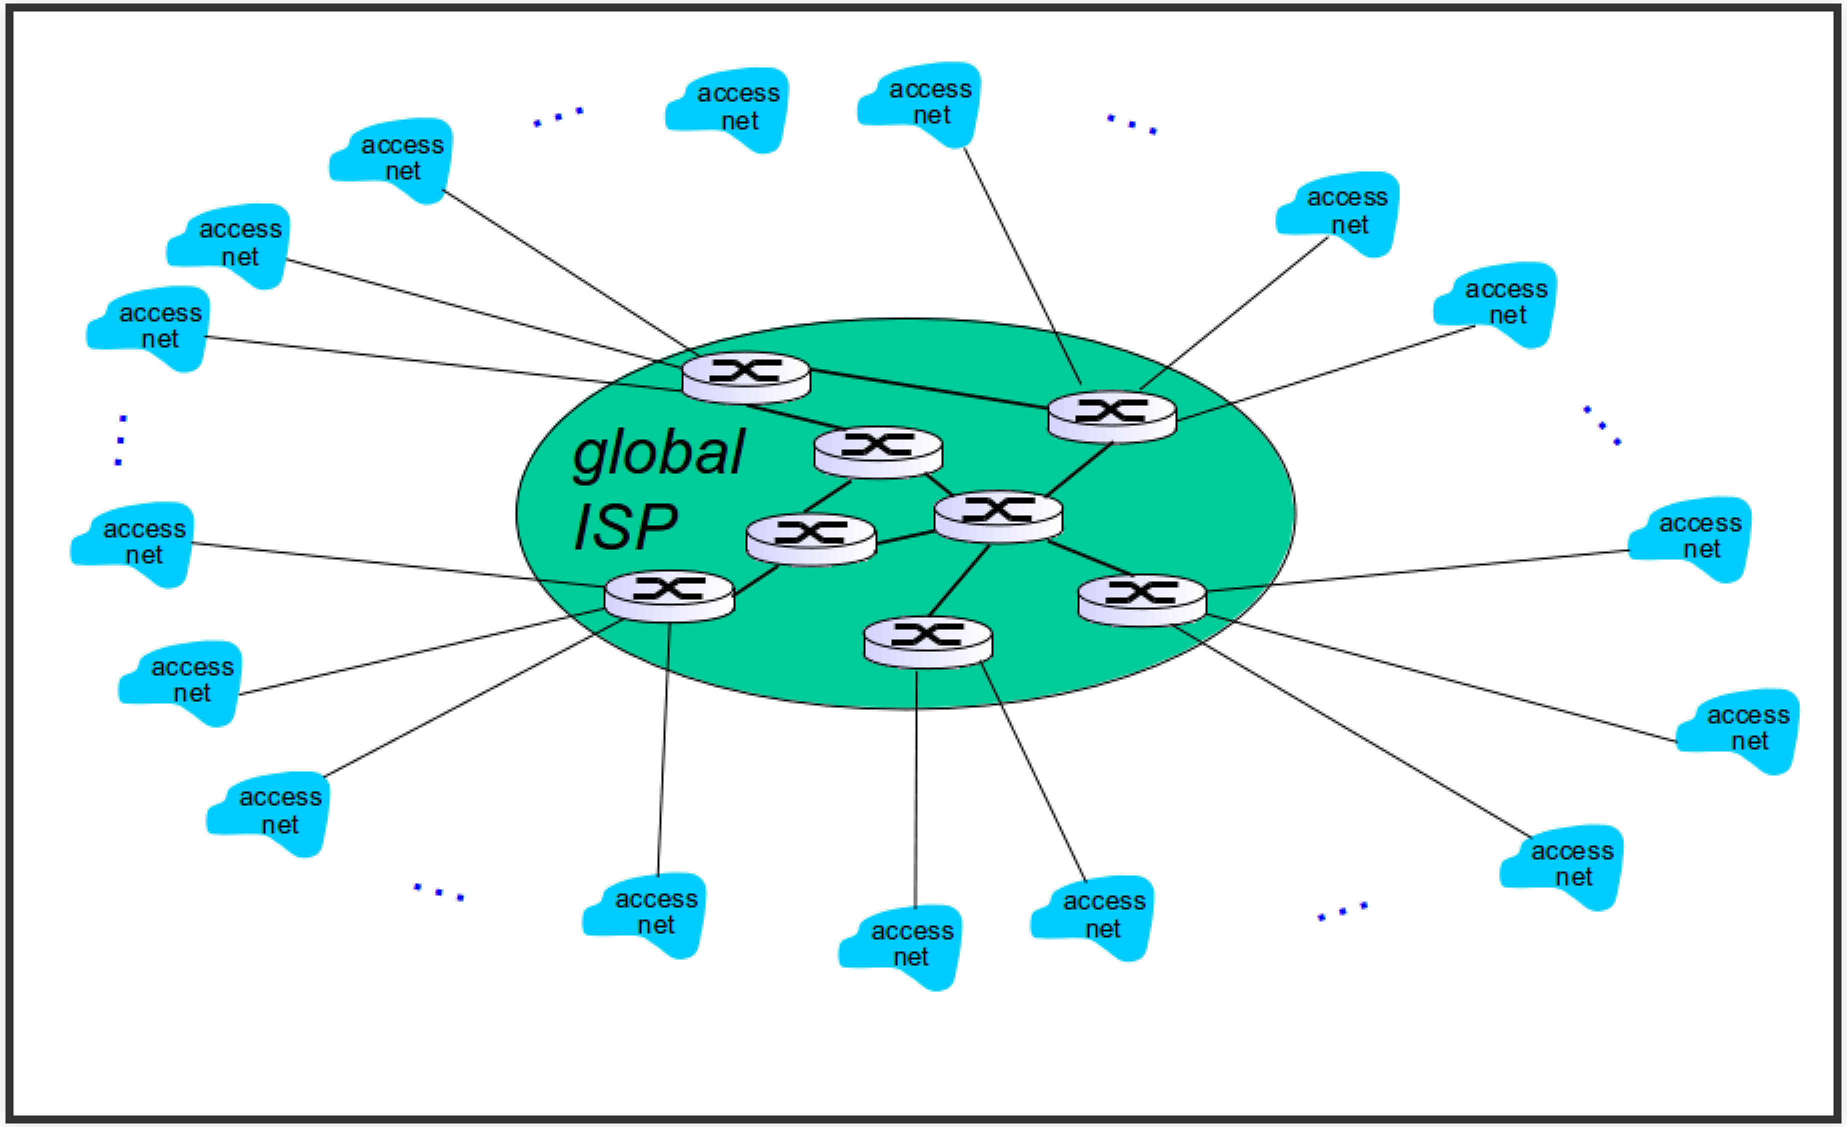

INTERNET STRUCTURE

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

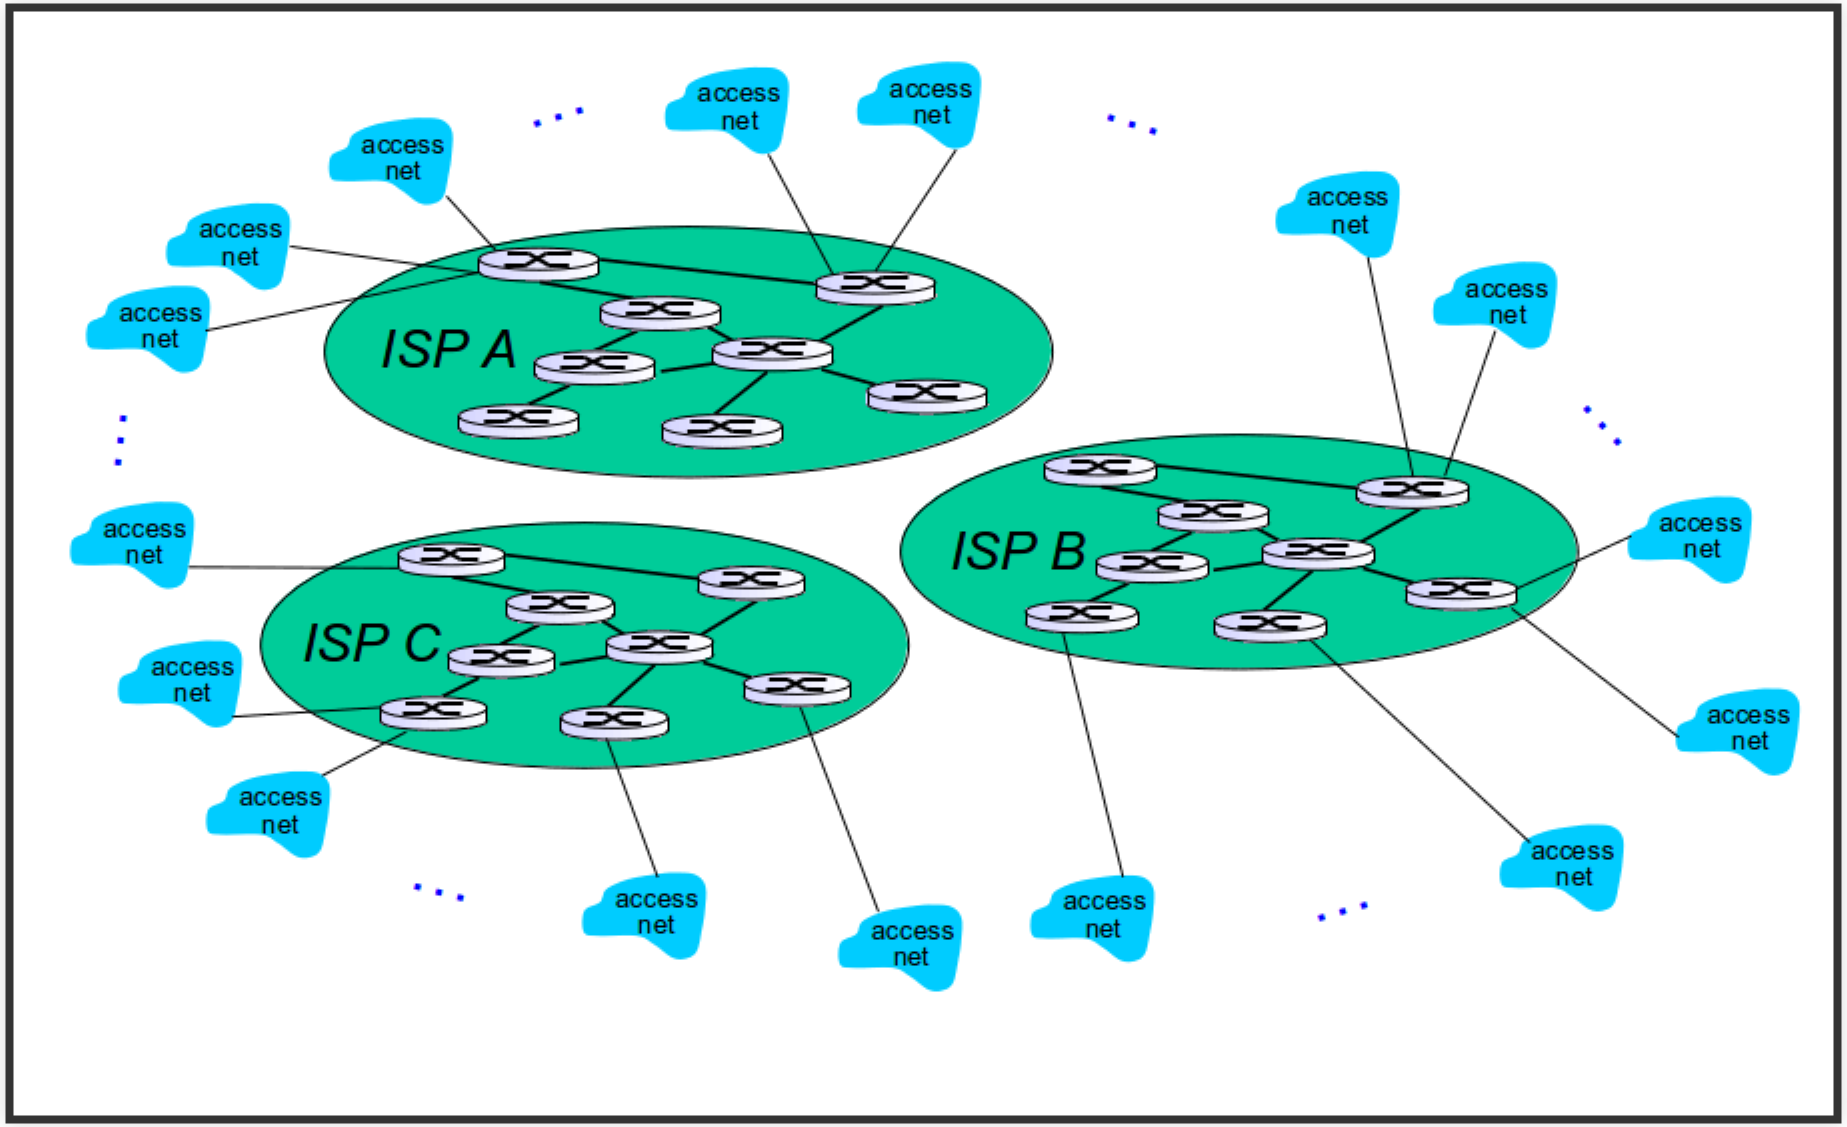

INTERNET STRUCTURE

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

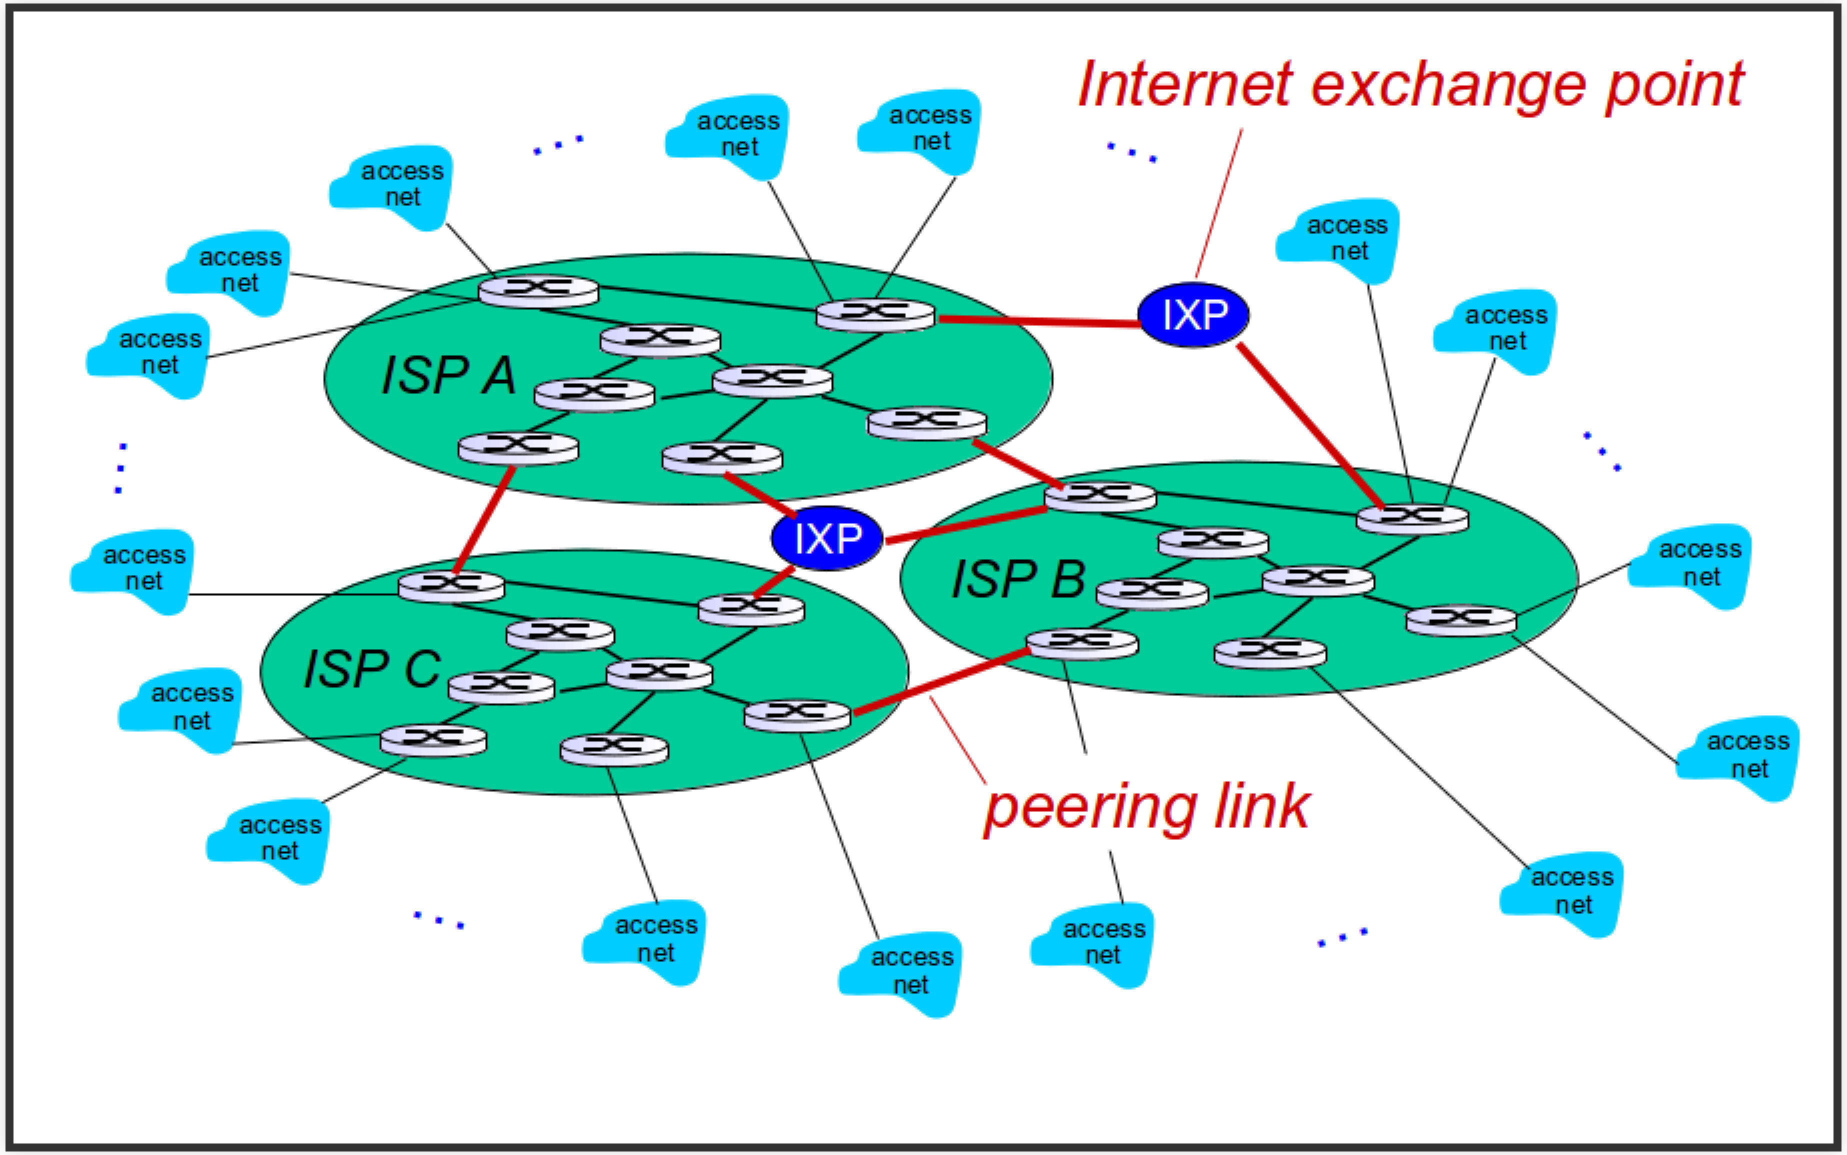

INTERNET STRUCTURE

❗ and regional networks may arise to connect access nets to ISPS

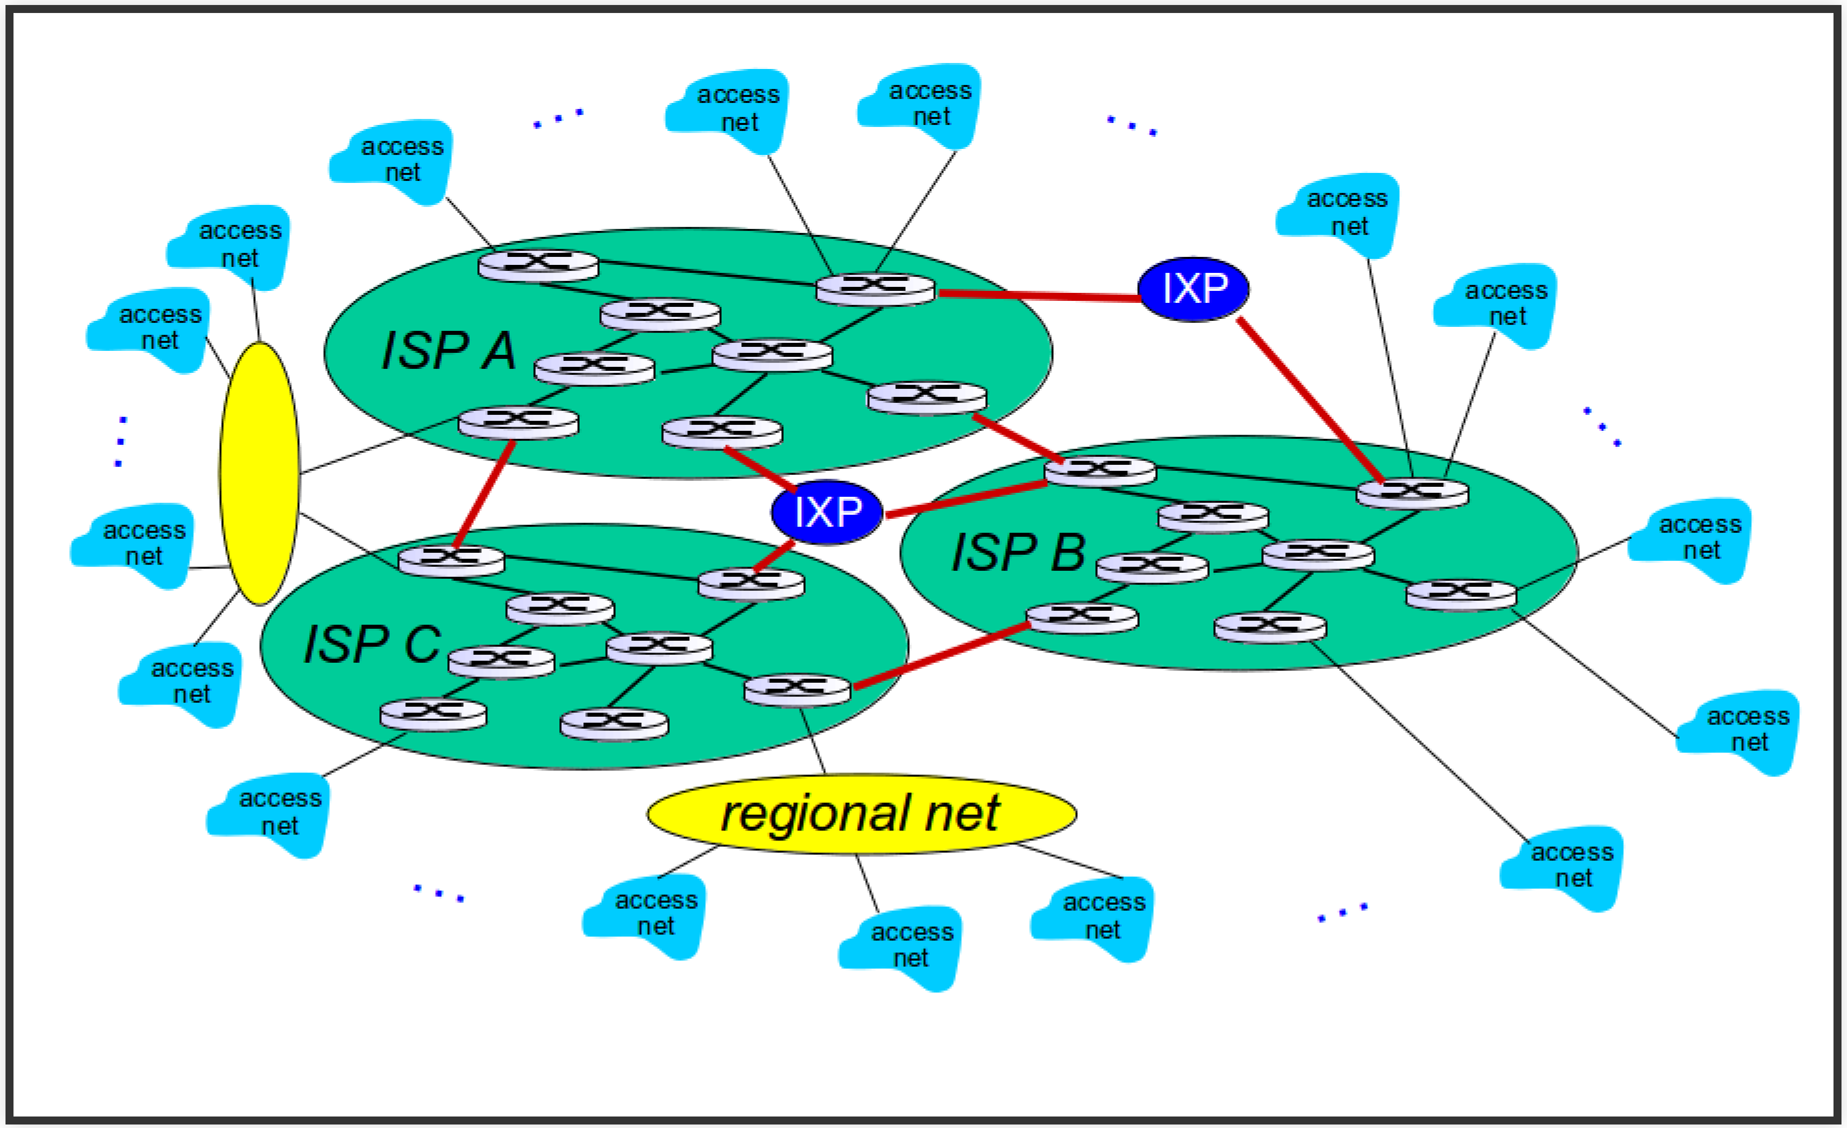

INTERNET STRUCTURE

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

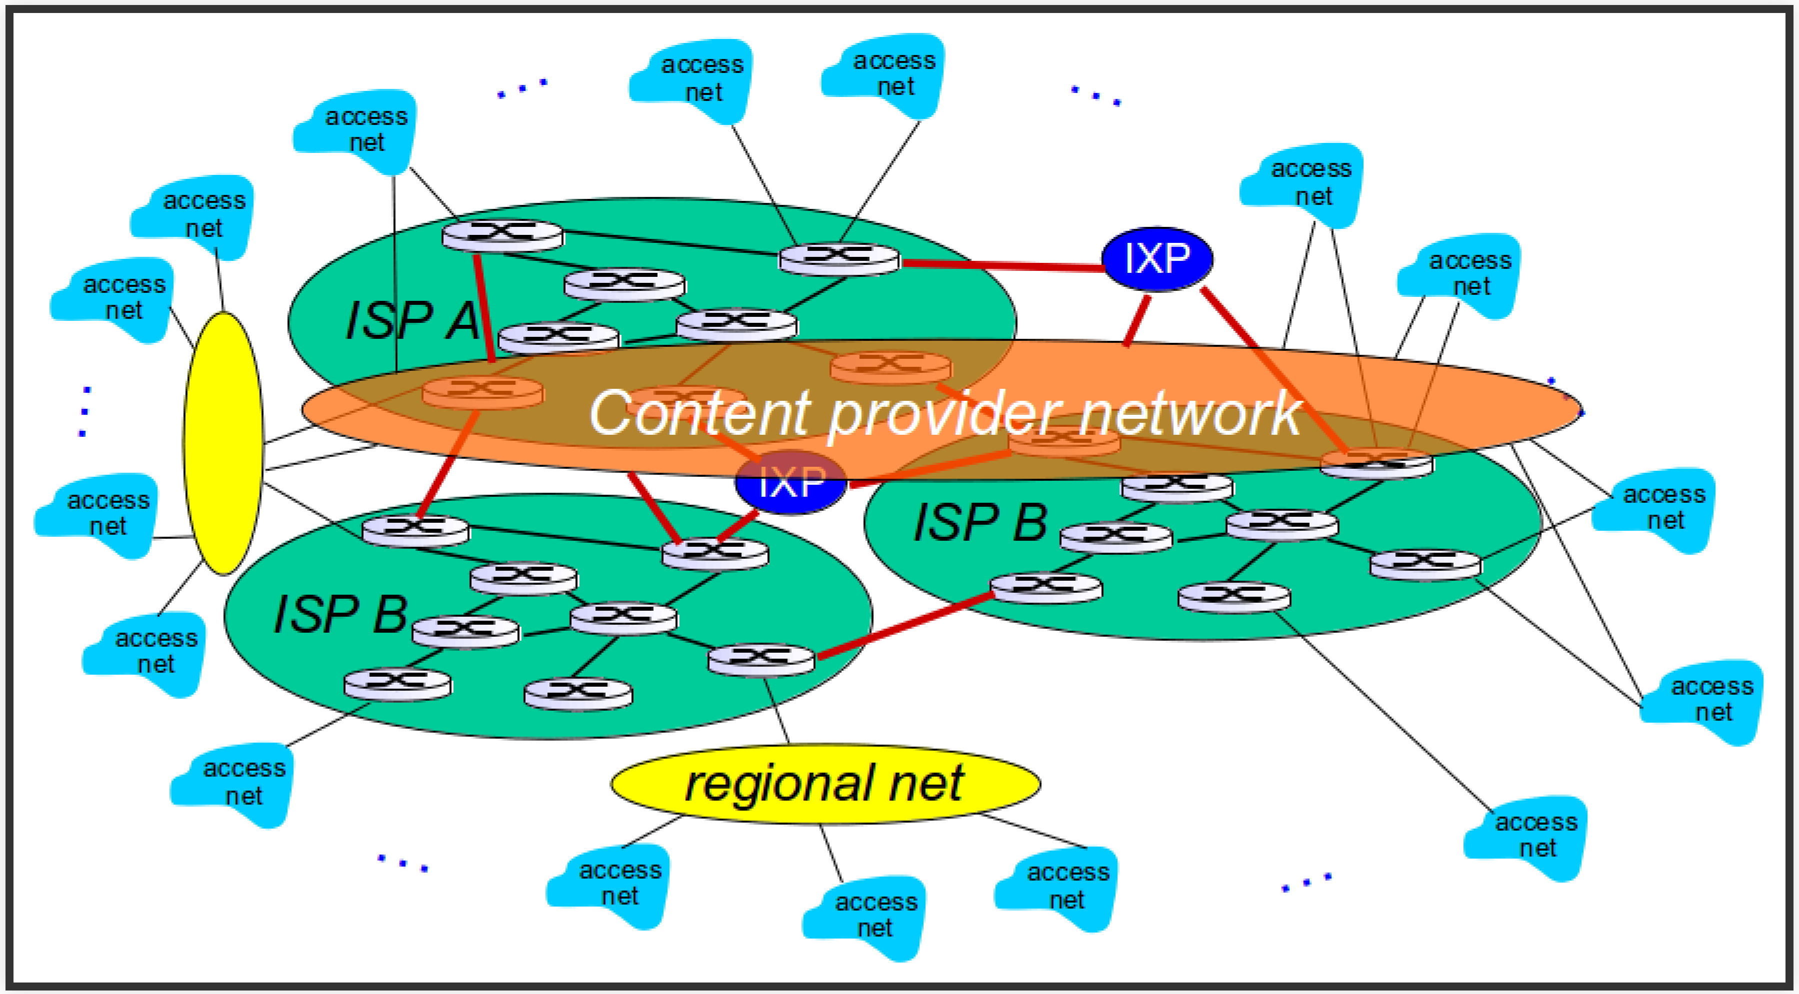

NETWORK OF NETWORKS

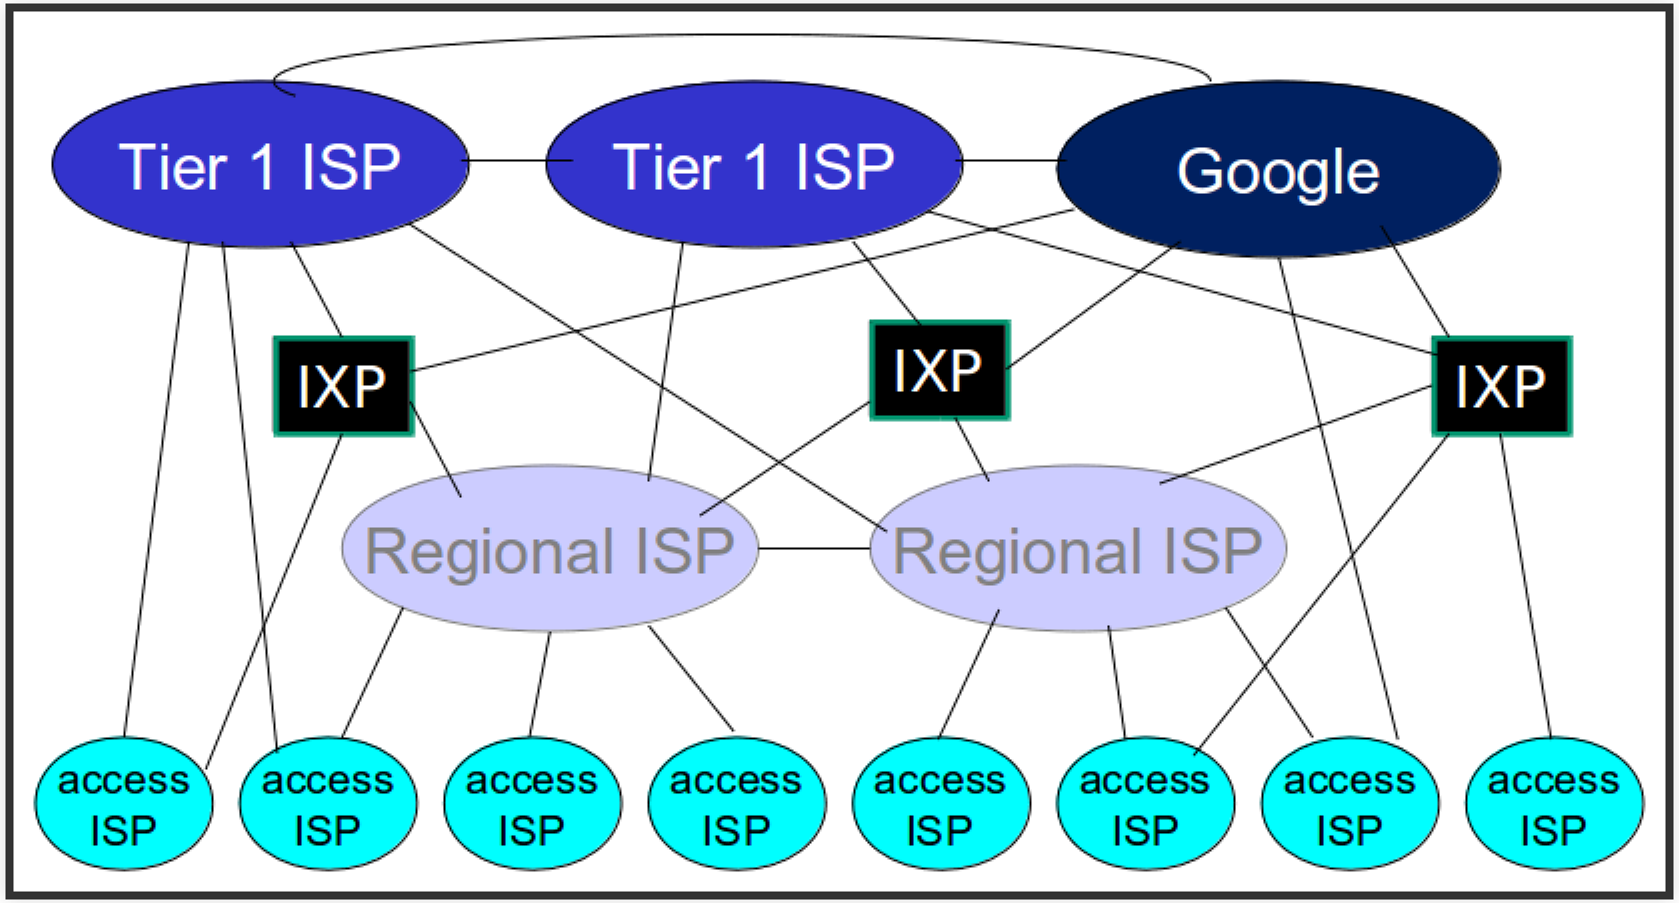

INTERNET STRUCTURE: TIER-1 ISP (SPRINT)

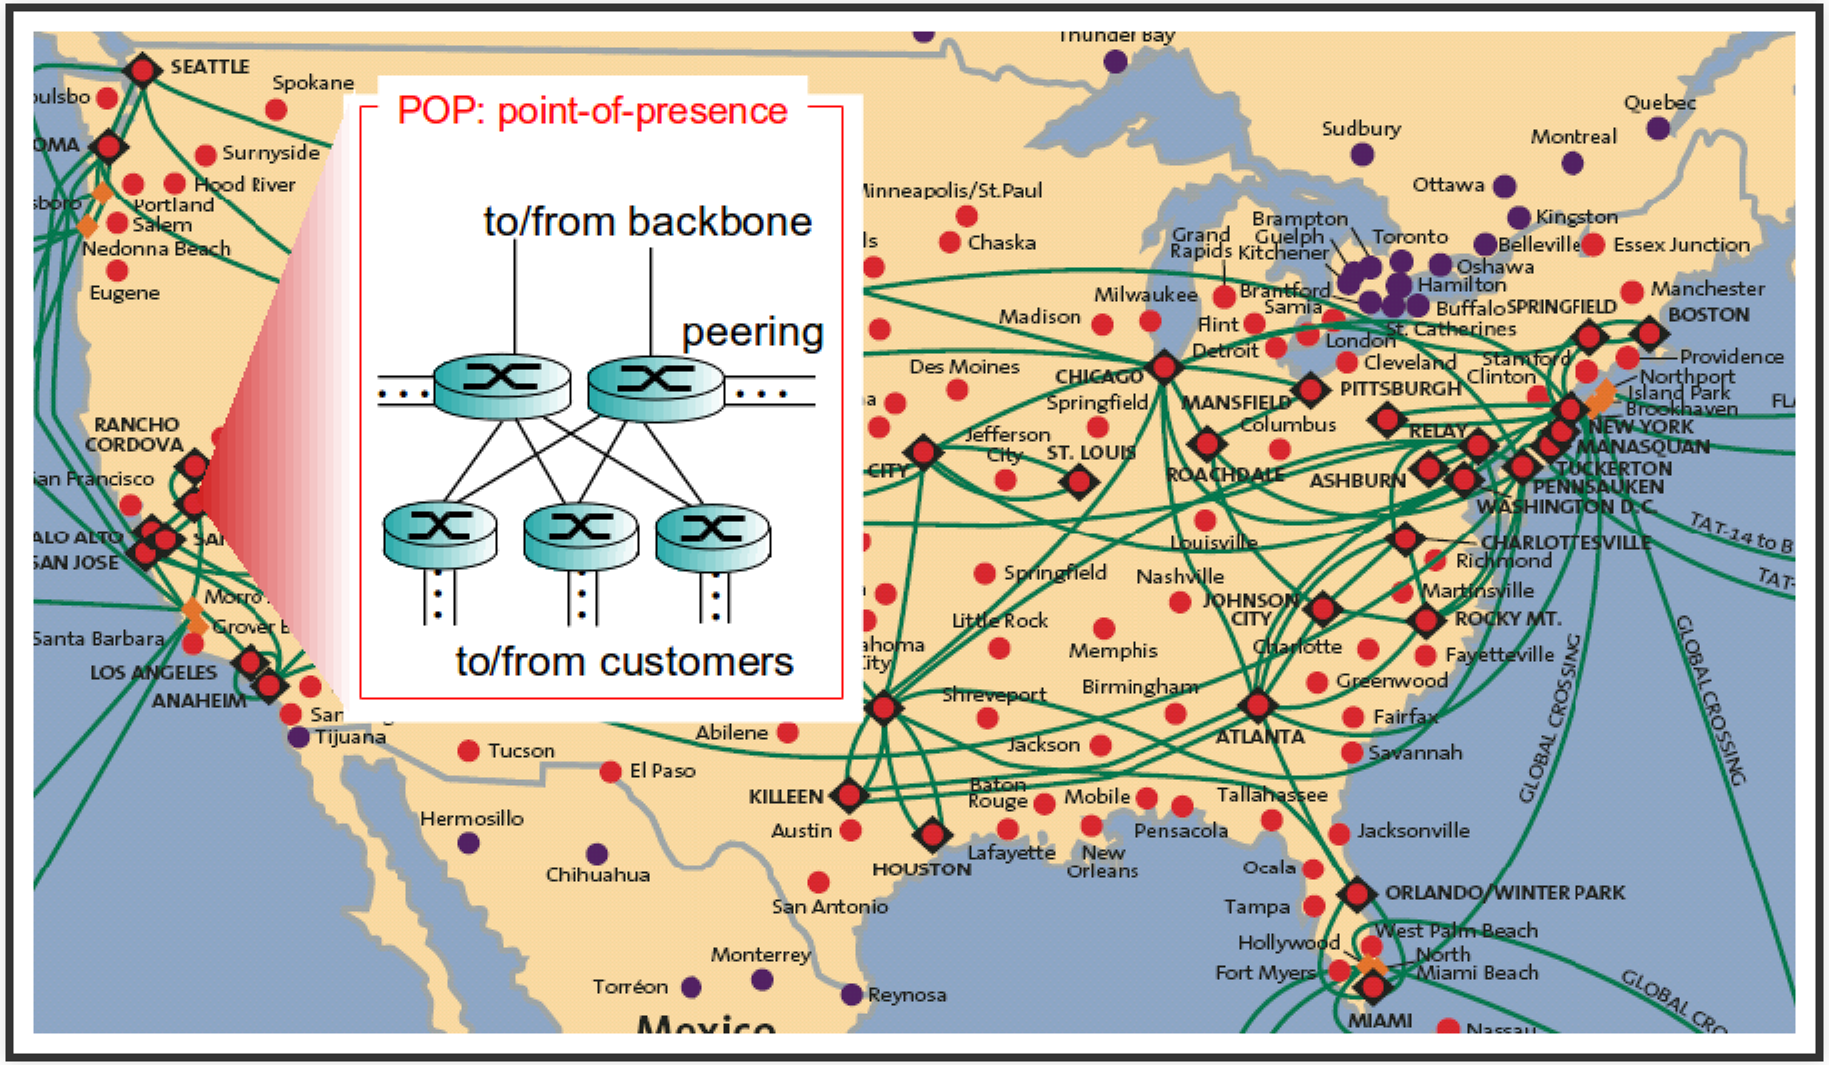

ORGANISATION

IP ADDRESSES

Authorities for IP addresses

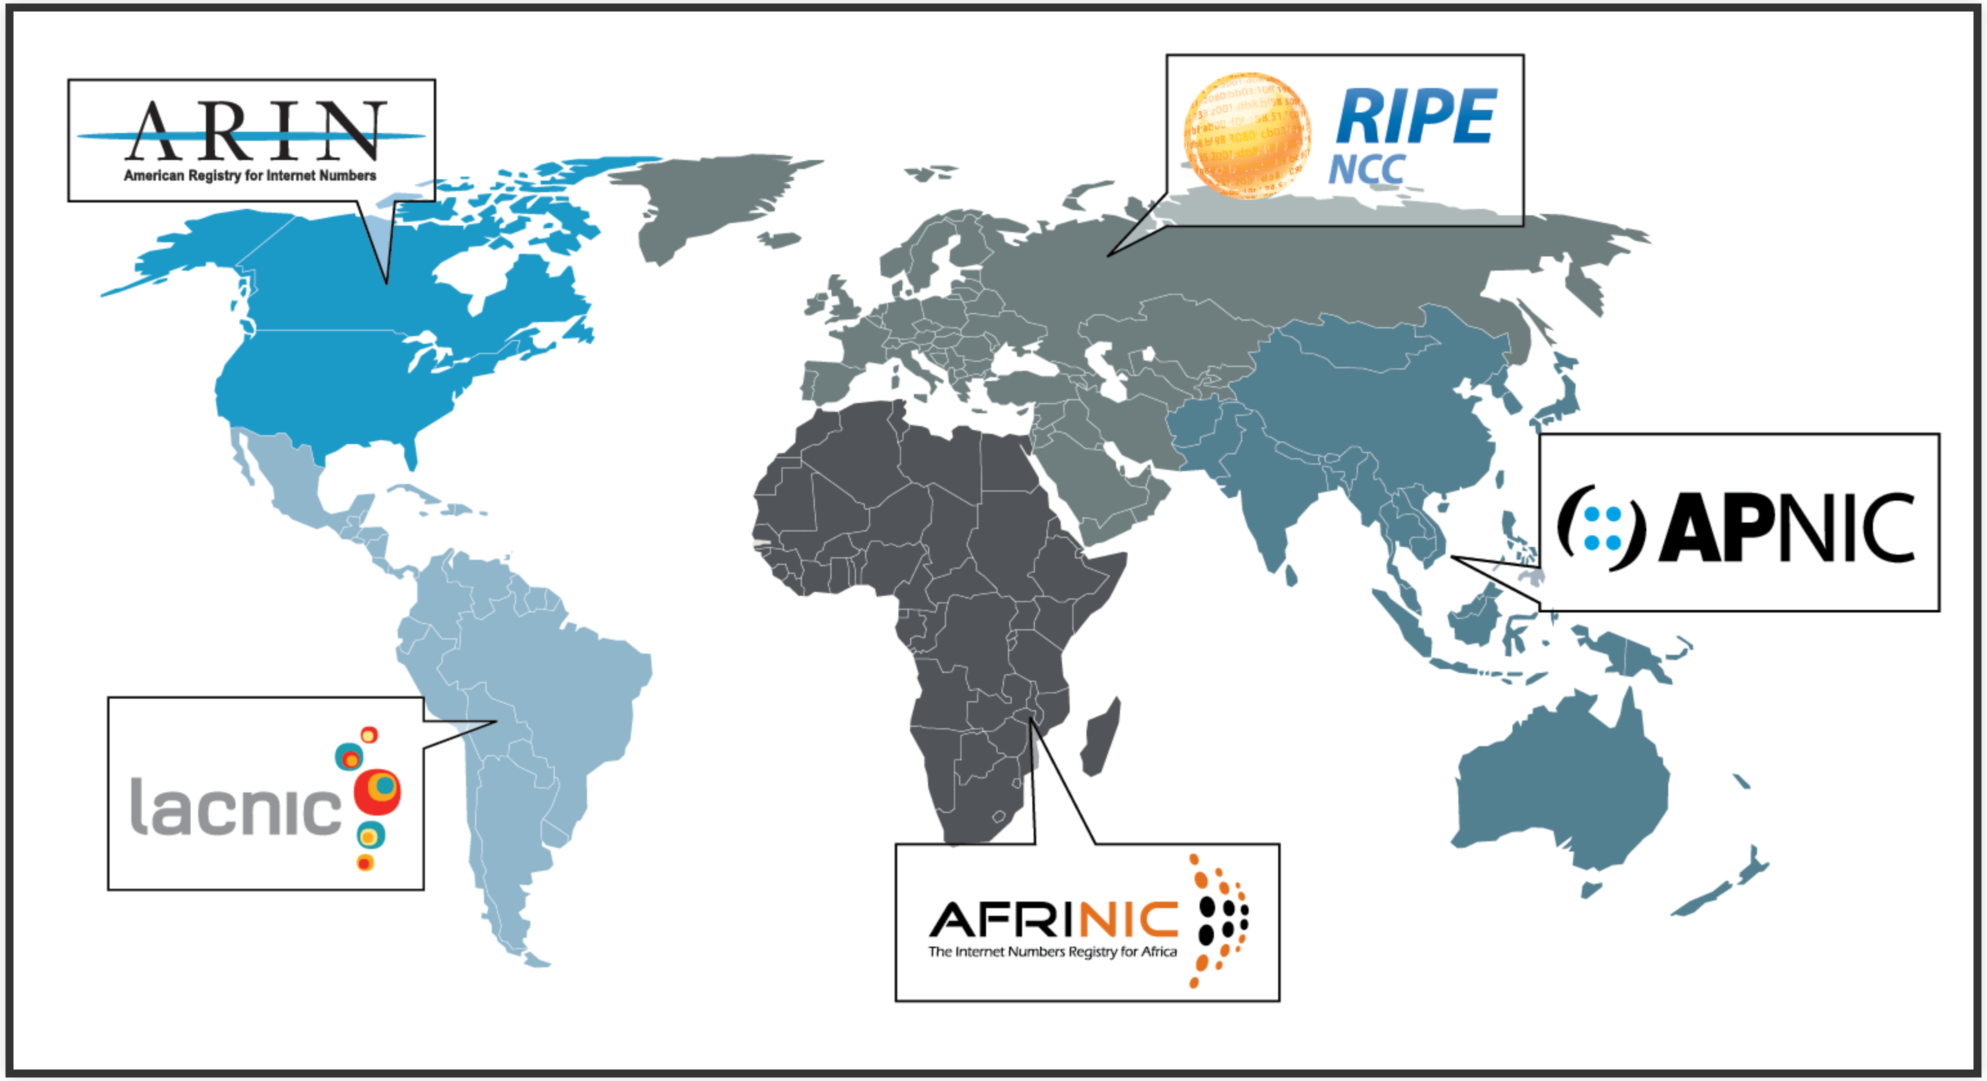

EUROPEAN DOMAIN NAMES

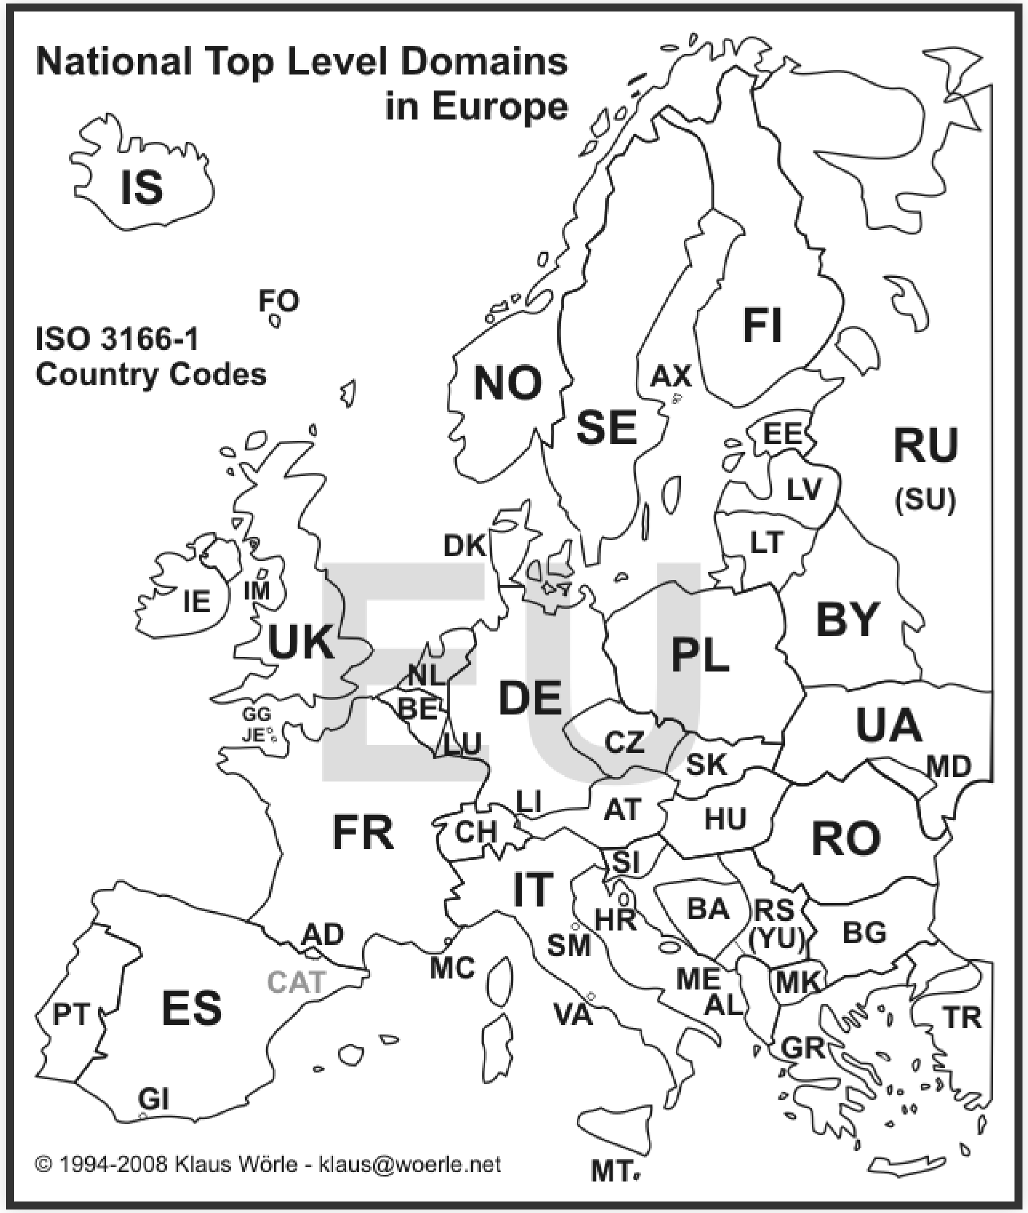

WORLD WIDE DOMAIN NAMES

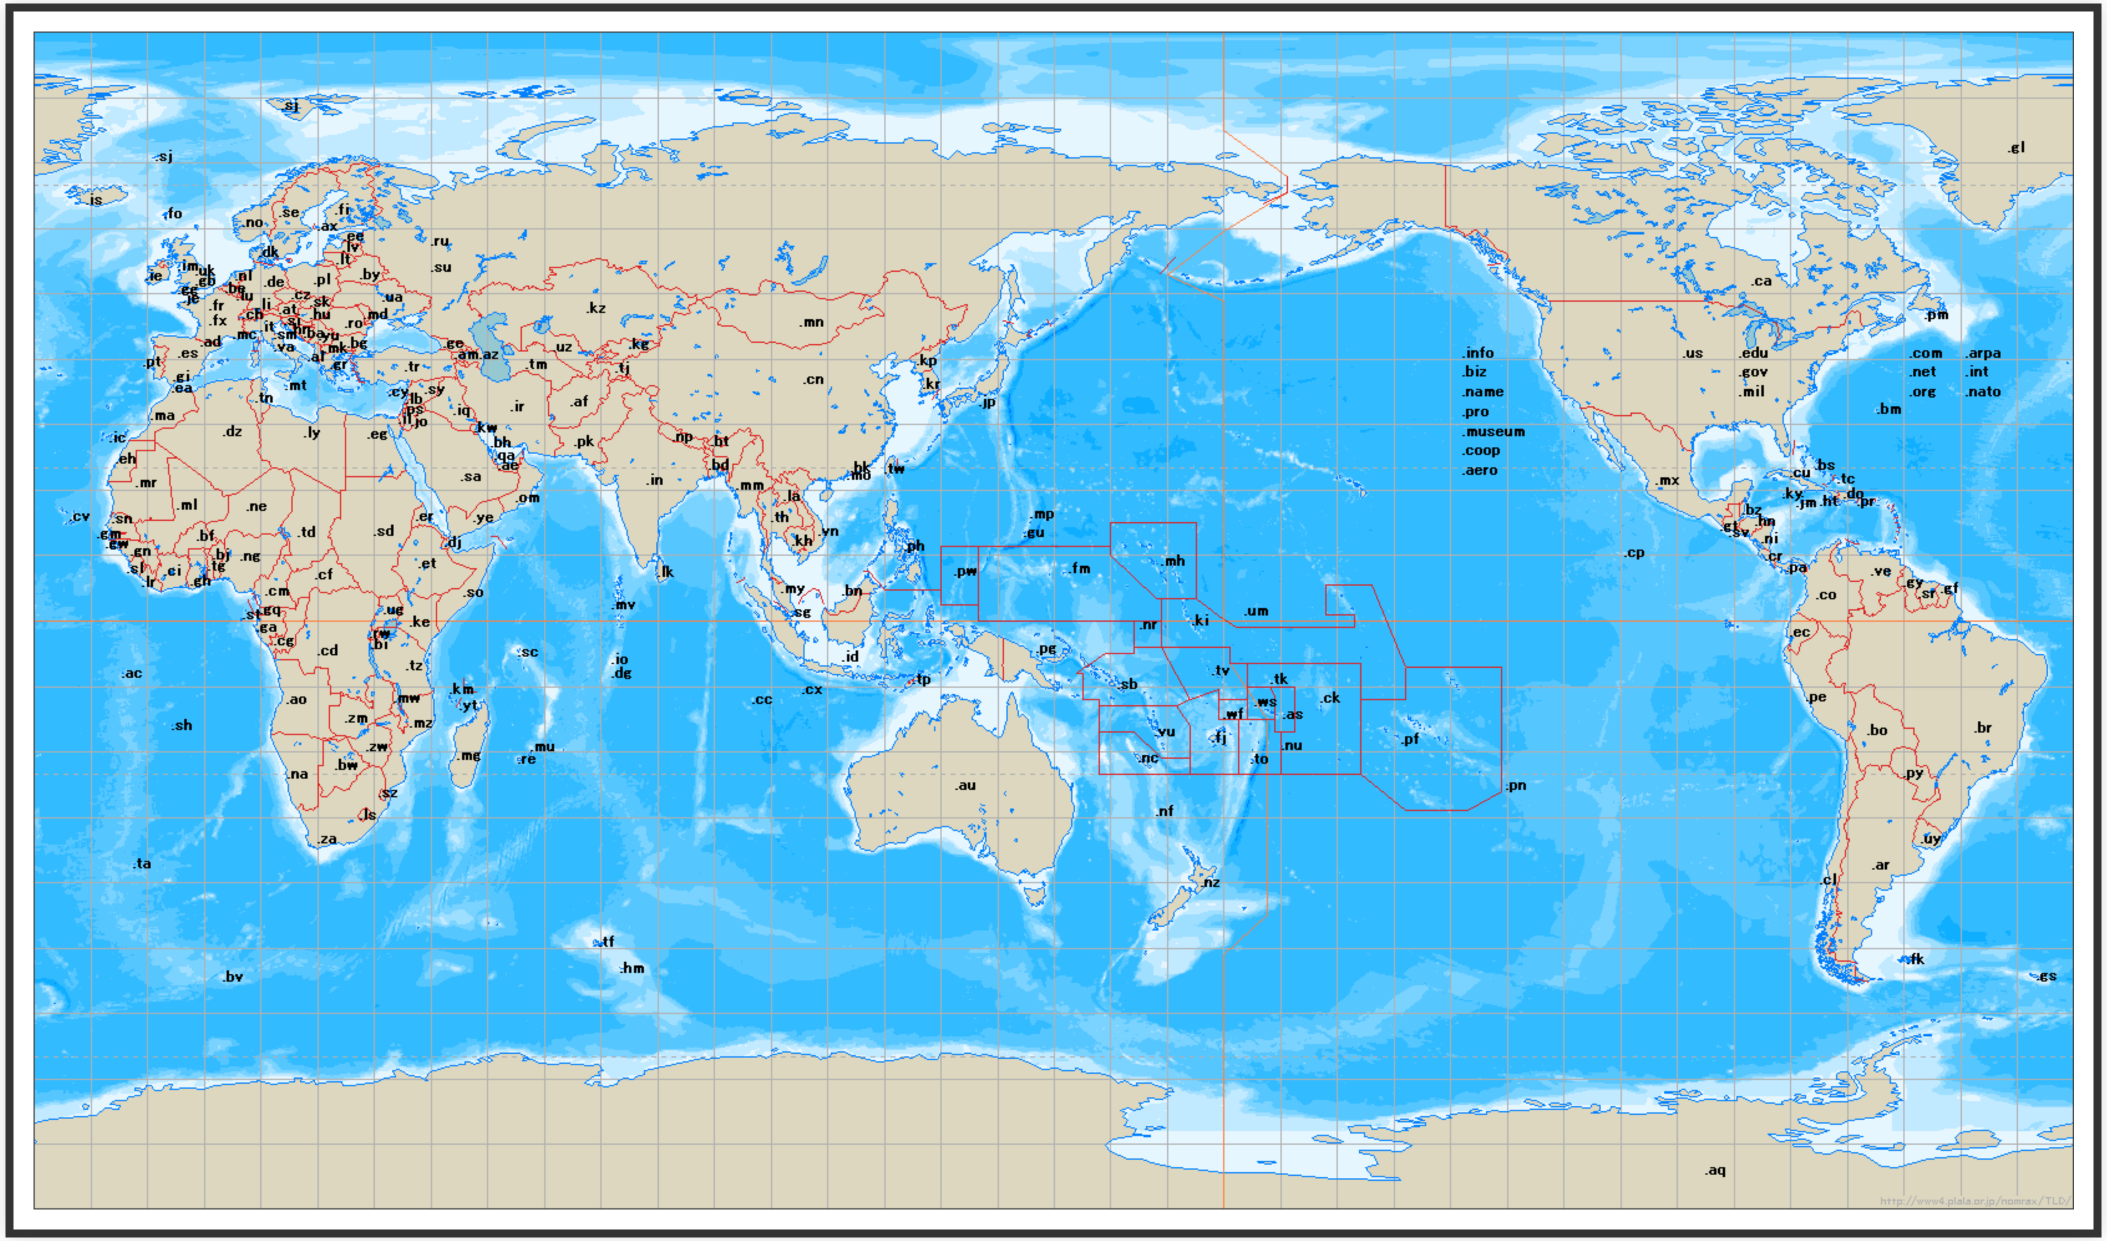

PERFORMANCE

Loss, Delay, Throughput

HOW DO LOSS AND DELAY OCCUR?

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

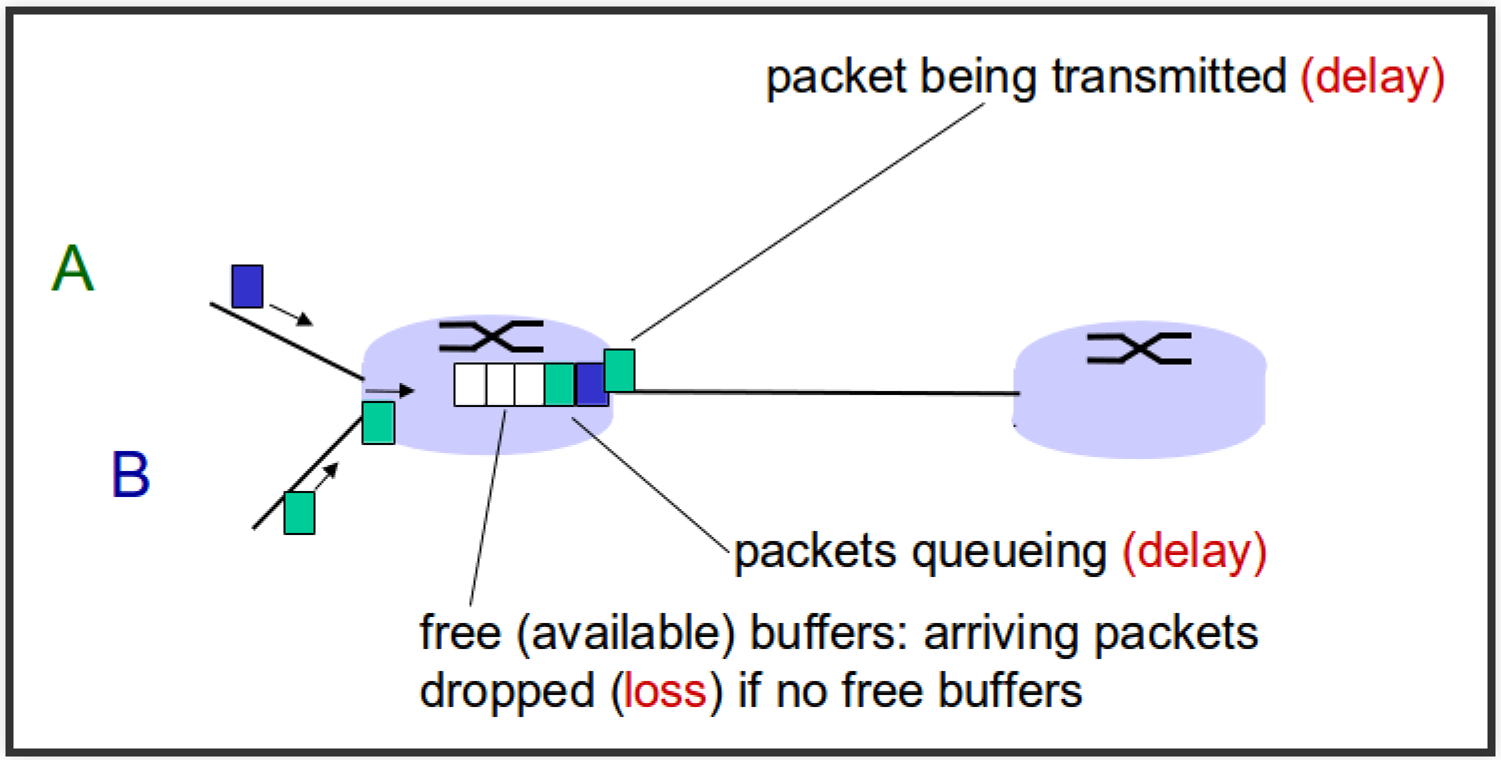

FOUR SOURCES OF PACKET DELAY

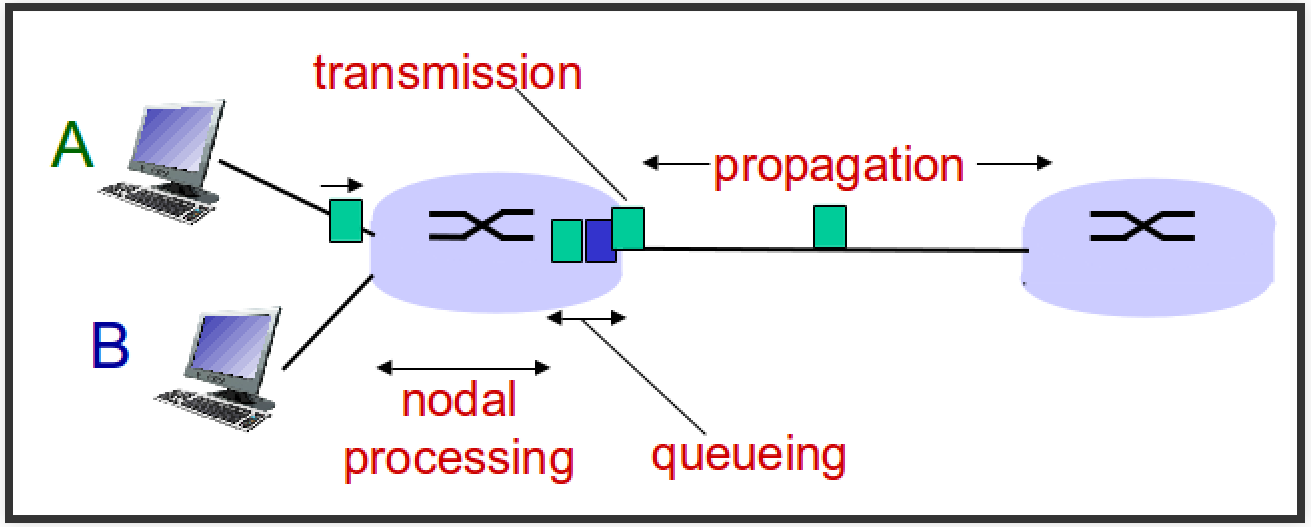

NODAL PROCESSING

- check bit errors
- determine output link
- typically < msec

QUEUEING DELAY

- time waiting at output link for transmission
- depends on congestion level of router

TRANSMISSION DELAY

- L: packet length (bits)
- R: link bandwidth (bps)
- Transmission delay = L/R

PROPAGATION DELAY

- d: length of physical link
- s: propagation speed in medium (~2x10^8 m/sec)
- Propagation delay = d/s

CARAVAN ANALOGY

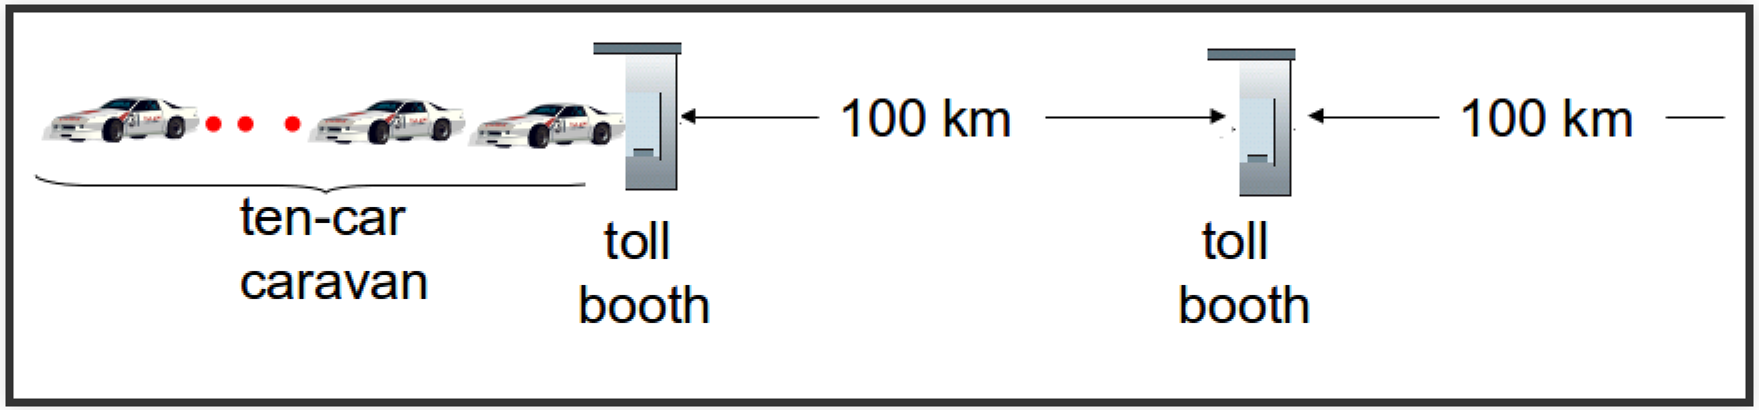

- cars "propagate" at 100 km/hr
- toll booth takes 12 sec to service car (bit transmission time)
- car ~ bit; caravan ~ packet

CARAVAN ANALOGY

Q: *How long until caravan is lined up before 2nd toll booth?*

- time to “push” entire caravan through toll booth onto highway = 12*10 = 120 sec
- time for last car to propagate from 1st to 2nd toll both: 100km/(100km/hr)= 1 hr
- **A:** *62 minutes*

CARAVAN ANALOGY (MORE)

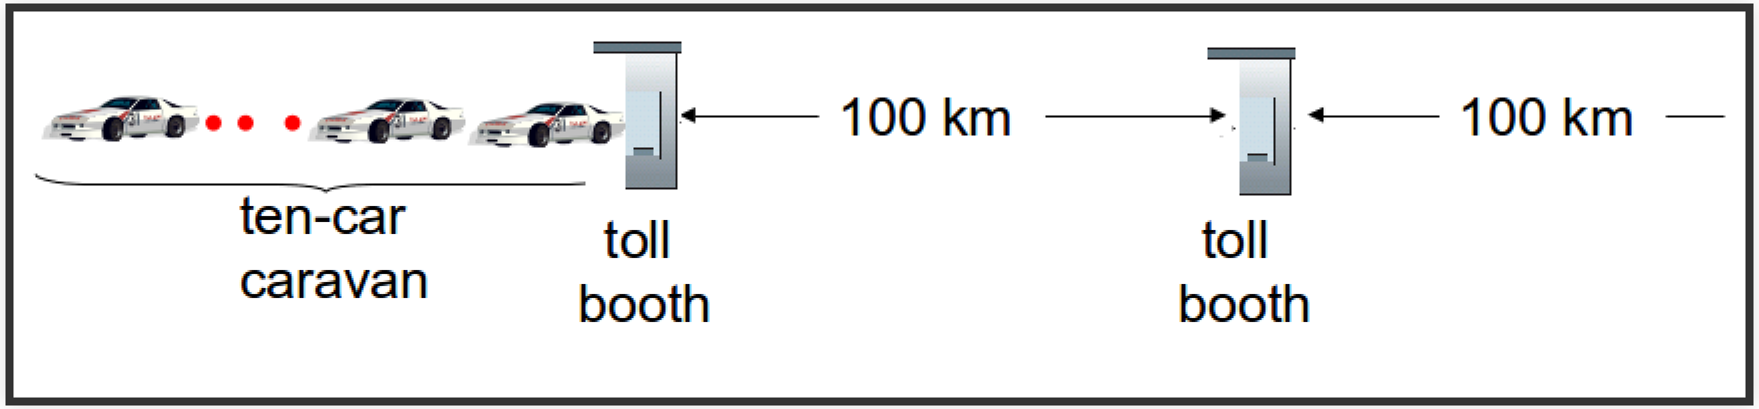

- Suppose cars now “propagate” at 1000 km/hr and suppose toll booth now takes one min to service a car
- **Q:** *Will cars arrive to 2nd booth before all cars serviced at first booth?*
- **A:** *Yes! after 7 min, 1st car arrives at second booth; three cars still at 1st booth.*

QUEUEING DELAY (REVISITED)

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

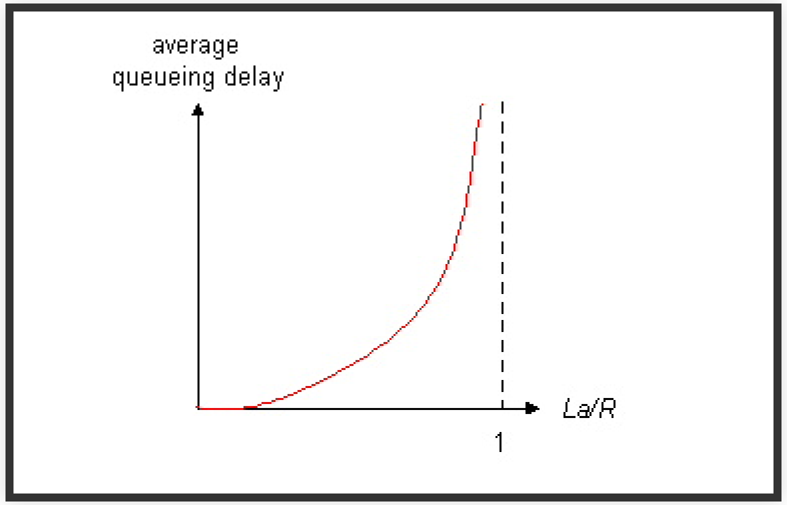

QUEUEING DELAY (REVISITED)

- 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!

"REAL" INTERNET DELAYS AND ROUTES

- what do “real” Internet delay and loss look like?
- *traceroute* program: provides delay measurement from source to router along end-end Internet path towards destination. For all i:
  - sends three packets that will reach router i on path towards destination
  - router i will return packets to sender
  - sender times interval between transmission and reply.

TRACEROUTE

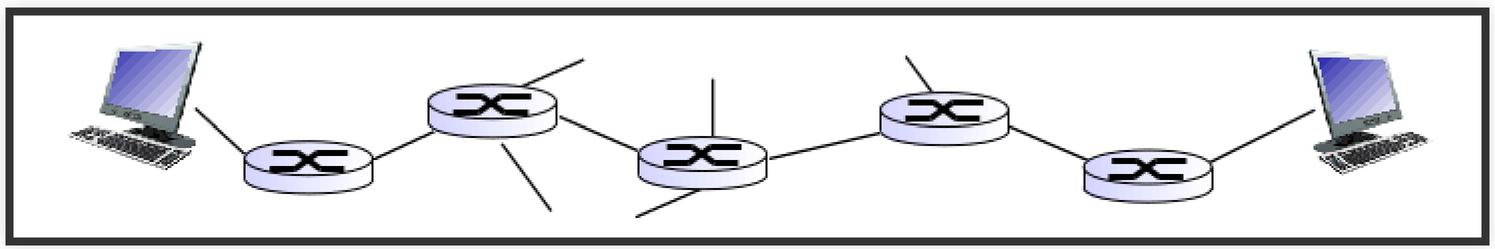

TRACEROUTE

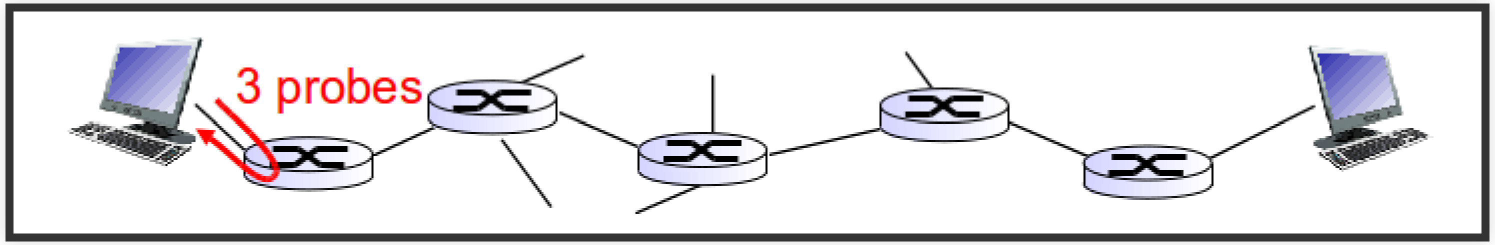

TRACEROUTE

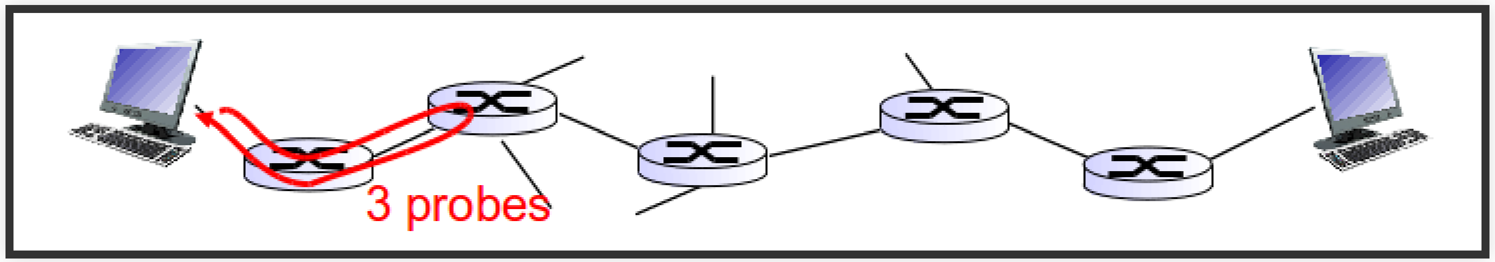

TRACEROUTE

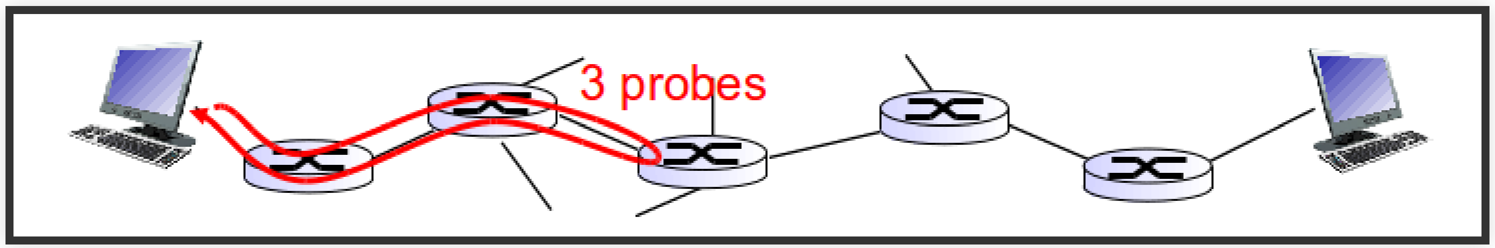

TRACEROUTE

```
$ traceroute cs.umass.edu
traceroute to cs.umass.edu (128.119.240.93), 30 hops max, 60 byte packets
 1  192.168.1.1 (192.168.1.1)  4.118 ms  5.399 ms  6.217 ms
 2  xe-2-0-0-1104.odnqe10.dk.ip.tdc.net (80.162.66.237)  27.879 ms  29.274 ms  29.625 ms
 3  as0-0.ashbnqp1.us.ip.tdc.net (83.88.31.141)  152.961 ms  152.947 ms  152.957 ms
 4  eeq-exchange.tr01-asbnva01.transitrail.net (206.126.236.45)  155.241 ms  154.874 ms  15
 5  ae-2.0.ny0.tr-cps.internet2.edu (64.57.20.197)  147.600 ms  151.800 ms  150.271 ms
 6  64.57.21.210 (64.57.21.210)  157.587 ms  137.632 ms  140.931 ms
 7  nox300gw1-peer--207-210-142-242.nox.org (207.210.142.242)  146.500 ms  150.512 ms  151.
 8  core1-rt-xe-0-0-0.gw.umass.edu (192.80.83.101)  153.933 ms  158.008 ms  160.034 ms
 9  lgrc-rt-106-8-po-10.gw.umass.edu (128.119.0.233)  167.644 ms  164.152 ms  164.443 ms
10  128.119.3.32 (128.119.3.32)  167.289 ms  169.114 ms  168.008 ms
11  loki.cs.umass.edu (128.119.240.93)  167.883 ms  169.216 ms  170.235 ms
```

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

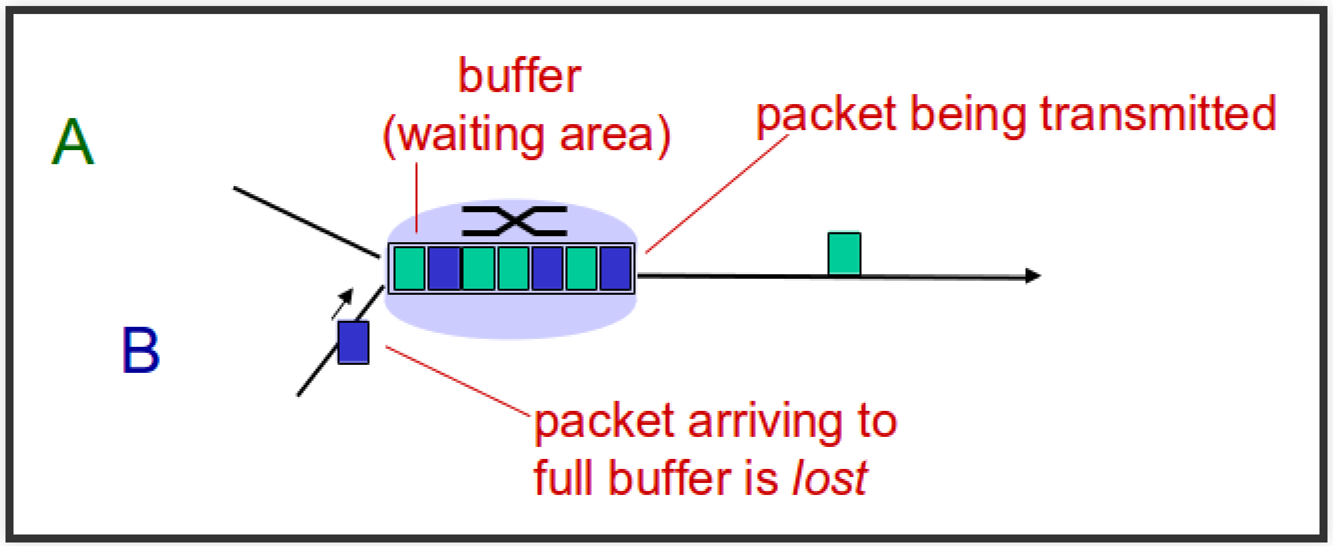

THROUGHPUT

- **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

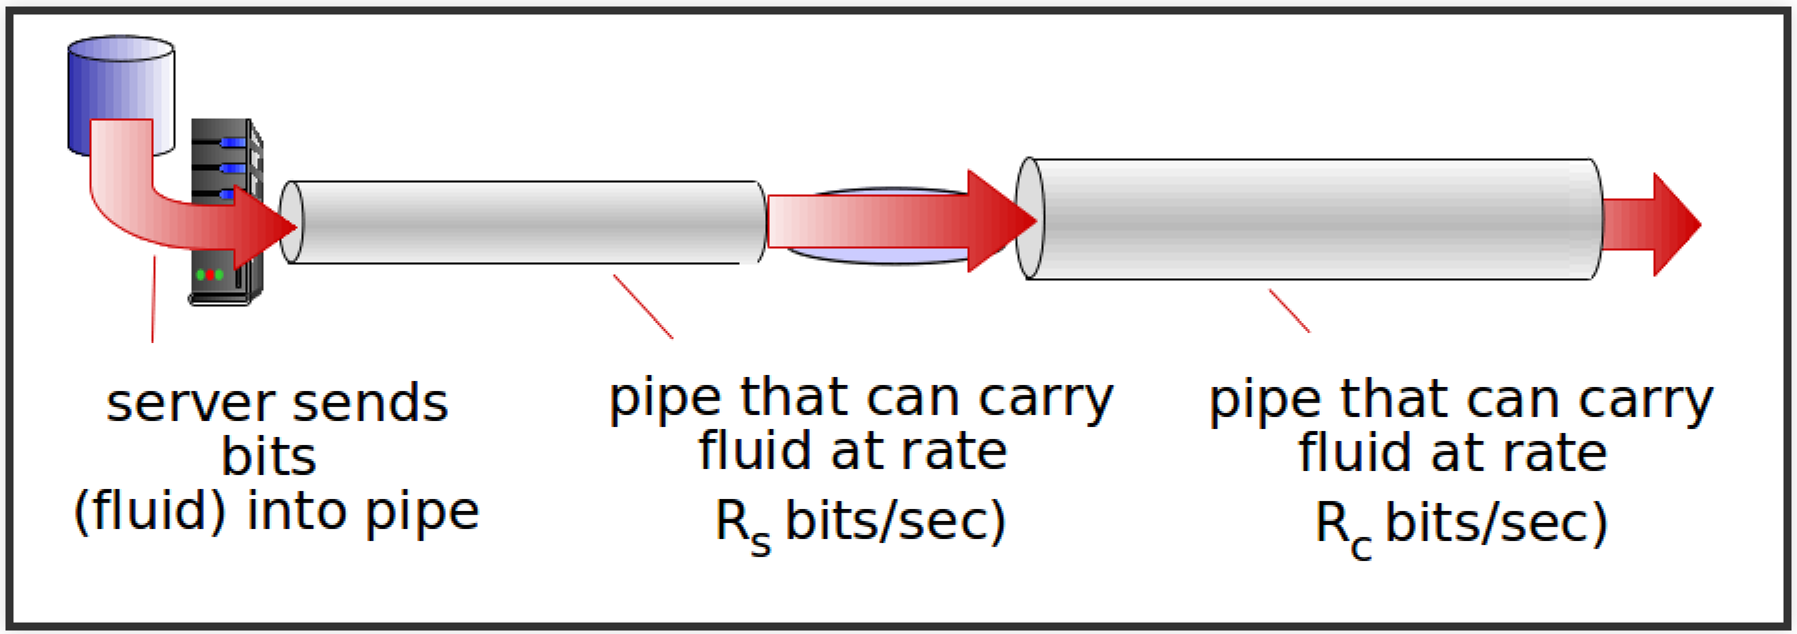

THROUGHPUT - MORE

- $R_S < R_C$ What is average end-end throughput?
- $R_S > R_C$ What is average end-end throughput?

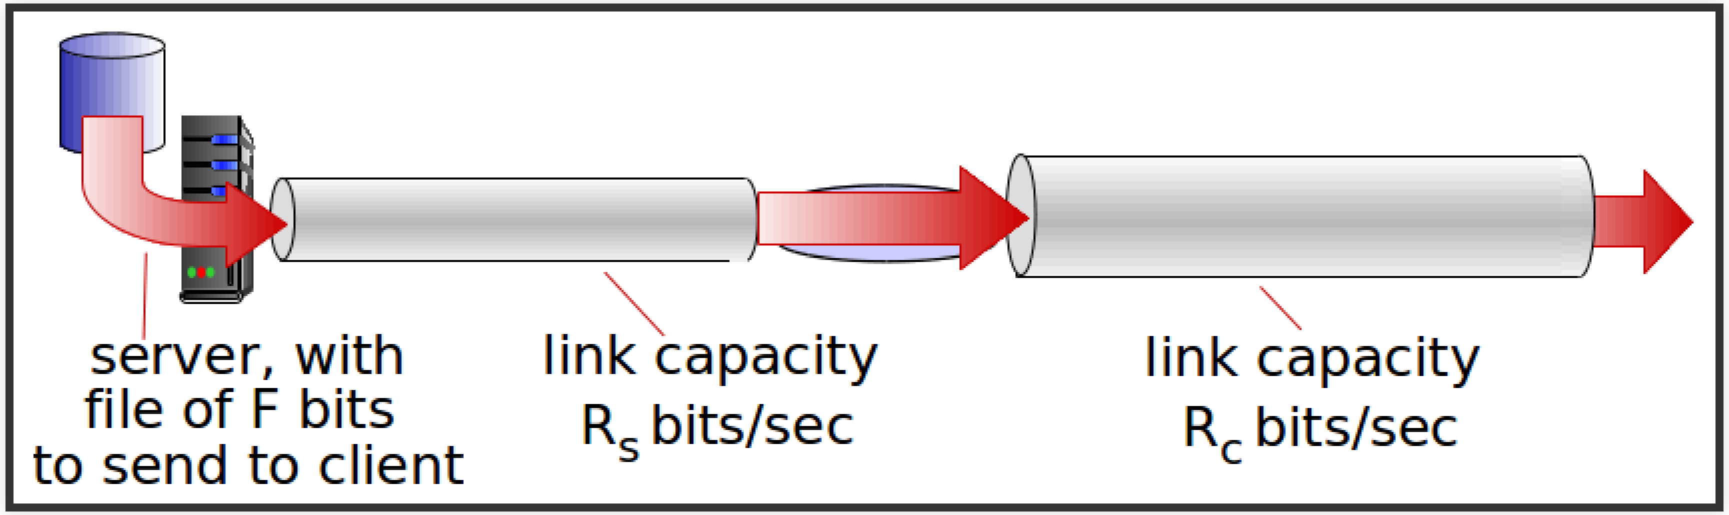

bottleneck link: link on end-end path that constrains end-end throughput

THROUGHPUT: INTERNET SCENARIO

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

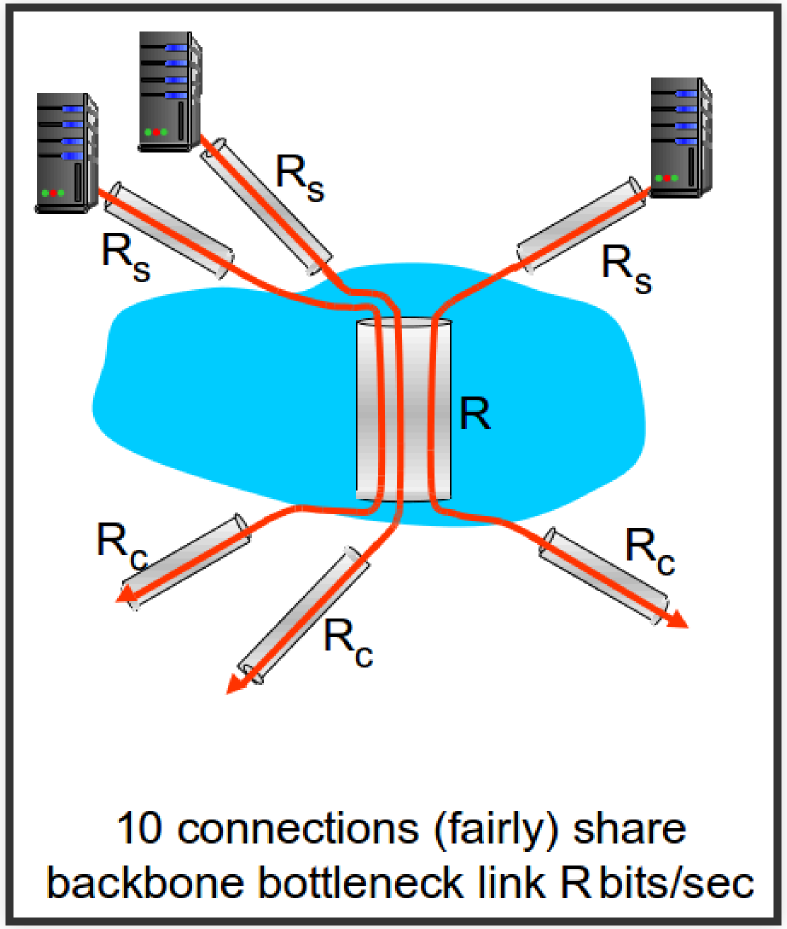

PROTOCOL LAYERS, SERVICE MODELS

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?

ORGANIZATION OF AIR TRAVEL

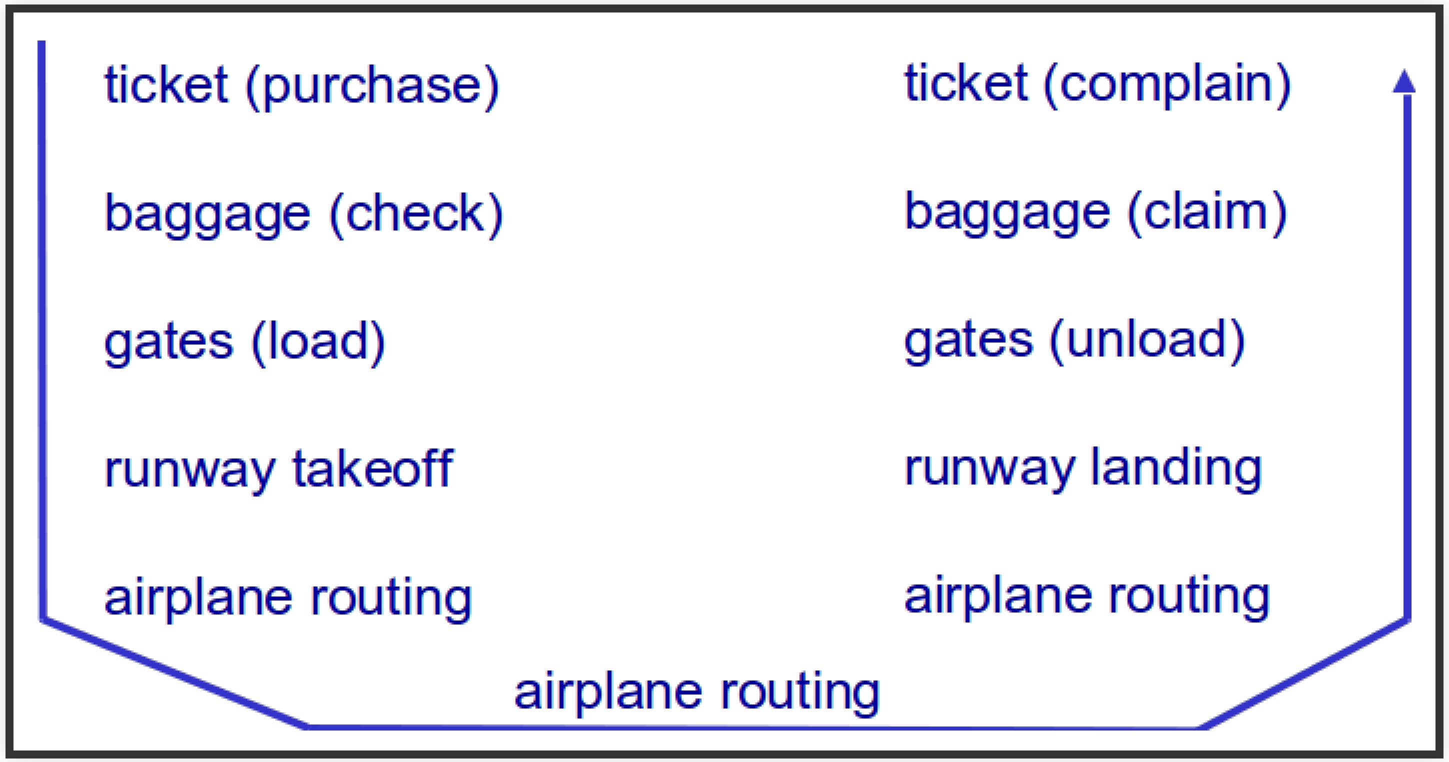

a series of steps

LAYERING OF AIRLINE FUNCTIONALITY

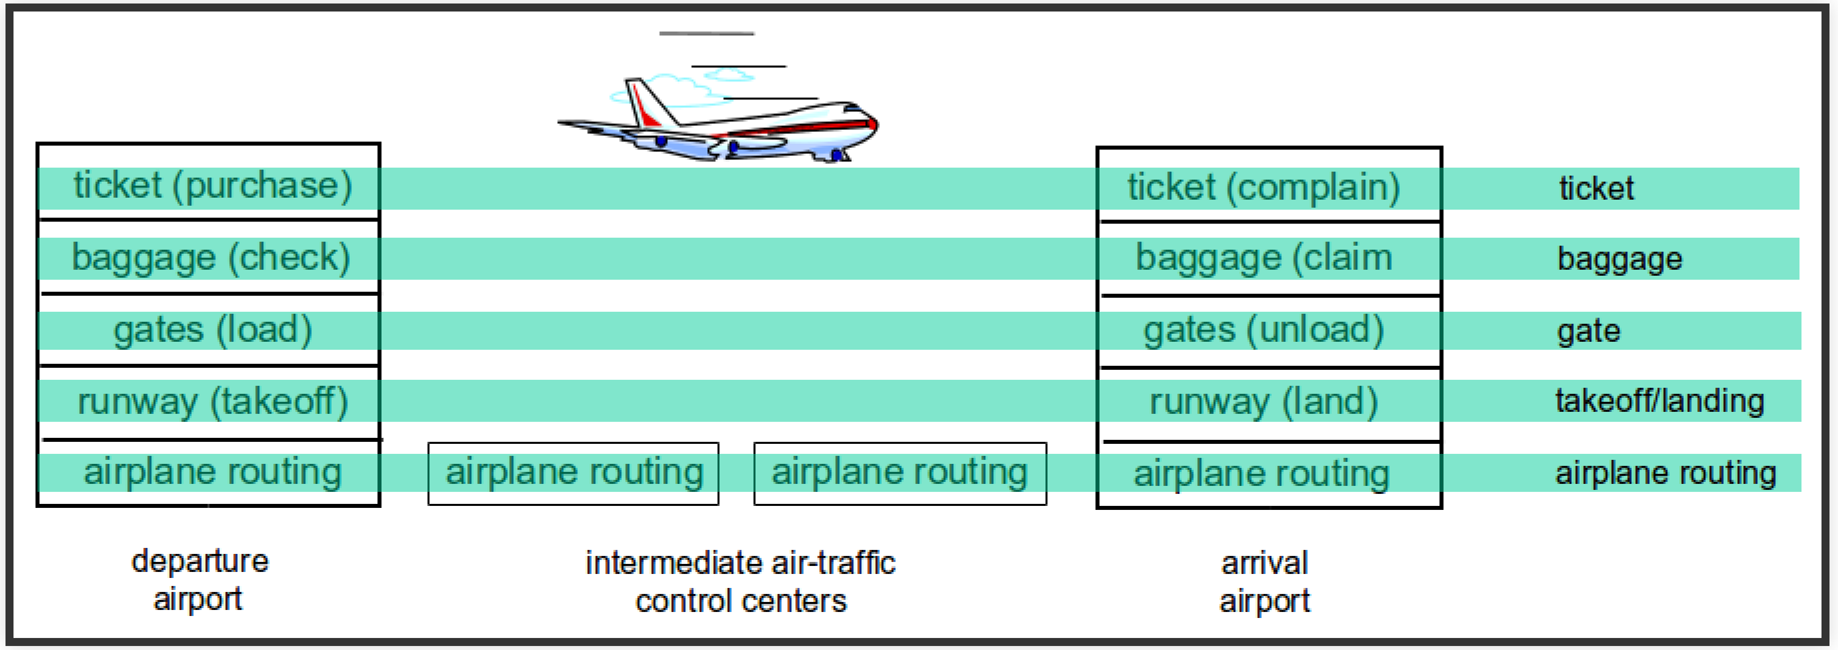

- layers: each layer implements a service
  - via its own internal-layer actions
  - relying on services provided by layer below

WHY LAYERING?

dealing with complex systems:

- explicit structure allows identification, relationship of complex system's pieces
  - layered reference model for discussion
- modularization eases maintenance, updating of system
  - change of implementation of layer's service transparent to rest of system
  - e.g., change in gate procedure doesn't affect 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”

INTERNET PROTOCOL STACK

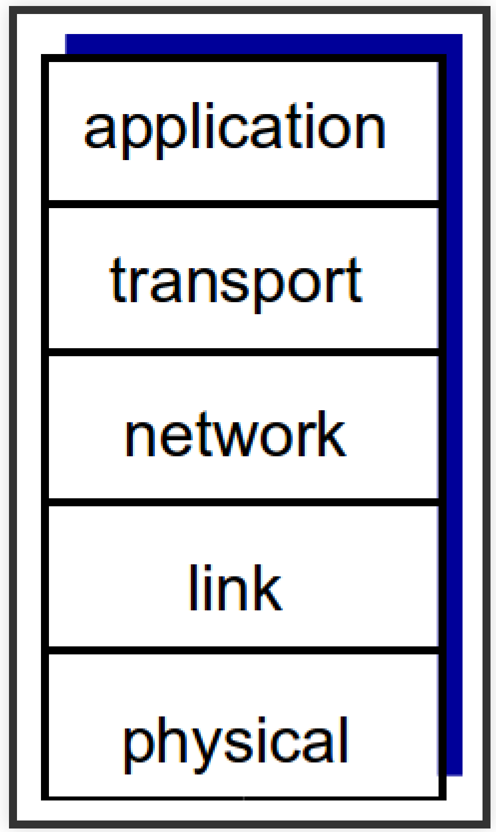

ISO/OSI REFERENCE MODEL

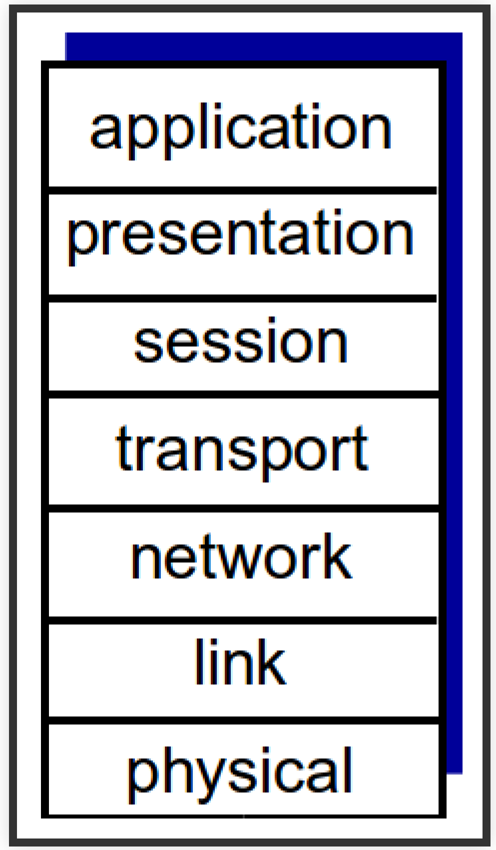

ISO/OSI REFERENCE MODEL

- 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?

ENCAPSULATION

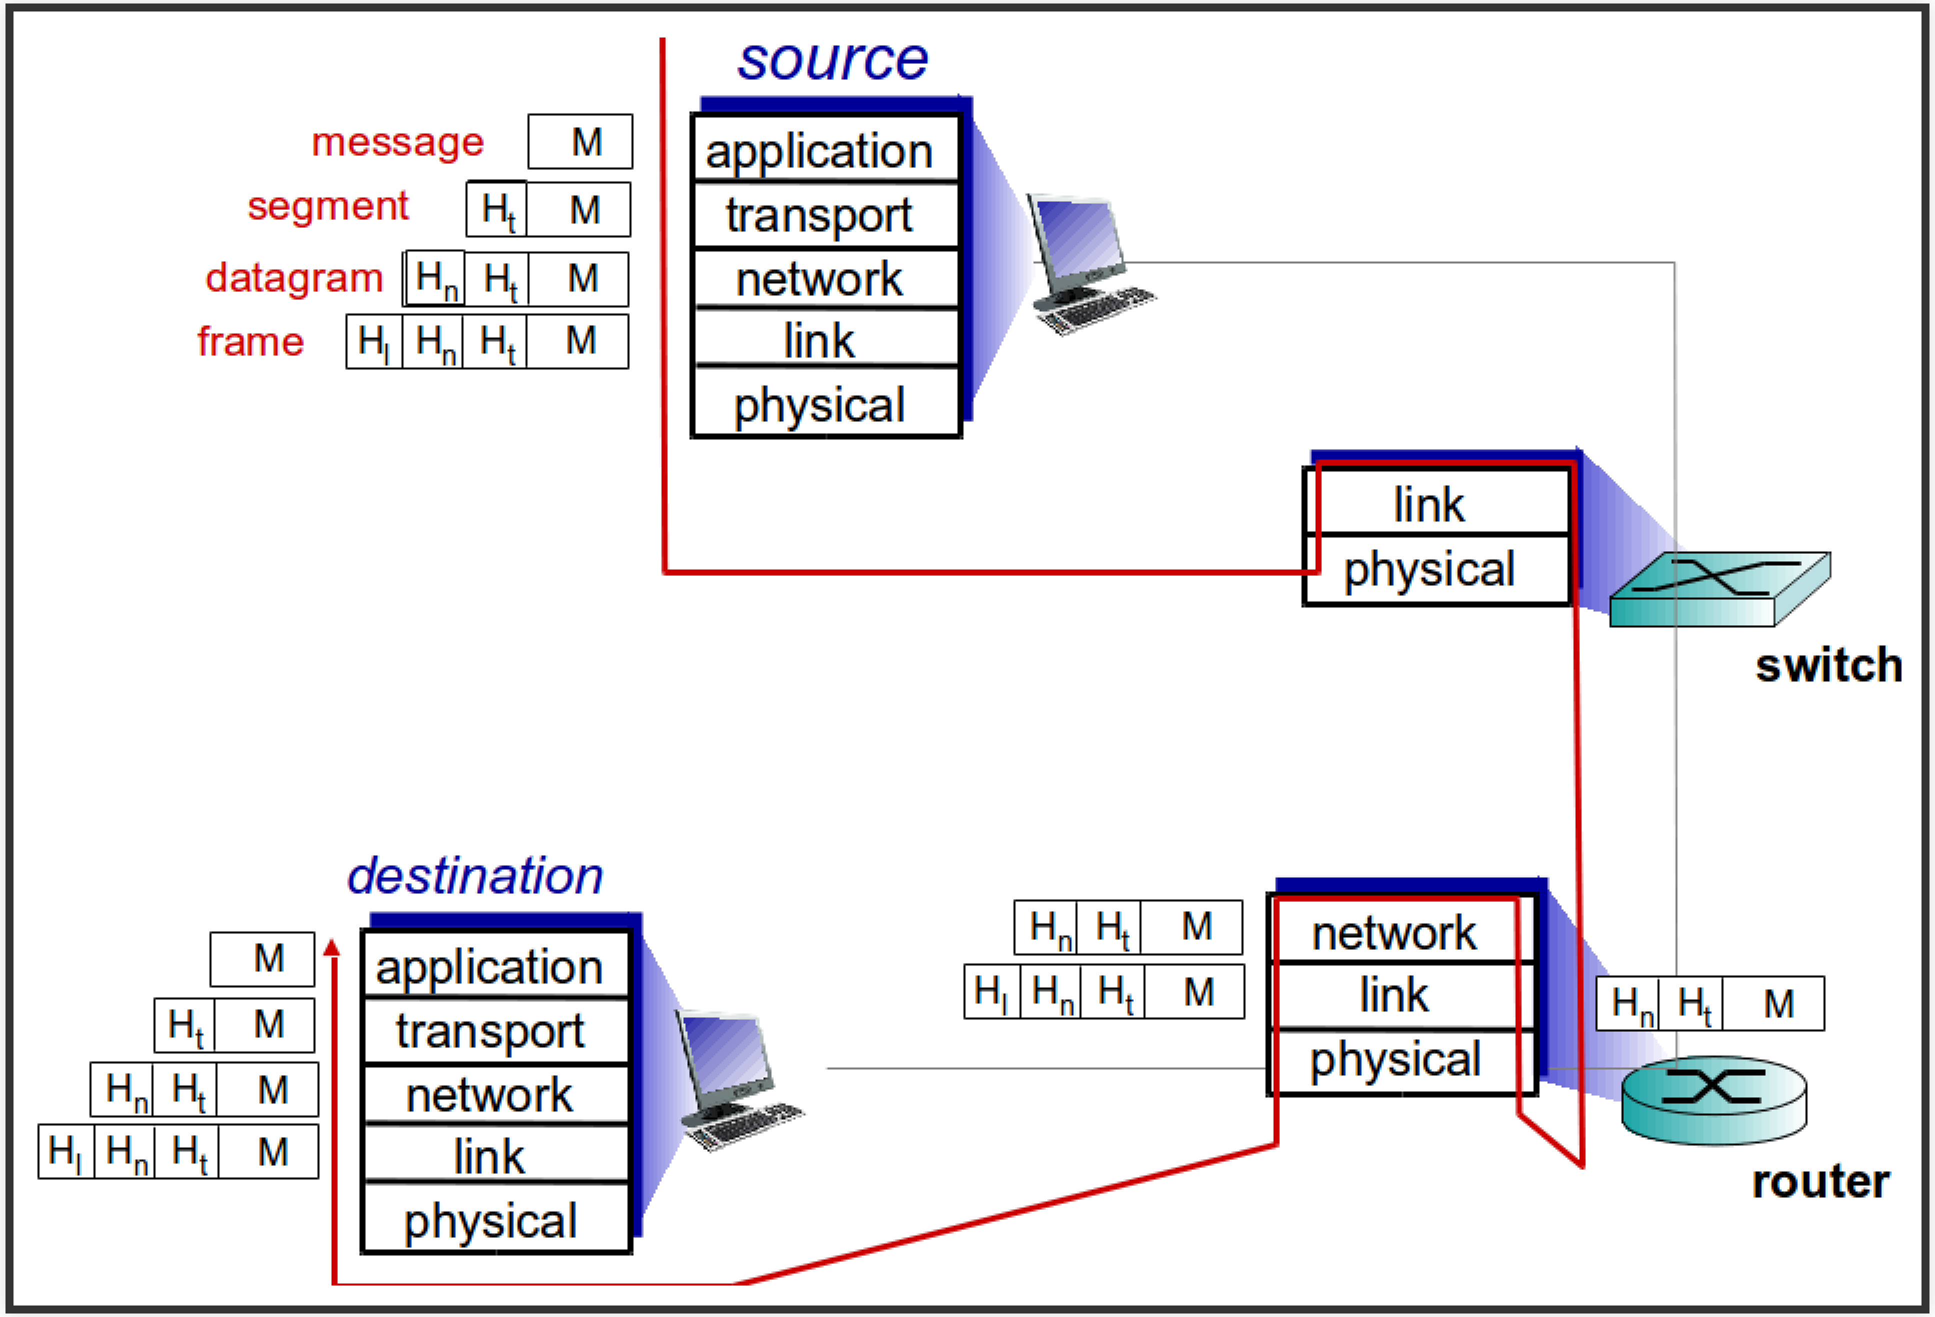

SECURITY THREADS

Threats to the Internet and Users of the internet

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

NETWORK SECURITY

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!

FRAGILE OPTIC CABLES

RISK OF DISCONNECTION

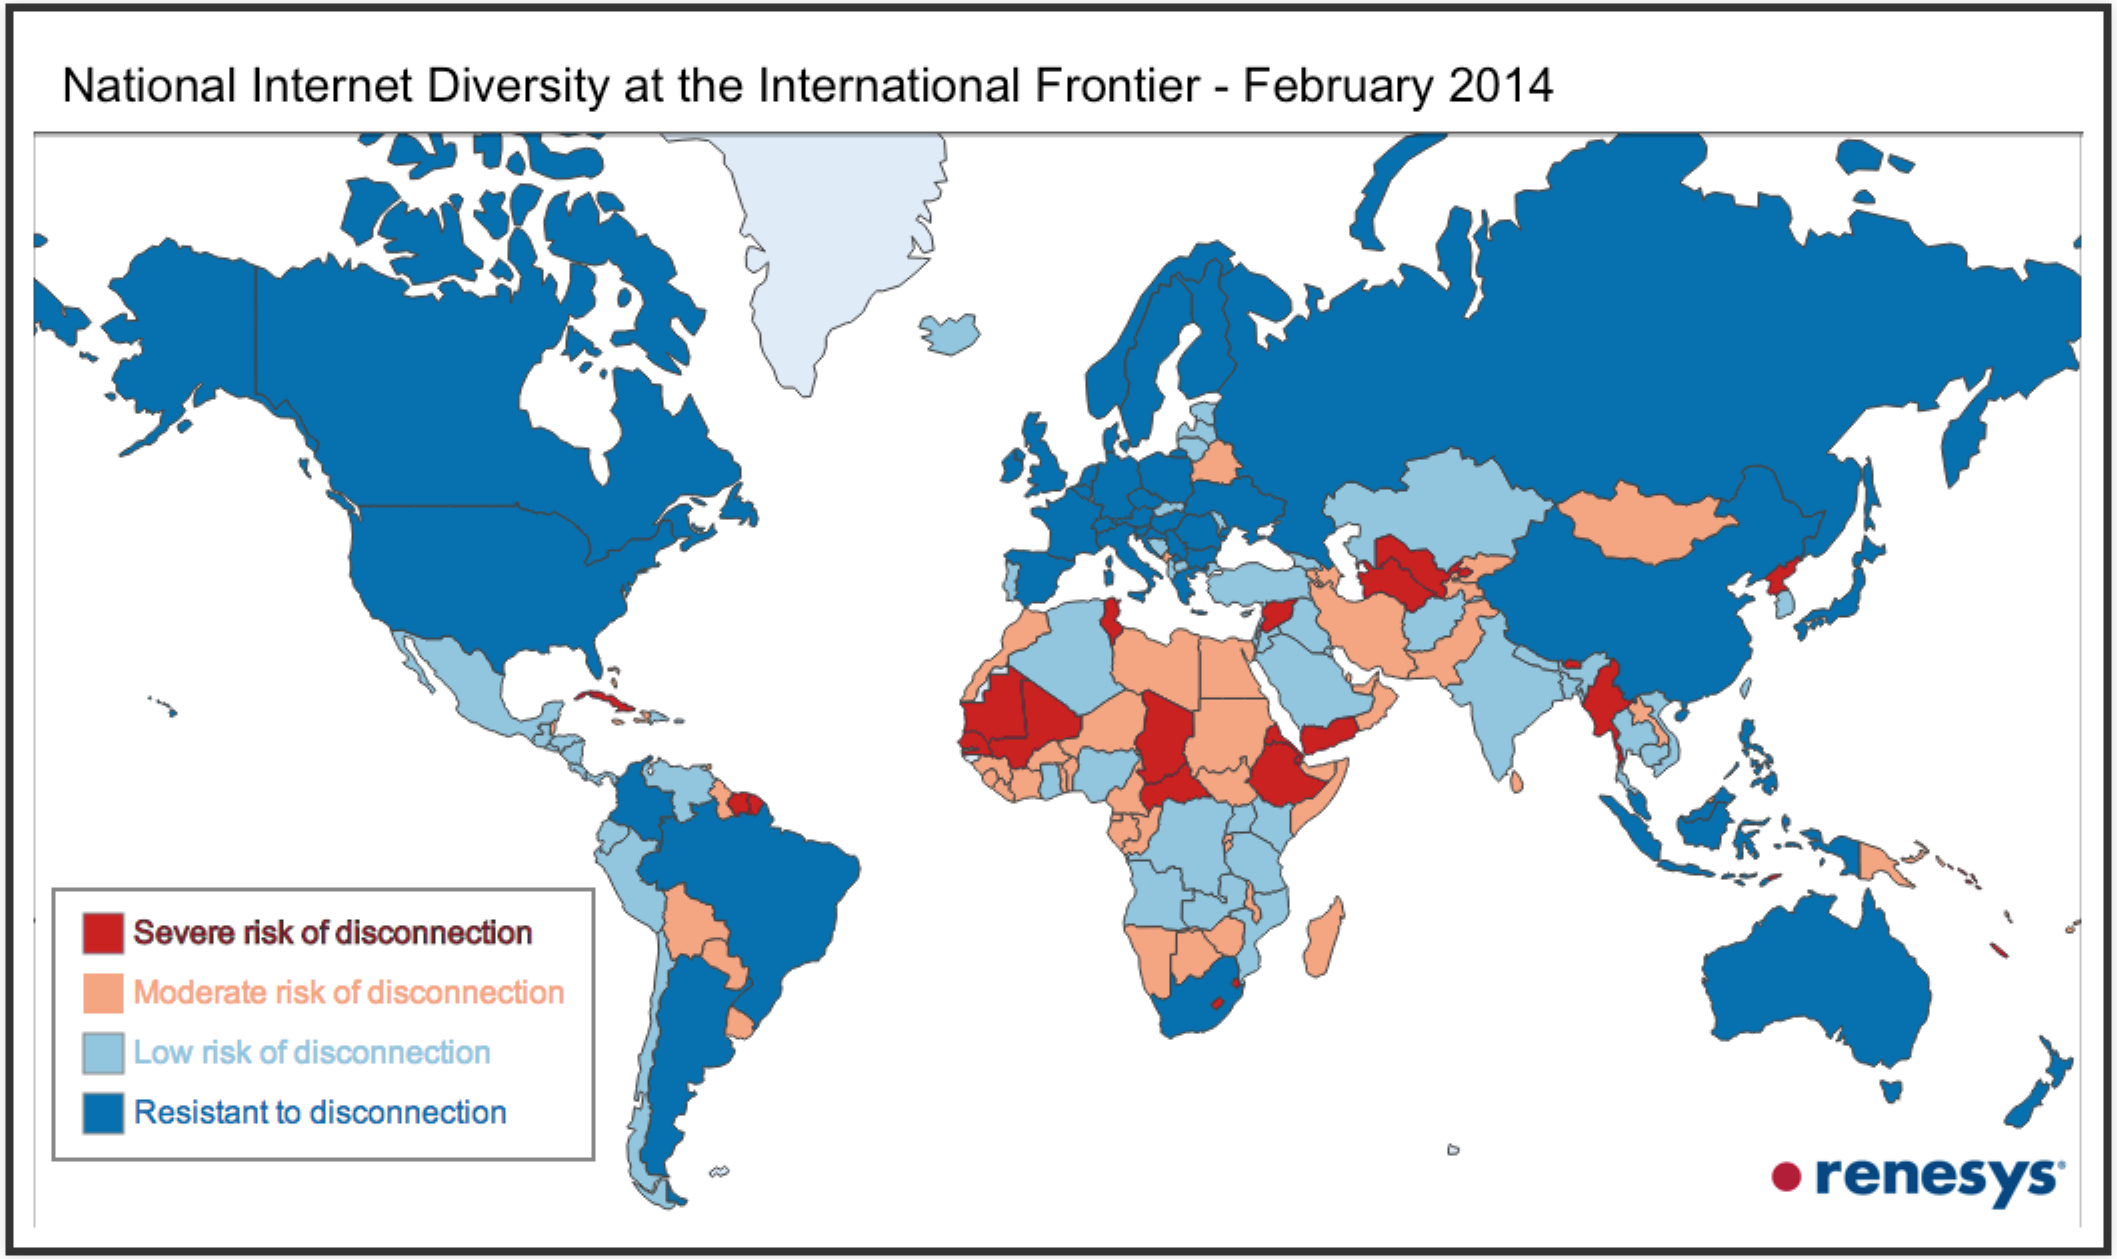

SYRIA'S PHYSICAL CONNECTION

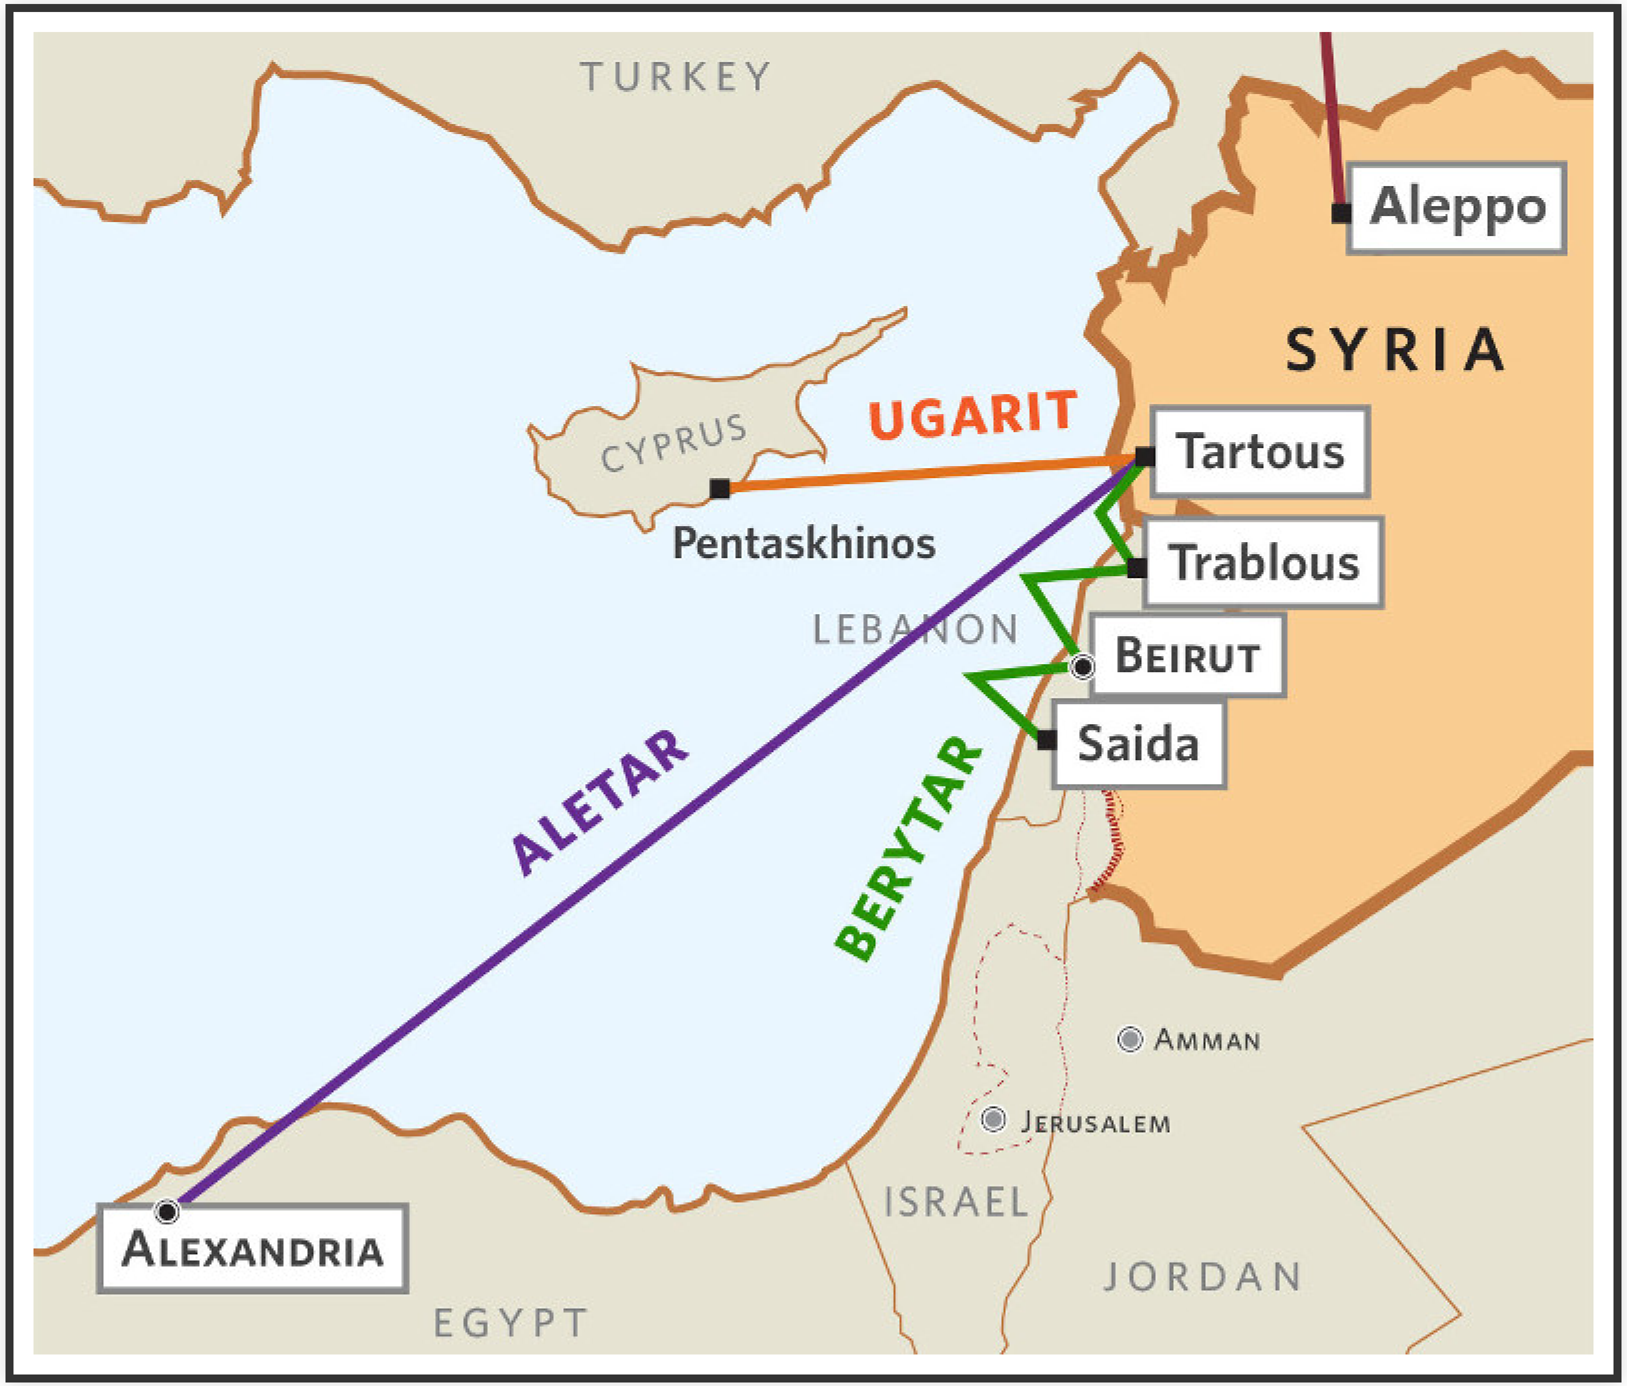

INTERNET CENSORSHIP

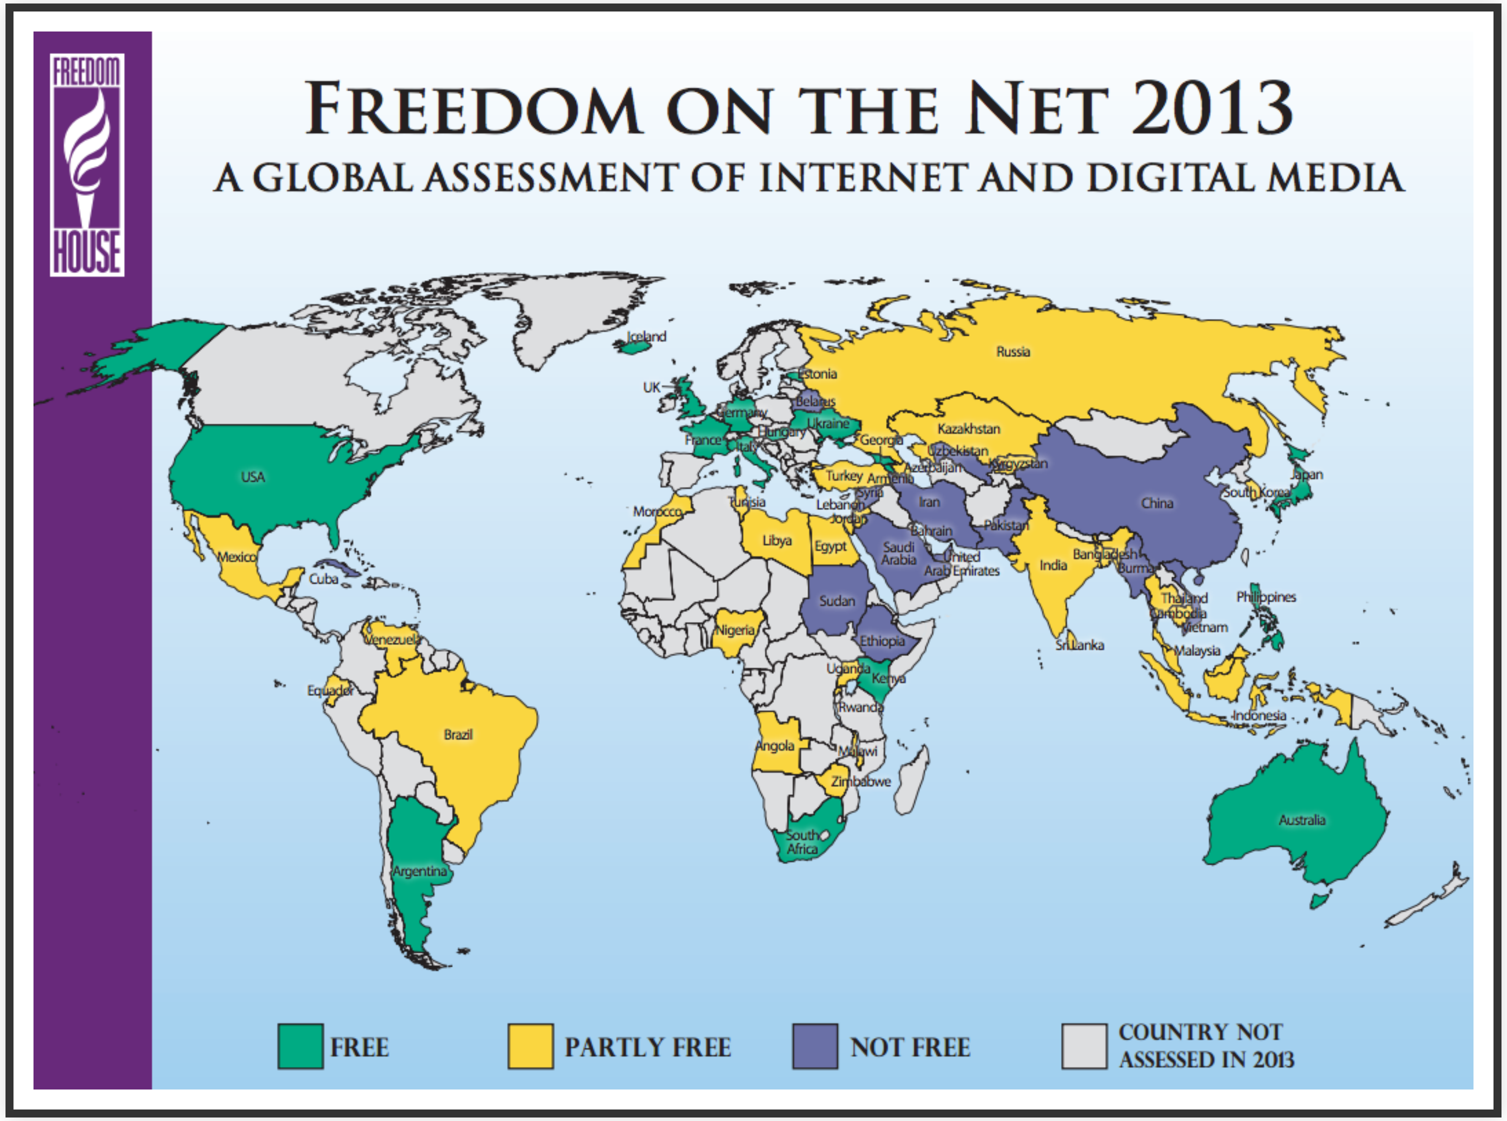

SOCIAL MEDIA

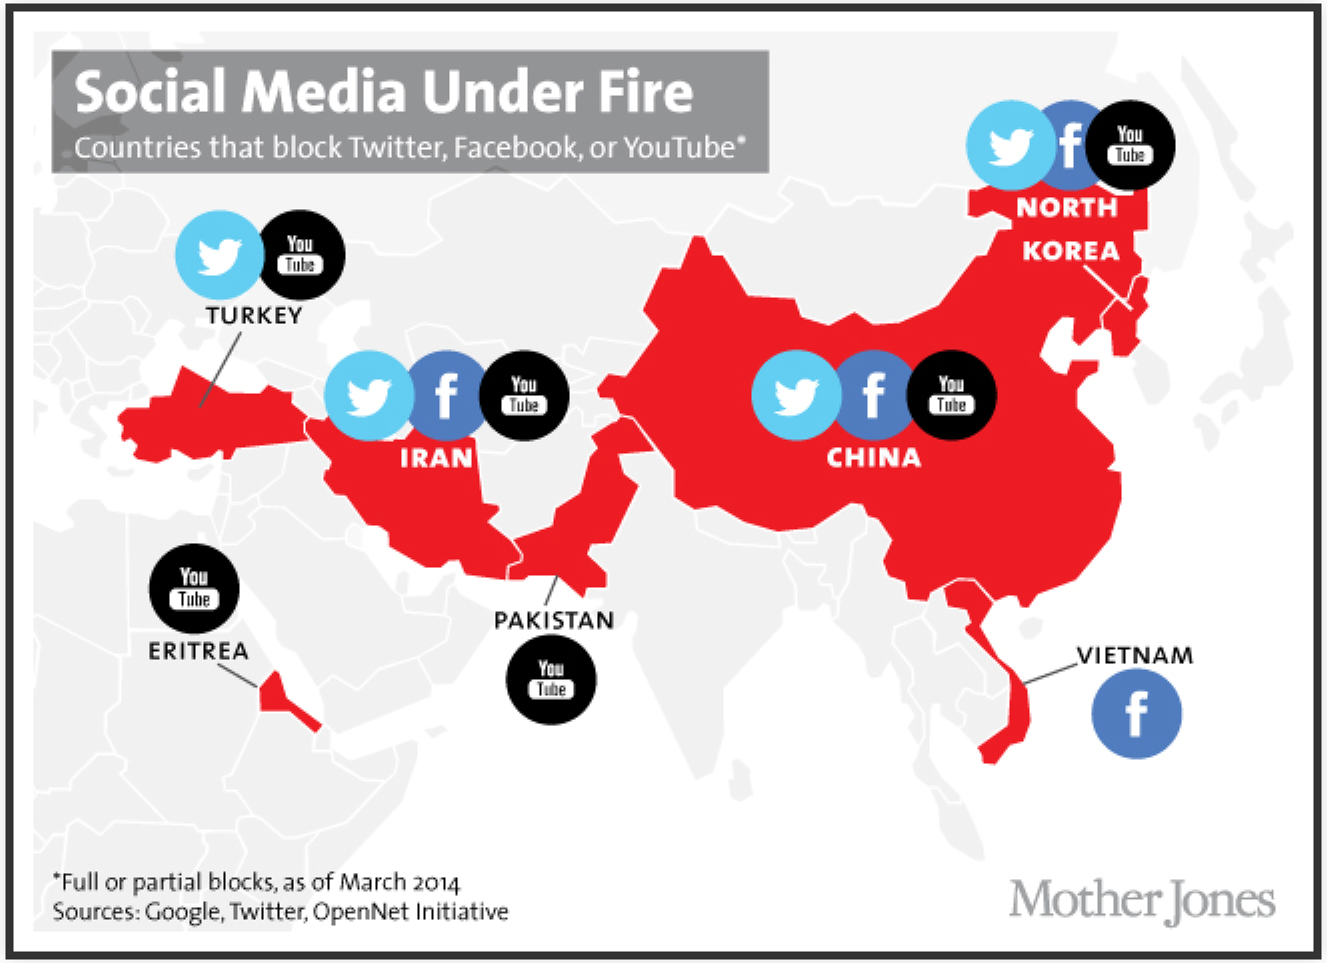

MALWARE

⚠ 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

DOS

⚠ Bad guys: attack server, network infrastructure: Denial of Service (DoS)

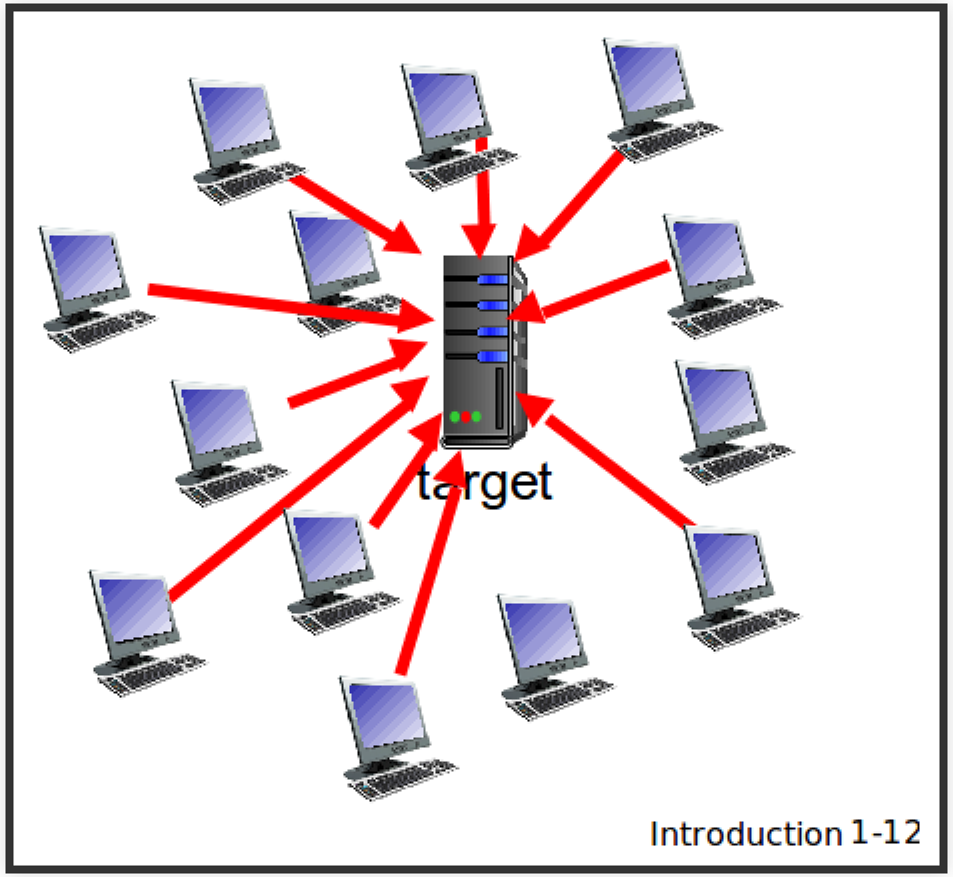

LISTEN ON THE LINE

⚠ 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

LISTEN ON THE LINE

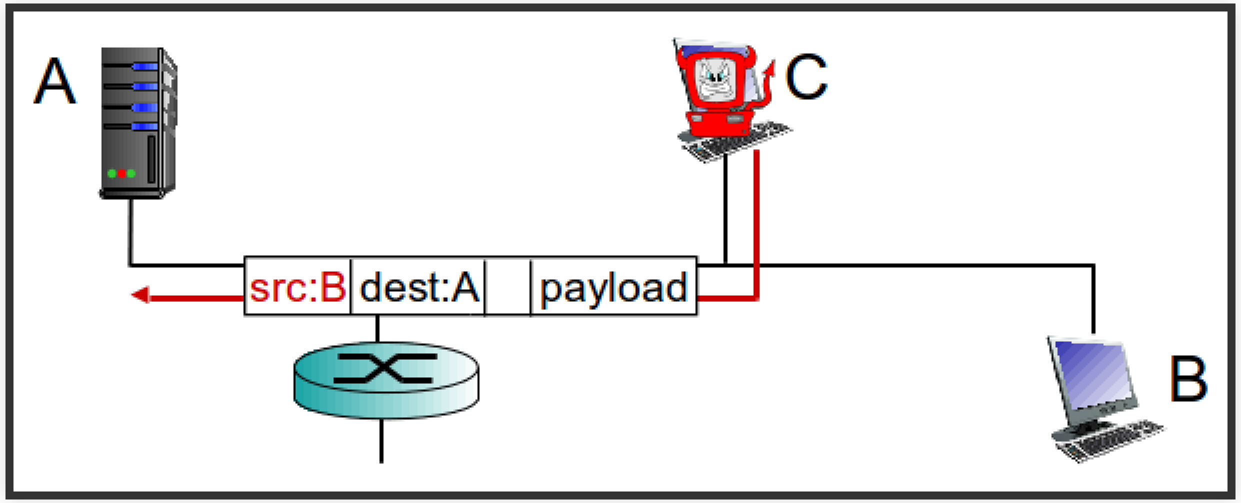

We will use a packagesniffing tool in this course, Wireshark, to inspect network trafic.

BAD GUYS CAN USE FAKE ADDRESSES

IP spoofing - send packet with false source address

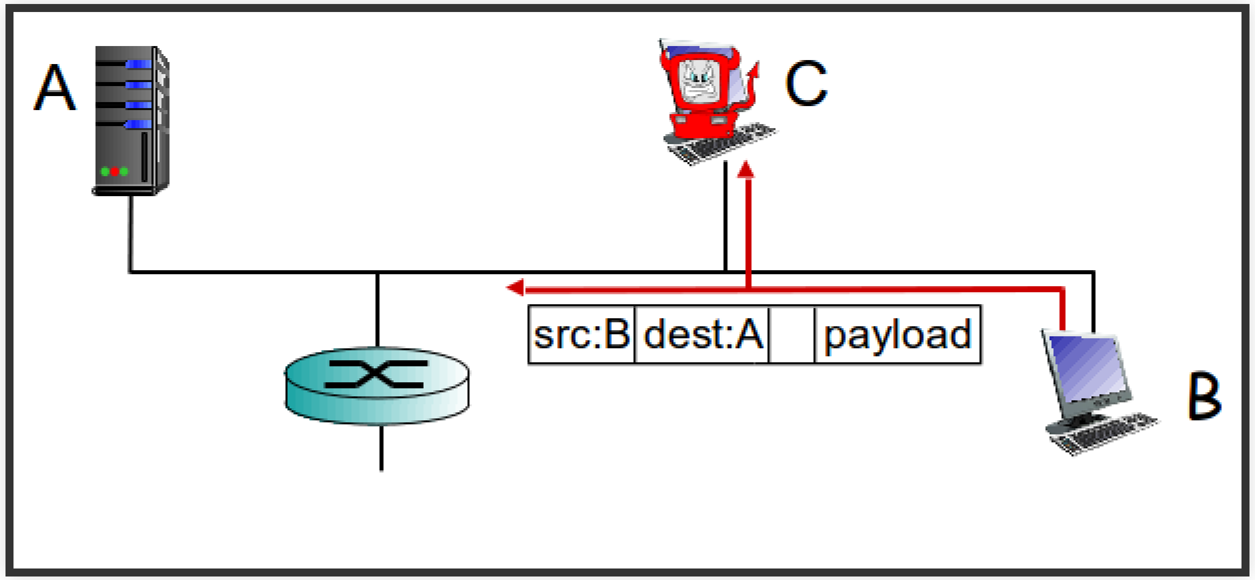

… lots more on security (later in the course)

SUMMARY 1

- History
- 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

REFERENCES

- Chapter 1 in textbook
- http://www.vox.com/a/internet-maps