Lecture 3: Internet Edge, Core, and Structure

COMP 332, Spring 2024

Victoria Manfredi

WESLEYAN
UNIVERSITY

Acknowledgements: materials adapted from Computer Networking: A Top Down Approach 7th edition: ©1996-2016, J.F Kurose and K.W. Ross, All Rights Reserved as well as from slides by Abraham Matta at Boston University and some material from Computer Networks by Tannenbaum and Wetherall.

Today

1. Announcements
   - Homework 1 posted, due Tuesday, Feb. 7 by 11:59p
   - Help sessions: Sun/Mon/Tues, 7-9p in Exley 638.
     - Also virtually: Mon/Tues/Wed
2. Internet organization
   - edge vs. core
   - Internetwork: network of networks
   - IP addresses
3. Network Applications
4. Network programming
   - TCP sockets

Internet Organization

CORE

vumanfredi@wesleyan.edu

How to move data through Internet core?

Internet core

- mesh of interconnected routers

Option 1: Packet-switching

- on-demand resource allocation
- best effort service
- good bandwidth use

Option 2: Circuit-switching

- reserved resources
- guaranteed service
- may waste bandwidth

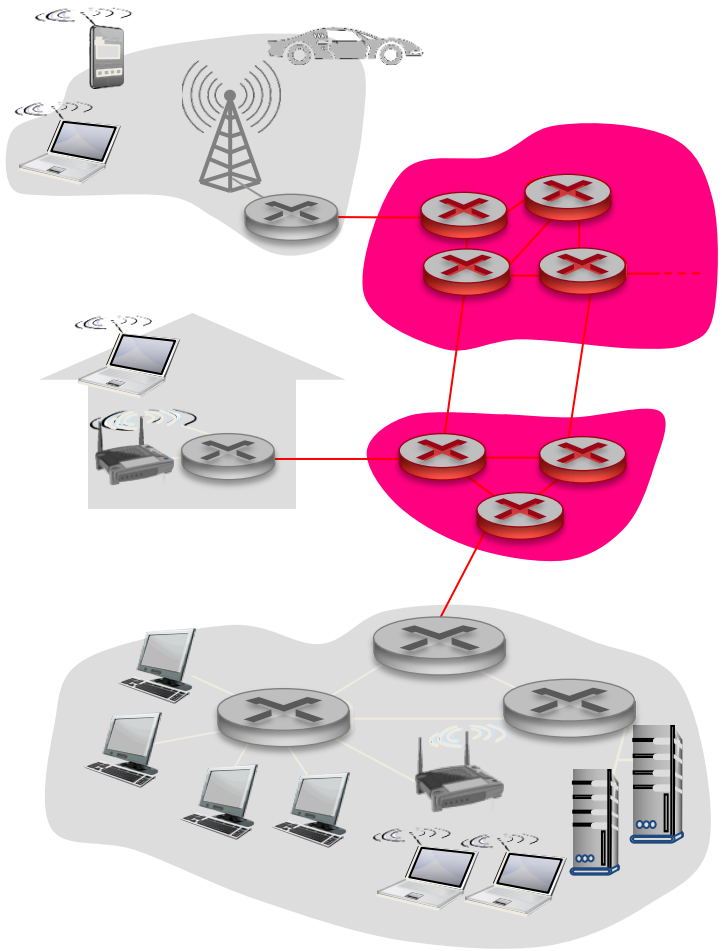

Q: human analogies of reserved versus on-demand allocation

Circuit switching

End-end resources allocated

- reserved for "call" between source & dest
- commonly used in traditional telephone networks

Dedicated resources

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

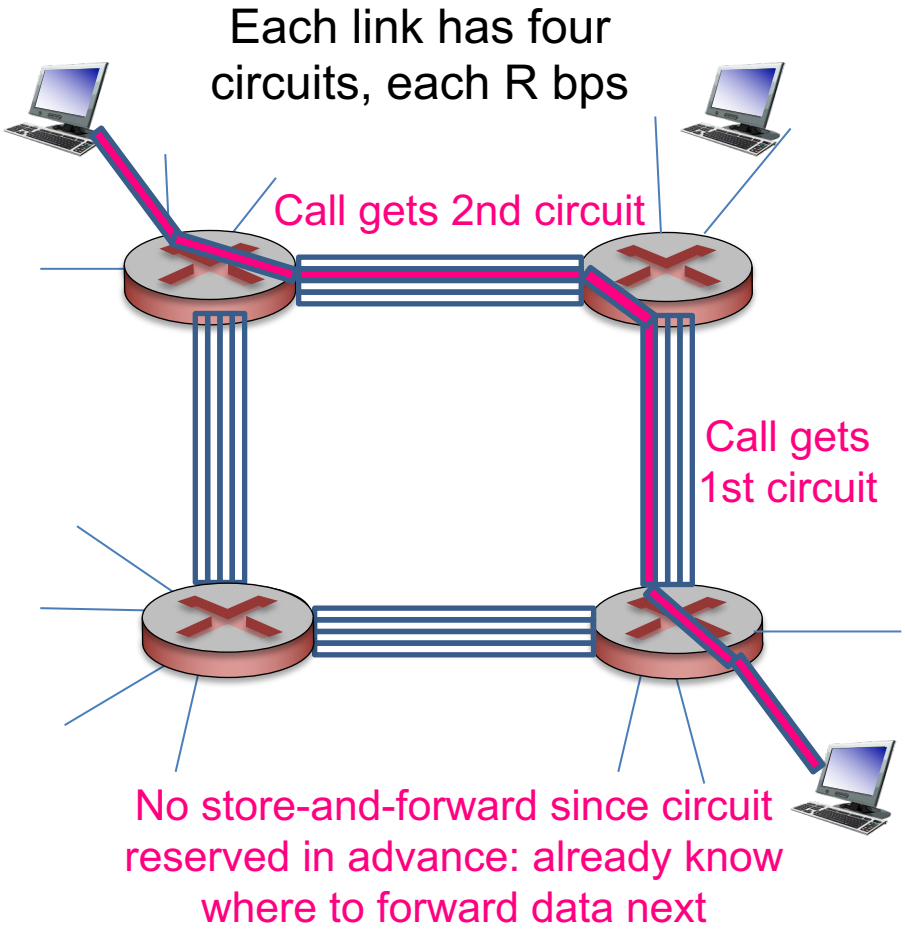

Q: what is end-end transmission delay for *L* bits? *L / R seconds*

Q: what happens if there is a lull in conversation?

Packet switching

1. Hosts break app-layer messages into packets

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

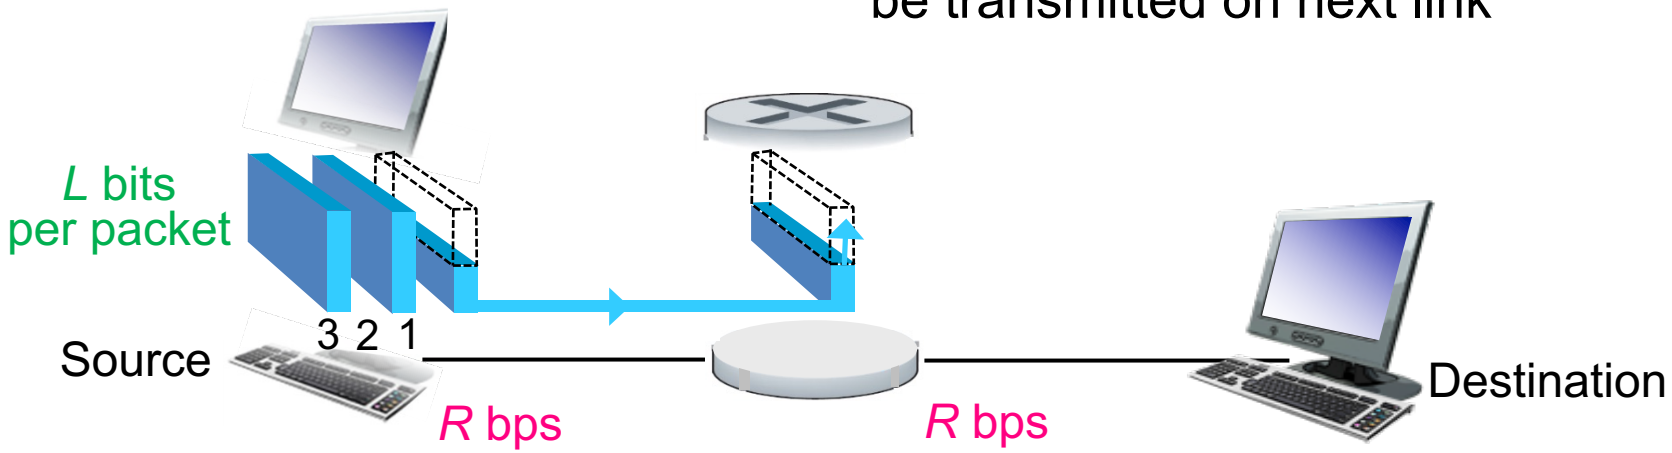

2. Time to transmit (push out) *L*-bit packet into *R* bps link: *L* / *R seconds*

4. Time to transmit (push out) *L*-bit packet into *R* bps link: *L* / *R seconds*

End-end transmission delay = *2 L* / *R seconds*

(assuming zero propagation, queuing, processing delay…)

Packet switching versus circuit switching

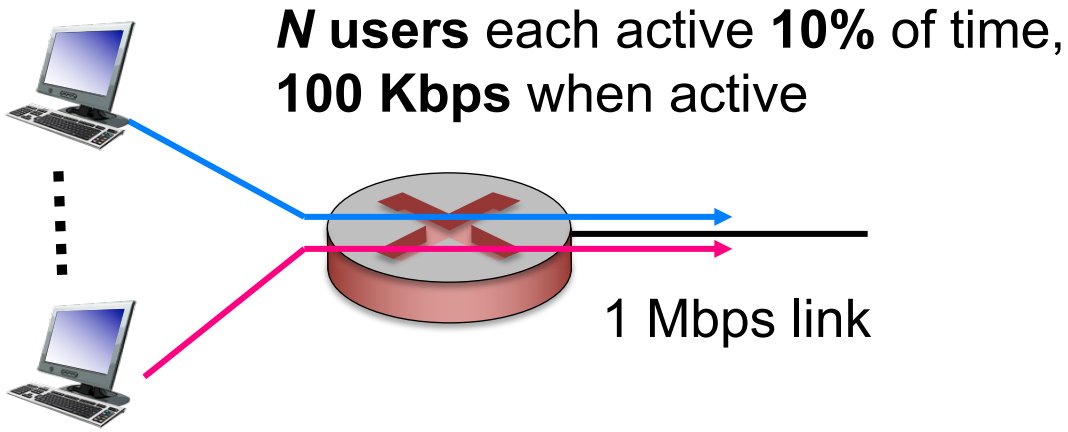

Q: How many users can be supported?

Circuit switching
- Link rate / User rate
- 1 Mbps / 100 Kbps = 10
- N = 10 users

Packet switching
- N = 35 users
- prob > 10 users active simultaneously < .0004

Q: how did we get value .0004?

Q: what happens if > 35 users?

Packet switching allows more users to use network!

Is packet switching always better?

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)

Binomial random variable (homework)

Suppose we do n independent experiments

- each experiment succeeds with probability p
- each experiment fails with probability 1-p

Independent experiment

- knowledge about one experiment occurring does not affect probability of other experiment occurring: e.g., coin toss

P(A and B) = P(A) x P(B)

P(A or B) = P(A) + P(B)

P(X=4 and X=5) = P(X=4) x P(X=5)

P(X=4 or X=5) = P(X=4) + P(X=5)

Binomial random variable (homework)

Suppose we do n independent experiments

- each experiment succeeds with probability p
- each experiment fails with probability 1-p

X = Random Variable indicating # of successes that occur in n experiments

Probability of i successes

$$P(X = i) = \binom{n}{i} p^i (1-p)^{n-i}$$

Probability of n-i failures

n choose i different ways to get i successful experiments: n! / (i! (n-i!))

Internet Organization

A NETWORK OF NETWORKS

vumanfredi@wesleyan

A network of networks

Hosts connect to Internet via access ISPs (e.g., Comcast)

– residential, company and university ISPs

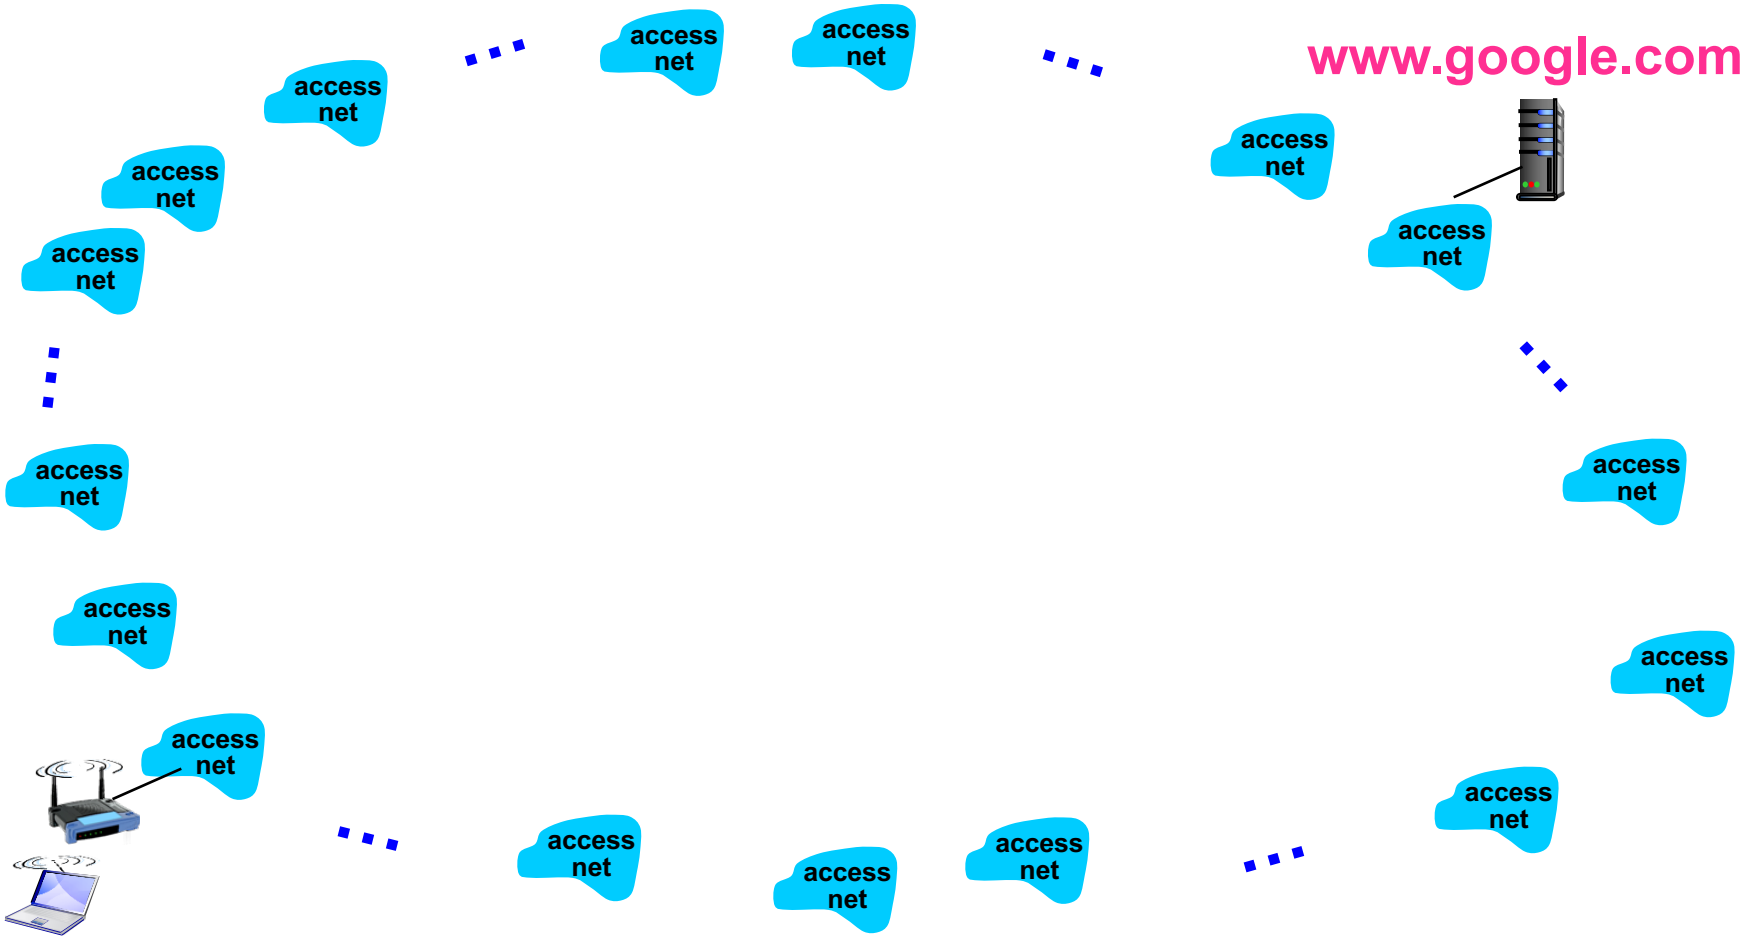

A network of networks

Q: given millions of access ISPs, how to connect together?

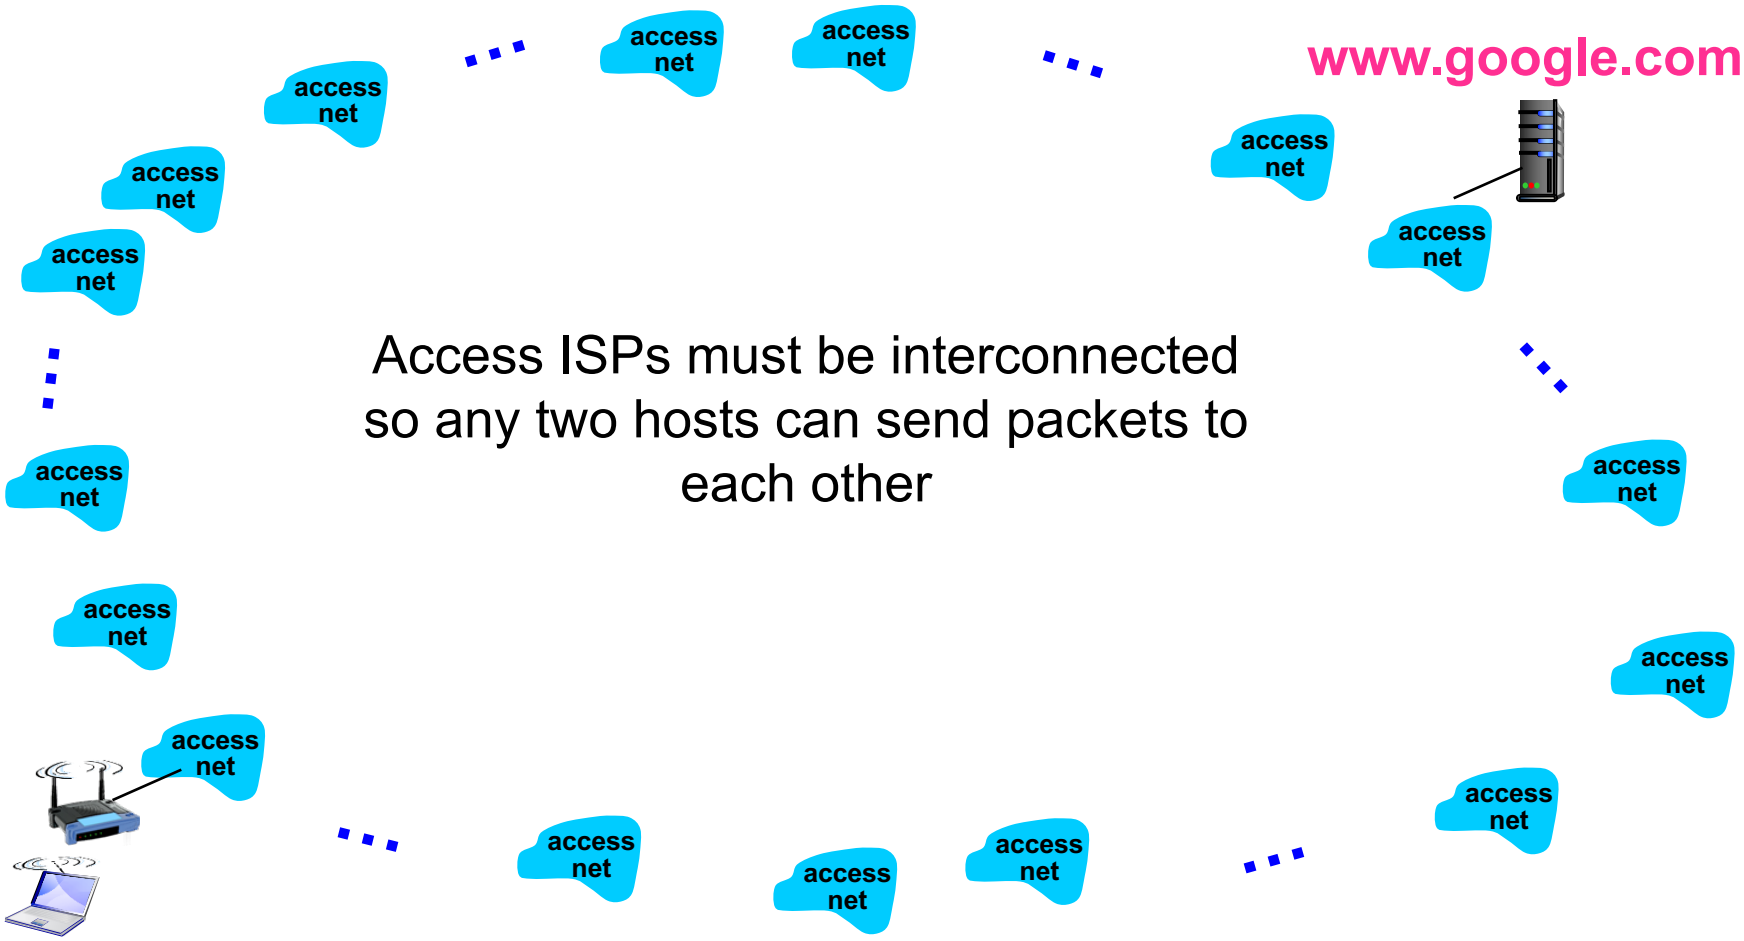

A network of networks

Option 1: connect each access ISP to every other access ISP?

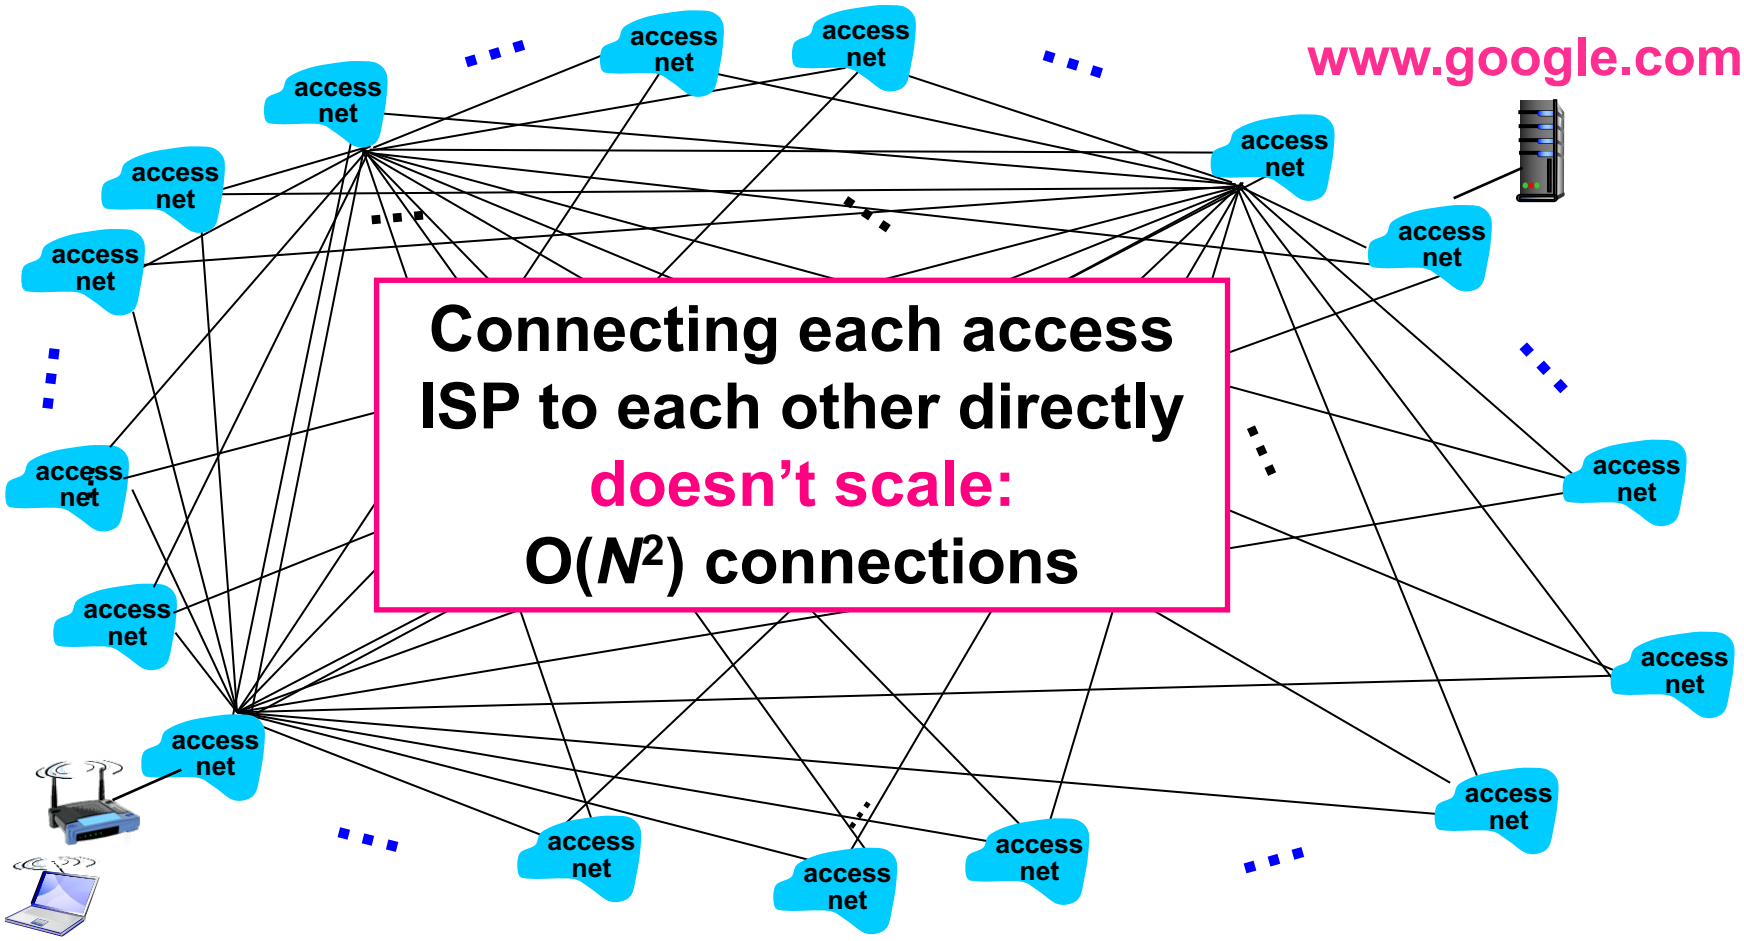

A network of networks

Option 2: connect each access ISP to one global transit ISP?
Customer and provider ISPs have economic agreement.

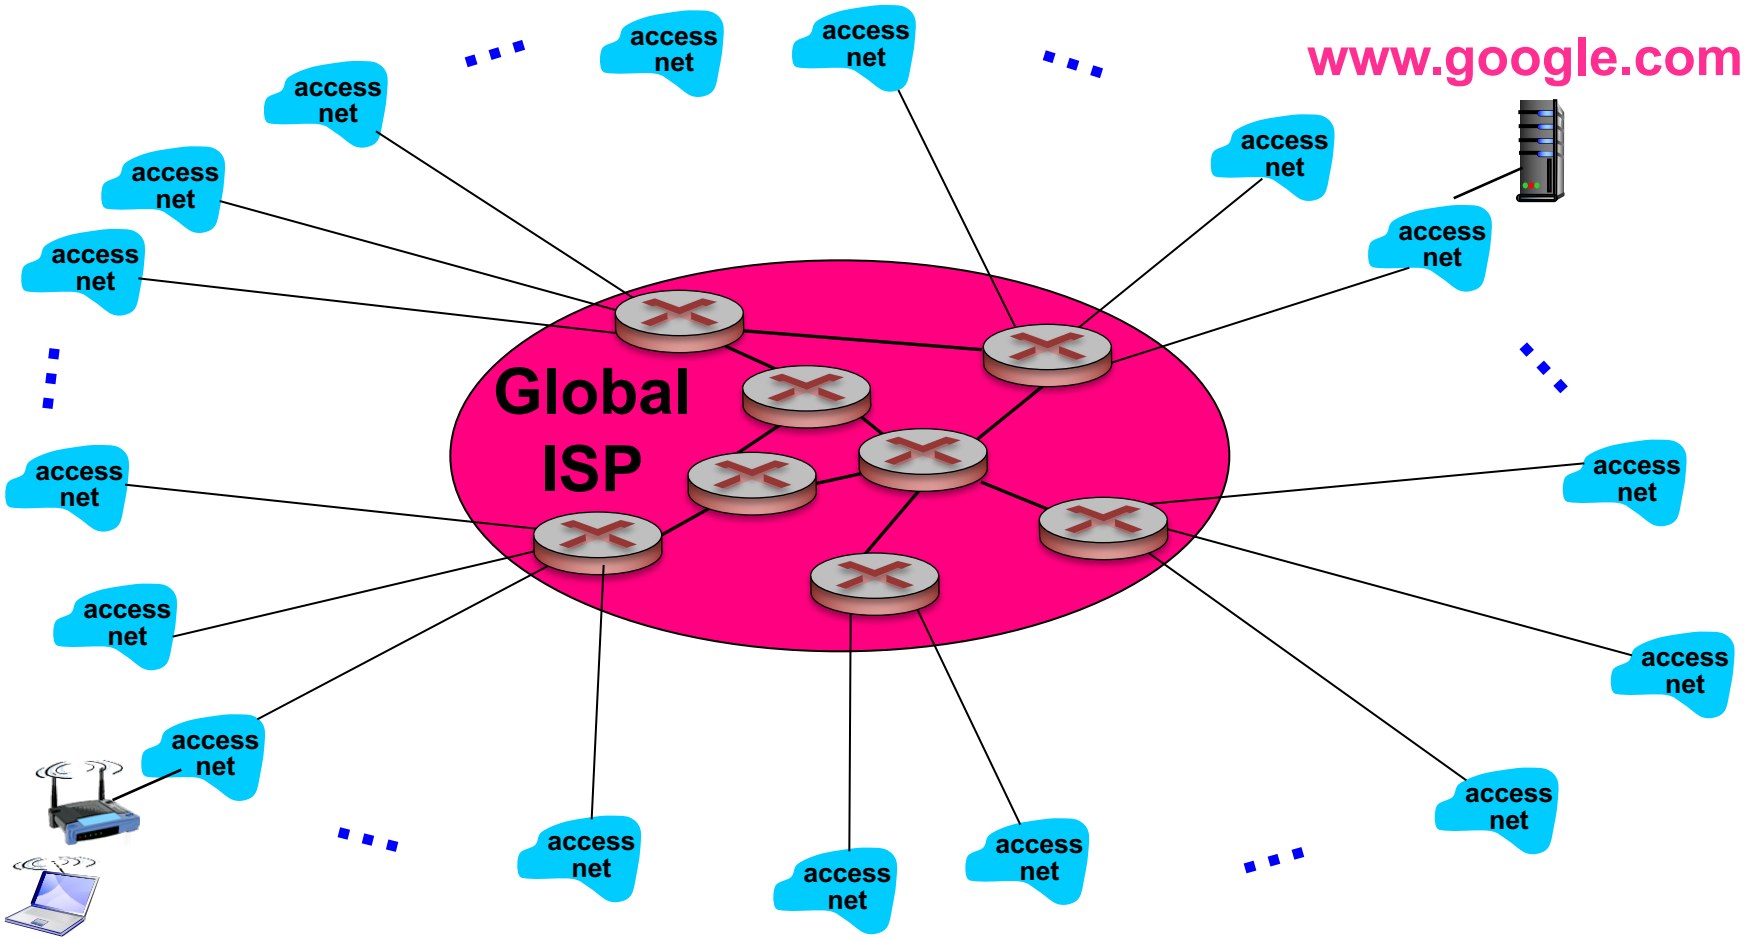

A network of networks

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

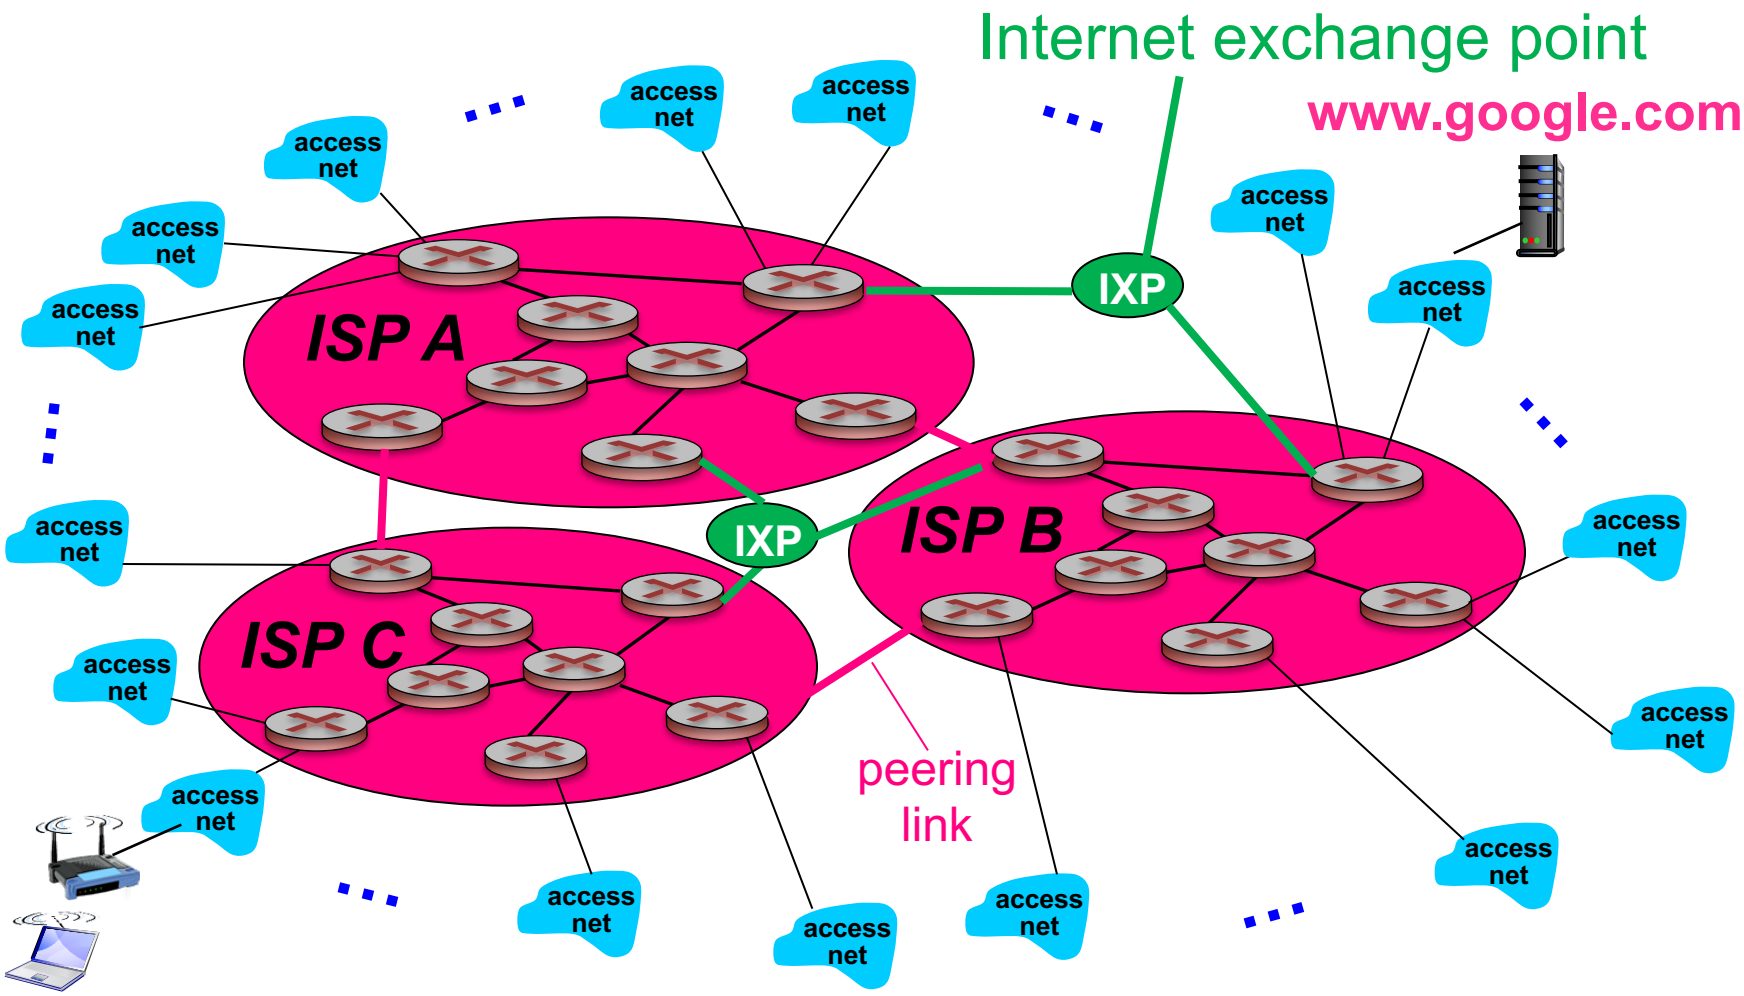

IXP: Internet Exchange Point

Room full of routers

– meeting point where multiple ISPs can peer together

Peer

– let Internet traffic cross/transit your computer network without fee

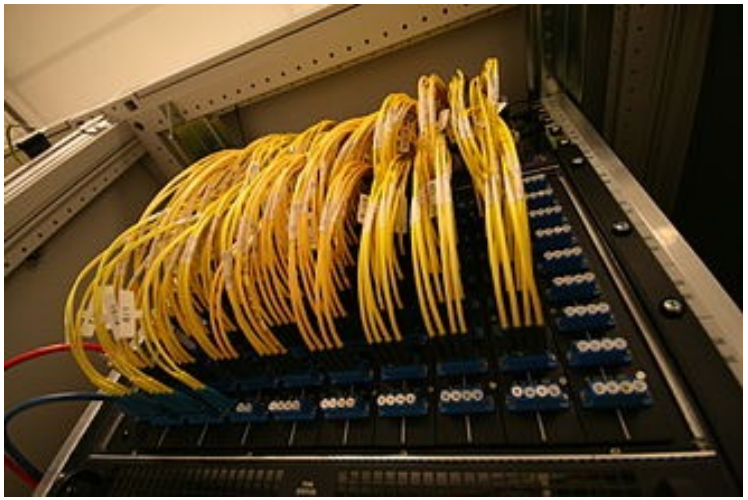

By Fabienne Serriere [CC BY-SA 3.0 (http://creativecommons.org/licenses/by-sa/3.0)], via Wikimedia Commons

Amsterdam IXP

- where 100s of ISPs connect
- optical fiber patch panel connecting different ISPs

A network of networks

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

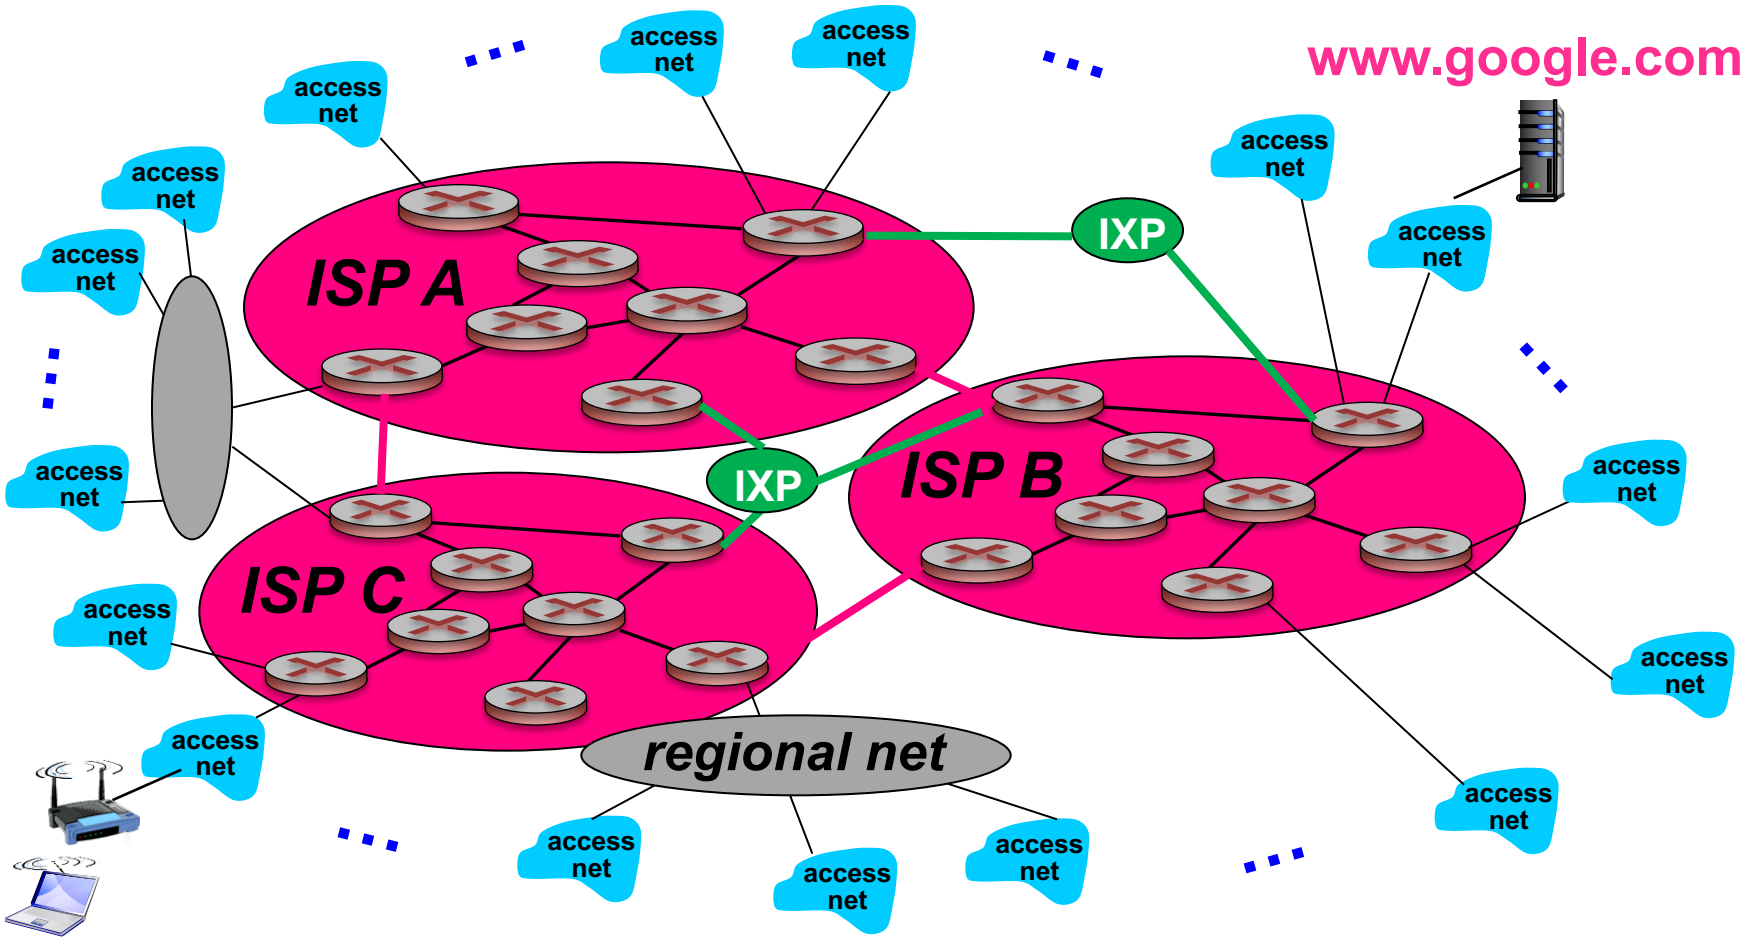

A network of networks

… and content provider networks may run their own network, to bring services, content close to end users

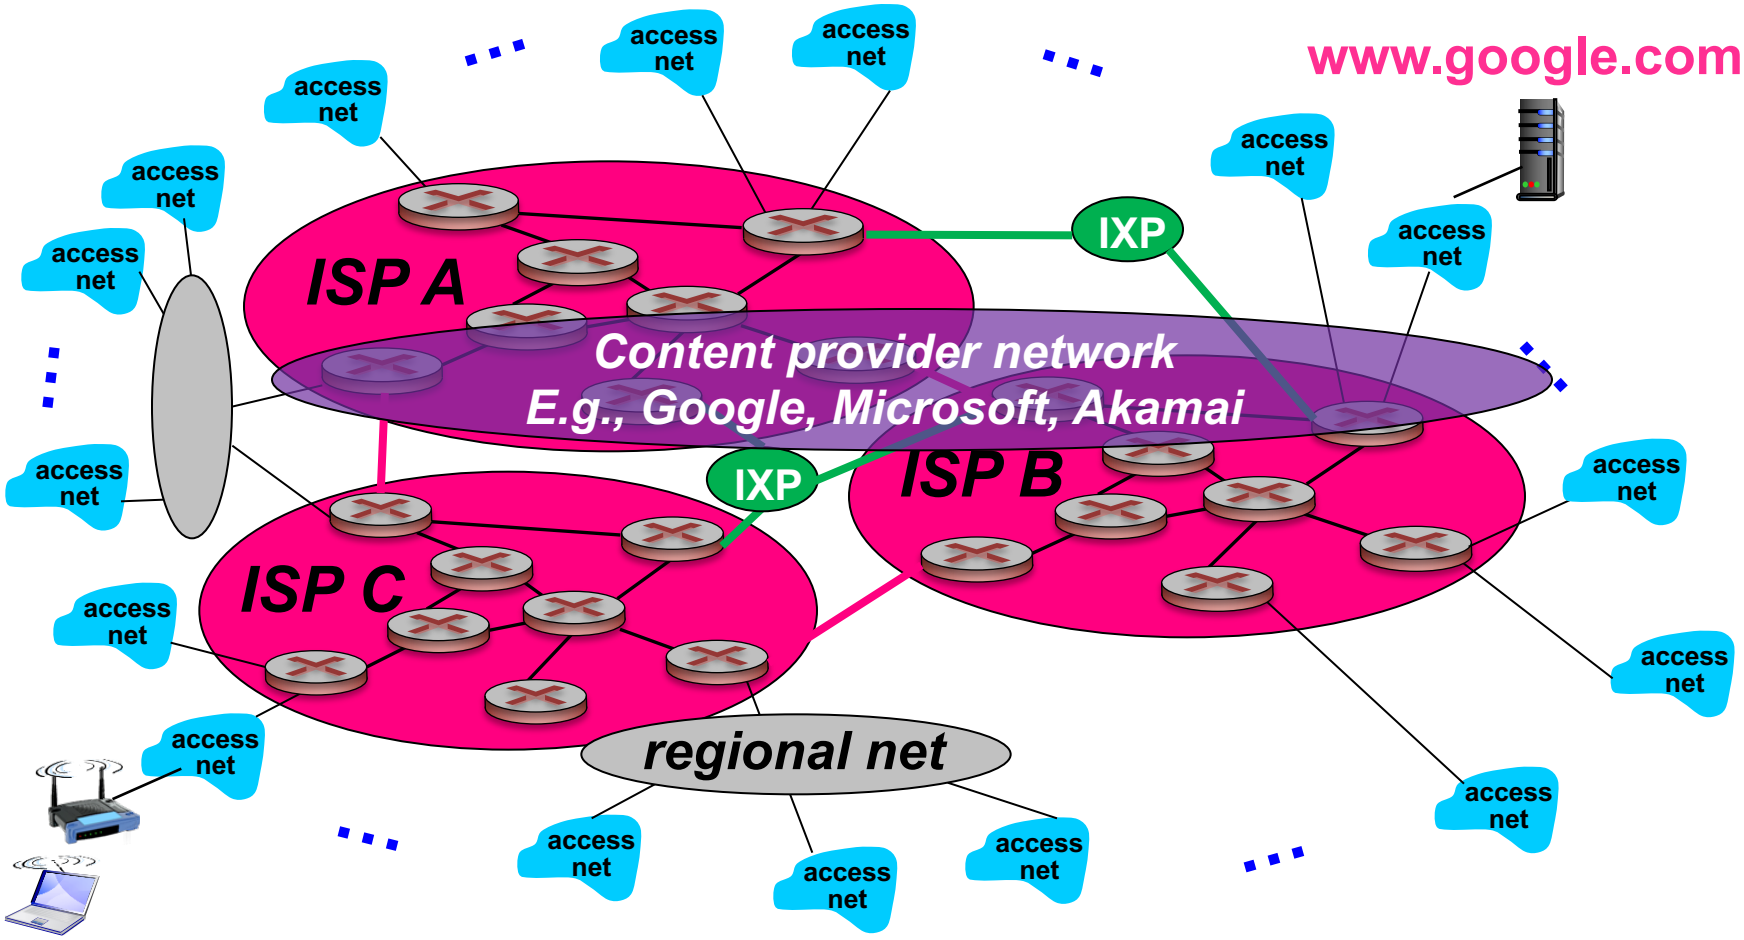

A network of networks

Resulting network of networks is very complex: evolution was driven by economics and national policies

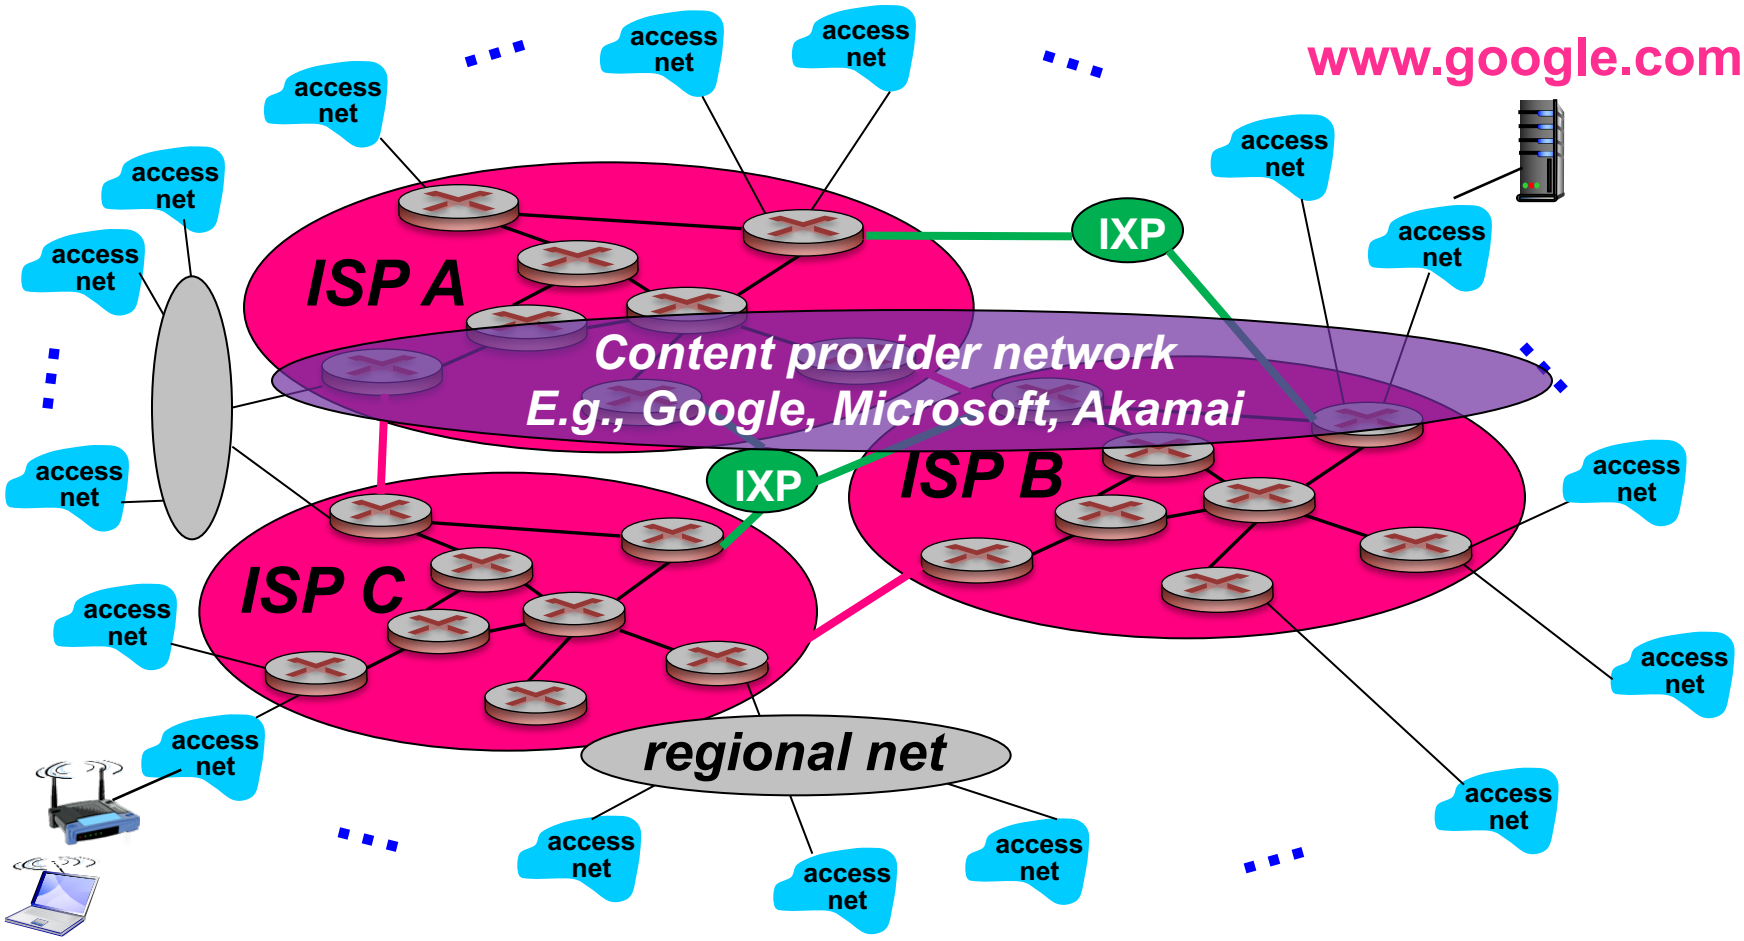

Different kinds of ISPs

Tier-1commercial ISPs: e.g., Level 3, Sprint, AT&T, NTT

- national & international coverage, peer with other tier 1 ISPs
- can reach all of Internet via peering only, peering typically payment free

Content provider network: e.g., Google (YouTube benefits)

- private network that connects its data centers to Internet
- often bypasses tier-1, regional ISPs, may buy transit
- by not paying provider ISP, save money and better control QoS for traffic

Regional ISP: e.g., Comcast

- customer ISP of Tier 1 ISP, provider ISP to access ISP
- peers with some networks but needs to purchase some IP transit to reach some parts of Internet

Access ISPs: connect end systems to Internet

- any of these could be access ISP
- company or Wesleyan may connect directly into Tier 1 or Regional ISP

Internet terminology

Internet transit

- service of letting Internet traffic cross or transit a computer network
- usually used to connect a smaller ISP to larger Internet

Multi-home

- customer ISP connects to 2 or more provider ISPs
- Why? For robustness. Can also multi-home your home network

PoP: Point-of-Presence

- where customer packets enter network
- e.g., you connecting to access ISP or ISP connecting to provider ISP

Tier-1 ISP has global reach

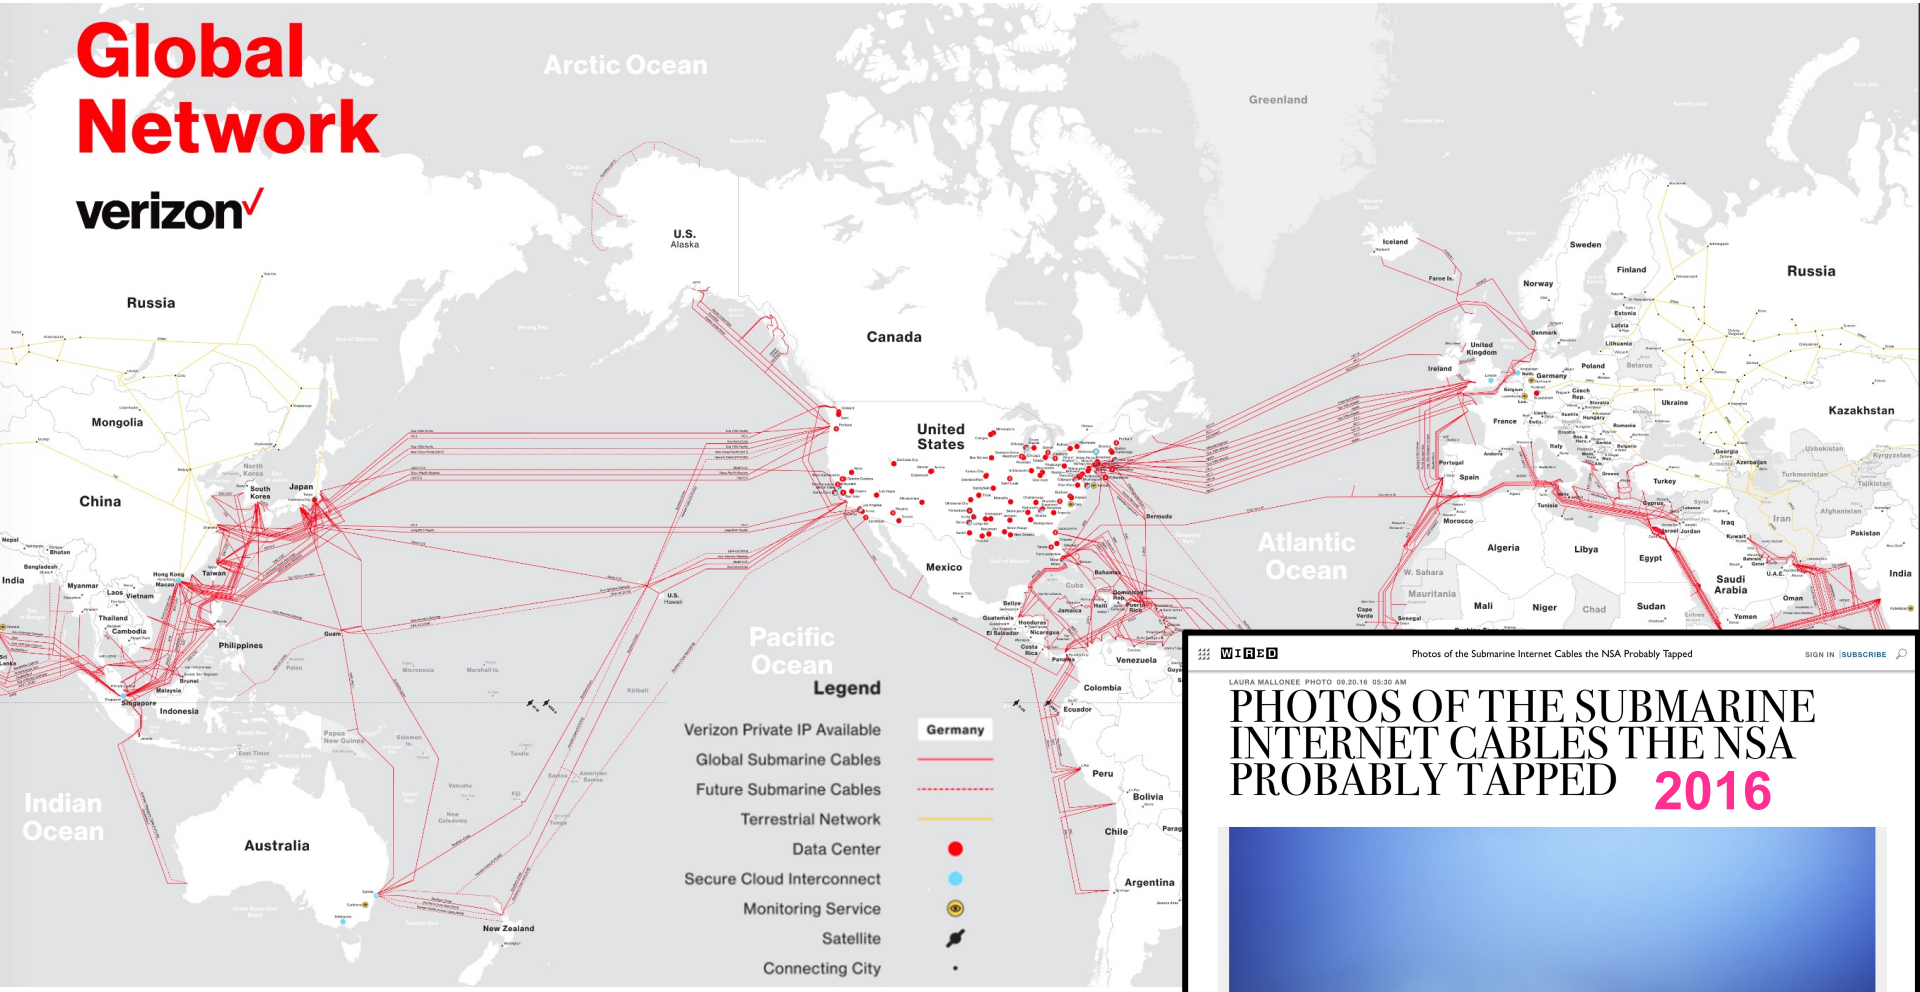

http://www.verizonenterprise.com/about/network/

Internet Organization

IP ADDRESSES

vumanfredi@wesleyan

Every device on Internet has an IP address

IPv4 addresses

- 4 bytes
  - space of addresses: 0-255 . 0-255 . 0-255 . 0-255
  - hostnames are human-readable, IP addresses are machine-readable
- Loopback address: send traffic to yourself
  - traffic sent here is “looped back” through network stack on machine on which sending process is running
  - 127 . * .* .*
  - typically 127.0.0.1, also called localhost
- Private subnet addresses
  - 10 .* .* .*
  - 172.16-31 .* .*
  - 192.168 .* .*

} Subnet: shared prefix portion of addr

IPv6 addresses

- 16 bytes: we’re running out of 4 byte addresses …

Who owns what address ranges?

Amazon

- 50.19.*.* → 256 x 256 = 65536 addresses
- 54.239.98.* → 256 addresses
- …

Google

- 64.233.160.0 to 64.233.191.255
- 66.102.0.0 to 66.102.15.255
- …

Facebook

- 57.240.0.0/17
- 157.240.10.0/24
- 157.240.1.0/24
- …

Wesleyan

- 129.133.21.*
- …

How are IP addresses assigned?

Your ISP or institution has block of IP addresses

- you are assigned one of those IP addresses
- (possible you will get NAT'd address …)

Static IP address

- manual configuration: set in network settings

Dynamic IP address

- using Dynamic Host Configuration Protocol (DHCP) in network-layer
- client (you) broadcasts request for IP address
- DHCP server on network assigns you address from address pool
  - typically get IP address for fixed period of time
  - router can be configured to act as DHCP server

Actually …

Many hosts have multiple IP addresses

How?

- IP address associated with network interface not host
- network interface card (NIC): connects computer to network

A host may have 1 or more network interfaces

- my laptop has (at least) 2 NICs: 1 wireless and 1 wired (via USB)
- router needs at least two interfaces
  - otherwise can't connect multiple networks together
- Cisco core router: can have up to 10,000 interfaces!
  - one interface per link: router has many IP addresses

VirtualBox Virtual Machine (VM)

- you can set the number and type of network interfaces for VM

What's my IP address?

ifconfig

- what network interfaces does my machine have?
- what are my IP and MAC # addresses?
- configure/enable/disable an interface

Linux

Ethernet 0

IPv4 address

IPv6 address

Loopback address

```
vmanfred@curveball-VirtualBox:~$ ifconfig
eth0      Link encap:Ethernet  HWaddr 08:00:27:e2:65:b0
          inet addr:129.133.178.53  Bcast:129.133.191.255  Mask:255.255.240.0
          inet6 addr: fe80::a00:27ff:fee2:65b0/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:102302 errors:0 dropped:0 overruns:0 frame:0
          TX packets:29698 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:141037591 (141.0 MB)  TX bytes:2394226 (2.3 MB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:1912 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1912 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:146886 (146.8 KB)  TX bytes:146886 (146.8 KB)
```

What's host's IP address?

Host

```
> host google.com
google.com has address 142.250.65.238
google.com has IPv6 address 2607:f8b0:4006:81e::200e
google.com mail is handled by 10 smtp.google.com.
```

What's host name for IP address?

```
> host 8.8.8.8
8.8.8.8.in-addr.arpa domain name pointer dns.google.
```

What's host's IP address?

dig

```
> dig google.com

; <<>> DiG 9.10.6 <<>> google.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 5897
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 512
;; QUESTION SECTION:
;google.com.                    IN      A

;; ANSWER SECTION:
google.com.             195     IN      A       142.250.72.110

;; Query time: 17 msec
;; SERVER: 2001:558:feed::1#53(2001:558:feed::1)
;; WHEN: Sun Feb 04 10:03:44 EST 2024
;; MSG SIZE  rcvd: 55
```

DNS resolver used

Is host up?

Ping

- sends ICMP echo request to host
- host sends ICMP echo reply back
- If no reply within timeout period, packet deemed lost

```
> ping stanford.edu
PING stanford.edu (171.67.215.200): 56 data bytes
64 bytes from 171.67.215.200: icmp_seq=0 ttl=237 time=94.951 ms
64 bytes from 171.67.215.200: icmp_seq=1 ttl=237 time=94.738 ms
64 bytes from 171.67.215.200: icmp_seq=2 ttl=237 time=95.525 ms
64 bytes from 171.67.215.200: icmp_seq=3 ttl=237 time=194.993 ms
64 bytes from 171.67.215.200: icmp_seq=4 ttl=237 time=97.139 ms
64 bytes from 171.67.215.200: icmp_seq=5 ttl=237 time=95.878 ms
64 bytes from 171.67.215.200: icmp_seq=6 ttl=237 time=95.667 ms
^C
--- stanford.edu ping statistics ---
7 packets transmitted, 7 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 94.738/109.842/194.993/34.770 ms
```

Is one IP address per machine enough?

What happens if you run multiple network applications?

- many processes running on computer
  - process is program in execution

How do messages received by computer get to right process?

- messages are addressed to (IP address, port number) pair
- different processes on computer will connect to network using same IP address but different port numbers

2 key functions of Internet core

How does Internet router determine outgoing link for packet?

1. Routing

- view Internet as giant graph
- run shortest path algorithms

2. Forwarding

- use paths to choose best output link for packet destination IP address
- if one link fails, chooses another

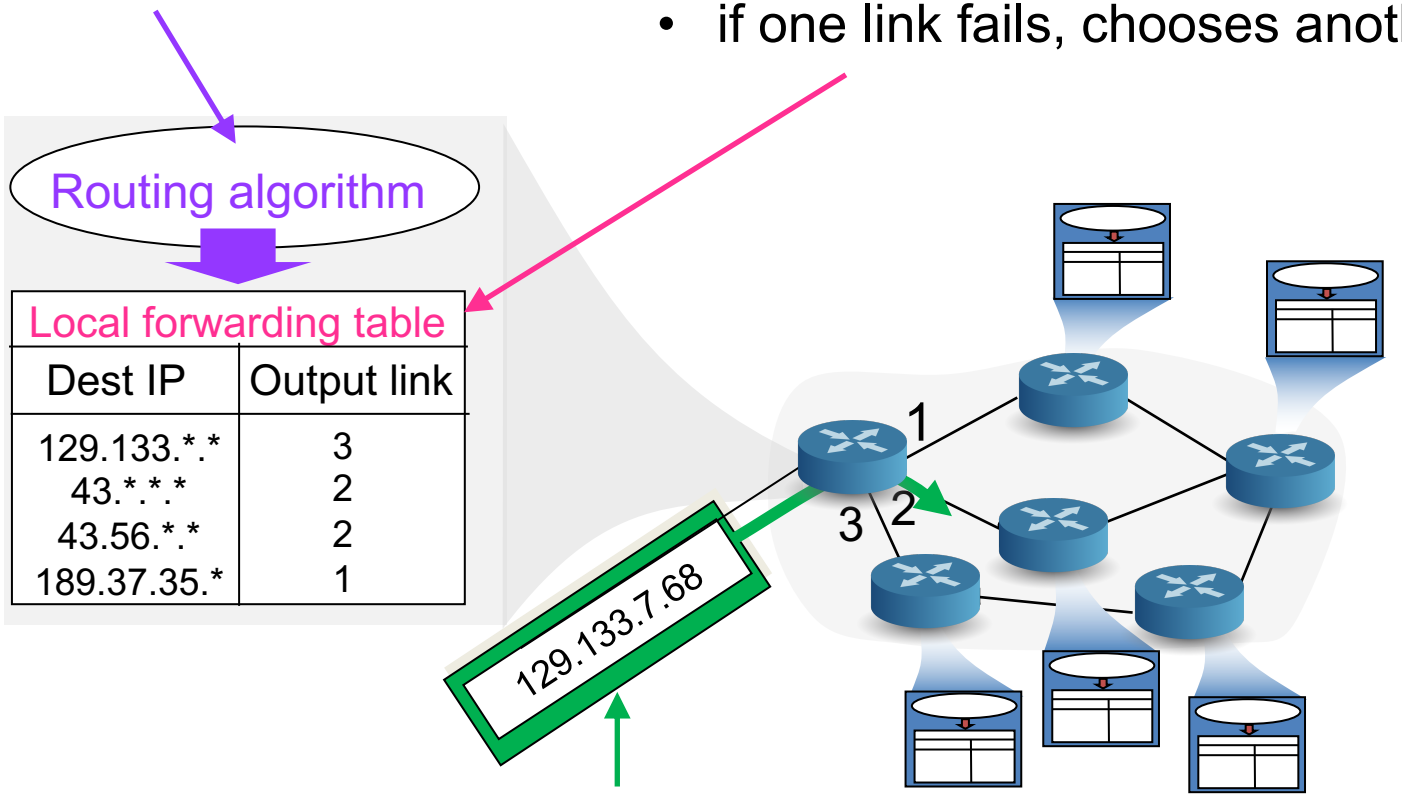

Dest IP addr in header of arriving packet (sequence of bits)

Routing of packets across Internet

Each router uses its forwarding table to choose outbound link based on packet's destination

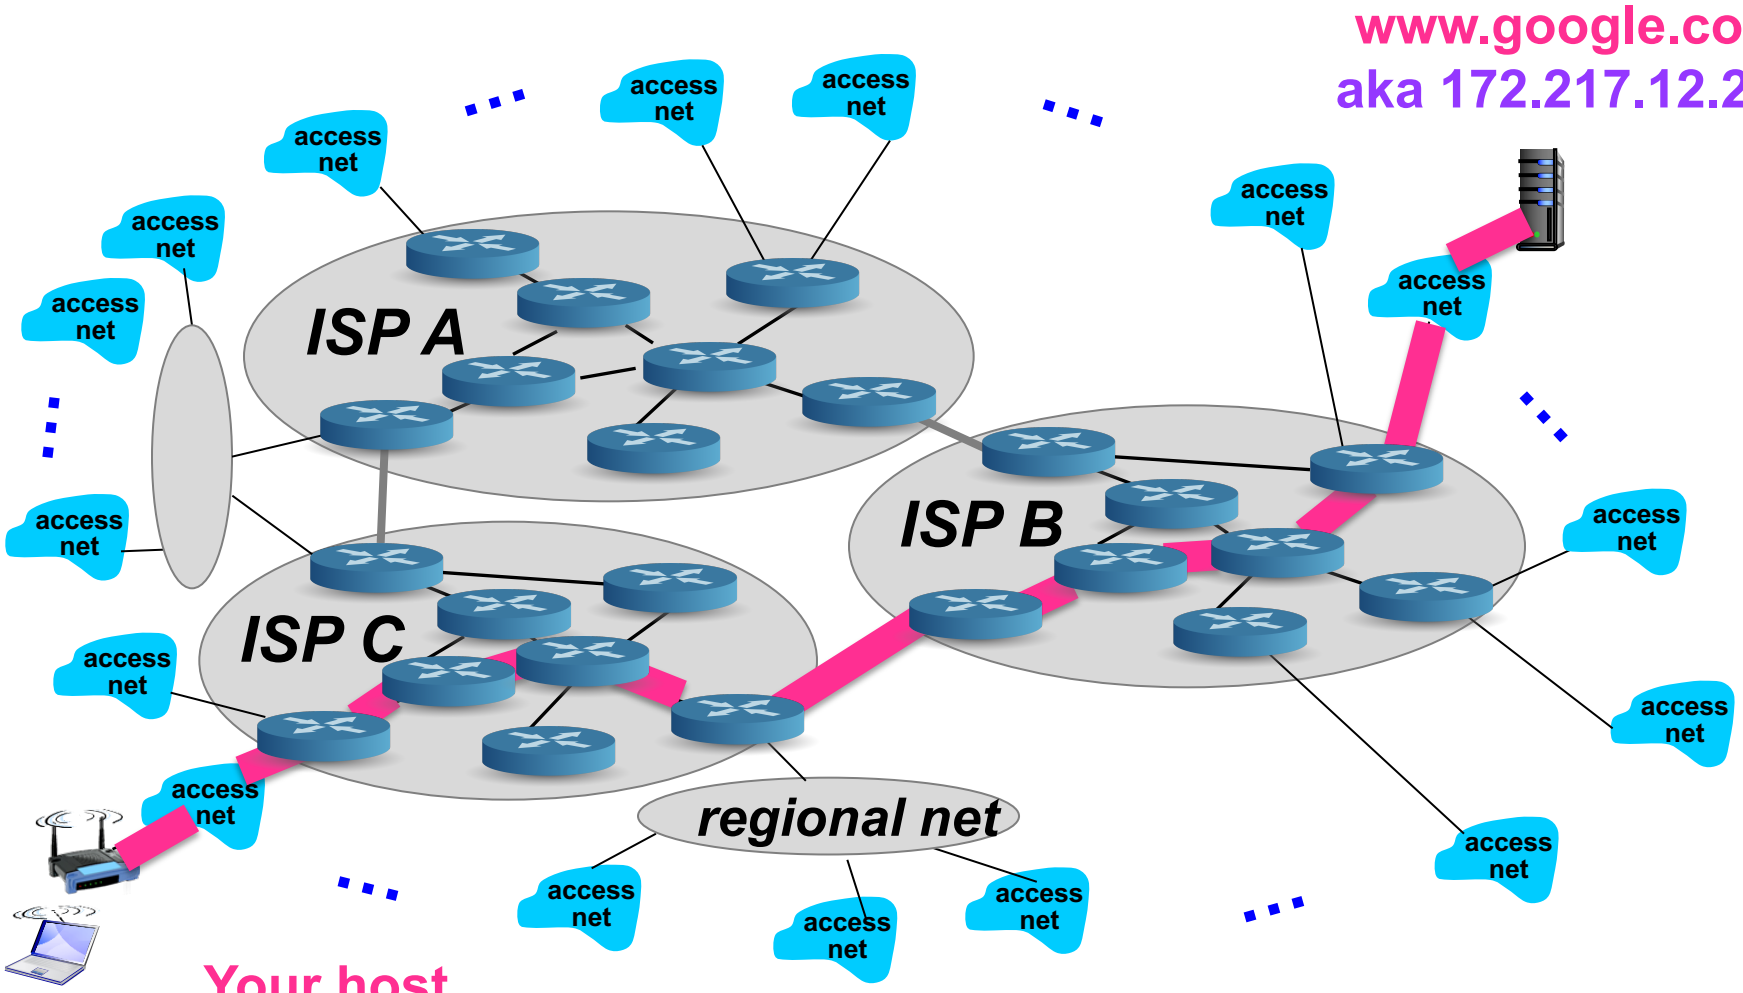

Network Applications

OVERVIEW

vumanfredi@wesleyan.edu

Creating a network app

Write programs that

- run on (different) end systems
- communicate over network
- e.g., web server software communicates with browser software

Q: Do we need to write software for network-core devices?

- No, network-core devices do not run user applications
- applications on end systems allows for rapid app development, propagation

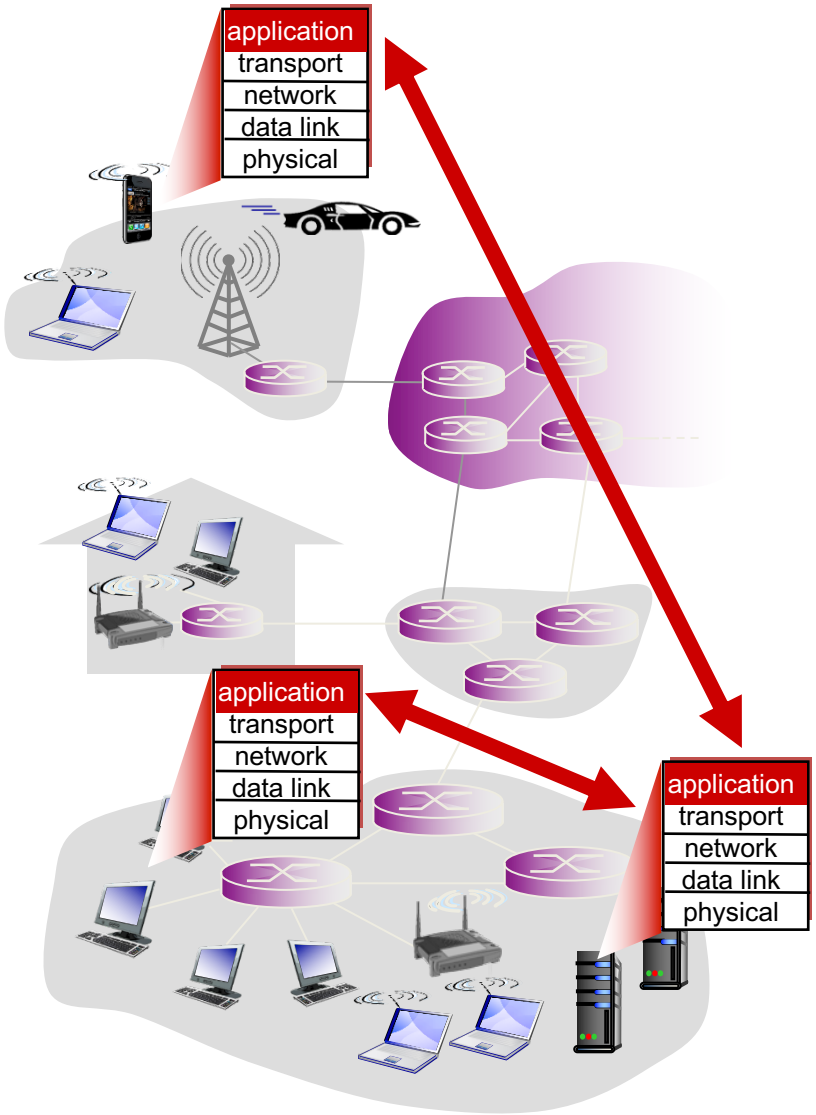

Client-server architecture

Client host requests and receives service from always on server host

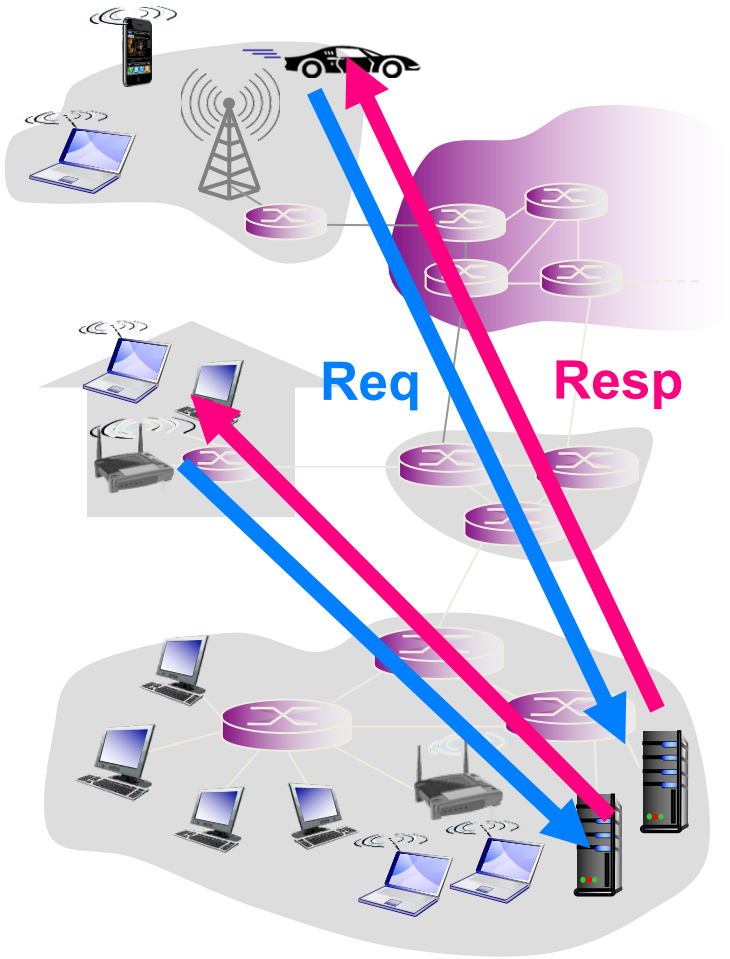

Server

- always-on, dedicated host
  - e.g., web server
- permanent IP address
- data centers for scaling

Clients

- communicate with server
- may be intermittently connected
- may have dynamic IP addresses
- do not communicate directly with other clients

Client and server devices are not equivalent

Peer-to-peer (P2P) architecture

Peers request service from other peers, provide service in return to other peers

End systems directly communicate

- self scalability – new peers bring new service capacity, as well as new service demands
- minimal/no use of always-on server
- E.g., Skype, BitTorrent

Complex management

- peers are intermittently connected and change IP addresses
- Q: why is this complex?

All devices are equivalent: a client can also be a server

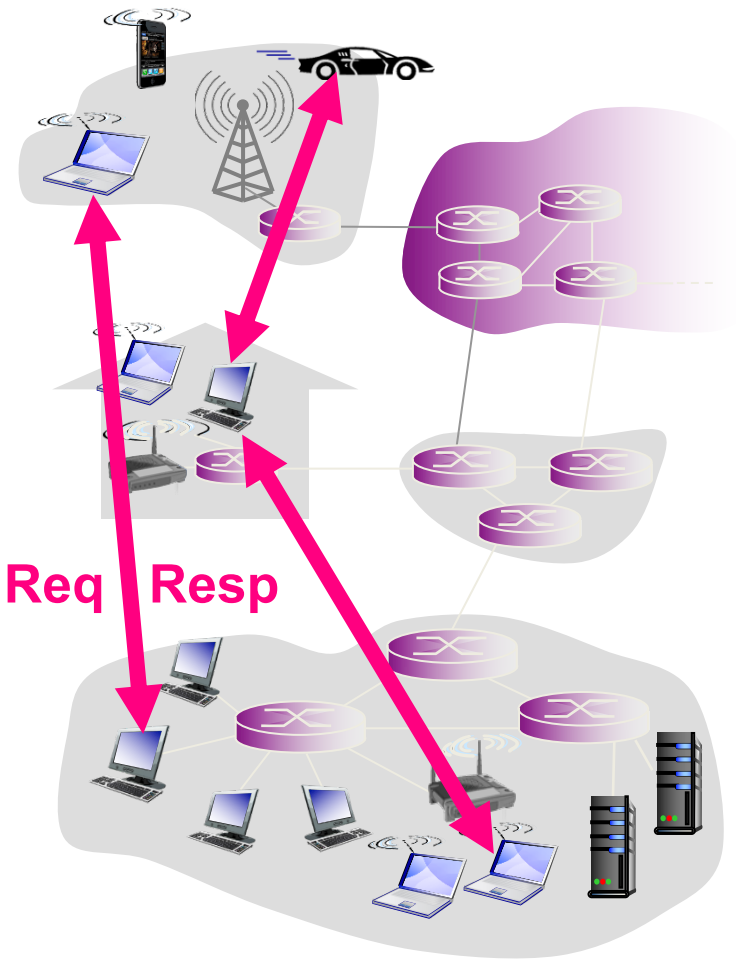

Processes communicating

Process

- program in execution, running within a host

Processes within same host

- communicate by using inter-process communication (defined by OS)

Processes on different hosts

- communicate by exchanging messages

Clients, servers

- client process
  - process that initiates communication
- server process
  - process that waits to be contacted

Aside

- applications with P2P architectures also have client & server processes

Our goal is to learn how to build client/server applications that use sockets to communicate