Lecture 3: Network Traffic and Performance Metrics

COMP 332, Spring 2018
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 due Wed.
   - no office hours today (CS candidate talk)
     - I have office hours tomorrow at 4p + I will be at help session tomorrow

2. Internet organization
   - Internetwork: network of networks
   - IP addresses

3. Network performance metrics
   - delay
   - loss
   - throughput

4. Network tools
   - Traceroute: measuring network delays and routes

Internet

A NETWORK OF NETWORKS

Internet structure: network of networks

End systems connect to Internet via access ISPs

- residential, company and university ISPs (Internet Service Providers)

Access ISPs in turn must be interconnected

- so any two hosts can send packets to each other

Resulting network of networks is very complex

- evolution was driven by economics and national policies

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

Internet structure: network of networks

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

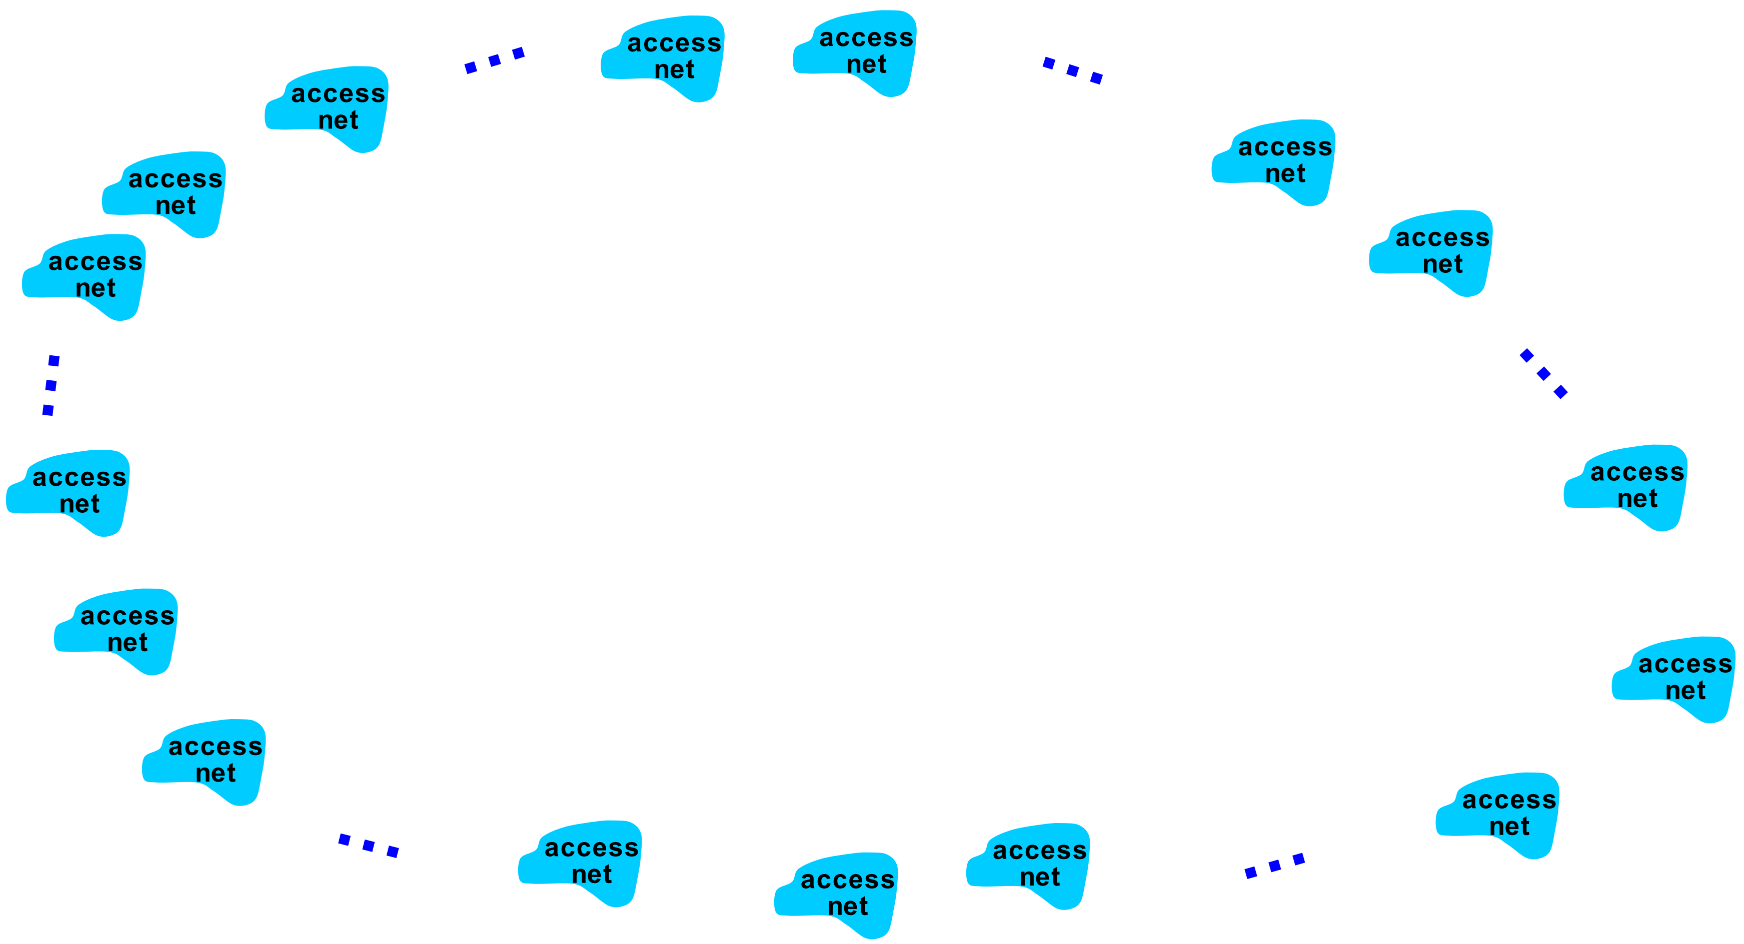

Internet structure: network of networks

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

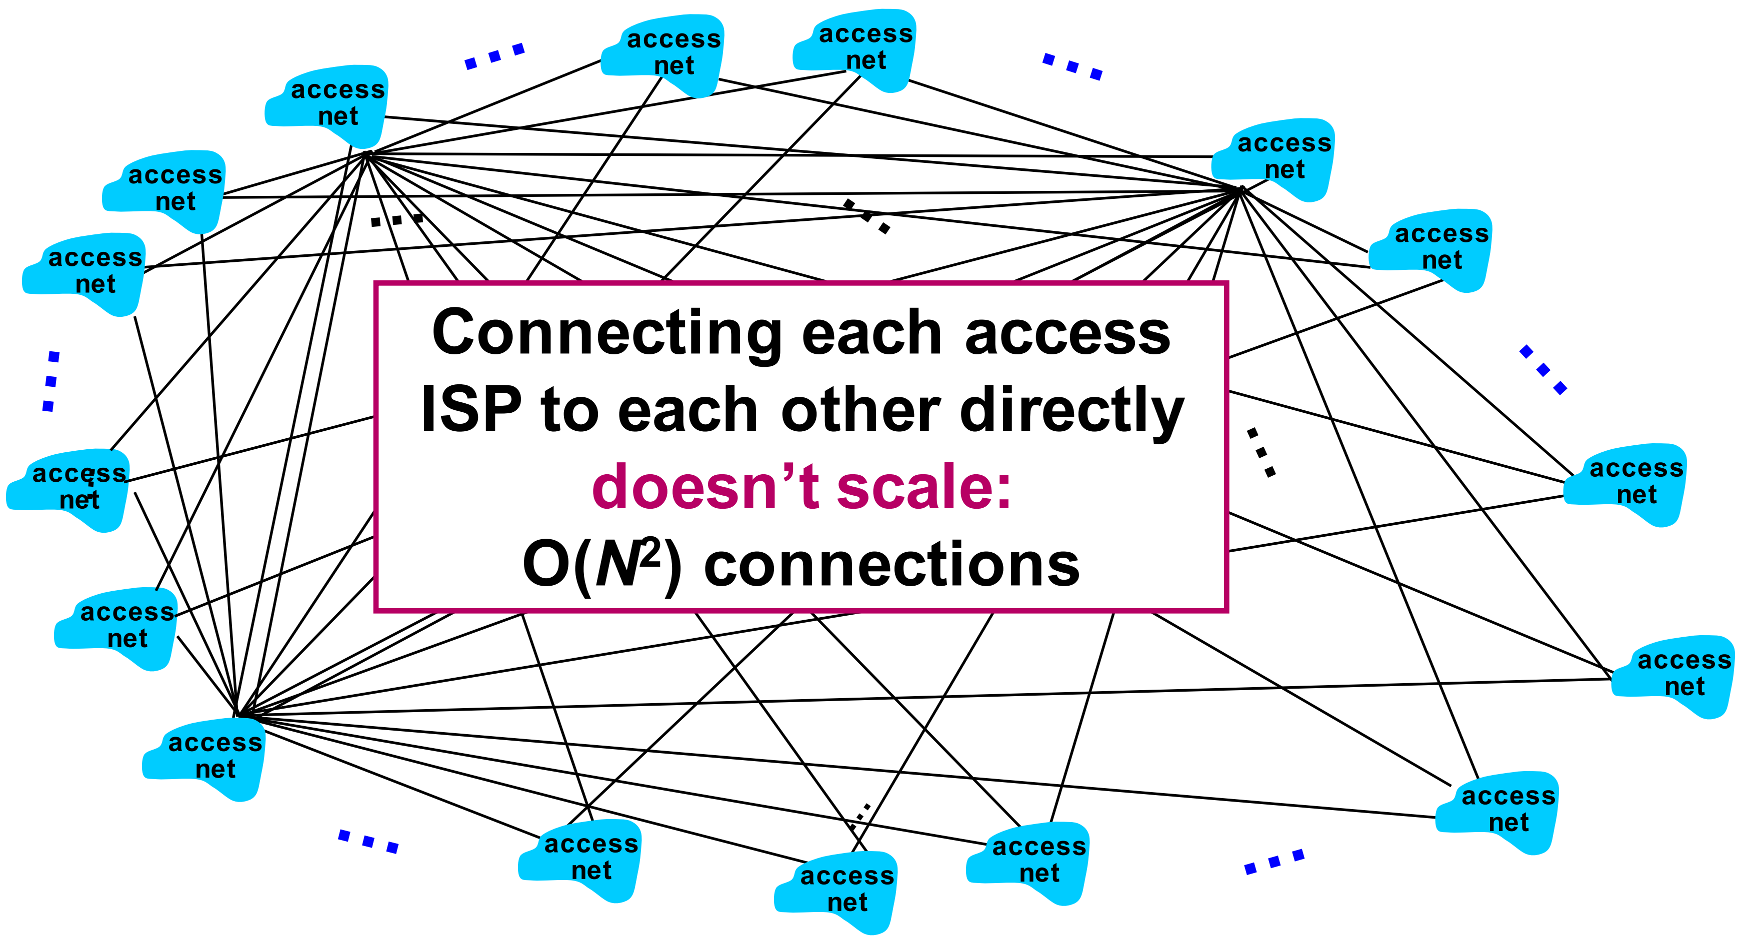

Internet structure: network of networks

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

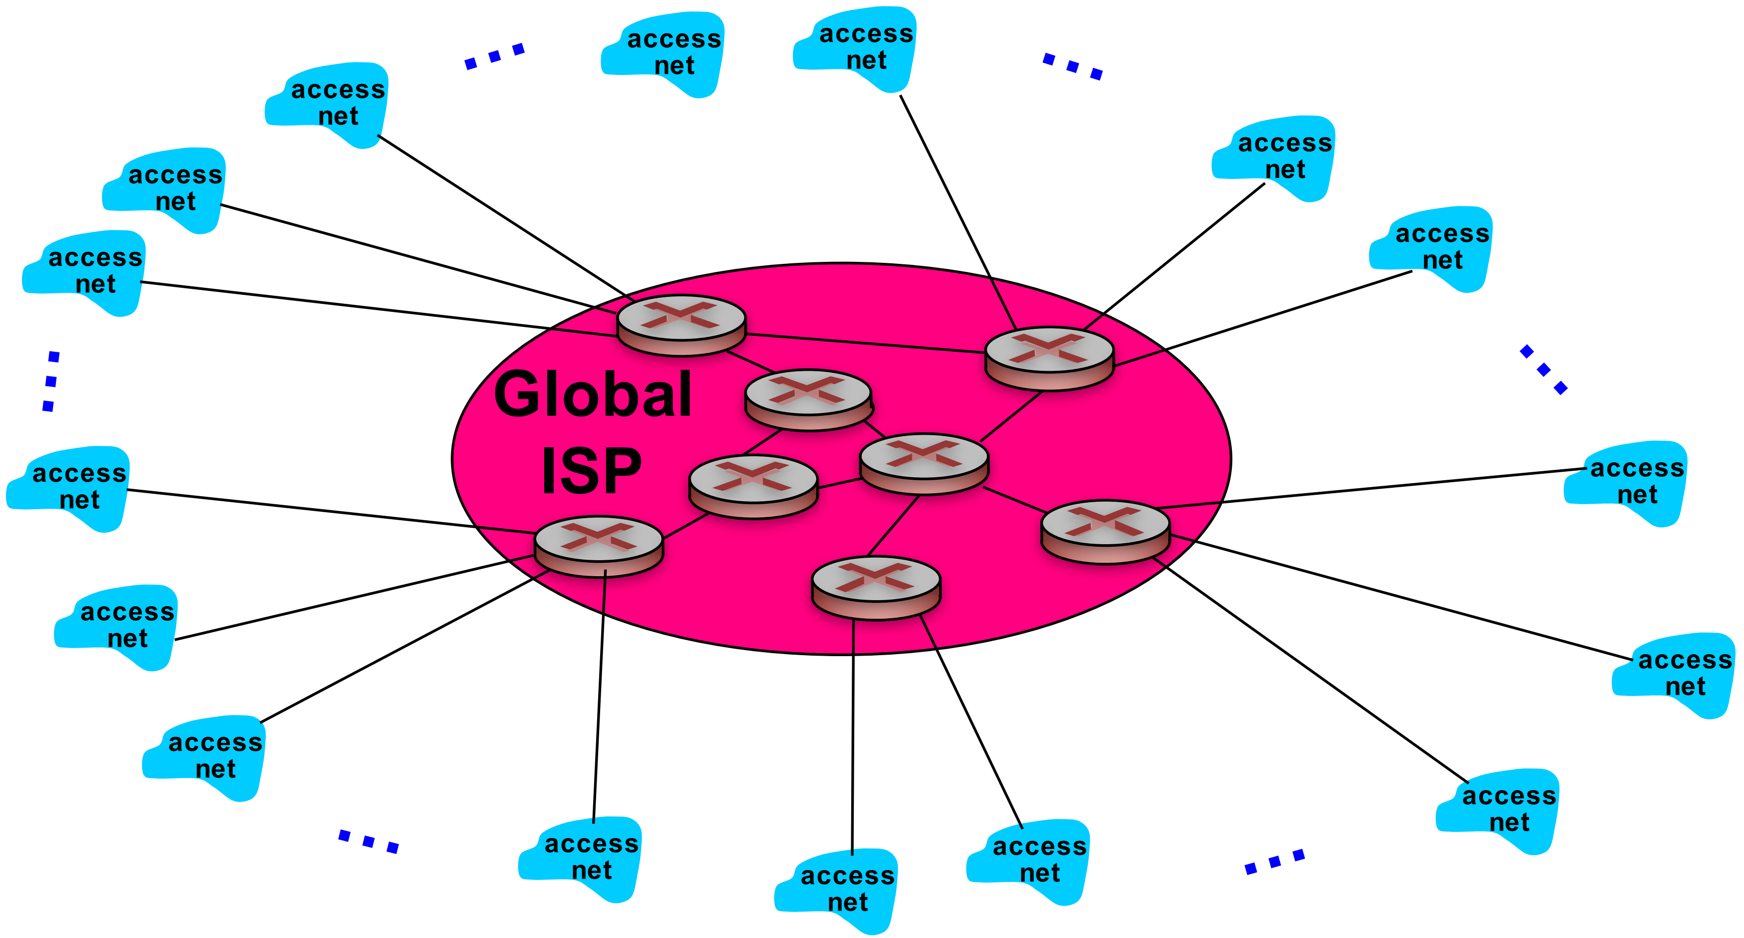

Internet structure: network of networks

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

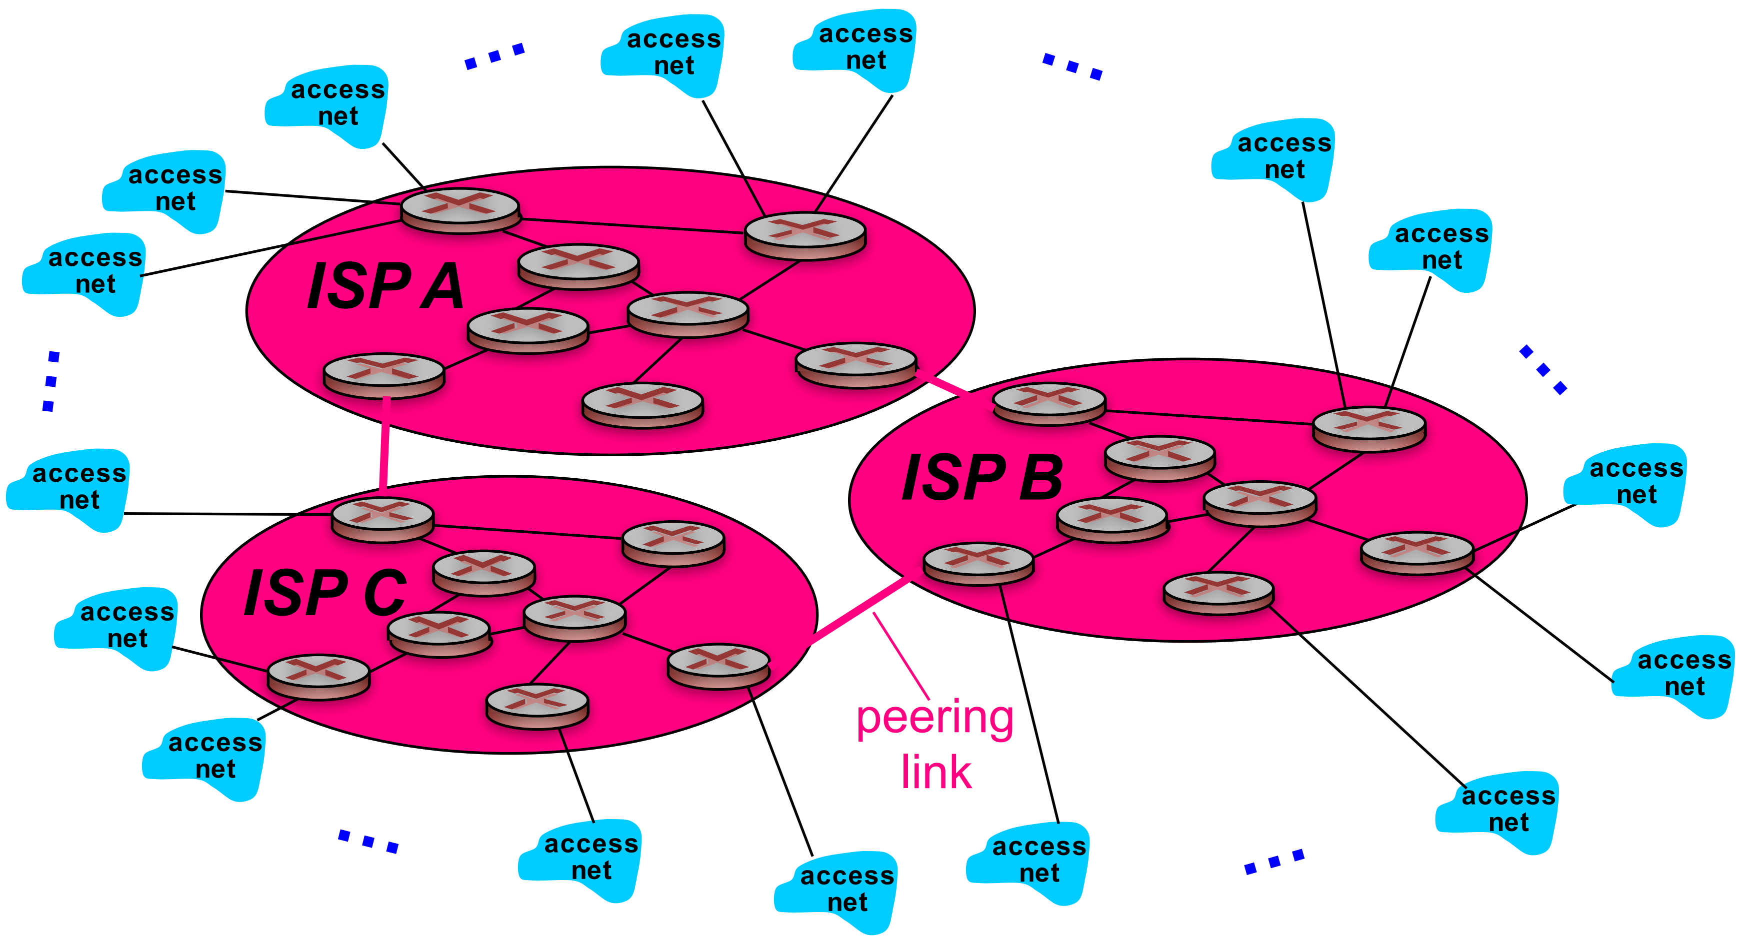

Internet structure: network of networks

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

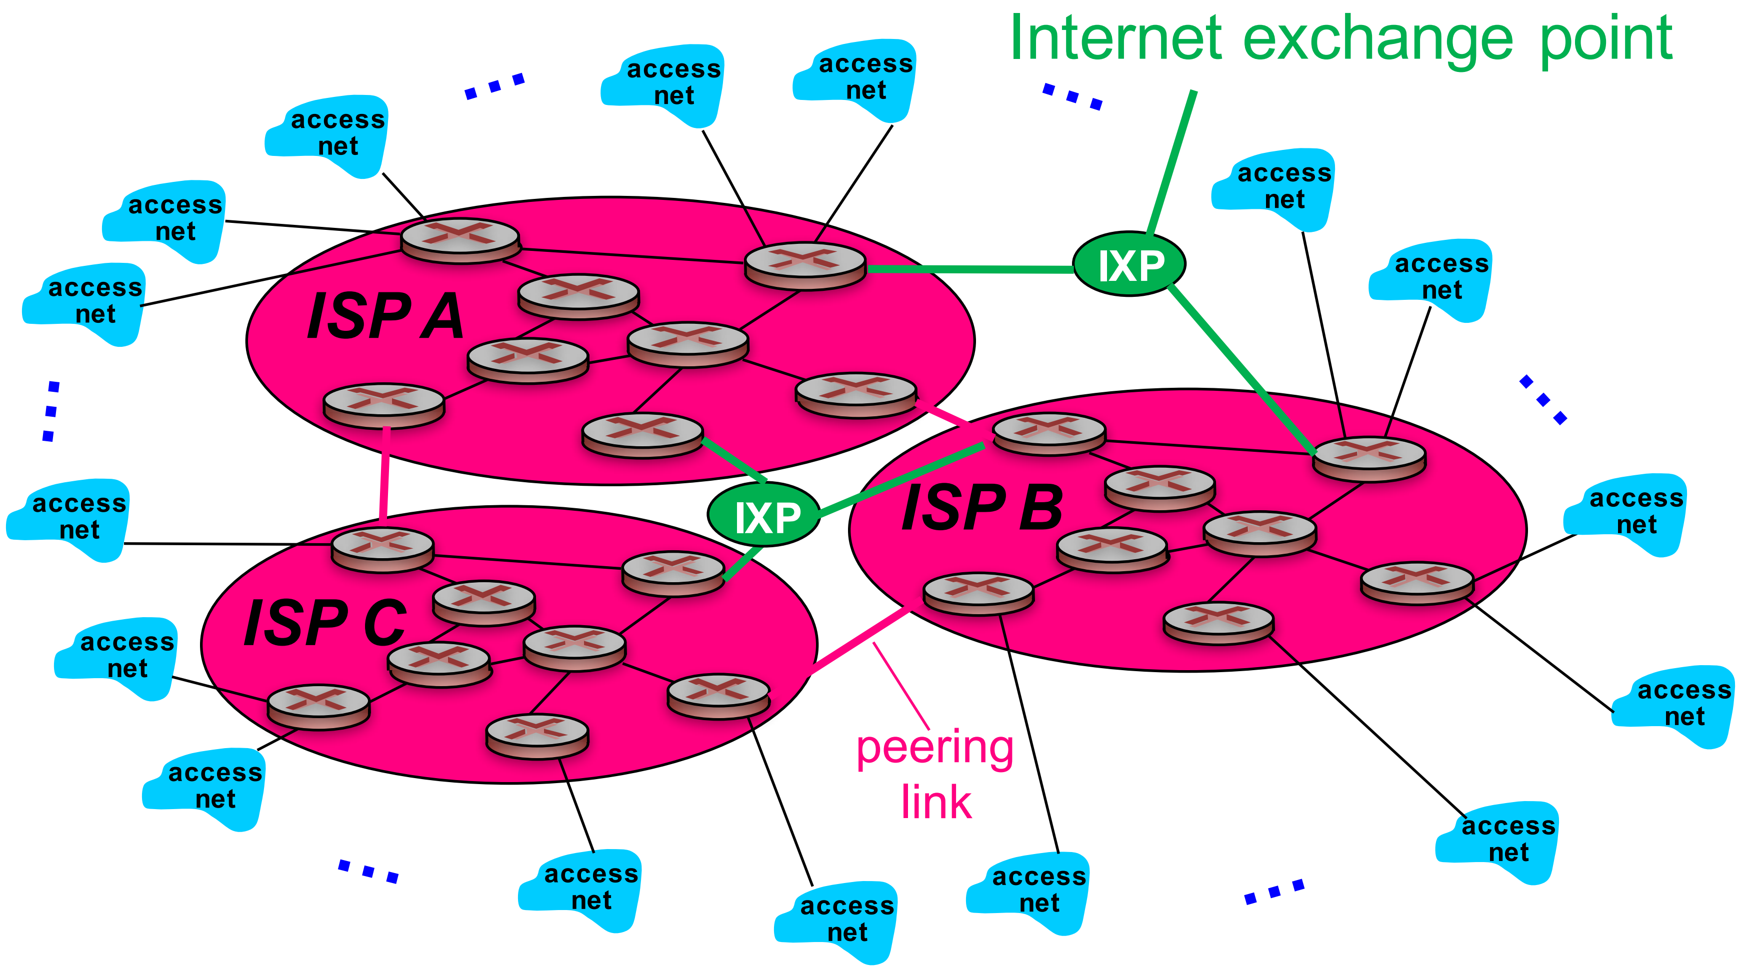

IXP: Internet Exchange Point

Room full of routers

- meeting point where multiple ISPs can peer together

Peer

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

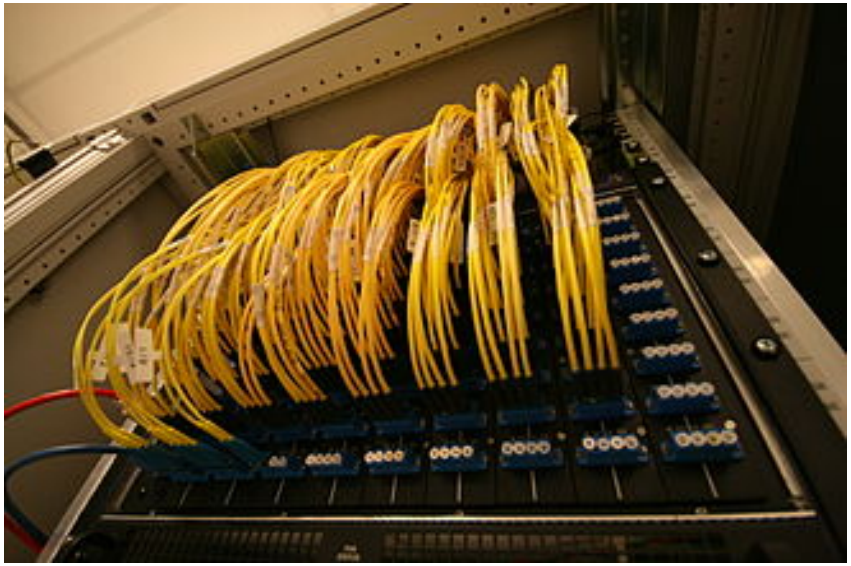

By Fabienne Serriere [CC BY-SA 3.0 (http://creativecommons.org/licenses/by-sa/3.0)], via Wikimedia Commons

Amsterdam IXP

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

Internet structure: network of networks

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

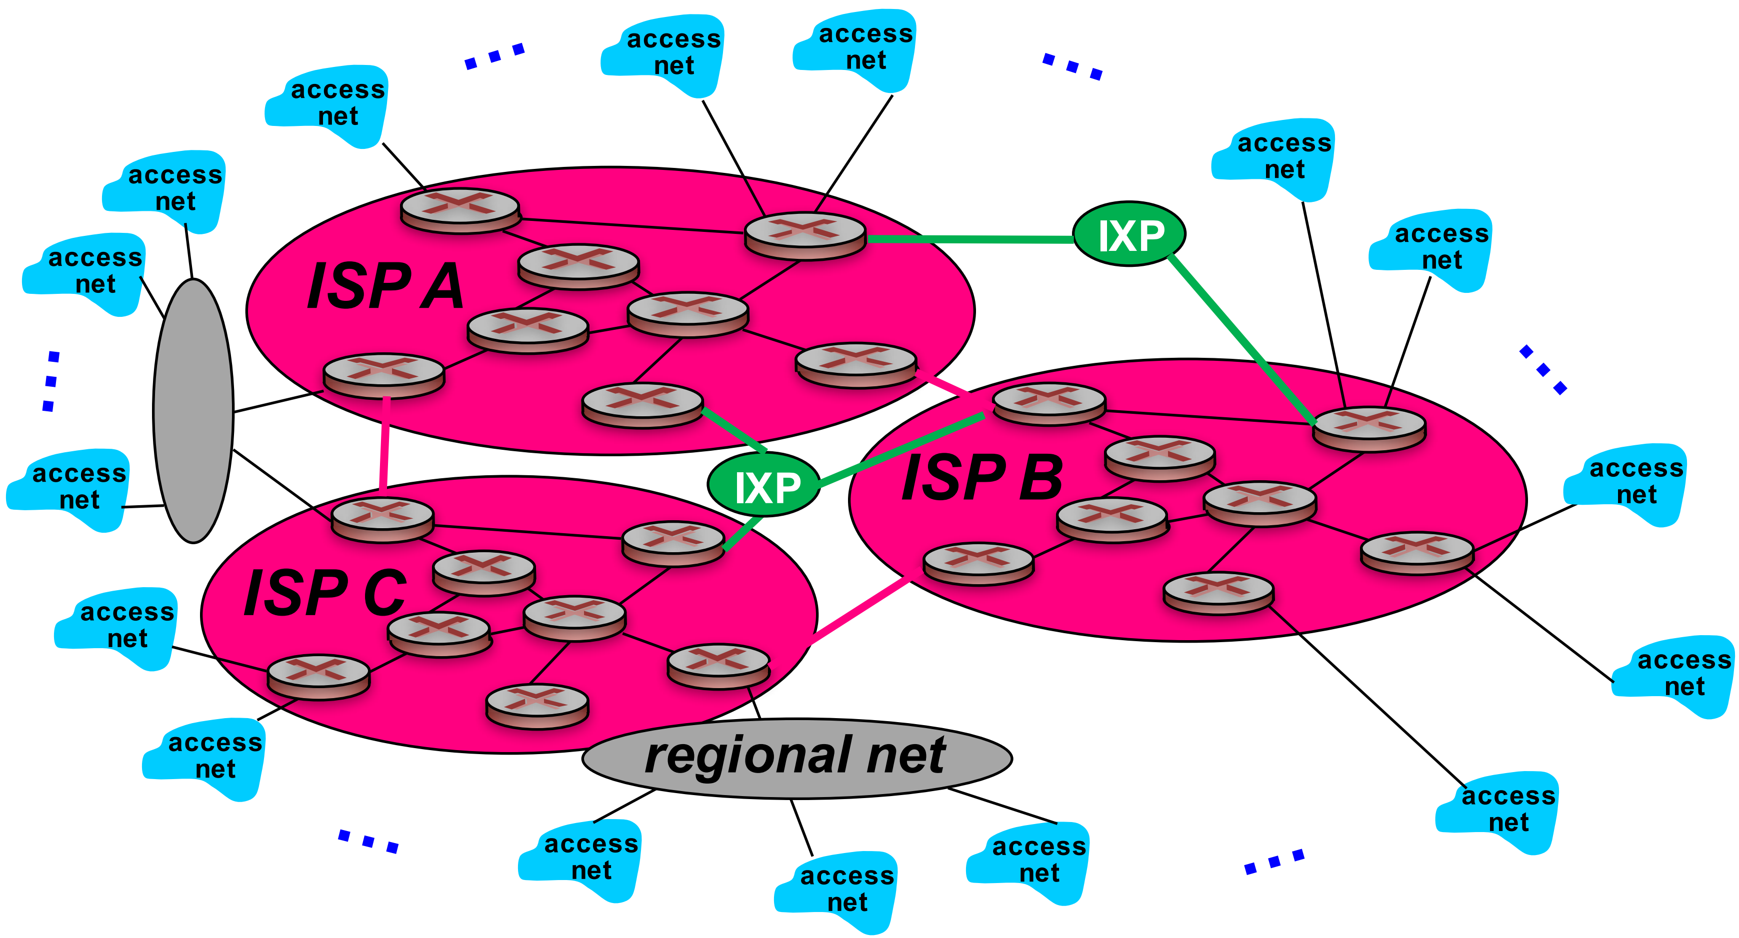

Internet structure: network of networks

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

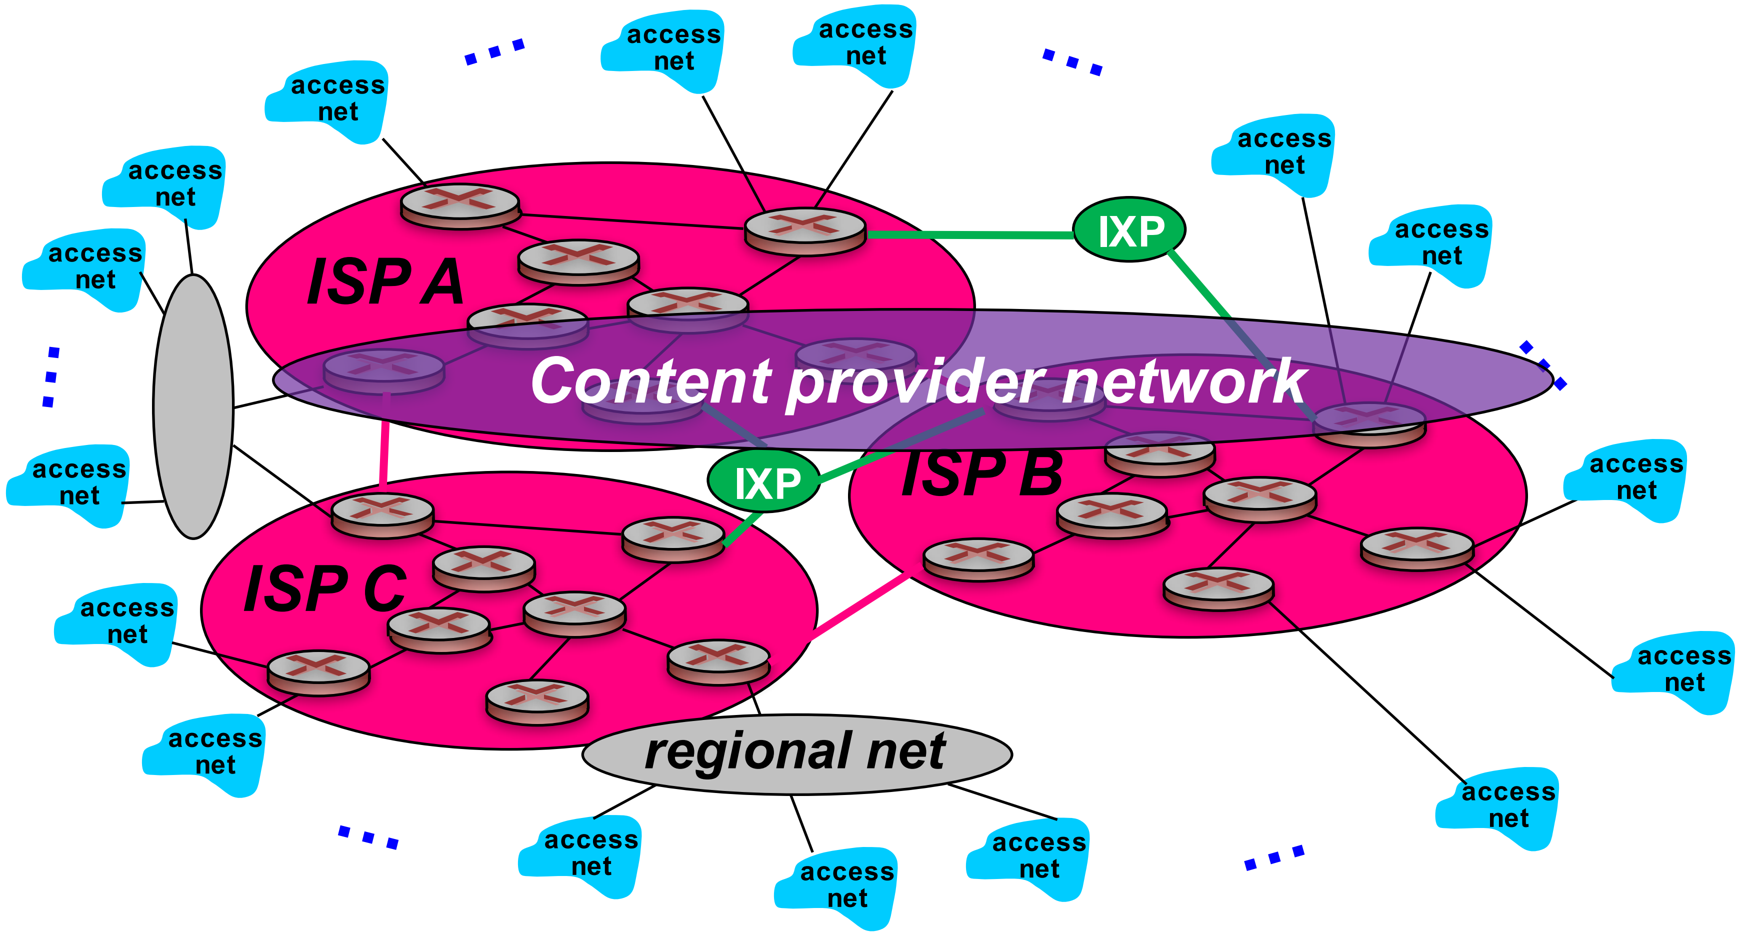

Internet structure: closer look

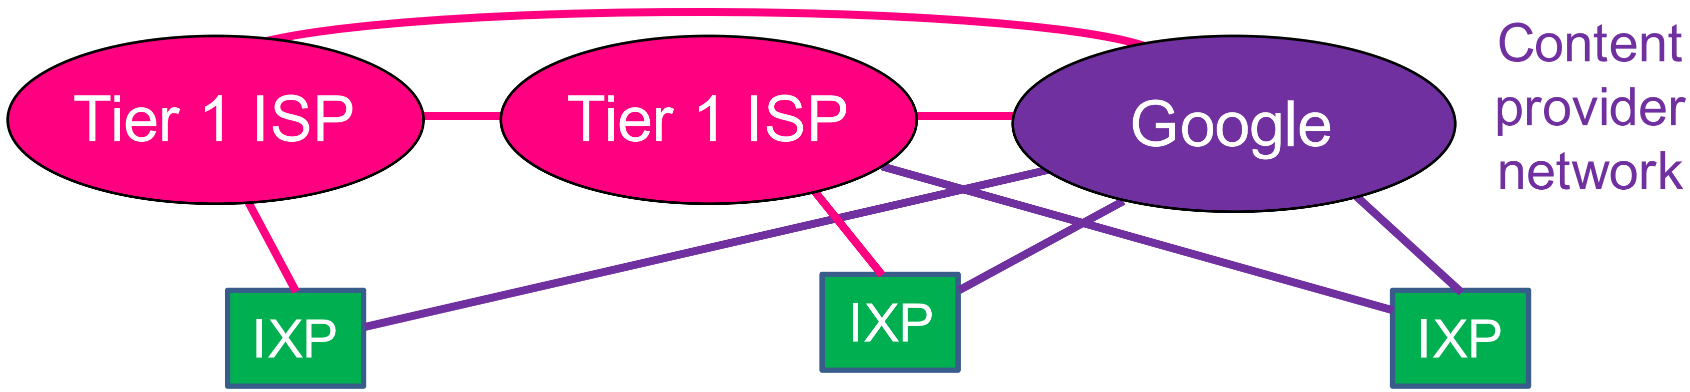

At center: small # of well-connected large networks

"Tier-1" commercial 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

Internet structure: closer look

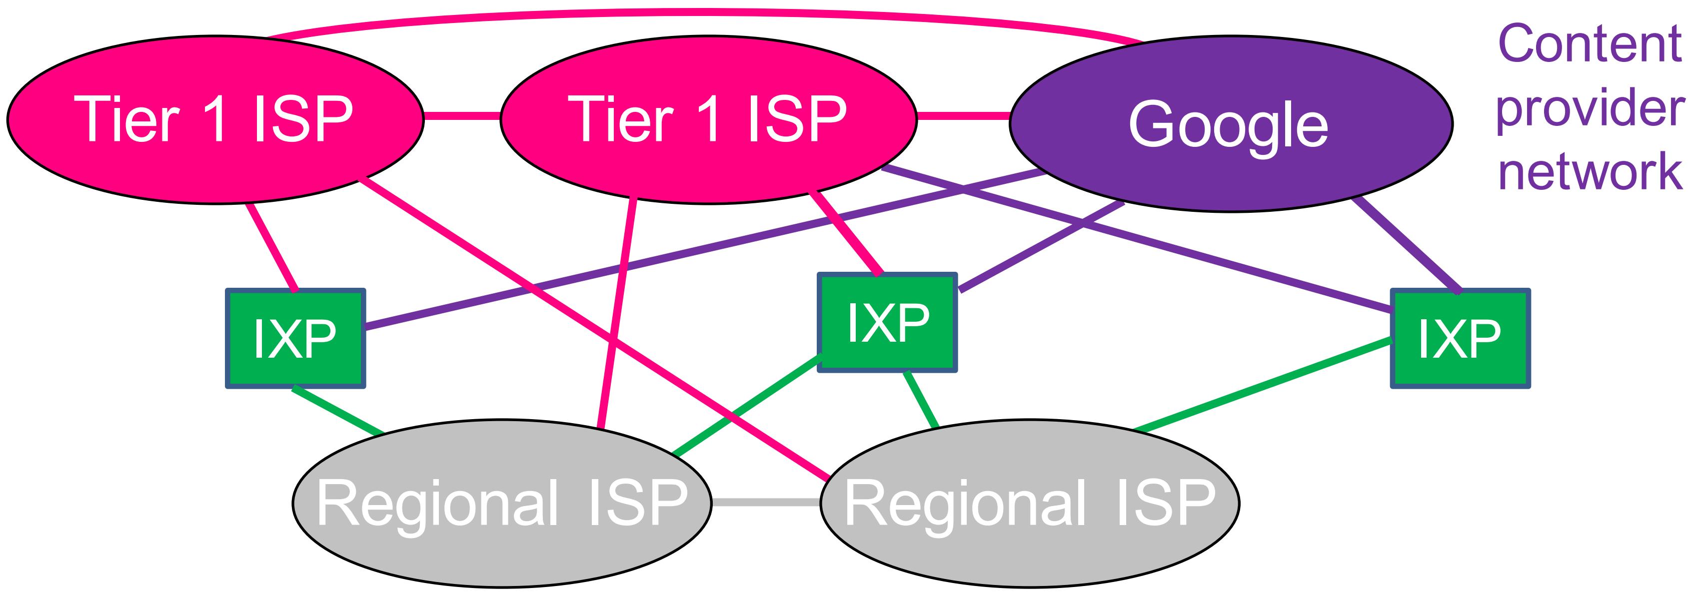

Regional ISP: e.g., Comcast

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

Internet transit

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

Internet structure: closer look

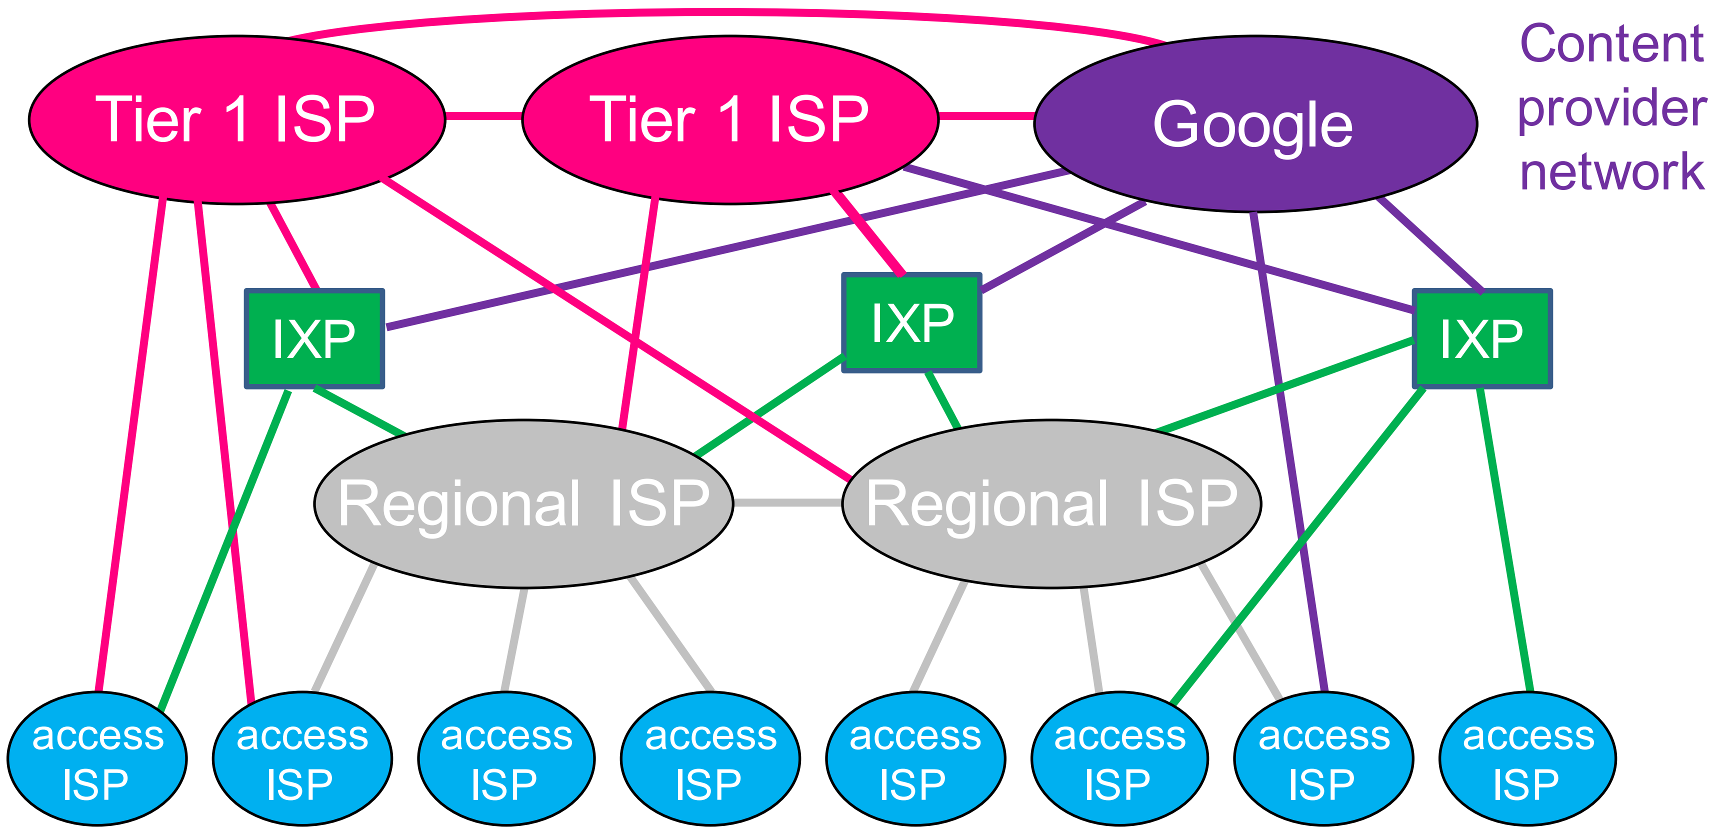

Access ISPs: connect end systems to Internet

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

Internet structure: closer look

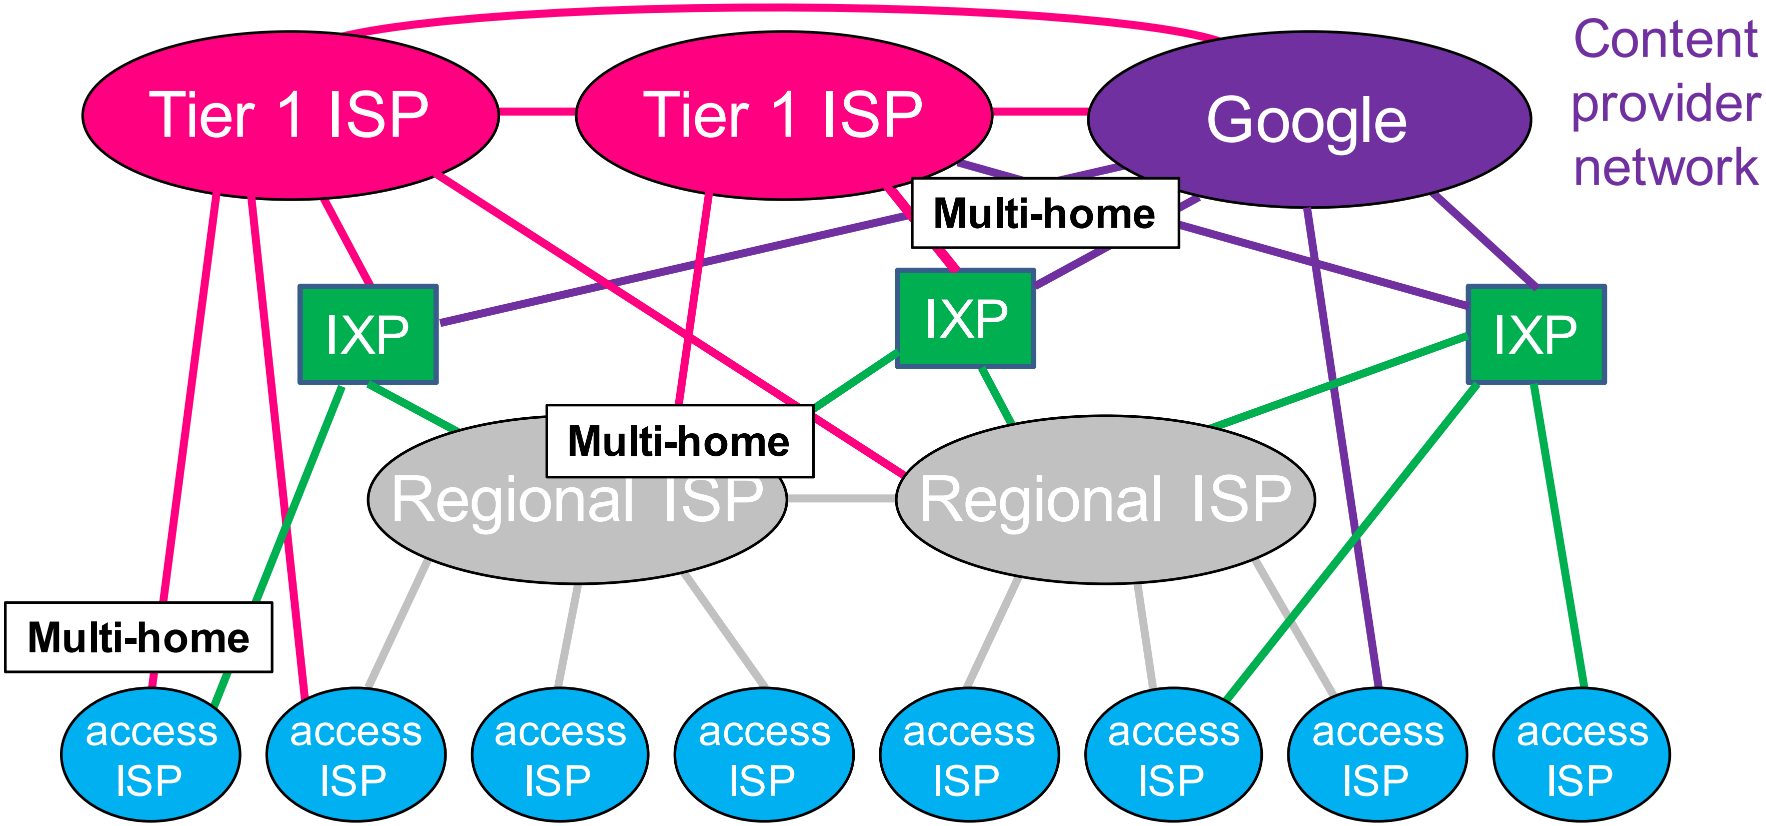

Multi-home

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

Tier-1 ISP: e.g., Sprint

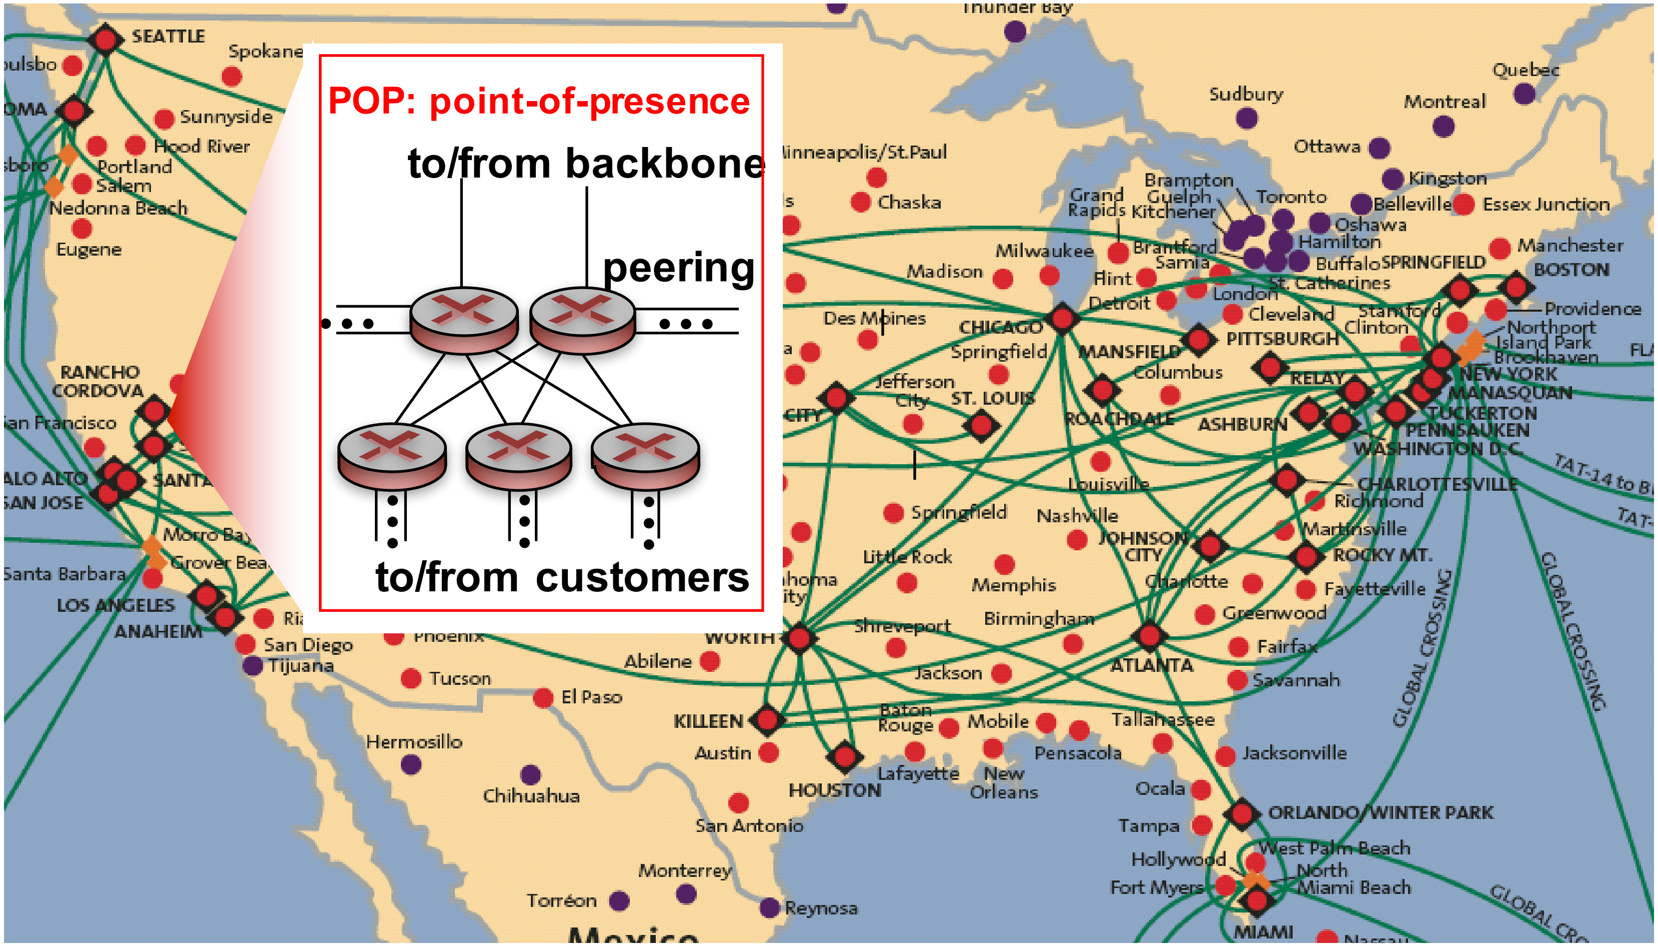

PoP: where customer packets enter network

– E.g., you connecting to access ISP or ISP connecting to provider ISP

Internet

IP ADDRESSES

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 www.google.com
www.google.com has address 74.125.141.99
www.google.com has address 74.125.141.103
www.google.com has address 74.125.141.105
www.google.com has address 74.125.141.147
www.google.com has address 74.125.141.104
www.google.com has address 74.125.141.106
www.google.com has IPv6 address 2607:f8b0:400c:c06::93
```

What's host name for IP address?

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

What's host's IP address?

dig

```
> dig www.google.com

; <<>> DiG 9.8.3-P1 <<>> www.google.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 4619
;; flags: qr rd ra; QUERY: 1, ANSWER: 6, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;www.google.com.                IN      A

;; ANSWER SECTION:
www.google.com.         56      IN      A       74.125.141.104
www.google.com.         56      IN      A       74.125.141.103
www.google.com.         56      IN      A       74.125.141.105
www.google.com.         56      IN      A       74.125.141.147
www.google.com.         56      IN      A       74.125.141.99
www.google.com.         56      IN      A       74.125.141.106

;; Query time: 7 msec
;; SERVER: 129.133.52.12#53(129.133.52.12)
;; WHEN: Mon Jan 22 14:06:38 2018
;; MSG SIZE  rcvd: 128
```

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 #) pair
- different processes on computer will connect to network using same IP address but different port numbers

How do two processes talk over a network?

Via sockets: connection endpoint with associated IP addr, port #

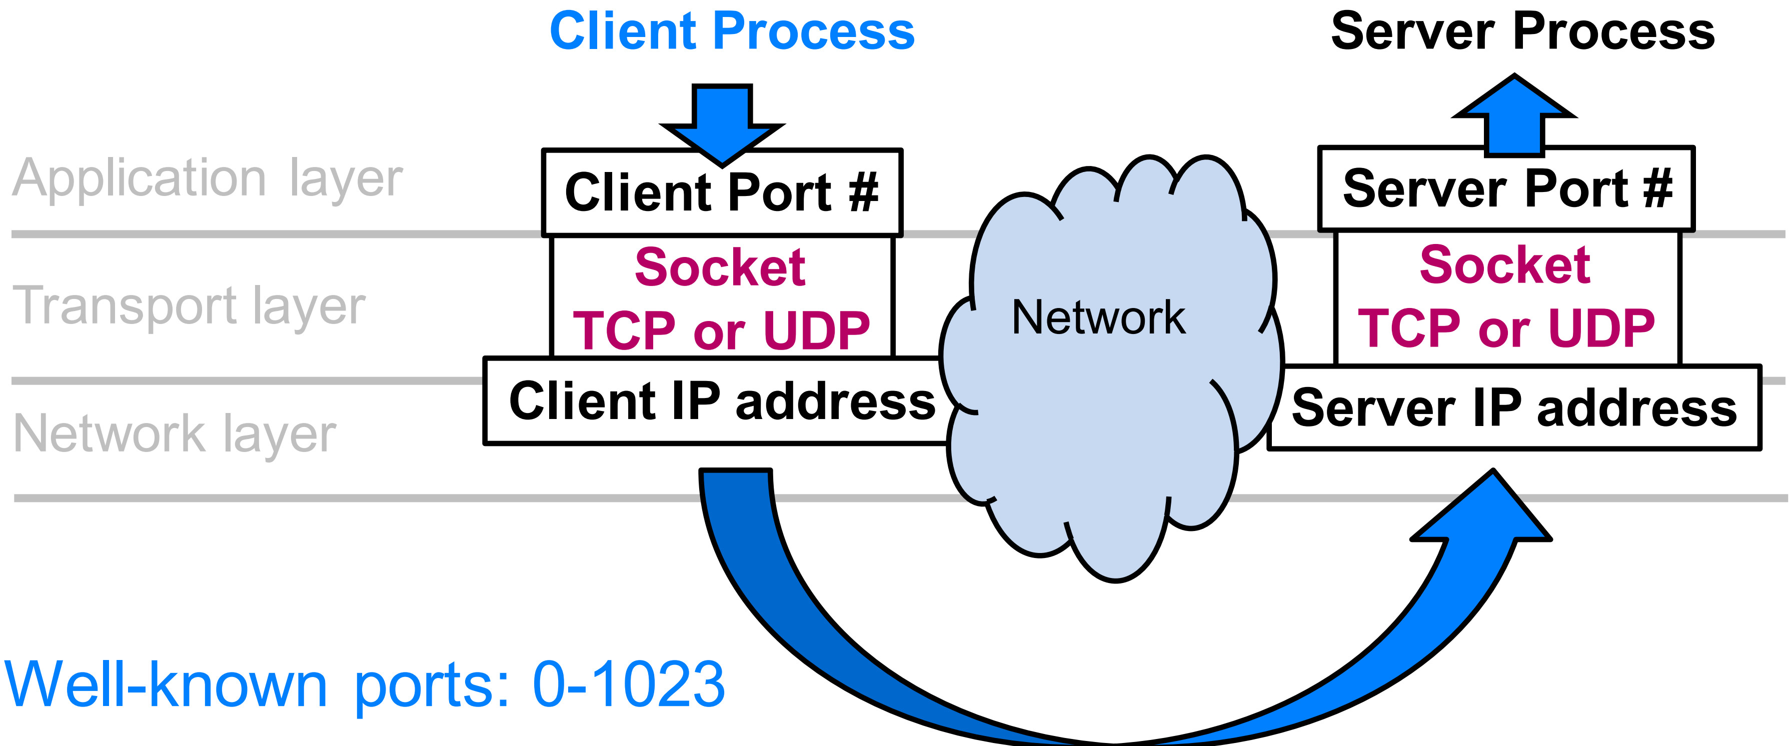

Well-known ports: 0-1023

- E.g., HTTP is port 80

Registered ports: 1023-49151

Available ports: 49152-65535

Internet

NETWORK PERFORMANCE METRICS

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

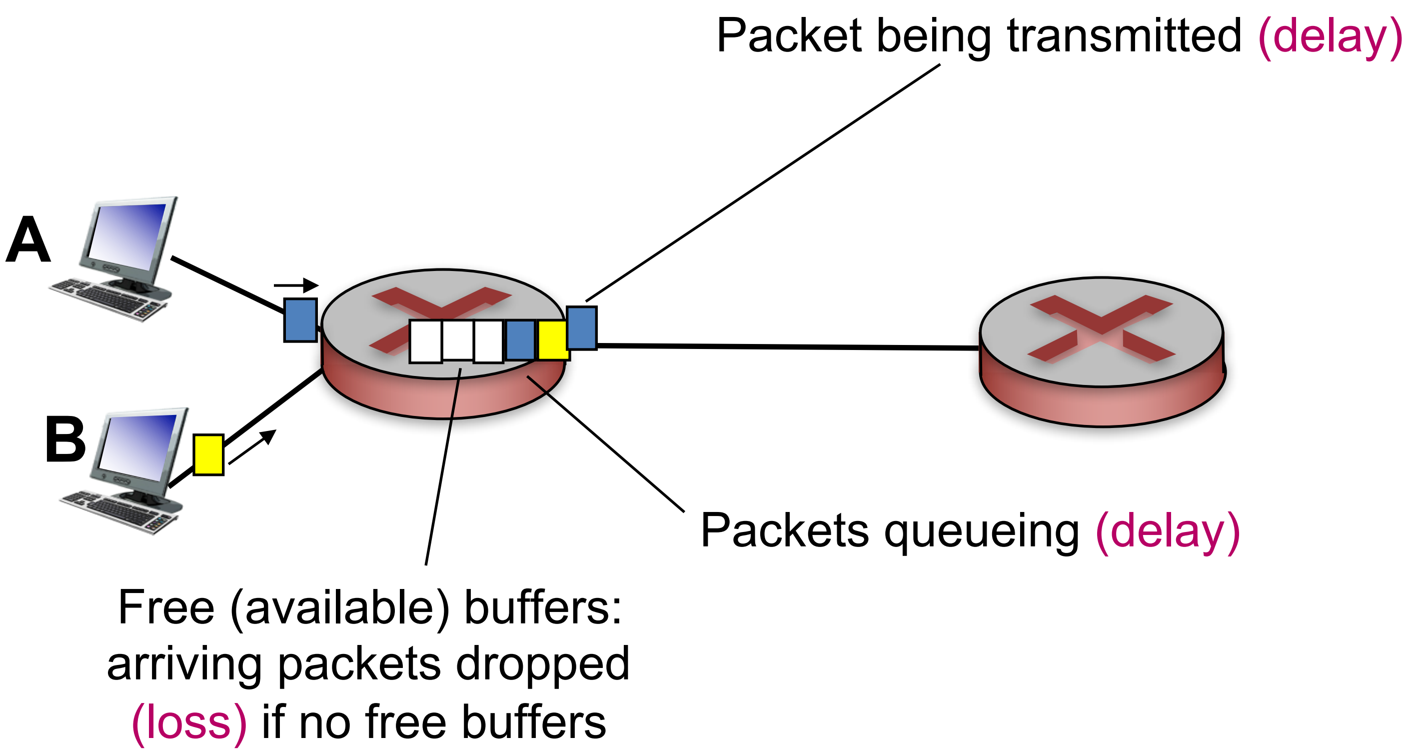

Four sources of packet delay

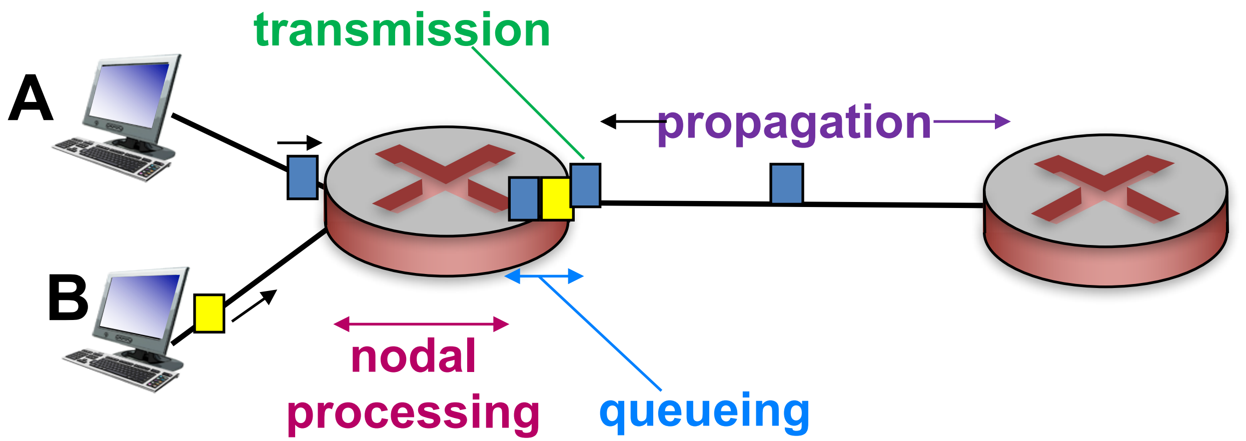

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

d_{proc}: processing delay

- check bit errors
- determine output link
- Fast: typically < msec
- Usually done in hardware not software

d_{queue}: queueing delay

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

Four sources of packet delay

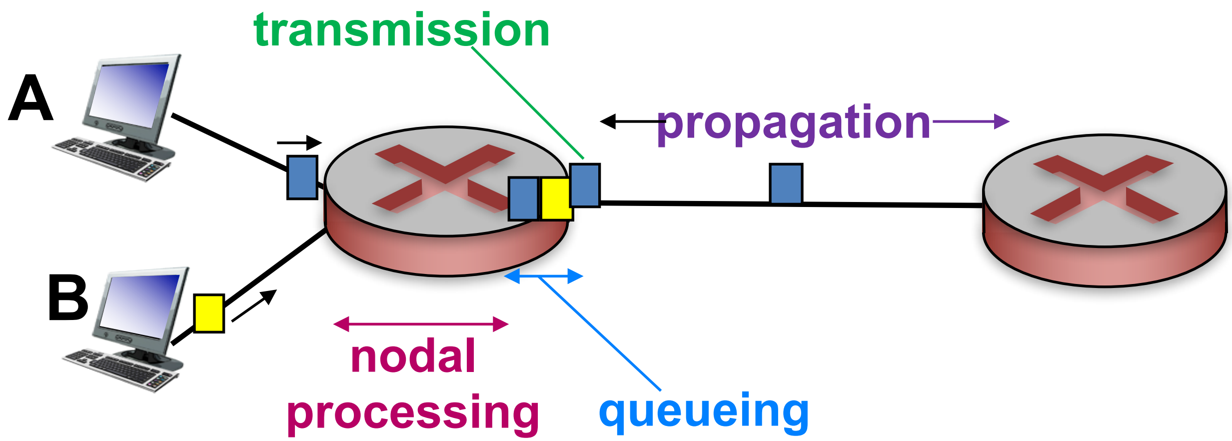

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

d_{trans}: transmission delay

- Depends on link bandwidth
- *L*: packet length (bits)
- *R*: link bandwidth (bps)
- $d_{trans} = L/R$

← d_{trans} and d_{prop} *very* different →

d_{prop}: propagation delay

- μsec (2 routers on same campus) to msec (satellite link)
- *d*: length of physical link
- *s*: propagation speed (~2×10^8 m/sec)
- $d_{prop} = d/s$

End-to-end delay

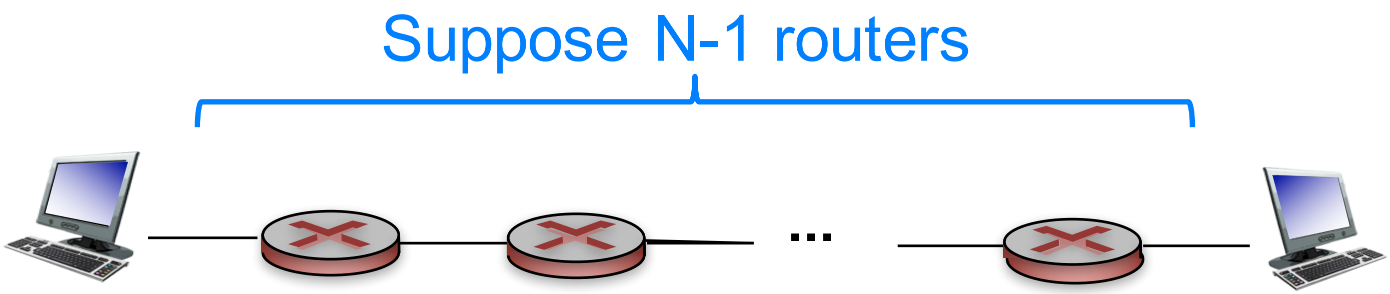

Q: what is end-end delay ignoring queuing delay?

End-end delay = $N * (d_{proc} + d_{trans} + d_{prop})$

Queueing delay (revisited)

L: packet length (bits)

R: link bandwidth (bps)

a: average packet arrival rate

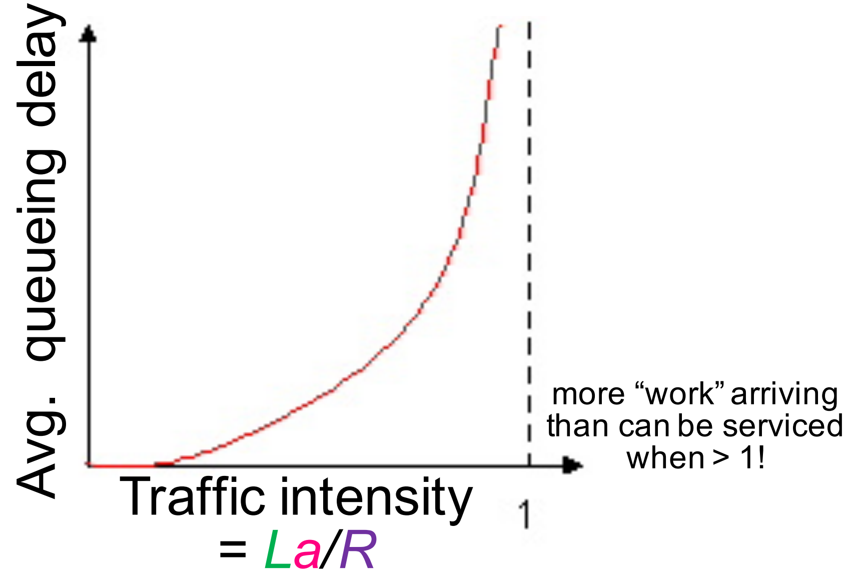

La/R ~ 0: small avg. queueing delay

La/R→*1*: large avg. queueing delay

La/R > *1*: infinite avg. queueing delay

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

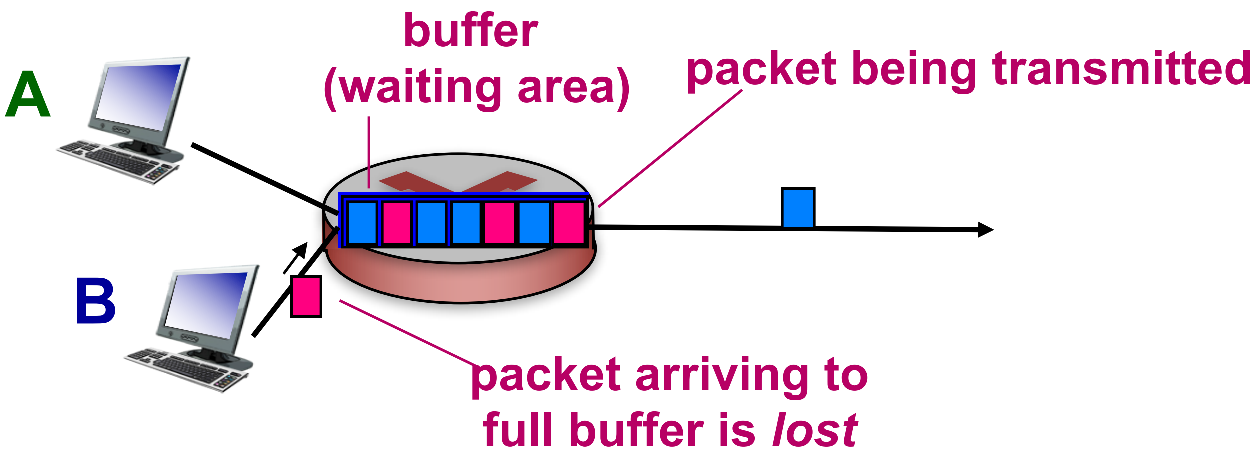

Throughput

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

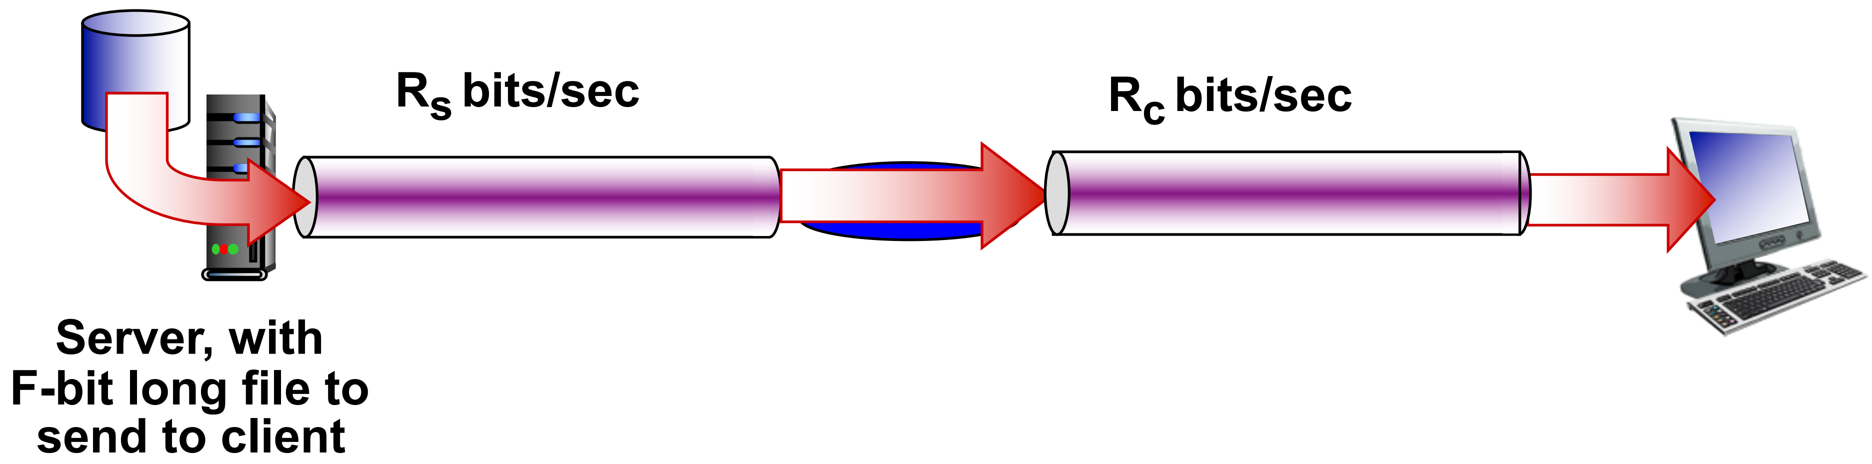

Throughput

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

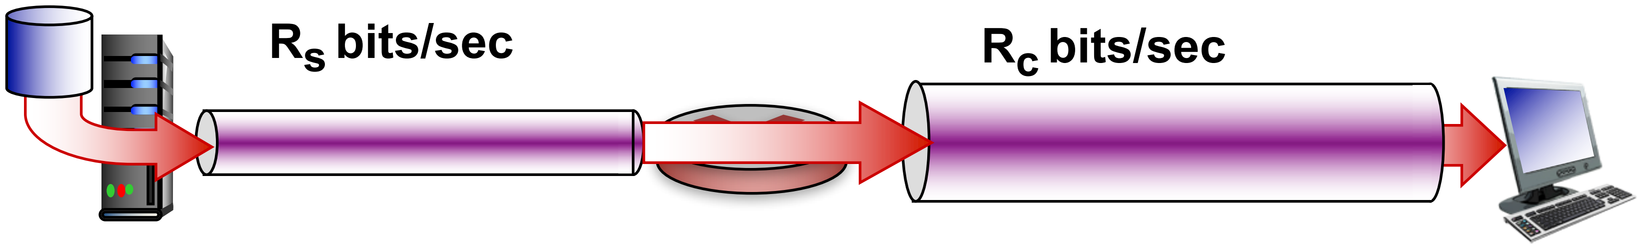

$R_s > R_c$ What is average end-end throughput?

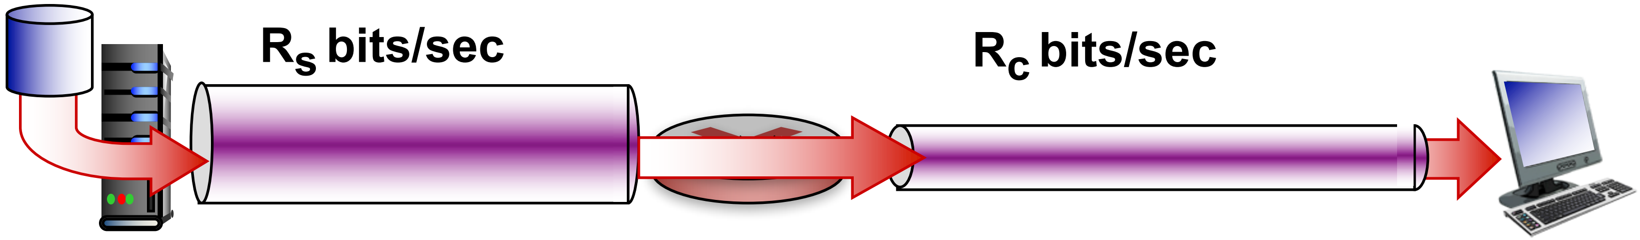

bottleneck link

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

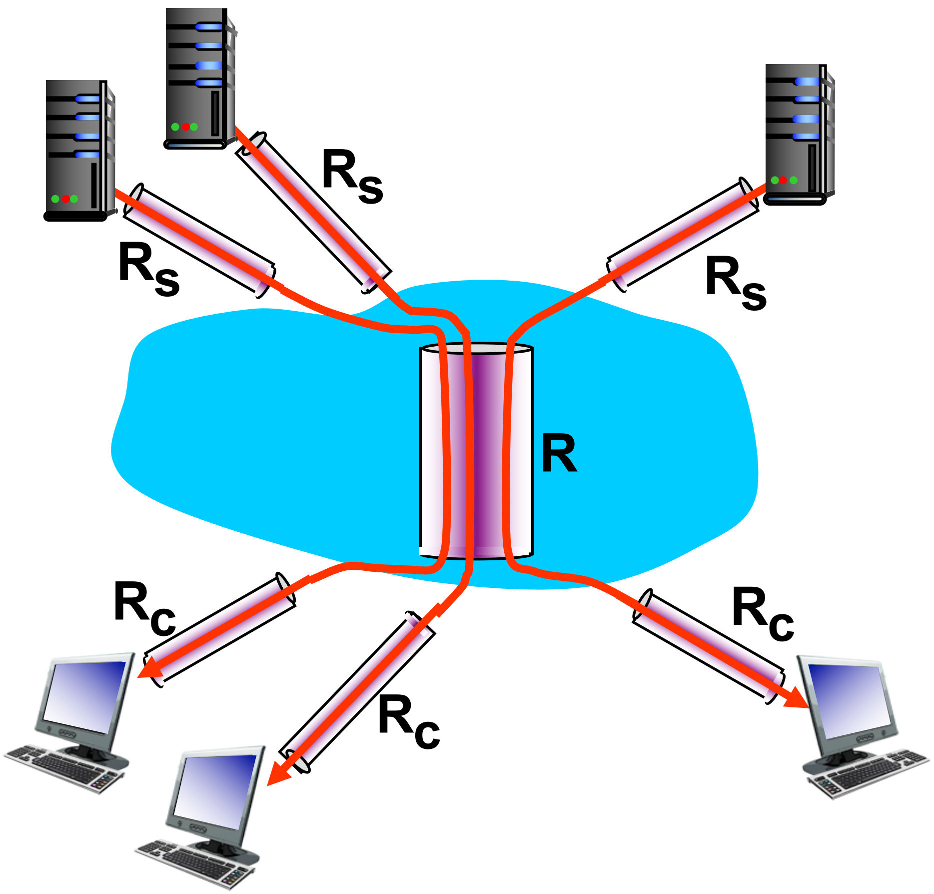

10 connections (fairly) share *R* bits/sec backbone bottleneck link

Internet

TRACEROUTE

Real Internet delays and routes

Traceroute program

- provides delay measurement from source to router along end-end Internet path towards destination

How?

- for all i:
  - sends three packets that will reach router i on path towards destination
    - sets packet time-to-live (TTL) to i
  - router i will return packets to sender
  - sender times interval between transmission and reply for each packet
    - measures Round Trip Time (RTT) delay

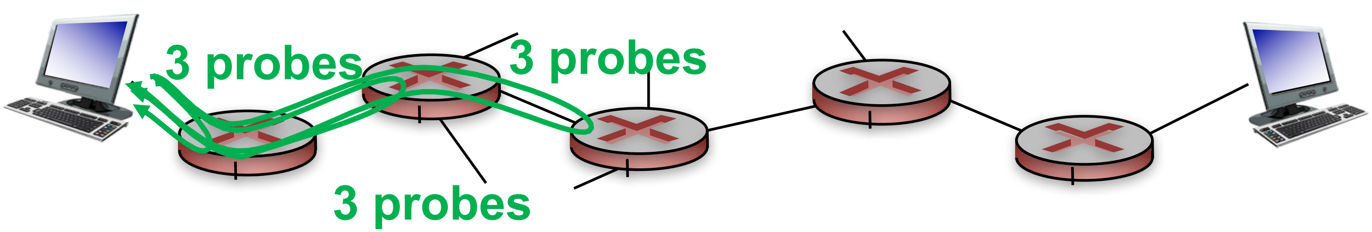

Note

- Different probe packets may take different paths, so delays can vary

Real Internet delays, routes

traceroute: from wesleyan network to cs.stanford.edu

```
> traceroute cs.stanford.edu
traceroute to cs.stanford.edu (171.64.64.64), 64 hops max, 52 byte packets
 1  129.133.176.1 (129.133.176.1)  6.138 ms  2.365 ms  3.913 ms
 2  172.16.100.1 (172.16.100.1)  3.857 ms  3.361 ms  5.545 ms
 3  129.133.2.5 (129.133.2.5)  1.958 ms  3.068 ms  1.906 ms
 4  129.133.4.11 (129.133.4.11)  3.865 ms  2.879 ms  3.850 ms
 5  72.10.111.129 (72.10.111.129)  1.984 ms  3.144 ms  3.761 ms
 6  64.251.60.122 (64.251.60.122)  5.928 ms  4.959 ms  4.910 ms
 7  enrt083h-9k-te0-3-0-7.net.cen.ct.gov (72.10.125.22)  7.003 ms  6.982 ms  5.593 ms
 8  enrt078h-9k-te-0-0-0-6-dwdm-1532-68.net.cen.ct.gov (67.218.83.185)  5.779 ms  6.966 ms  6.300 ms
 9  cen-re-nox300gw1.nox.org (192.5.89.202)  7.421 ms  5.172 ms  5.948 ms
10  nox300gw1-cen-re.nox.org (192.5.89.201)  8.196 ms  8.217 ms  8.240 ms
11  192.5.89.22 (192.5.89.22)  9.766 ms  7.964 ms  7.810 ms
12  i2-re-nox1sumgw1.nox.org (192.5.89.18)  12.955 ms  7.642 ms  8.033 ms
13  et-7-0-0.4079.sdn-sw.alba.net.internet2.edu (162.252.70.96)  11.953 ms  10.251 ms  12.146 ms
14  et-3-1-0.4079.rtsw.clev.net.internet2.edu (162.252.70.93)  21.406 ms  20.401 ms  21.959 ms
15  ae-1.4079.sdn-sw.eqch.net.internet2.edu (162.252.70.131)  29.059 ms  30.883 ms  29.264 ms
16  ae-2.4079.rtsw.chic.net.internet2.edu (162.252.70.132)  29.075 ms  30.298 ms  29.413 ms
17  ae-3.4079.rtsw.kans.net.internet2.edu (162.252.70.141)  40.831 ms  40.250 ms  41.068 ms
18  ae-5.4079.rtsw.salt.net.internet2.edu (162.252.70.145)  60.625 ms  61.459 ms  60.568 ms
19  ae-1.4079.rtsw.losa.net.internet2.edu (162.252.70.114)  72.171 ms  73.579 ms  74.209 ms
20  hpr-lax-hpr2--i2-r&e.cenic.net (137.164.26.200)  73.938 ms  73.487 ms  72.439 ms
21  hpr-svl-hpr3--lax-hpr3-100ge.cenic.net (137.164.25.74)  83.925 ms  84.645 ms  83.688 ms
22  hpr-stan-ge--svl-hpr2.cenic.net (137.164.27.162)  86.215 ms  86.925 ms  84.094 ms
23  csmx-west-rtr.sunet (171.64.255.214)  109.002 ms  144.984 ms  94.379 ms
24  cs.stanford.edu (171.64.64.64)  84.106 ms  84.984 ms  83.928 ms
```

3 delay measurements from cs.stanford.edu

cross-country links

* means no response (probe lost, router not replying)