Chapter 1
Introduction

Yaxiong Xie

Department of Computer Science and Engineering
University at Buffalo, SUNY

Adapted from the slides of the book's authors

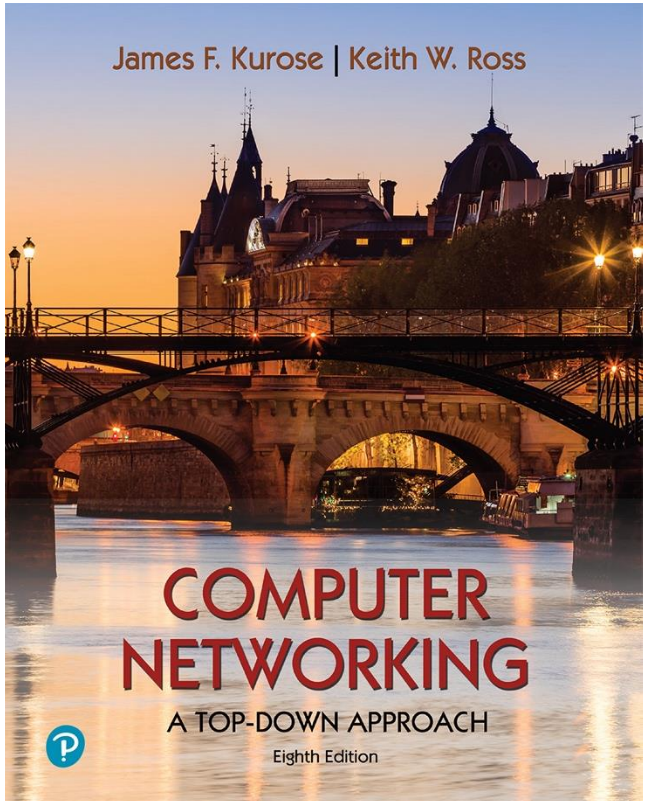

Computer Networking: A Top-Down Approach
8th edition
Jim Kurose, Keith Ross
Pearson, 2020

Chapter 1: introduction

Chapter goal:

- Get "feel," "big picture," introduction to terminology
  - more depth, detail *later* in course

Overview/roadmap:

- What *is* the Internet? What *is* a protocol?
- Network edge: hosts, access network, physical media
- Network core: packet/circuit switching, internet structure
- Performance: loss, delay, throughput
- Protocol layers, service models
- Security

What’s Internet? An analogy: USPS

USPS system

USPS network

What’s Internet? An analogy: USPS

USPS system

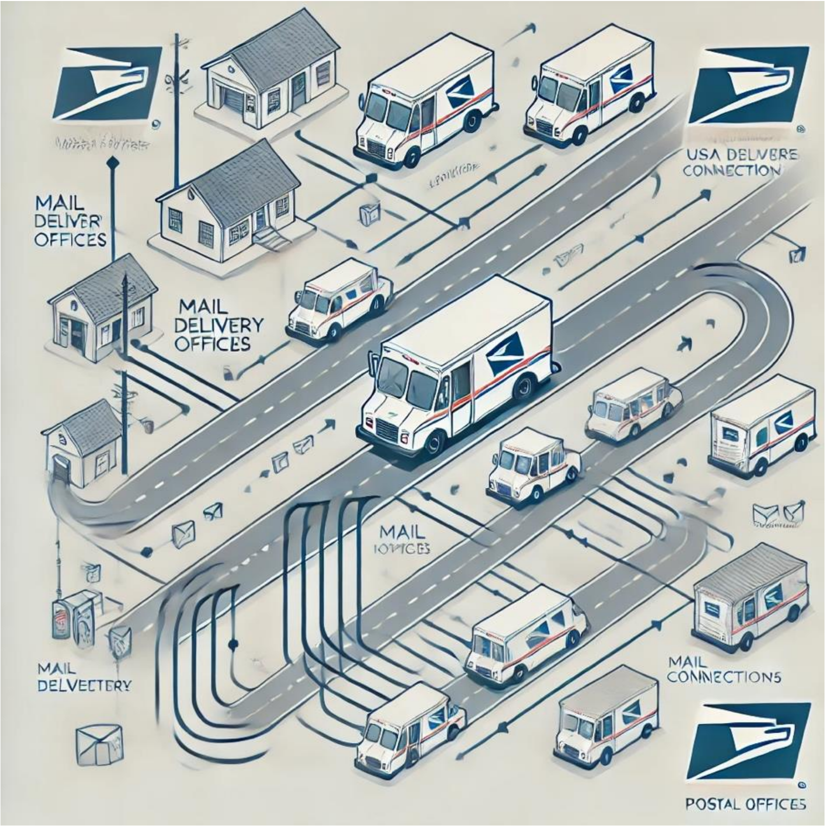

USPS network

What’s Internet? An analogy: USPS

USPS system

USPS edge

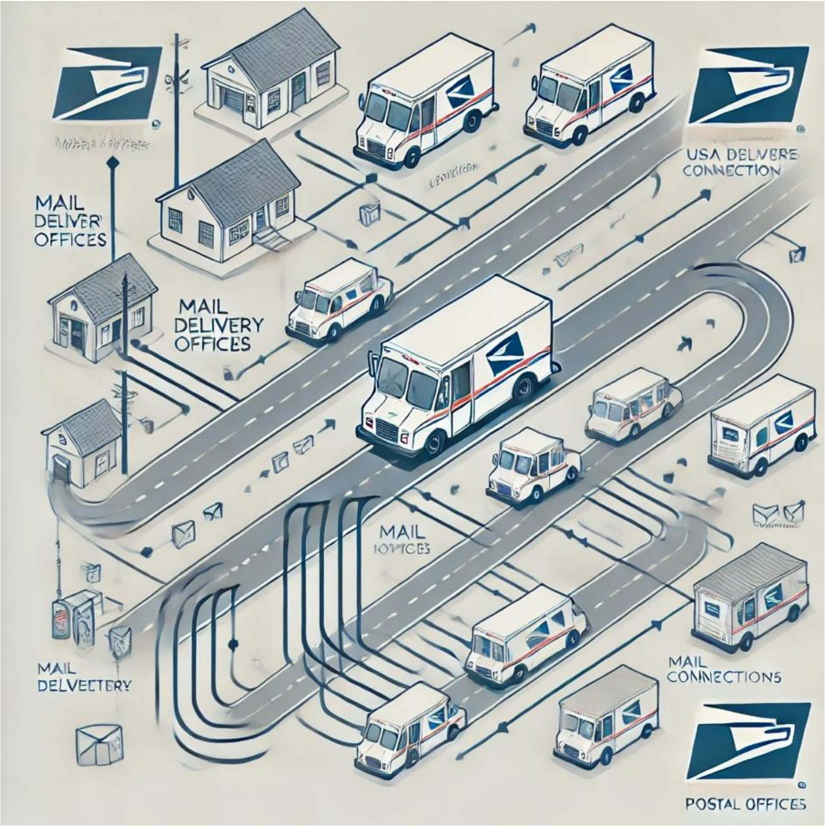

USPS core network

What’s Internet?

Internet System

Packets

Internet edge

Packets

Internet core network

The Internet: a “nuts and bolts” view

Billions of connected computing *devices*:
- *hosts = end systems*
- running *network apps* at Internet’s “edge”

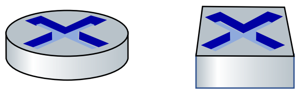

Packet switches: forward packets (chunks of data)
- *routers, switches*

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

Networks
- collection of devices, routers, links: managed by an organization

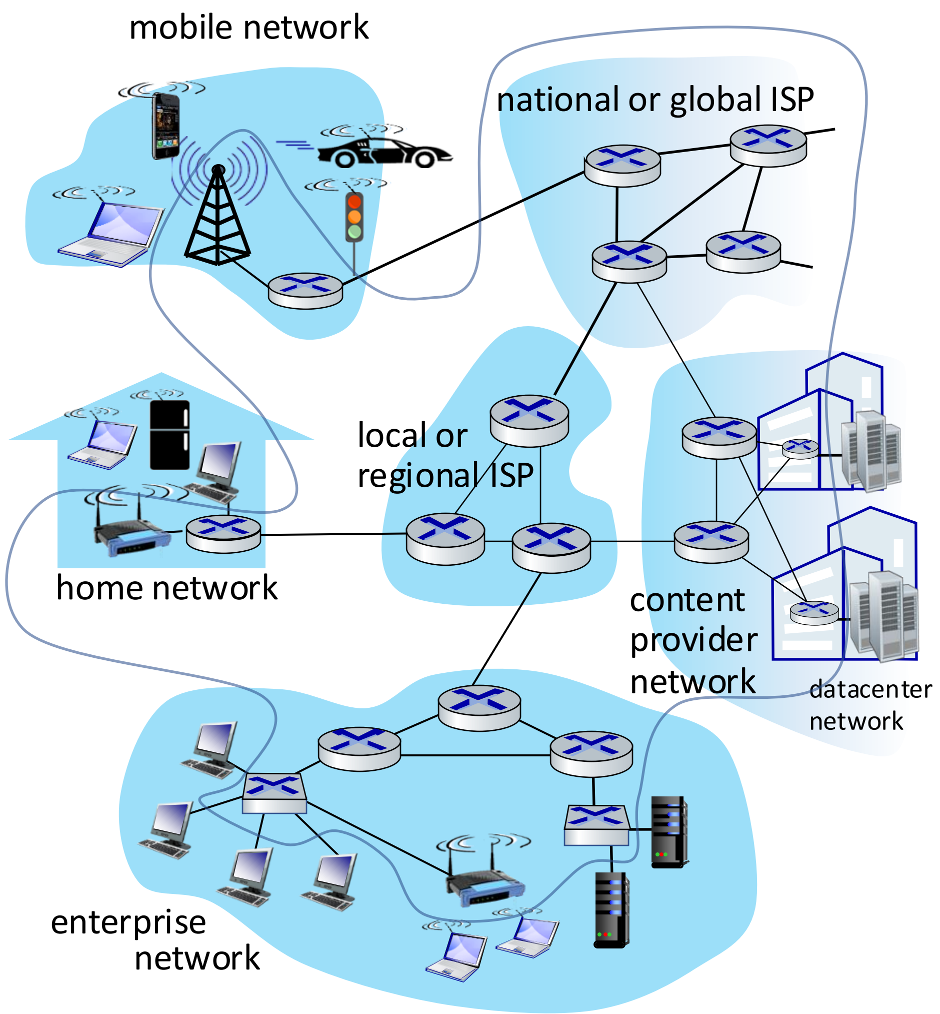

What's the Internet: "nuts and bolts" view -continued

- *Software: protocols* control sending, receiving of msgs
  - e.g., HTTP (web), SMTP (for email server),
  - Wifi /BT (802.x) for wireless devices,
  - Ethernet (for local area networks),
  - TCP/UDP (for hosts on the internet)
  - IP (for the routers in the core networks)
- Internet standards define these protocols
  - RFC: Request for comments
  - IETF: Internet Engineering Task Force

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

What's a protocol?

A human protocol and a computer network protocol:

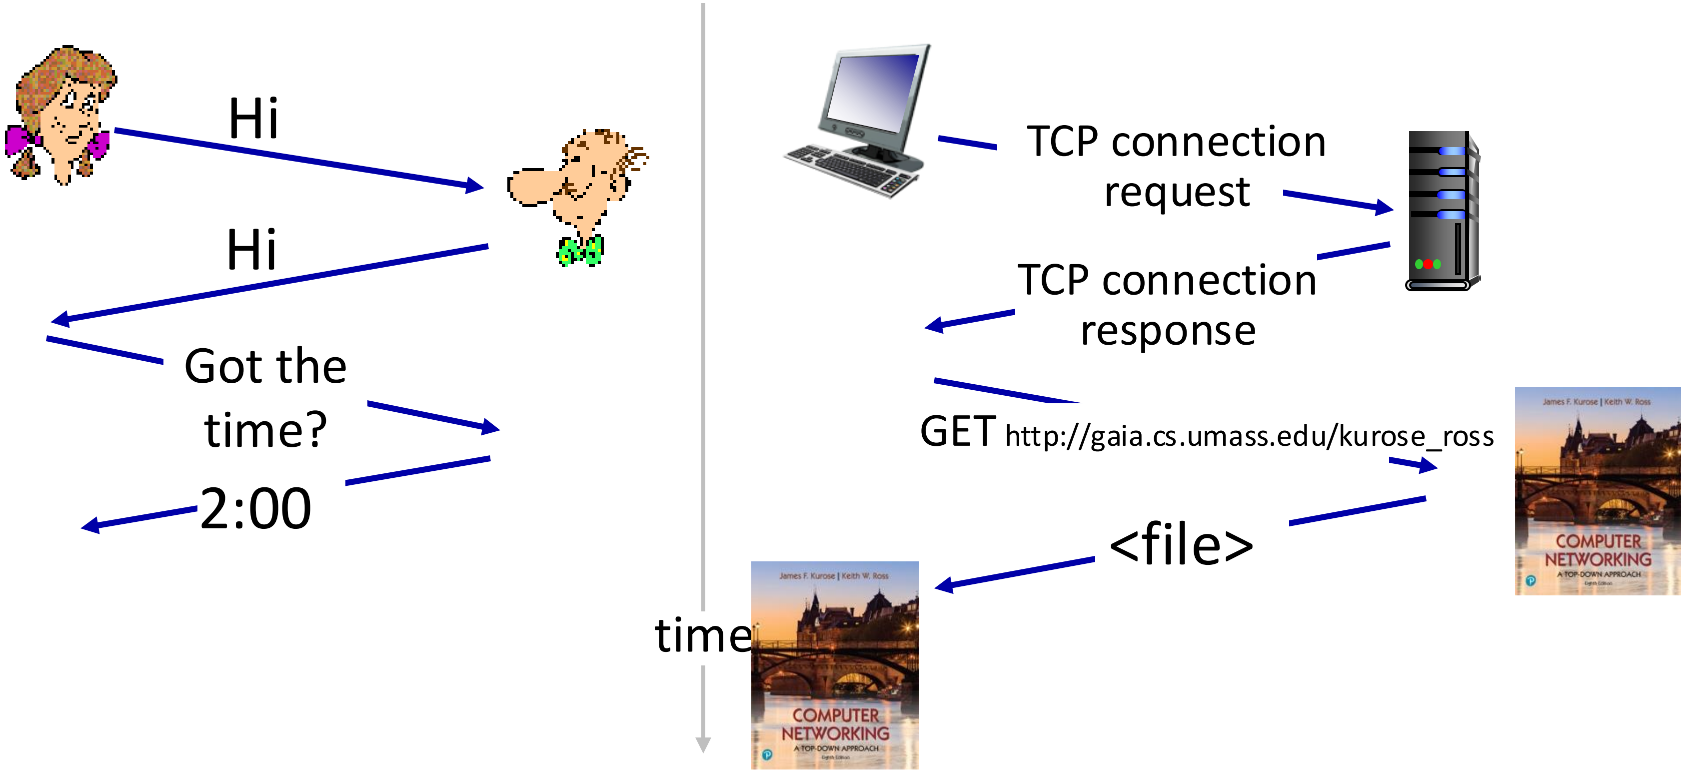

What’s a protocol?

A human protocol and a computer network protocol:

1011011010101010011011111........

TCP

Connection Request

Format of messages

What's a protocol?

A human protocol and a computer network protocol:

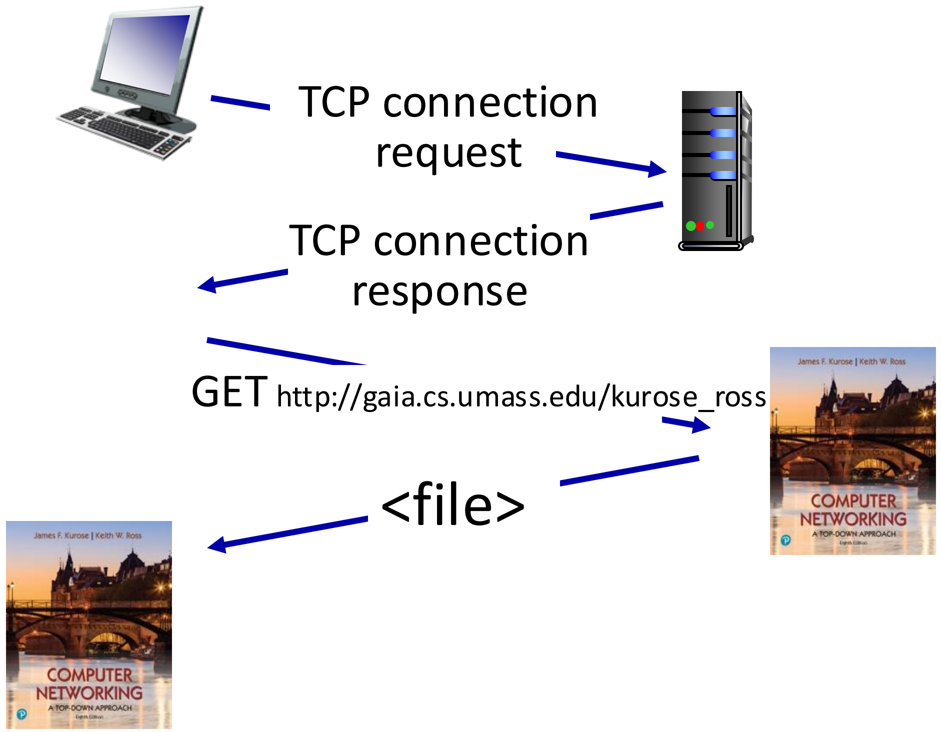

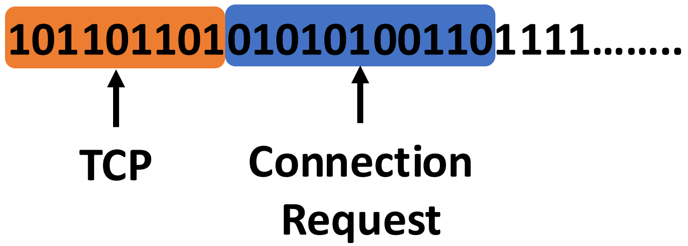

Format of messages

Order of messages

What’s a protocol?

A human protocol and a computer network protocol:

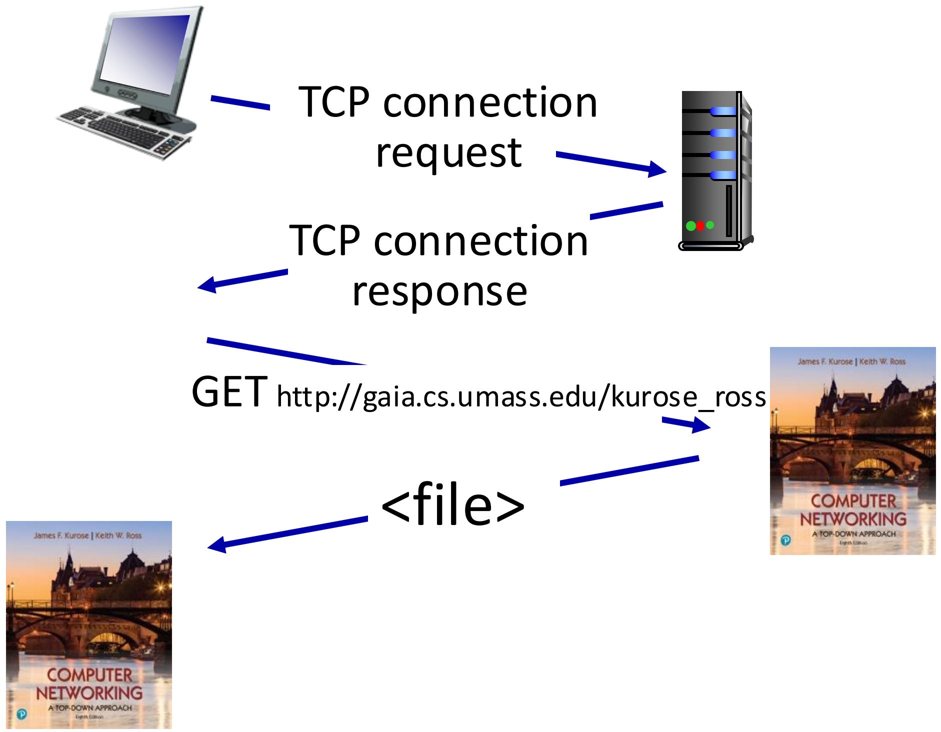

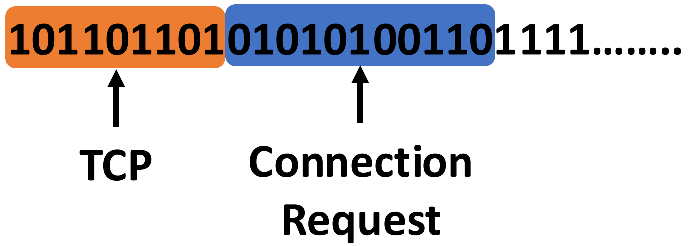

Format of messages

Order of messages

Actions actions taken on message transmission, receipt

The Internet: a "services" view

- *As an Infrastructure* that provides services to applications:
  - Web, streaming video, multimedia teleconferencing, social media,…
  - provided by hardware and software (*protocols*)
- provides *programming interface* to distributed applications:
  - "hooks" allowing sending/receiving apps to "connect" to, use Internet transport service
  - provides service options, analogous to postal service

Infrastructure

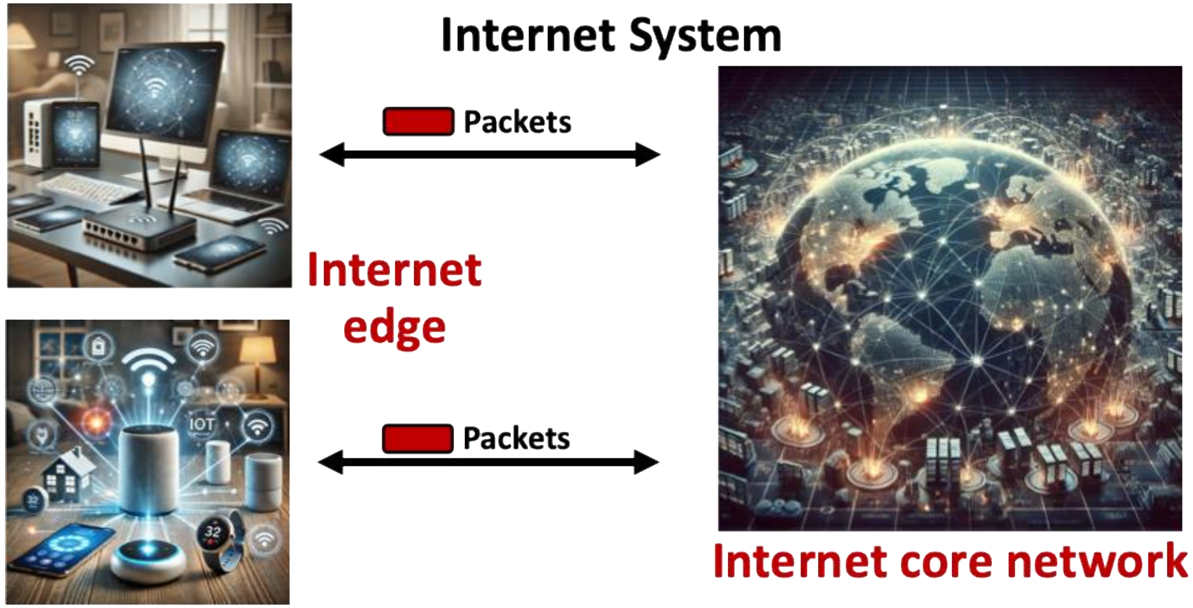

Programming Interface

What's the Internet: a service view

- services provided by protocols
  - running on hosts and routers.
- two types of services provided to apps:
  - Connectionless (UDP)
    - faster/quicker delivery (no need to set up any connection)
    - less reliable, no orderly packets delivered
    - Suitable for real-time streaming
  - Connection-oriented (TCP)
    - Suitable for file/email transfers

Chapter 1: roadmap

- What *is* the Internet?
- What *is* a protocol?
- Network edge: hosts, access network, physical media
- Network core: packet/circuit switching, internet structure
- Performance: loss, delay, throughput
- Security
- Protocol layers, service models
- History

A closer look at Internet structure

Network edge:

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

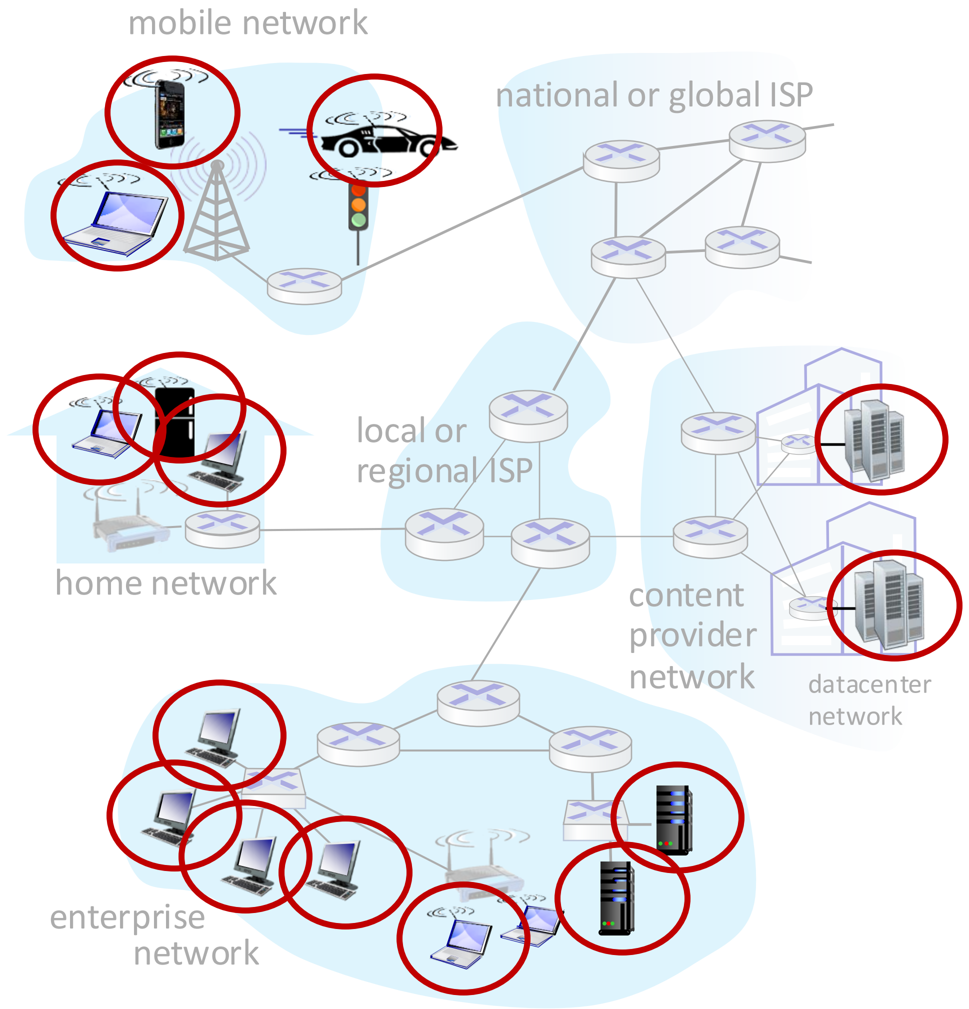

A closer look at Internet structure

Network edge:

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

Access networks, physical media:

- wired, wireless communication links

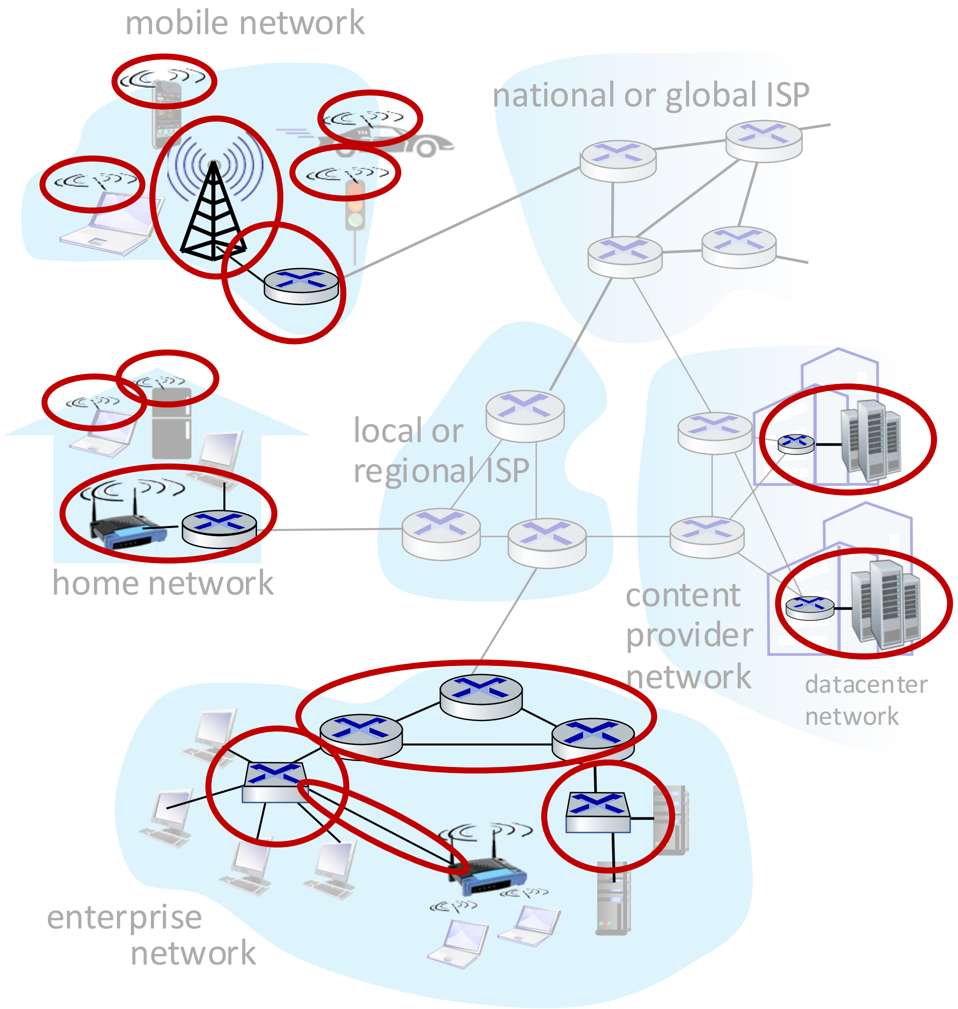

A closer look at Internet structure

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

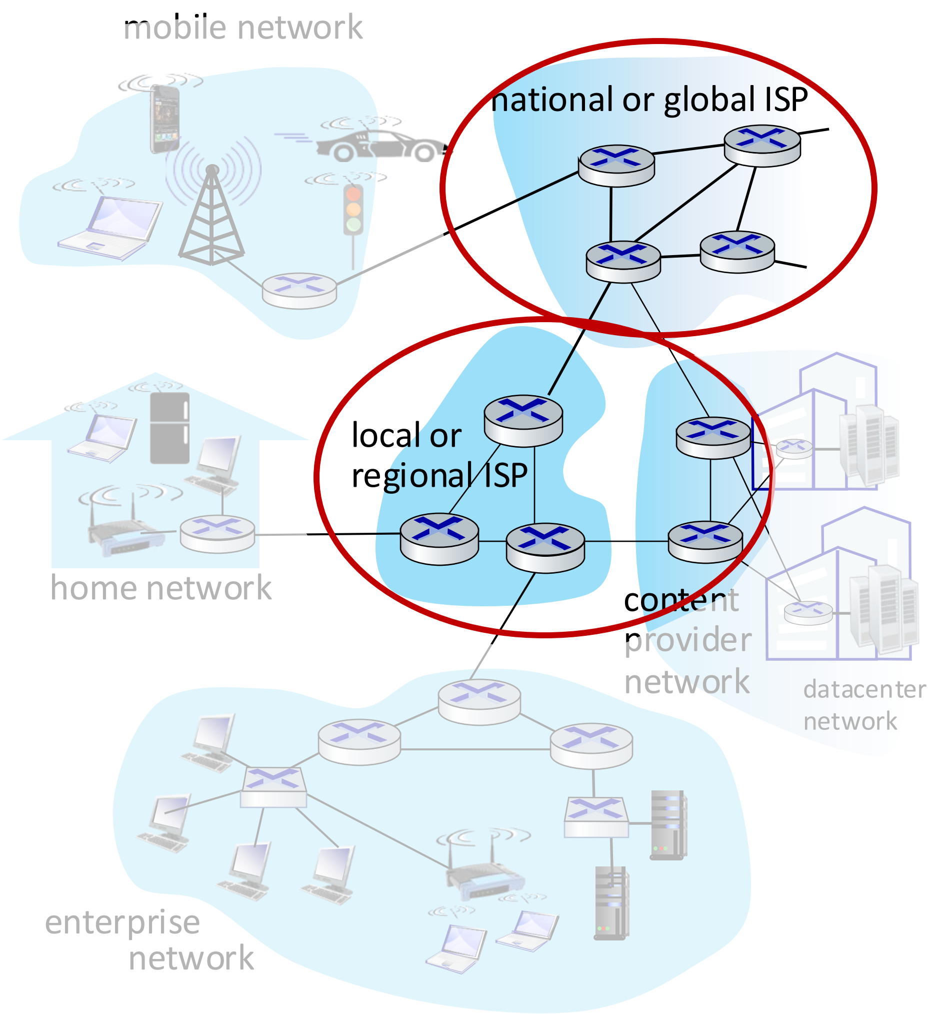

Access networks and physical media

Q: How to connect end systems to edge router?

- residential access nets
- institutional access networks (school, company)
- mobile access networks (WiFi, 4G/5G)

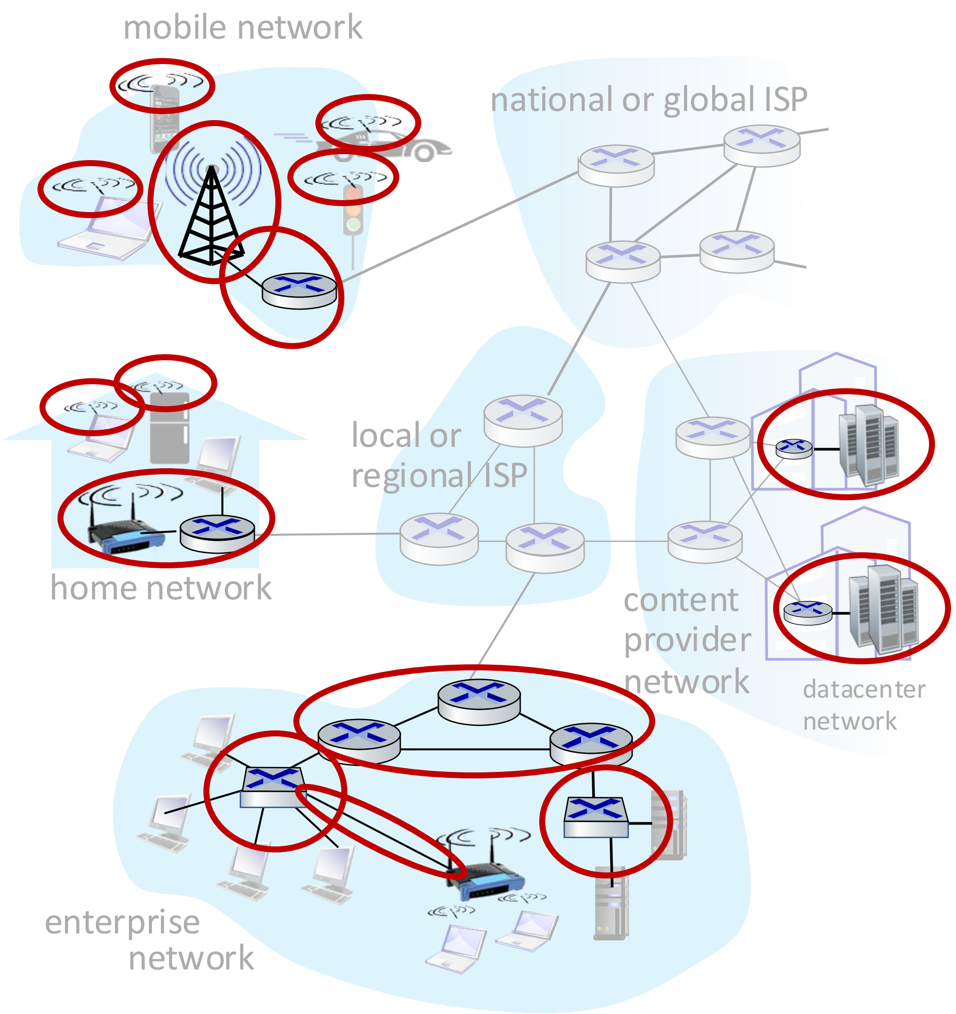

Access networks: cable-based access

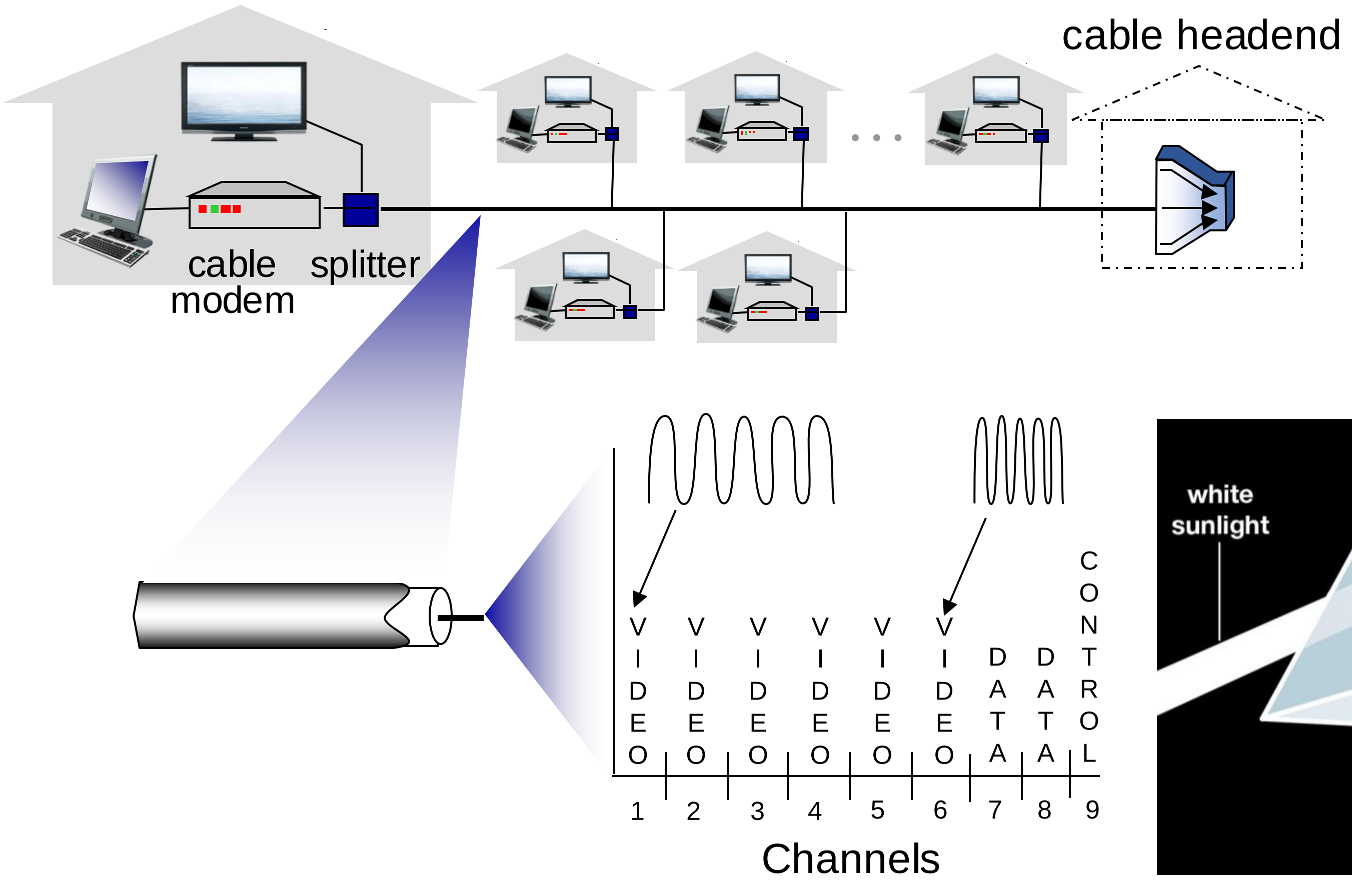

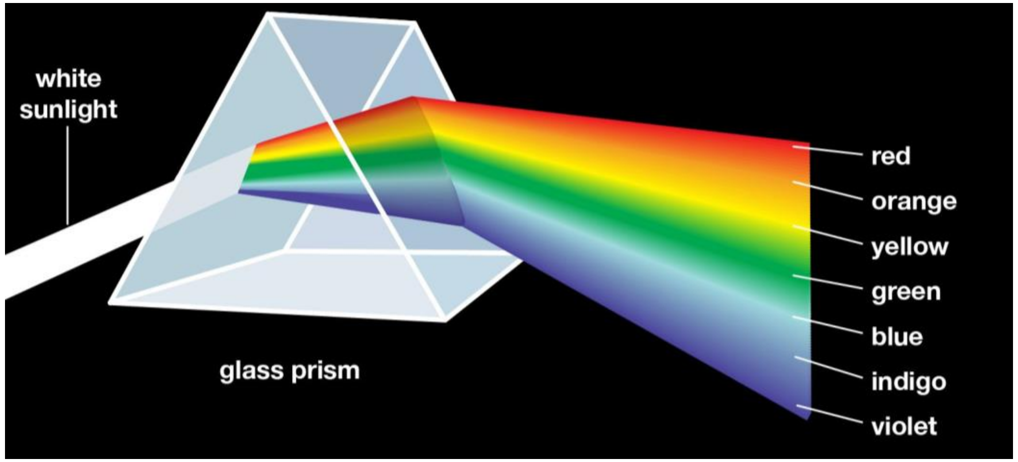

frequency division multiplexing (FDM): different channels transmitted in different frequency bands

Access networks: home networks

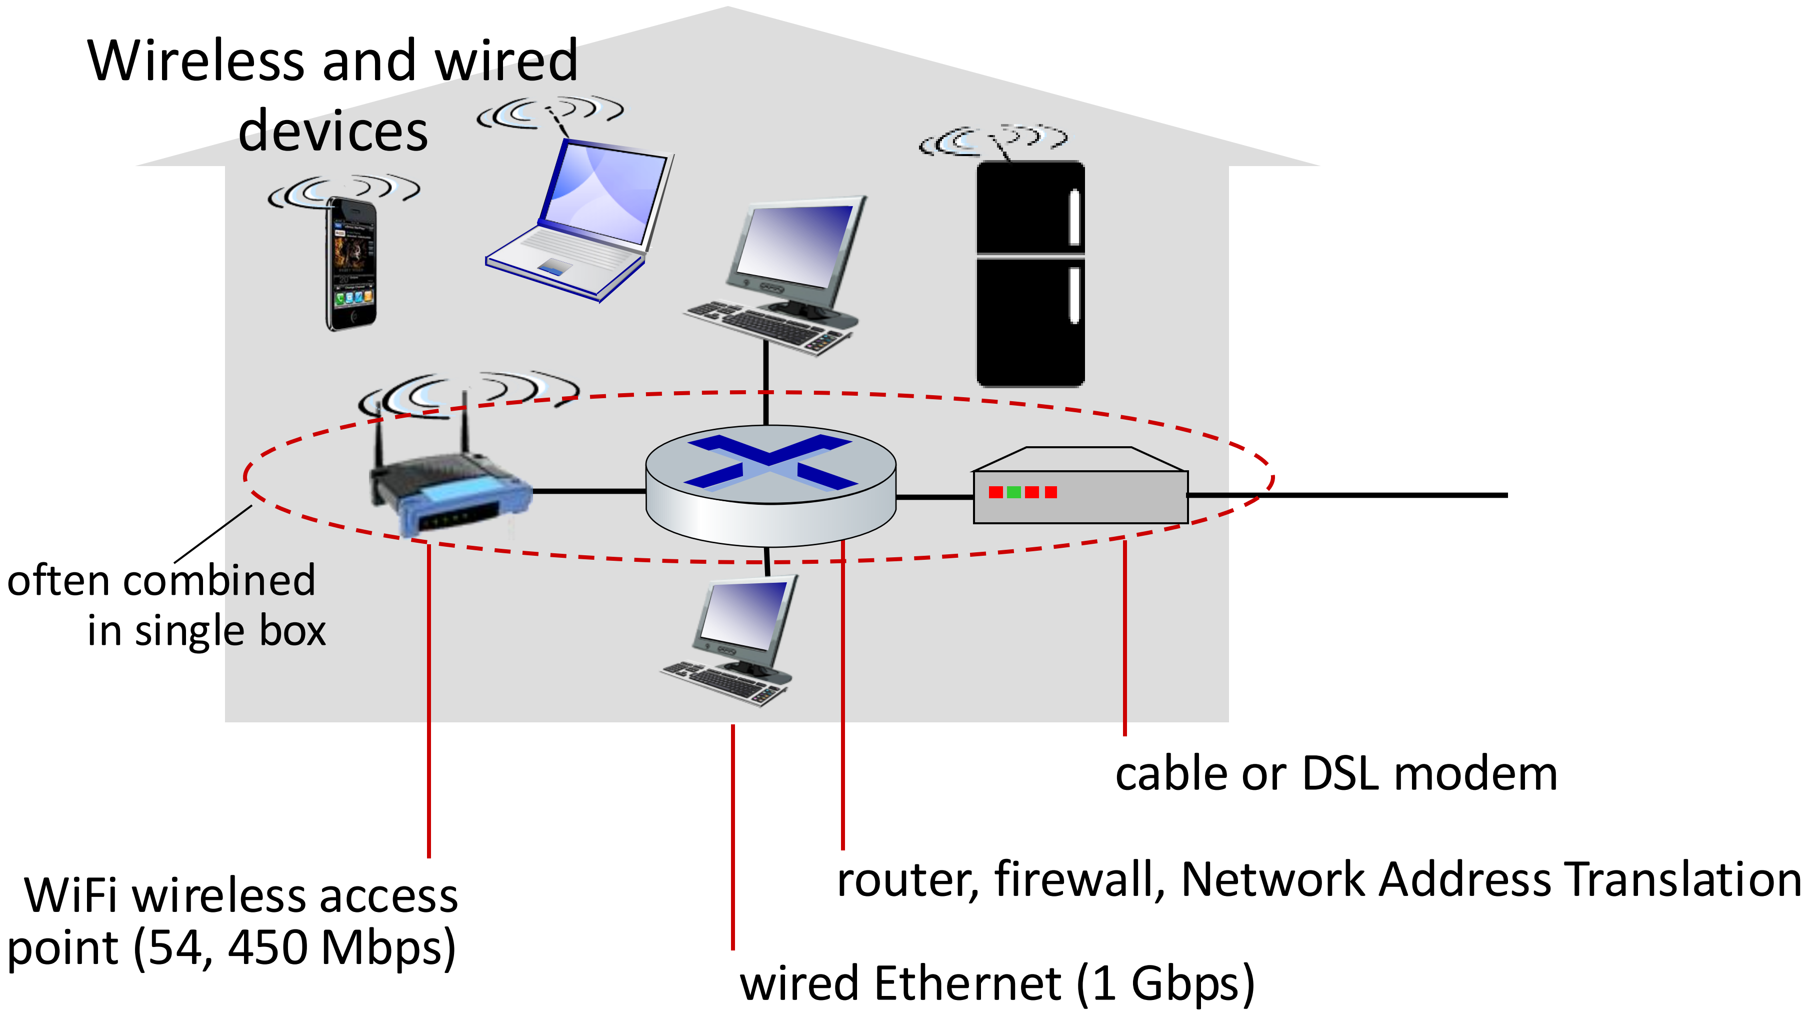

Wireless access networks

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

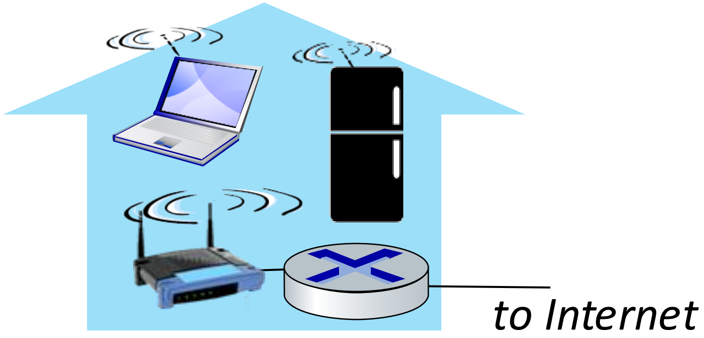

Wide-area cellular access networks

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

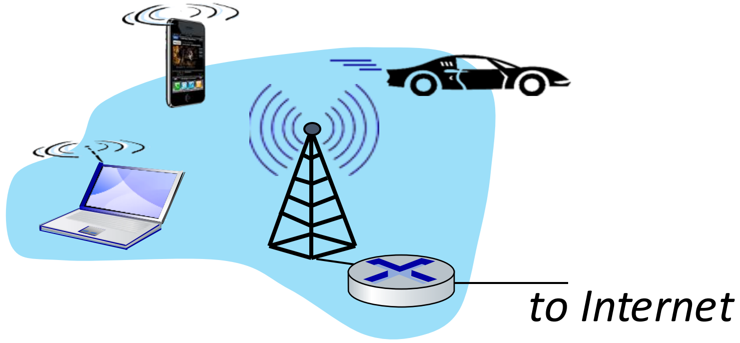

Access networks: enterprise networks

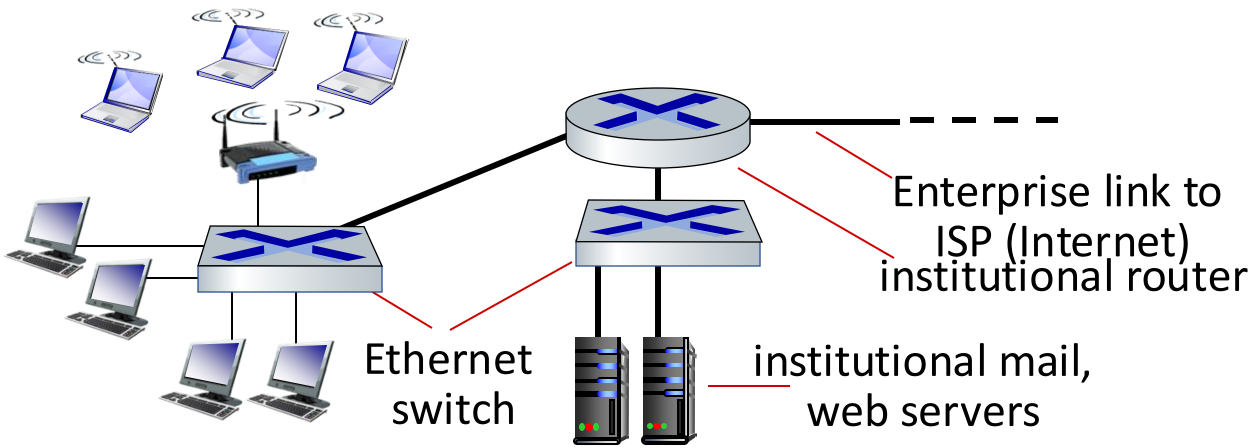

- companies, universities, etc.
- mix of wired, wireless link technologies, connecting a mix of switches and routers (we'll cover differences shortly)
  - Ethernet: wired access at 100Mbps, 1Gbps, 10Gbps
  - WiFi: wireless access points at 11, 54, 450 Mbps

Access networks: data center networks

- high-bandwidth links (10s to 100s Gbps) connect hundreds to thousands of servers together, and to Internet

Courtesy: Massachusetts Green High Performance Computing Center (mghpcc.org)

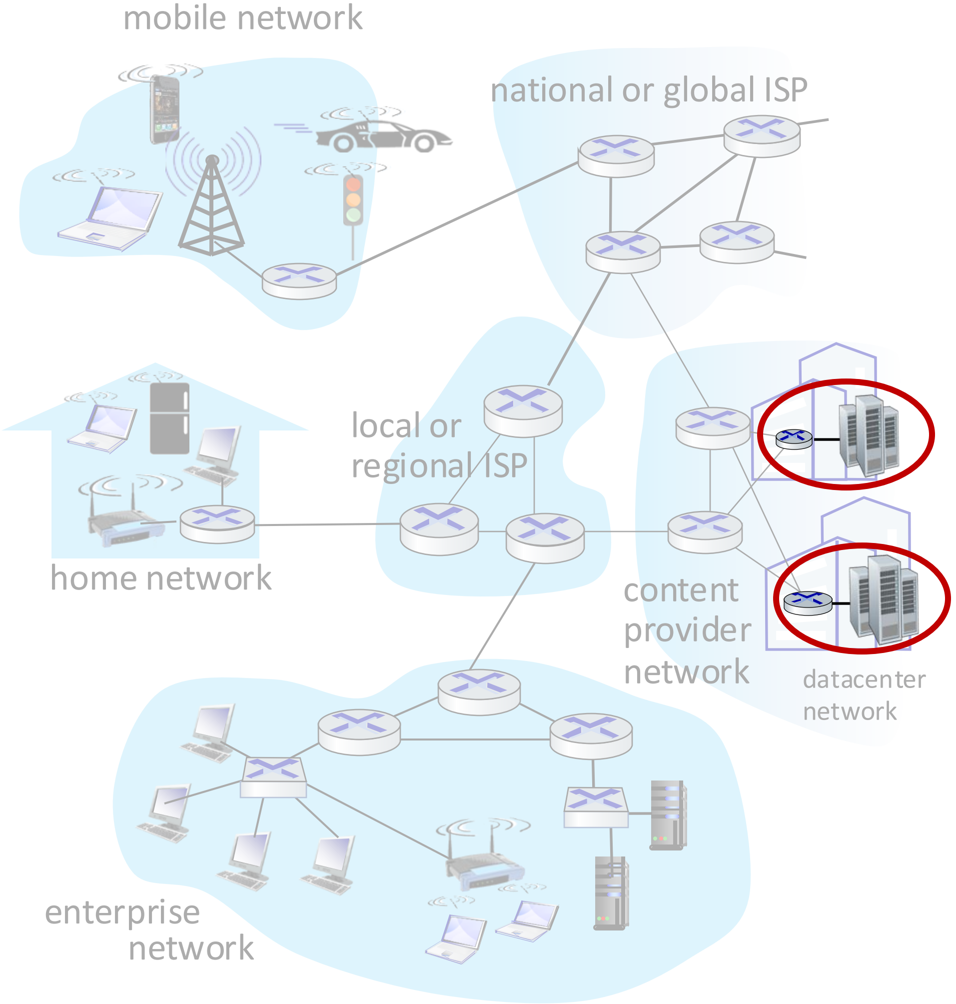

Host: sends *packets* of data

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*

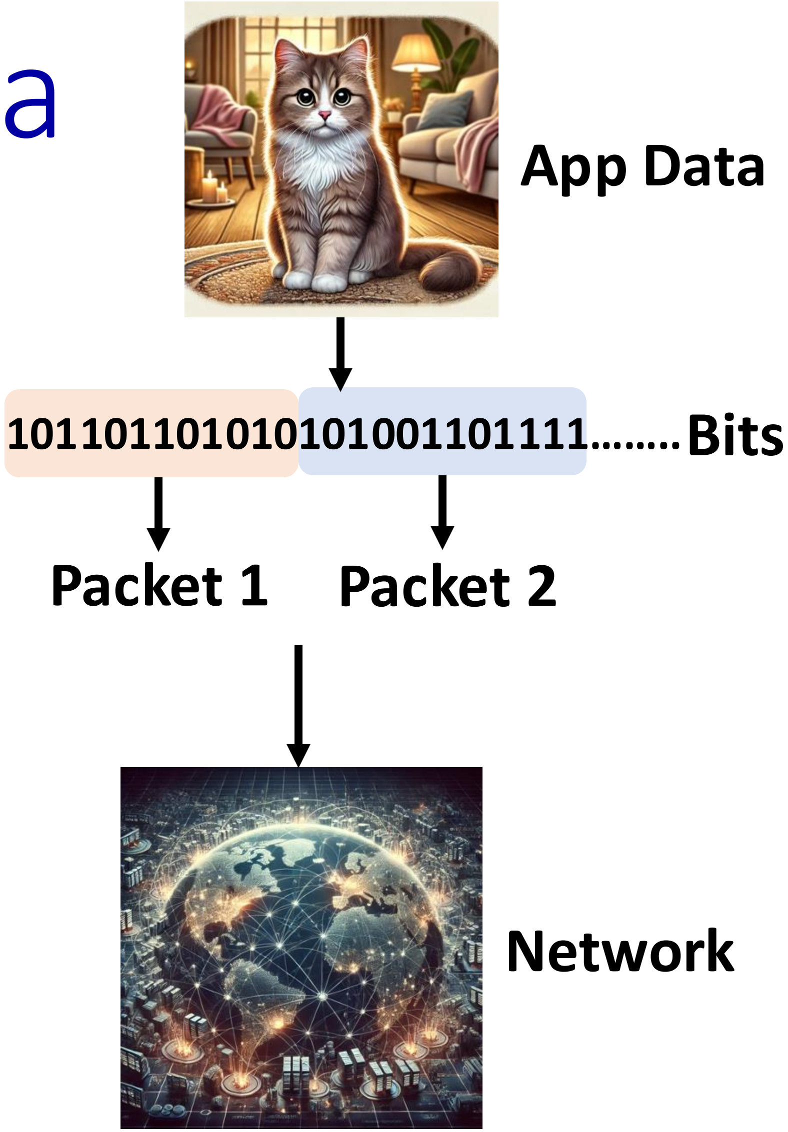

Host: sends *packets* of data

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*

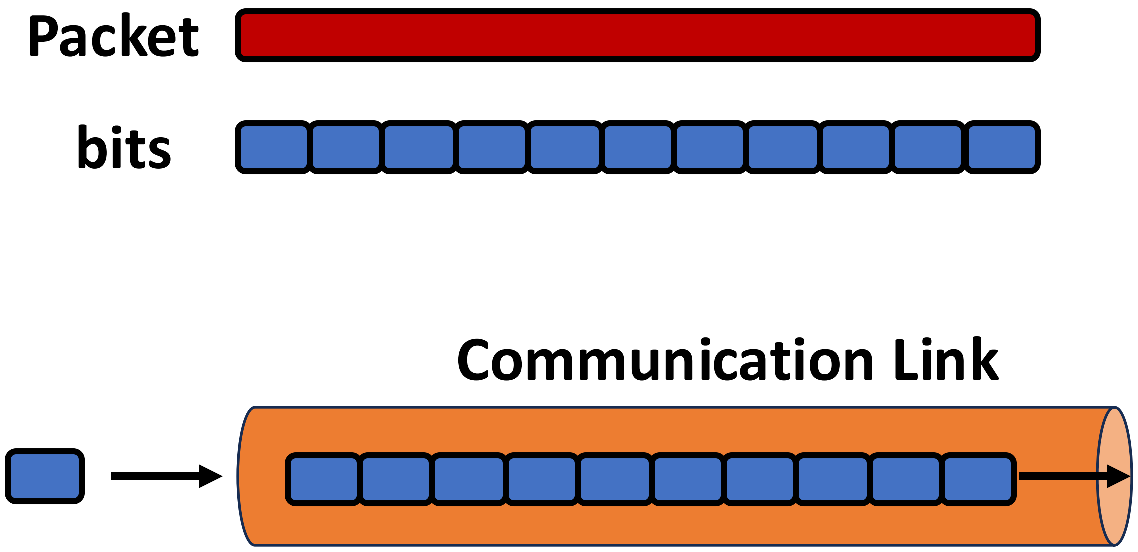

What's the transmission rate R, link capacity or link bandwidth?

Host: sends *packets* of data

Link *transmission rate R*, *aka Link Capacity*, *aka link bandwidth*

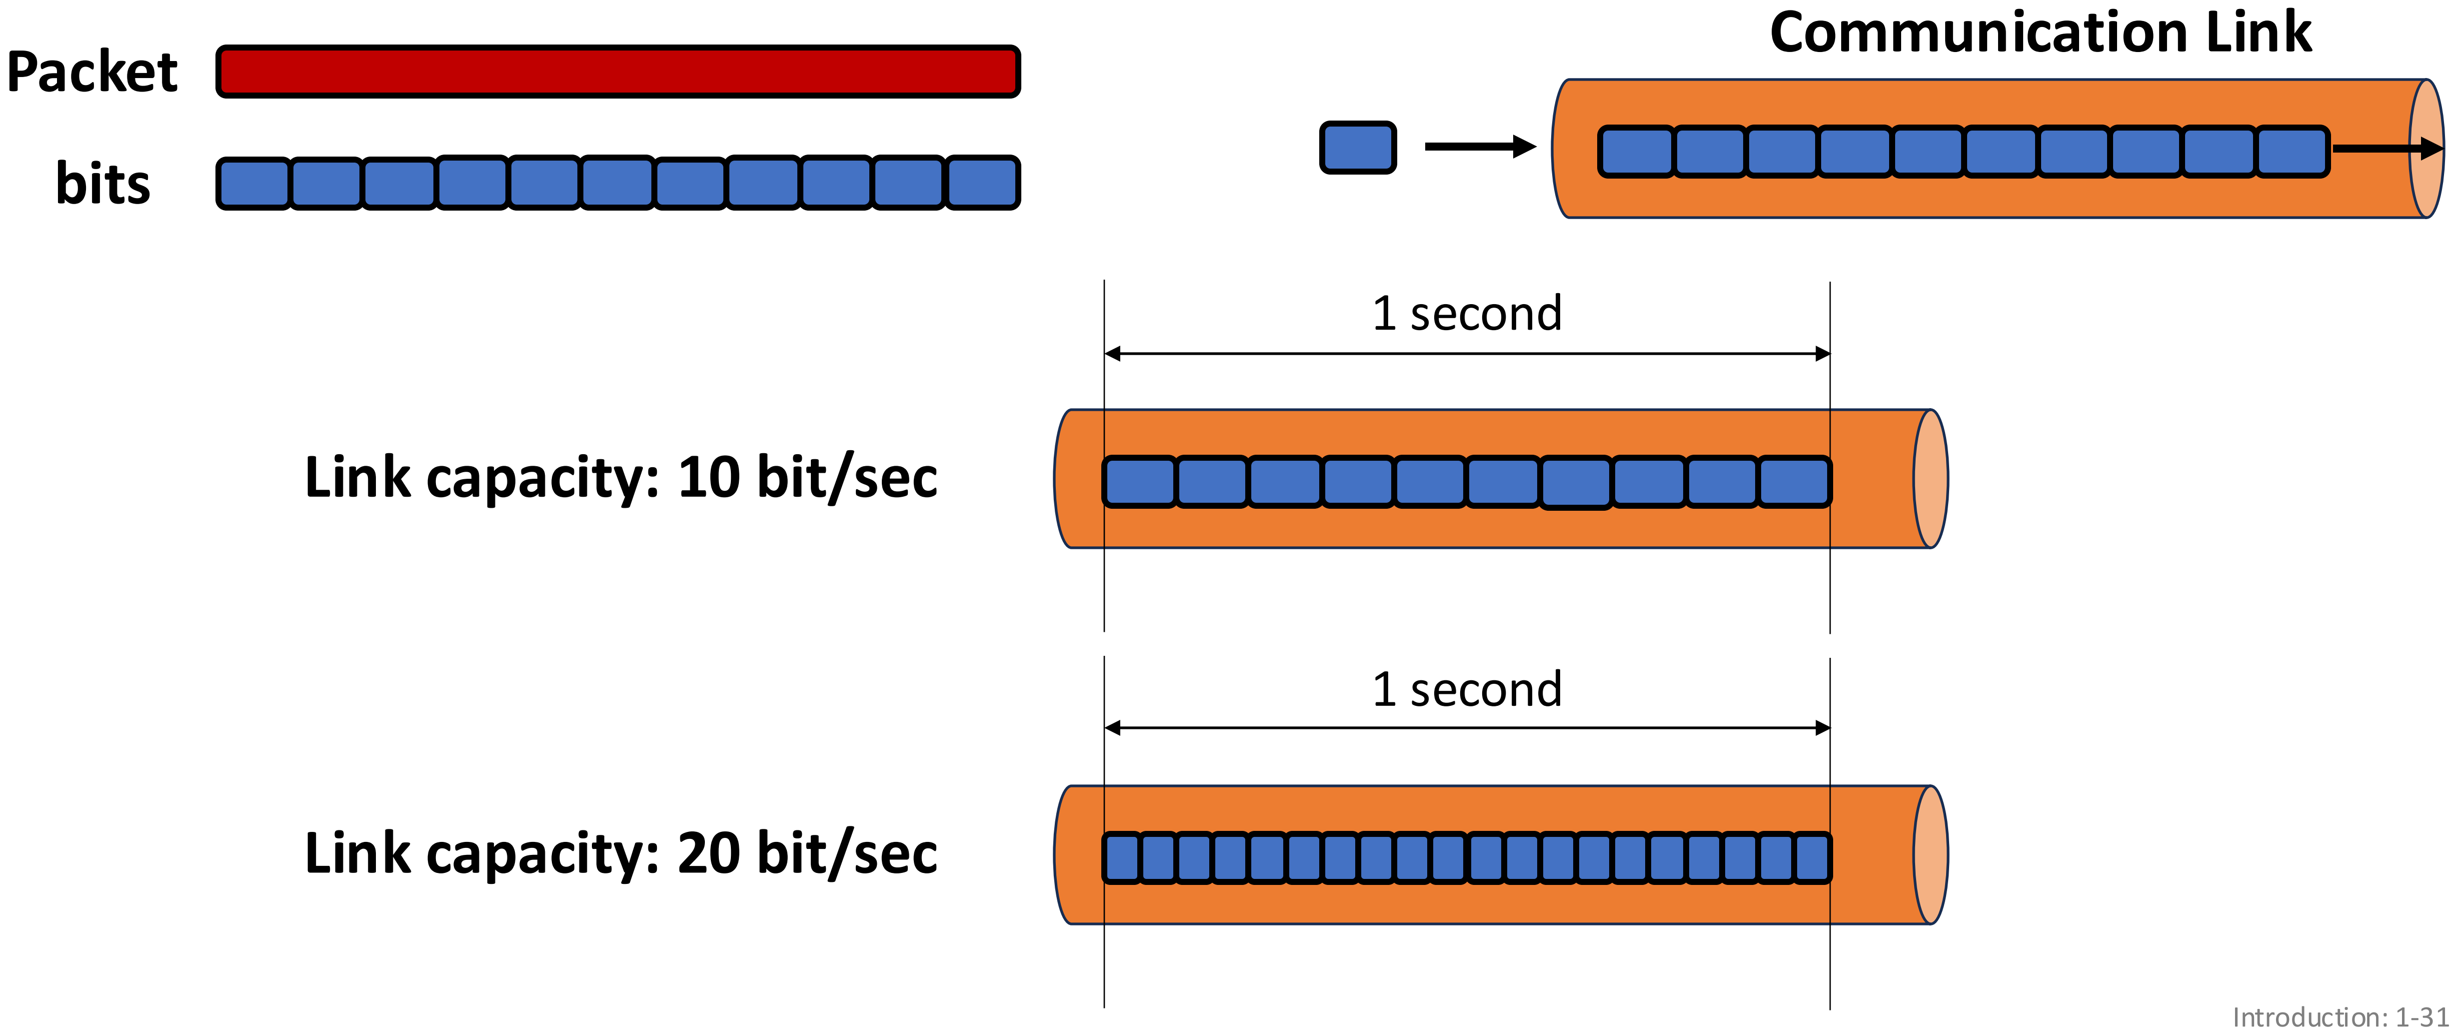

Host: sends *packets* of data

Link *transmission rate R, aka Link Capacity, aka link bandwidth*

Packet transmission delay

- How long it takes for transmitting all the bits into the network or communication link

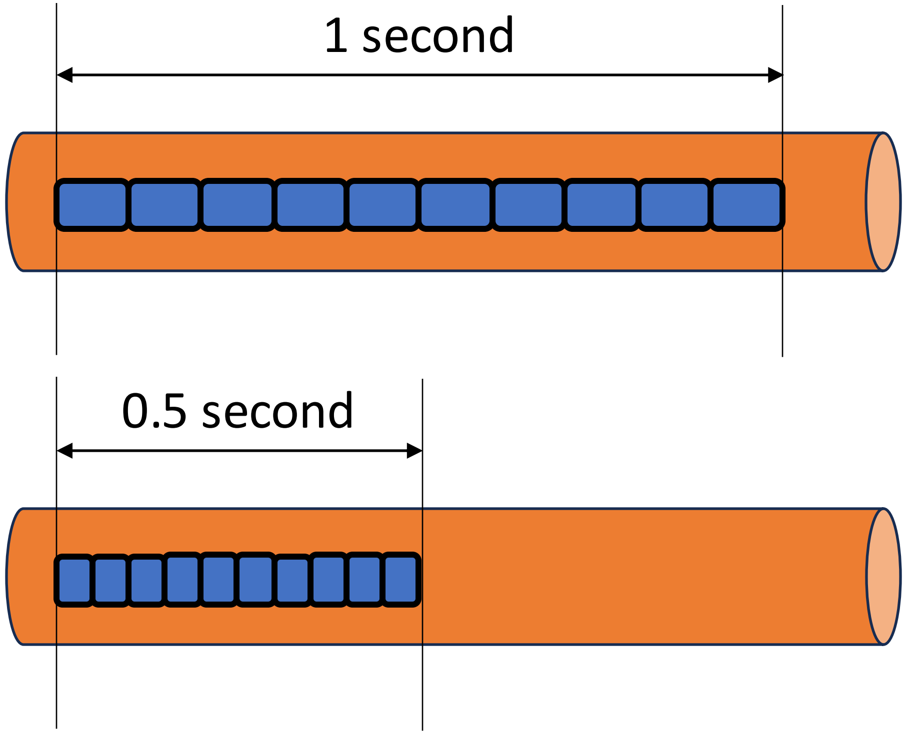

A packet with 10 bits

10 bits

$$\text{packet transmission delay} = \text{time needed to transmit } L\text{-bit packet into link} = \frac{L \text{ (bits)}}{R \text{ (bits/sec)}}$$

Communication Links: physical media

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

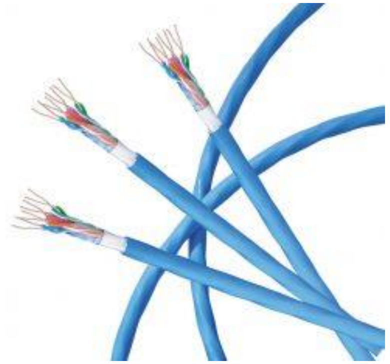

Communication Links: physical media

Coaxial cable:

- two concentric copper conductors
- bidirectional
- broadband:
  - multiple frequency channels on cable
  - 100's Mbps per channel

Fiber optic cable:

- glass fiber carrying light pulses, each pulse a bit
- high-speed operation:
  - high-speed point-to-point transmission (10's-100's Gbps)
- low error rate:
  - repeaters spaced far apart
  - immune to electromagnetic noise

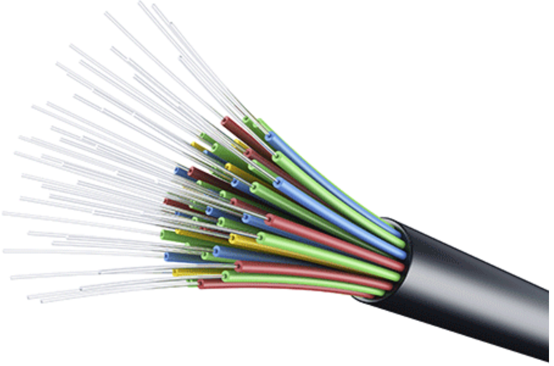

Links: physical media

Wireless radio

- signal carried in various "bands" in wireless spectrum
- no physical "wire"
- broadcast, "half-duplex" (sender to receiver)
- propagation environment effects:
  - reflection
  - obstruction by objects
  - Interference/noise

Links: physical media

Wireless radio

- signal carried in various "bands" in wireless spectrum
- no physical "wire"
- broadcast, "half-duplex" (sender to receiver)
- propagation environment effects:
  - reflection
  - obstruction by objects
  - Interference/noise

Radio link types:

- Wireless LAN (WiFi)
  - 10-100's Mbps; 10's of meters
- wide-area (e.g., 4G cellular)
  - 10's Mbps over ~10 Km
- Bluetooth: cable replacement
  - short distances, limited rates
- satellite
  - up to 45 Mbps per channel
  - 270 msec end-end delay

Wireless Spectrum

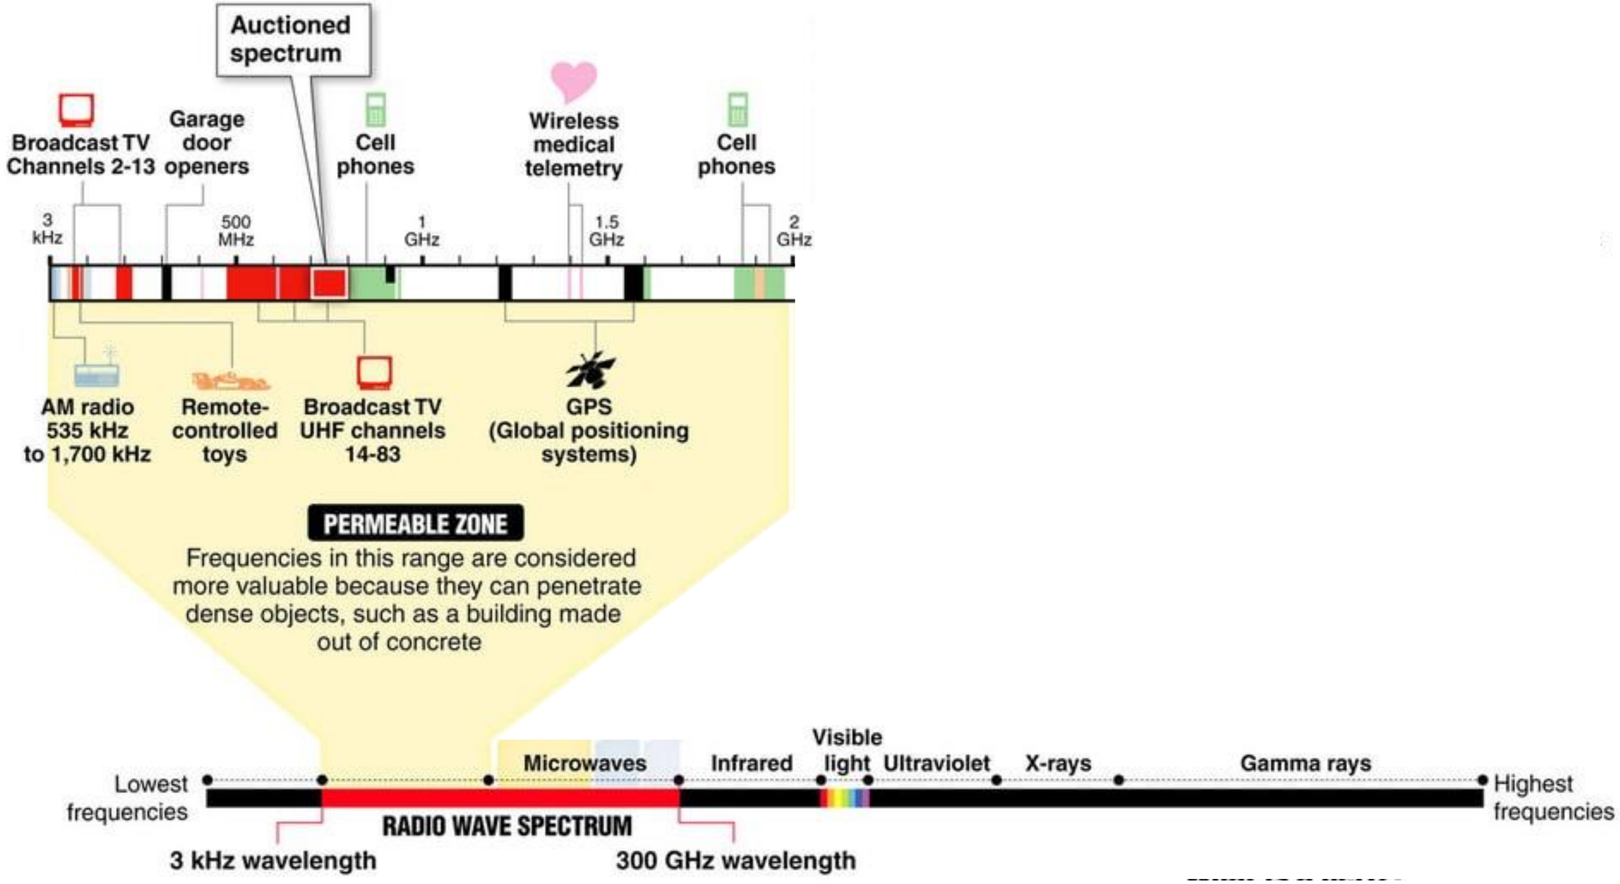

Wireless Spectrum

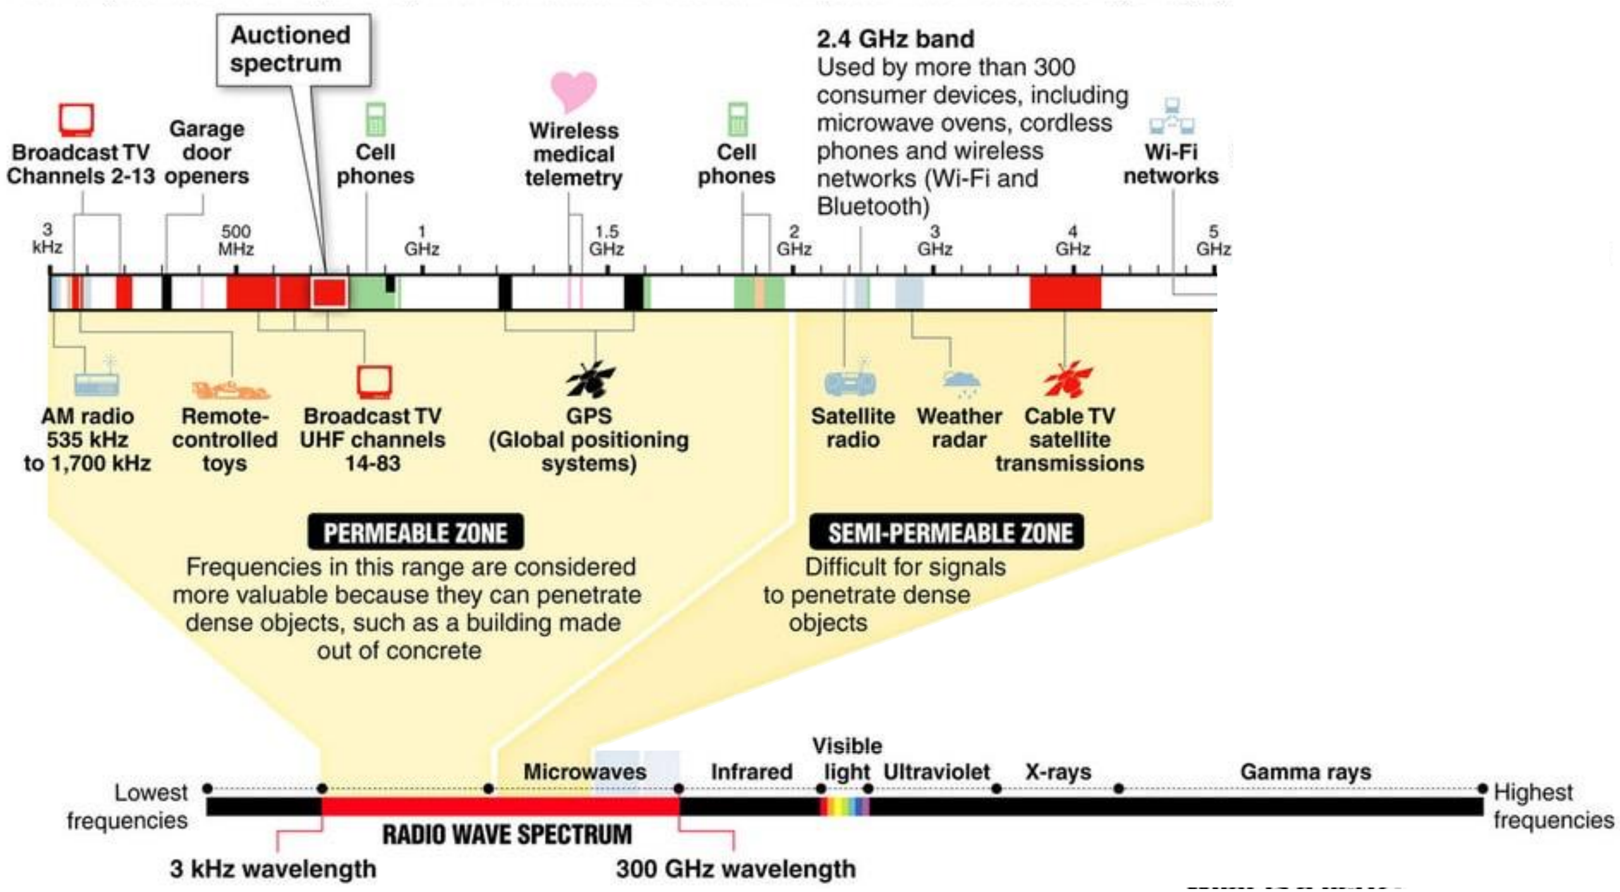

Wireless Spectrum

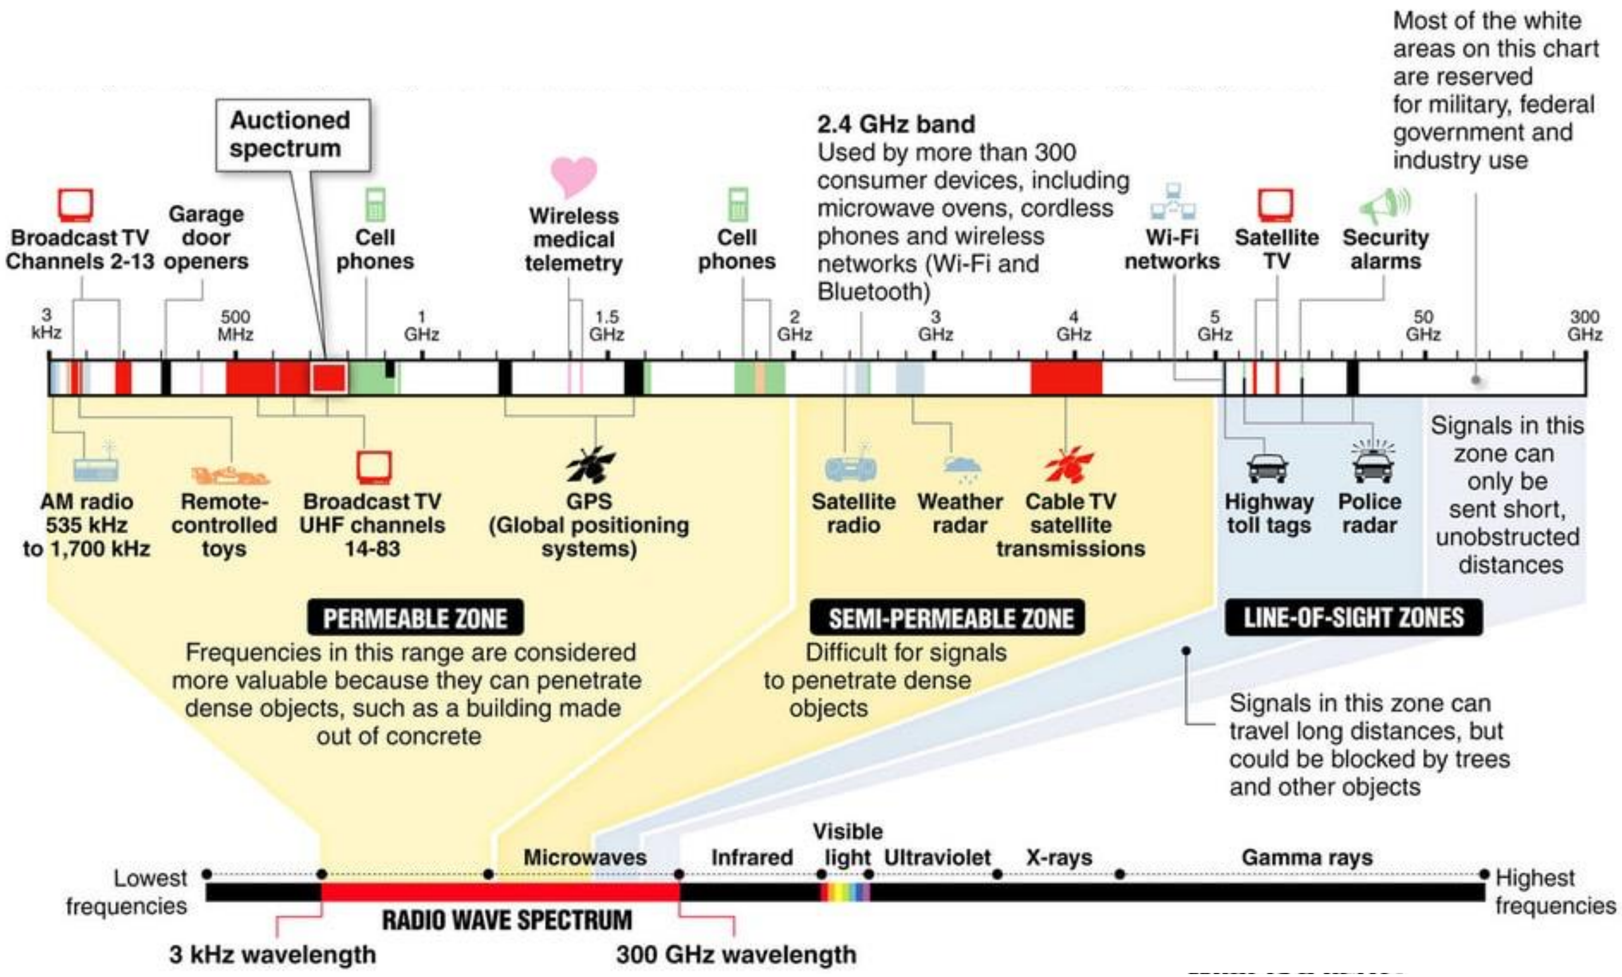

Wireless Spectrum

MID-BAND SPECTRUM AUCTION —

Verizon and AT&T dominate spectrum auction, spending combined $69 billion

Top two carriers buy licenses nationwide, outspending T-Mobile and US Cellular.

JON BRODKIN - 2/25/2021, 11:14 AM

Verizon and AT&T dominated the US government's latest spectrum auction, spending a combined $68.9 billion on licenses in the upper 3GHz band.

Verizon's winning bids totaled $45.45 billion, while AT&T's came in at $23.41 billion. T-Mobile was third with $9.34 billion as the three biggest wireless carriers accounted for the vast majority of the $81.17 billion in winning bids, the Federal Communications Commission said in results released yesterday. US Cellular, a regional carrier, was a distant fourth in spending, at $1.28 billion, but came in third, ahead of T-Mobile, in the number of licenses won.

Chapter 1: roadmap

- What *is* the Internet?
- What *is* a protocol?
- Network edge: hosts, access network, physical media
- Network core: internet structure, routing and forwarding
- Performance: loss, delay, throughput
- Security
- Protocol layers, service models
- History

The network core

- mesh of interconnected routers

Ok, but WHY?

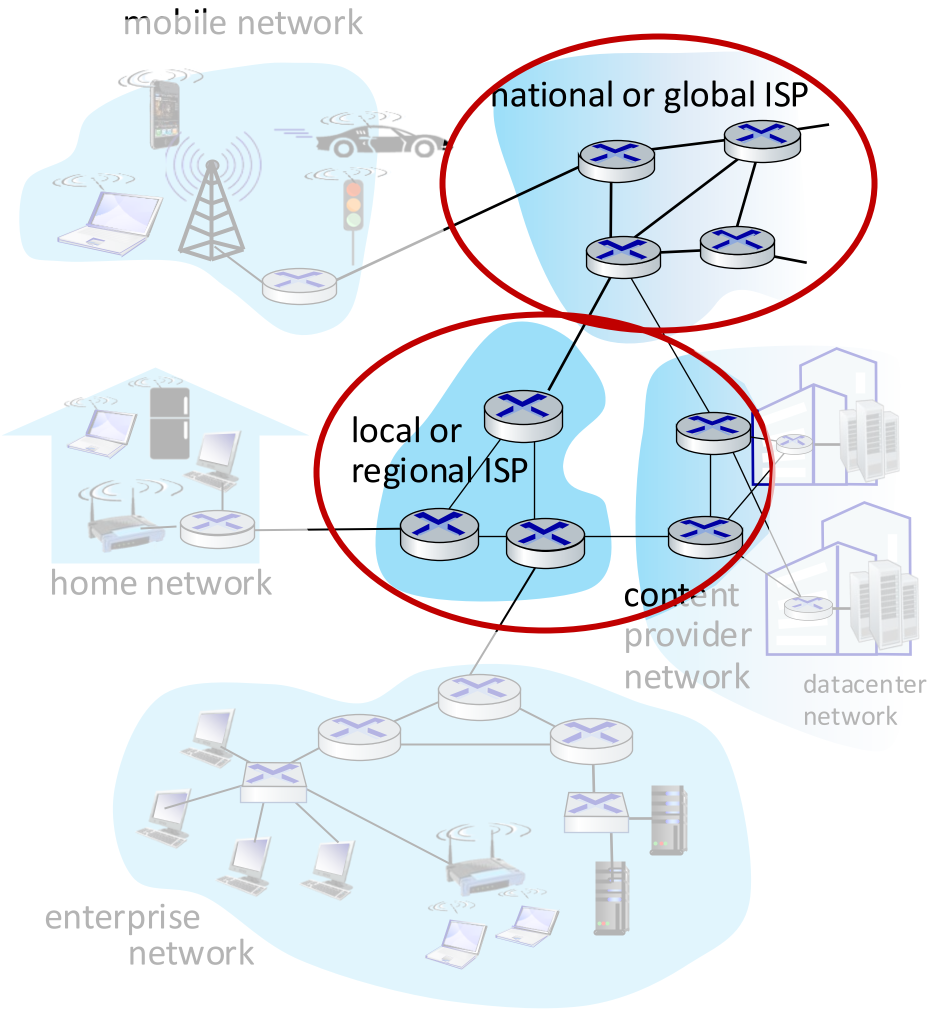

The network core

- mesh of interconnected routers

The network core

- mesh of interconnected routers

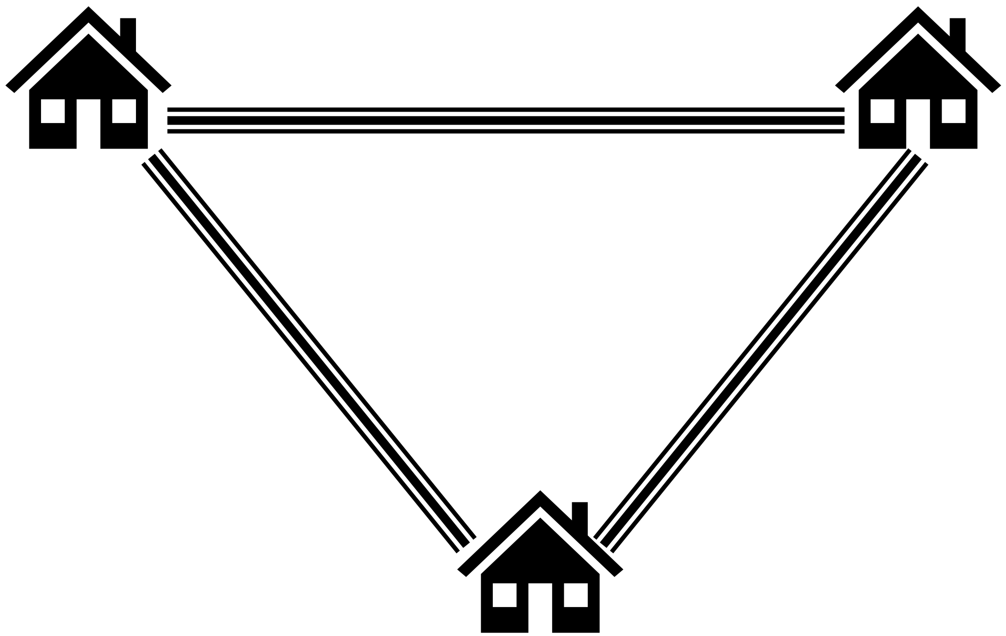

The network core

- mesh of interconnected routers

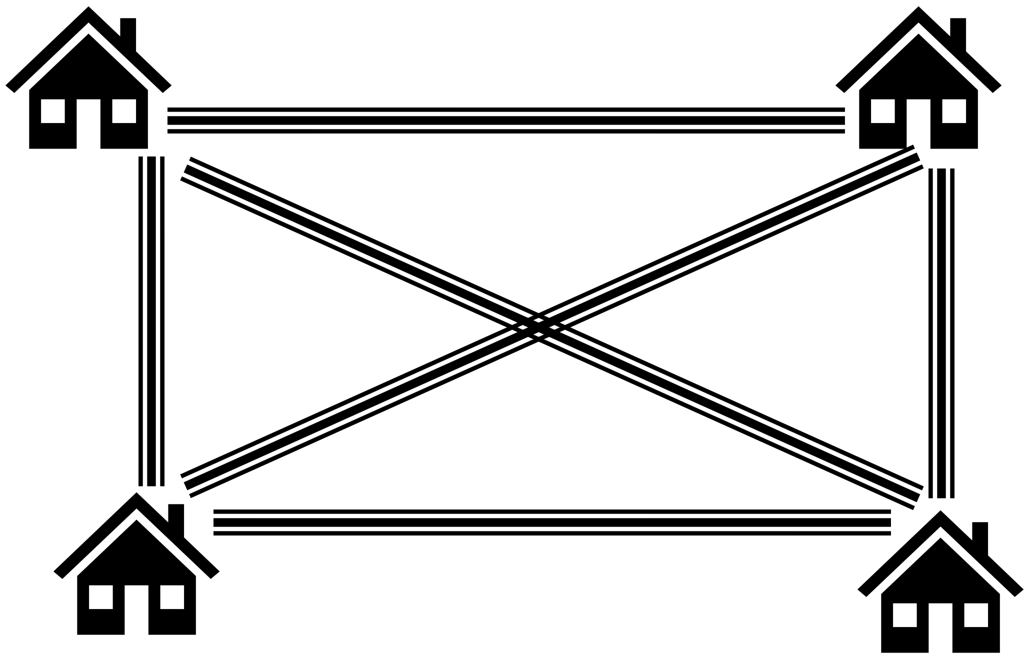

The network core

- mesh of interconnected routers

Router

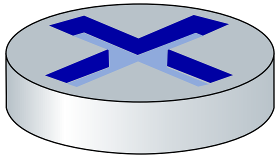

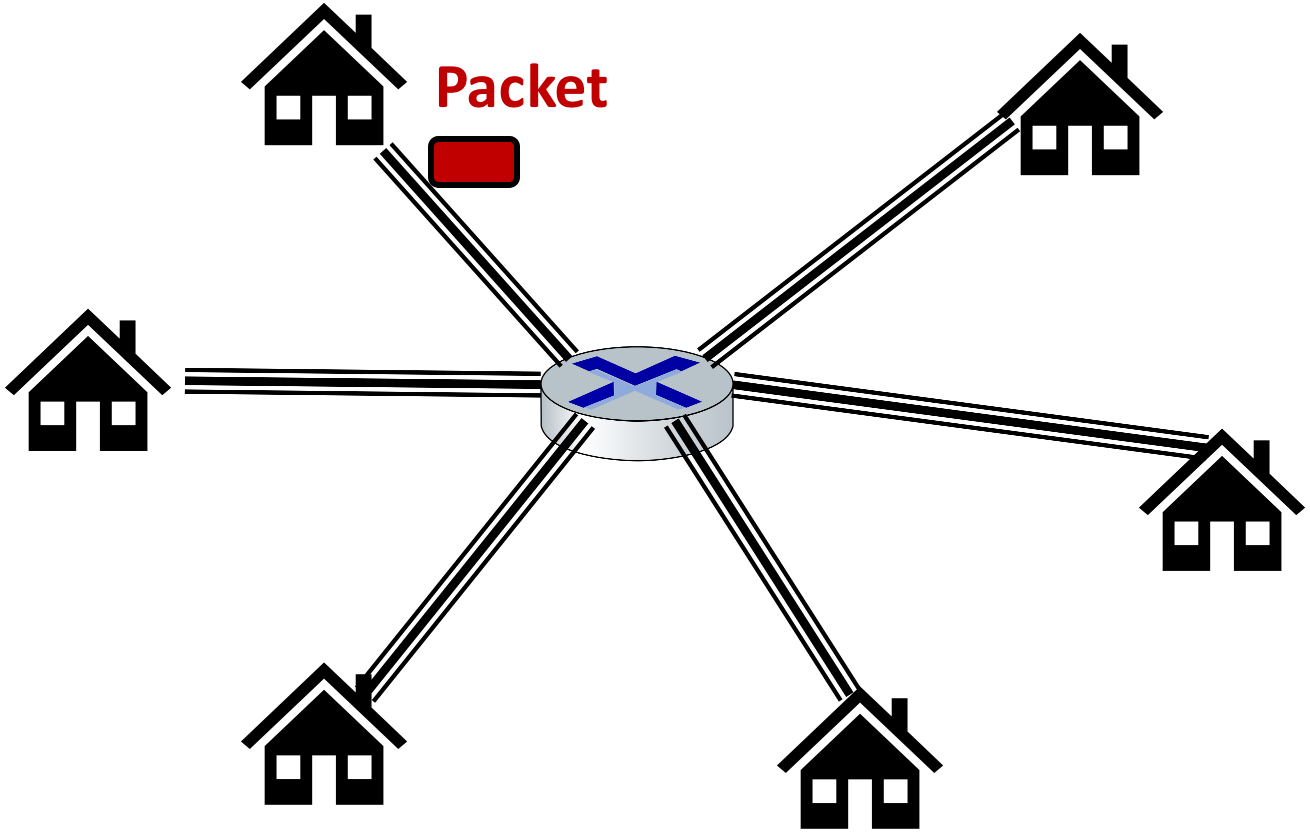

The network core

- mesh of interconnected routers

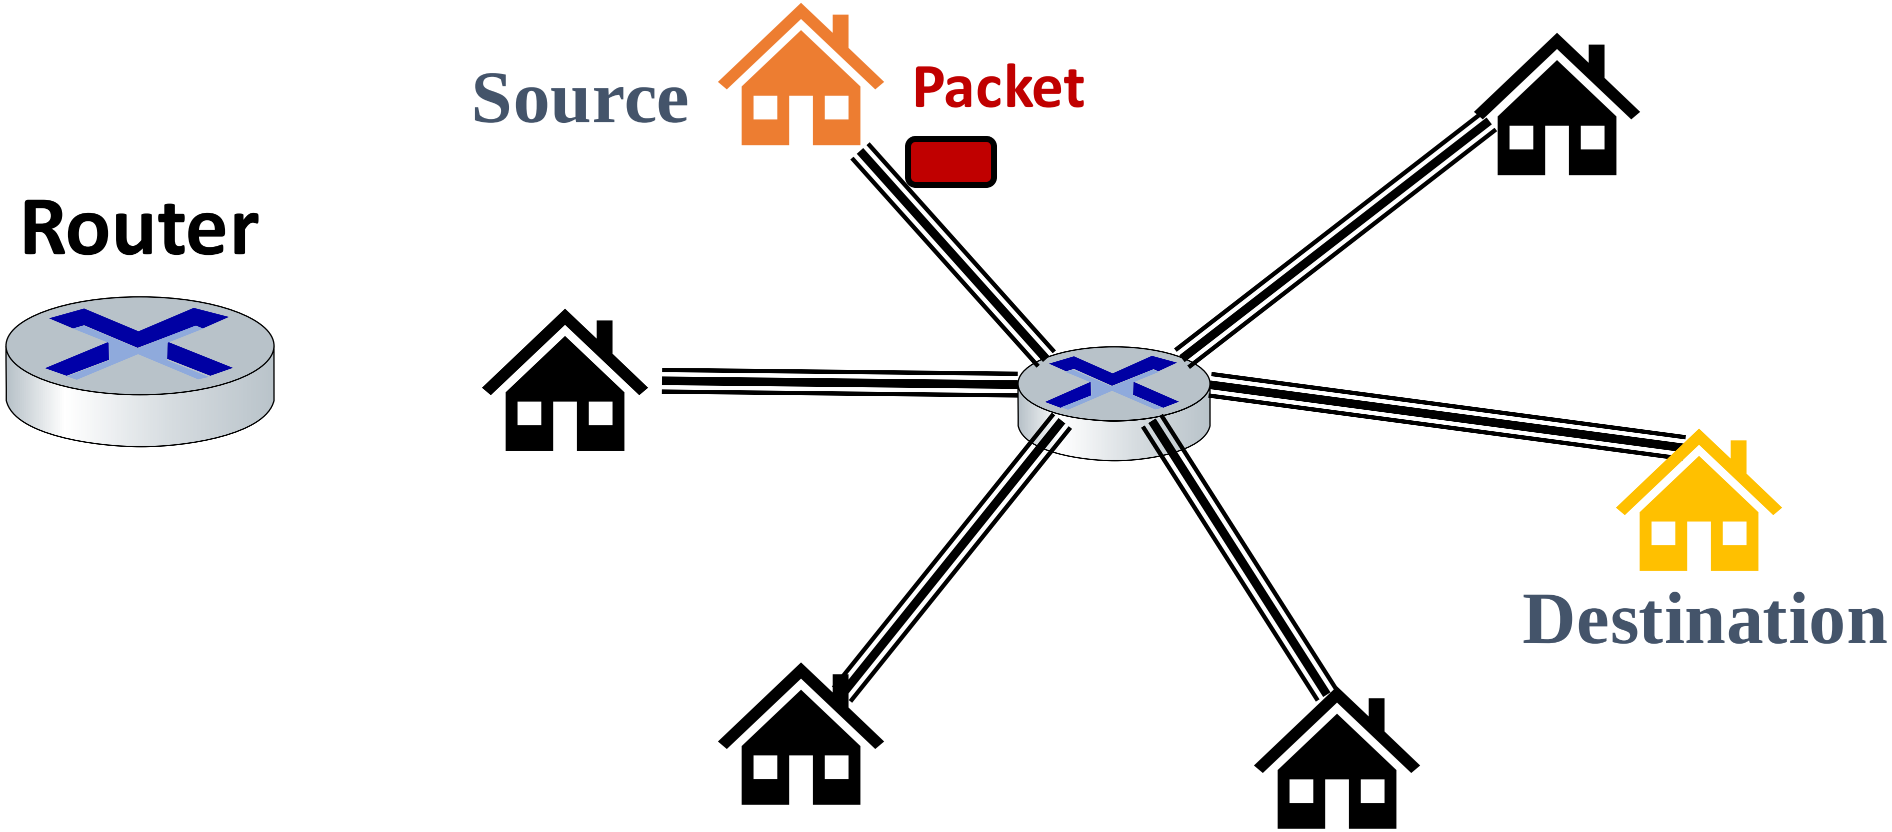

The network core

- mesh of interconnected routers

Router

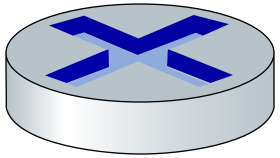

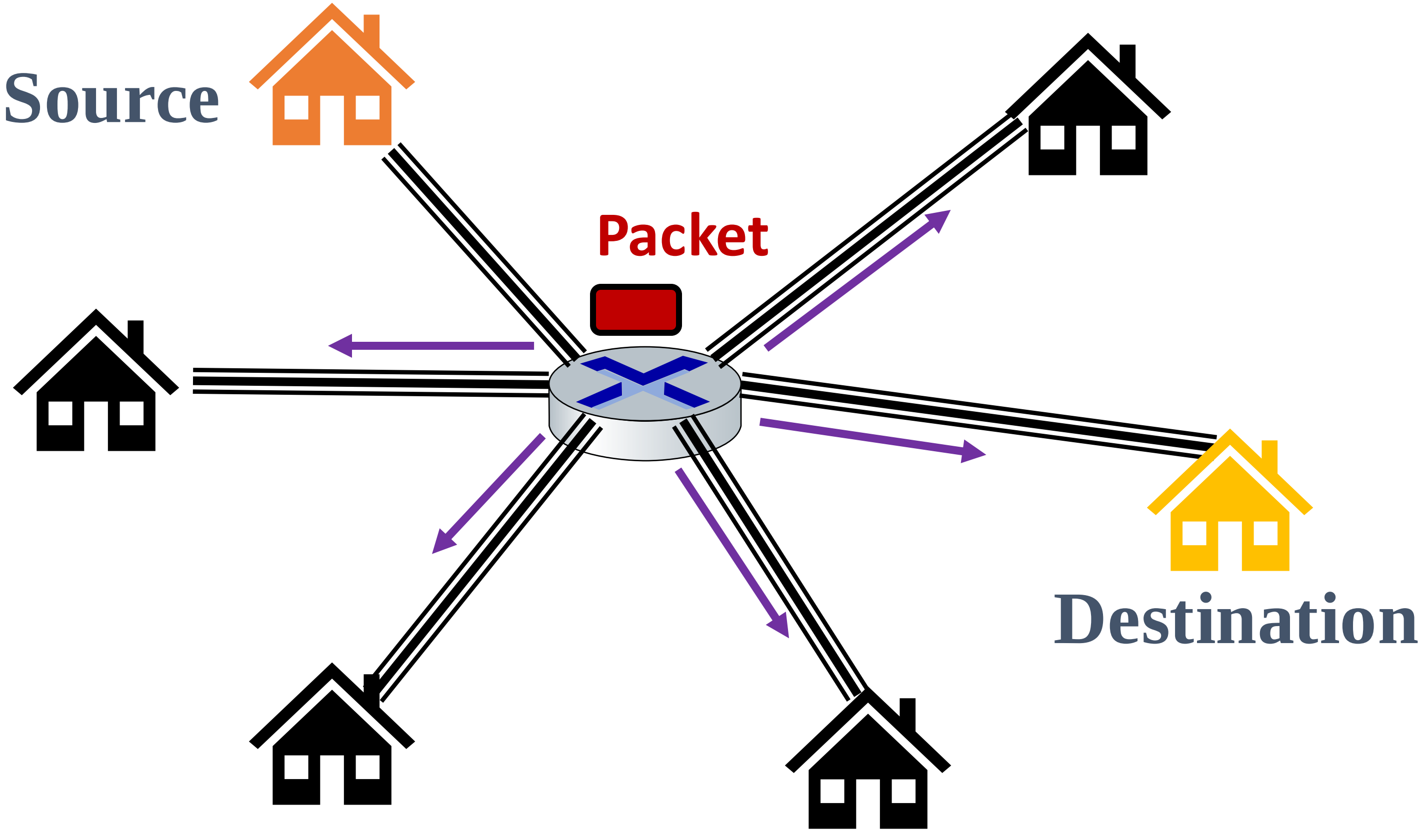

Forwarding:

- aka “switching”
- *local* action: move arriving packets from router’s input link to appropriate router output link

The network core

- mesh of interconnected routers

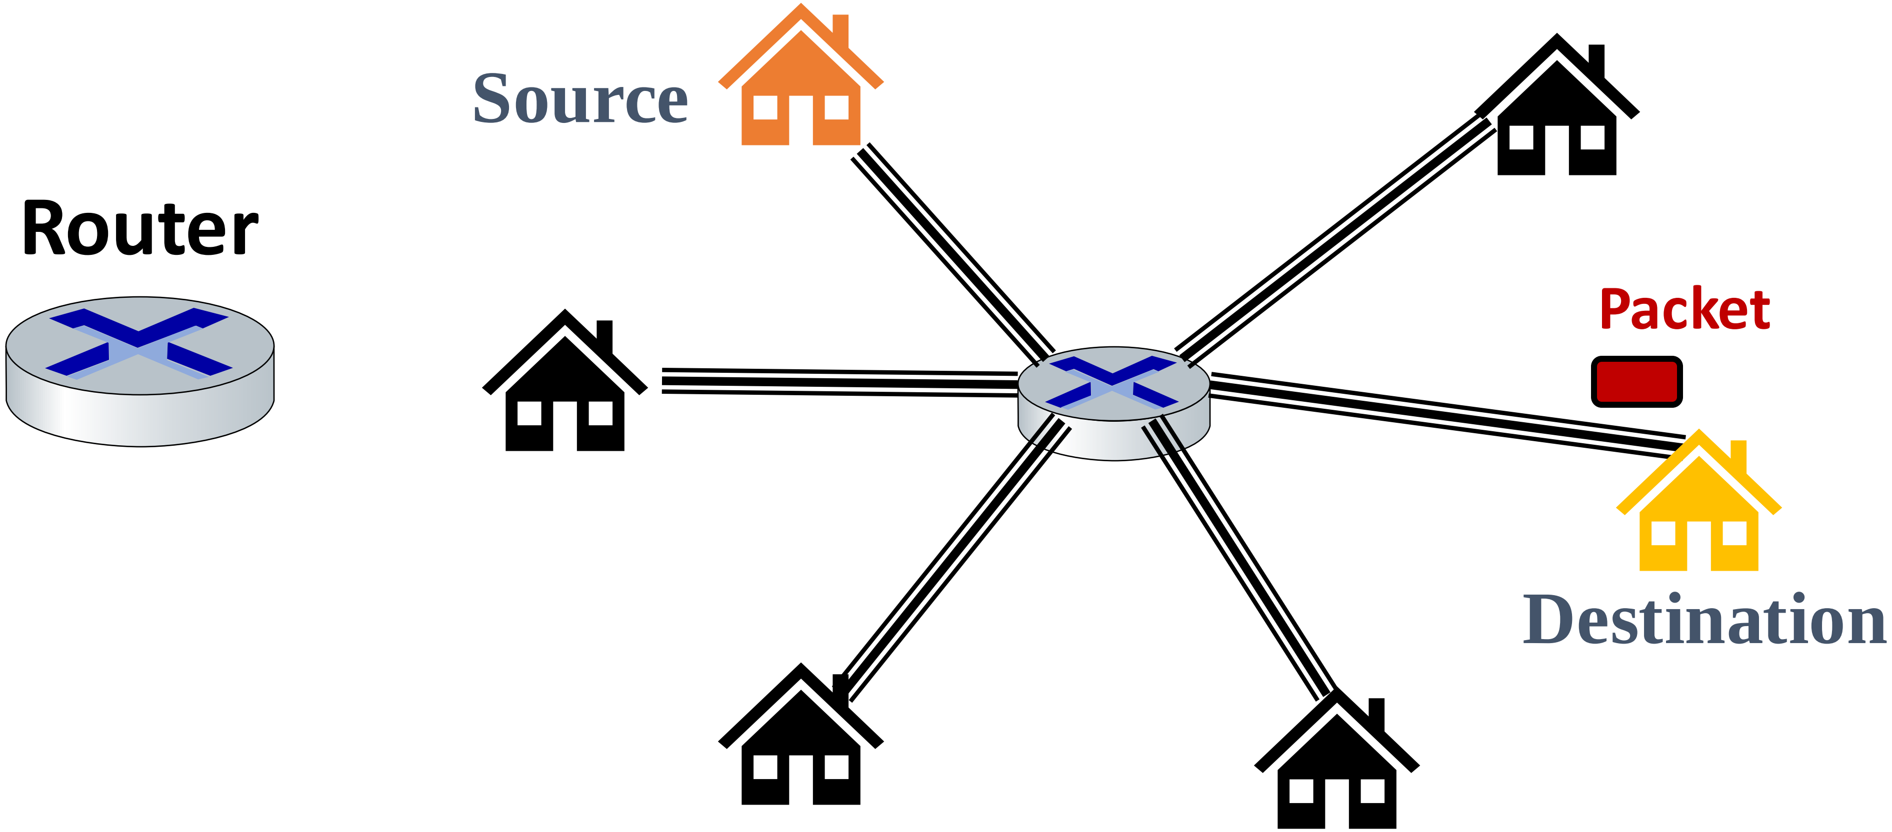

Forwarding:
- aka "switching"
- *local* action: move arriving packets from router's input link to appropriate router output link

The network core

- mesh of interconnected routers

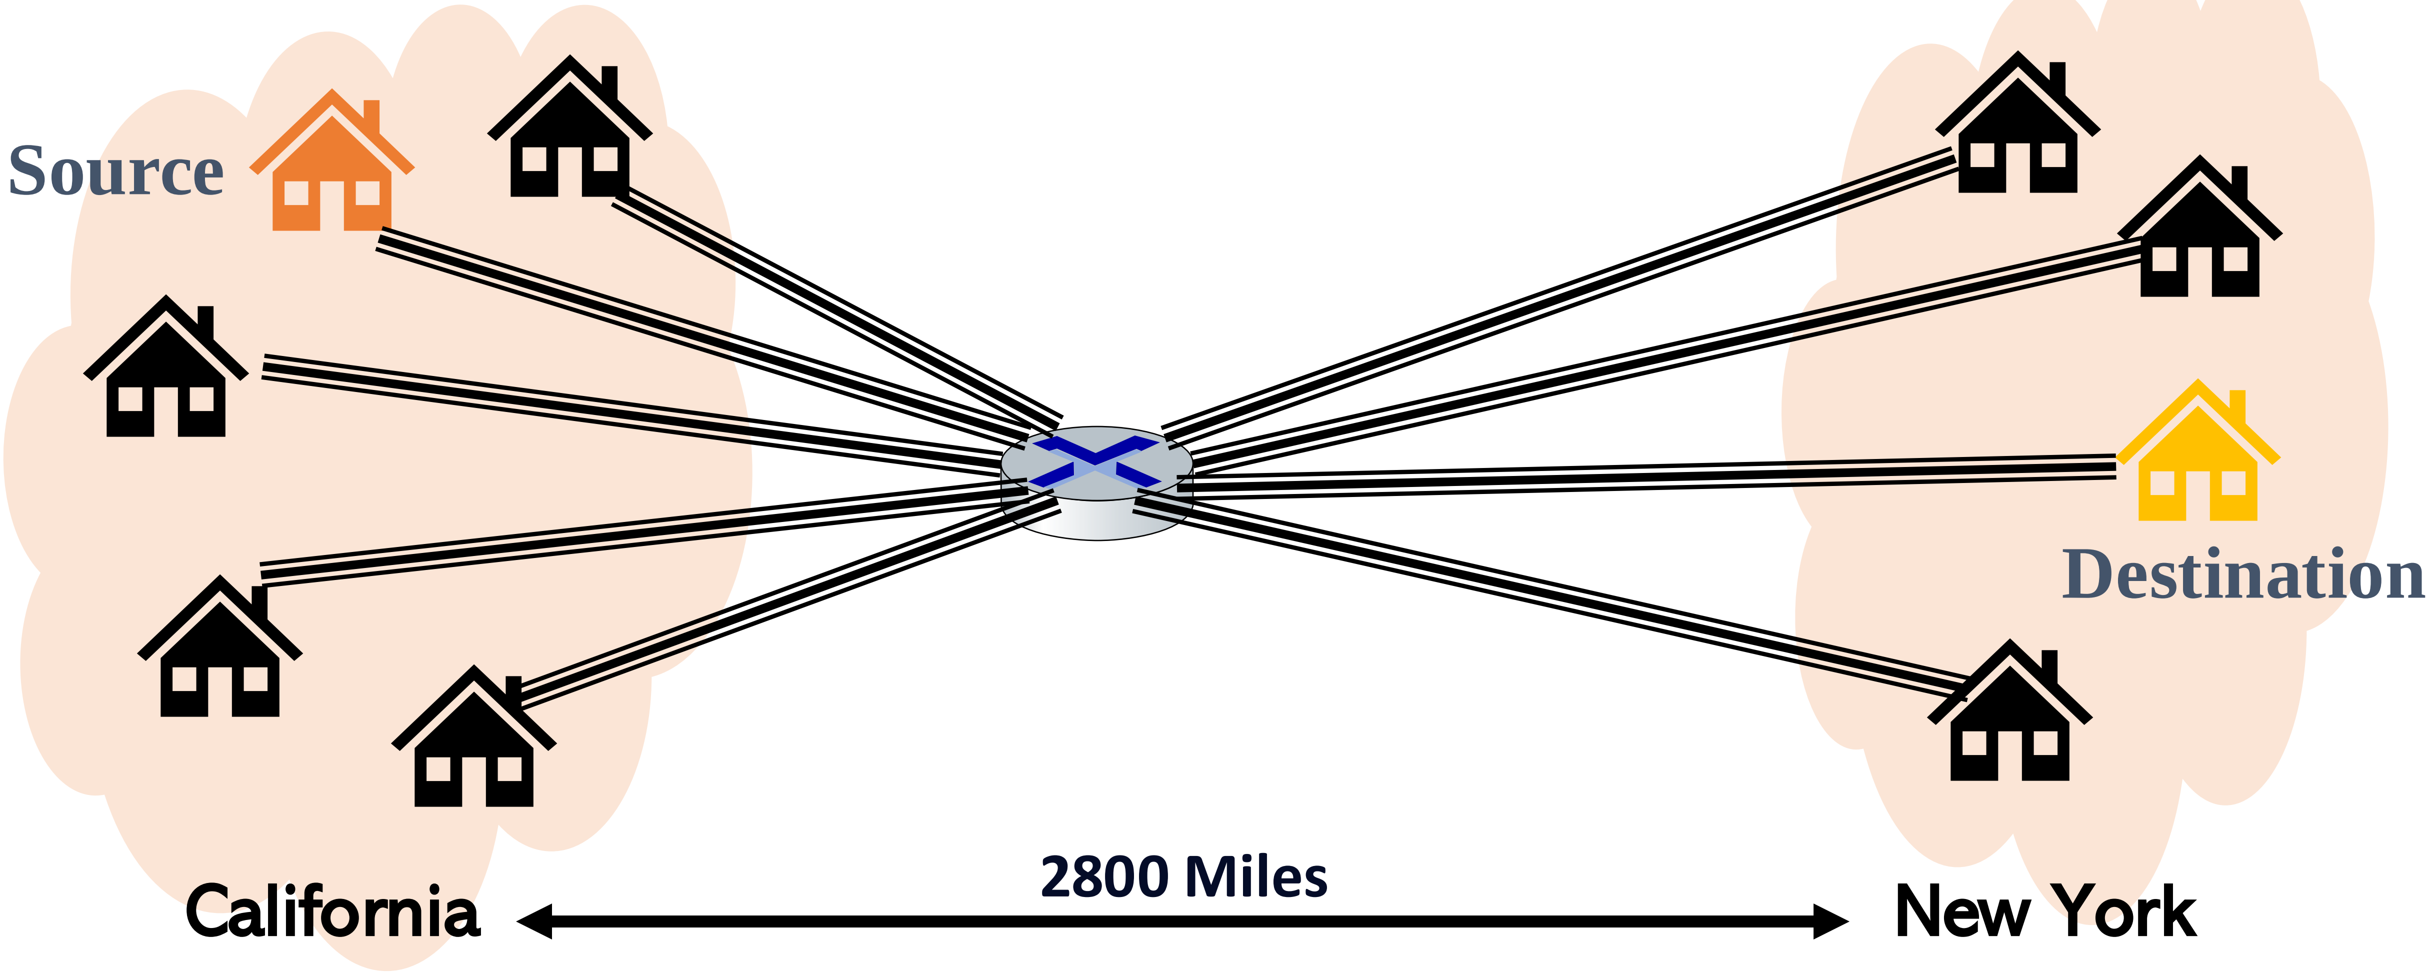

The network core

- mesh of interconnected routers

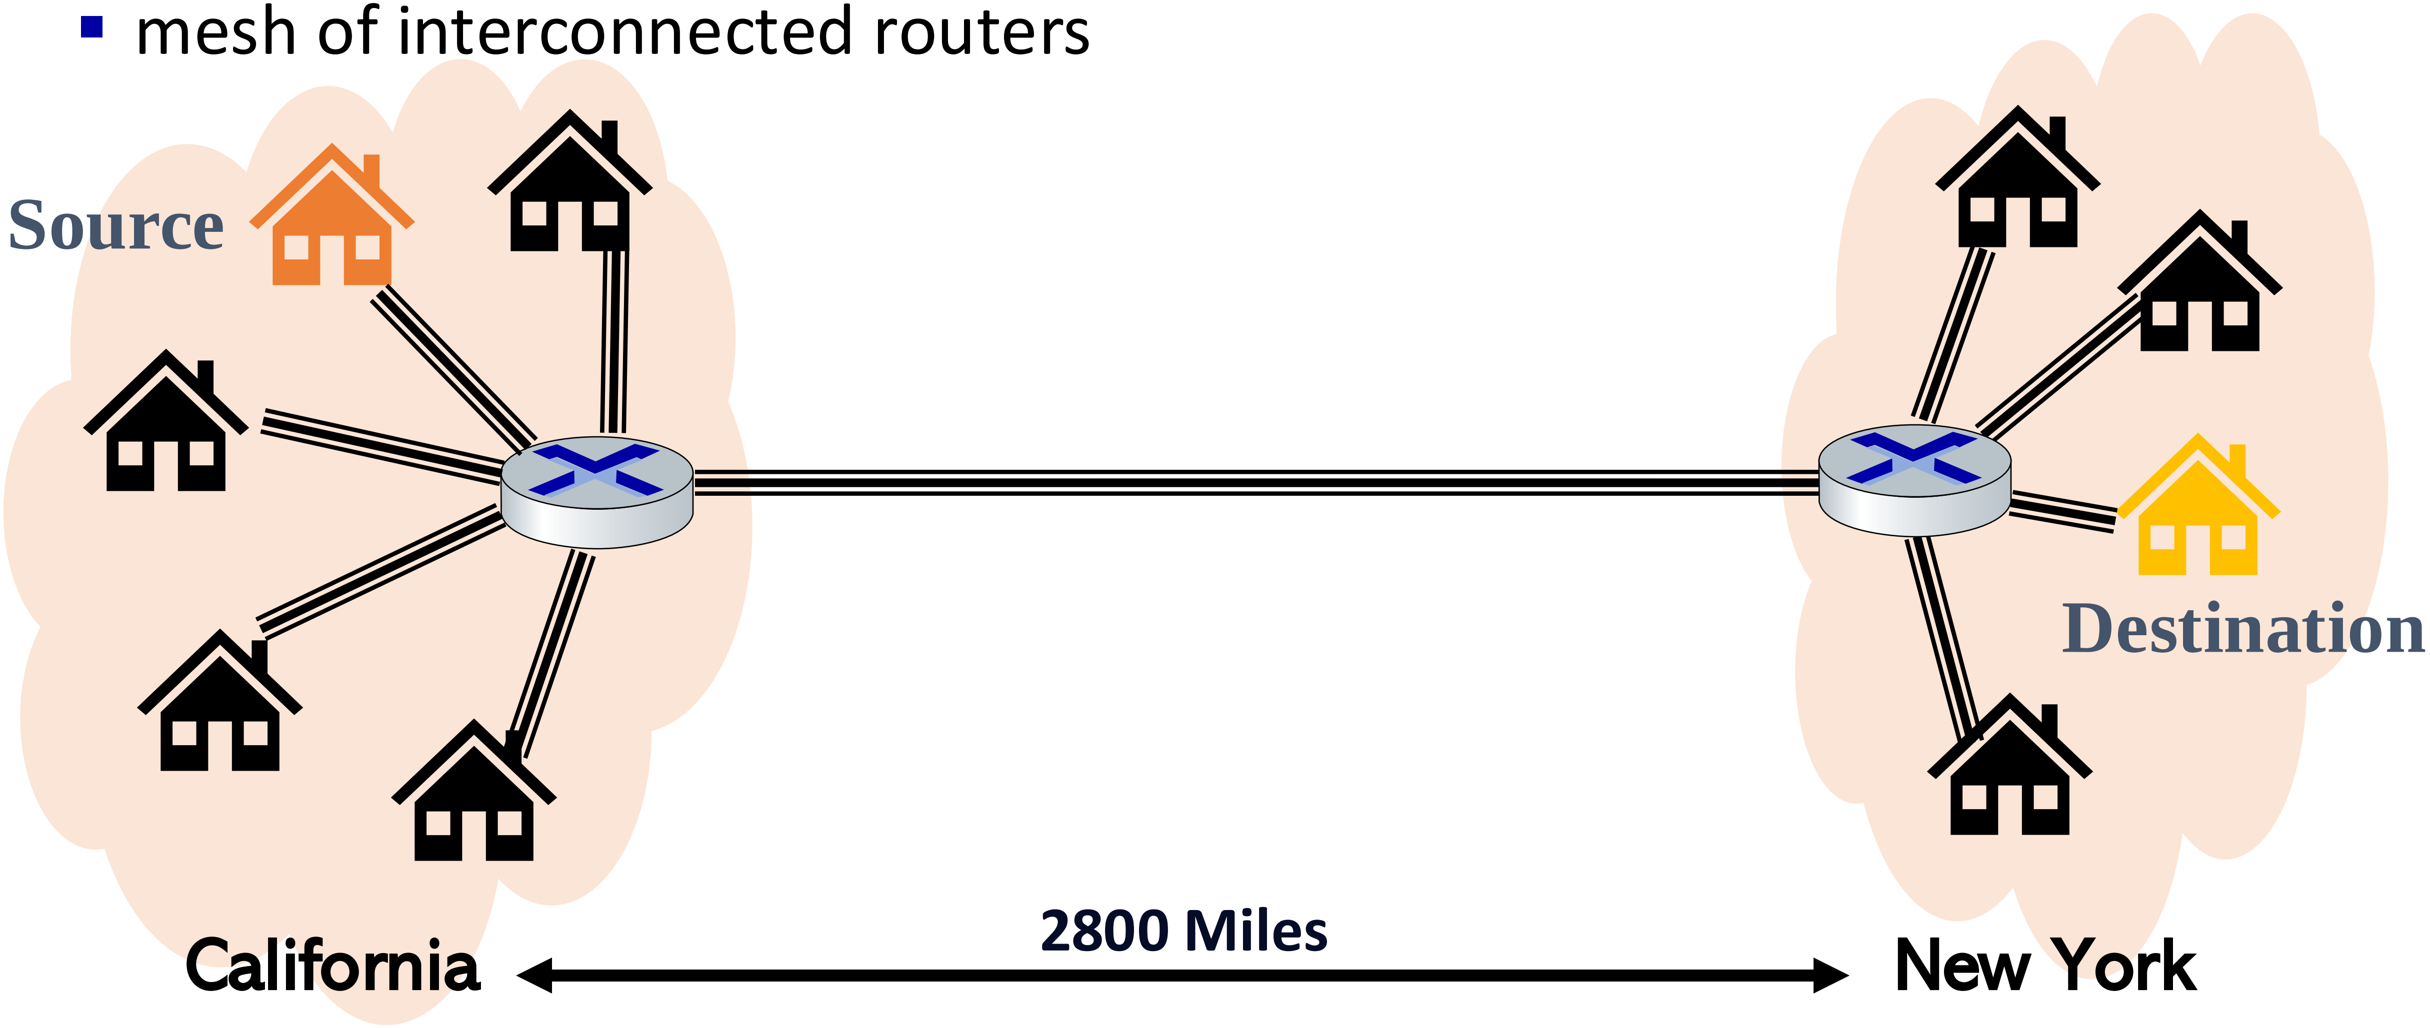

The network core

- mesh of interconnected routers

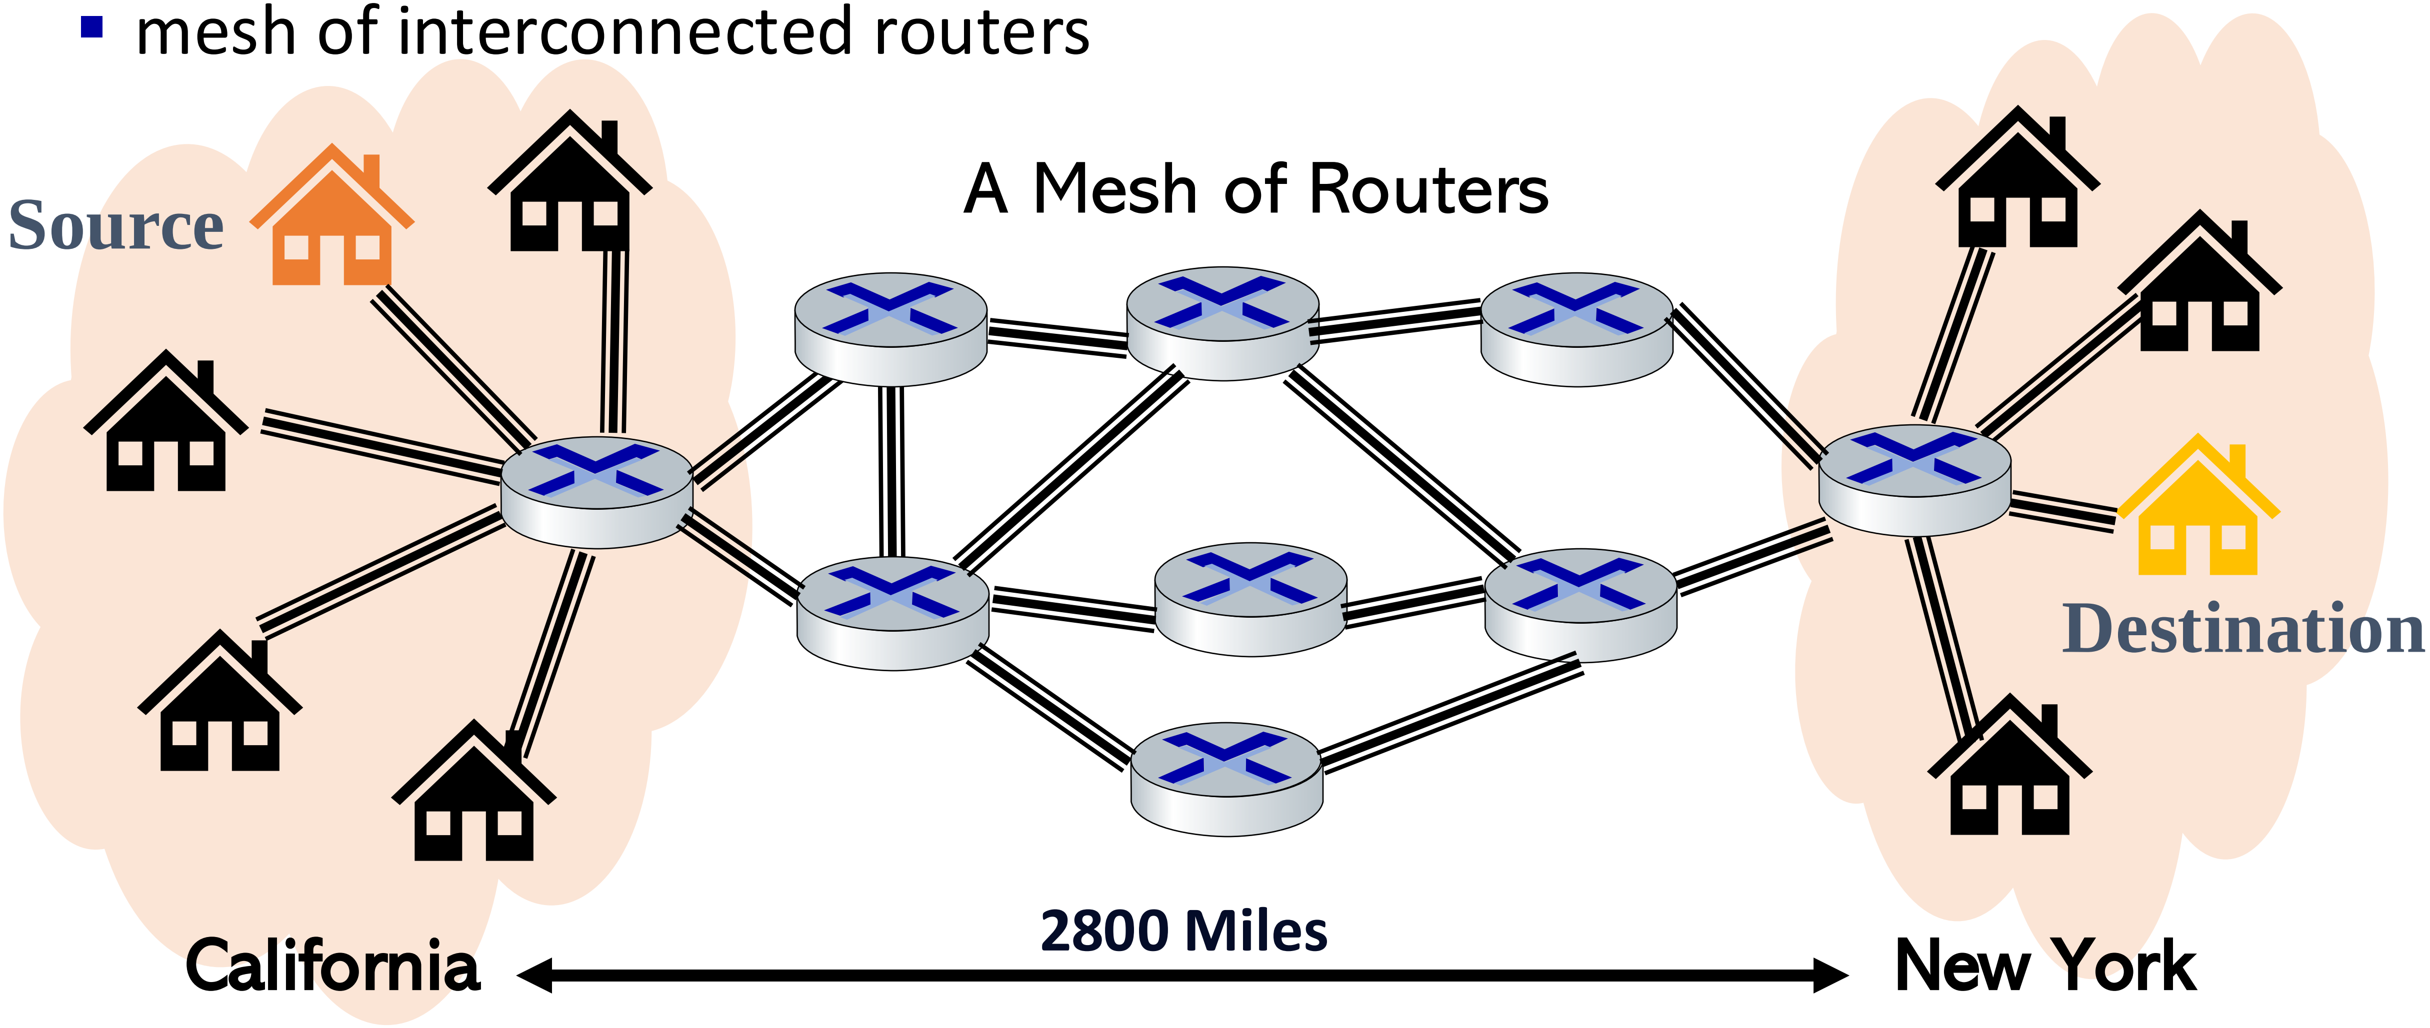

The network core

- mesh of interconnected routers

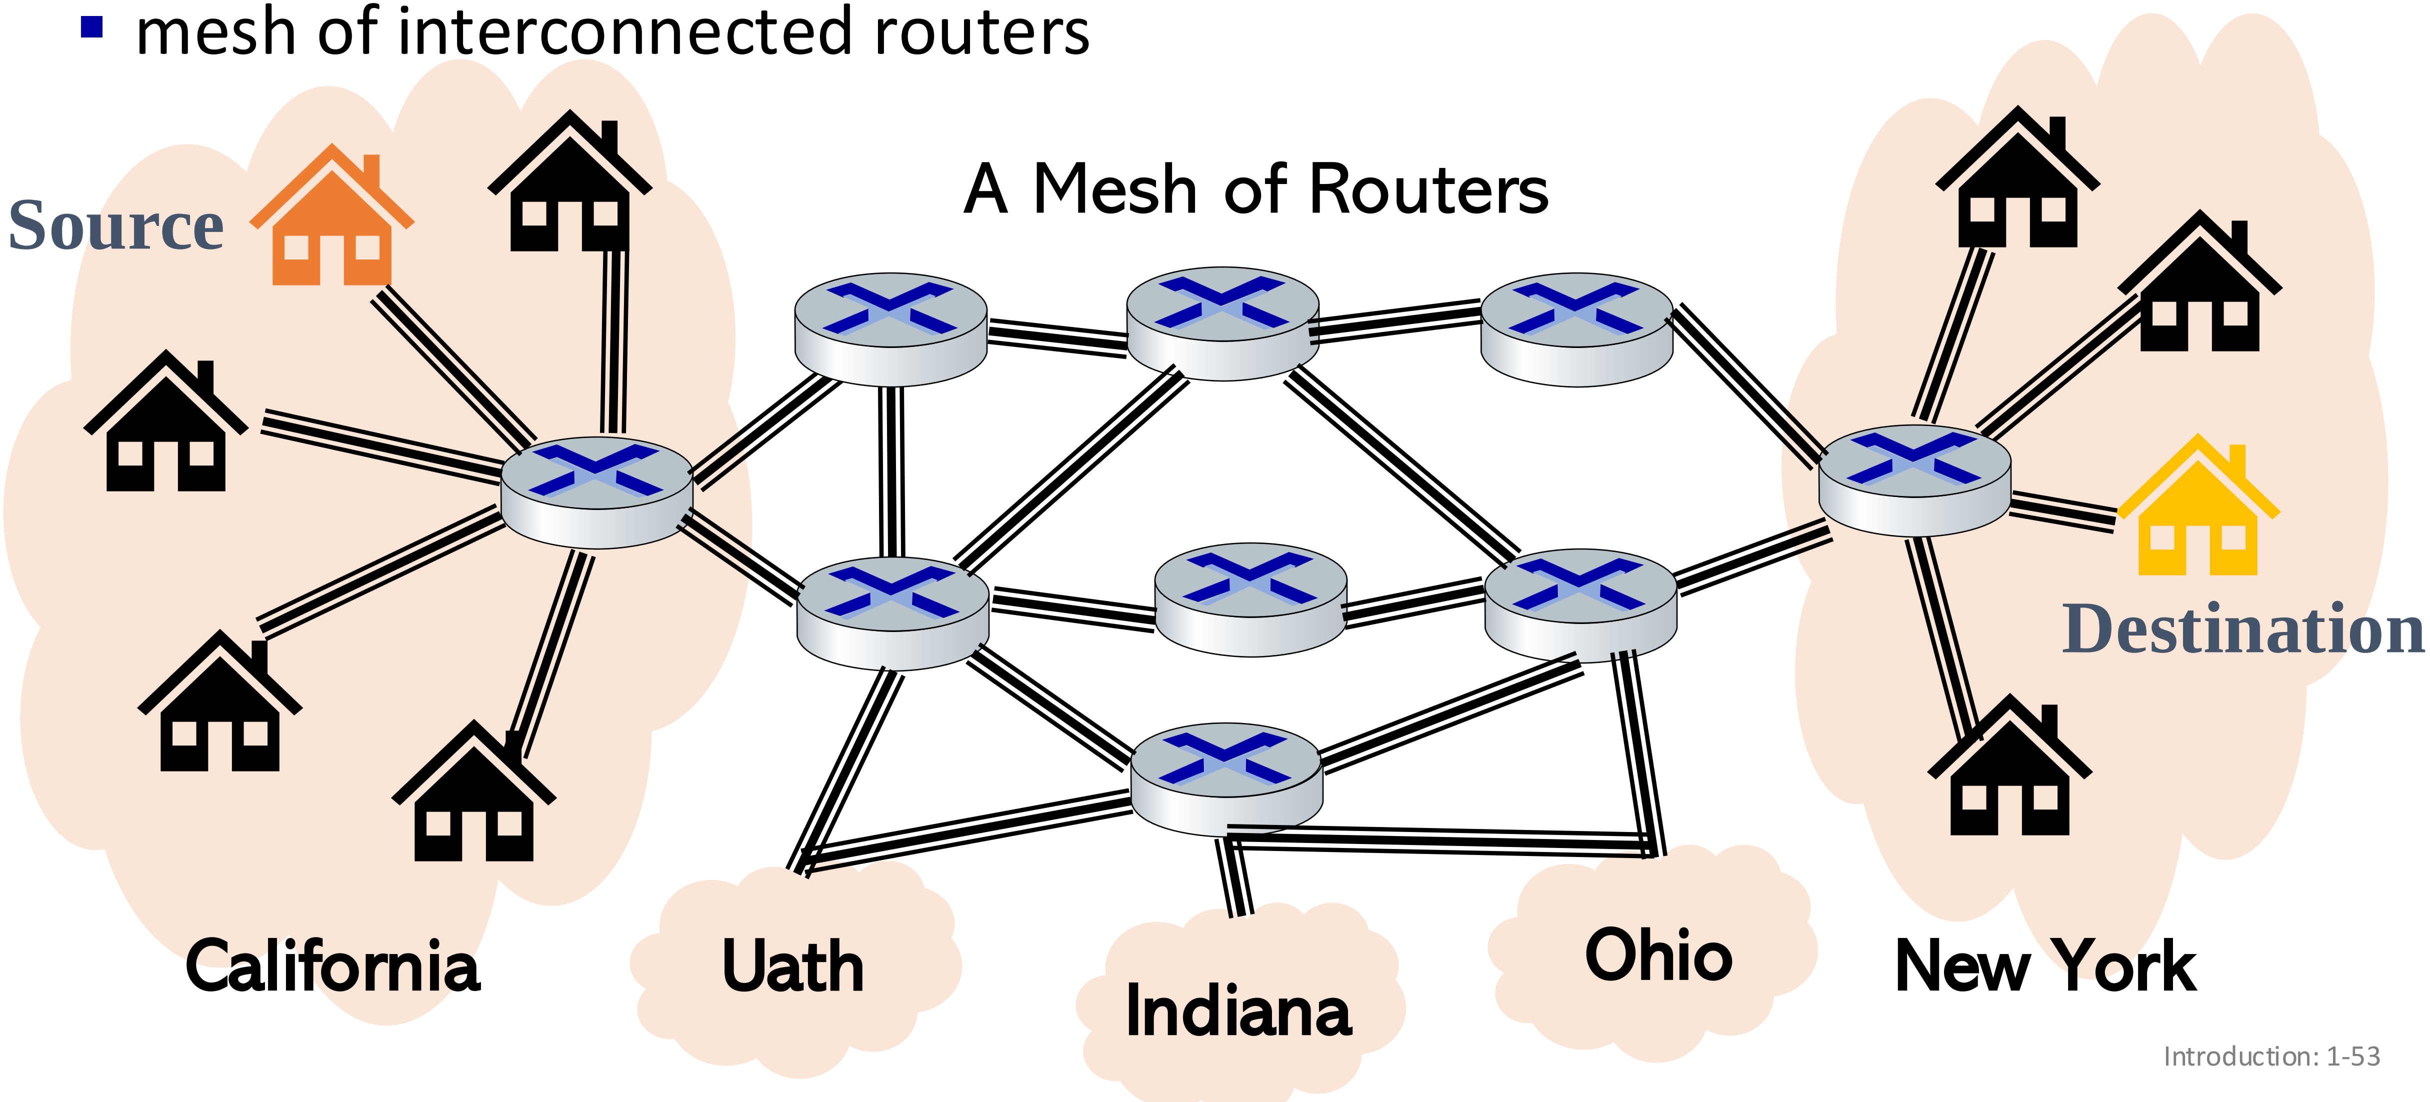

The network core

- mesh of interconnected routers

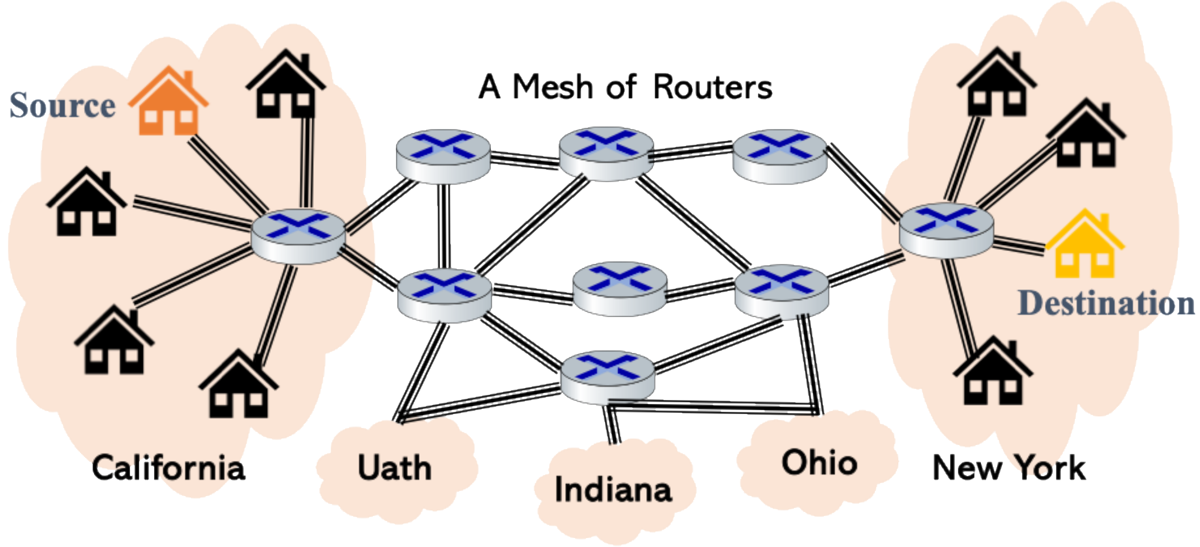

Aren't they very similar to each other?

The network core – Routing

- Routing: Finding the Correct/Optimal path from source to destination

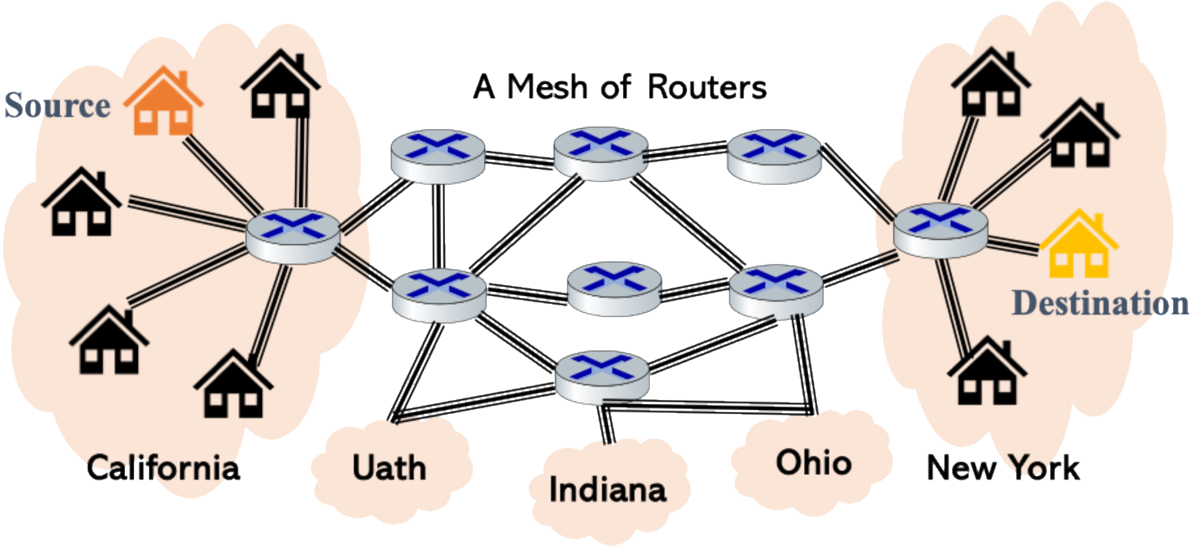

What's a correct/optimal path?

The network core – Routing

- Routing: Finding the Correct/Optimal path from source to destination

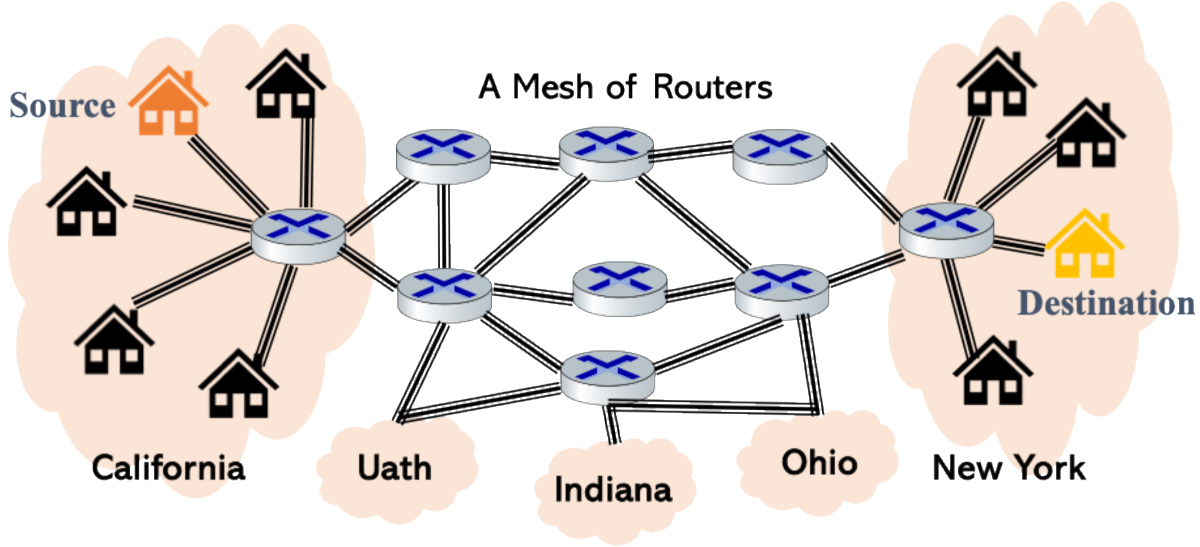

Routing Algorithm

- ❖ Shortest distance?

- ❖ Shortest distance?
- ❖ Cheapest without tolls?
- ❖ Best views?

The network core – Routing

- Routing: Finding the Correct/Optimal path from source to destination

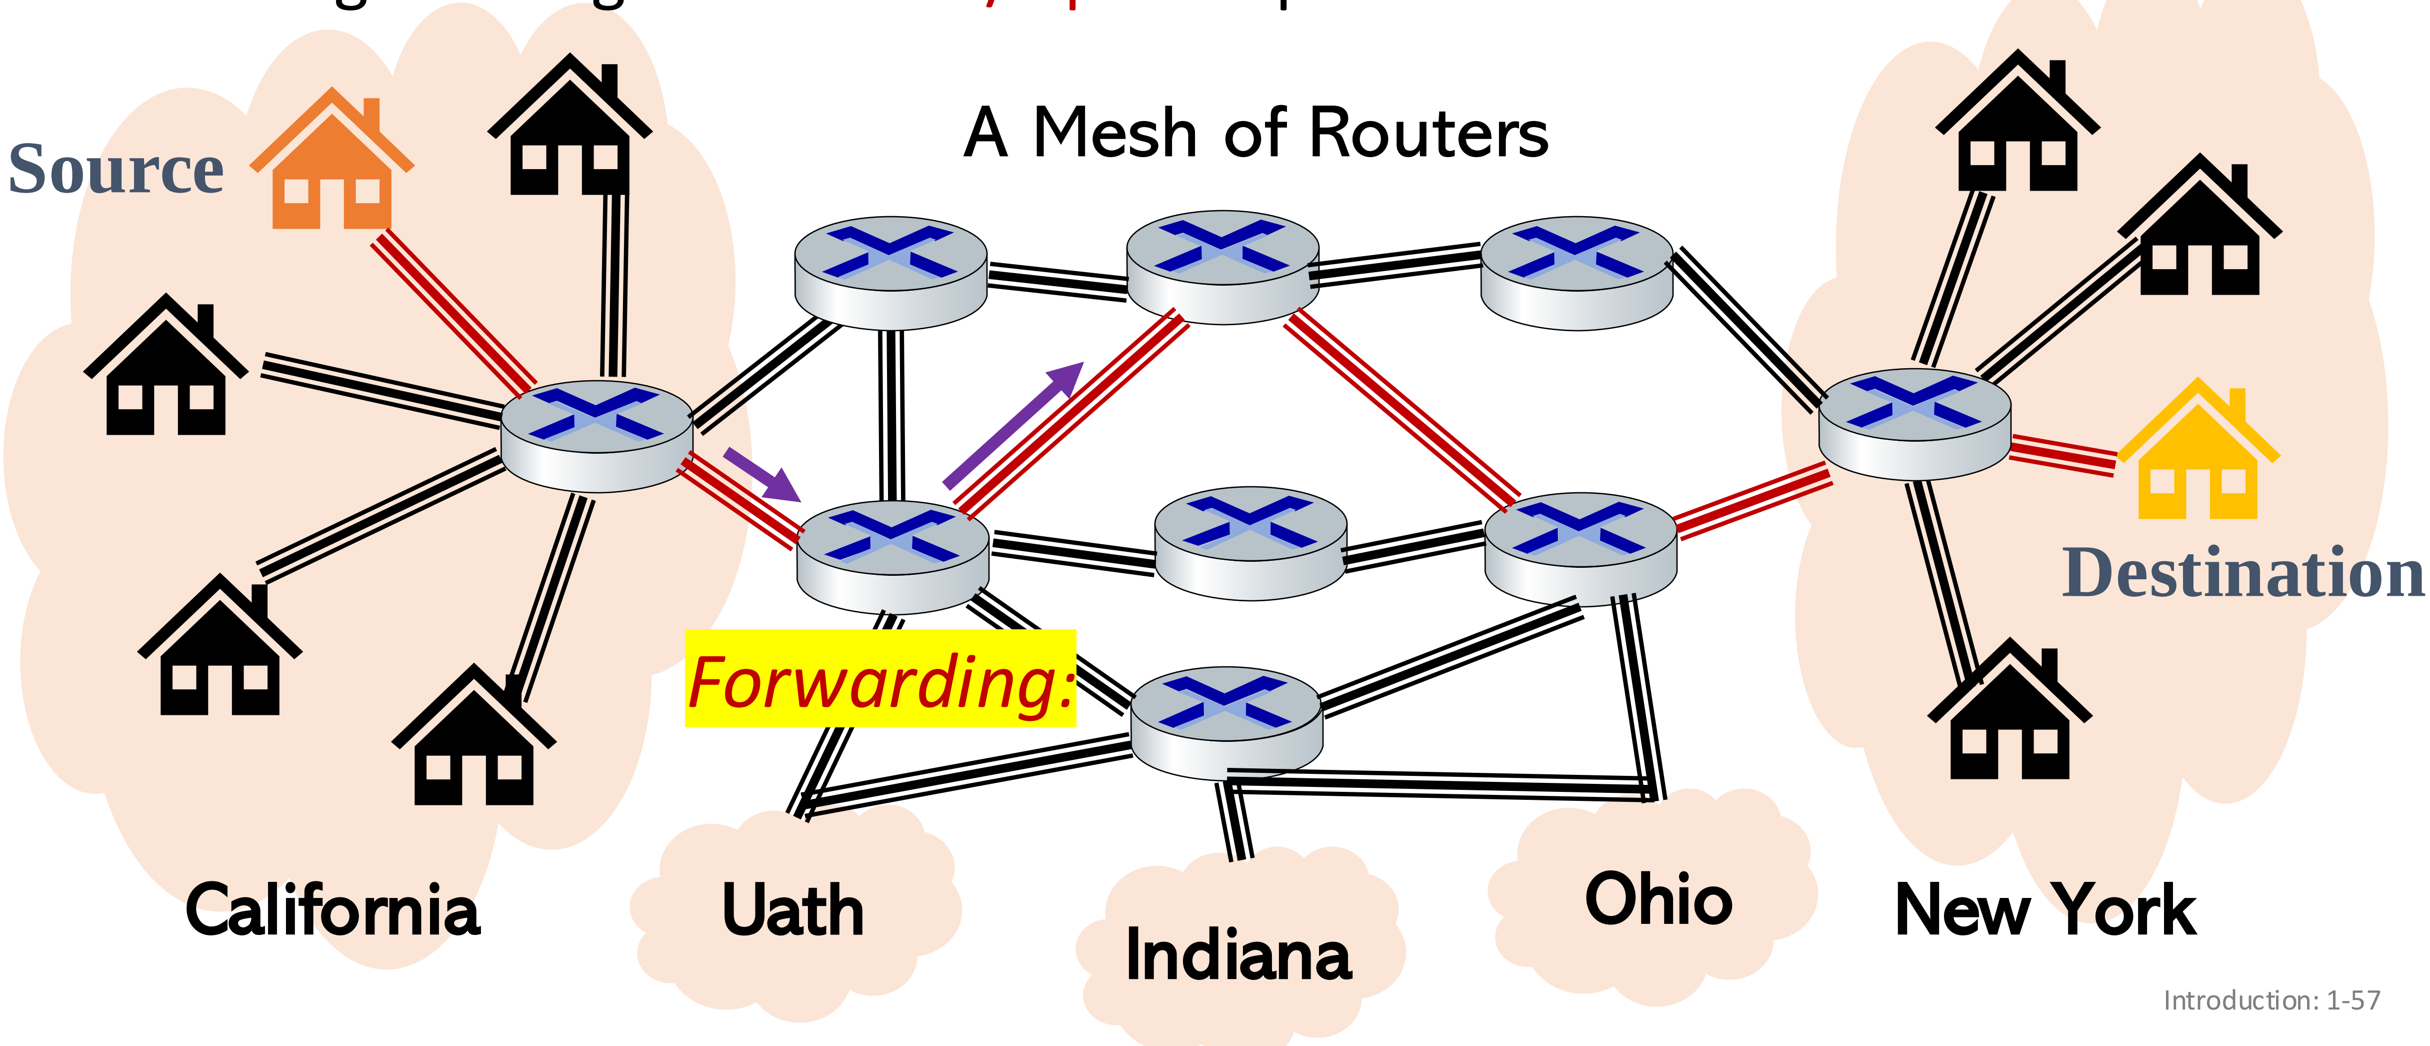

The network core – Routing

- Routing: Finding the Correct/Optimal path from source to destination

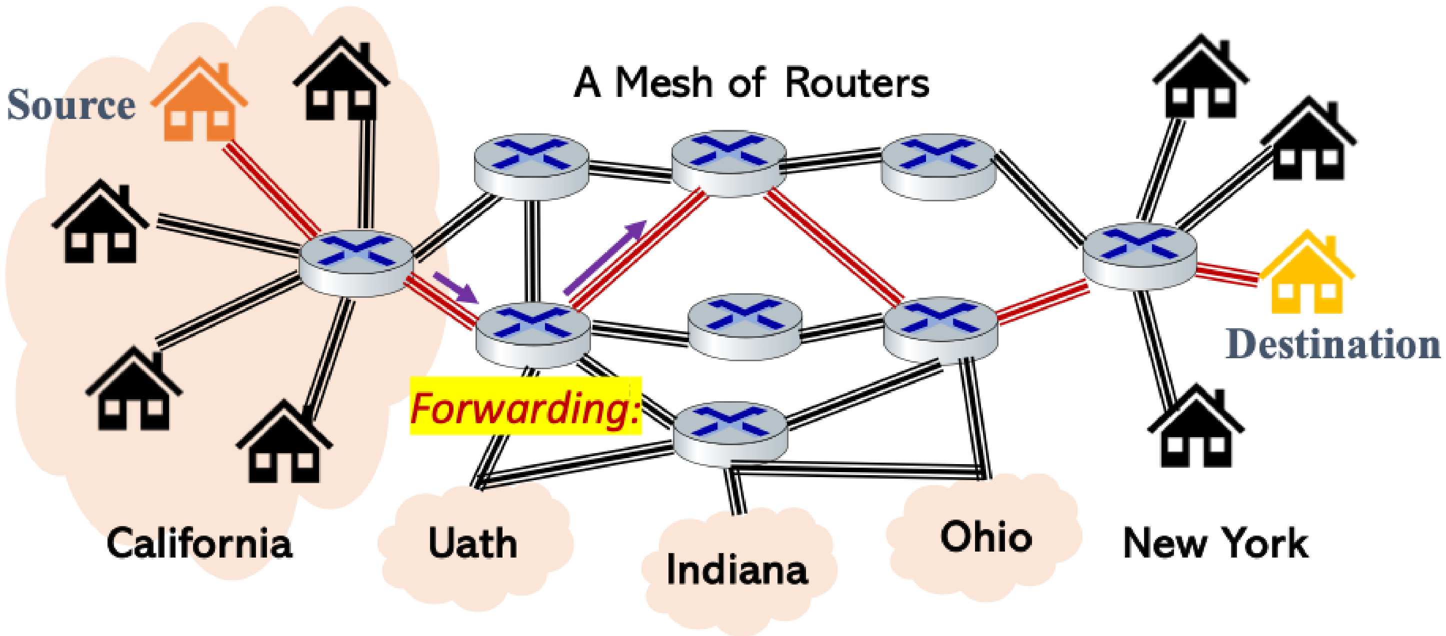

Forwarding:

- *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

Packet Switching VS Circuit Switching

- Forward is also called switching

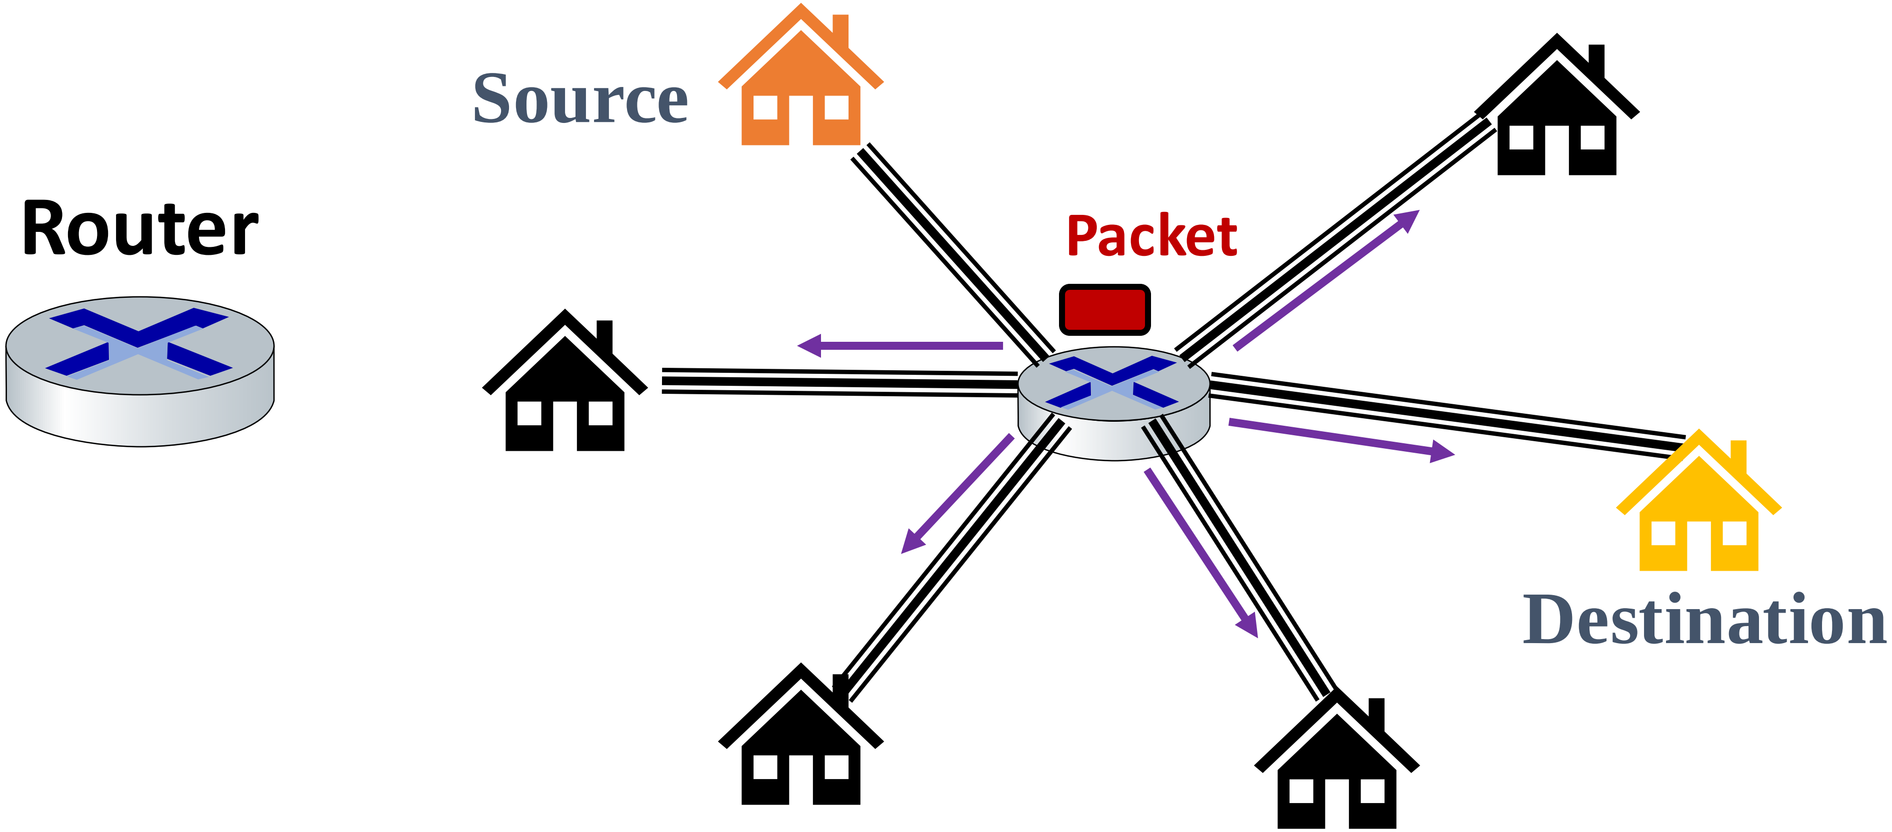

Forwarding:

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

Packet Switching

- Forward is also called switching

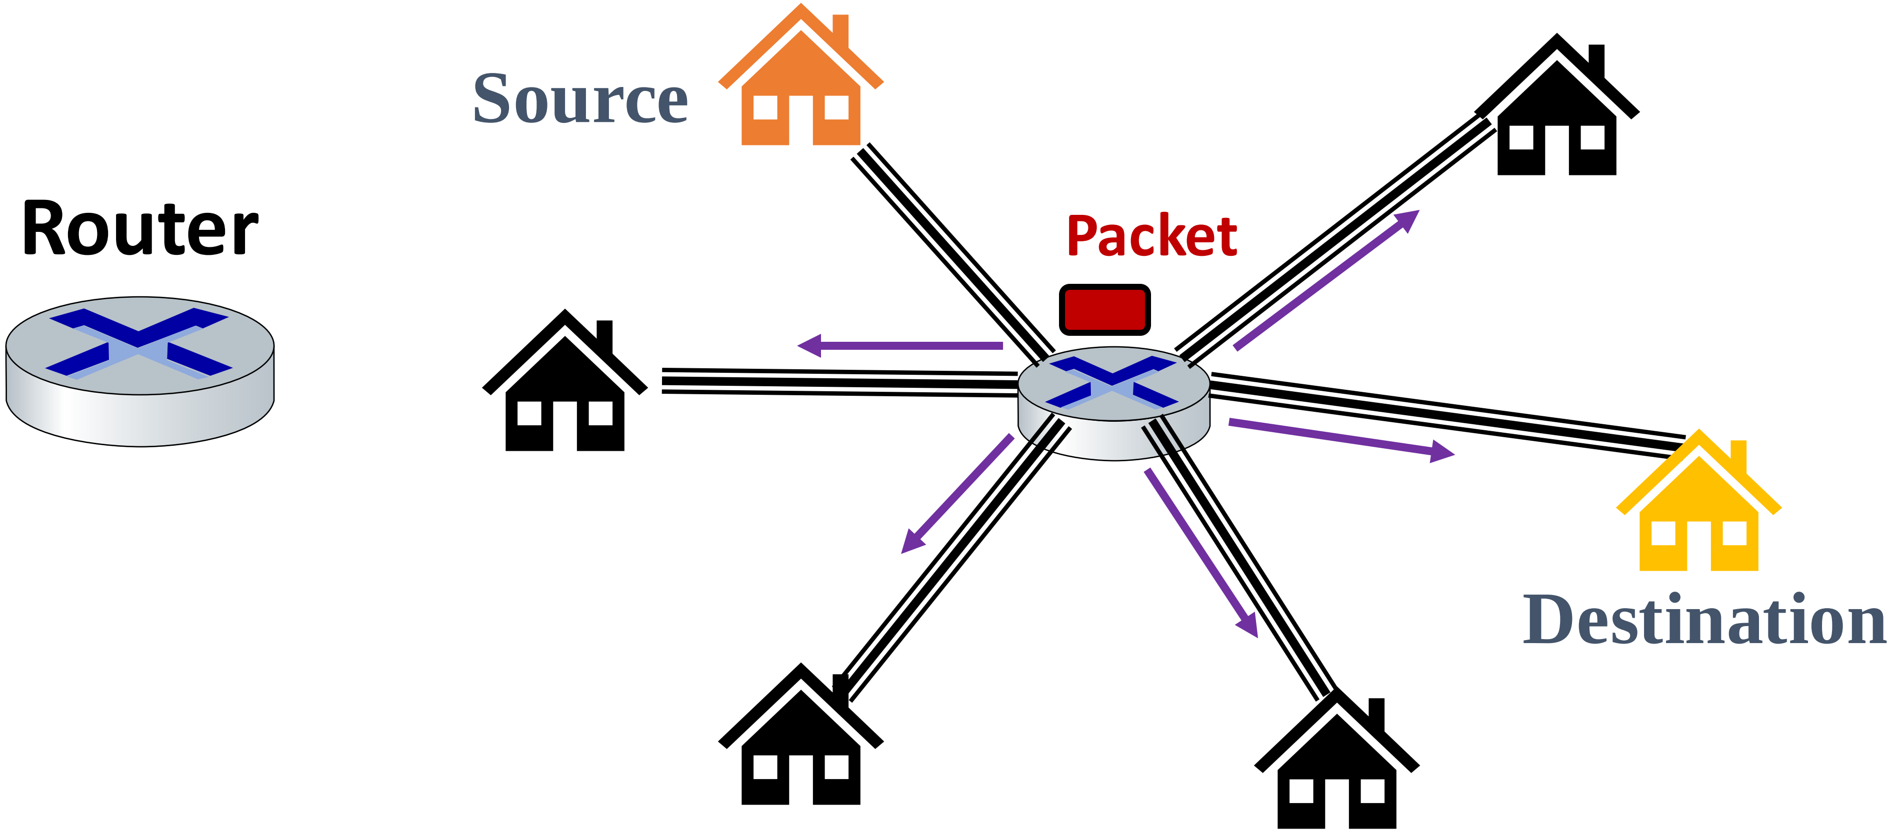

Forwarding:

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

Packet Switching

- Forward is also called switching

Router

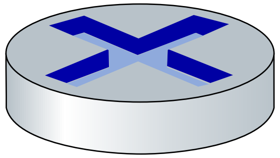

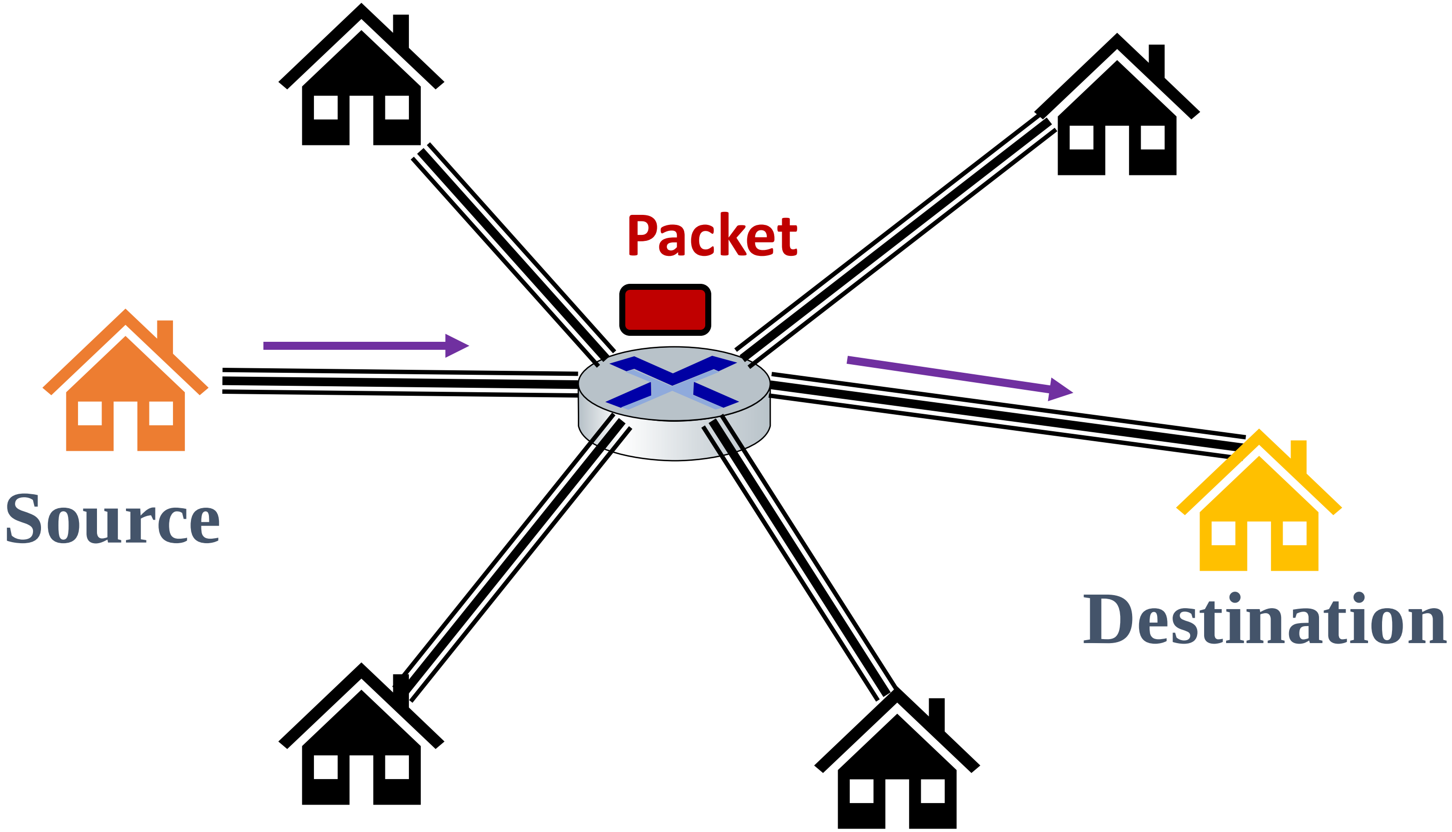

Forwarding:

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

Packet-switching: store-and-forward

- Forward is also called switching
  - *store and forward: entire* packet must arrive at router before it can be transmitted on next link

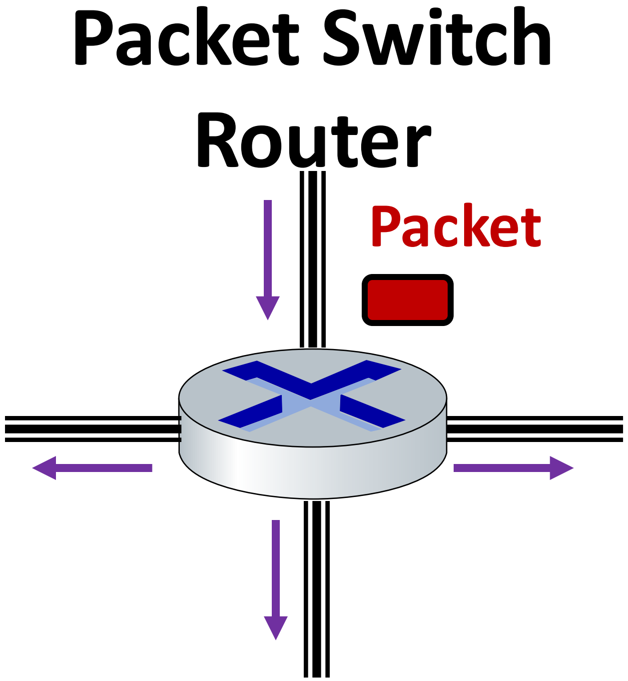

Forwarding:

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

Packet-switching: store-and-forward

- Forward is also called switching
  - *store and forward: entire* packet must arrive at router before it can be transmitted on next link

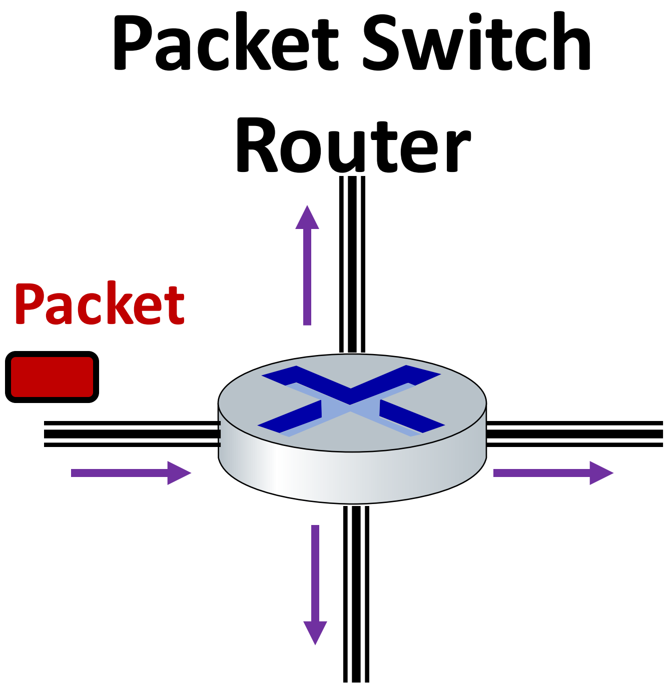

Forwarding:

- aka “switching”
- *local* action: move arriving packets from router’s input link to appropriate router output link

Circuit switching

- Forward is also called switching

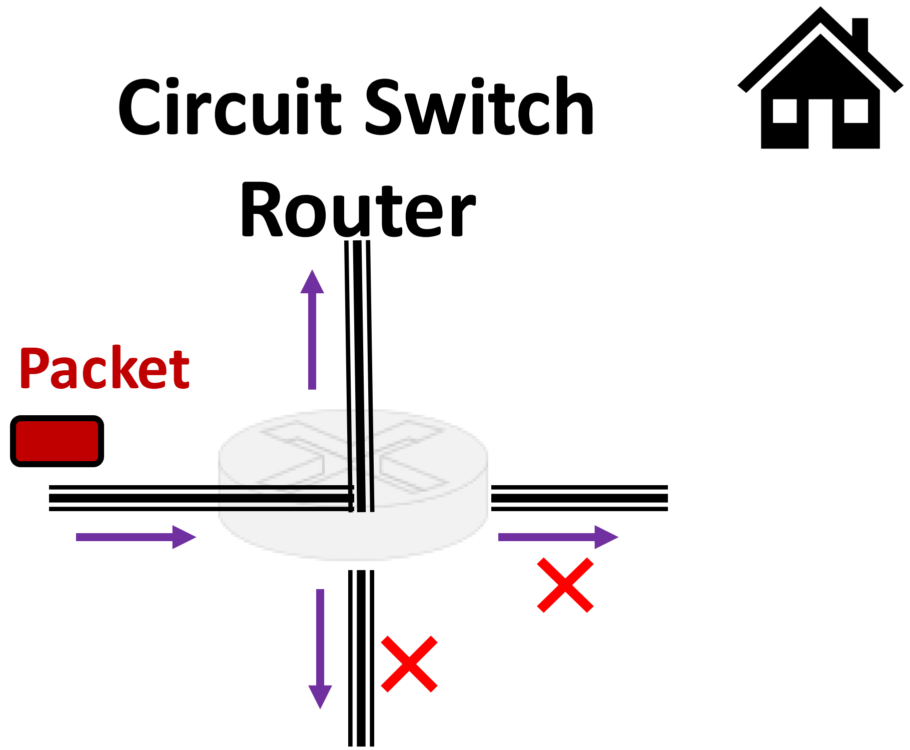

Forwarding:

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

Circuit switching

- Forward is also called switching

Circuit Switch

Router

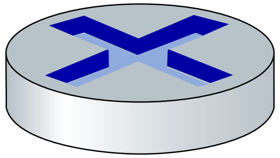

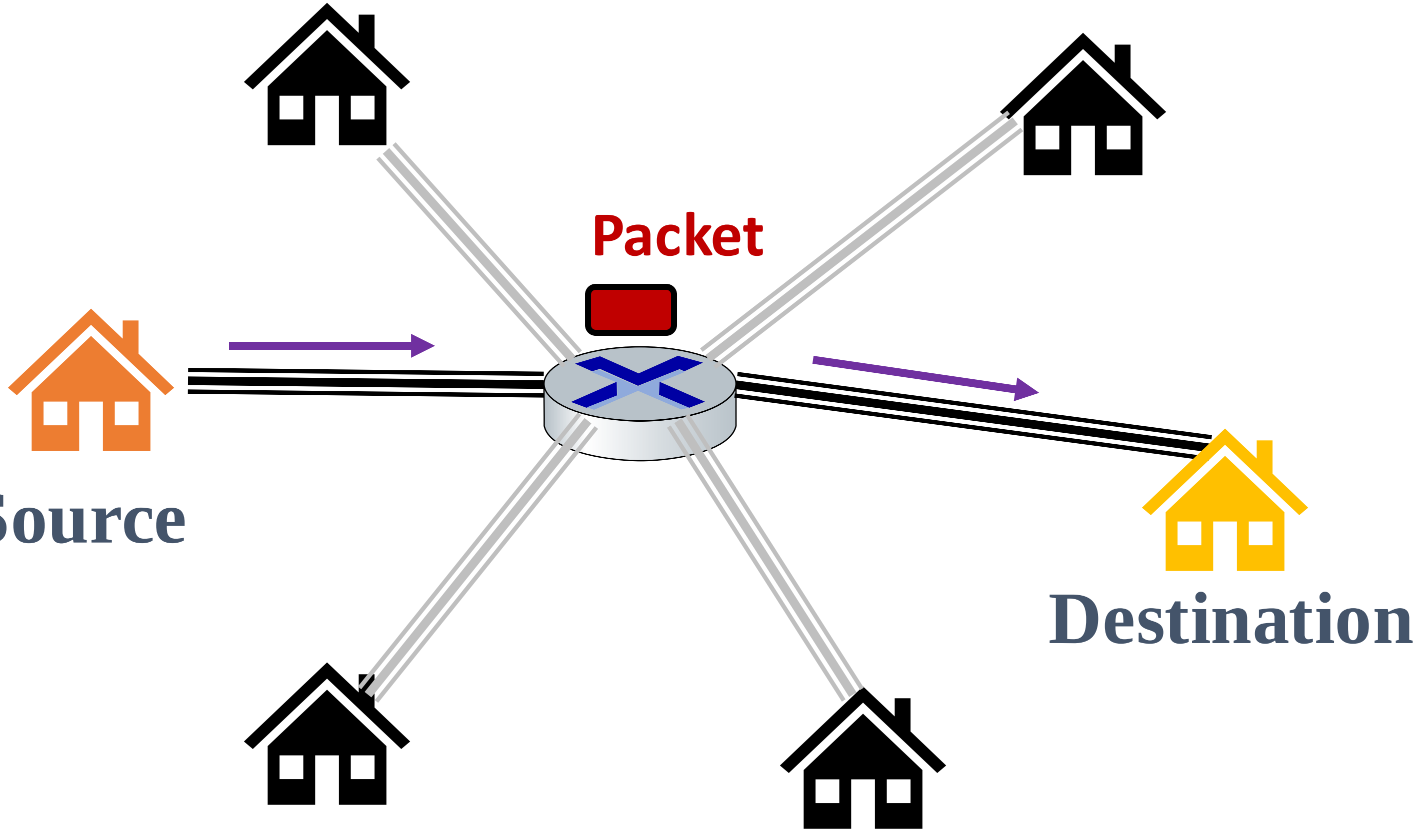

Forwarding:

- aka “switching”
- *local* action: move arriving packets from router’s input link to appropriate router output link

Internet Core: Packet Switching

each end-end data stream divided into *packets*

- users A through C packets *share* network resources
- each packet uses full link bandwidth
- resources used *as needed*

- each packet has a "header" (containing e.g., destination address) in addition to "payload" (data)
- Store and Forward (requires buffer and introduces delay)

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

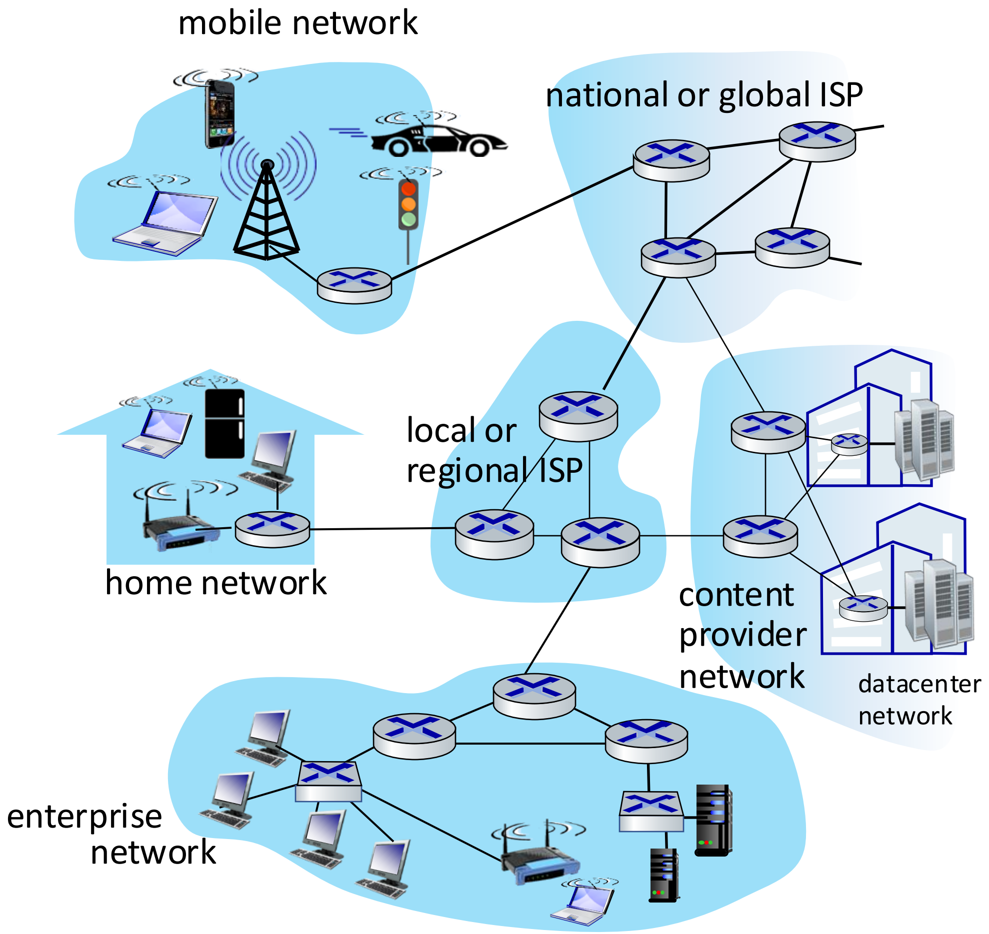

Let's take a stepwise approach to describe current Internet structure

Internet structure: a "network of networks"

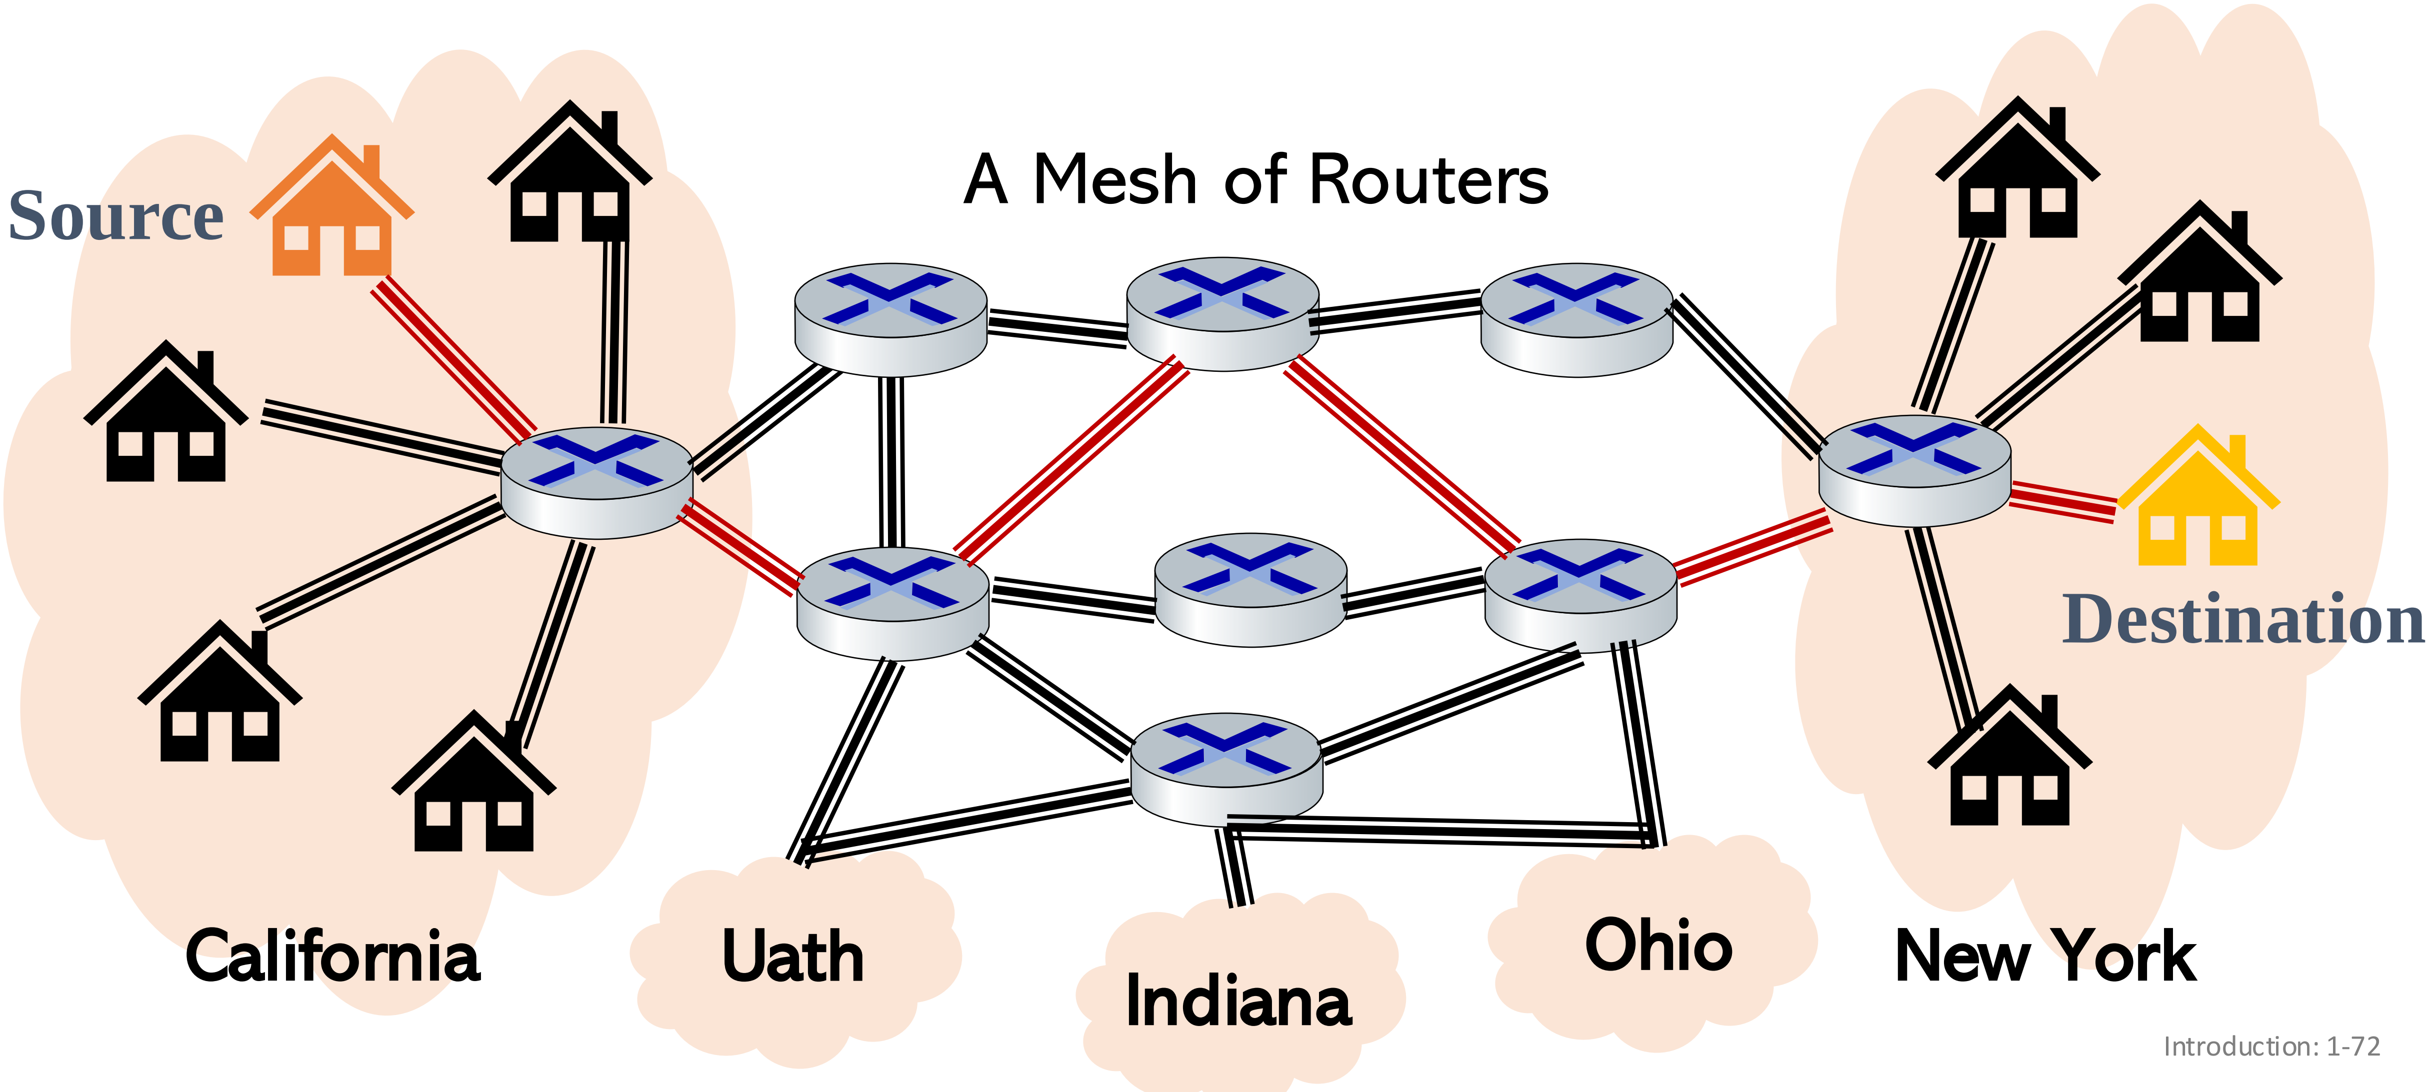

Internet structure: a "network of networks"

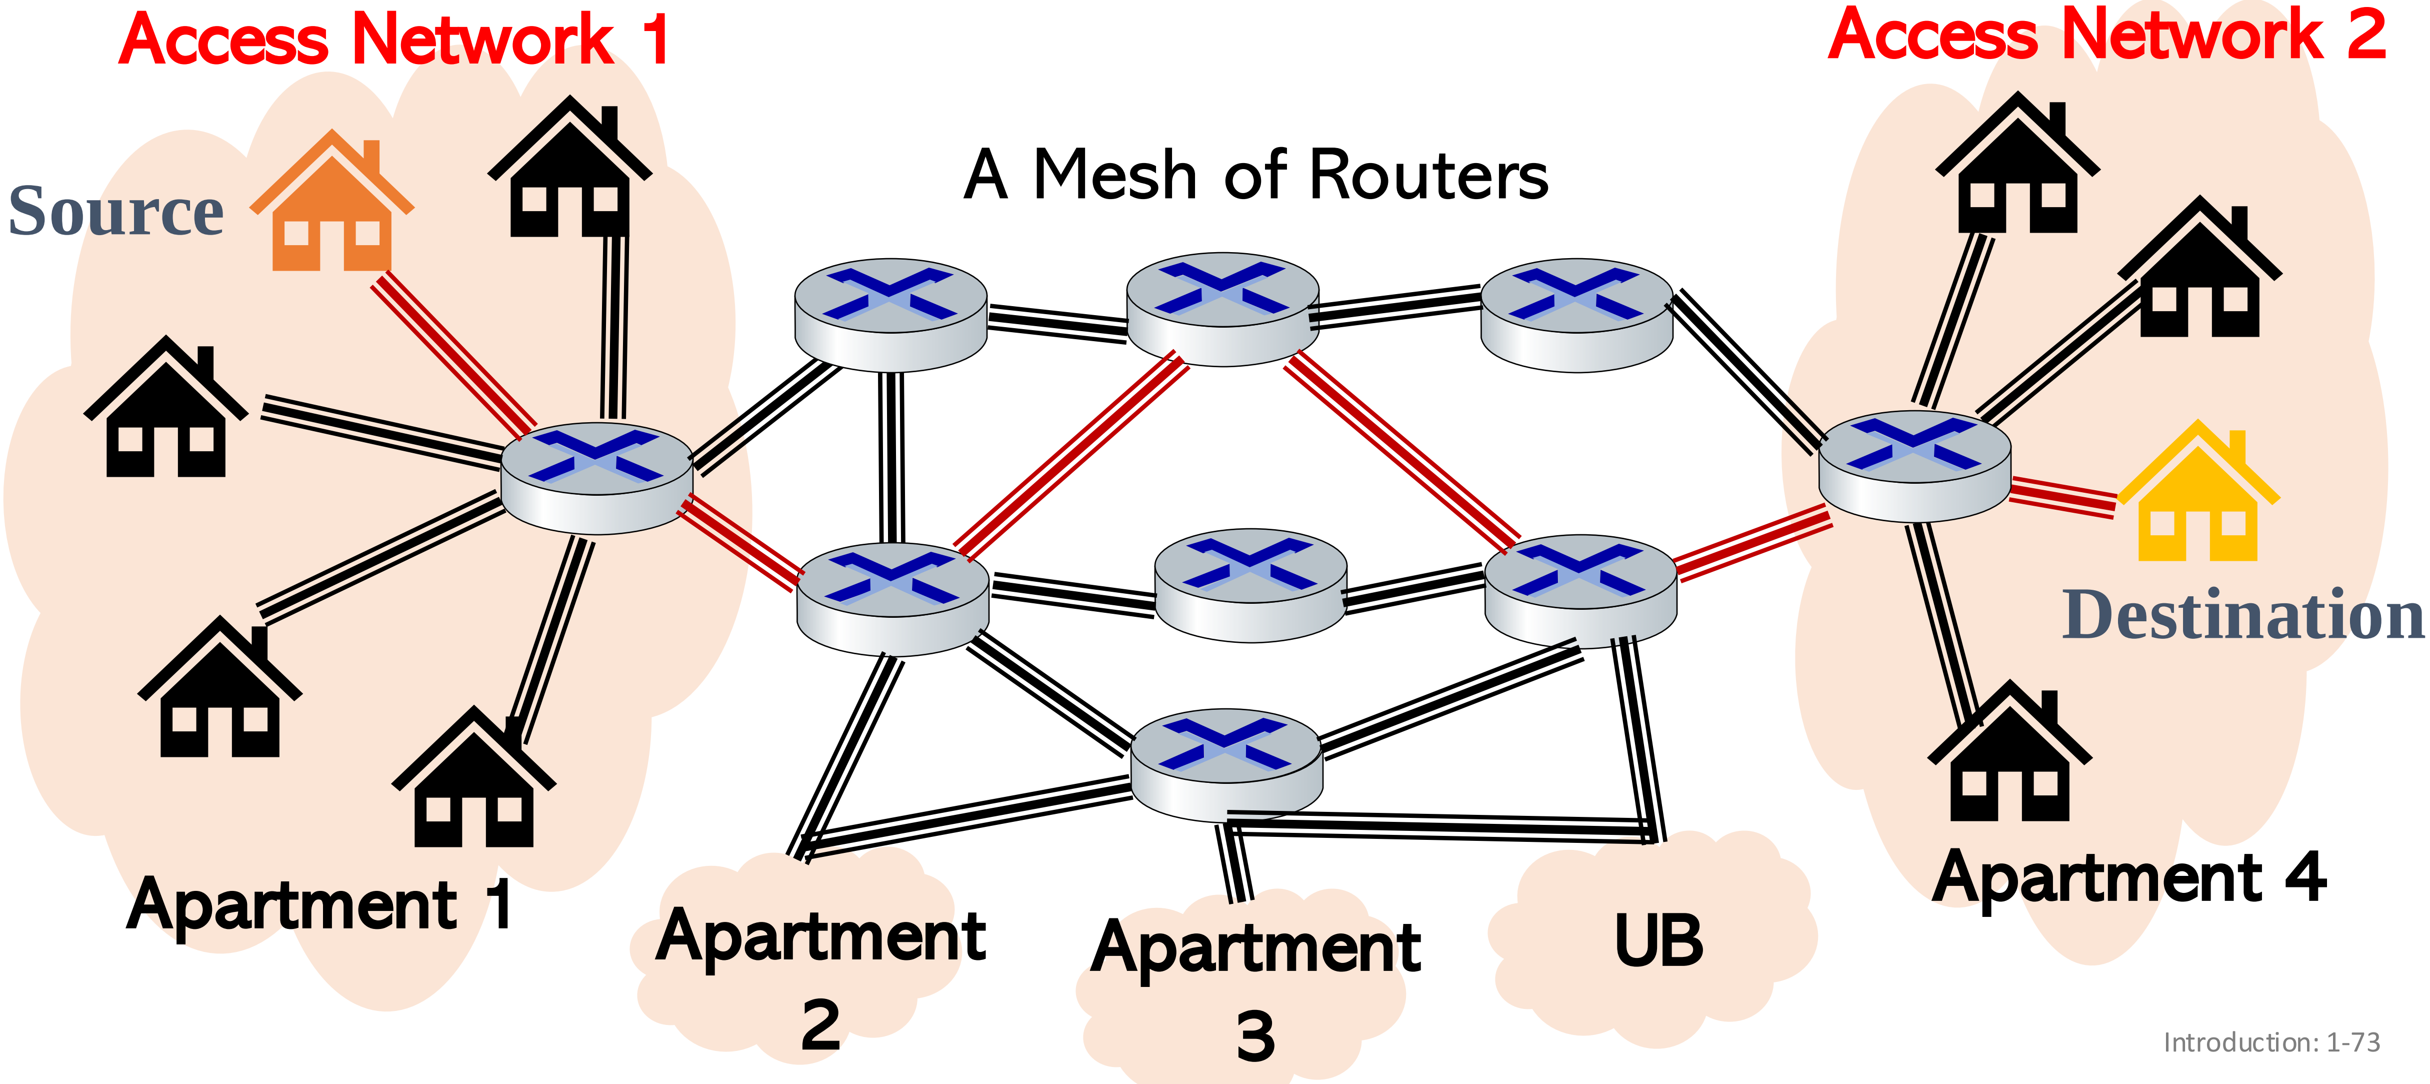

Internet structure: a "network of networks"

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

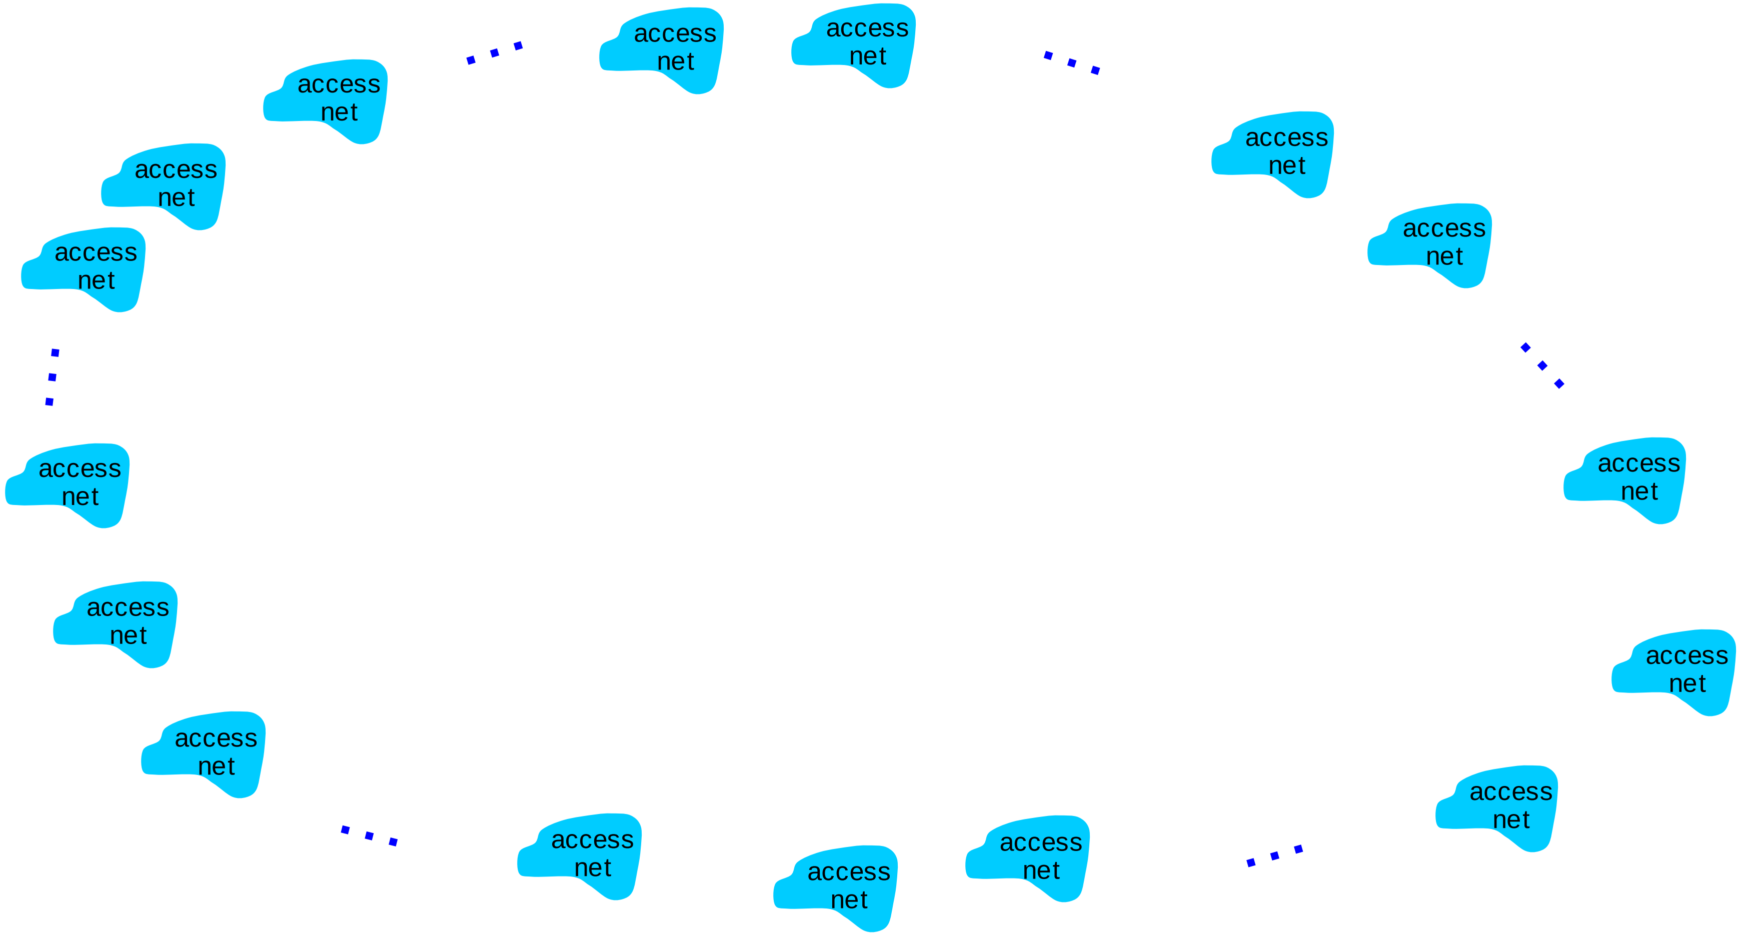

Internet structure: a "network of networks"

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

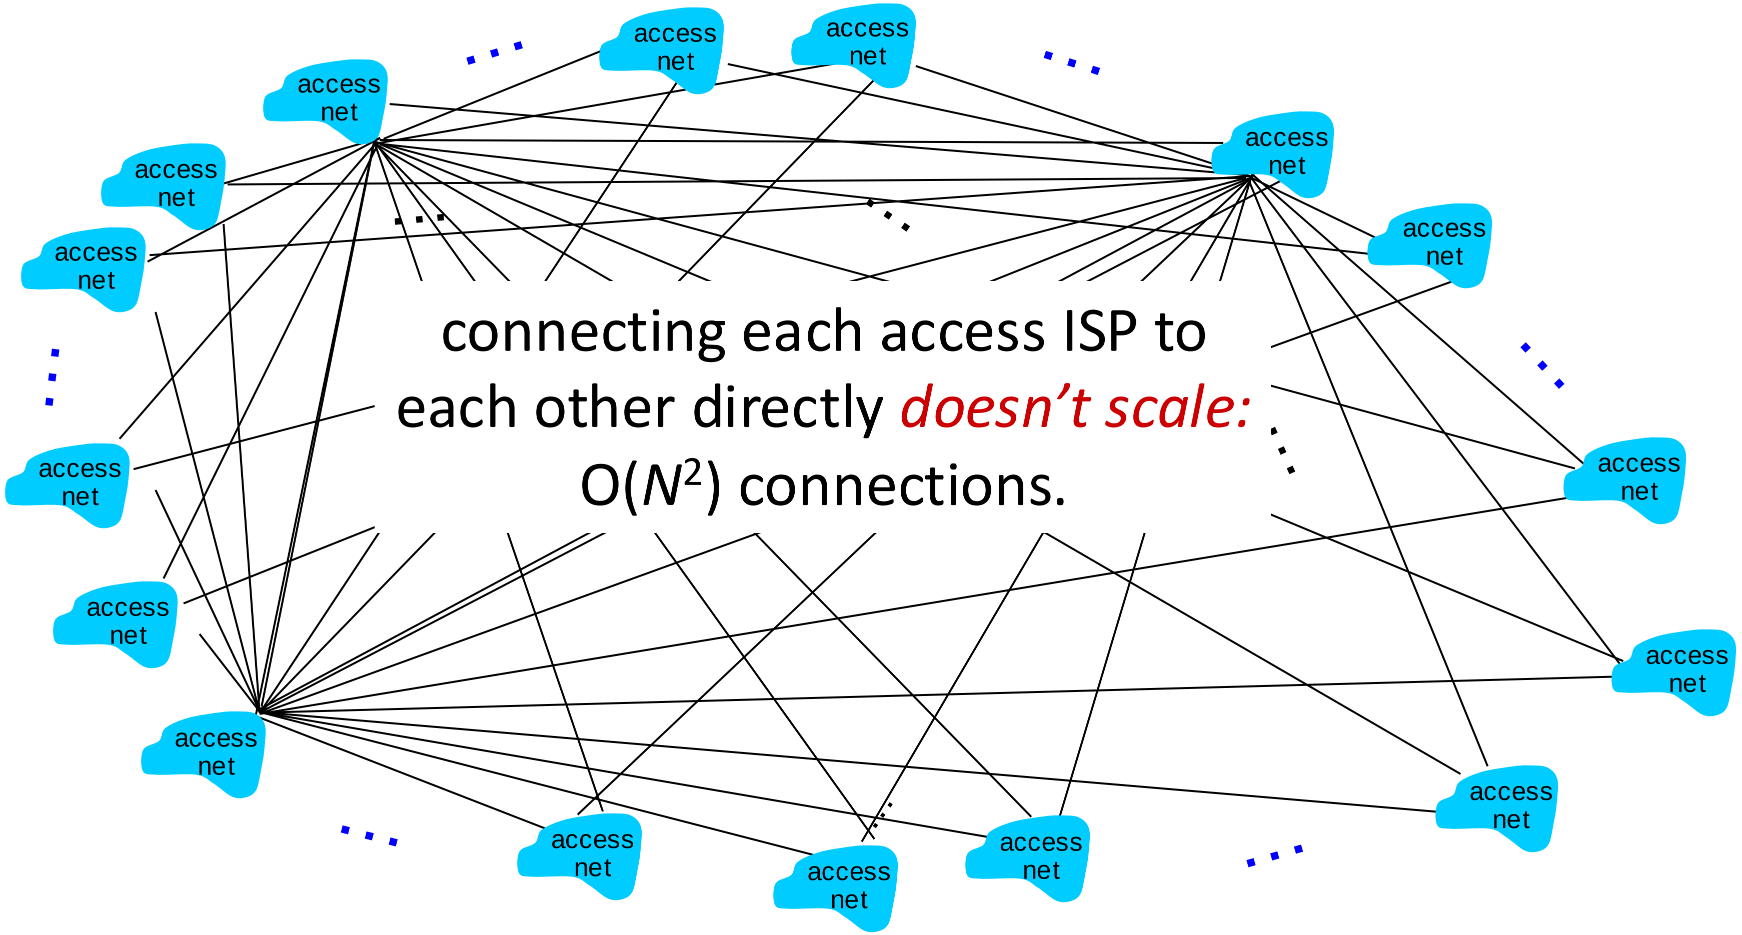

connecting each access ISP to each other directly *doesn't scale:* O(N^2) connections.

Internet structure: a "network of networks"

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

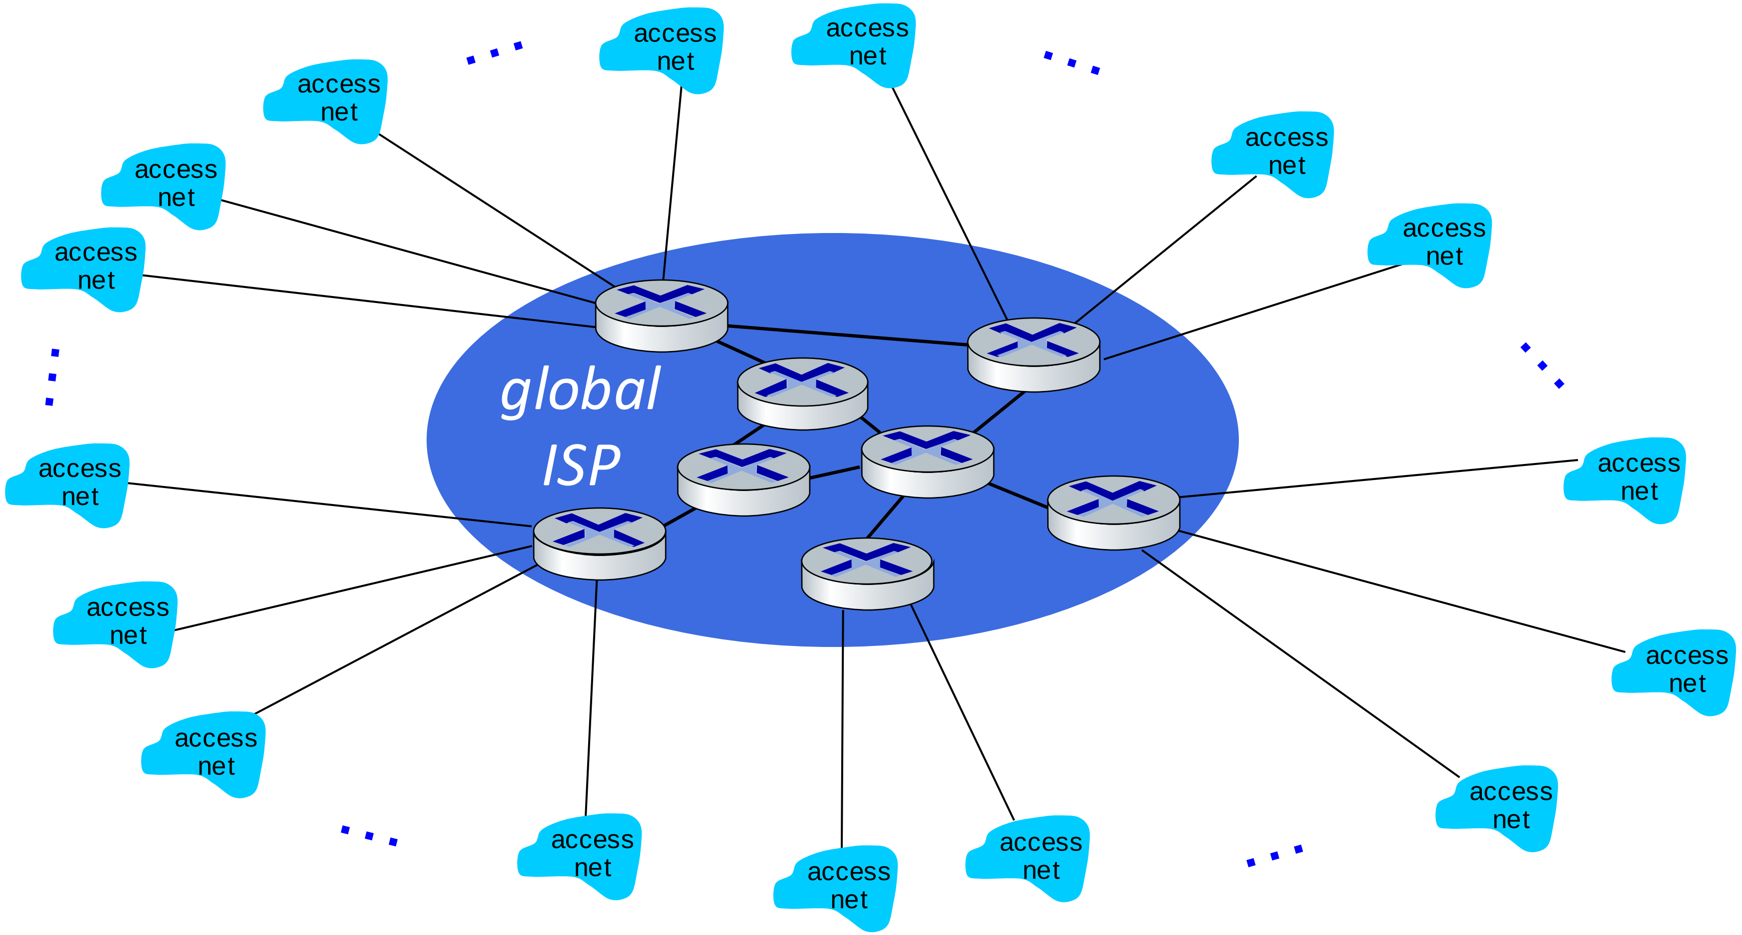

Internet structure: a "network of networks"

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

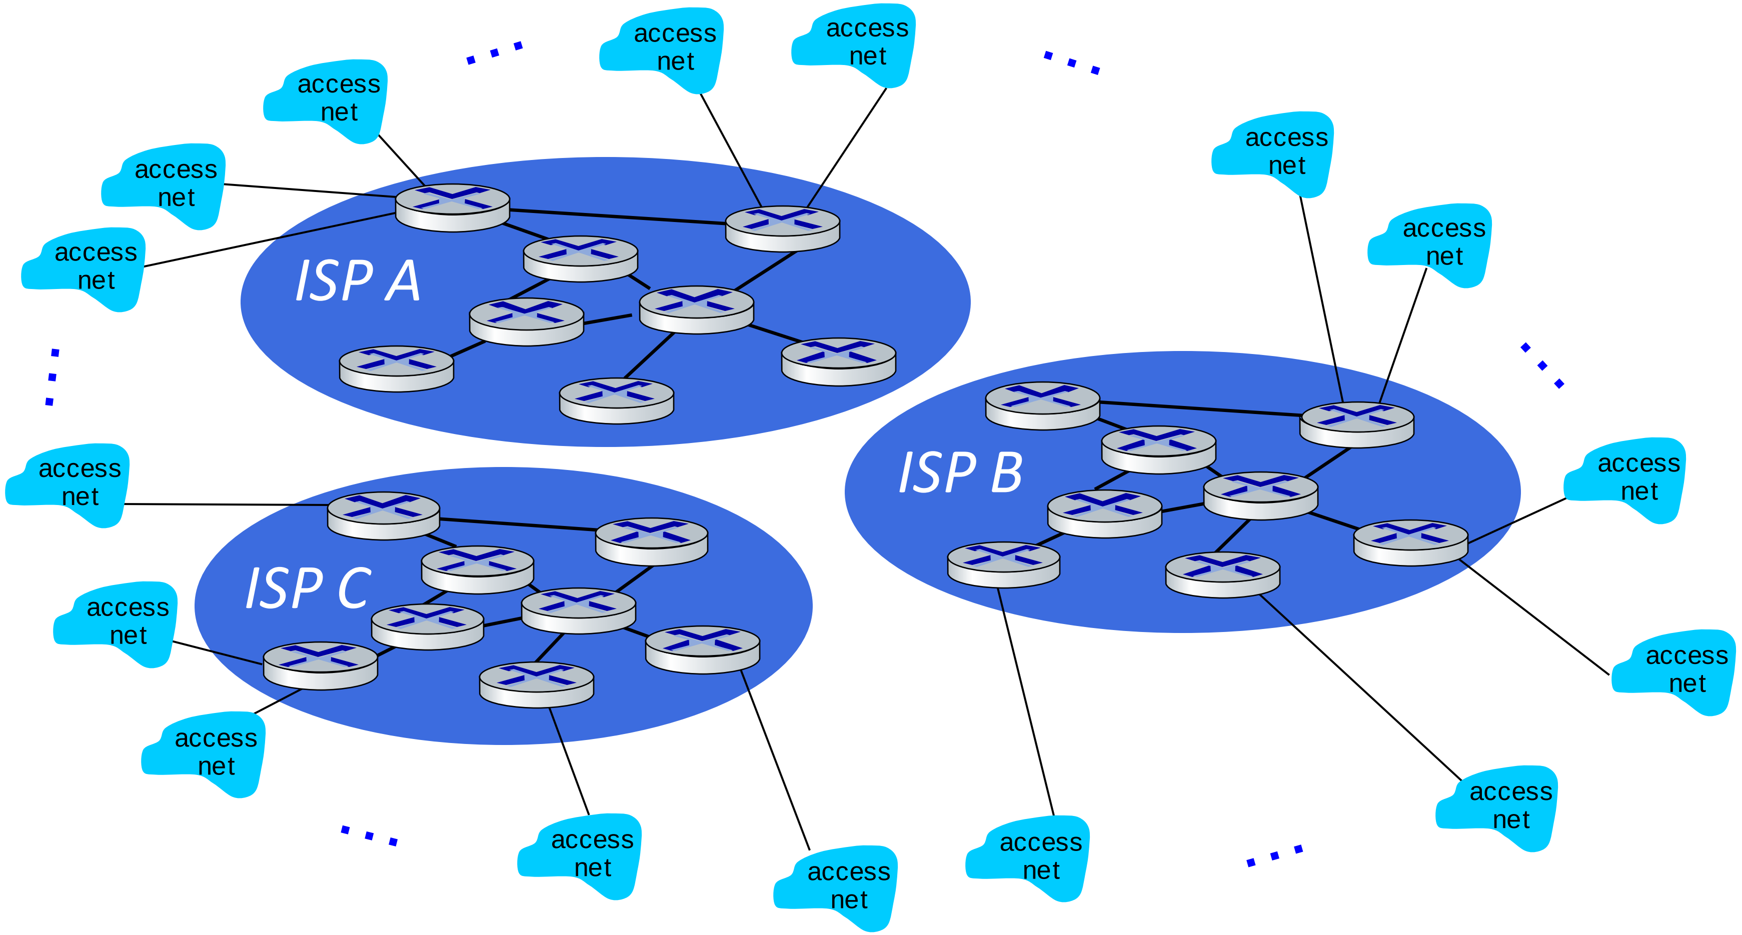

Internet structure: a "network of networks"

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

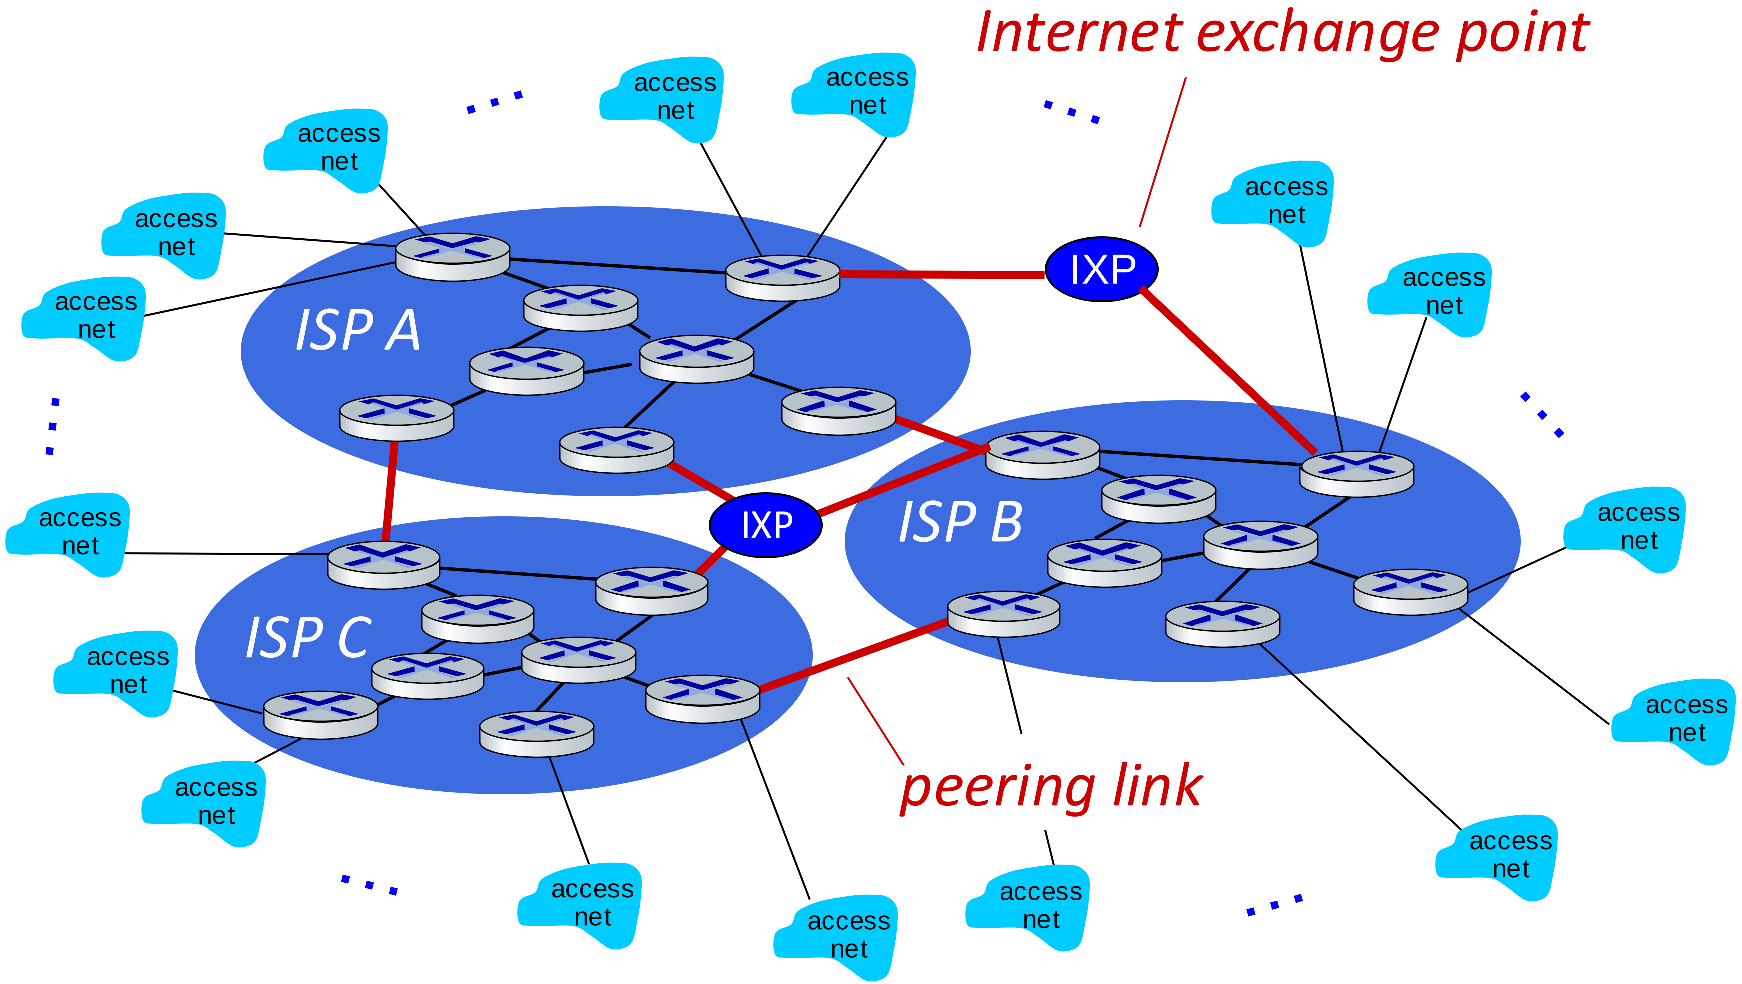

Internet structure: a "network of networks"

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

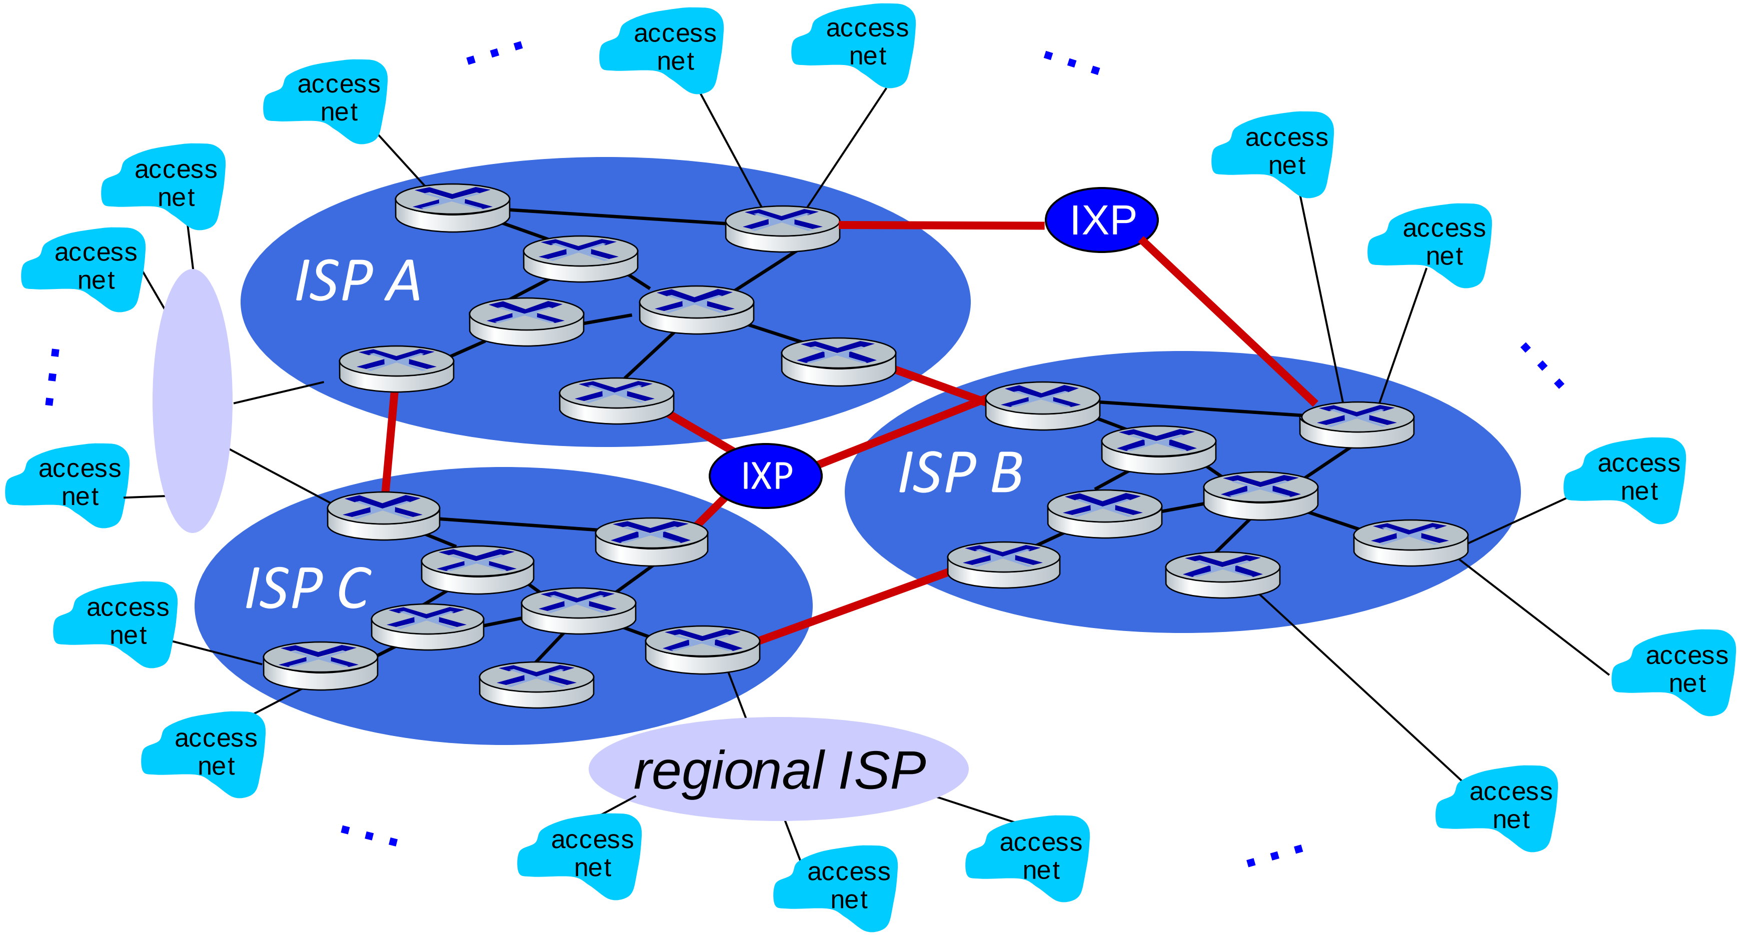

Internet structure: a "network of networks"

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

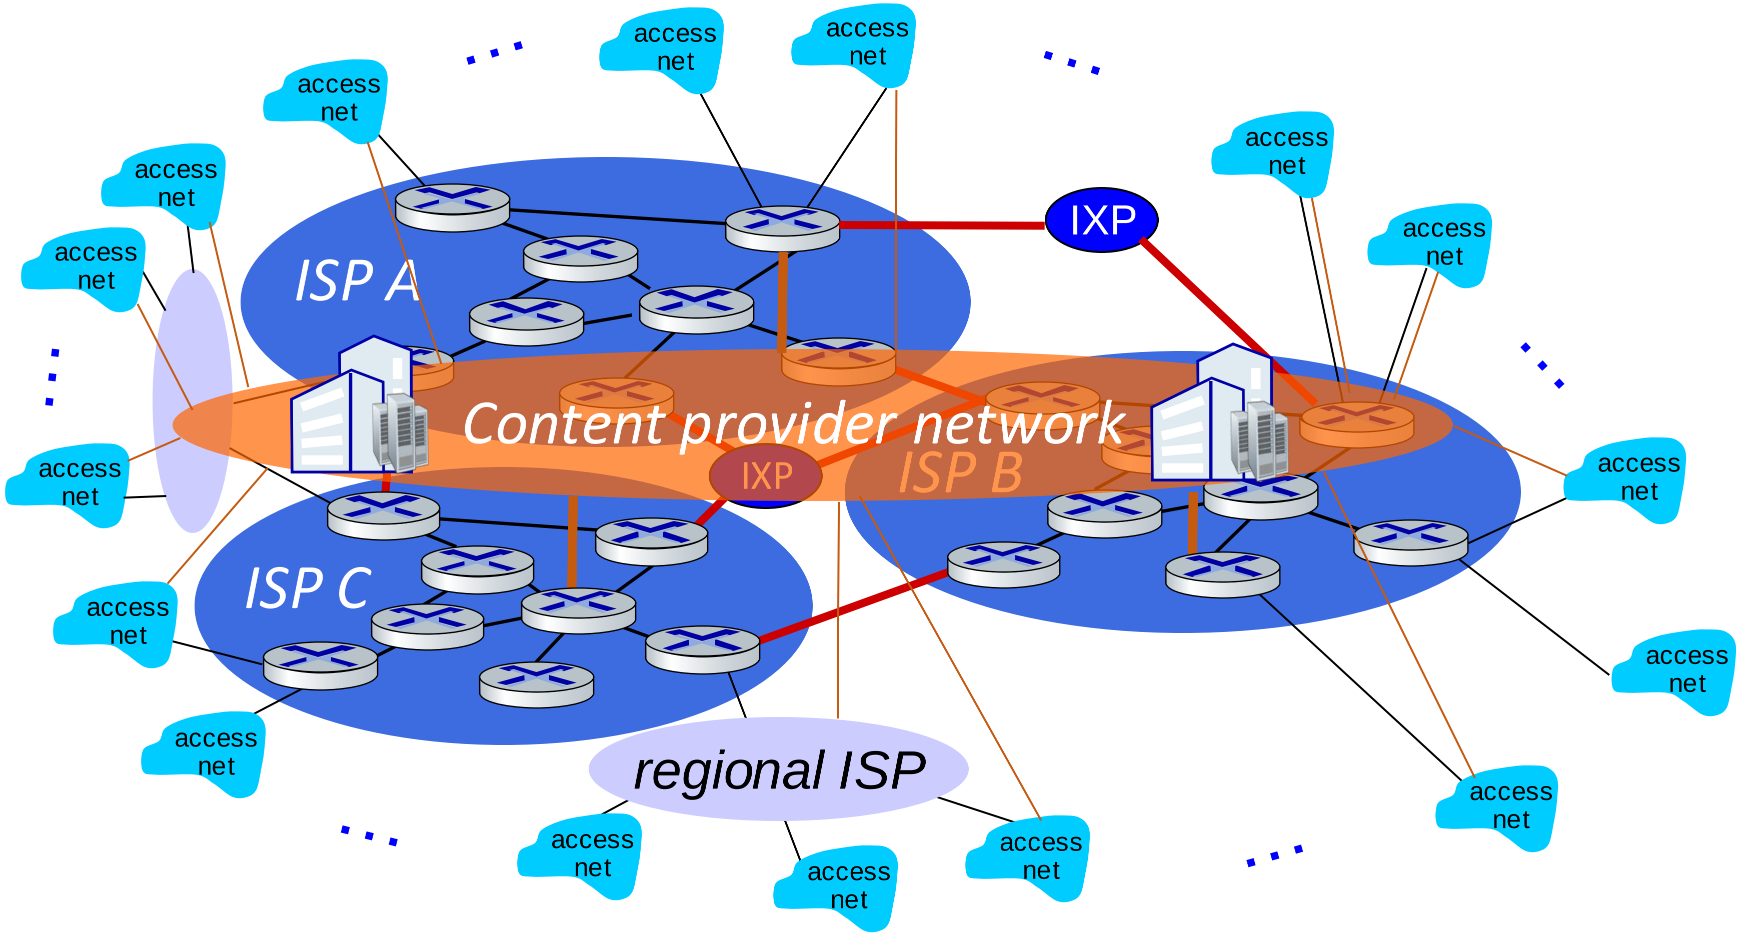

Internet structure: a "network of networks"

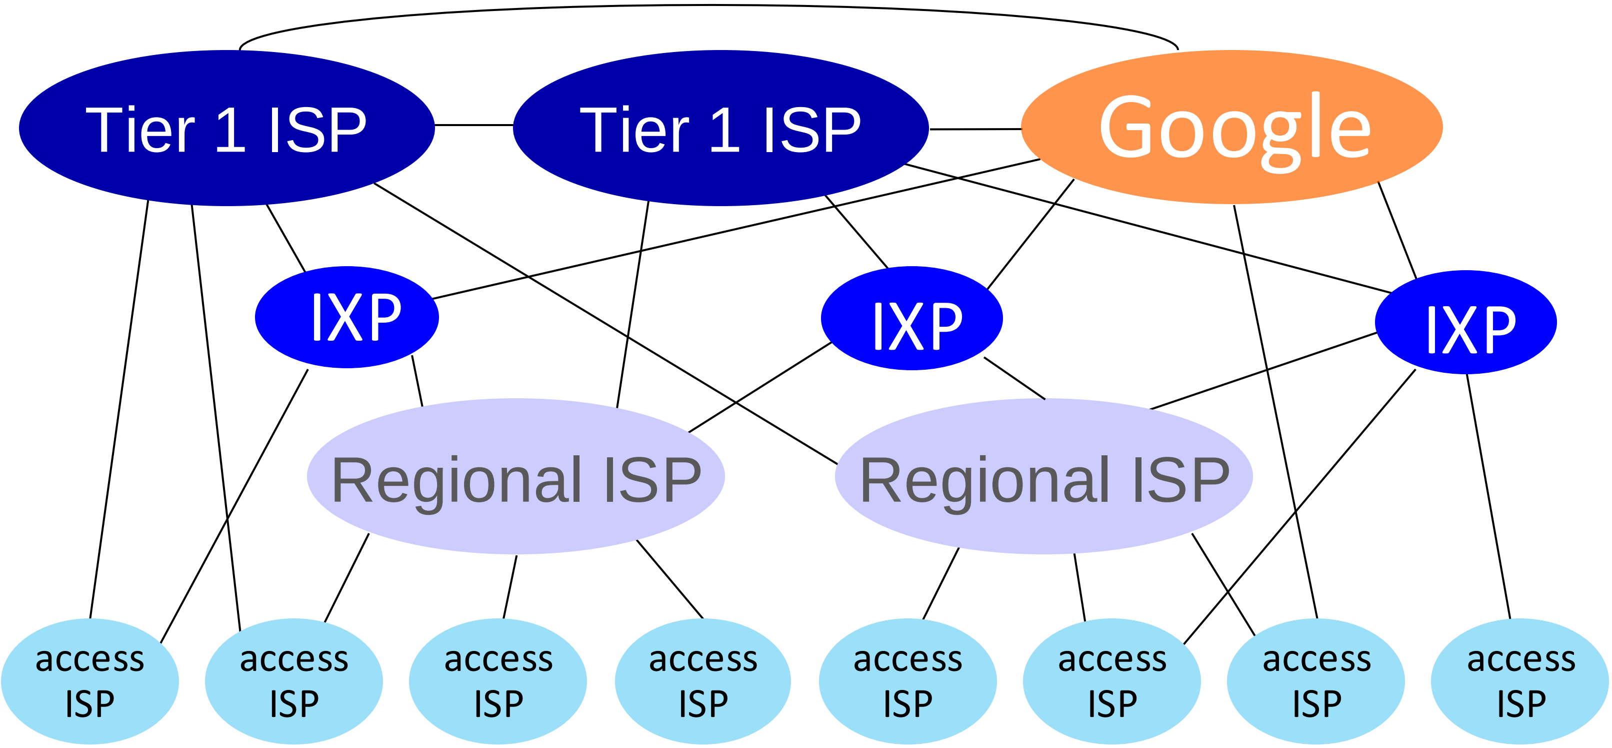

At "center": small # of large but well-connected networks

- "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

Tier-1 ISP: e.g., Sprint/T-Mobile

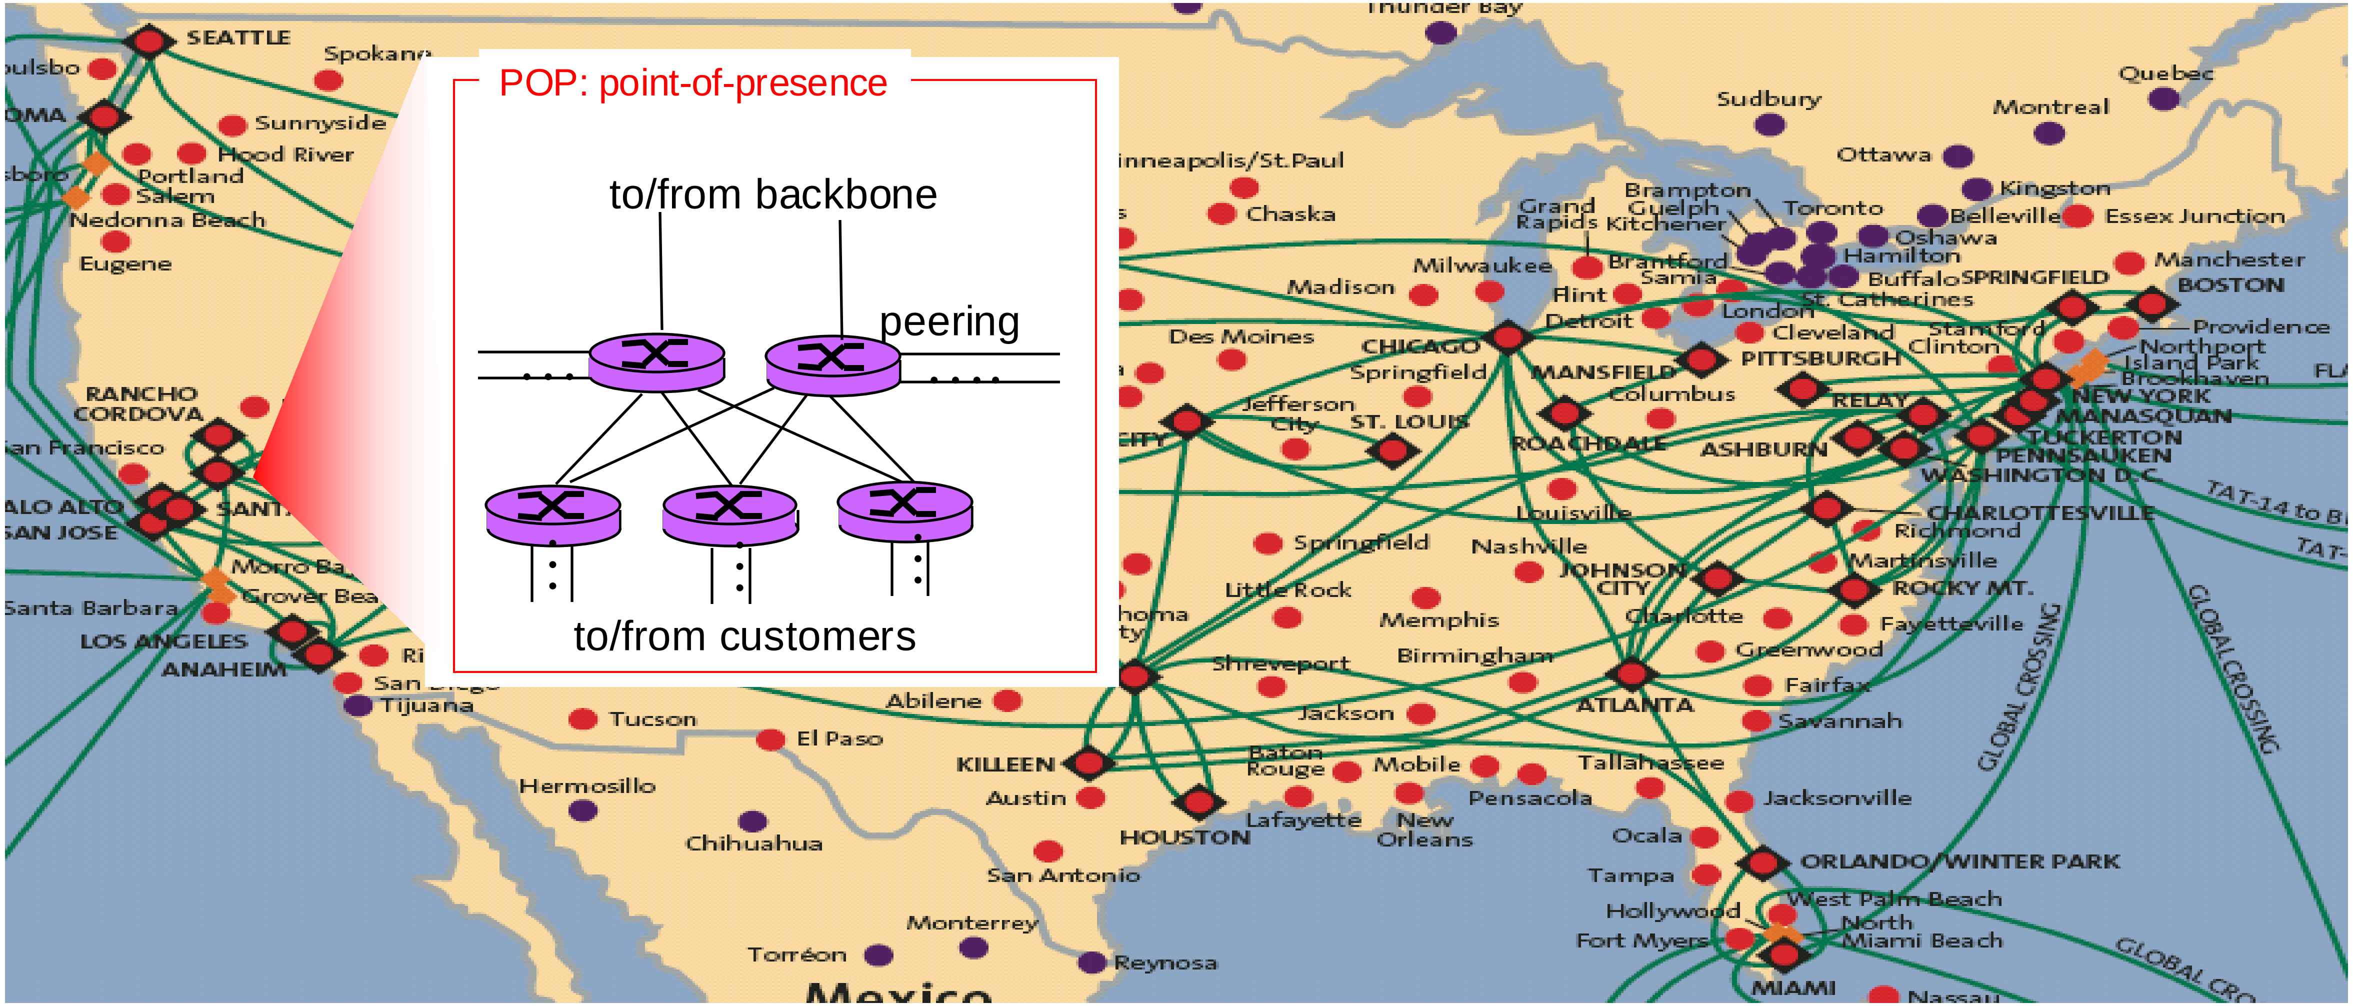

POPs from different Tier-1 ISP connect to each other at IXPs – residing at a building like this in London

Internet Core Routers (including those at POPs/IXP)

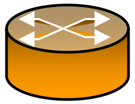

Router on "paper"

Chapter 1: roadmap

- What *is* the Internet?
- What *is* a protocol?
- Network edge: hosts, access network, physical media
- Network core: internet structure, routing and forwarding
- Performance: loss, delay, throughput
- Security
- Protocol layers, service models
- History

How to send a packet via network

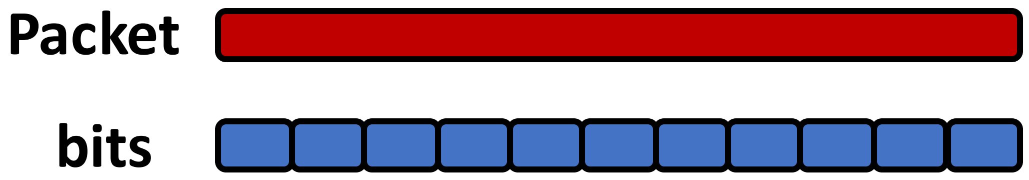

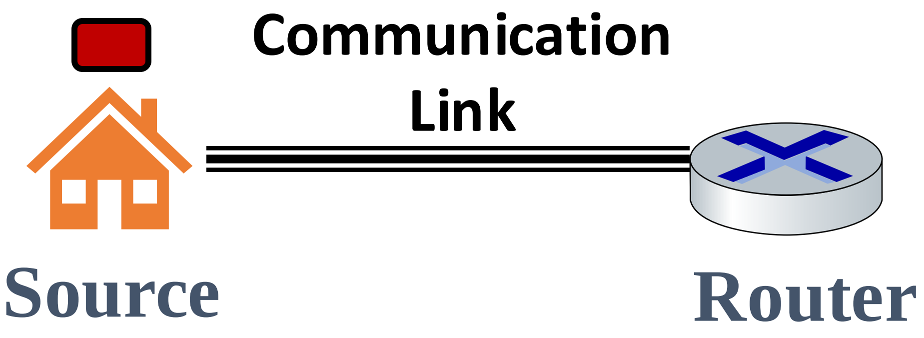

Step 1: Transmit the packets into the link

How to send a packet via network

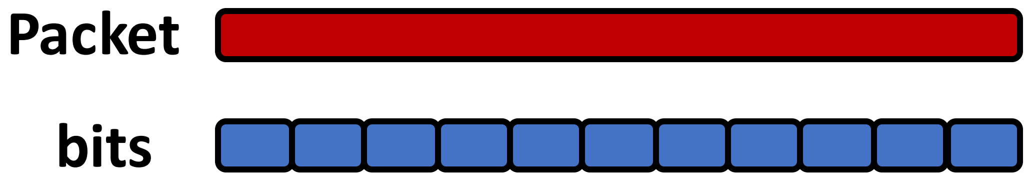

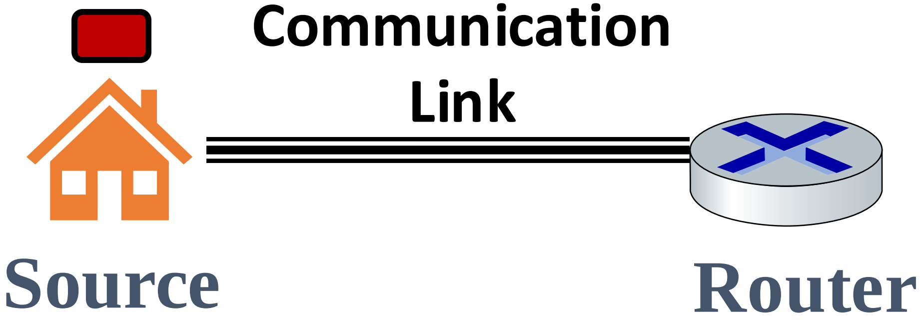

Step 1: Transmit the packets into the link

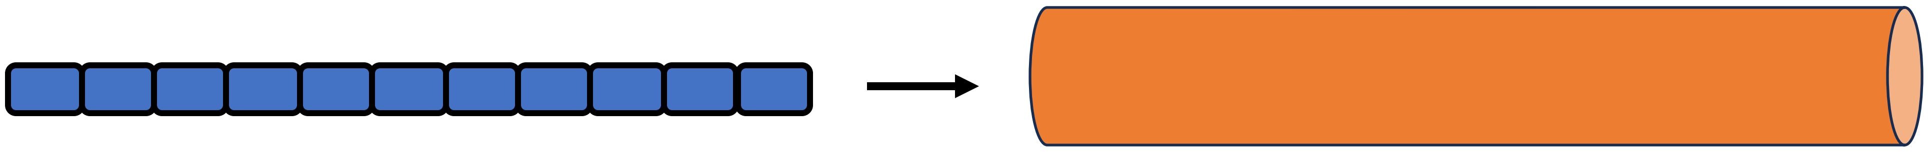

How to send a packet via network

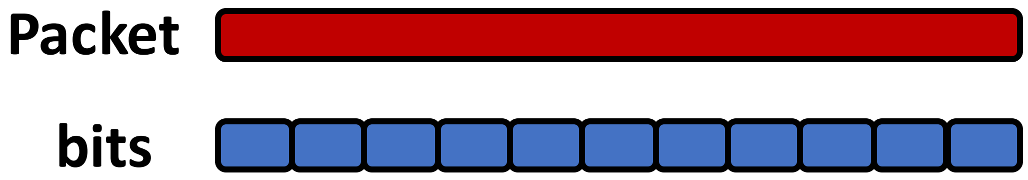

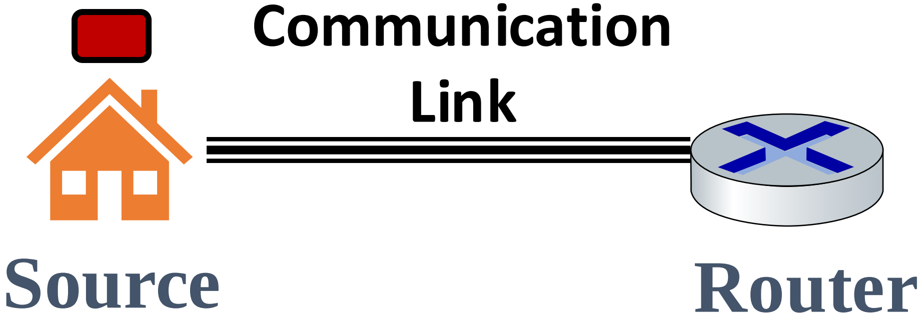

Step 1: Transmit the packets into the link

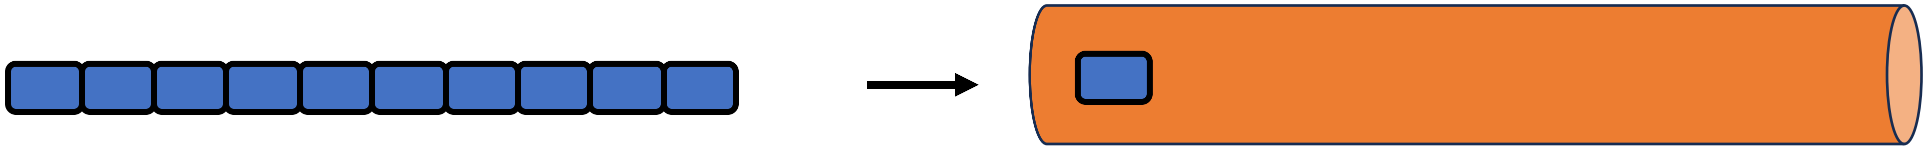

Link capacity: 10 bit/sec

How to send a packet via network

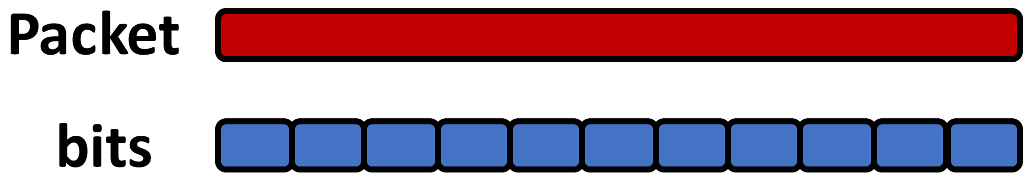

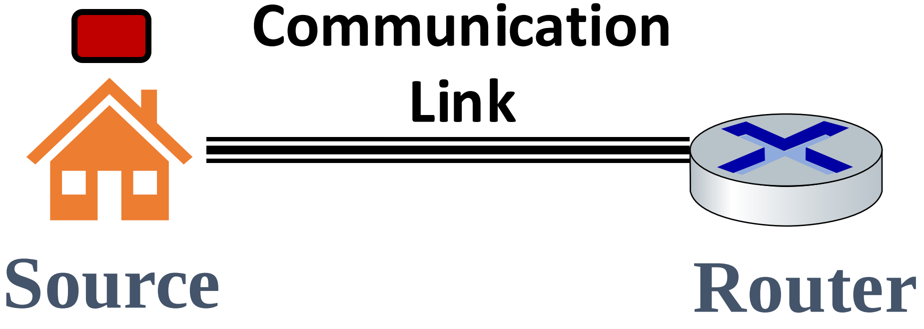

Step 1: Transmit the packets into the link

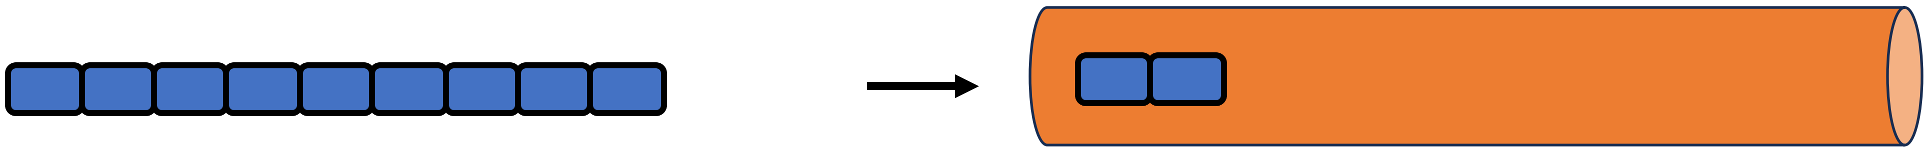

Link capacity: 10 bit/sec

How to send a packet via network

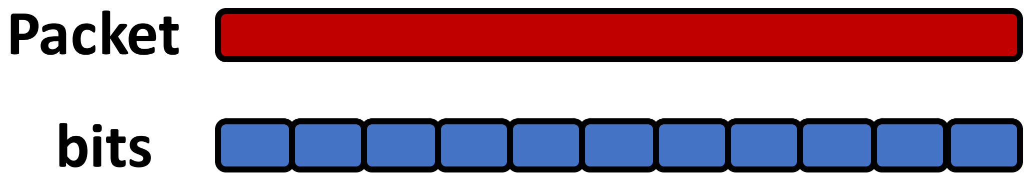

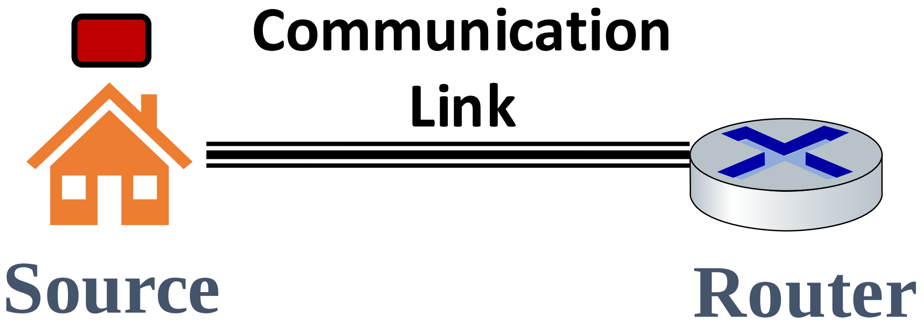

Step 1: Transmit the packets into the link

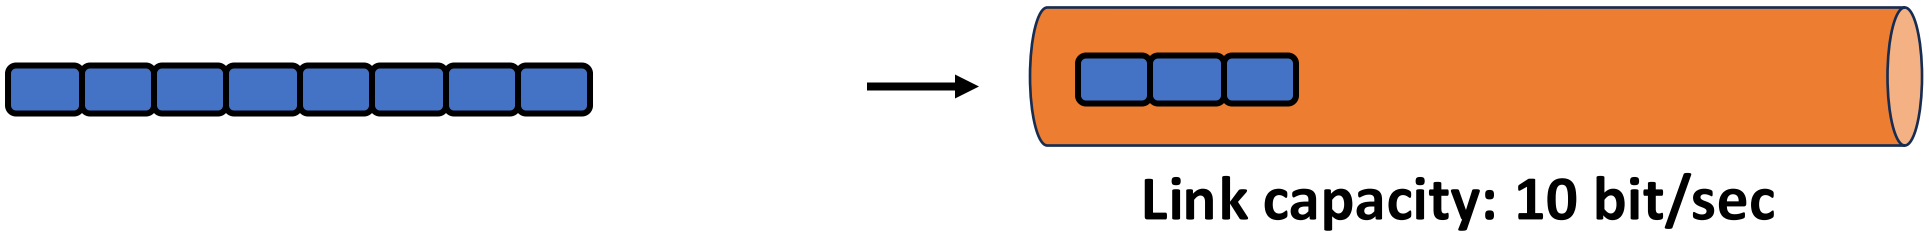

How to send a packet via network

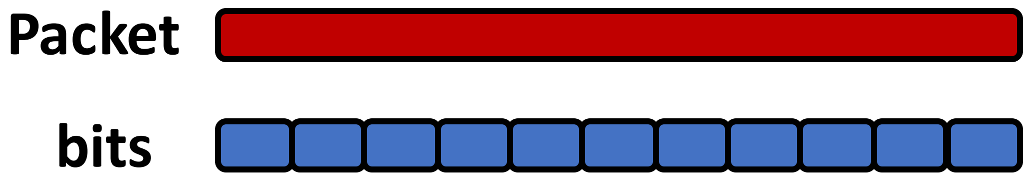

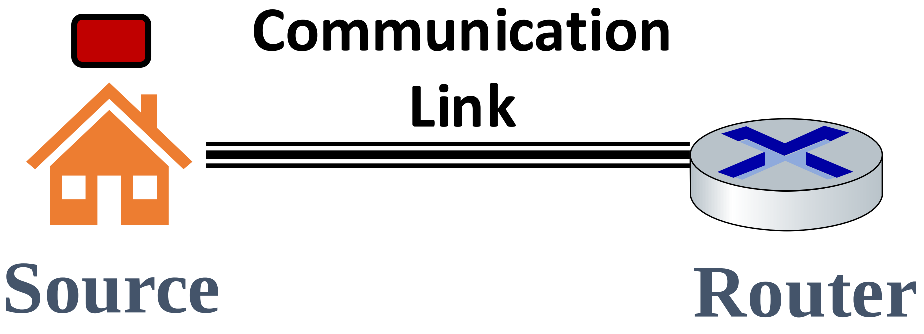

Step 1: Transmit the packets into the link

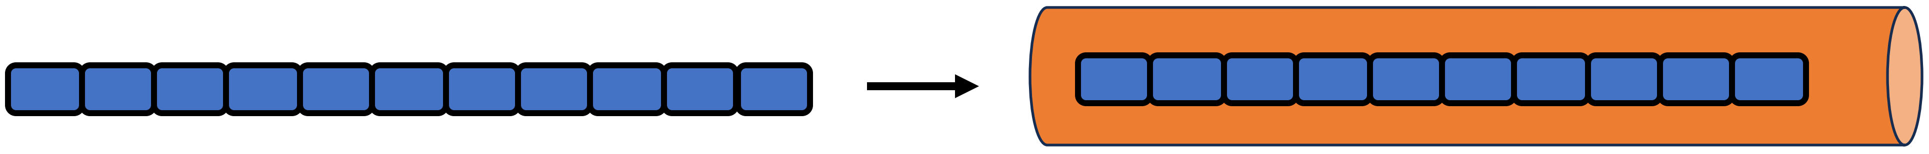

Link capacity: 10 bit/sec

How to send a packet via network

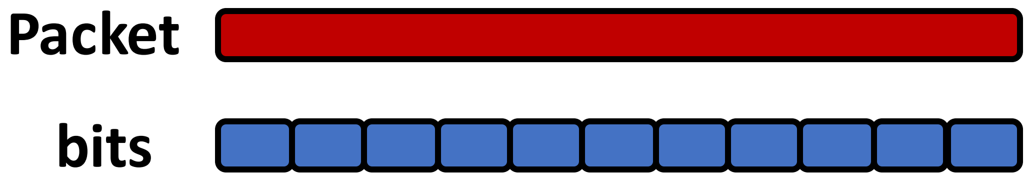

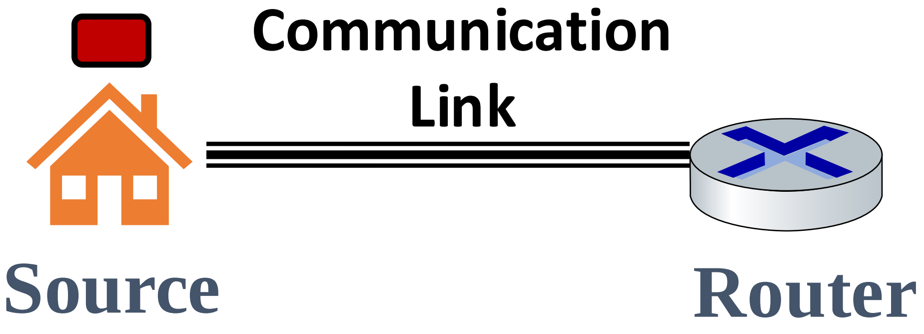

Step 1: Transmit the packets into the link

d_{trans}: transmission delay:

- L: packet length (bits)
- R: link *transmission rate (bps)*
- $d_{trans} = L/R$

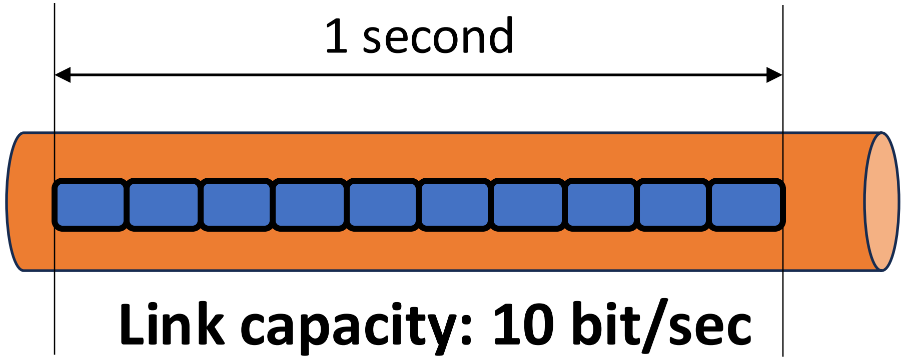

How to send a packet via network

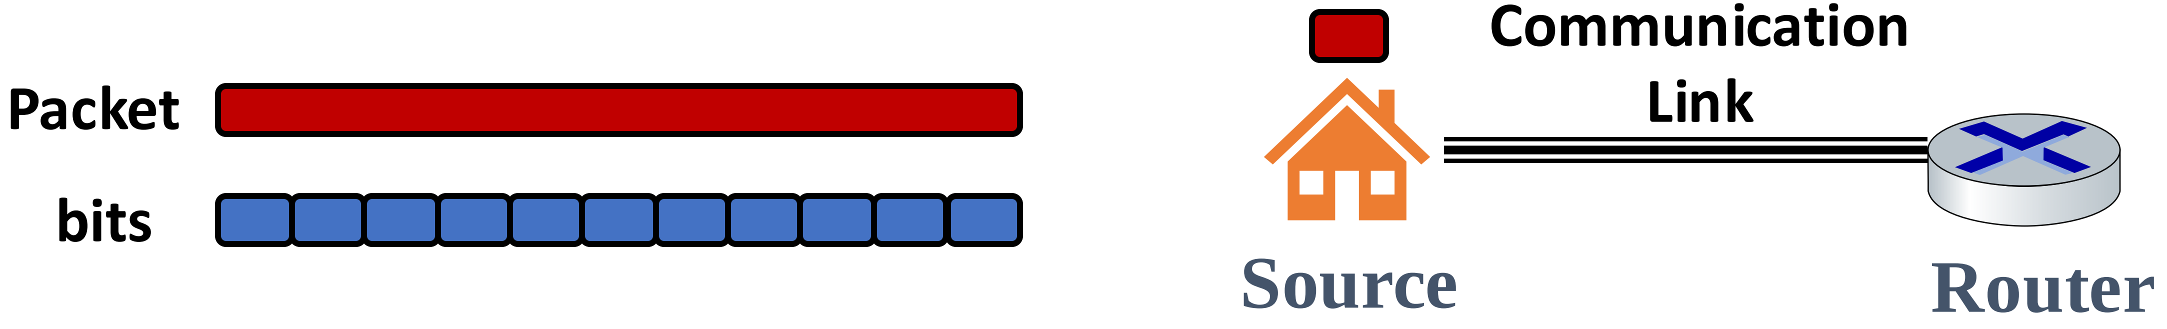

Step 1: Transmit the packets into the link

Step 2: The packet bits propagates to the router

How to send a packet via network

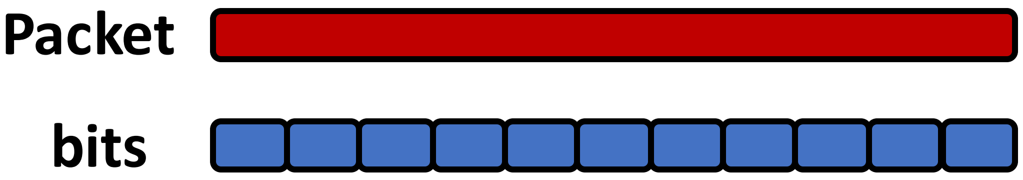

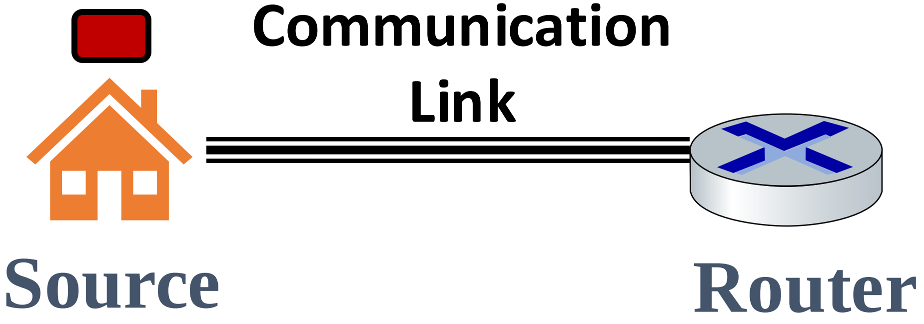

Step 1: Transmit the packets into the link

Step 2: The packet bits propagates to the router

d_{prop}: propagation delay:

- d: length of physical link
- s: propagation speed (~2×10^8 m/sec
- $d_{prop} = d/s$

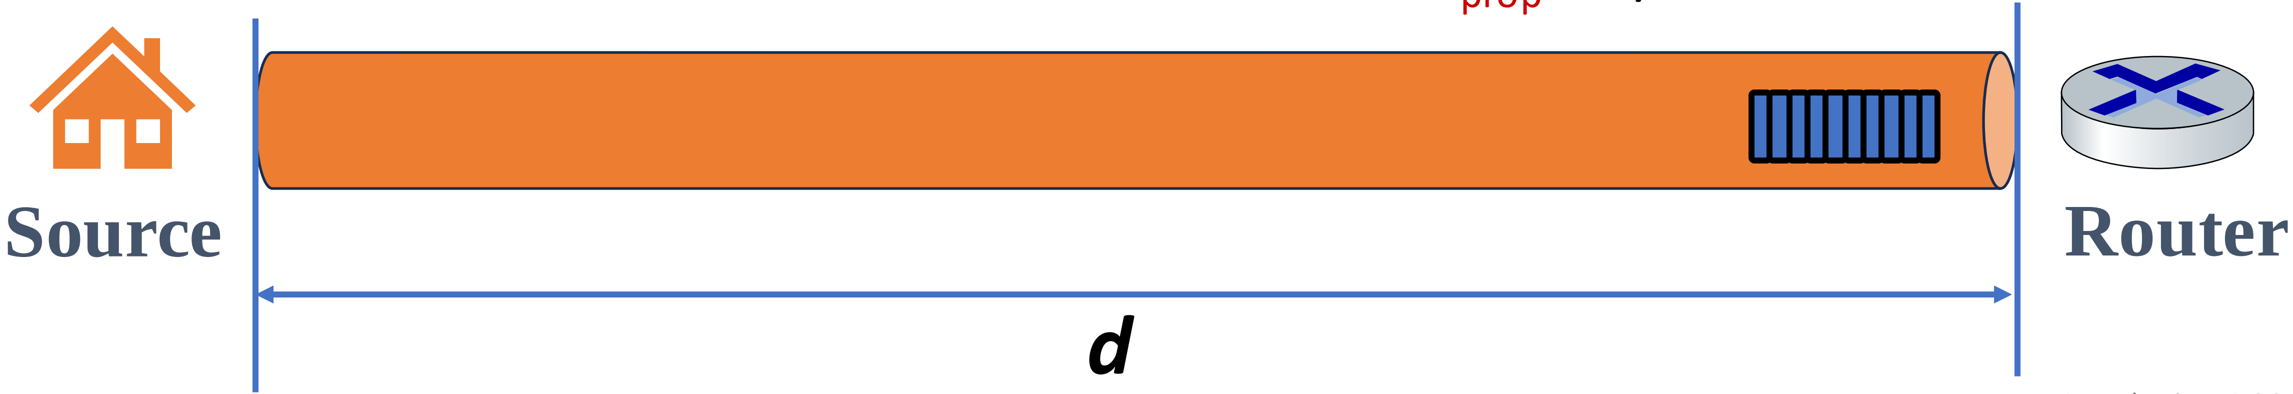

How to send a packet via network

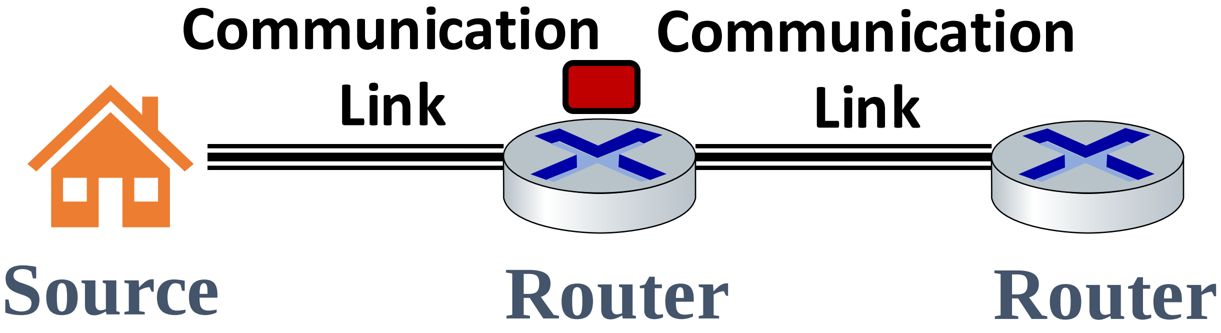

Step 1: Transmit the packets into the link

Step 2: The packet bits propagates to the router

d_{trans}: transmission delay:

- L: packet length (bits)
- R: link *transmission rate (bps)*
- $d_{trans} = L/R$

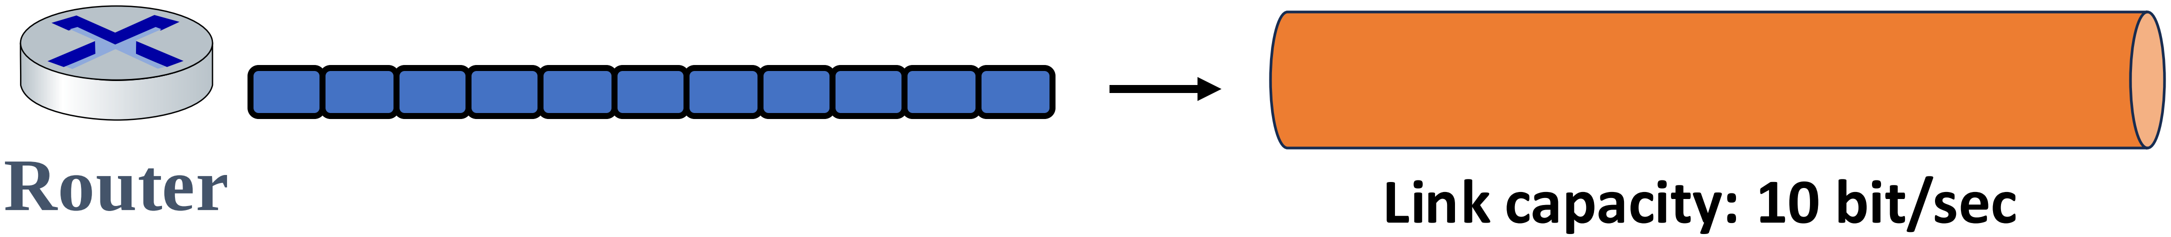

Link capacity: 10 bit/sec

How to send a packet via network

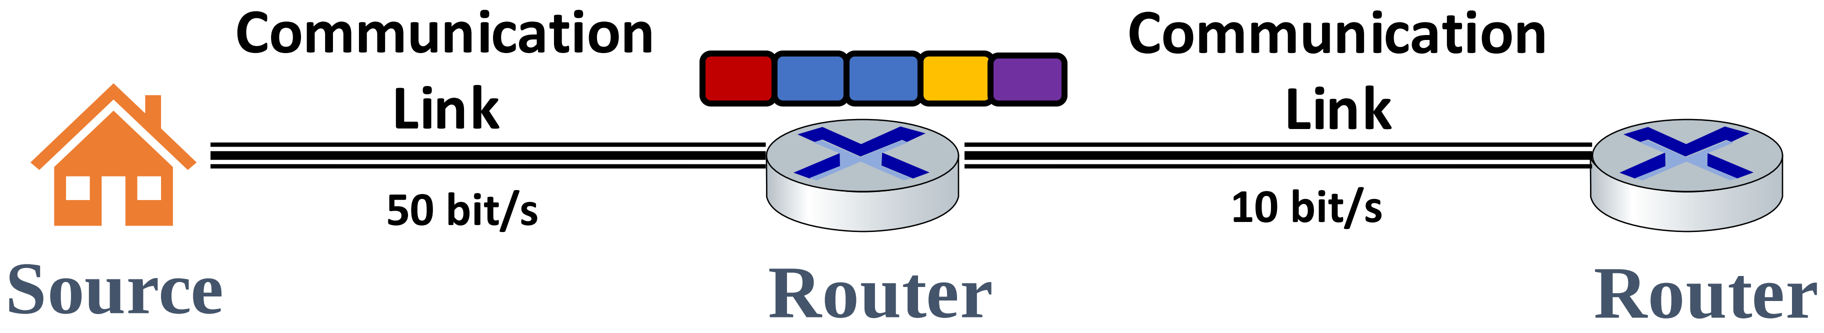

Key point:

- Router takes transmission delay to transmit a packet to the link
- The packet may arrive faster than the packets get out of the router
- The later arrived packets must wait at the router until all the packets arriving before it are transmitted into the link

d_{queue}: queueing delay

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

How to send a packet via network

Various reasons of queuing inside the router

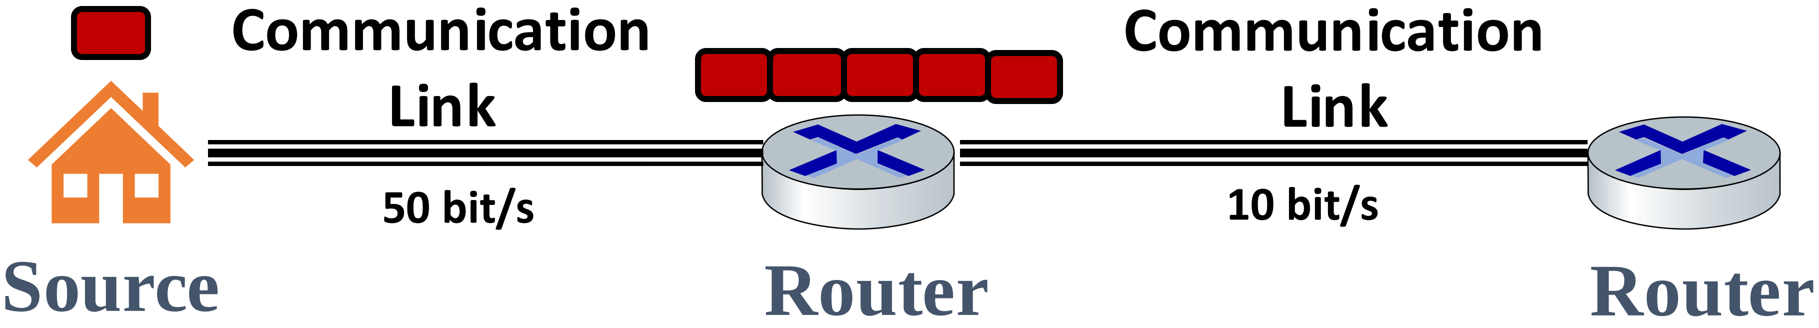

How to send a packet via network

Various reasons of queuing inside the router

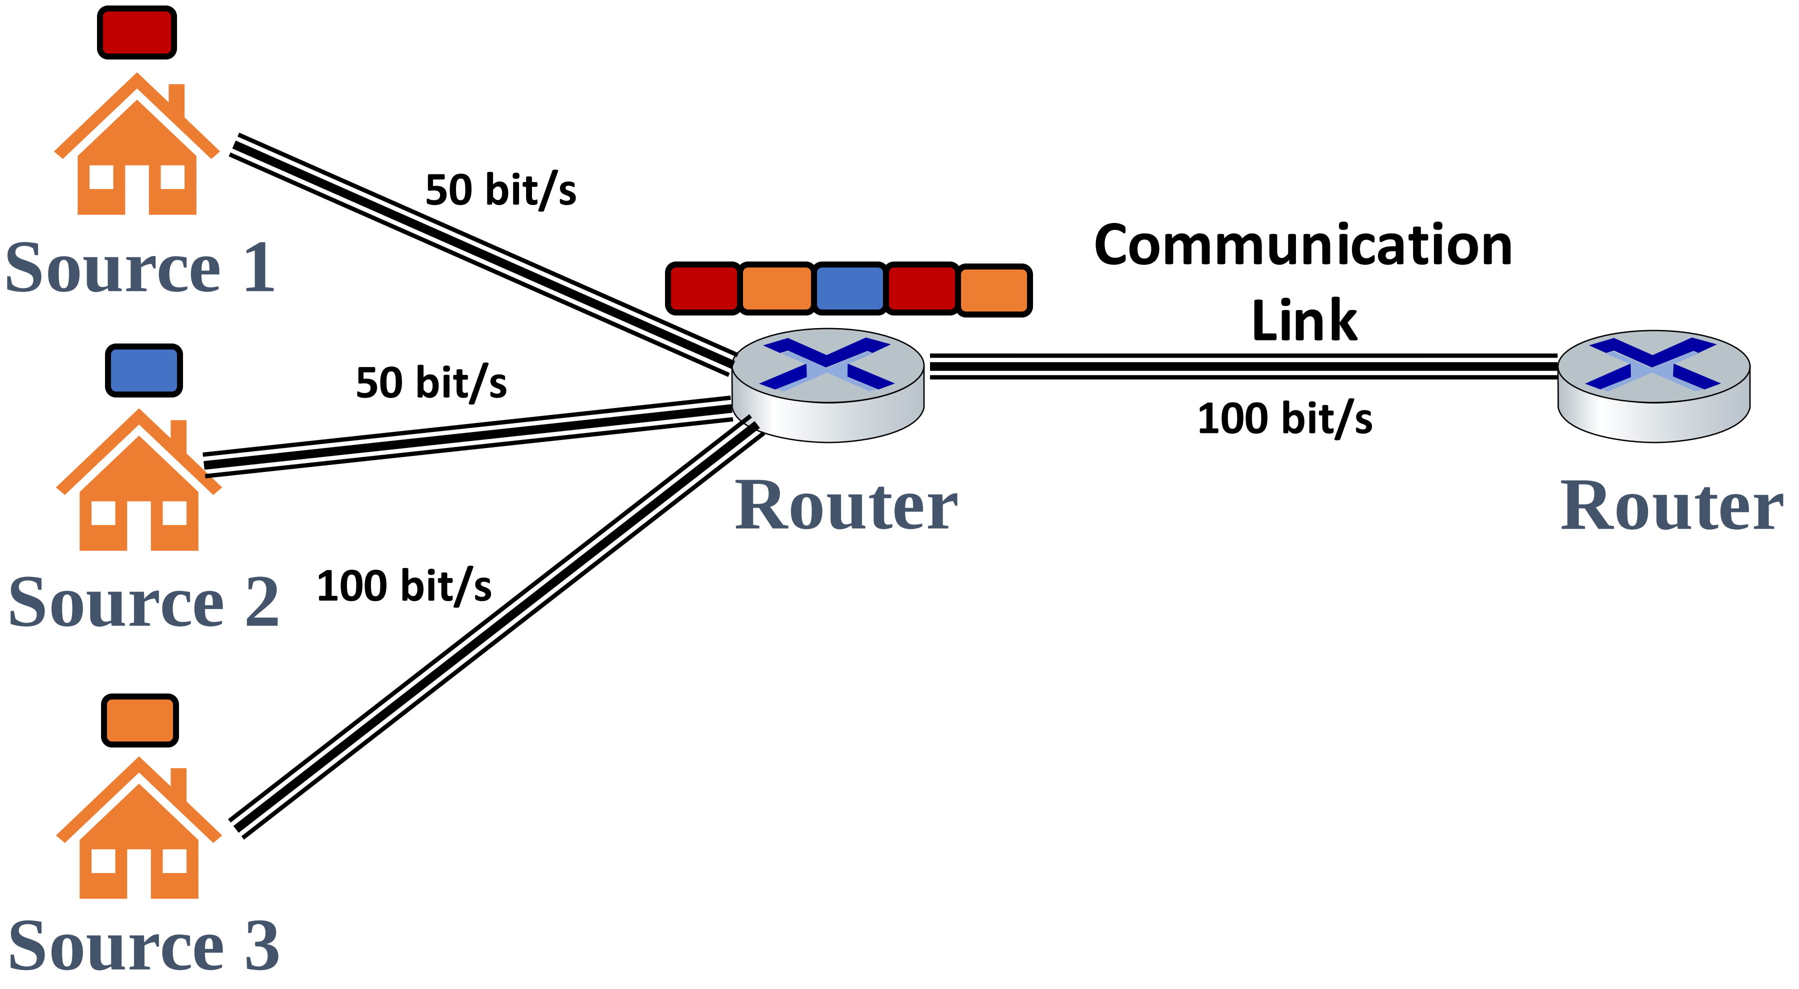

Packet queueing delay (revisited)

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

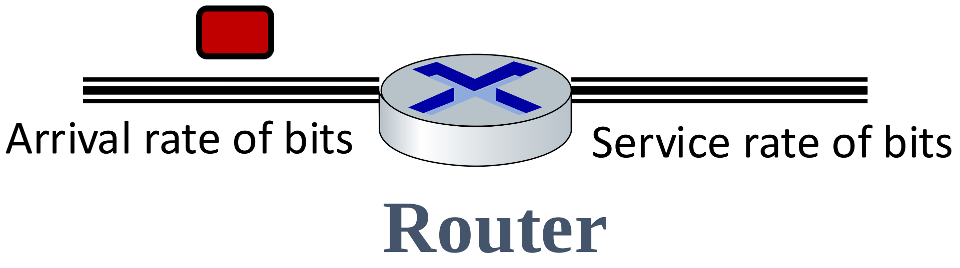

$$\frac{L \cdot a}{R} : \frac{\text{arrival rate of bits}}{\text{service rate of bits}}$$ *"traffic intensity"*

- *La/R* ~ 0: avg. queueing delay small
- *La/R* -> 1: avg. queueing delay large
- *La/R* > 1: more "work" arriving is more than can be serviced - average delay infinite!

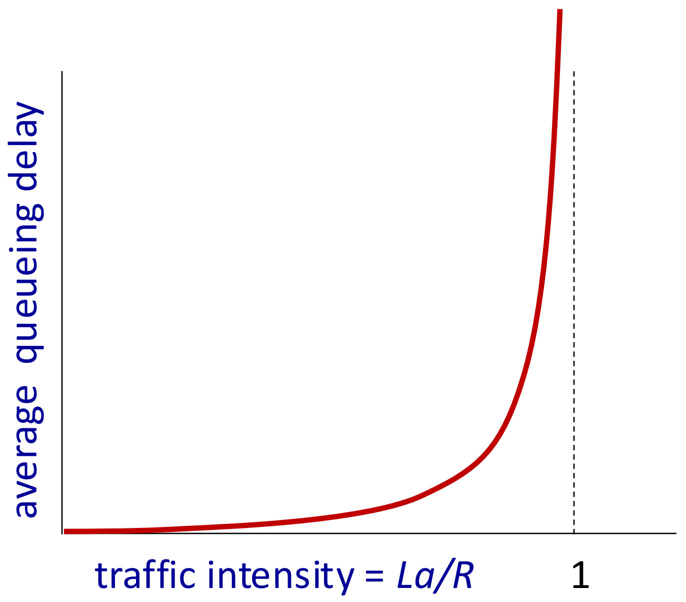

"Real" Internet delays and routes

- what do "real" Internet delay & 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 (with time-to-live field value of *i*)
  - router *i* will return packets to sender
  - sender measures time interval between transmission and reply

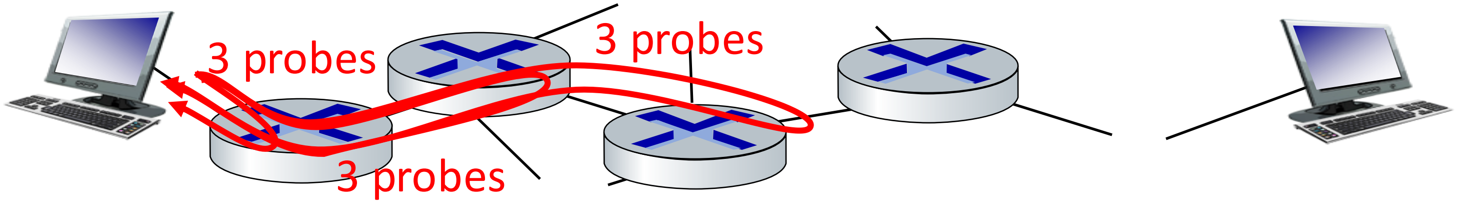

Real Internet delays and routes

traceroute: gaia.cs.umass.edu to www.eurecom.fr

3 delay measurements from
gaia.cs.umass.edu to cs-gw.cs.umass.edu

3 delay measurements
to border1-rt-fa5-1-0.gw.umass.edu

```
1  cs-gw (128.119.240.254)  1 ms  1 ms  2 ms
2  border1-rt-fa5-1-0.gw.umass.edu (128.119.3.145)  1 ms  1 ms  2 ms
3  cht-vbns.gw.umass.edu (128.119.3.130)  6 ms 5 ms 5 ms
4  jn1-at1-0-0-19.wor.vbns.net (204.147.132.129)  16 ms 11 ms 13 ms
5  jn1-so7-0-0-0.wae.vbns.net (204.147.136.136)  21 ms 18 ms 18 ms
6  abilene-vbns.abilene.ucaid.edu (198.32.11.9)  22 ms  18 ms  22 ms
7  nycm-wash.abilene.ucaid.edu (198.32.8.46)  22 ms  22 ms  22 ms
8  62.40.103.253 (62.40.103.253)  104 ms 109 ms 106 ms
9  de2-1.de1.de.geant.net (62.40.96.129)  109 ms 102 ms 104 ms
10  de.fr1.fr.geant.net (62.40.96.50)  113 ms 121 ms 114 ms
11  renater-gw.fr1.fr.geant.net (62.40.103.54)  112 ms  114 ms  112 ms
12  nio-n2.cssi.renater.fr (193.51.206.13)  111 ms  114 ms  116 ms
13  nice.cssi.renater.fr (195.220.98.102)  123 ms  125 ms  124 ms
14  r3t2-nice.cssi.renater.fr (195.220.98.110)  126 ms  126 ms  124 ms
15  eurecom-valbonne.r3t2.ft.net (193.48.50.54)  135 ms  128 ms  133 ms
16  194.214.211.25 (194.214.211.25)  126 ms  128 ms  126 ms
17  * * *
18  * * *
19  fantasia.eurecom.fr (193.55.113.142)  132 ms  128 ms  136 ms
```

trans-oceanic link

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

* Do some traceroutes from exotic countries at www.traceroute.org

Packet loss

- queue (aka buffer) preceding link in buffer has finite capacity
- packet arriving at a full queue dropped (aka lost)
- lost packet may be retransmitted by previous node, by source end system, or not at all

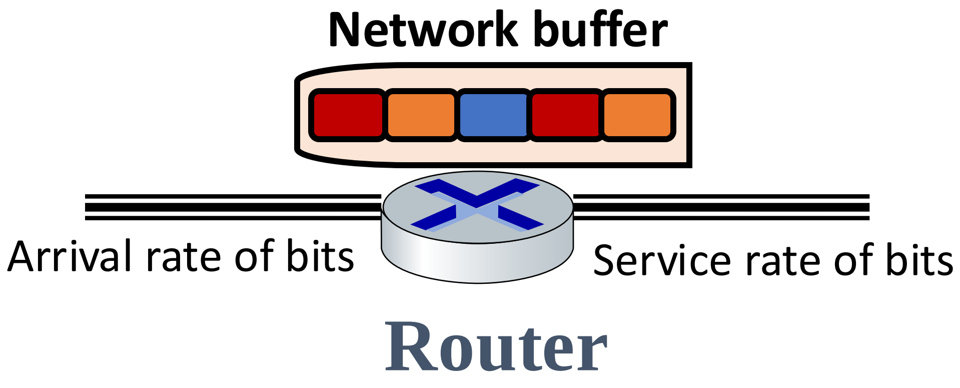

* Check out the Java applet for an interactive animation (on publisher's website) of queuing and loss

Throughput

- *throughput:* rate (bits/time unit) at which bits are being sent from sender to receiver
  - *instantaneous:* rate at a given point in time
  - *average:* rate over longer period of time

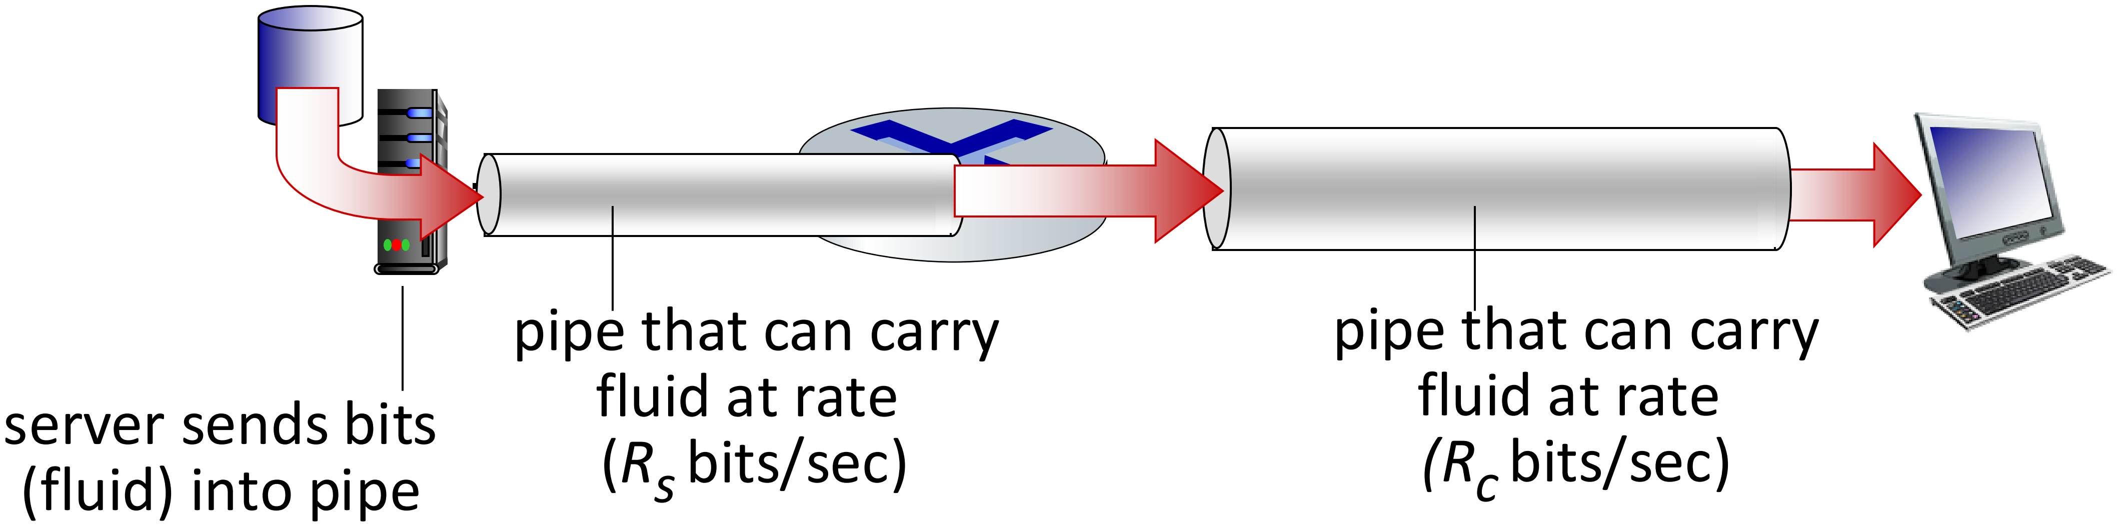

Throughput

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

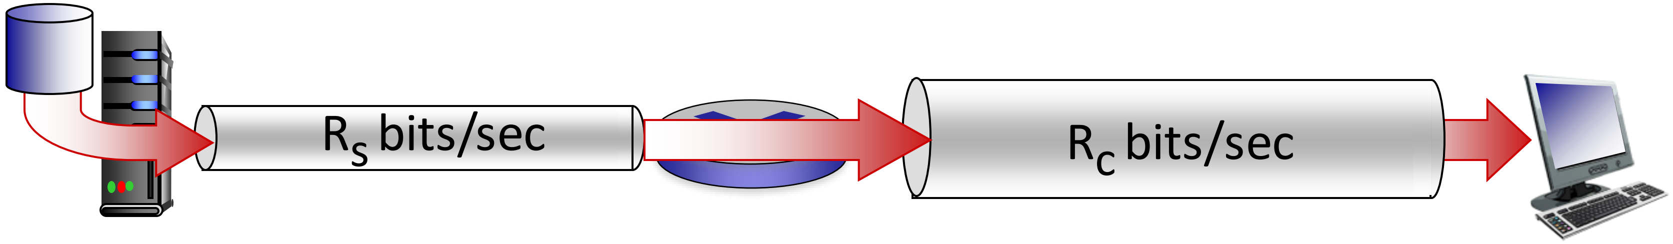

$R_s > R_c$ What is average end-end throughput?

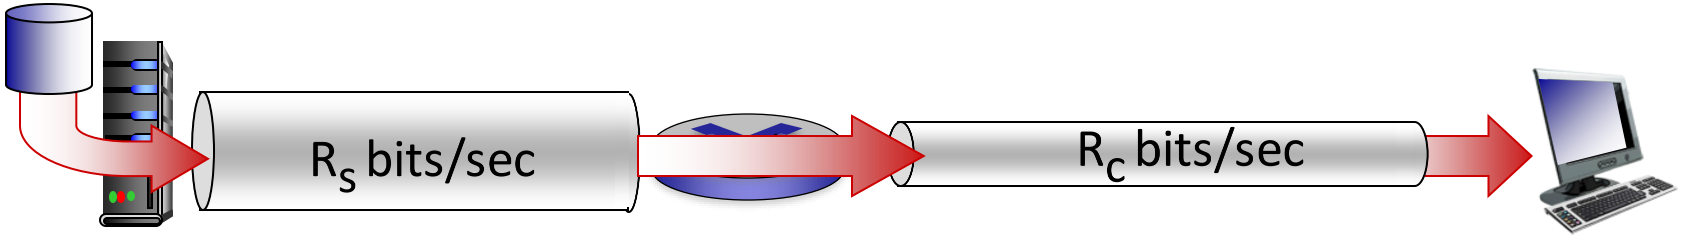

bottleneck link

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

Throughput: network scenario

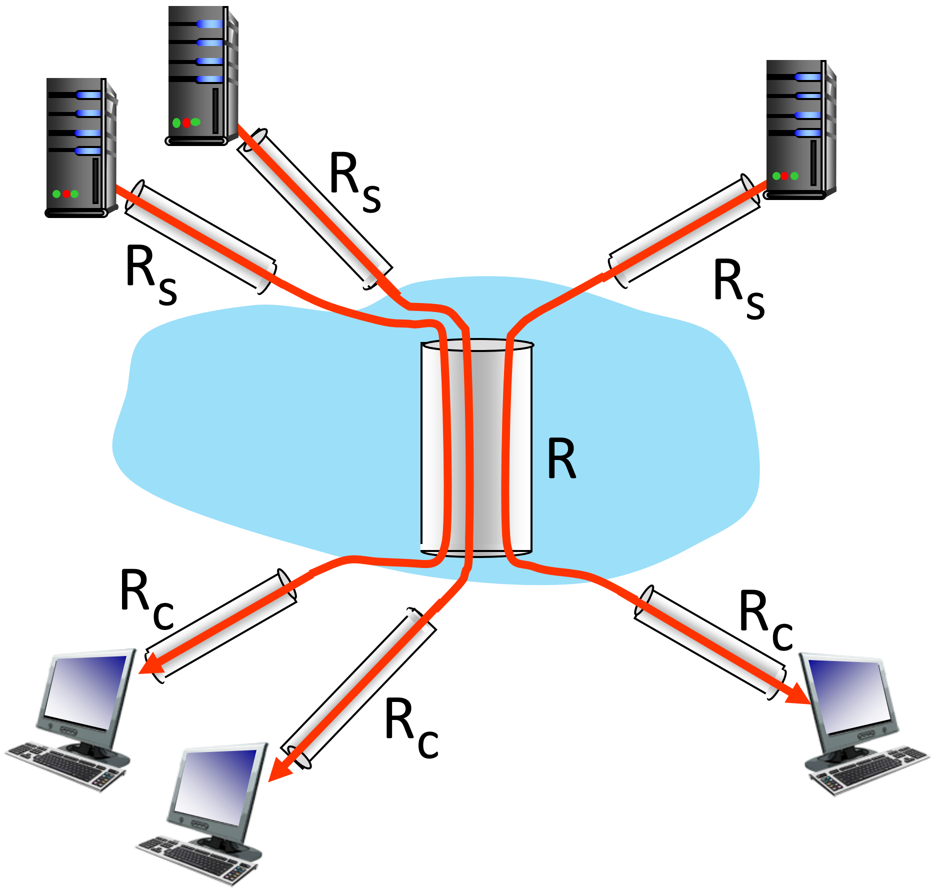

10 connections (fairly) share backbone bottleneck link *R* bits/sec

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

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