Lecture 02: Packets, Routing, and Performance

CS 356 Computer Networks

Mikyung Han

mhan@cs.utexas.edu

Please, interrupt and ask questions AT ANY TIME !

Perusall

- Add your UTEID to student ID in your Perusall setting
- Make 4+ “high-quality” comments
  - Upvotes
  - Getting responses from other students
  - Reading till the end
  - Active engagement time

Recap: Layers and protocols are basic building blocks

Note: Protocols are horizontal and layers are vertical

5

Example Protocols	Layers	Responsible for
FTP, HTTP, SMTP	Application	application specific needs
TCP, UDP	Transport	process to process data transfer
IP	Network	host to host data transfer across different network
Ethernet, WiFi	Link	data transfer between physically adjacent nodes
802.3 PHY	Physical	bit-by-bit or symbol-by-symbol delivery

Layer N provides service to Layer N+1
is serviced by Layer N-1

Protocols provides ways for peers to communicate horizontally

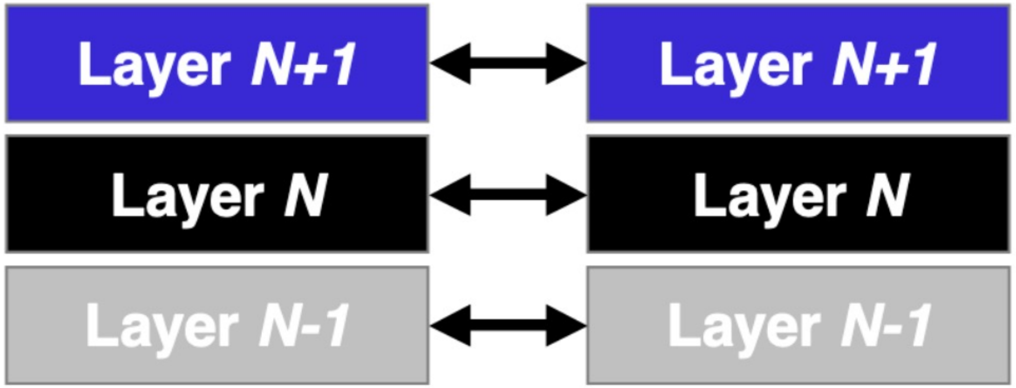

Protocol in layer N only interacts with peers in the same layer N

7

Outline

0. Administrivia and recap
1. Goals
2. Packet forwarding vs routing
3. Packet loss and delay
4. Packet switching vs circuit switching

Goals

Upon completing this lesson, you will be able to

- Explain the difference between forwarding and routing
- Explain four sources of packet delay and cause of packet loss
- Understand pros/cons of packet switching and circuit switching

Outline

1. Goals
2. Recap: One interface, two addresses

10

One network interface has two addresses

- Network layer address
  - IP address
  - IPv4 looks like 192.168.86.250 (32 bit)
  - IPv6 looks like 2001:db8:3333:4444:5555:6666:7777:8888 (128 bit)

- Link layer address
  - MAC address
  - Looks like 14:7d:da:d9:eb:fb (48 bit)

Why 2 types of addresses?

Routing to another subnet with addressing

Sending a datagram from A to B via R

- assume that:
  - A knows B's IP address (how?)
  - A knows IP address of first hop router, R (how?)
  - A knows R's MAC address (how?)

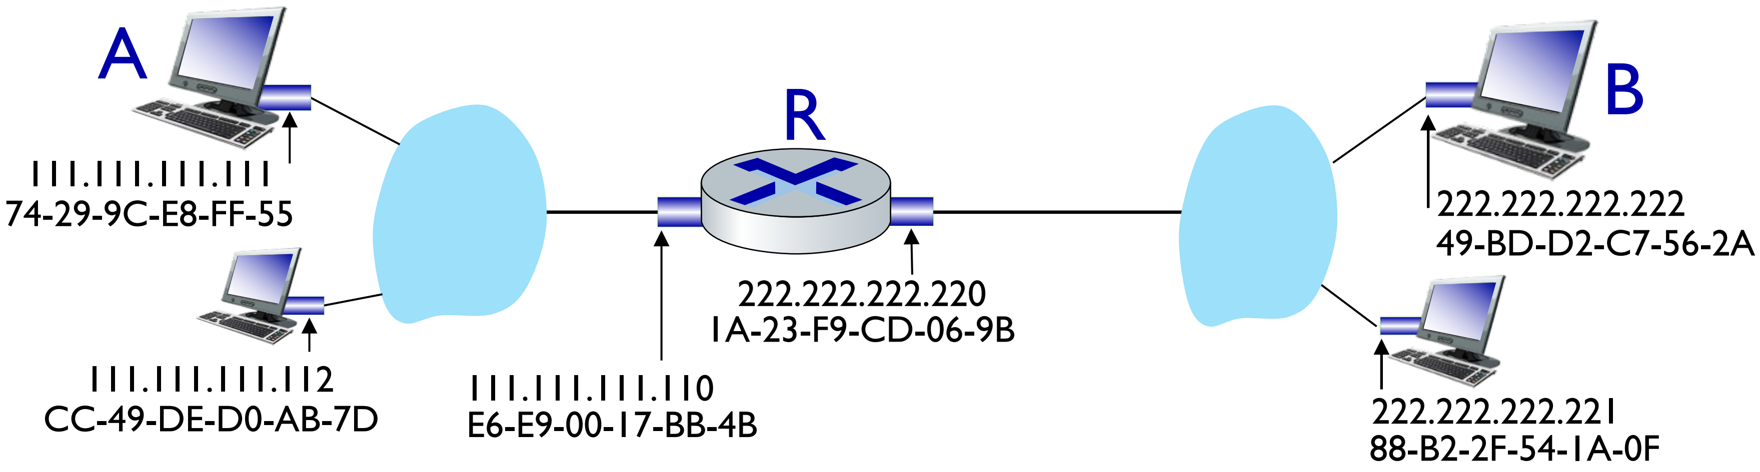

Routing to another subnet with addressing

- A creates IP datagram with IP source A, destination B
- A creates link-layer frame containing A-to-B IP datagram
  - R's MAC address is frame’s destination

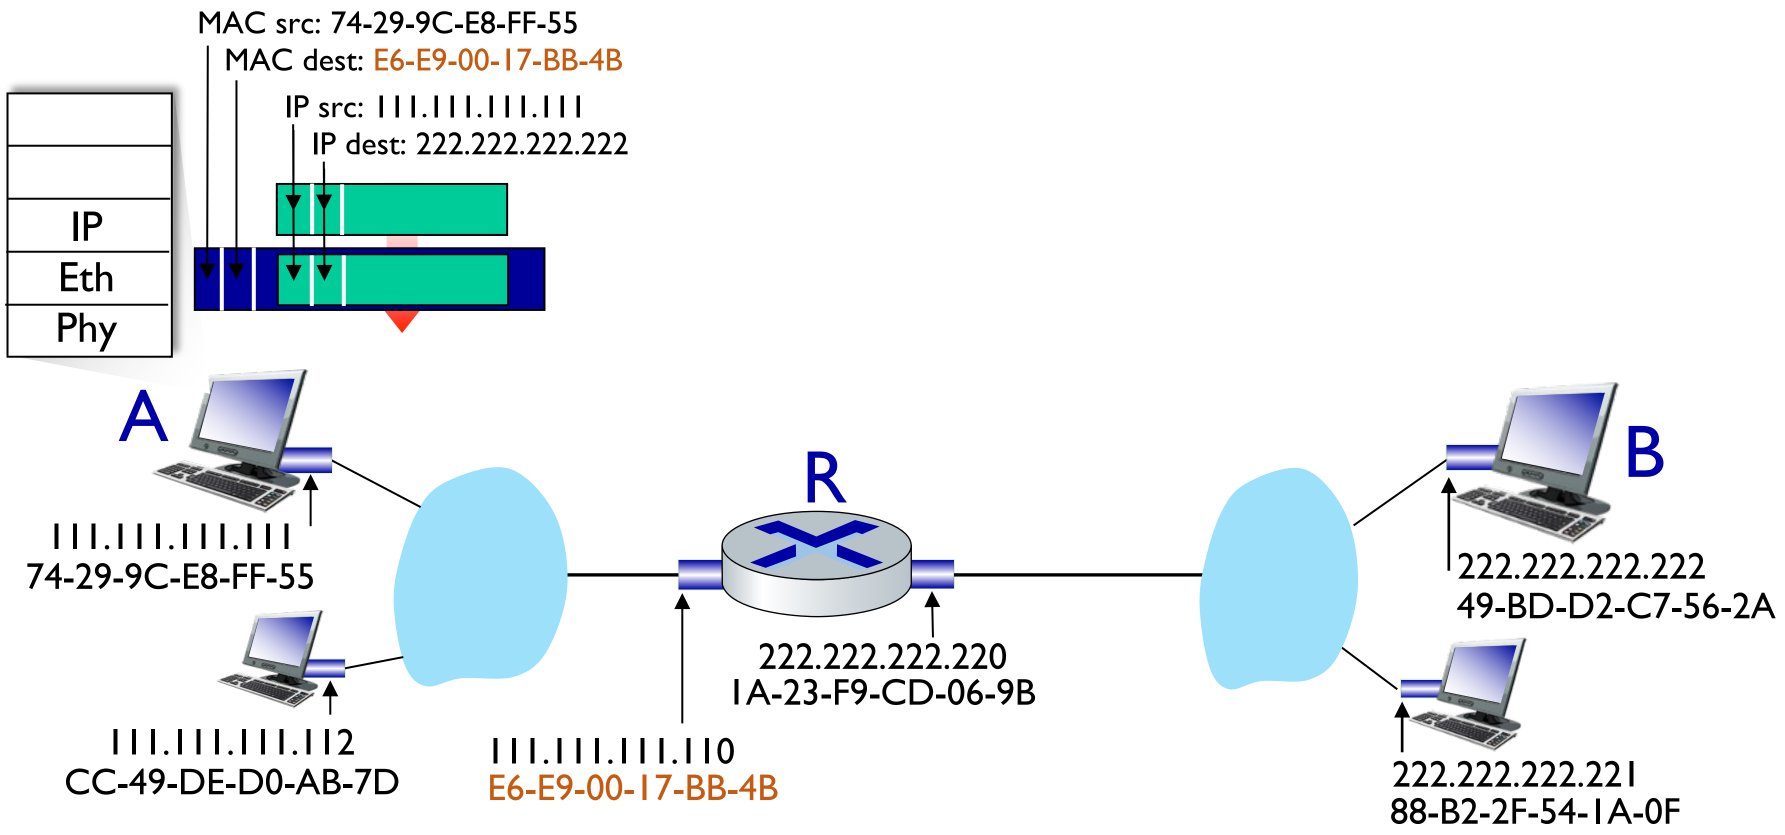

Routing to another subnet with addressing

- frame sent from A to R
- frame received at R, datagram removed, passed up to IP

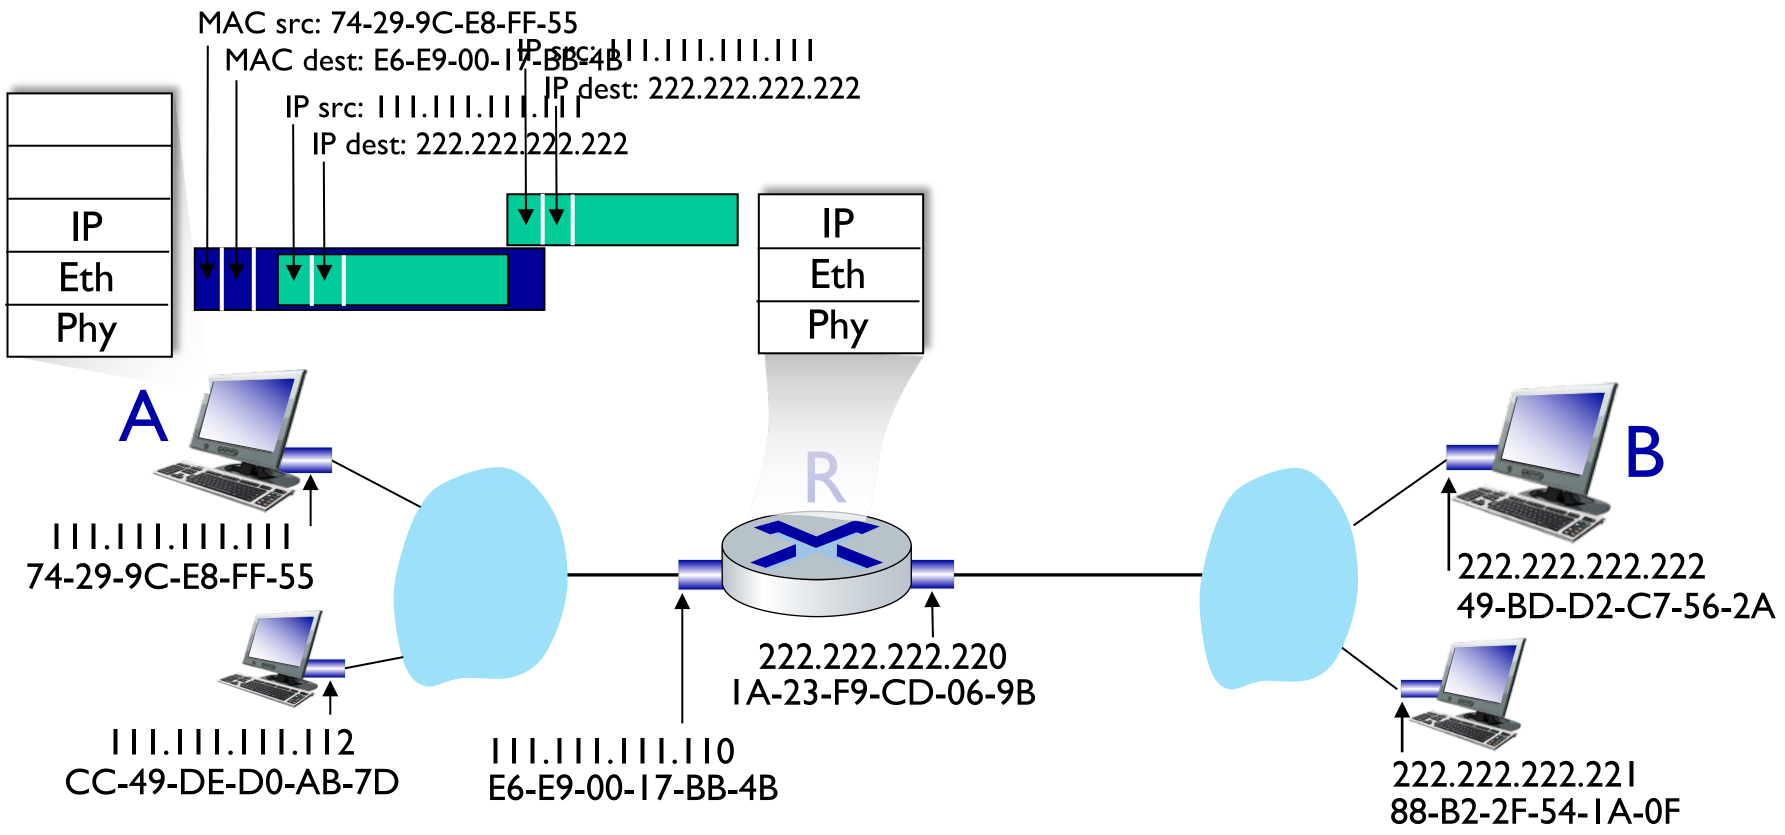

Routing to another subnet with addressing

- R determines outgoing interface, passes datagram with IP source A, destination B to link layer
- R creates link-layer frame containing A-to-B IP datagram. Frame destination address: B's MAC address

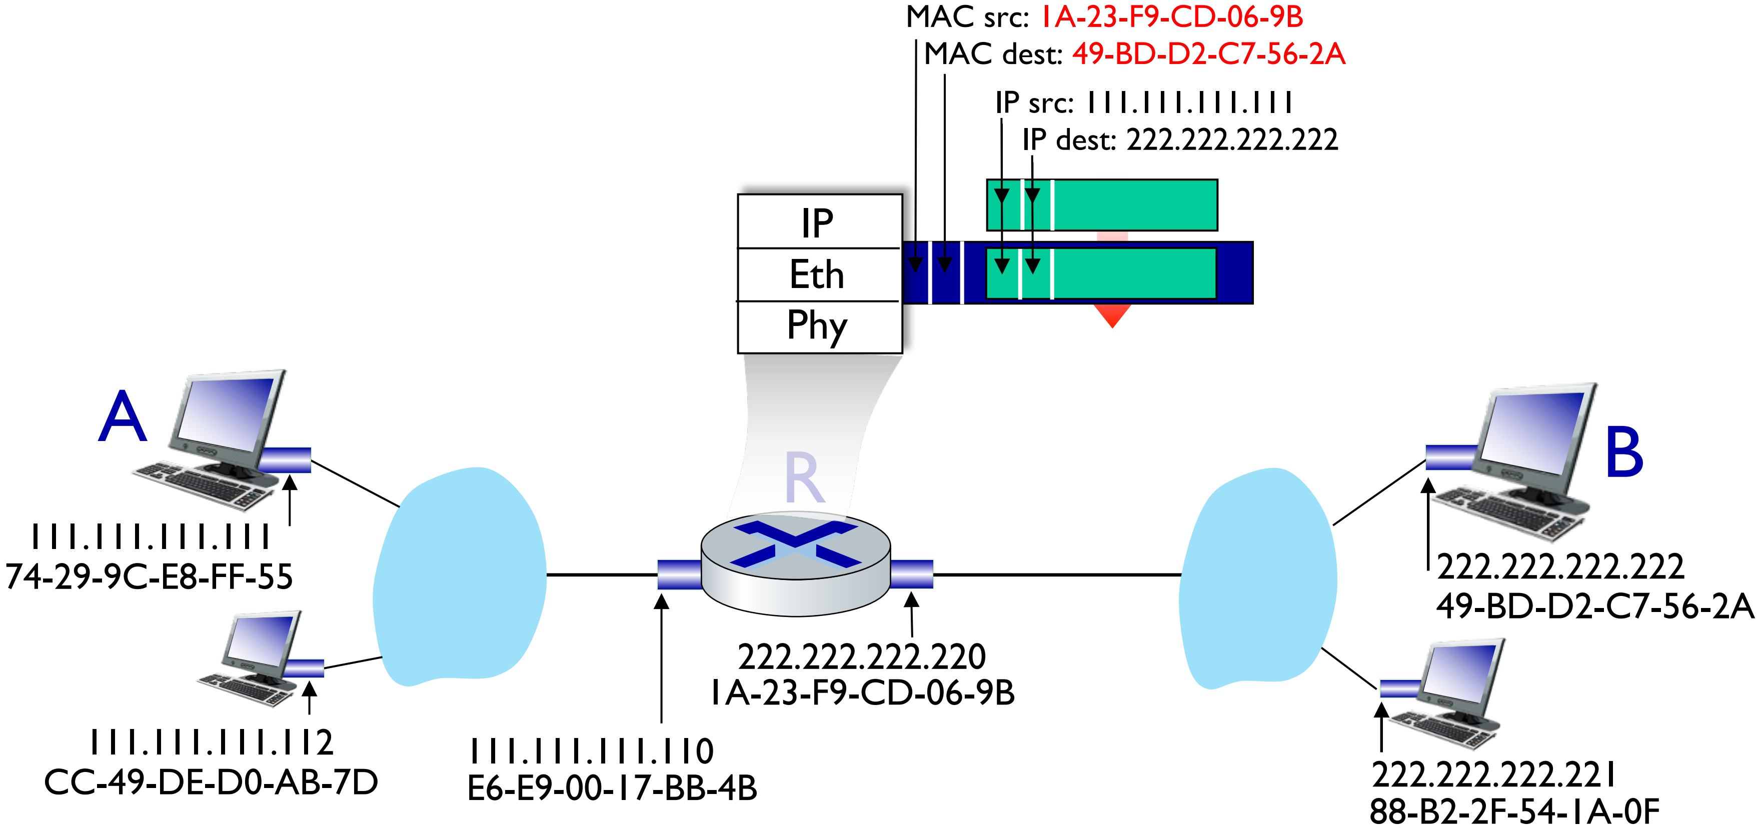

Routing to another subnet with addressing

- R determines outgoing interface, passes datagram with IP source A, destination B to link layer
- R creates link-layer frame containing A-to-B IP datagram. Frame destination address: B's MAC address
- transmits link-layer frame

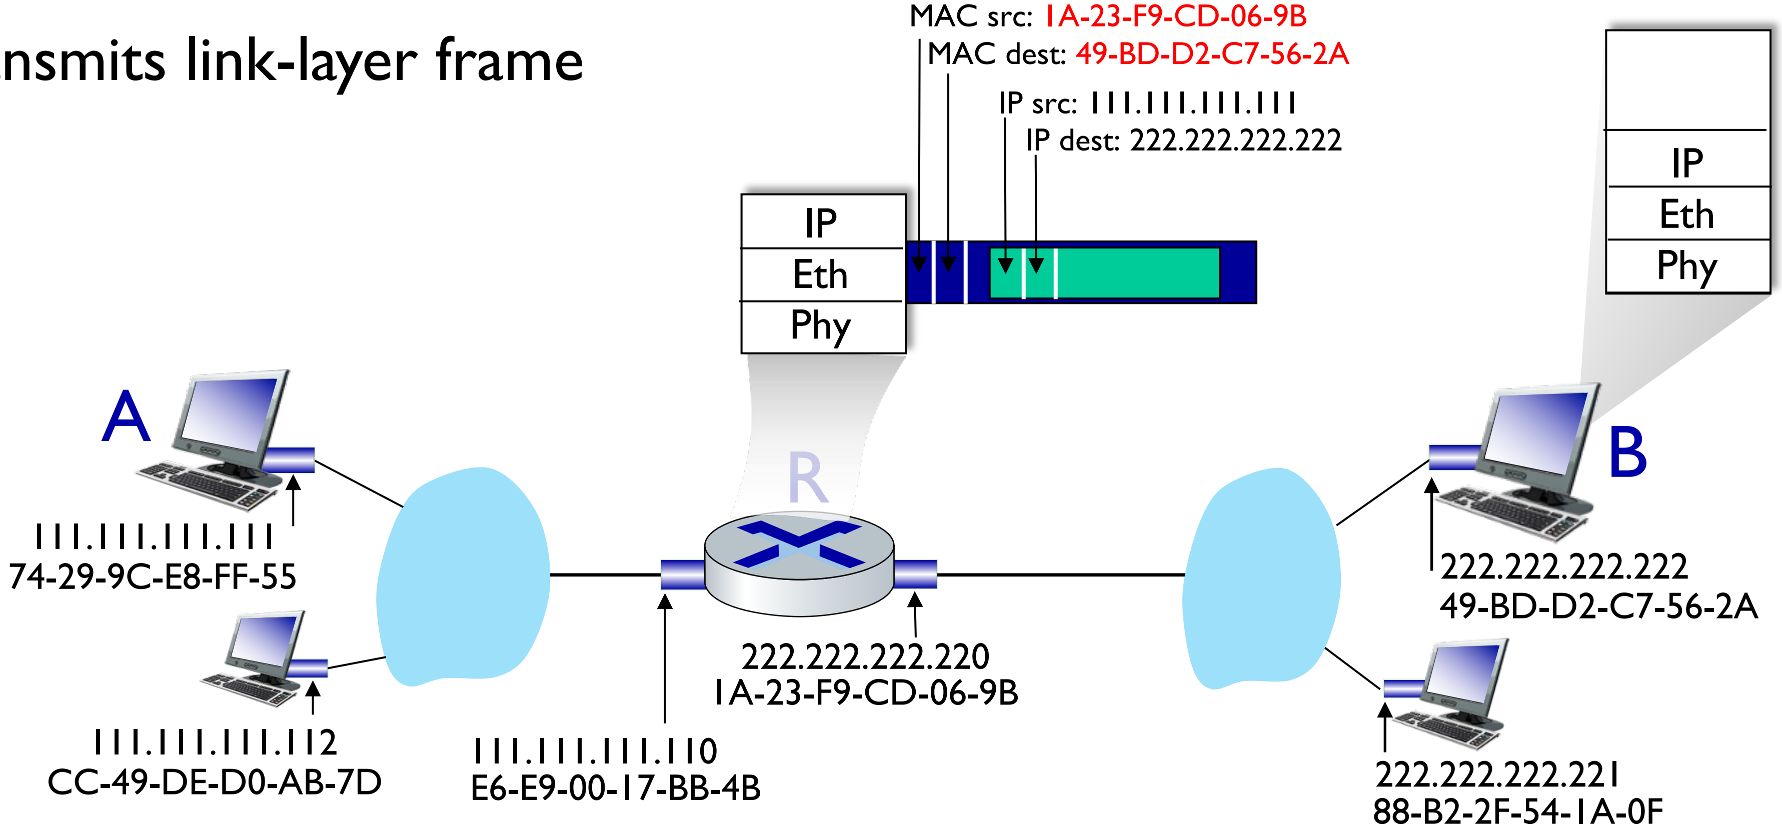

Routing to another subnet with addressing

- B receives frame, extracts IP datagram destination B
- B passes datagram up protocol stack to IP

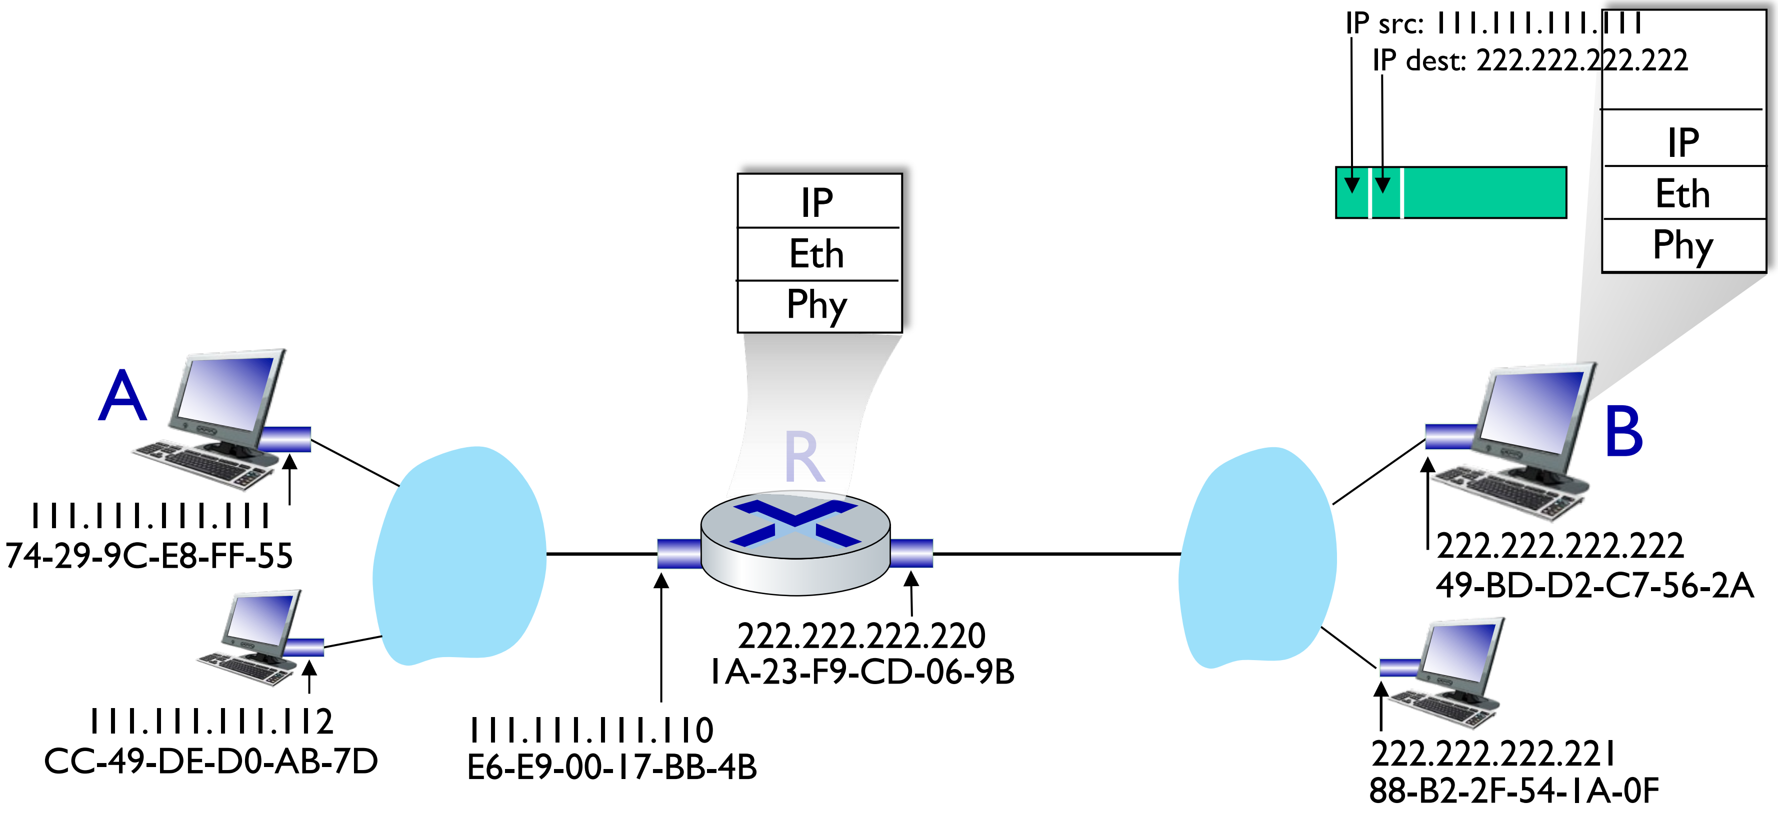

Outline

1. Goals
2. Recap
3. Packet forwarding vs routing

The network core consists of switches and routers which forwards packets from src to dst

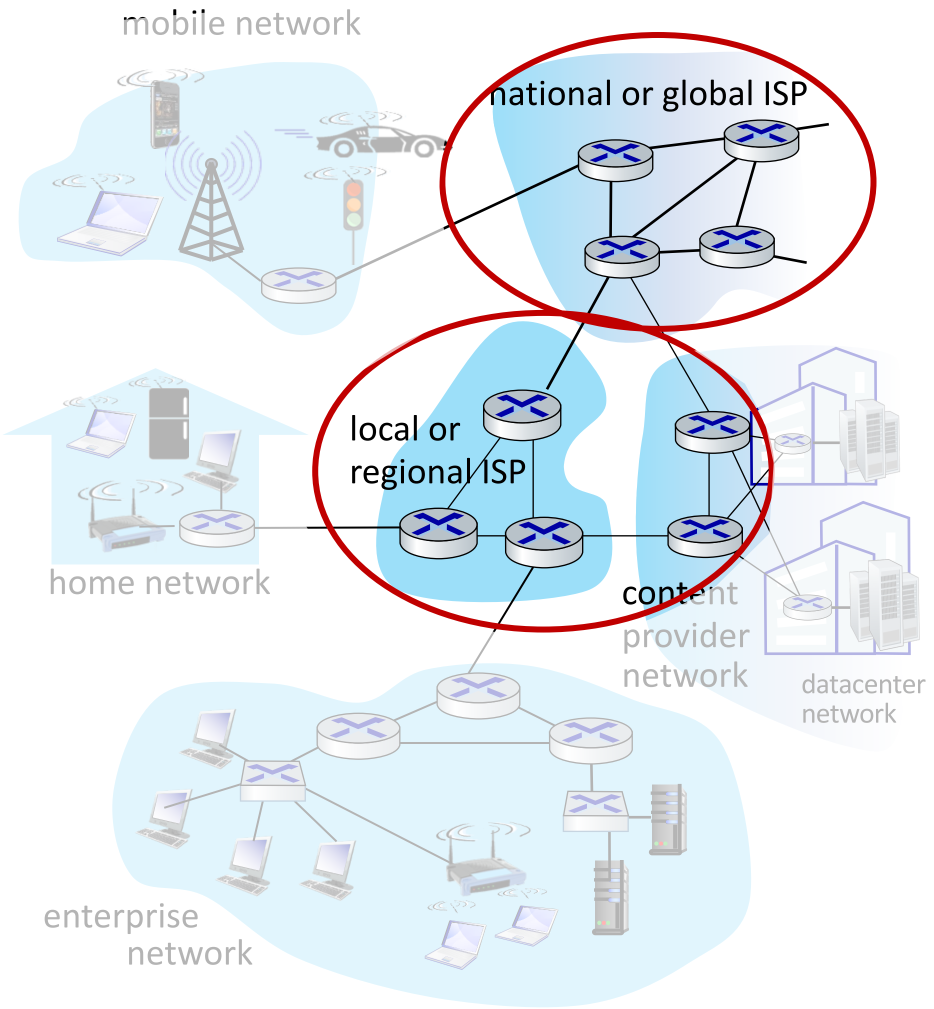

Switches vs routers: Both do packet forwarding!

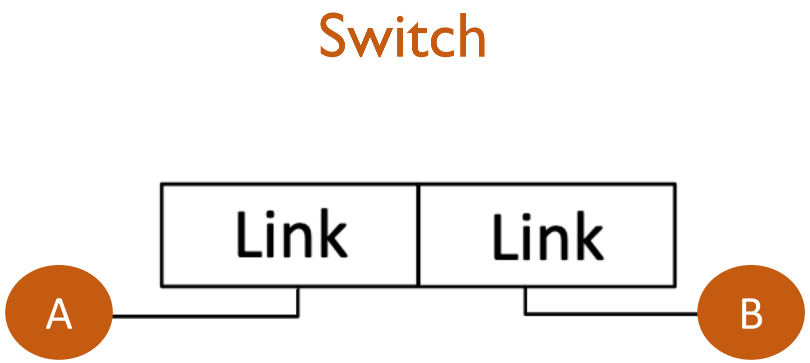

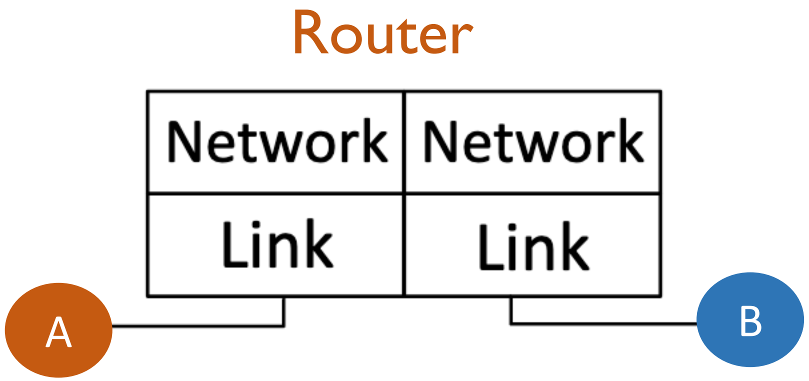

Switch forwards within the same network,
whereas routers forwards across different network

Forwarding vs routing

Forwarding:

- local action: move arriving packets to appropriate output link

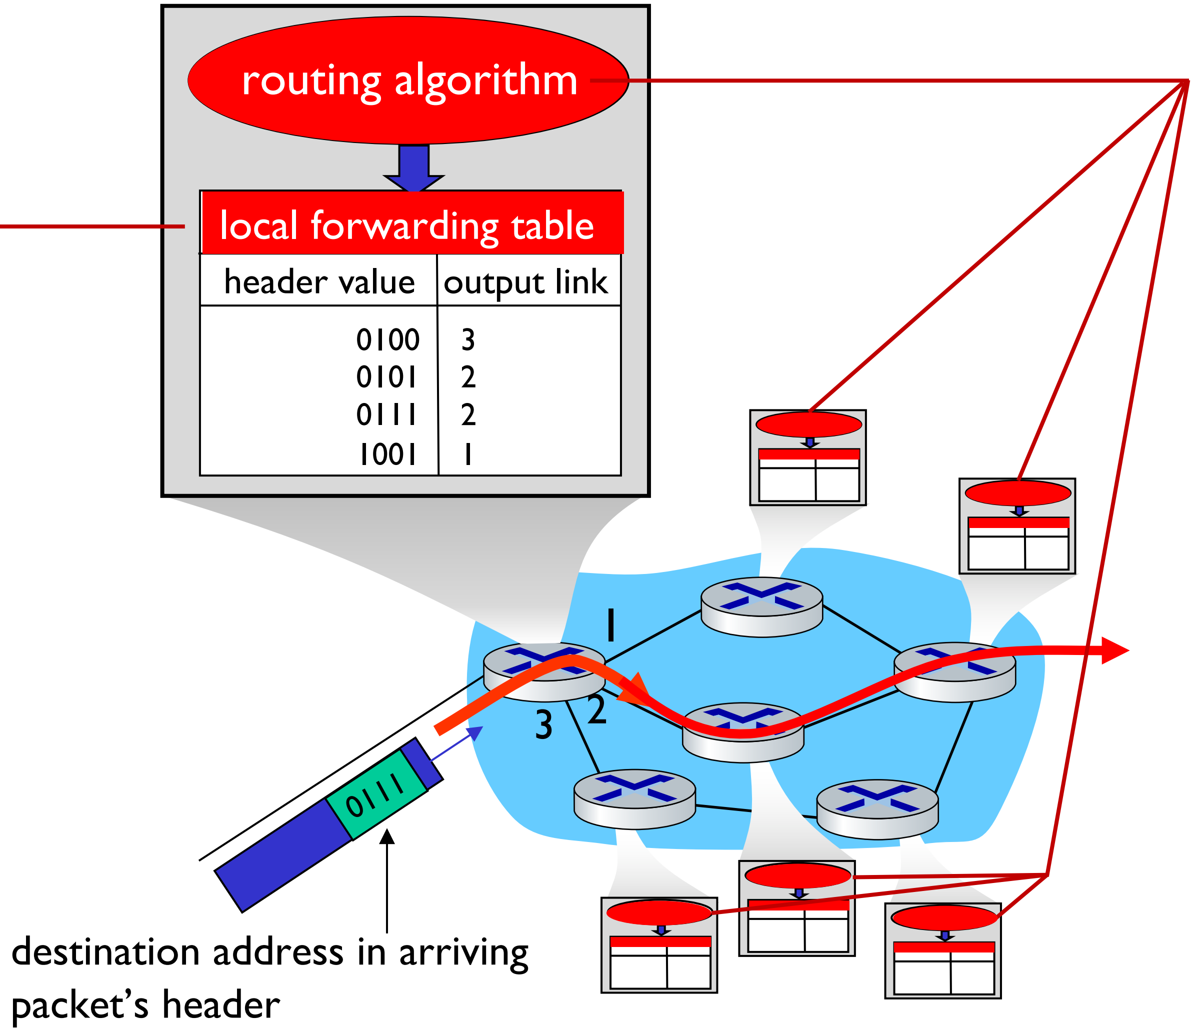

Routing:

- global action: determine src-dst paths taken by packets

routing
45 h
3,053 miles
48 h
3,309 miles
United States
San Jose
Northampton


forwarding
forwarding
45 h
3,053 miles
48 h
3,309 miles
San Jose
Northampton
United States
Sacramento
Cleveland
Seattle
Portland
San Francisco
Los Angeles
San Diego
Las Vegas
Phoenix
Tucson
El Paso
Dallas
Austin
Houston
San Antonio
New Orleans
Kansas City
Minneapolis
Chicago
Detroit
Indianapolis
Nashville
Atlanta
Toronto
Ottawa
Montreal
Quebec City
New York
Tampa

Outline

Store-and-forward: entire packet must arrive before it can be transmitted on to next link!

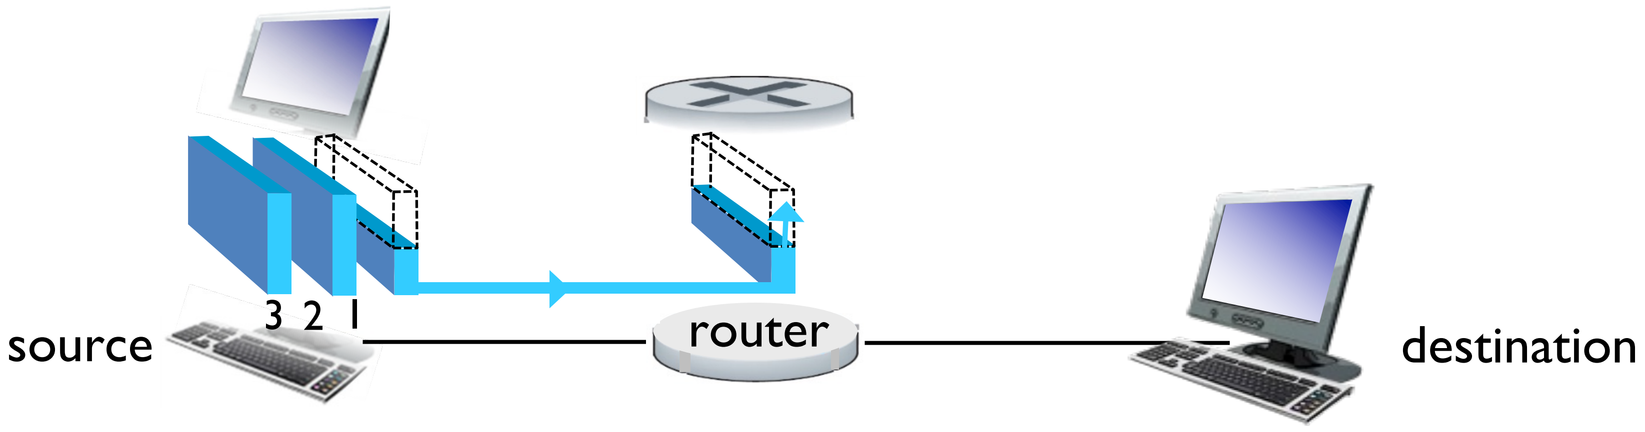

What happens if we don't?

How do packet delay occur?

Packet delay: four sources

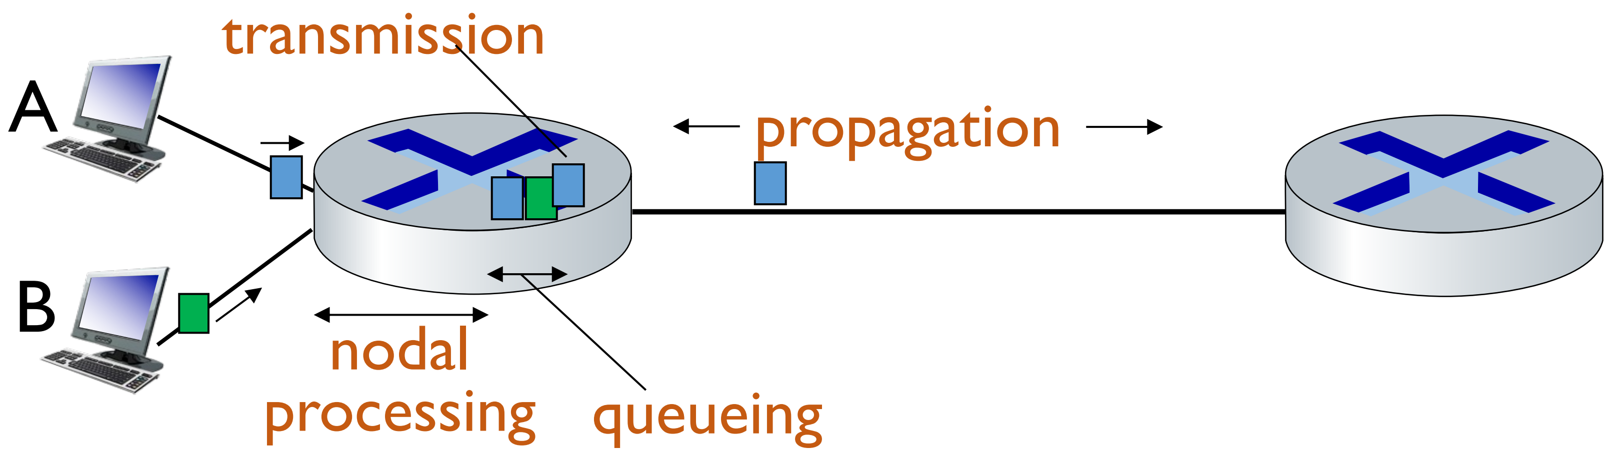

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

d_{proc}: nodal processing

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

d_{queue}: queueing delay

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

Packet delay: four sources

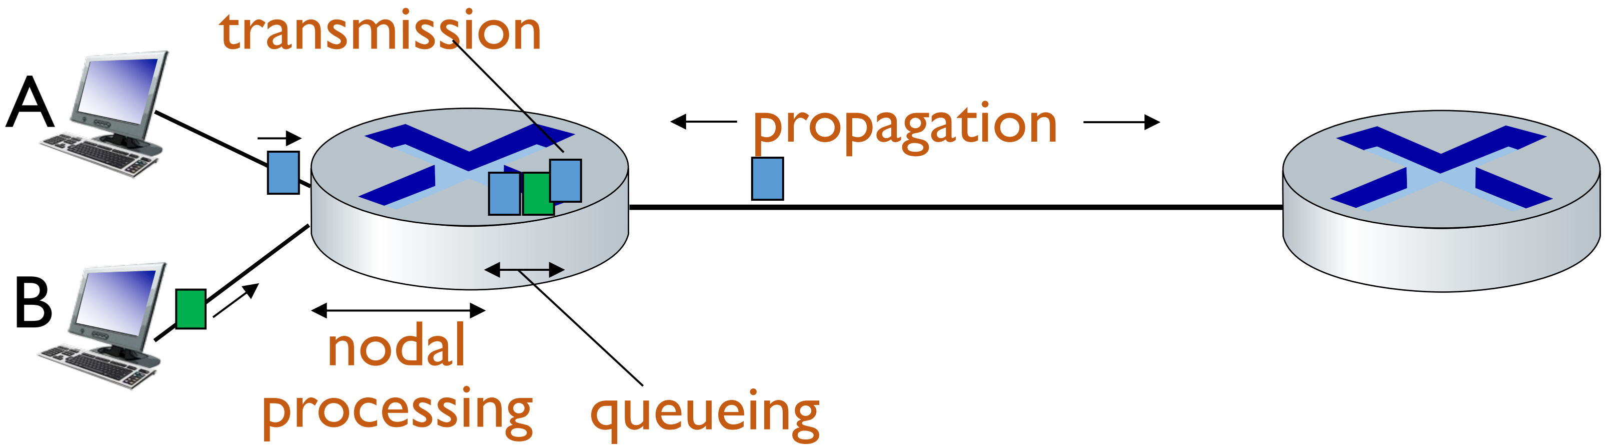

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

d_{trans}: transmission delay:

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

d_{prop}: propagation delay:

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

d_{trans} and d_{prop} very different

Transmission delay vs Propagation Delay

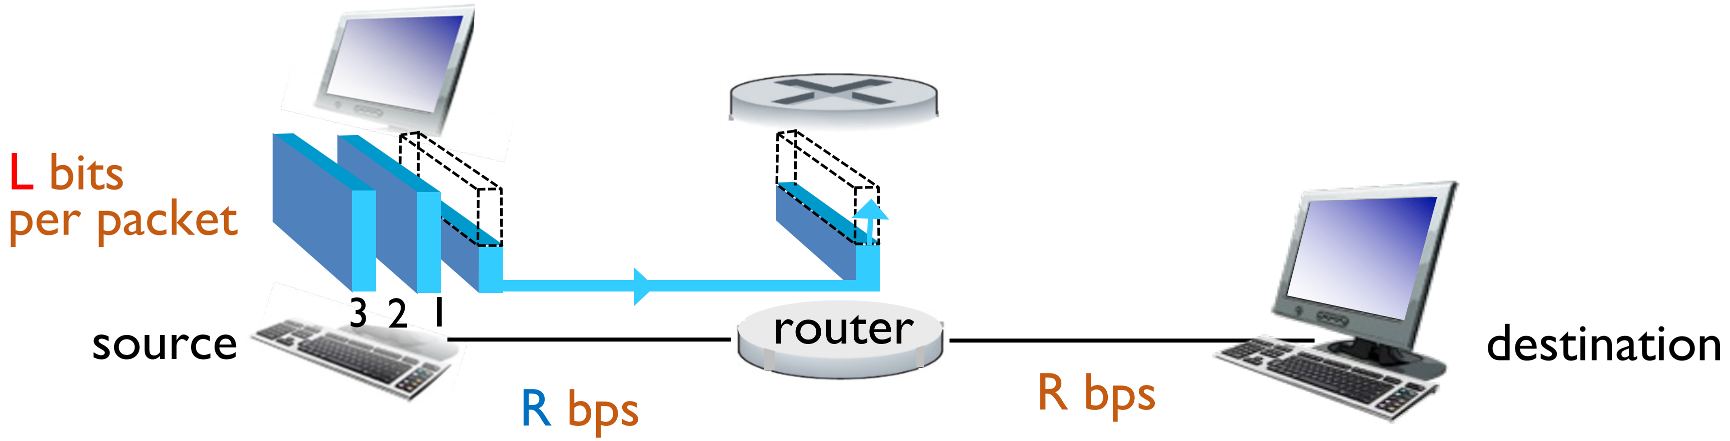

- transmission rate: how fast data is pushed onto a link (in bits per sec)
- transmission delay: time to take to push all bits in the packet to the output link
- d_{trans}: L/R sec

One-hop example:
- L = 10 Kbits
- R = 100 Mbps
- one-hop d_{trans} = 0.1 msec

Transmission delay vs Propagation Delay

- Depends on the propagation speed (s meters/sec) of the physical medium
  - Fastest is the speed of light (optical fiber)
- Depends on the distance of travel (m meters)
- Propagation delay: Time it takes for a bit in the beginning of the link to get to the next hop
- d_{prop}: m/s sec

Caravan analogy

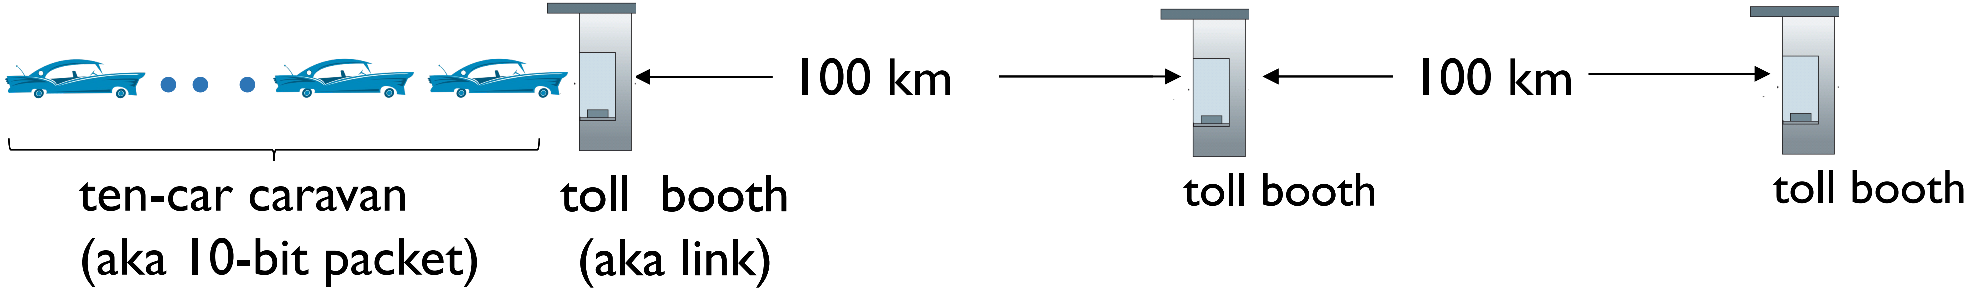

- car ~ bit; caravan ~ packet; toll service ~ link transmission
- toll booth takes 12 sec to service a car (bit transmission time)
- "propagate" at 100 km/hr
- Q: How long until caravan is lined up before 2nd toll booth?
- time to "push" entire caravan through toll booth onto highway = 12*10 = 120 sec
- time for last car to propagate from 1st to 2nd toll both: 100km/(100km/hr) = 1 hr
- A: 62 minutes

Caravan analogy

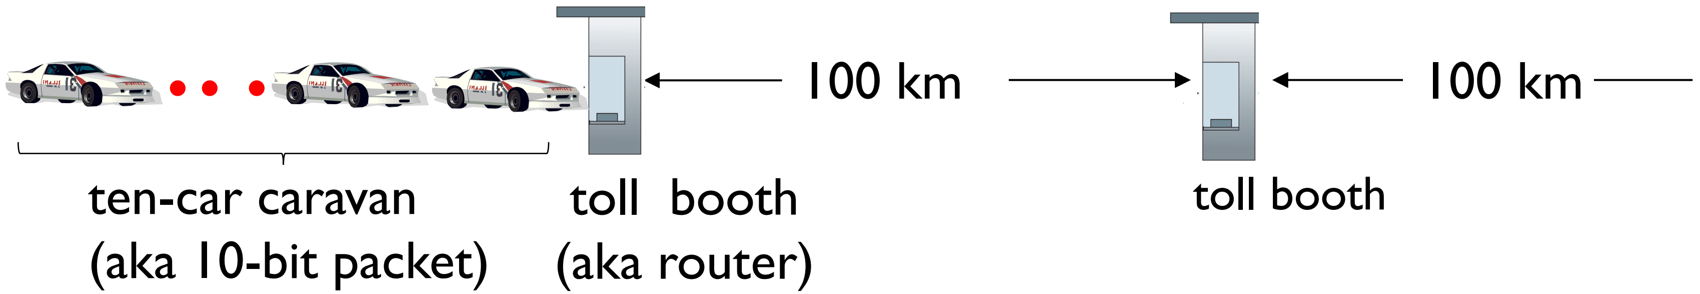

- suppose cars now "propagate" at 1000 km/hr
- and suppose toll booth now takes one min to service a car
- Q: Will cars arrive to 2nd booth before all cars serviced at first booth?

A: Yes! after 7 min, first car arrives at second booth; three cars still at first booth

In-class activity

Take a screenshot of your Kahoot score

- Submit to Canvas!

Queueing delays

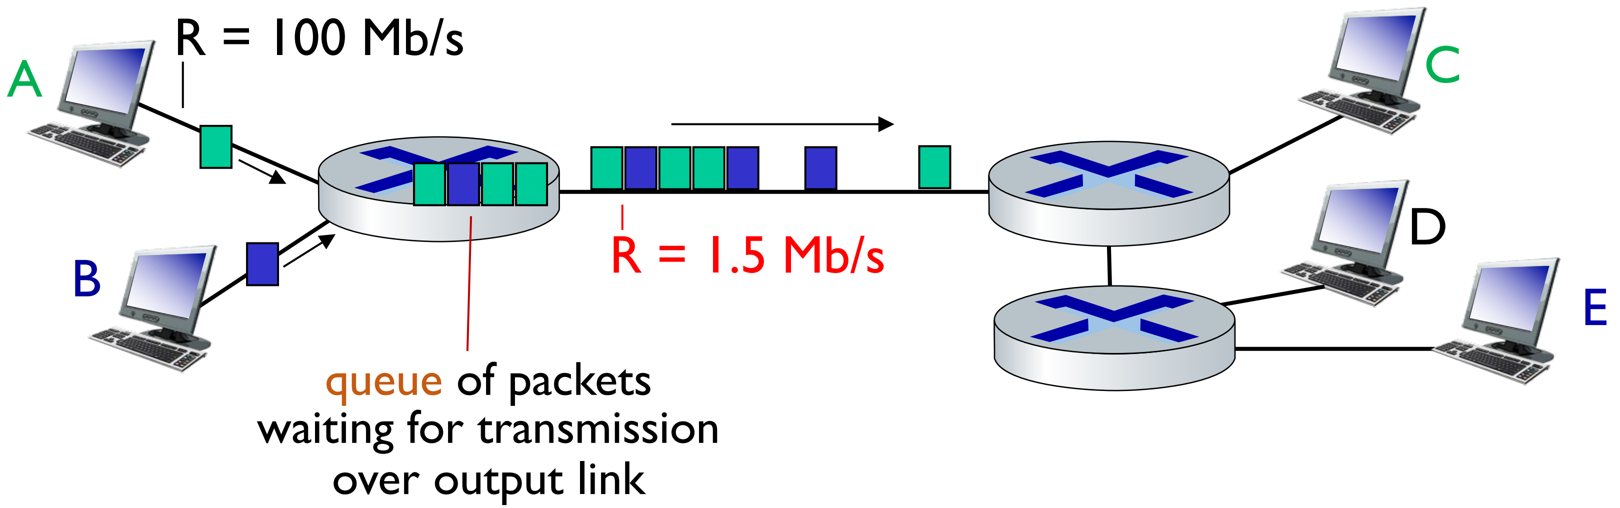

Queueing occurs when work arrives faster than it can be serviced:

Queueing delay analysis

- a: average packet arrival rate
- L: packet length (bits)
- R: 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!


La/R ~ 0


La/R => 1

How about packet loss? Why/where does it happen?

Packet loss happens

- queue (buffer) has finite capacity
- packets arriving to full queue dropped (lost)
- Lost packets may be retransmitted by src, previous hop, or not at all!

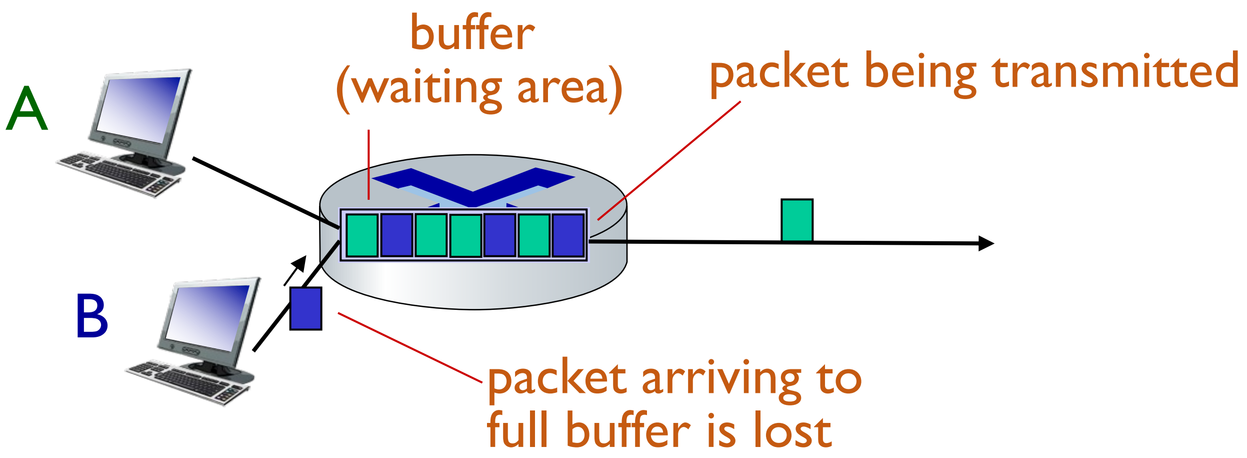

Where else can packet loss happen?

Outline

1. Goals
2. Recap
3. Packet forwarding vs routing
4. Performance: loss, delay, and throughout
5. Sharing is caring: Packet switching vs circuit switching

How to share a link between multiple users?

Circuit switching is an alternative approach

end-end resources are allocated to, reserved for "call" btw src and dst

- in diagram, each link has four circuits.
  - call gets 2nd circuit in top link and 1st circuit in right link.
- dedicated resources: no sharing
  - circuit-like (guaranteed) performance
- circuit segment idle if not used
- commonly used in traditional telephone networks

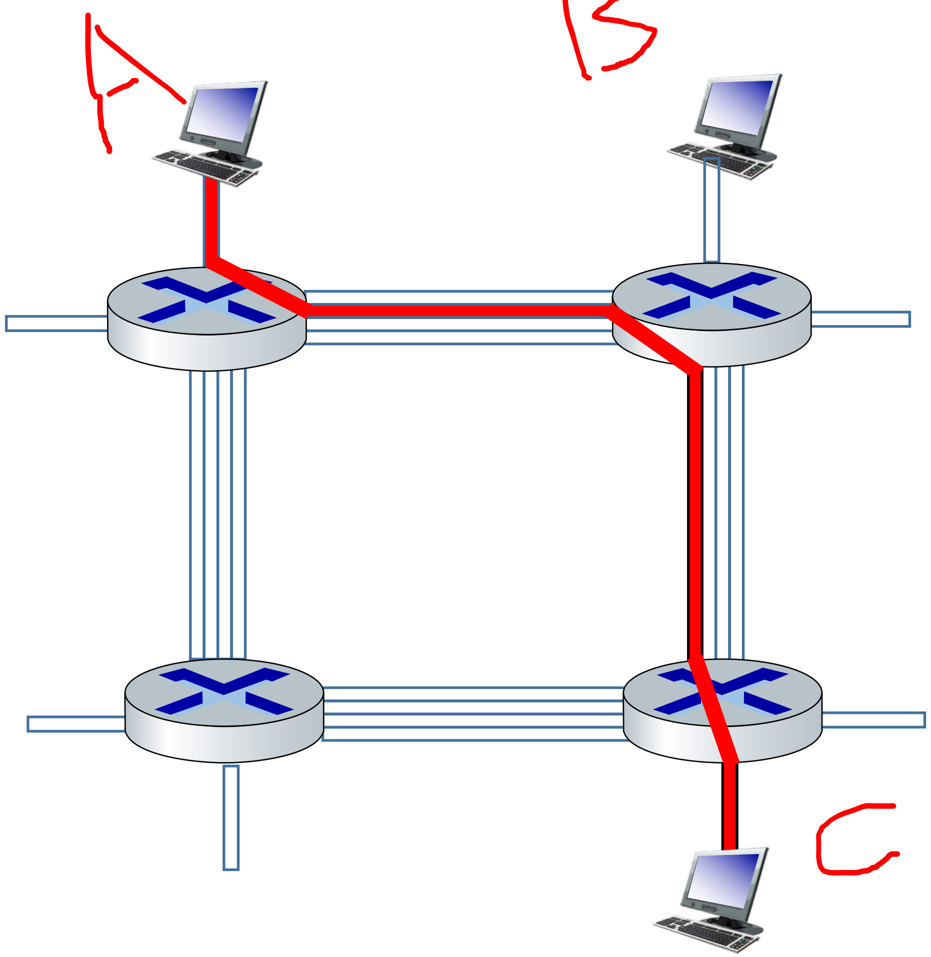

Circuit switching: FDM and TDM

Frequency Division Multiplexing (FDM)

- optical, electromagnetic frequencies divided into (narrow) frequency bands
- each call allocated its own band, can transmit at max rate of that narrow band

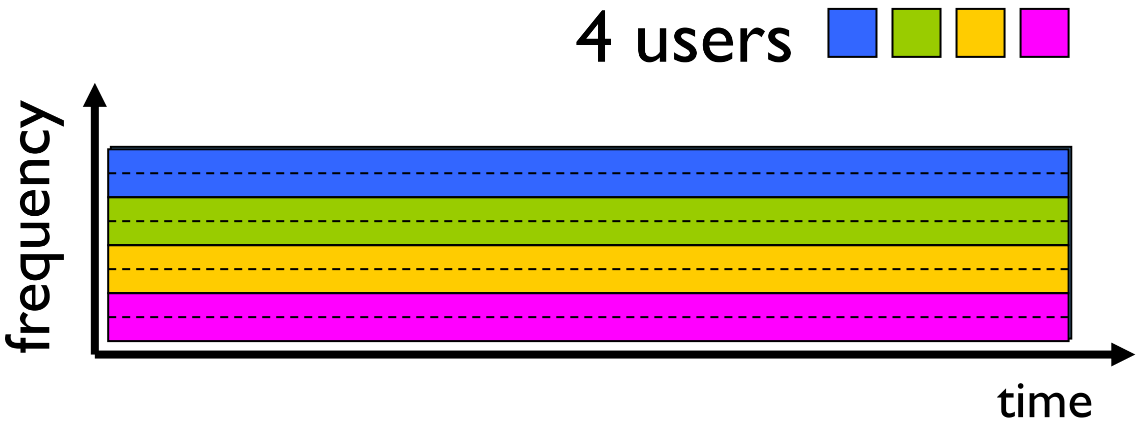

Time Division Multiplexing (TDM)

- time divided into slots
- each call allocated periodic slot(s), can transmit at maximum rate of (wider) frequency band (only) during its time slot(s)

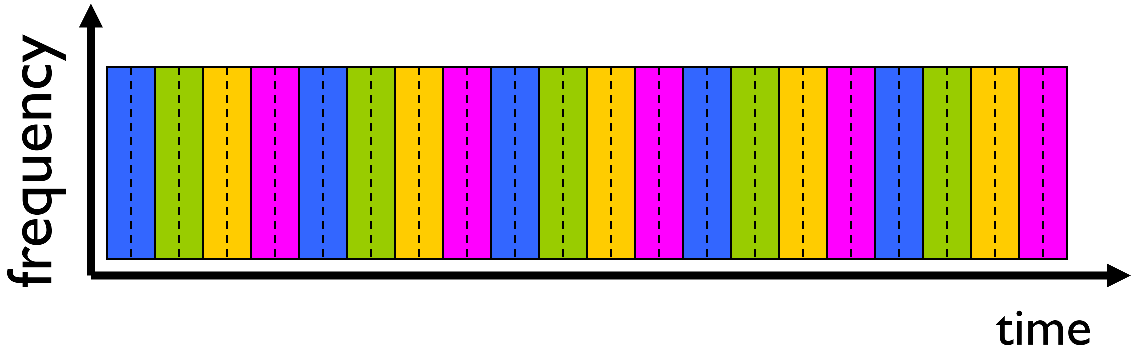

Packet switching versus circuit switching

example:

- 1 Gb/s link
- each user:
  - 100 Mb/s when "active"
  - active 10% of time (happens randomly)

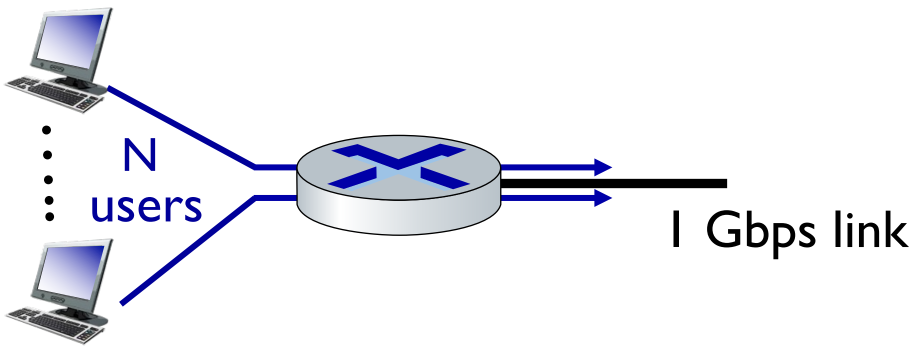

Q: What is the max number of users that can share this network?

- circuit-switching: 10 users
  - $10^9/(100\times10^6)=1000/100 = 10$

Packet switching versus circuit switching

example:

- 1 Gb/s link
- each user:
  - 100 Mb/s when "active"
  - active 10% of time (happens randomly)

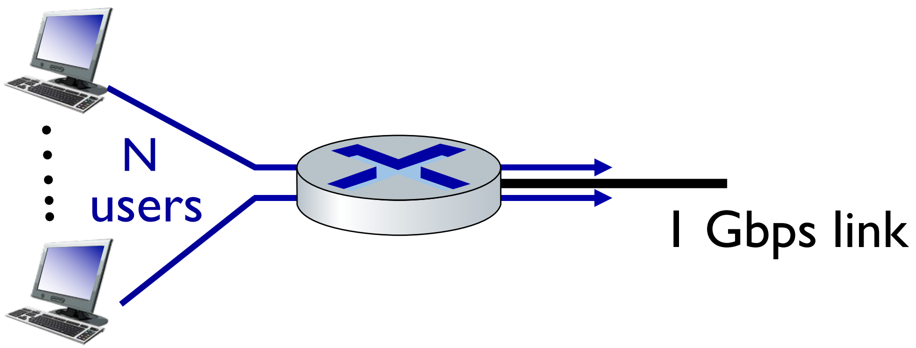

Q: What is the max number of users that can share this network?

- packet switching: ???

A: Need some assumption on link availability guarantee.

Say we guarantee 99.99% link availability for each user.

Failure rate < 0.0001 (When does failure happen?)

What is the number of users such that the probability of more than 10 users being active simultaneously < 0.0001?

N total number of users in the system k active users

- N = 4 (A, B, C, D), k = 2,
- What are the possible scenarios/cases that exactly 2 people out of 4 speaking?
  - AB, AC, AD, BC, BD, CD
  - Combination = 4 Choose 2 = 4 x 3 / 2 x 1
    →there are duplicates (AB/BA, AC/CA, need to divide by 2)
- N = 5 (A, B, C, D, E), k =3
- Possible cases of picking 3 out of 5
  - (5 x 4 x 3) / (3 x 2 x 1)
  - includes duplicates, ABC/BAC/BCA….. For each case, 3 x 2 x1
- assume if more than 2 people out of 5 speak there is error
  - P(3/5) + P(4/5) + P(5/5)

N total number of users in the system
k active users

- assume if more than 2 people out of 5 speak there is error
  - P(3/5) + P(4/5) + P(5/5)
  - P(k/N) : the probability that exactly k out of N users are active, when each user can be active 10%
- P(3/5) = 0.1^3 x $0.9^{(5-3)}$x (5 Choose 3)
- P(4/5) = 0.1^4 x $0.9^{(5-4)}$x (5 Choose 4)
- P(5/5) = 0.1^5 x 0.90 x (5 Choose 5)
- Going back to 1Gbps with 100 Mbps with 10 max users with no error
- Total Error =$\sum_{k=11}^{N} 0.1^k\ x\ 0.9^{(N-k)} x(N\ choose\ k)\ \ < 0.0001$

Hint: What's the probability to flip exactly 2 heads in 5 coin tosses?

- One possible outcome: HHTTT
- What's the probability to get above?

  $P(H) \times P(H) \times P(T) \times P(T) \times P(T) = P(H)^2 \times P(T)^3 = 0.5^2 \times (1\text{-}0.5)^3$
- Another possible outcome: TTTHH

  $P(T) \times P(T) \times P(T) \times P(H) \times P(H) = P(H)^2 \times P(T)^3 = 0.5^2 \times (1\text{-}0.5)^3$
- All outcomes have the same prob of $P(H)^2$ x $P(T)^3$
- How many possible outcome?

  $${}_5C_2 = \frac{5!}{(5-2)!2!} = 10$$
- Putting it together

  ${}_5C_2 \times P(T)^2 \times P(H)^3 = 10 \times 0.5^2 \times (1\text{-}0.5)^3 = 0.3125$

This is an example of binomial distribution

Error rate calculation Example

- N=11
  - $P(11/11) = 1 \times 0.1^{11} \times 0.9^{0}$
  - Is it smaller than 0.0001? Then add 1 more user in the system
- N=12
  - $P(11/12) + P(12/12) = {}_{12}C_{11} \times 0.1^{11} \times 0.9^{1} + {}_{12}C_{12} \times 0.1^{12} \times 0.9^{0}$
- N=13
  - $P(11/13) + P(12/13) + P(13/13) = {}_{13}C_{11} \times 0.1^{11} \times 0.9^{2} + {}_{13}C_{12} \times 0.1^{12} \times 0.9^{1} + {}_{13}C_{13} \times 0.1^{13} \times 0.9^{0}$
- …

What is max N such that it will still have the error rate guarantee of 0.0001?

Group activity

- Make sure to submit to Canvas as well!

Is packet switching a "slam dunk winner"?

- great for "bursty" data – sometimes has data to send, but at other times not
  - resource sharing
  - simpler, no call setup
- excessive congestion possible: packet delay and loss due to buffer overflow
  - protocols needed for reliable data transfer, congestion control

Best of both worlds:
How to provide "circuit-like" behavior with packet-switching?

Questions?

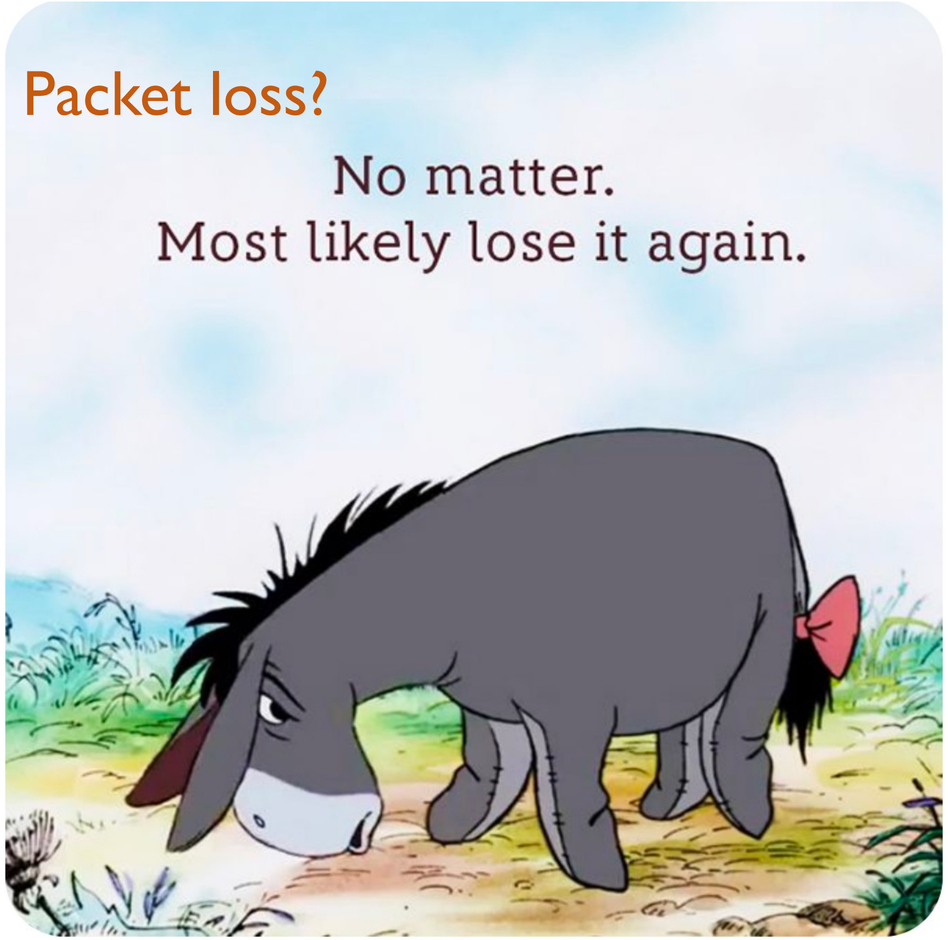

Backup Slides

Why La = 1 results in infinitely long queue?

- See [this example]
- Read [this paper]

We now compute the mean number in queue from (4). The most convenient way to do this is using generating functions. We have

$$G_N(z) = E[z^N] = \sum_{n=0}^{\infty}(1-\rho)\rho^n z^n = \frac{1-\rho}{1-\rho z},$$

provided $|z| < 1/\rho$. From this, we obtain

$$E[N] = G'_N(1) = \frac{\rho(1-\rho)}{(1-\rho z)^2} \mid |_{z=1} = \frac{\rho}{1-\rho}. \qquad (5)$$

Observe that the mean queue length increases to infinity as ρ increases to 1,

Throughput is the rate at which bits are being sent from sender to receiver

- **instantaneous:** rate at given point in time
- **average:** rate over longer period of time

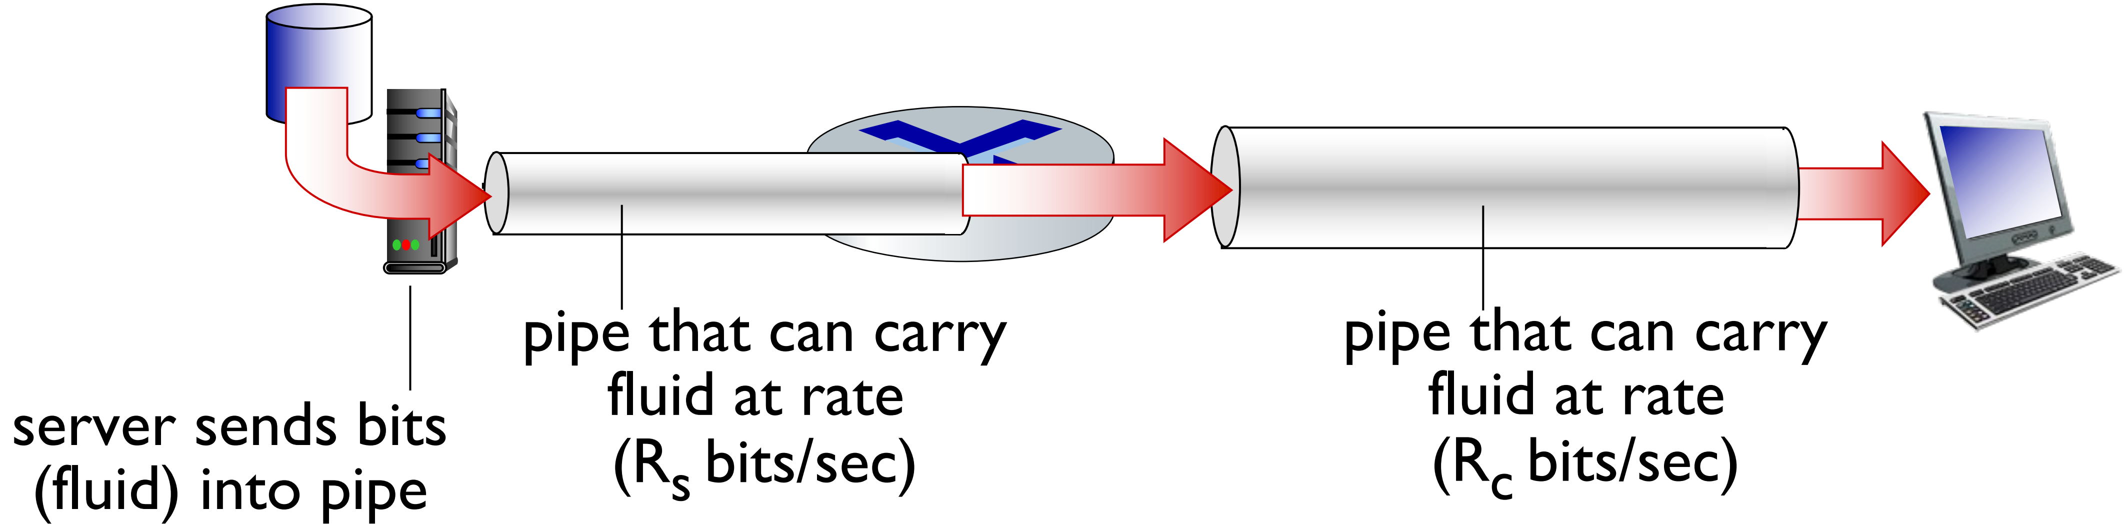

Throughput

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

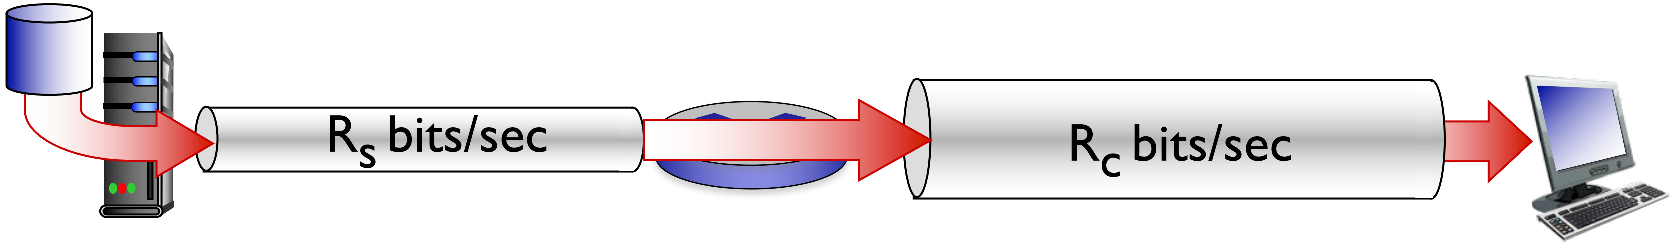

$R_s > R_c$ What is average end-end throughput?

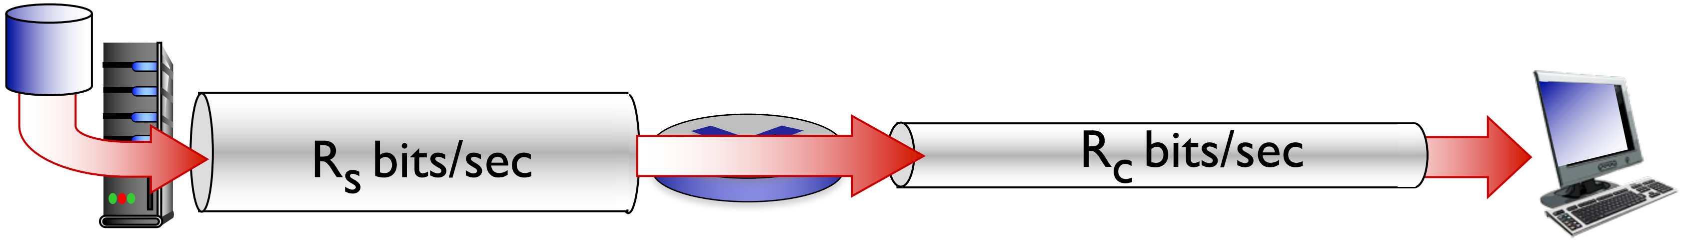

bottleneck link

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

Throughput: network scenario

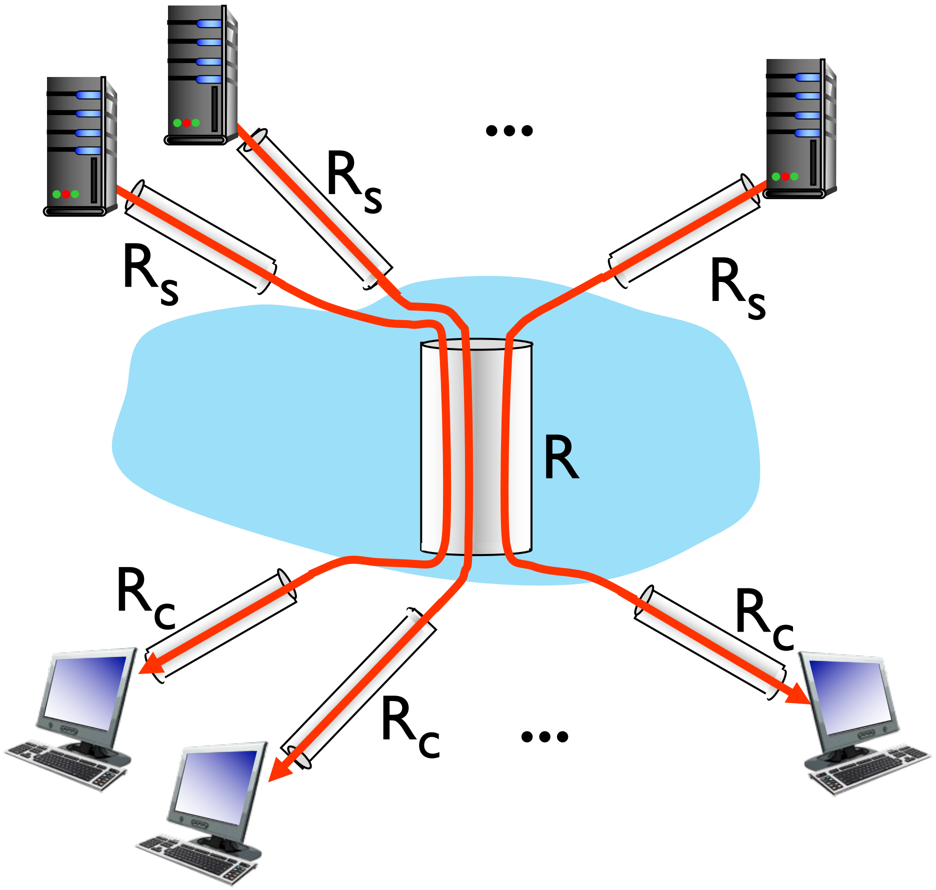

Say10 connections fairly share backbone bottleneck link R bits/sec

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

Acknowledgements

Slides are adopted from Kurose' Computer Networking