Lecture 02:
Packets, Routing, and Performance

CS 326E Computer Networks

Mikyung Han

mhan@cs.utexas.edu

Please, interrupt and ask questions AT ANY TIME !

Reminders

- Hands-on 0 Env Setup: Deadline extended Due tonight!
- Office hours: Tues/Thurs/Fri 4-5 PM & right after our class
- Lab 1 Wireshark Intro assigned

Recap questions

Why do we NEED layers?

What are the benefits of having layers?

What are the downside of having layers?

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

Sender pushes a packet top-down

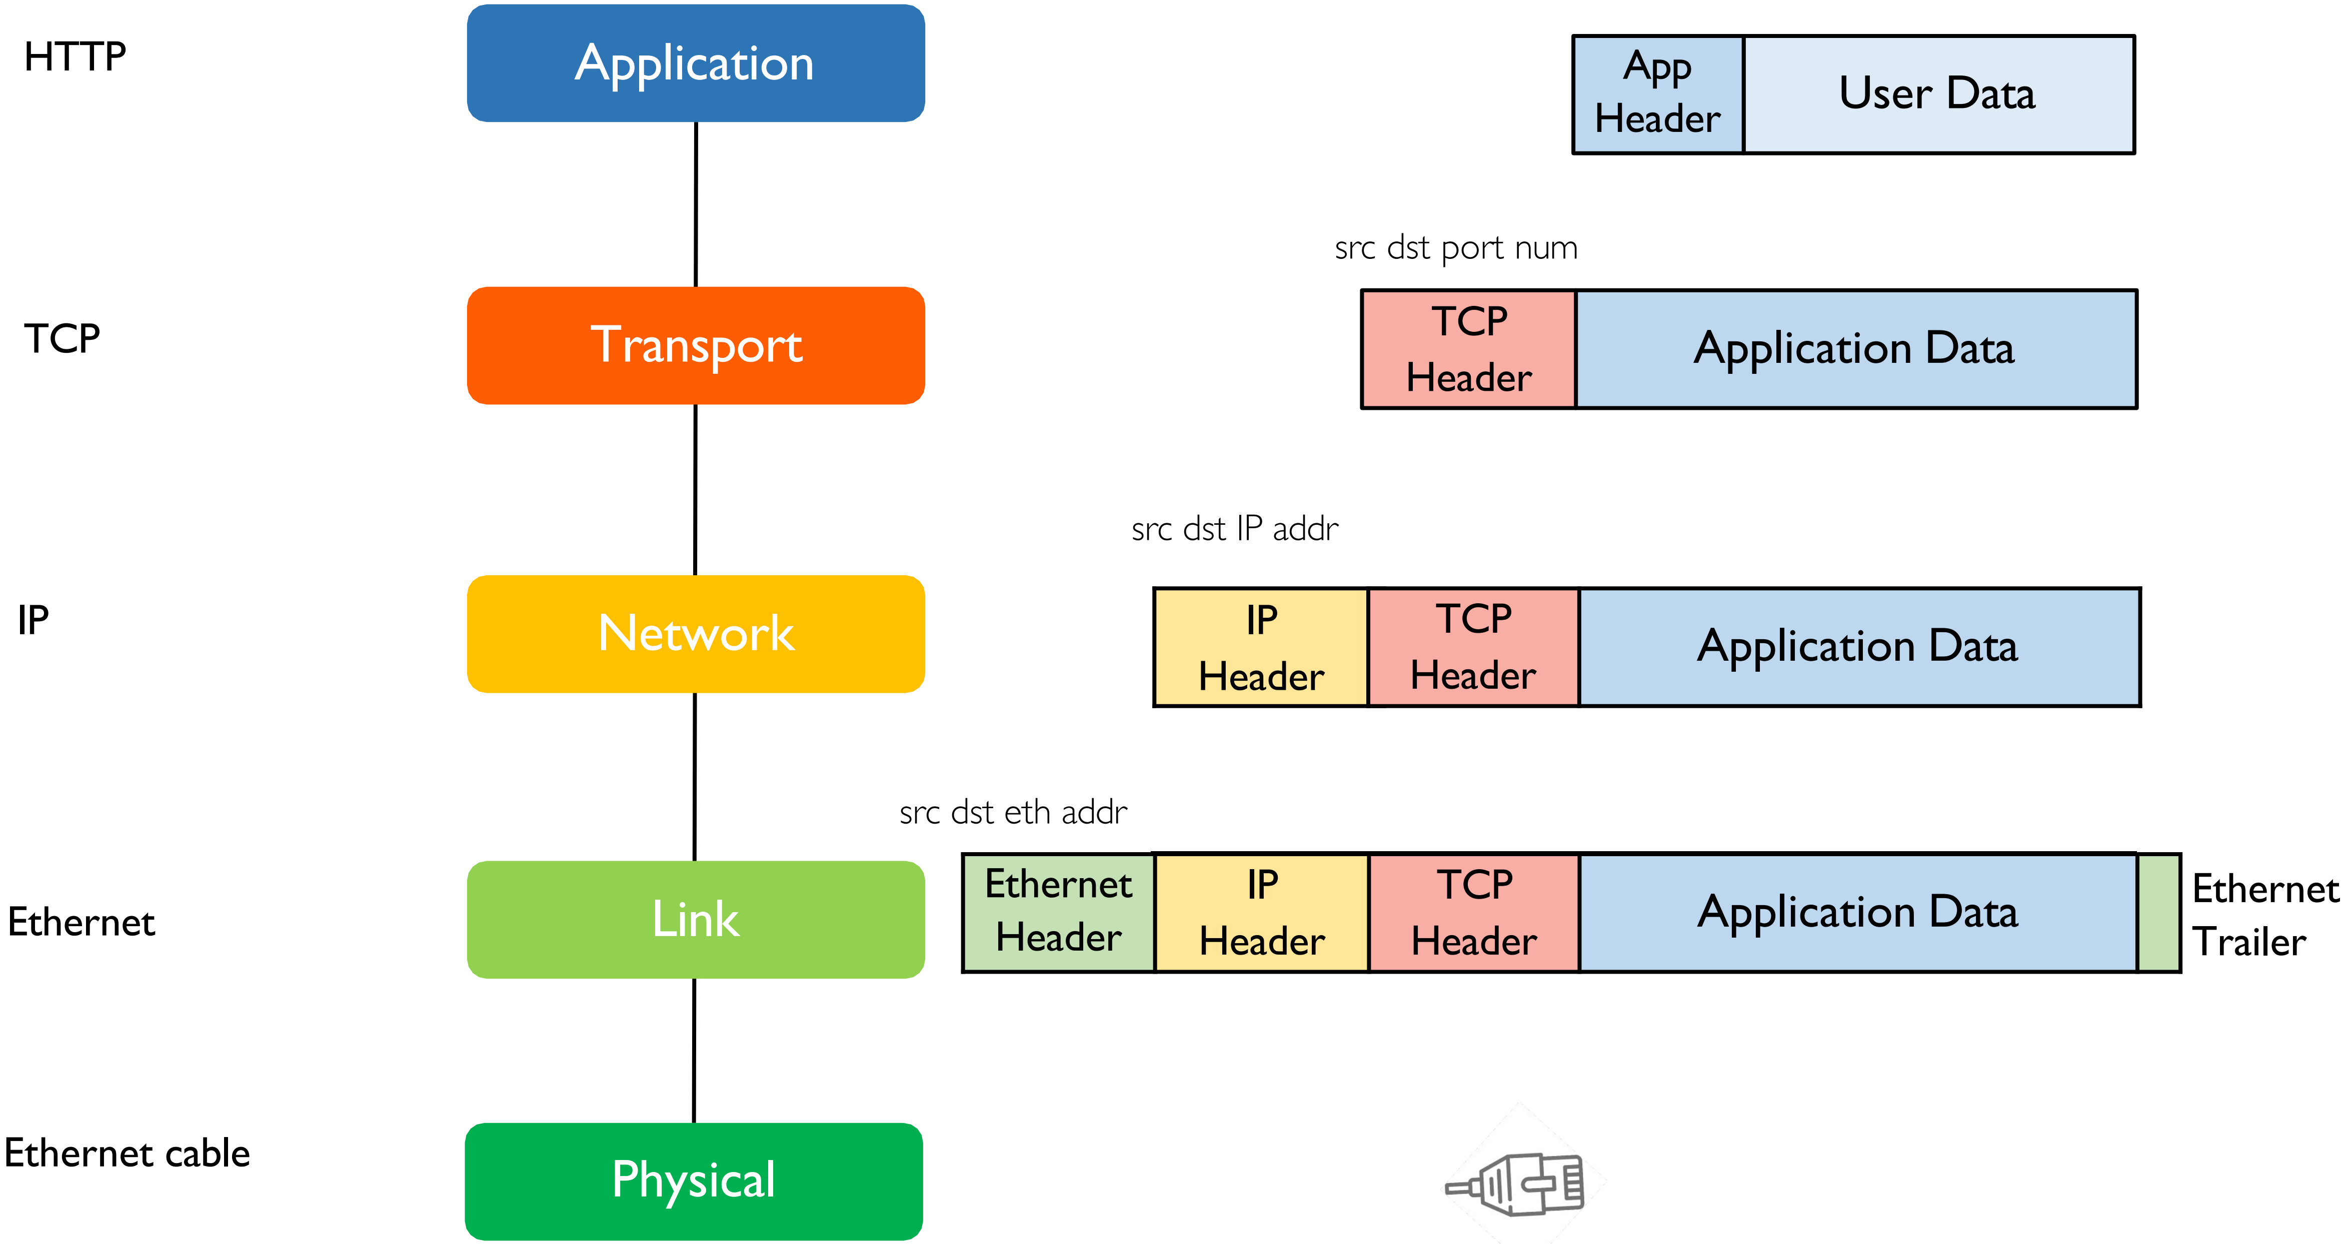

Receiver pushes a packet bottom-up

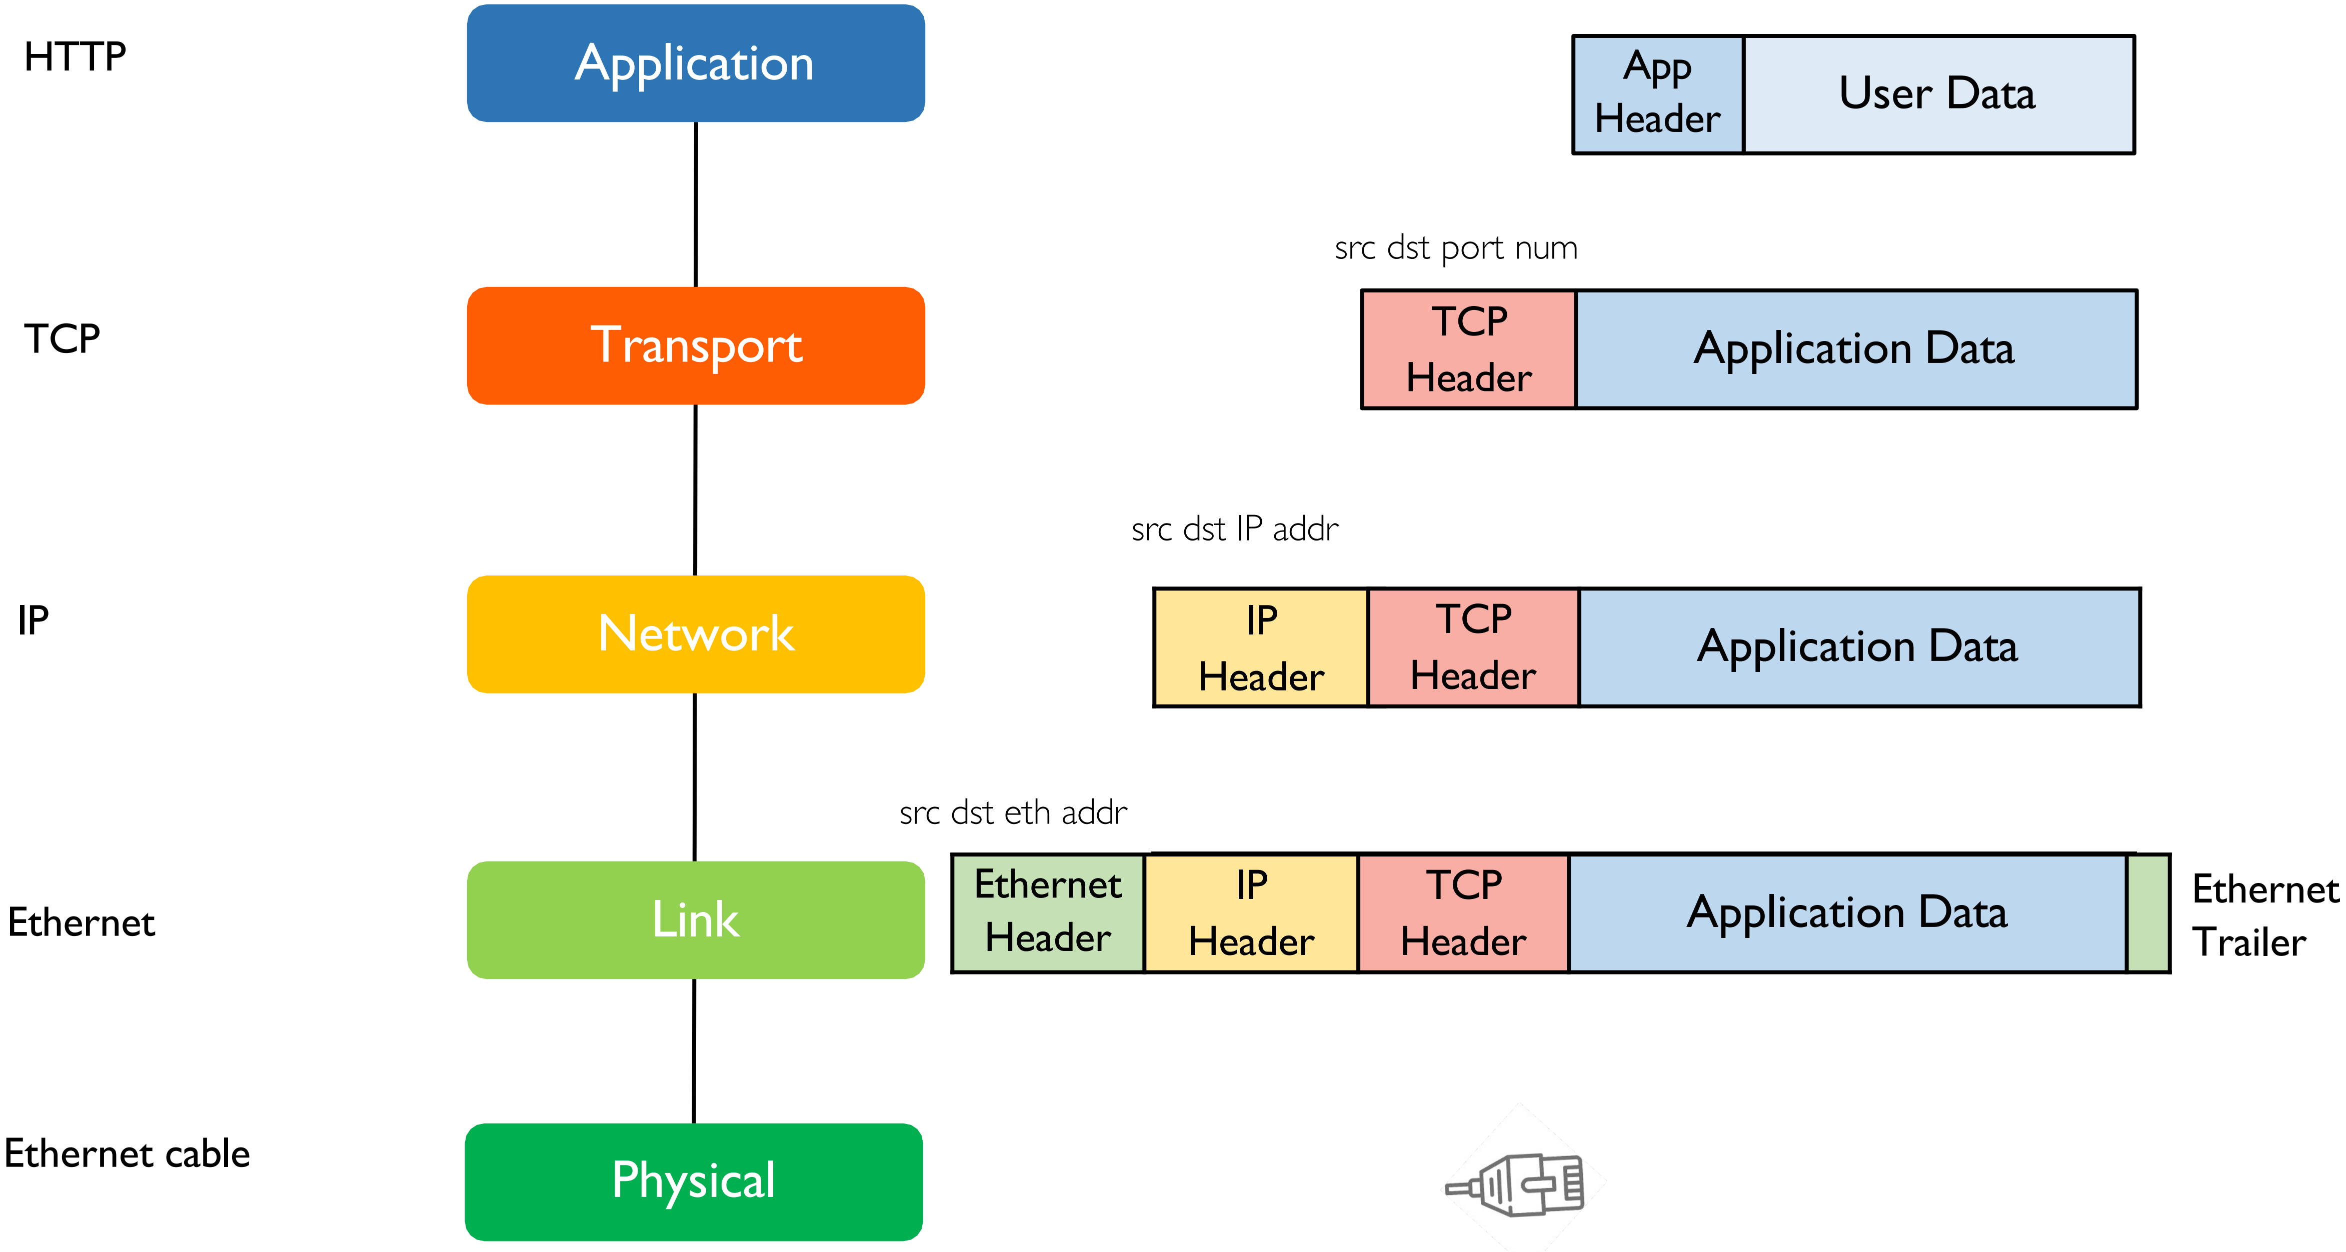

Recap questions

What are two different addresses for one network interface card?
Why have both?

Which header gets changed at each hop?

Which header remains the same from end-to-end?

Outline

1. Administrative and recap
2. Packet forwarding vs routing
3. Packet loss and delay
4. Packet switching vs circuit switching

Switches vs routers: Both do packet forwarding!

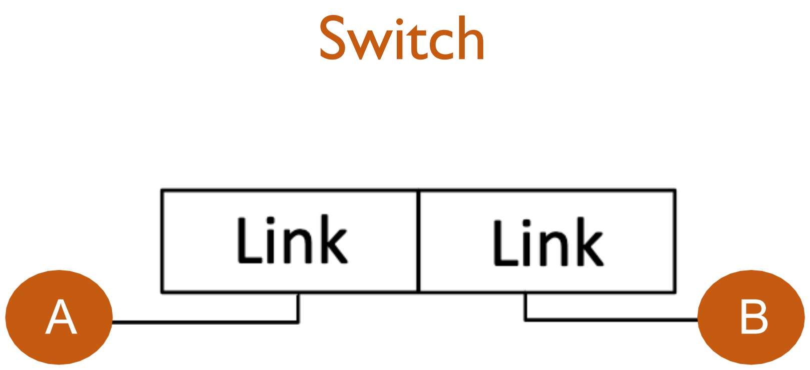

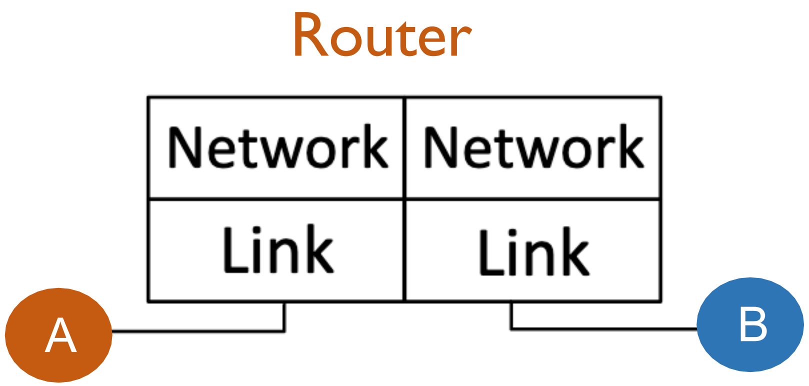

Switch forwards within the same network,
whereas routers forwards across different network

Forwarding vs routing

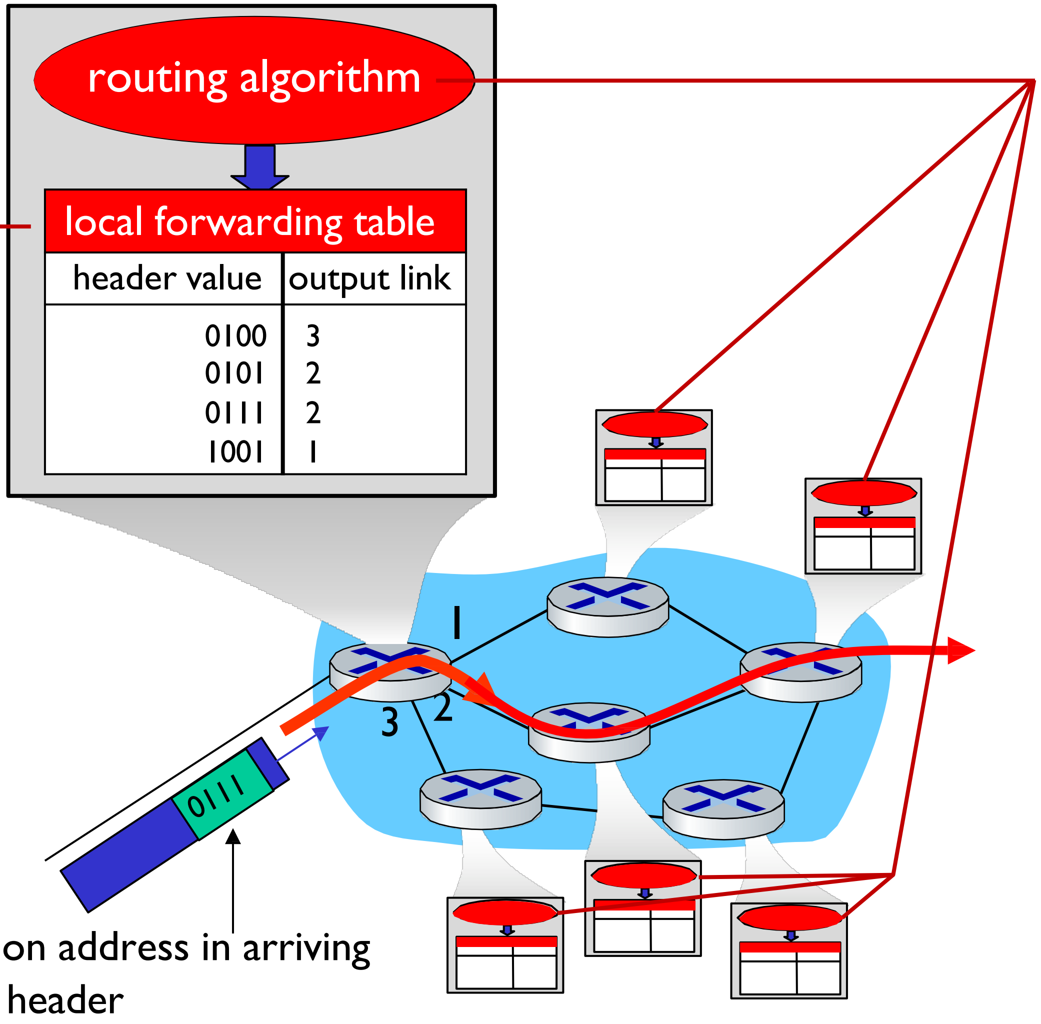

Forwarding:

- local action: move arriving packets to appropriate output link

Routing:

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


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


forwarding
forwarding
45 h
3,053 miles
48 h
3,309 miles
United States
San Jose
Northampton
Sacramento
Cleveland

Outline

1. Goals
2. Packet forwarding vs routing
3. Packet loss and delay
4. Sharing is caring: Packet switching vs circuit switching

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

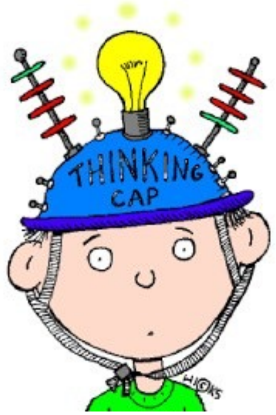

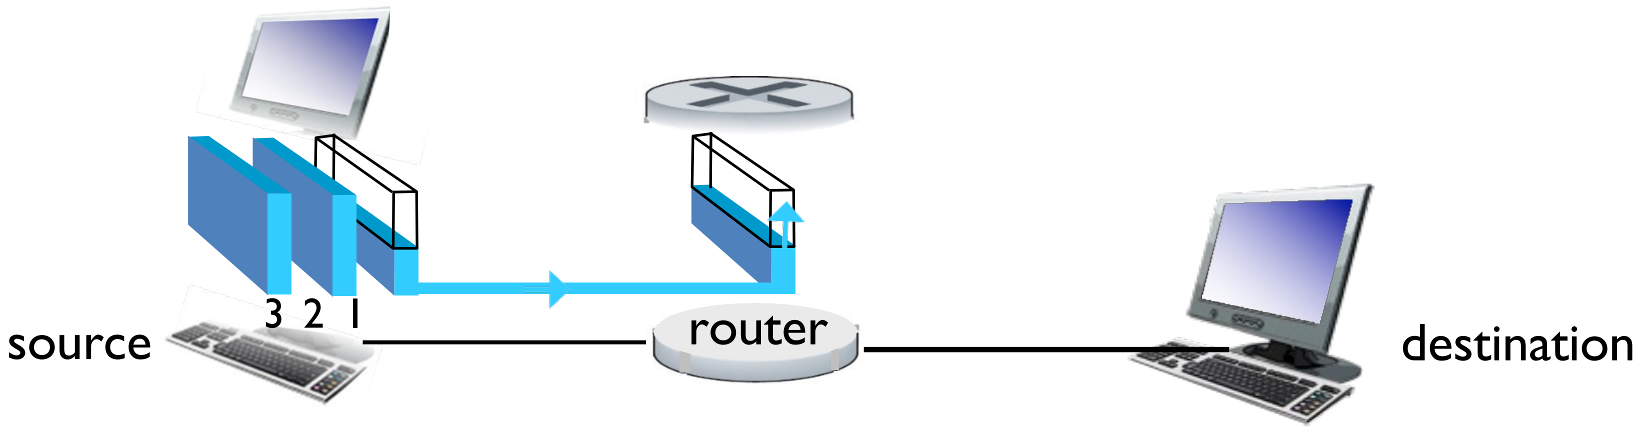

What happens if we don't?

Recap: Pair Activity

Pair and share: Talk to your neighbor about it

- 1 What is the difference between routing vs forwarding?
- 2 What is the difference between a switch and a router?
- 3 What is store-and-forward?
- 4 Is it necessary to do store-and-forward? Why? Or why not?

Network performance

What do you care about your network?

We will look at network delay and loss!

Where do packet delay occur?

- Packet delay:

  The time it took for a sender to send out a packet and till it is received by a receiver

What could be the source of delay?

Imagine…

You are a packet trying to crawl from one node to the next node

What should happen to move from O to U?

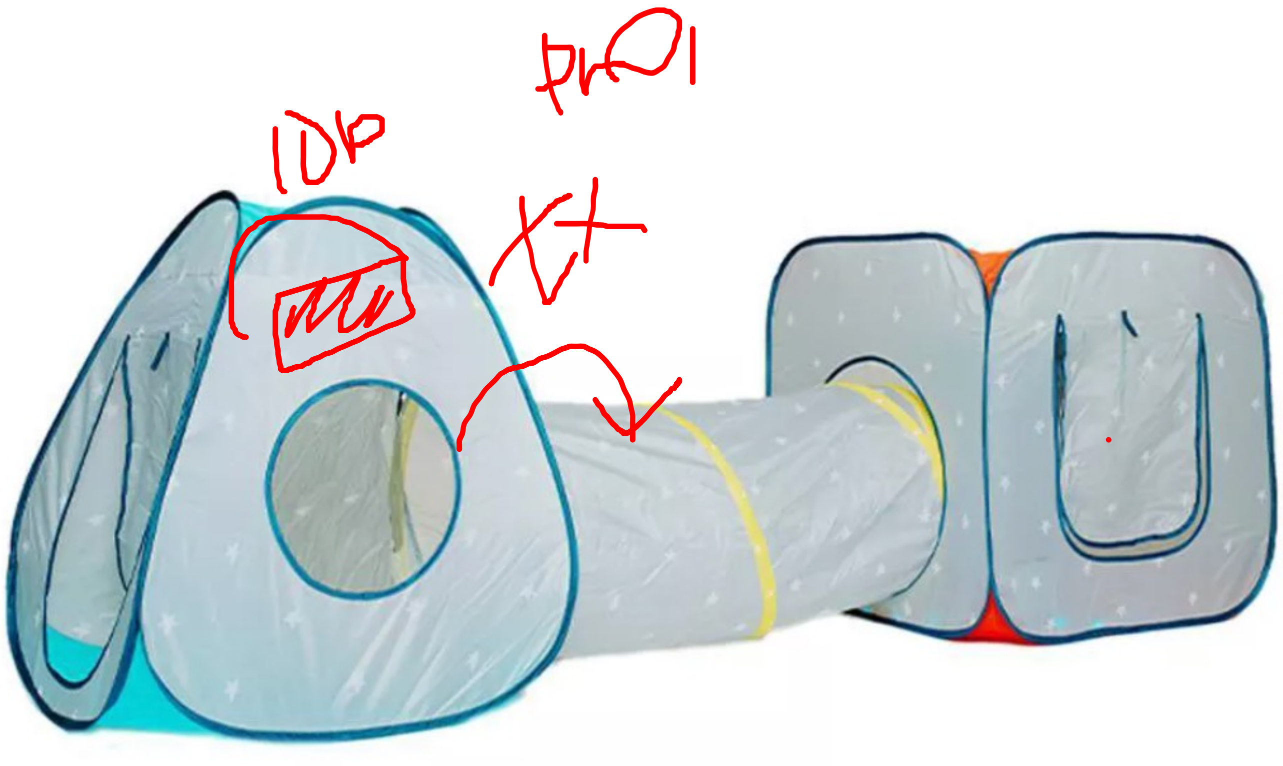

Packet delay: four sources

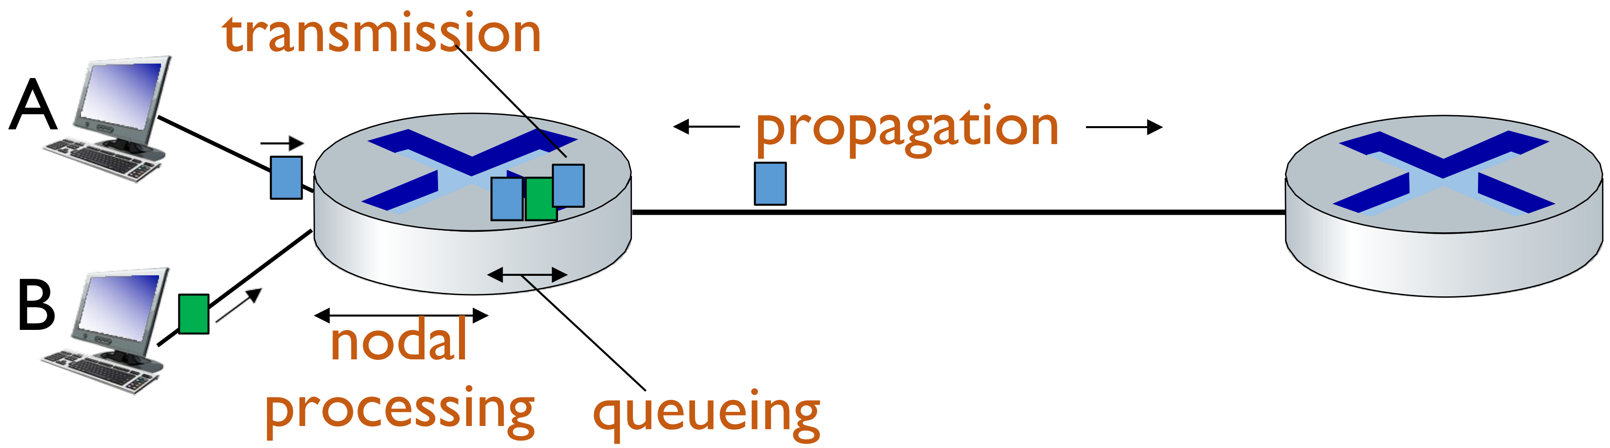

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

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

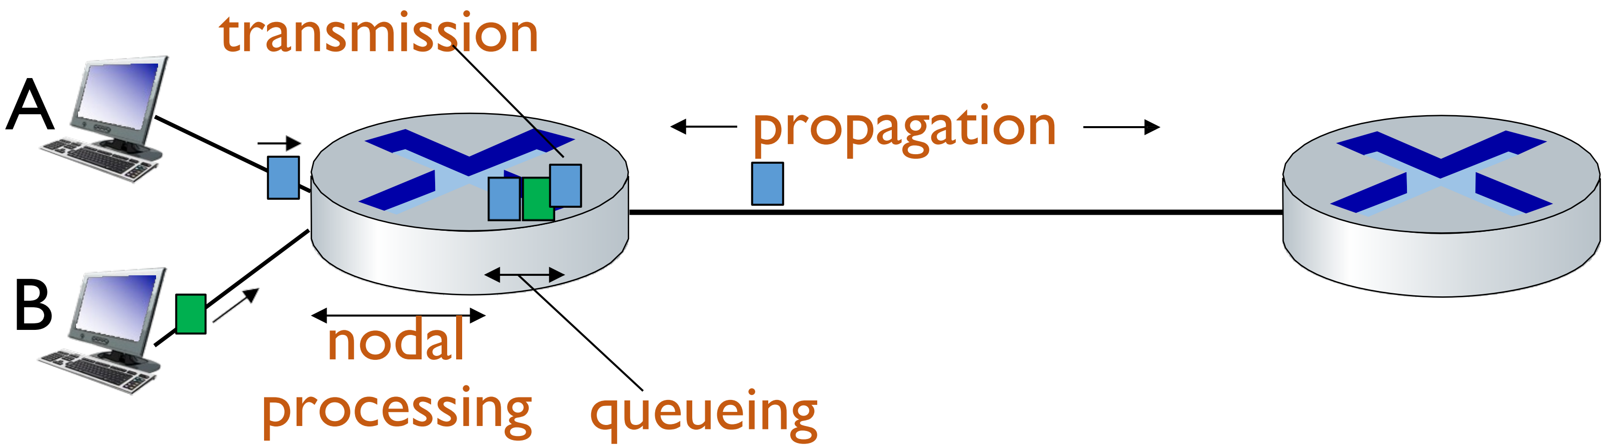

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

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

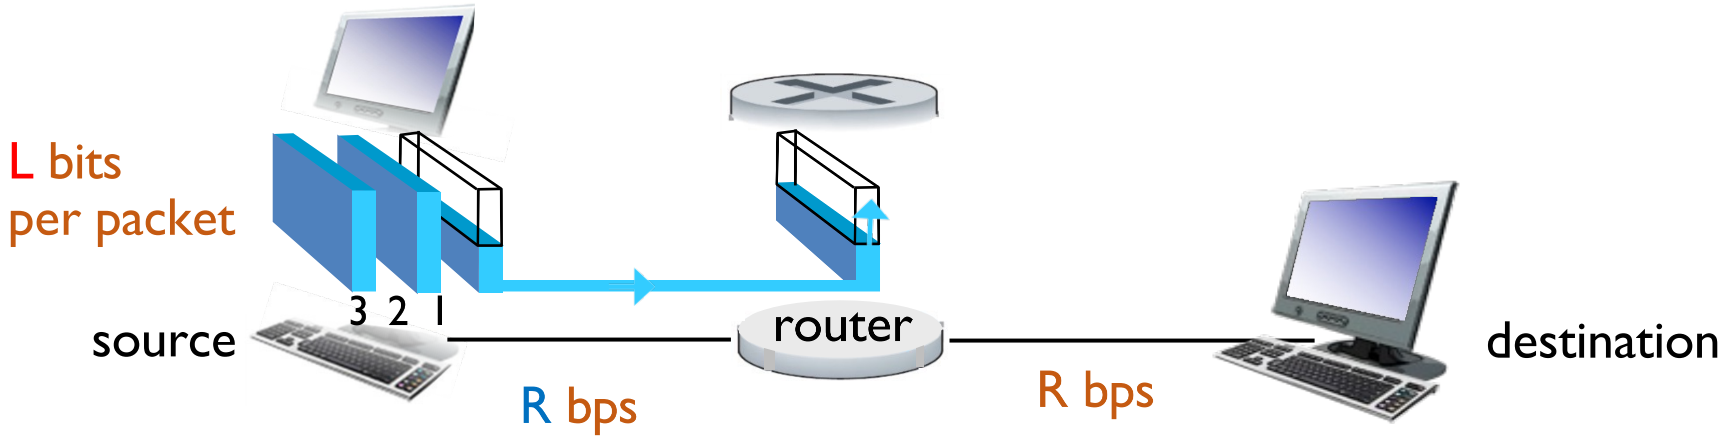

- transmission rate: how fast data is pushed onto a link (in bits per sec)
- transmission delay: time to take to push all bits in the packet to the output link
- d_{trans}: L/R sec

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

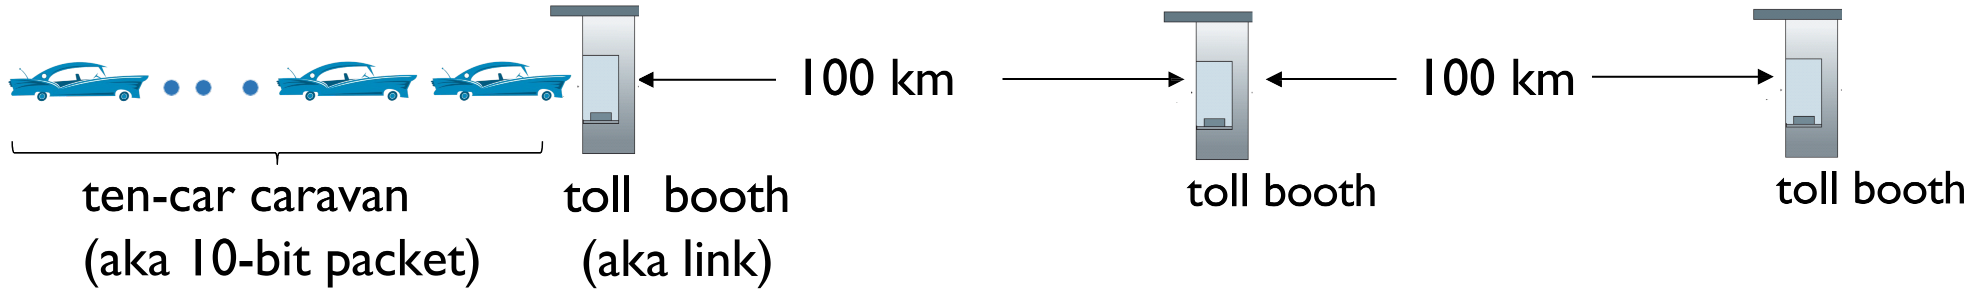

- car ~ bit; caravan ~ packet; toll service ~ link transmission
- toll booth takes 12 sec to service a car (bit transmission time)
- "propagate" at 100 km/hr
- Q: How long until caravan is lined up before 2nd toll booth?
- time to "push" entire caravan through toll booth onto highway = 12*10 = 120 sec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

In-class activity

Recap: Kahoot Exercise

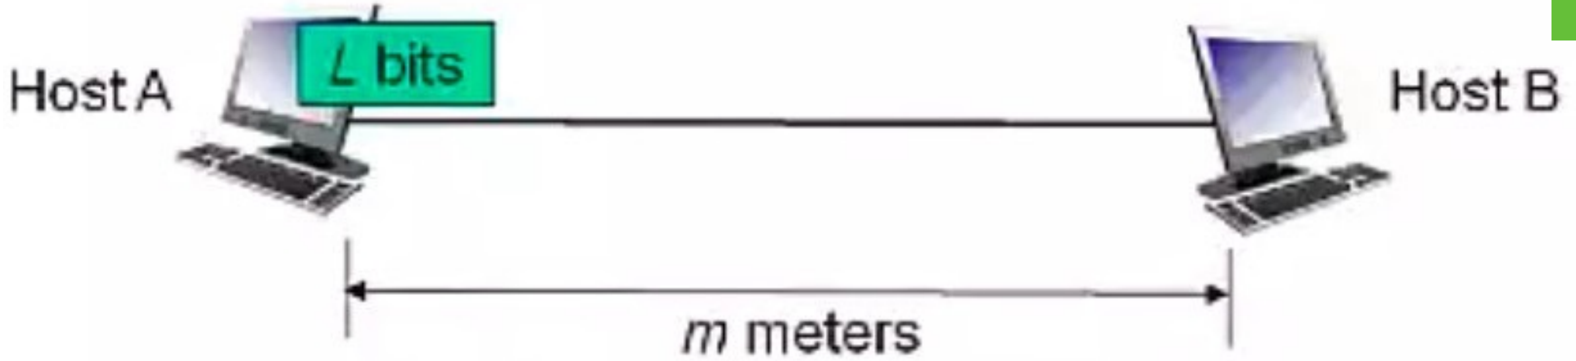

transmission rate = R bits/sec

propagation speed = s meters/sec

At time t=d_trans, where is the last bit of the packet?
(d_trans is the trasmission delay)

Recap: Kahoot Exercise

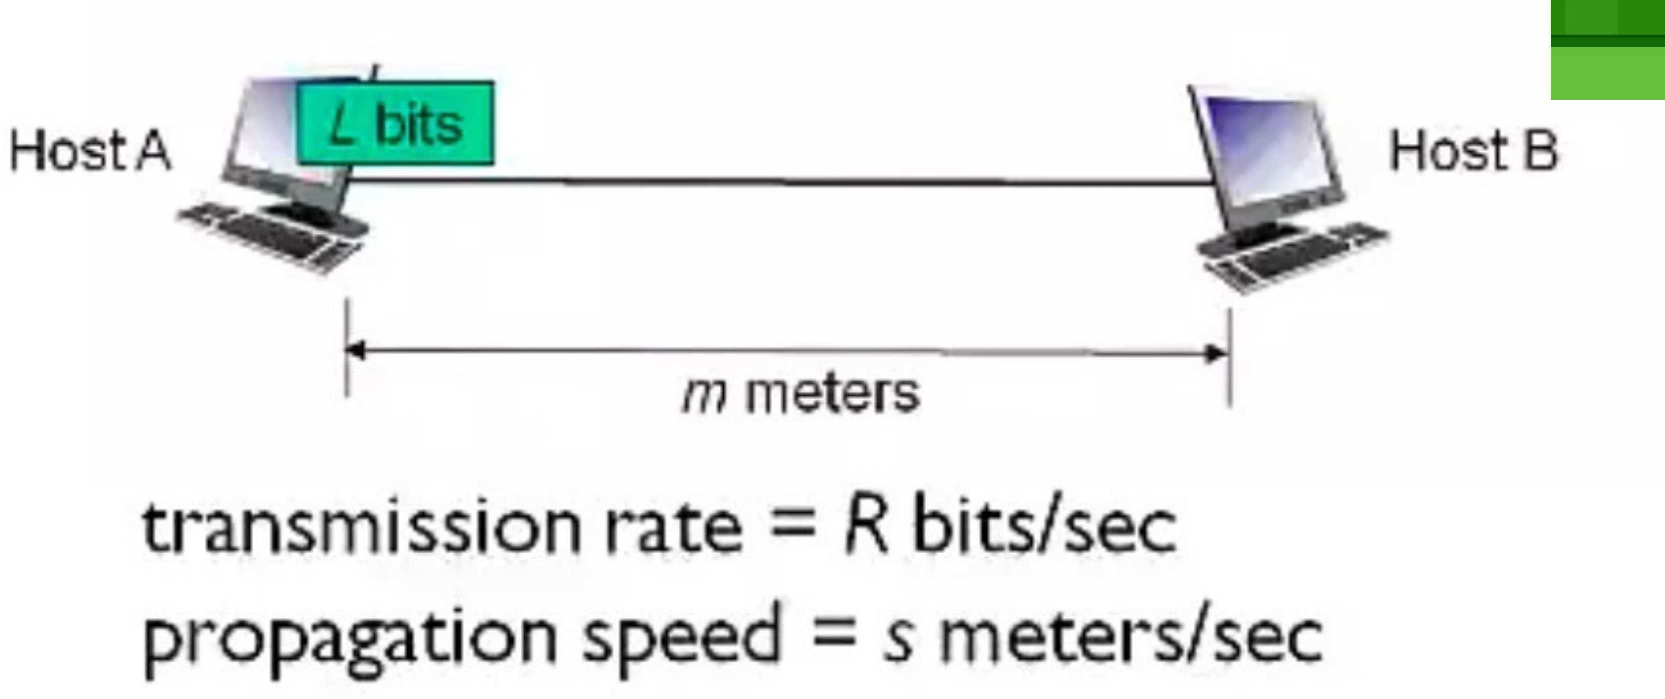

Suppose d_prop > d_trans, at t=d_trans where is the first bit?

Recap: Kahoot Exercise

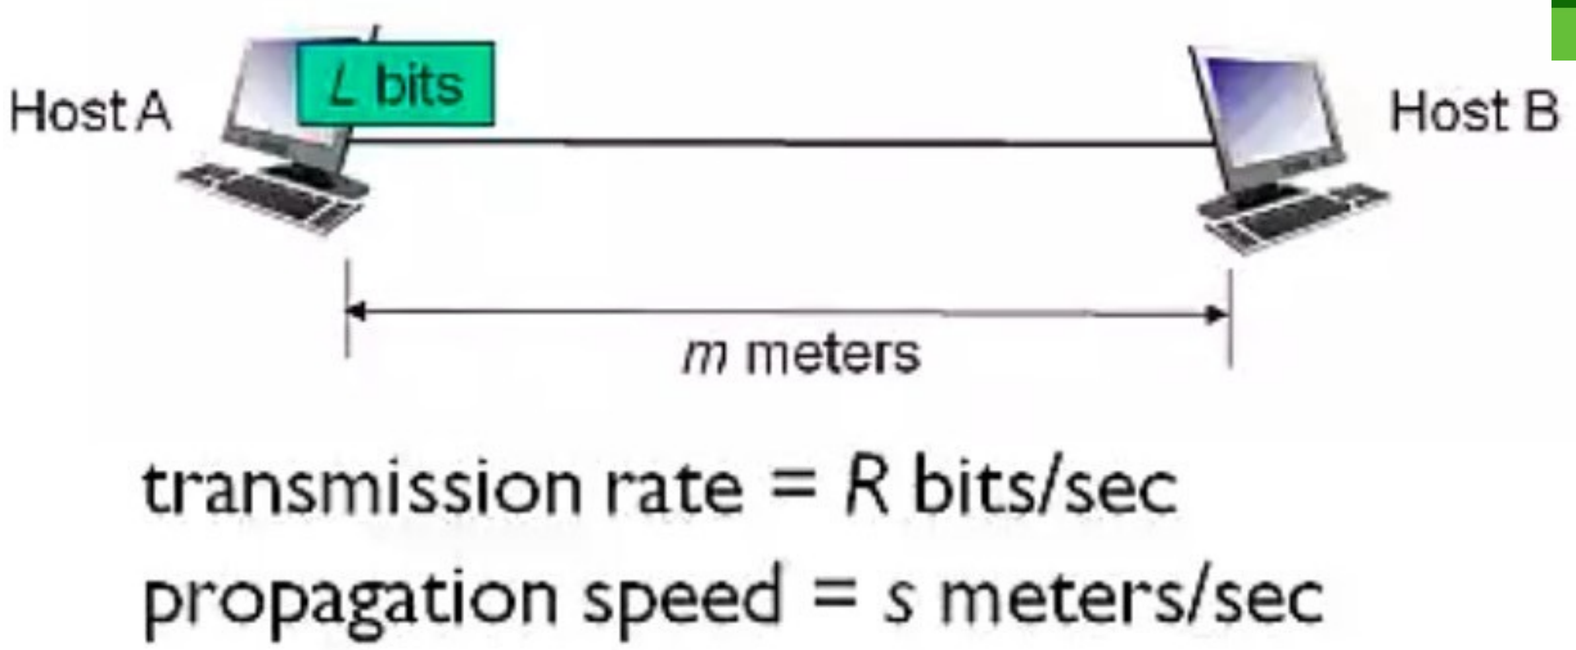

Suppose d_prop < d_trans, at t=d_trans where is the first bit?

Take a screenshot of your Kahoot score

- Submit to Canvas!

Packet delay: four sources

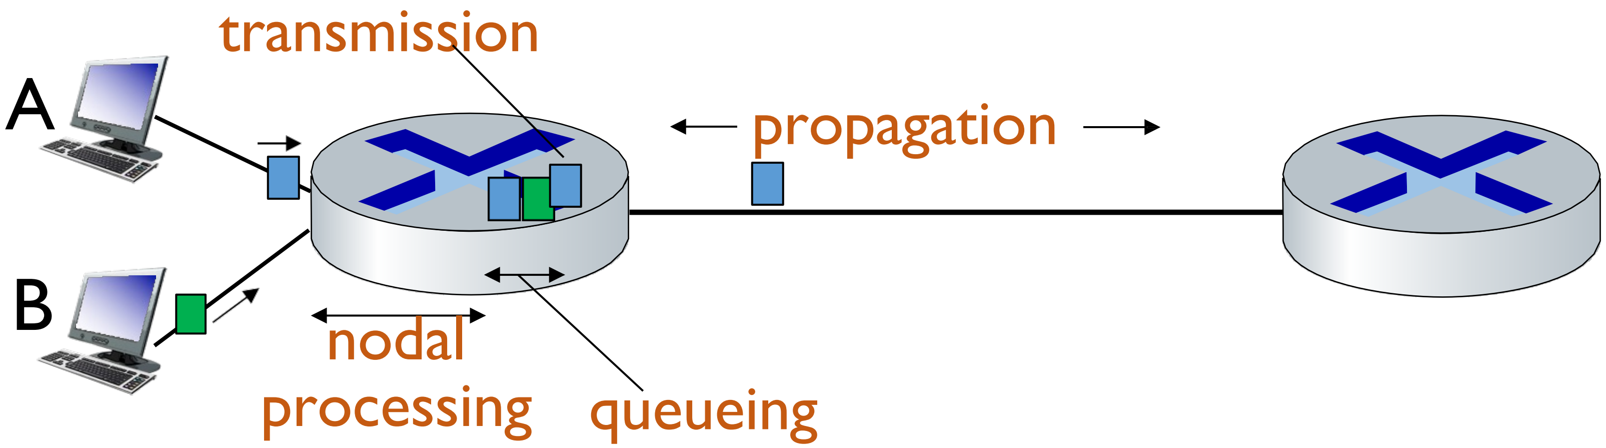

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

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

Queueing delays

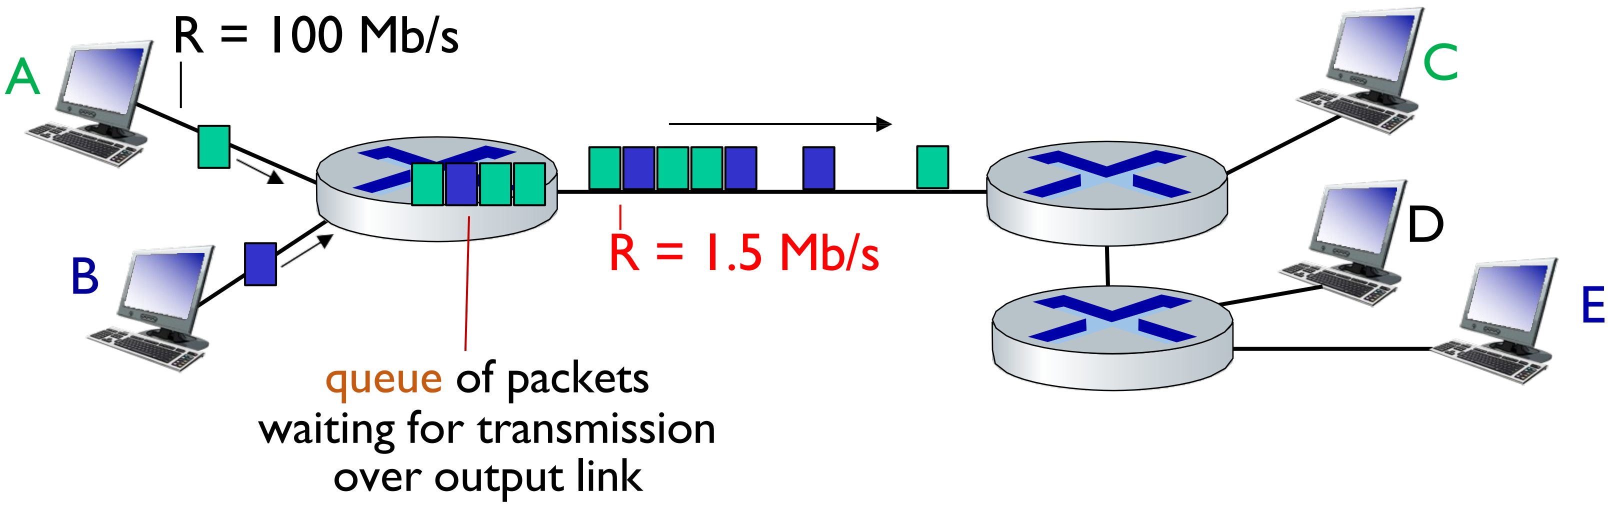

Queueing occurs when work arrives faster than it can be serviced:

Queueing delay analysis

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

$\frac{L \cdot a}{R}$: <u>arrival rate of bits</u> / service rate of bits — "traffic intensity"

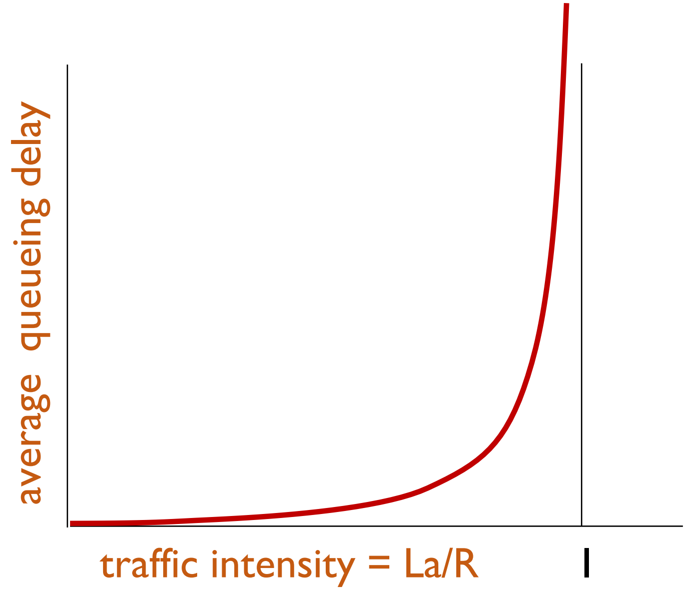

- La/R ~ 0: avg. queueing delay small
- La/R = 1: avg. queueing delay large
- La/R > 1: more "work" arriving is more than can be serviced - average delay infinite!

La/R ~ 0

La/R => 1

Where else can packet loss 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!

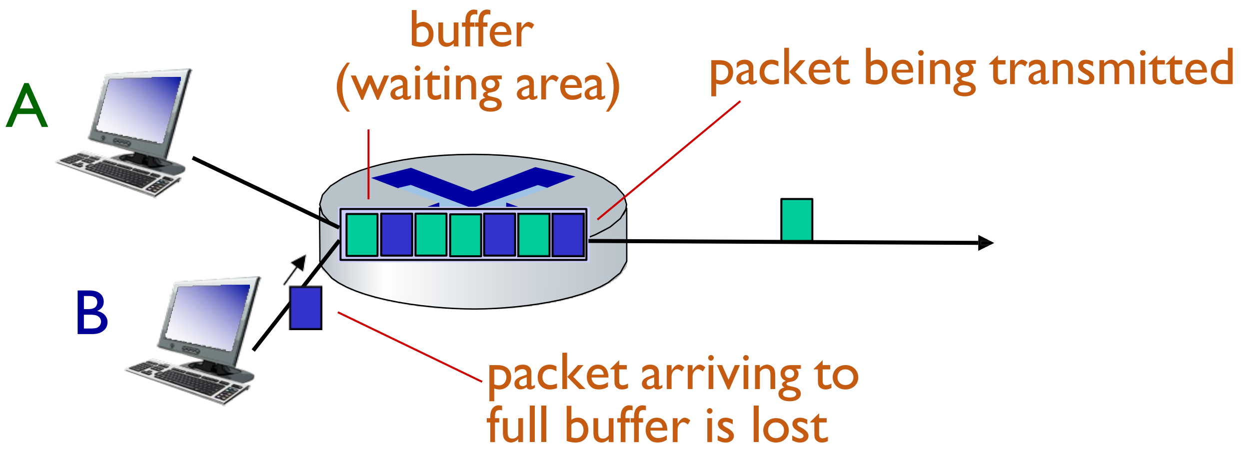

What are other possible ways for a packet to get lost?

Outline

1. Goals
2. Packet forwarding vs routing
3. Packet loss and delay
4. 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

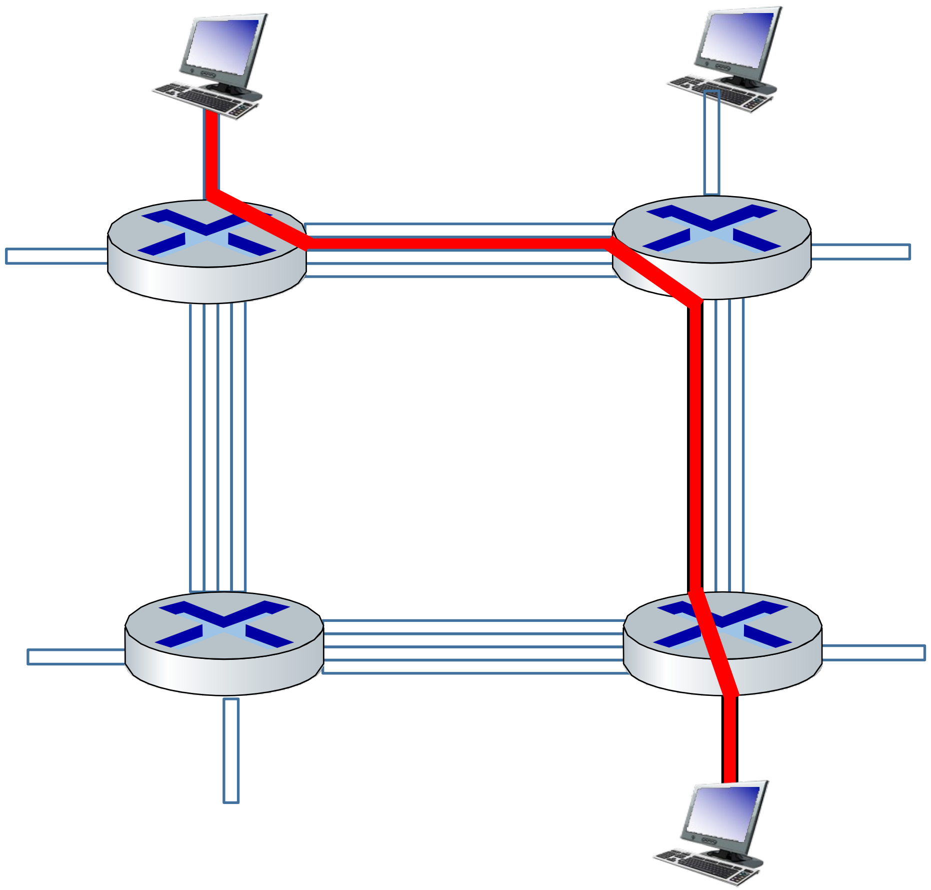

Circuit switching: FDM and TDM

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

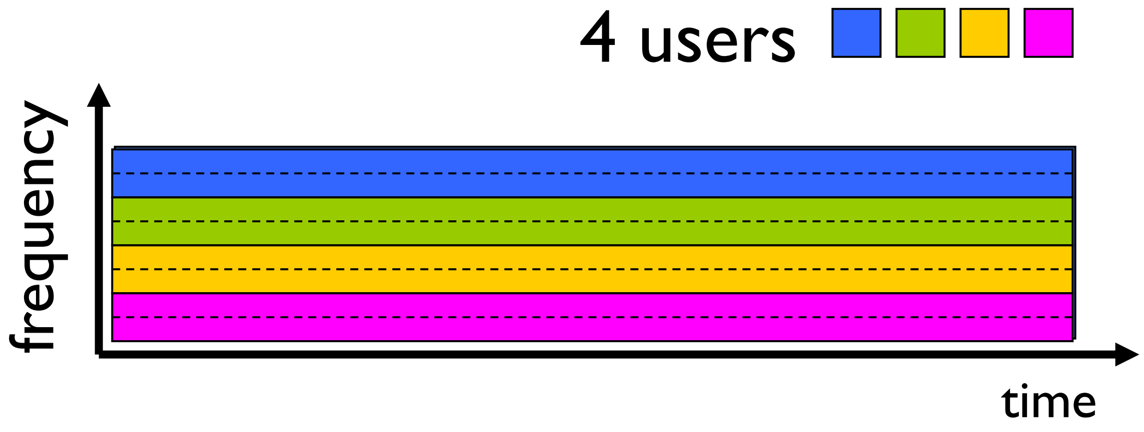

Time Division Multiplexing (TDM)

- 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)

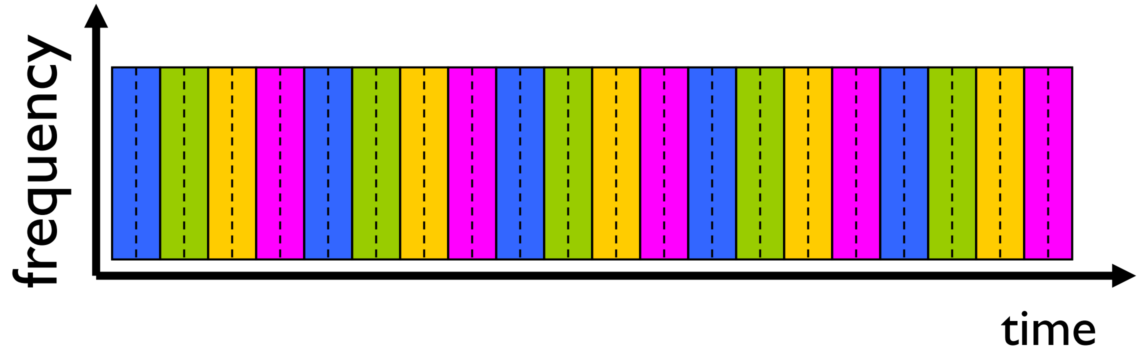

Circuit switching reserves 100Mbps per user

example:

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

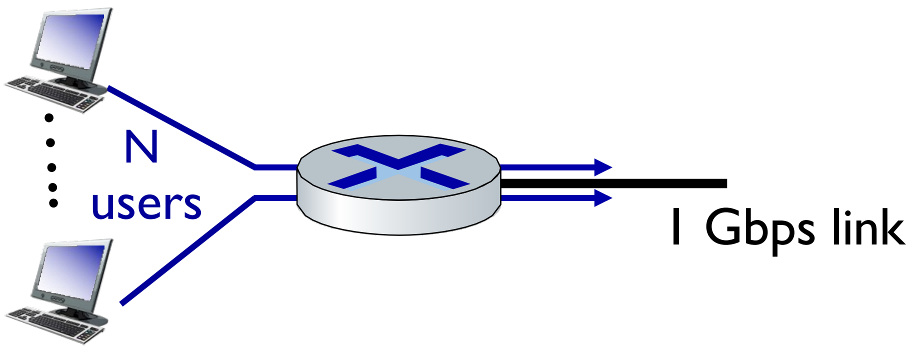

Q: What is the max number of users that can share this network?

- circuit-switching:

Only 10 users (=1 Gbps/100 Mbps)

Given each users are active only 10% of the time
Can we allow more number of users?

Packet switching allows more users to share with some probability of failure

example:

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

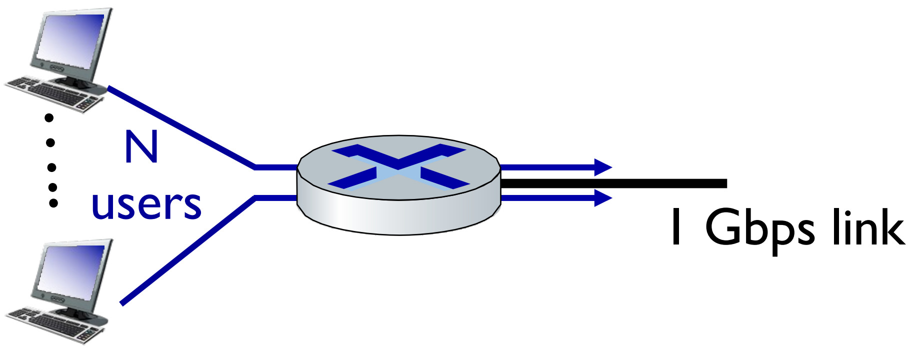

Q: What is the max number of users that can share this network?

- packet switching:

A: Need some assumption on link availability guarantee.

Say we guarantee 99.99% link availability for each user.

That is failure rate < 0.01% == 0.0001

When does "failure" happen?

When does failure happen?

example:

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

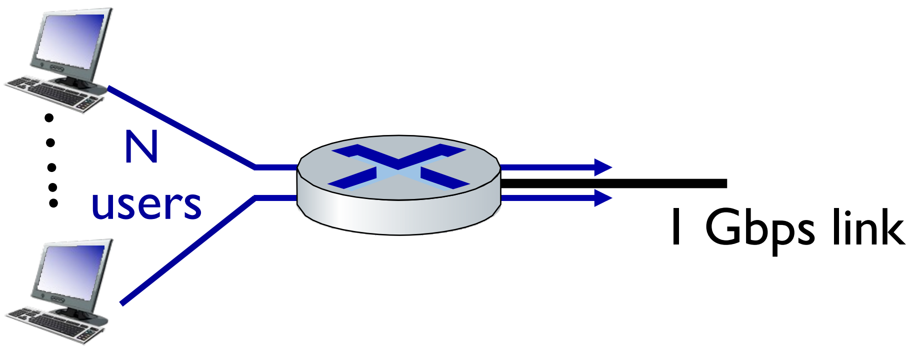

Whenever more than 10 users happen to be active simultaneously!

Say total number of users is 11 (N=11)

example:

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

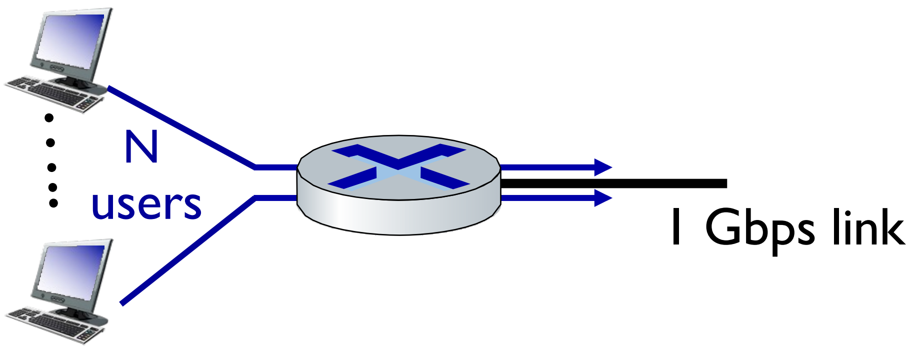

What is the probability of failure?

- Case 1: When 11 out of 11 users send simultaneously

$\text{FailureProb}_{N=11}$ = Prob(case 1)

If $\text{FailureProb}_{N=11} < 0.0001$, we can add one more user!

Say total number of users is 12 (N=12)

example:

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

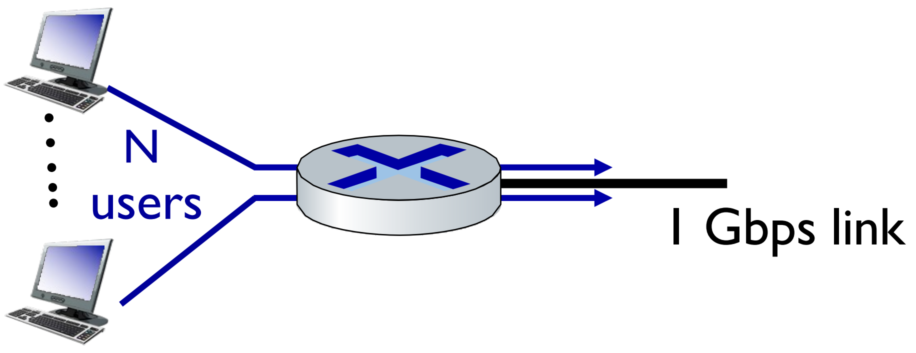

What is the probability of failure?

- Case 1: When 11 out of 12 users send simultaneously
- Case 2: When 12 out of 12 users send simultaneously

$FailureProb_{N=12}$ = Prob(case 1) + Prob(case 2)

If $FailureProb_{N=12} < 0.0001$, we can add one more user!

Say total number of users is 13 (N=13)

example:

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

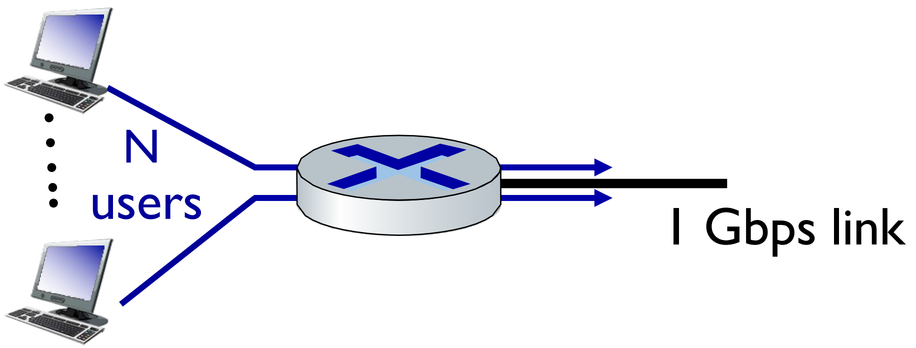

What is the probability of failure?

- Case 1: When 11 out of 13 users send simultaneously
- Case 2: When 12 out of 13 users send simultaneously
- Case 3: When 13 out of 13 users send simultaneously

$\text{FailureProb}_{N=13}$ = Prob(case 1) + Prob(case 2) + Prob(case 3)

If $\text{FailureProb}_{N=13}$ < 0.0001, we can add one more user!

Here is the algorithm to find max N

1 Find out the num of users in case of circuit switching. Let's call this m

2 Set initial N = m + 1

3 Calculate failure probability given N.
If it is smaller than THRESHOLD, increment N and repeat Step 3

What is the probability of k users out of N users are active simultaneously?

- Let's scope it down
- Say, N = 5
- There are total 5 users in the system.
  What is the probability that 2 of them are active simultaneously?

The probability of exactly 2 users are active among 5 users

- This means all other users must be silent (non-active)
- One possible outcome: AANNN
- What's the probability to get above?

  $P(A) \times P(A) \times P(N) \times P(N) \times P(N) = P(A)^2 \times P(N)^3 = 0.1^2 \times (1\text{-}0.1)^3$
- Another possible outcome: NNNAA

  $P(N) \times P(N) \times P(N) \times P(A) \times P(A) = P(A)^2 \times P(N)^3 = 0.1^2 \times (1\text{-}0.1)^3$
- Note all outcomes have the same prob of happening: $P(A)^2 \times P(N)^3$
- How many possible outcome?

  ${}_5C_2 = 10$
- Putting it together

  ${}_5C_2 \times P(A)^2 \times P(N)^3 = 10 \times 0.1^2 \times (1\text{-}0.1)^3 = 0.0729$

This is an example of binomial distribution

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?

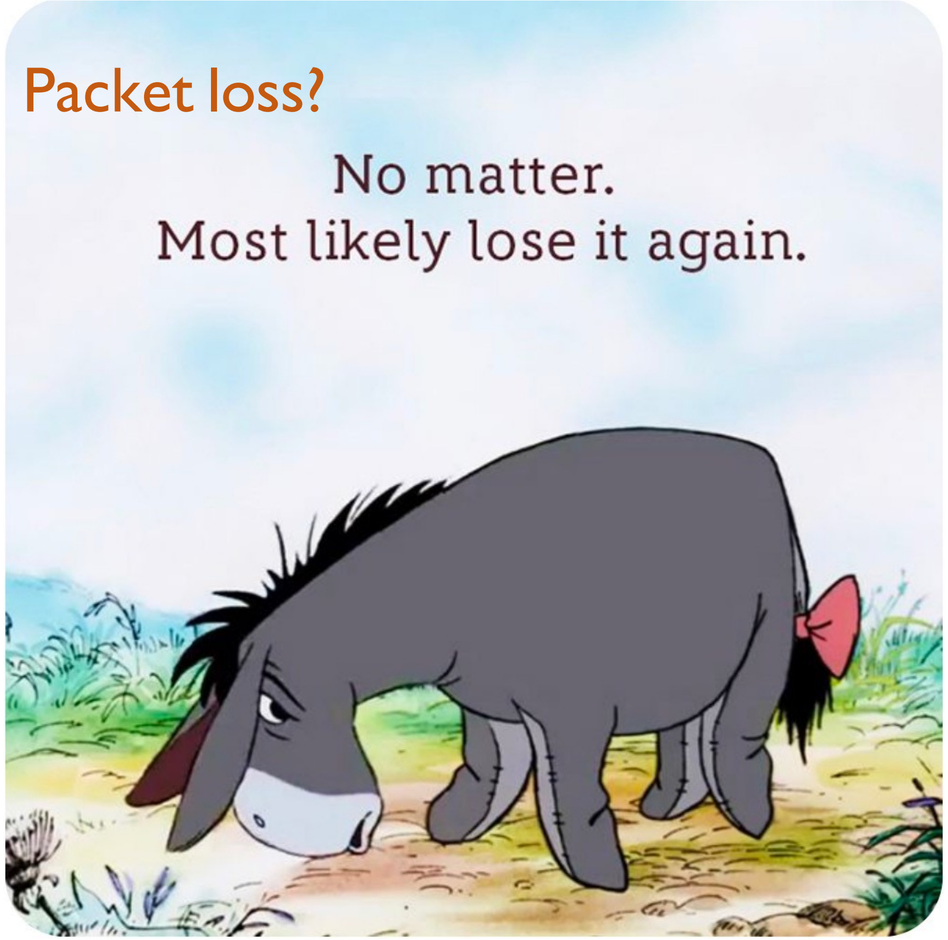

Backup Slides

Caravan analogy

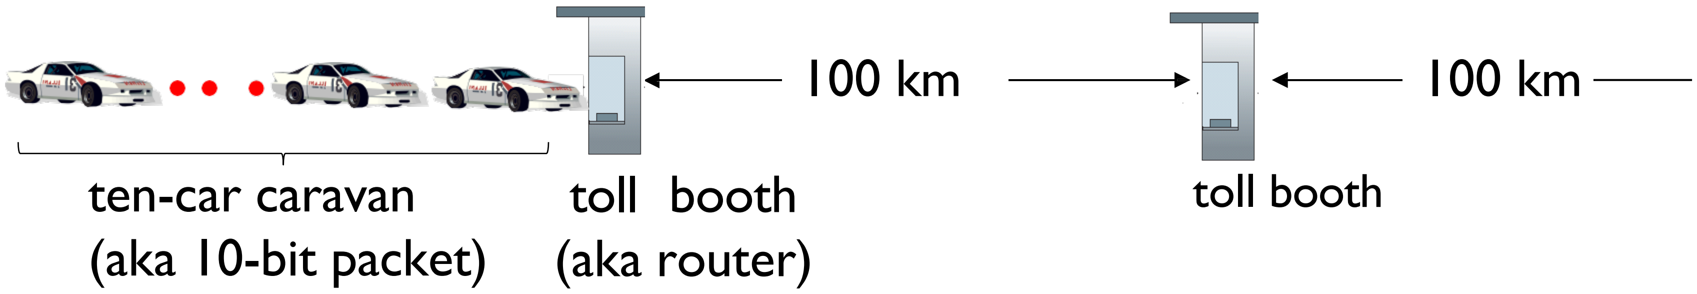

- suppose each car now "propagates" at 1000 km/hr
- and suppose toll booth now takes one min to service a car
- Q: Will cars arrive to 2nd booth before all cars serviced at first booth?

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

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 original sender to final receiver

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

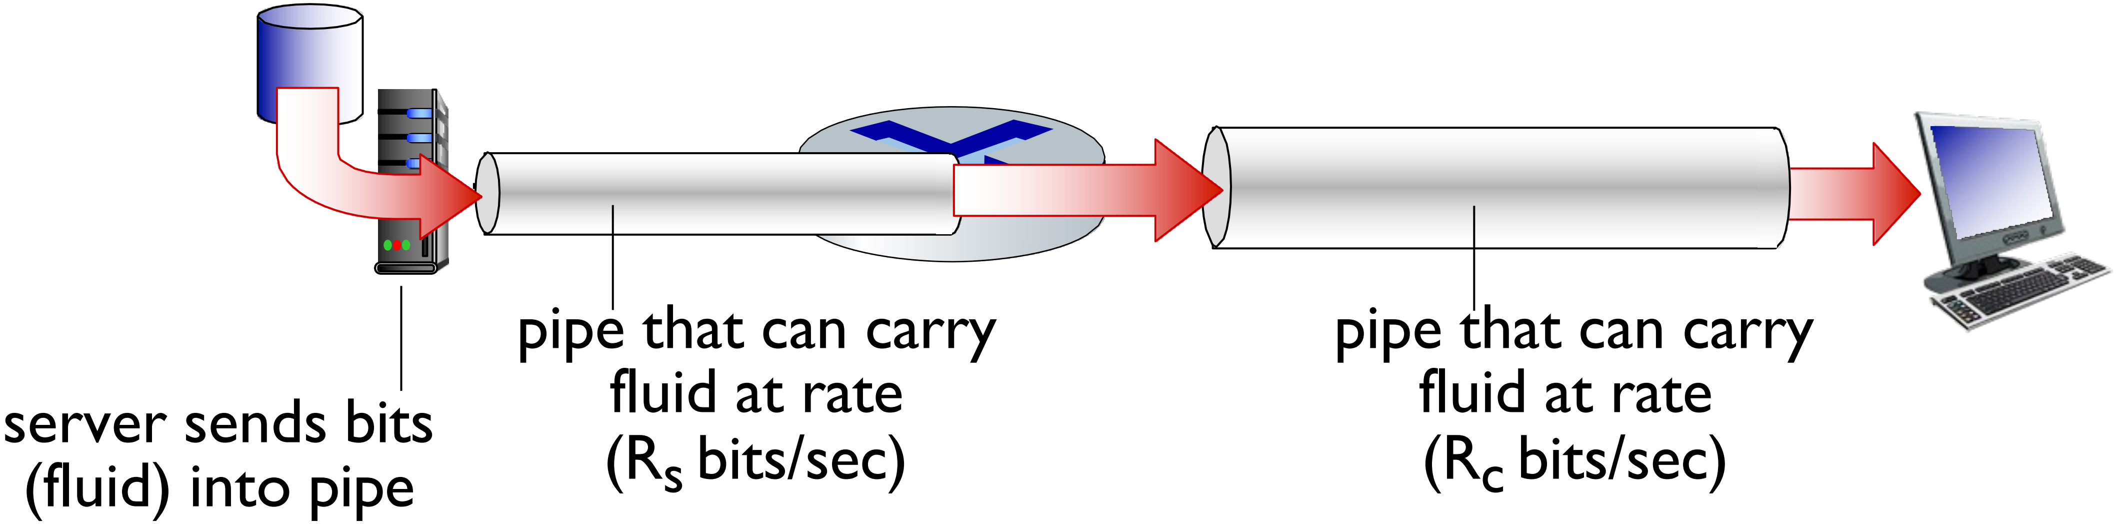

Throughput

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

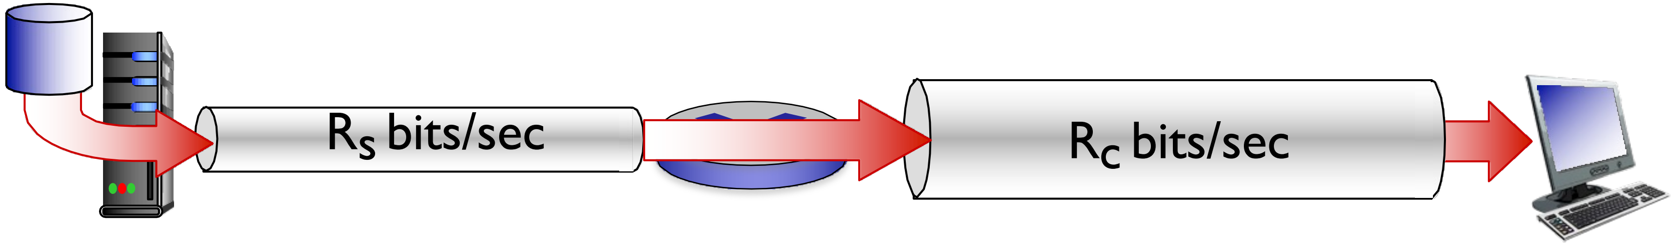

$R_s > R_c$ What is average end-end throughput?

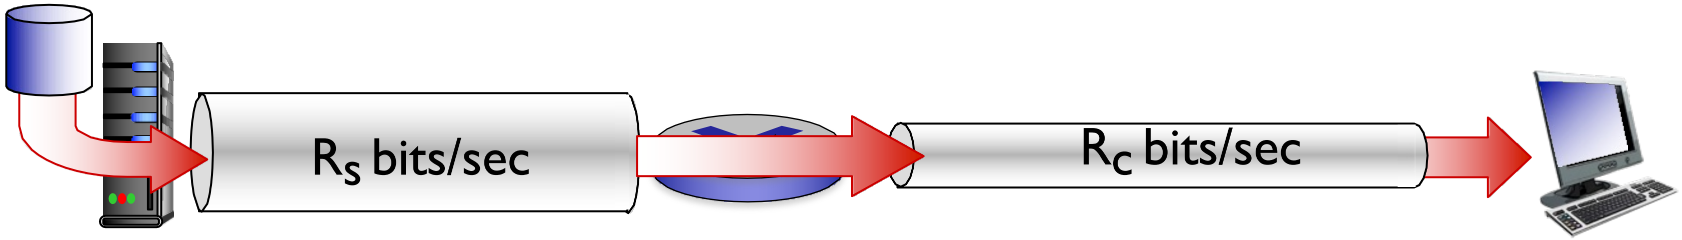

bottleneck link

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

Throughput: network scenario

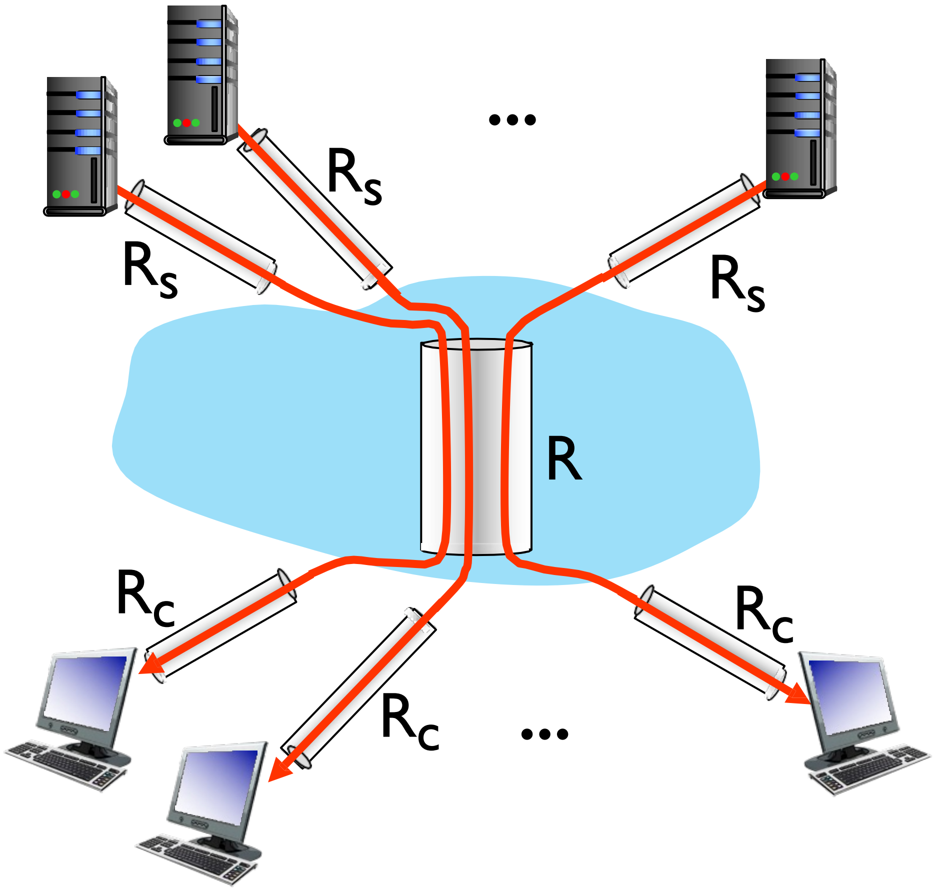

Say 10 connections fairly share backbone bottleneck link R bits/sec

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

Acknowledgements

Slides are adopted from Kurose' Computer Networking