<u>This lecture will be recorded!</u>

Welcome to

CS 3516:
Computer Networks

Prof. Yanhua Li

Time: 9:00am –9:50am M, T, R, and F
Zoom Lecture
Fall 2020 A-term

Lab assignment 1

We are grading it.

To be done by tomorrow (Tuesday)!

Quiz 2

Tomorrow (Tuesday)

- Topics: Network Performances, such as delay, loss, throughput

Chapter 1: roadmap

1.3 network core

- packet switching, circuit switching, **network structure**

1.4 network performance in packet-switched networks

delay

loss

throughput

Chapter 1: roadmap

1.3 network core

- packet switching, circuit switching, network structure

1.4 network performance in packet-switched networks

Cons:

delay

loss

throughput

Pros:

Easy to implement

Support more users

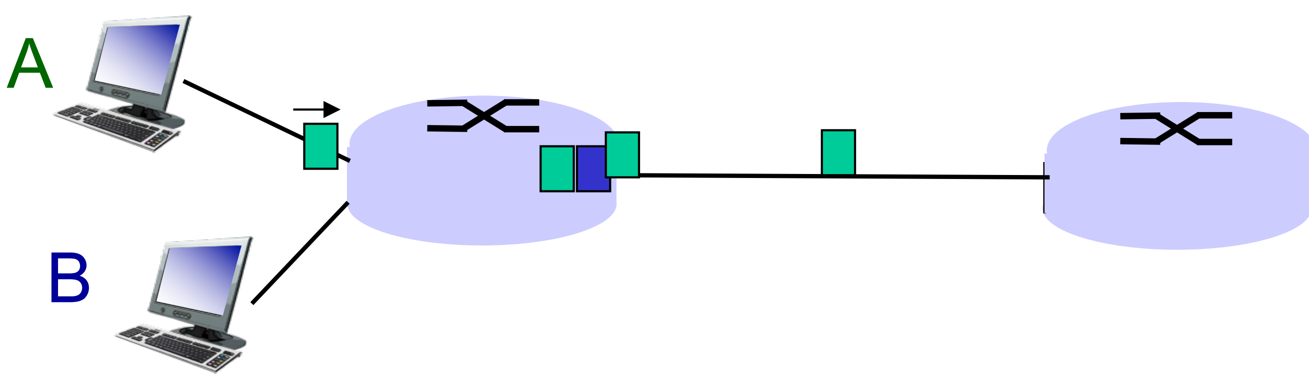

Four sources of packet delay

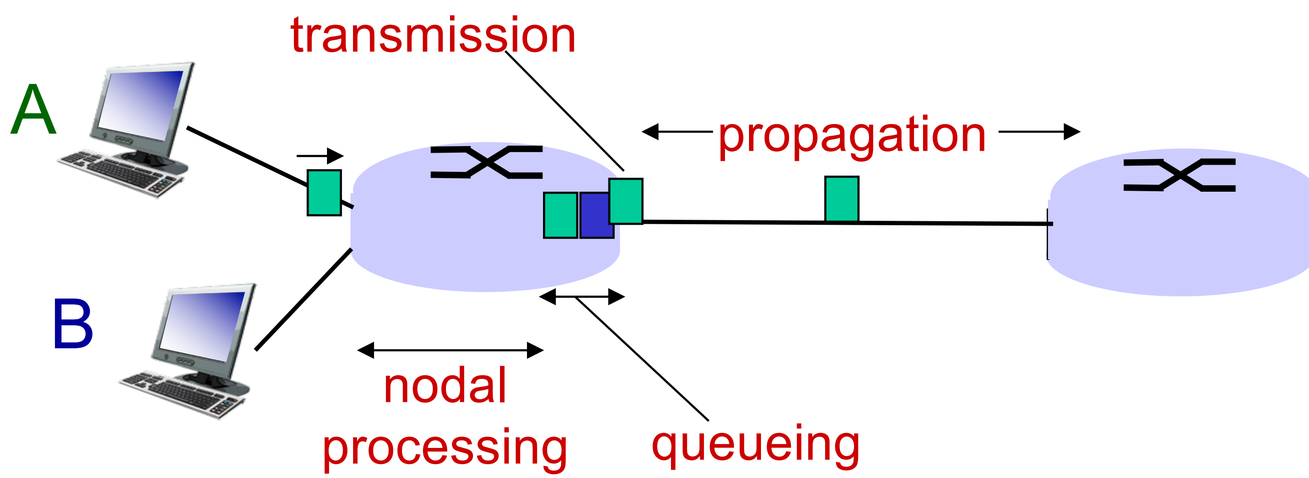

$$d_{nodal} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

When a packet can be transmitted on a link?

- 1) no other pkt transmitted on the link (in practice)
- 2) no other pkt preceding it in the queue

bandwidth-delay product

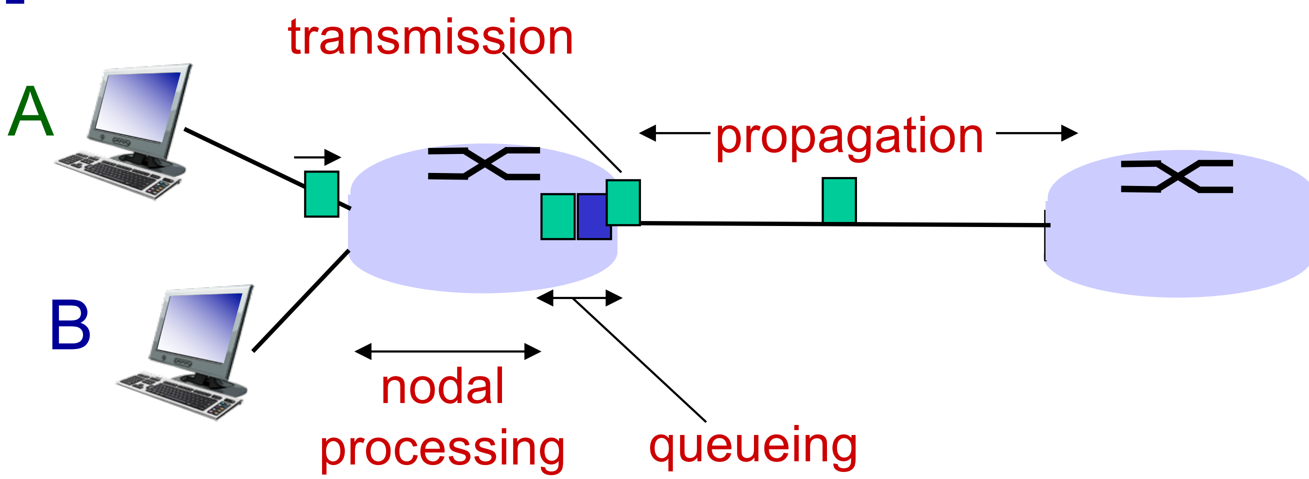

bandwidth-delay product= RT_{prop}*

- *R*: link *bandwidth (bps)*
- *Propagation delay:*T_{prop}=*d/s*

Example:

R = 10Mbps, d=10,000km, s=10^9 m/s, what is the bandwidth-delay product?

T_{prop}=d/s=10^7m / 10^9 m/s =0.01s

Bandwidth-delay product = R*T_{prop}=10Mbps * 0.01s =0.1Mb=100kb.

It represents the maximum number of bits that can be in the link at any given time.

Queueing delay (revisited)

- *R:* link bandwidth (bps)
- *L:* packet length (bits)
- a: average packet arrival rate (pkt/s)
- *La*: the packet arrival rate in bps

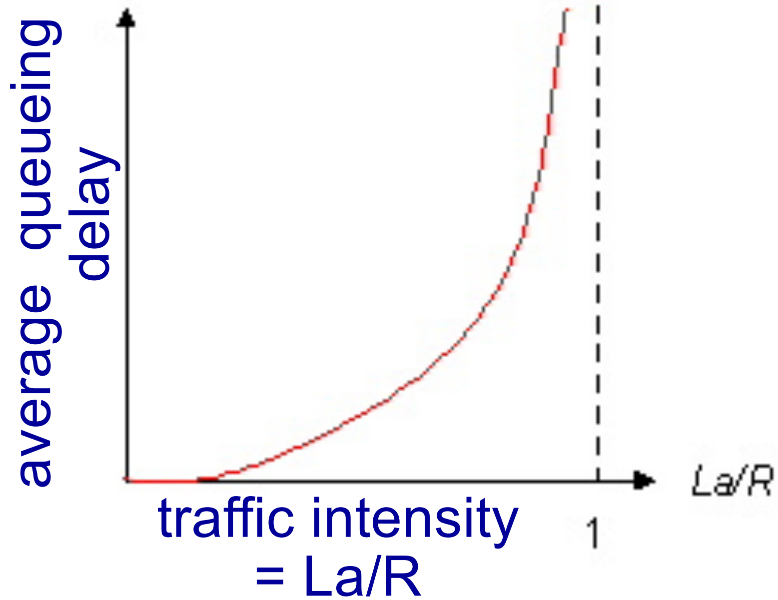

- *La/R: Traffic Intensity*
- *La/R* ~ 0: avg. queueing delay small
- *La/R* -> 1: avg. queueing delay large
- *La/R* > 1: more "work" arriving than can be serviced, average delay infinite!

La/R ~ 0

La/R -> 1

End-to-end delay

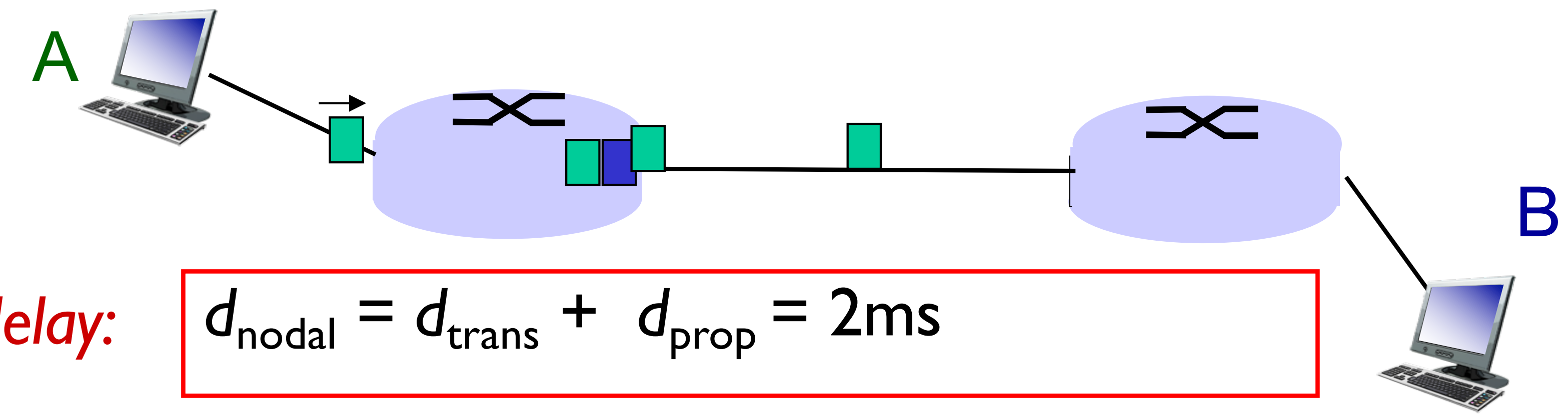

Nodal delay: $d_{nodal} = d_{trans} + d_{prop} = 2ms$

Ten packets are sent from A to B:

Assumption: the time at which the 2nd packet is received at the first router is equal to the time at which the 1st packet is received at the second router.

Nodal delay: single packet end-to-end delay.

$D_{end\text{-}to\text{-}end,\ single} = 3*d_{nodal} = 6\ ms$

Nodal delay: ten packet end-to-end delay.

$D_{end\text{-}to\text{-}end,\ 10} = 3*d_{nodal} + 9 * d_{nodal} = 24\ ms$

Chapter 1: roadmap

1.4 network performance in packet-switched networks

- delay
- **loss**
- throughput

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

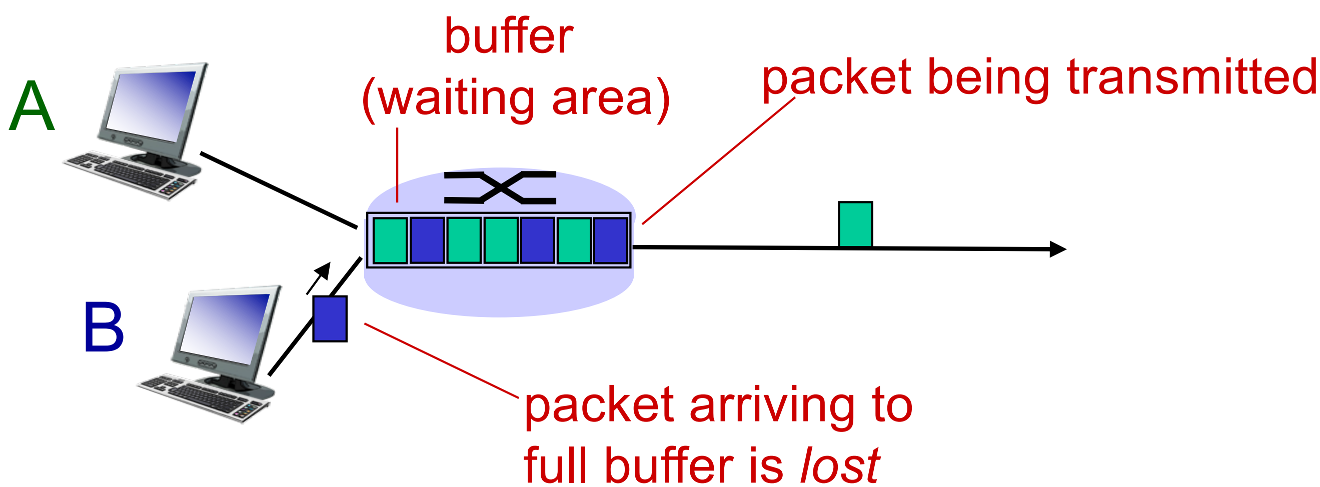

Chapter 1: roadmap

1.4 network performance in packet-switched networks

- delay
- loss
- **throughput**

Throughput (an end-to-end measure)

- *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

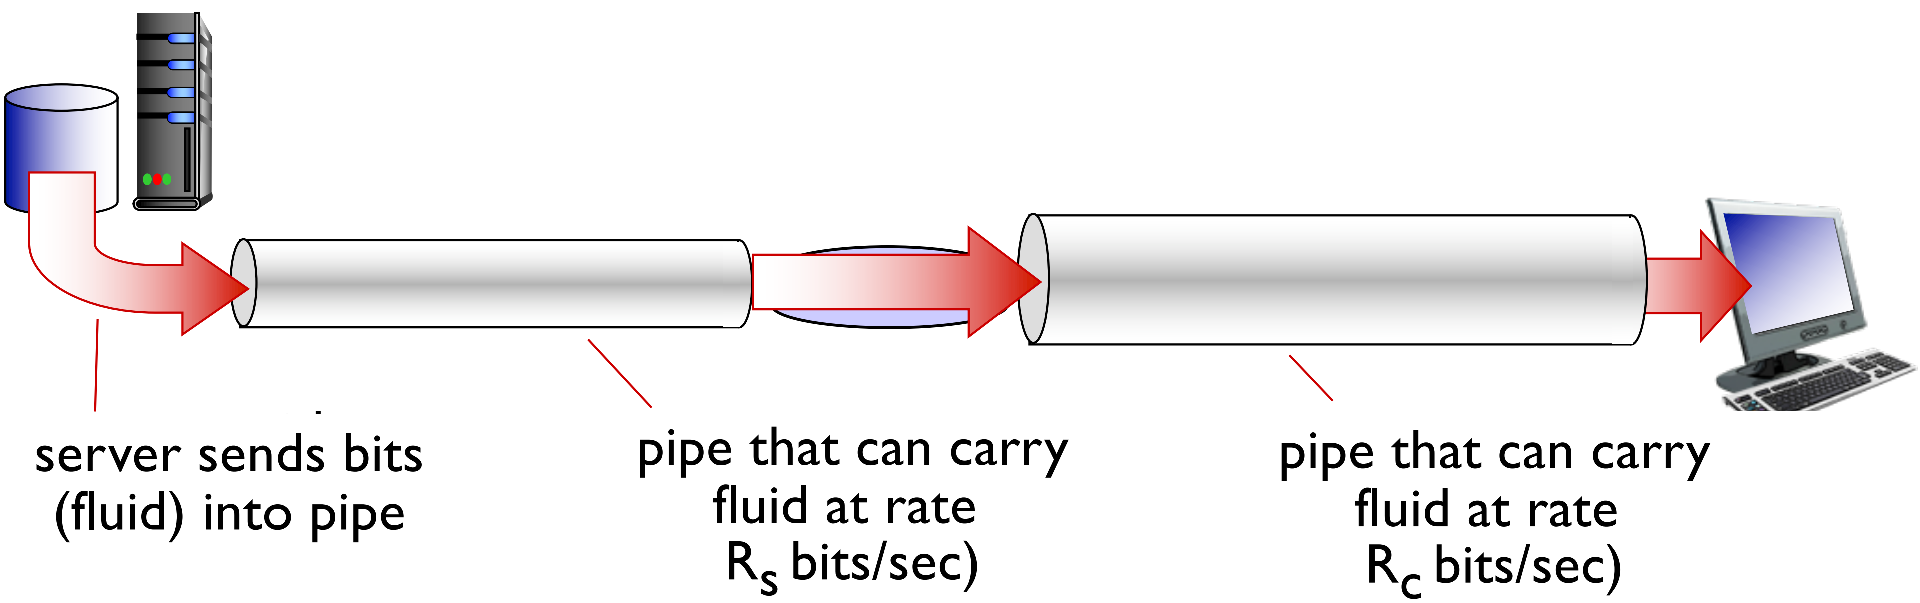

Throughput (more)

$\min(R_c, R_s)$

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

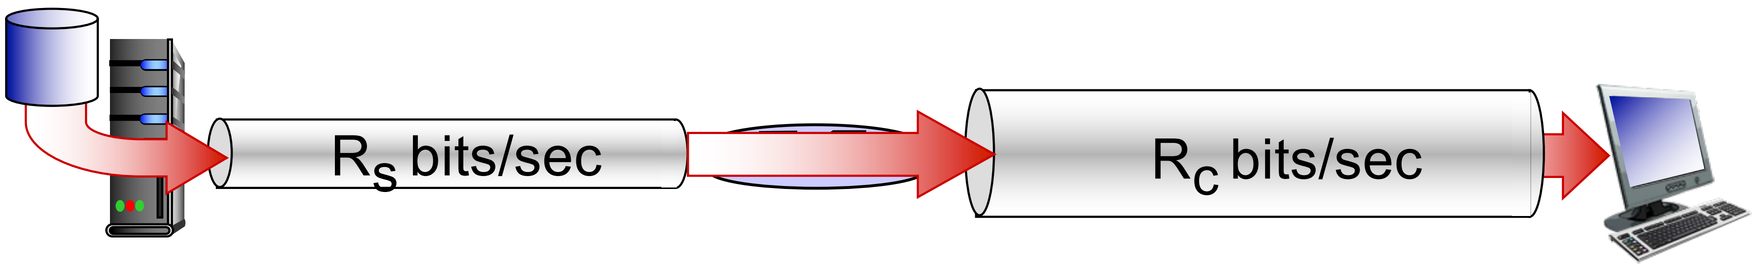

- $R_s > R_c$ What is average end-end throughput? R_c

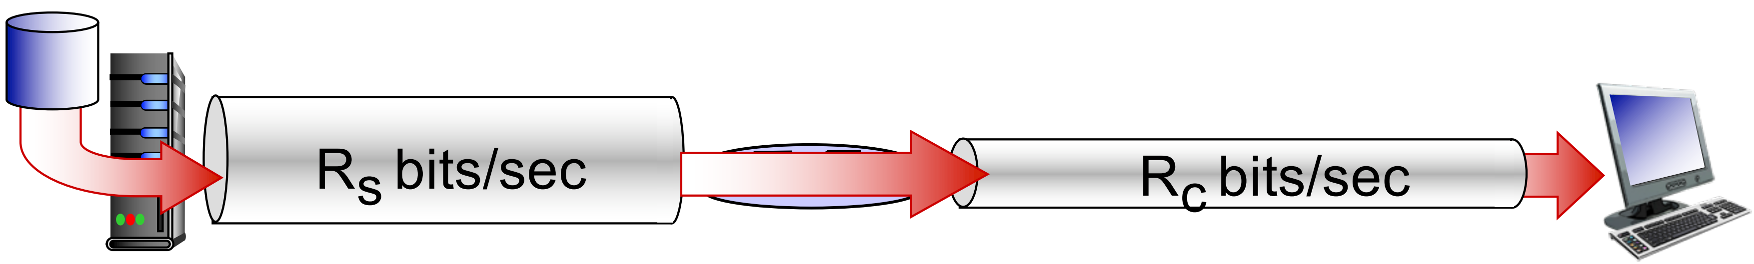

bottleneck link

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

Throughput

- per-connection end-end throughput:

$max(R_1,R_2,R_3)=R_3=3Mbps$

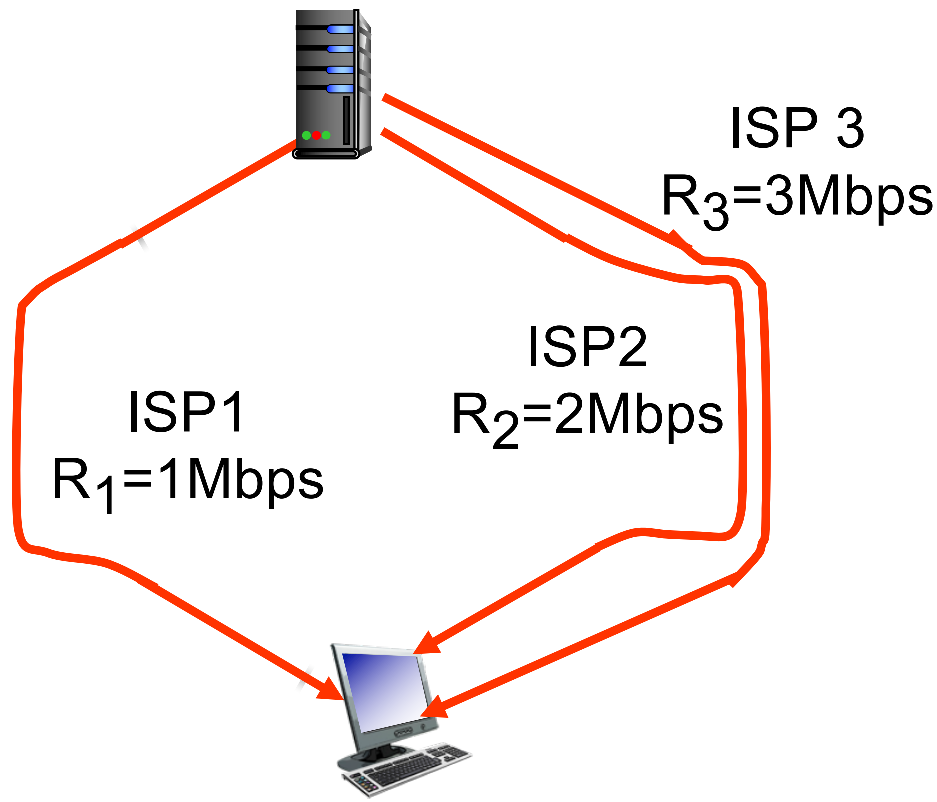

Throughput: Internet scenario

- per-connection end-end throughput:

$\min(R_c, R_s, R/10) = 500\text{kbps}$

- in practice: R_c or R_s is often bottleneck

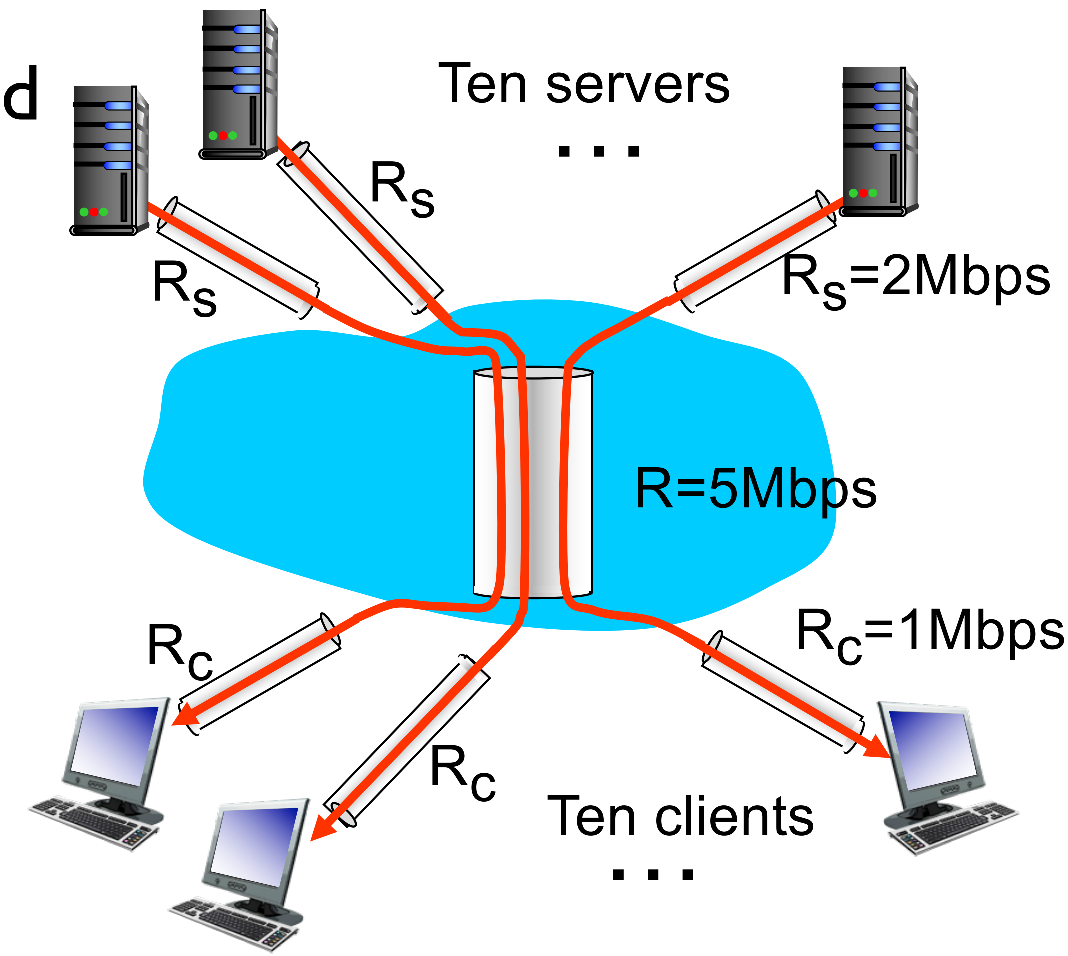

10 connections (fairly) share backbone bottleneck link R bits/sec