Chapter 7: Network Impairments

Magda El Zarki
Prof. of CS
Univ. of CA, Irvine
Email: elzarki@uci.edu
http://www.ics.uci.edu/~magda

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

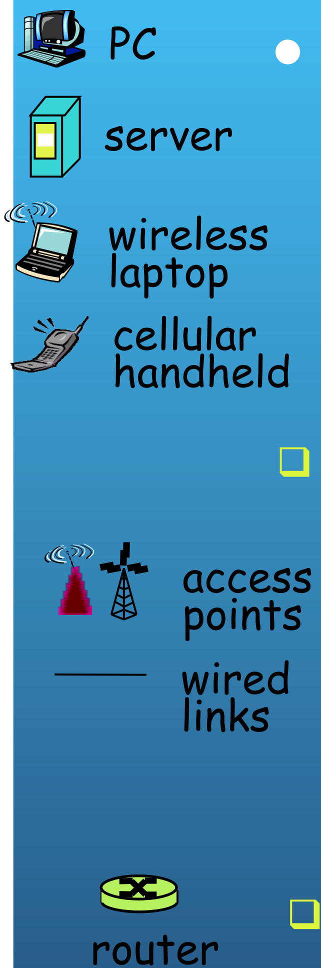

- millions of connected computing devices: *hosts = end systems*
  - running *network apps*
- *communication links*
  - fiber, copper, radio, satellite
  - transmission rate = *bandwidth*
- *routers:* forward packets (chunks of data)

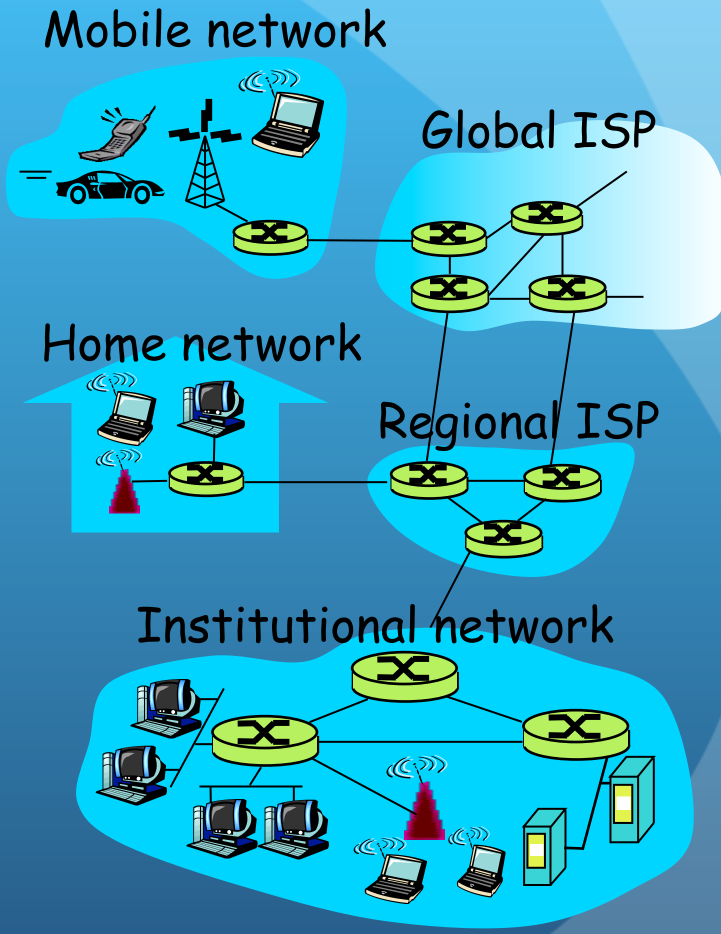

What's the Internet: a service view

- communication *infrastructure*
  - enables distributed applications:
  - Web, VoIP, email, games, e-commerce, file sharing
- communication services provided to apps:
  - reliable data delivery from source to destination
  - "best effort" (unreliable) data delivery

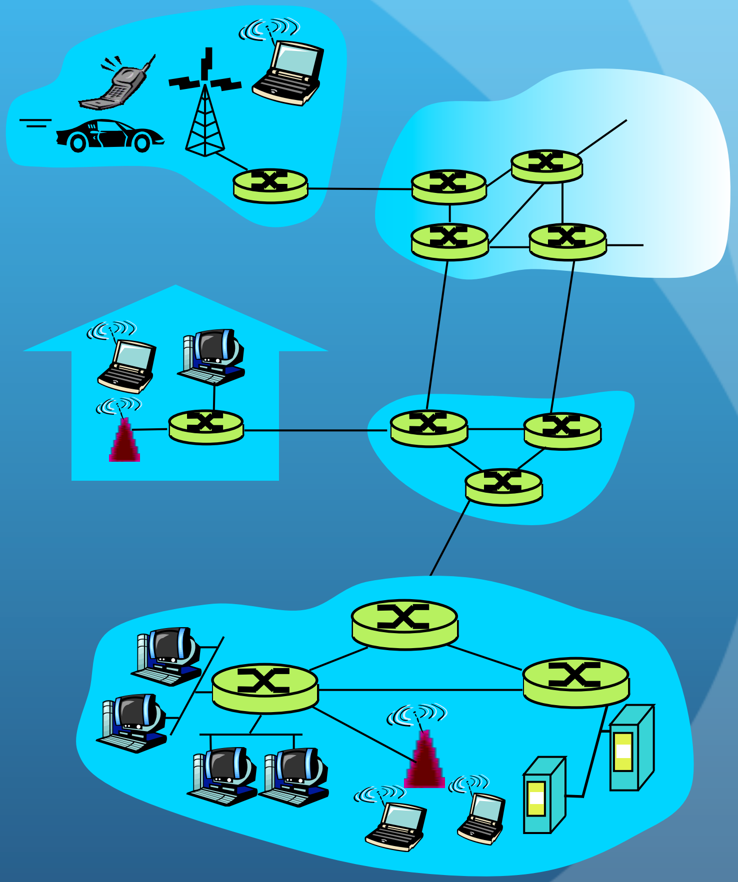

The Network Core

- mesh of interconnected routers
- *the* fundamental question: how is data transferred through net?
  - circuit switching: dedicated circuit per call: telephone net
  - packet-switching: data sent thru net in discrete "chunks"

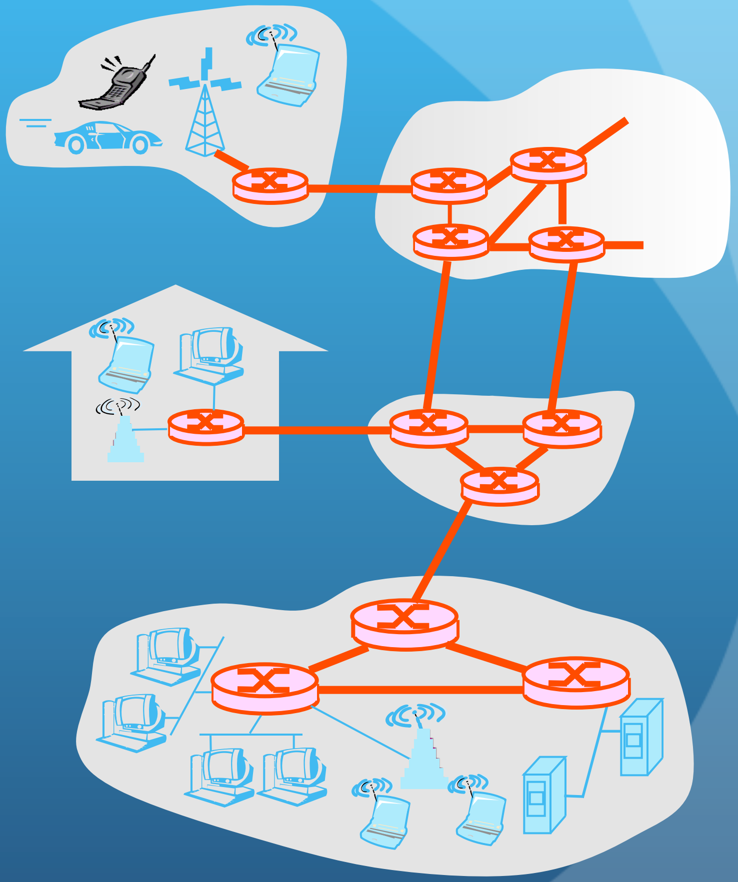

Network Core: Circuit Switching

- End-end resources reserved for "call"
- link bandwidth, switch capacity
- dedicated resources: no sharing
- circuit-like (guaranteed) performance
- call setup required

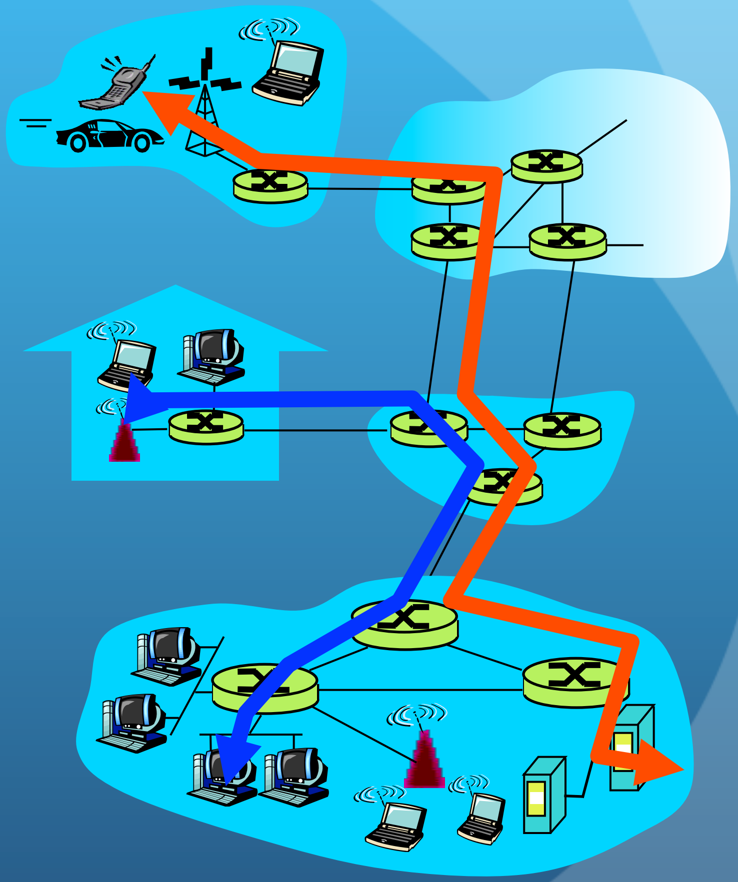

Network Core: Circuit Switching

- network resources (e.g., bandwidth) divided into "pieces"
  - dividing link bandwidth into "pieces"
    - frequency division
    - time division
- pieces allocated to calls
- resource piece *idle* if not used by owning call *(no sharing)*

Circuit Switching: FDM and TDM

Example:

4 users

FDM

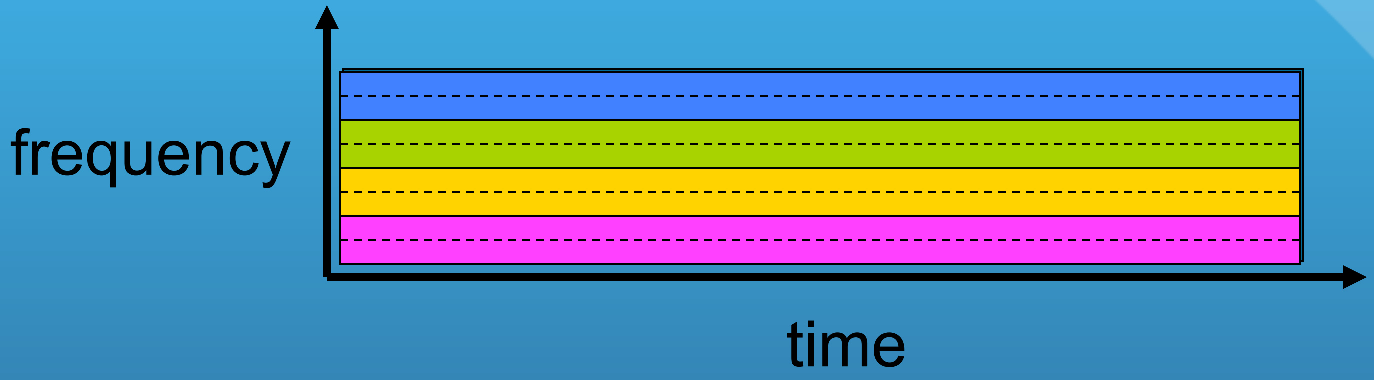

TDM

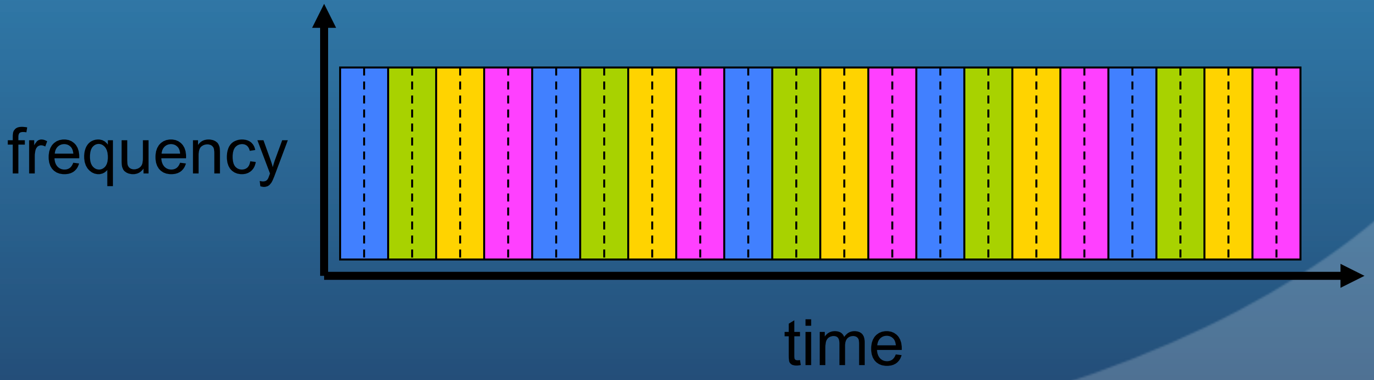

Network Core: Packet Switching

- each end-end data stream divided into *packets*
- user A, B packets *share* network resources
- each packet uses full link bandwidth
- resources used *as needed*

Packet Switching: Statistical Multiplexing

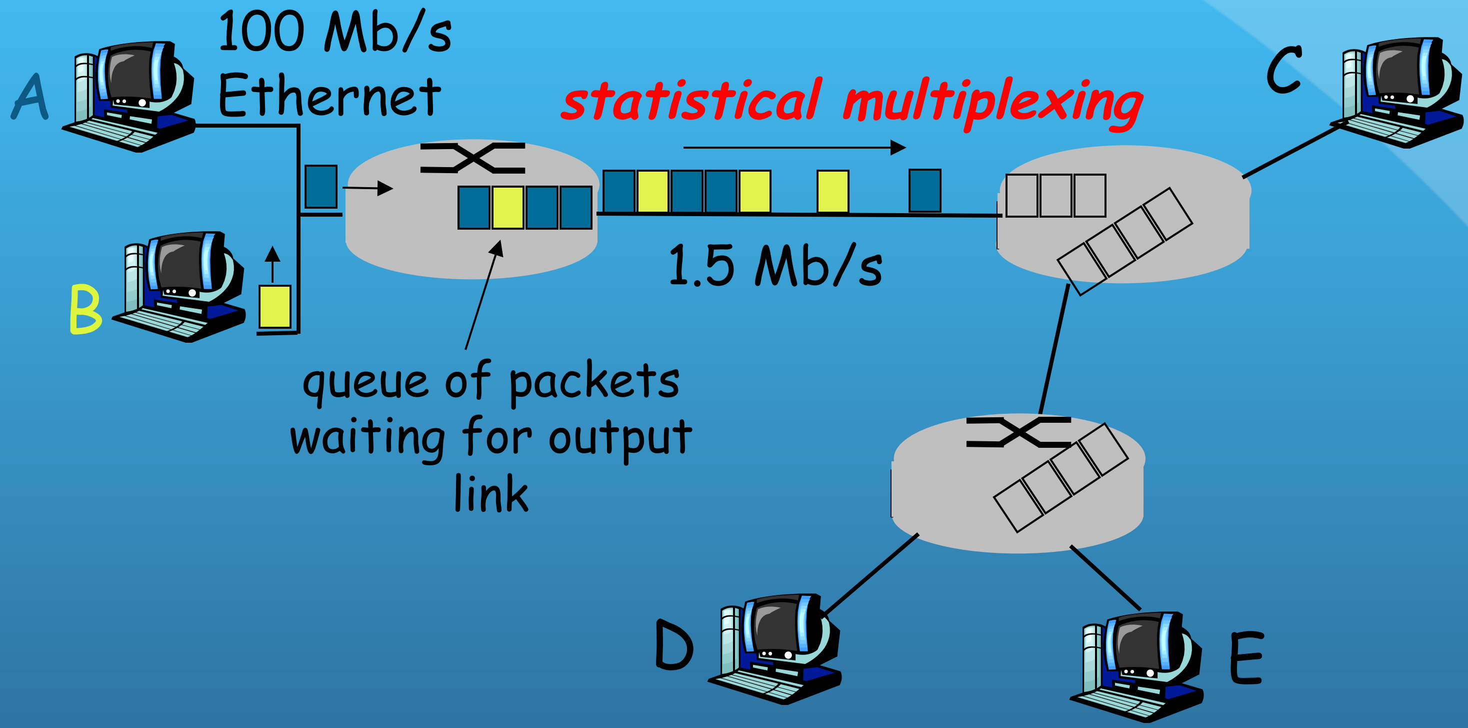

Sequence of A & B packets does not have fixed pattern, bandwidth shared on demand ➡ ***statistical multiplexing***.

TDM: each host gets same slot in revolving TDM frame.

Resource Contention:

- aggregate resource demand can exceed amount available
- congestion: packets queue, wait for link use
- store and forward: packets move one hop at a time
  - Node receives complete packet before forwarding

Packet-switching: store-and-forward

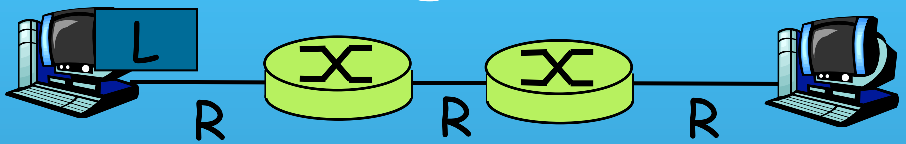

- takes L/R seconds to transmit (push out) packet of L bits on to link at R bps
- *store and forward:* entire packet must arrive at router before it can be transmitted on next link
- delay = 3L/R (assuming zero propagation delay)

Example:

- L = 7.5 Mbits
- R = 1.5 Mbps
- transmission delay = 15 sec

Packet switching versus circuit switching

Packet switching allows more users to use network!

- 1 Mb/s link
- each user:
  - 100 kb/s when "active"
  - active 10% of time

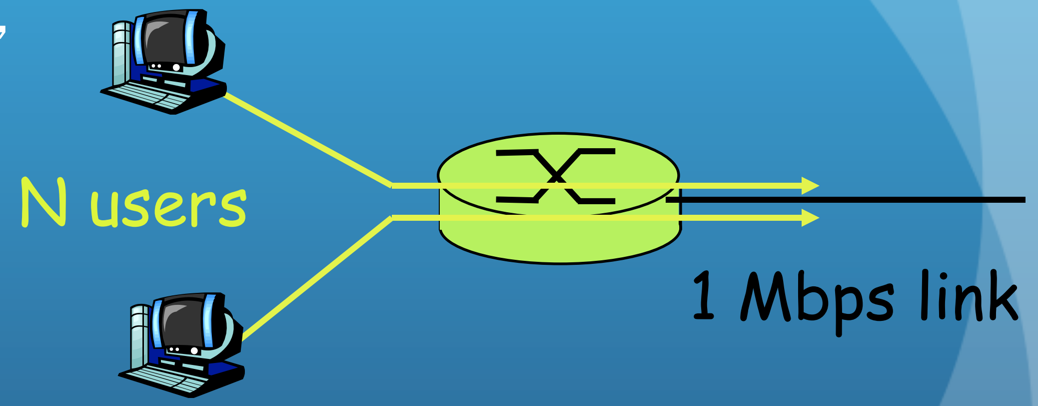

- *circuit-switching:*
  - 10 users
- *packet switching:*
  - with 35 users, probability > 10 active at same time is less than .0004

Packet switching versus circuit switching

Is packet switching a "slam dunk winner?"

- great for bursty data
  - resource sharing
  - simpler, no call setup
- excessive congestion:
  - packet delay (avg., max and jitter) and loss
  - protocols needed for reliable data transfer, congestion control
- Q: How to provide circuit-like behavior?
  - bandwidth guarantees needed for audio/video apps
  - still an unsolved problem

How do loss and delay (latency/lag) occur?

- packets *queue* in router buffers
  - packet arrival rate to link exceeds output link capacity
  - packets queue, wait for turn

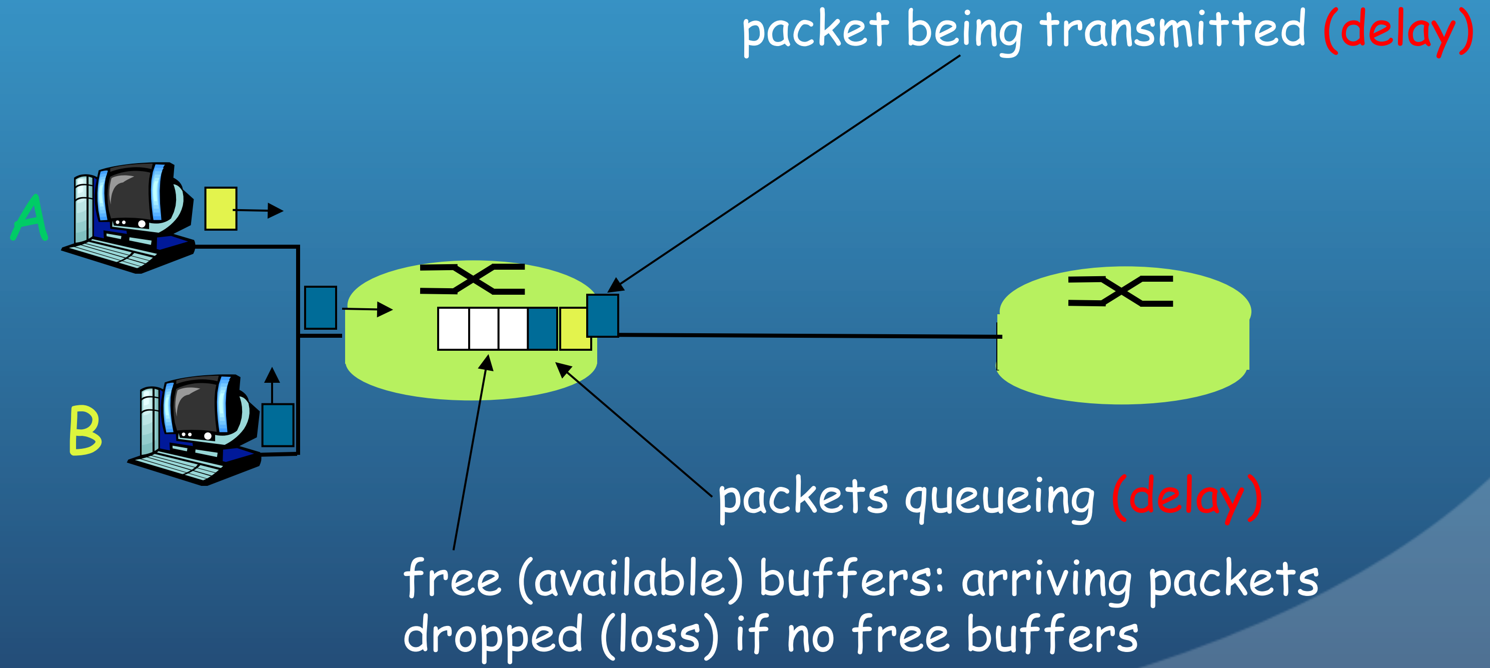

Four sources of packet delay

1. nodal processing:

- check bit errors
- determine output link

2. queueing:

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

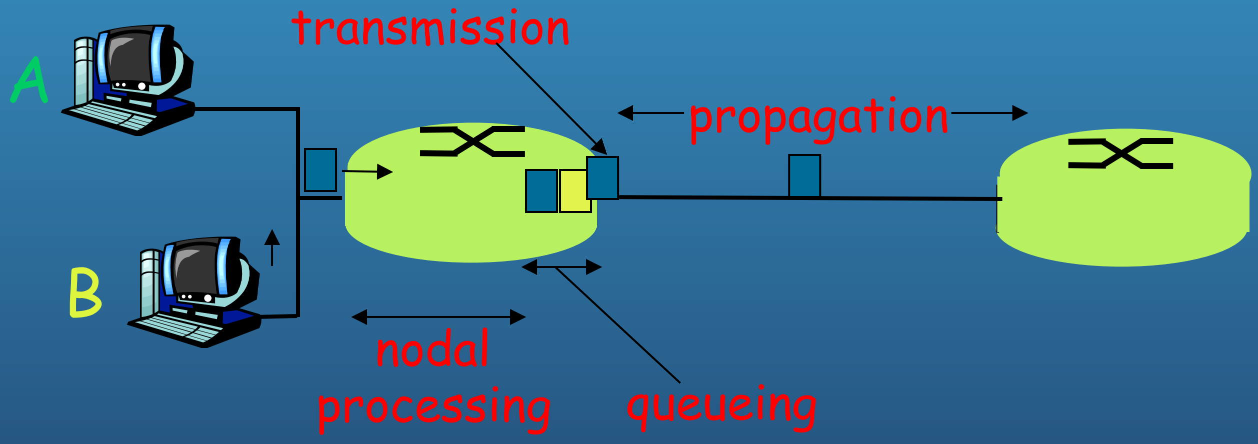

Delay in packet-switched networks

3. Transmission delay:

- R=link bandwidth (bps)
- L=packet length (bits)
- time to send bits into link = L/R

4. Propagation delay:

- d = length of physical link
- s = propagation speed in medium (~2x10^8 m/sec)
- propagation delay = d/s

Note: s and R are *very* different quantities!

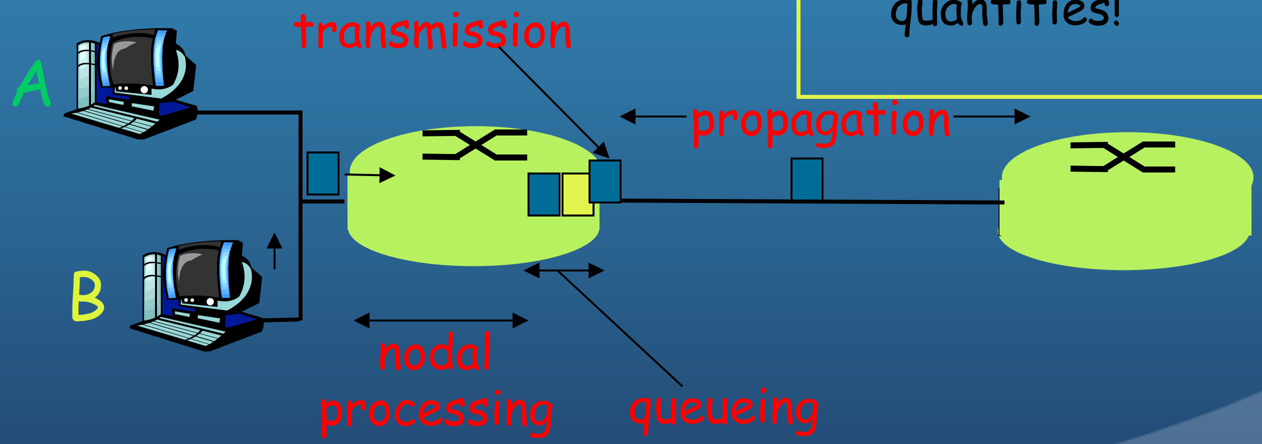

Caravan analogy

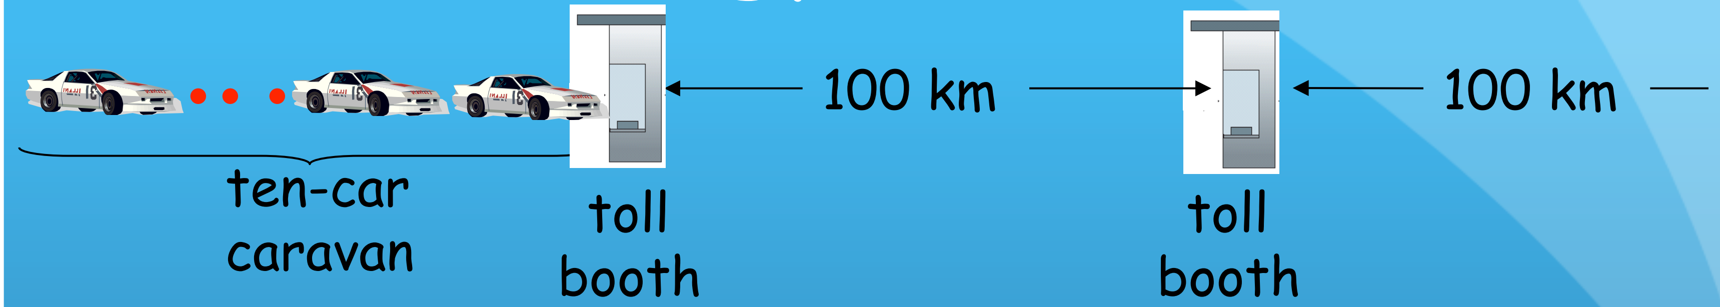

- cars "propagate" at 100 km/hr
- toll booth takes 12 sec to service car (transmission time)
- car~bit; caravan ~ packet
- Q: How long until caravan is lined up before 2nd toll booth?

- 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 (more)

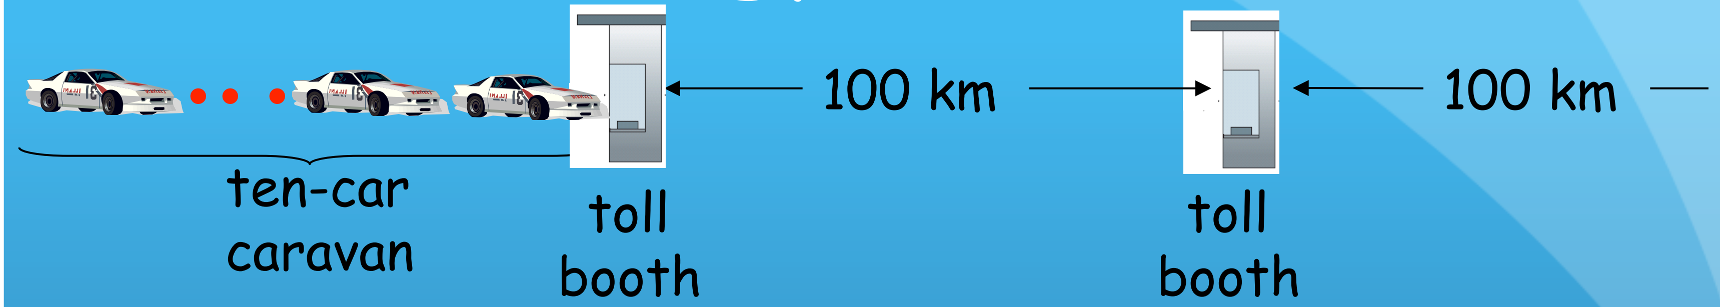

- Cars now "propagate" at 1000 km/hr
- Toll booth now takes 1 min to service a car
- Q: Will cars arrive to 2nd booth before all cars serviced at 1st booth?

- Yes!
  - After 7 min, 1st car at 2nd booth and 3 cars still at 1st booth. (takes cars 6mins to cross dist. bet. 2 booths)
- 1st bit of packet can arrive at 2nd router before packet is fully transmitted at 1st router!

Total delay

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

- d_{proc} = processing delay
  - typically a few microsecs or less
- d_{queue} = queuing delay
  - depends on congestion
- d_{trans} = transmission delay
  - = L/R, significant for low-speed links
- d_{prop} = propagation delay
  - a few microsecs to hundreds of msecs

Queueing delay (revisited)

- R=link bandwidth (bps)
- L=packet length (bits)
- a=average packet arrival rate

traffic intensity = La/R

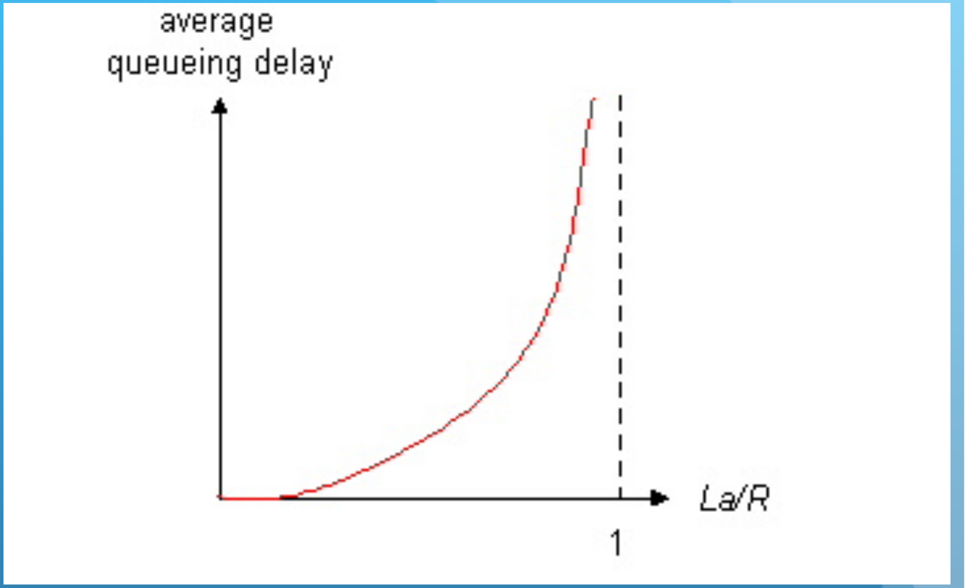

- La/R ~ 0: average queueing delay small
- La/R -> 1: delays become large
- La/R > 1: more "work" arriving than can be serviced, average delay infinite!

"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
  - router *i* will return packets to sender
  - sender times interval between transmission and reply.

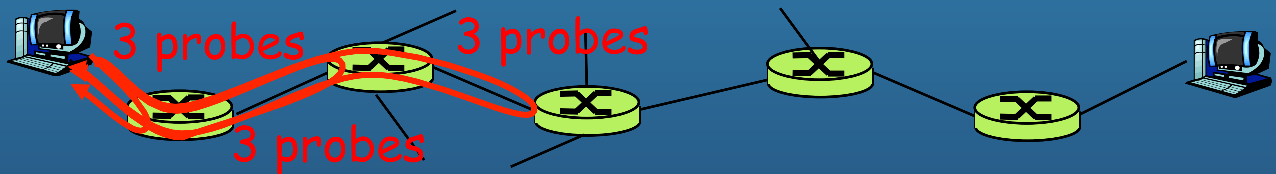

Real Internet delays and routes

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

Three delay measurements from gaia.cs.umass.edu to cs-gw.cs.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 (hop 8)

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

Packet Jitter

- Variation in packet delay
- Causes
  - Variation in packet lengths -> different transmission times
  - Variation in path lengths -> no fixed paths in the Internet

Difference: Jitter and Latency

Latency and Jitter affect streams of packets travelling across the network

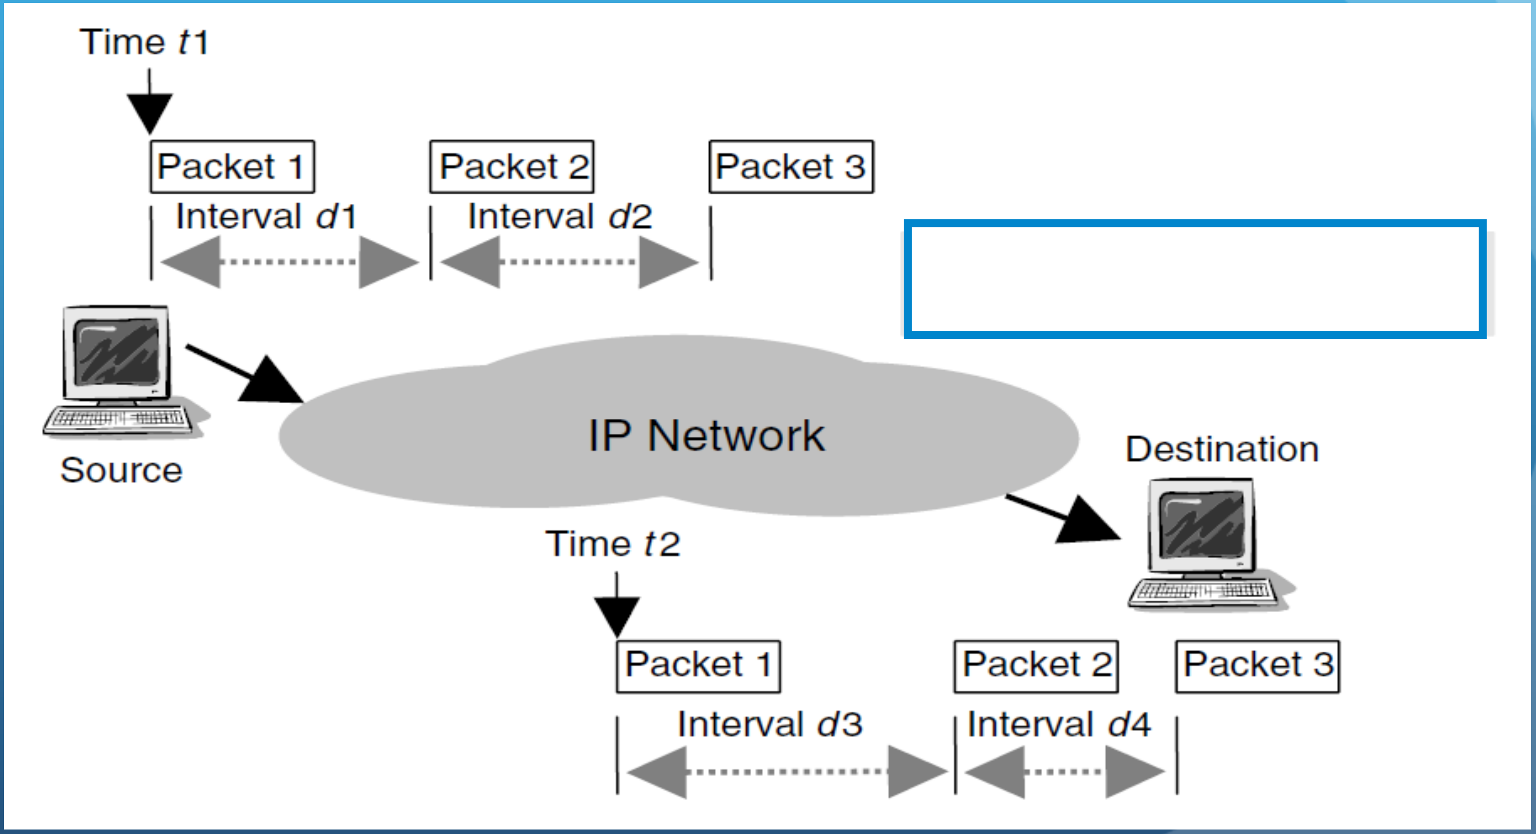

Packet loss

- queue (aka buffer) preceding link 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

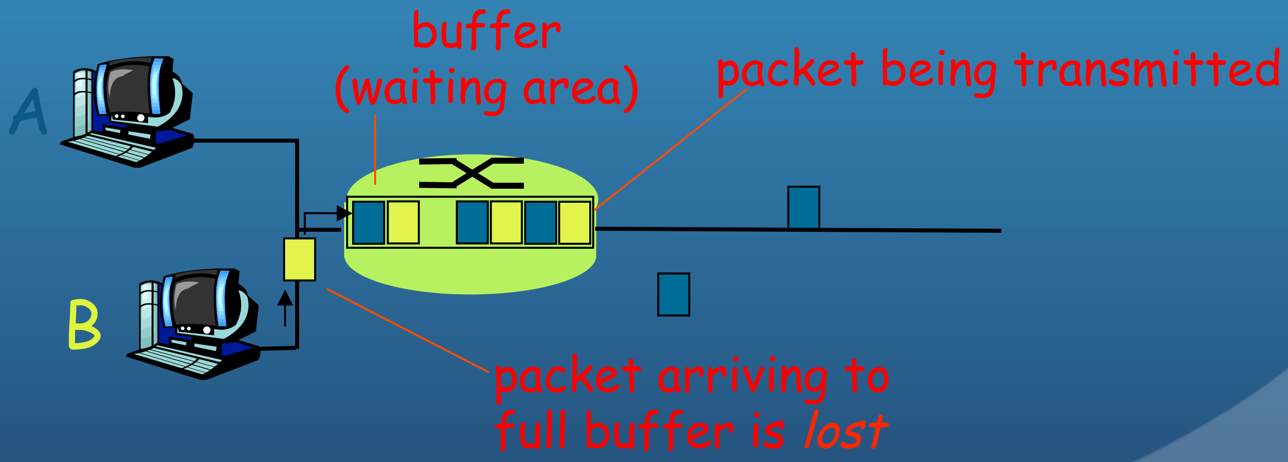

Throughput

- *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

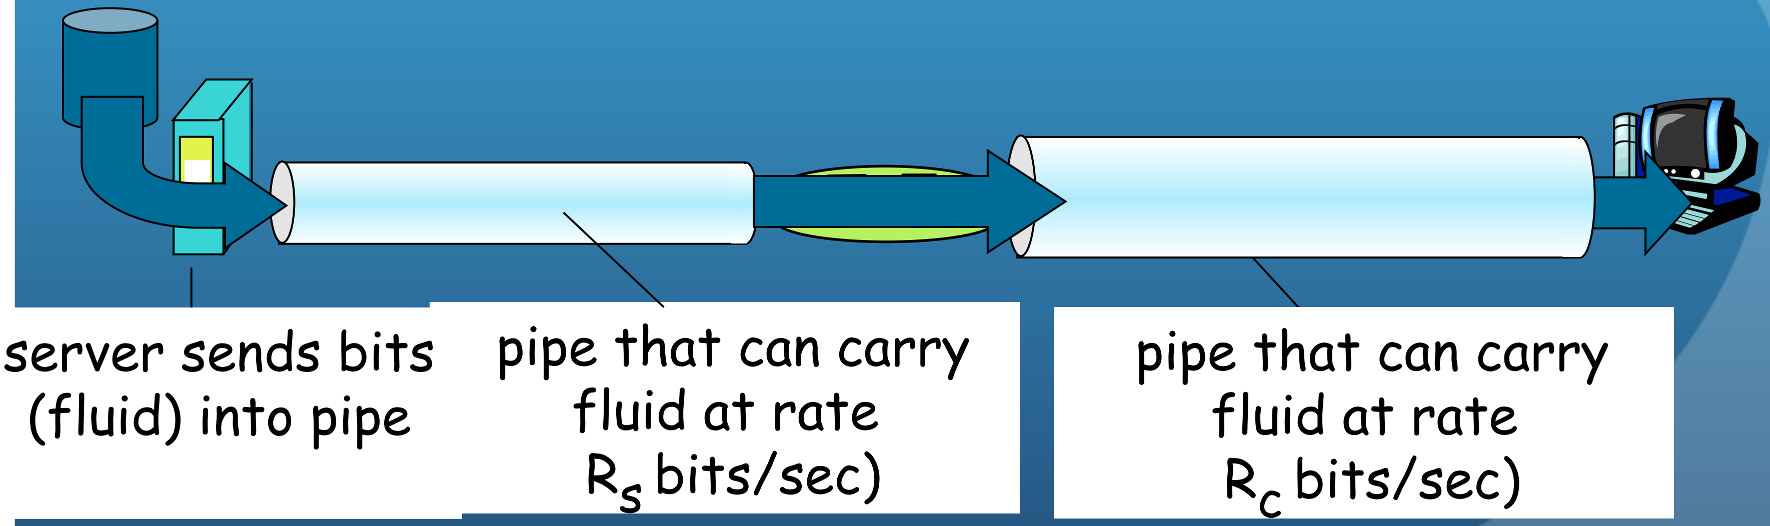

Throughput (more)

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

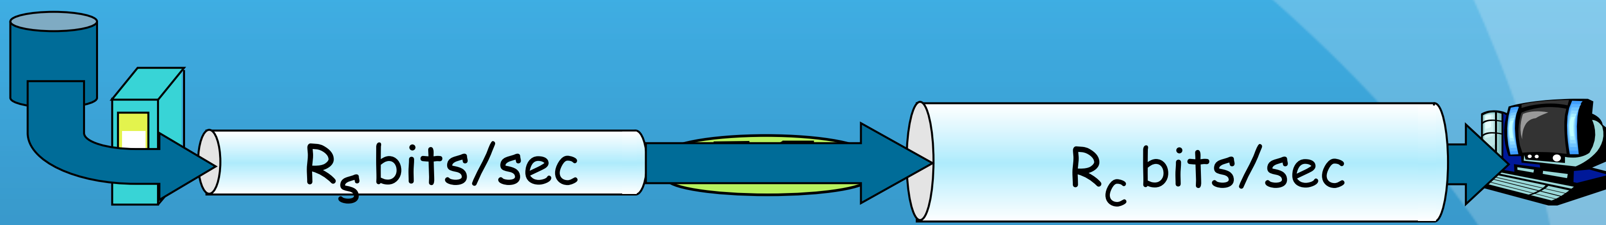

- $R_s > R_c$ What is average end-end throughput?

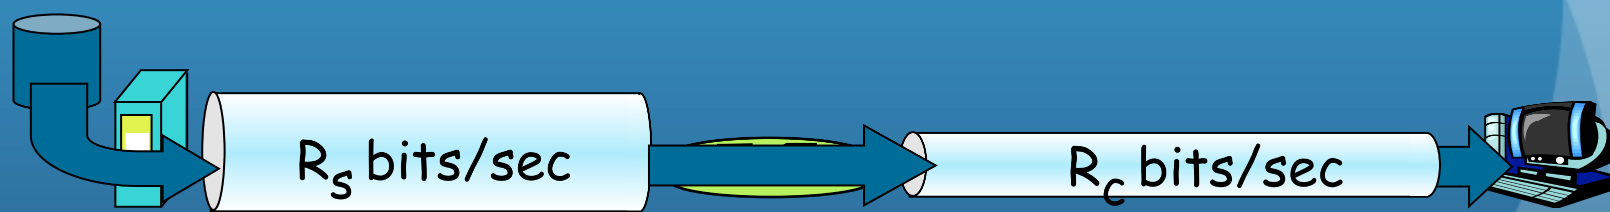

bottleneck link

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

Network Control

- Ensure that the network impairments (delay, jitter and packet losses) do not exceed a certain value
  - Configure/size of links, routers and buffers to avoid congestion
  - Prioritize traffic with hard limits on delay, jitter and losses
    - By source/destination IP address pairs
    - By protocol type
    - By source/destination port numbers

Preferential Queuing

Preferential queuing and scheduling can allow priority packets to ‘jump the queue’

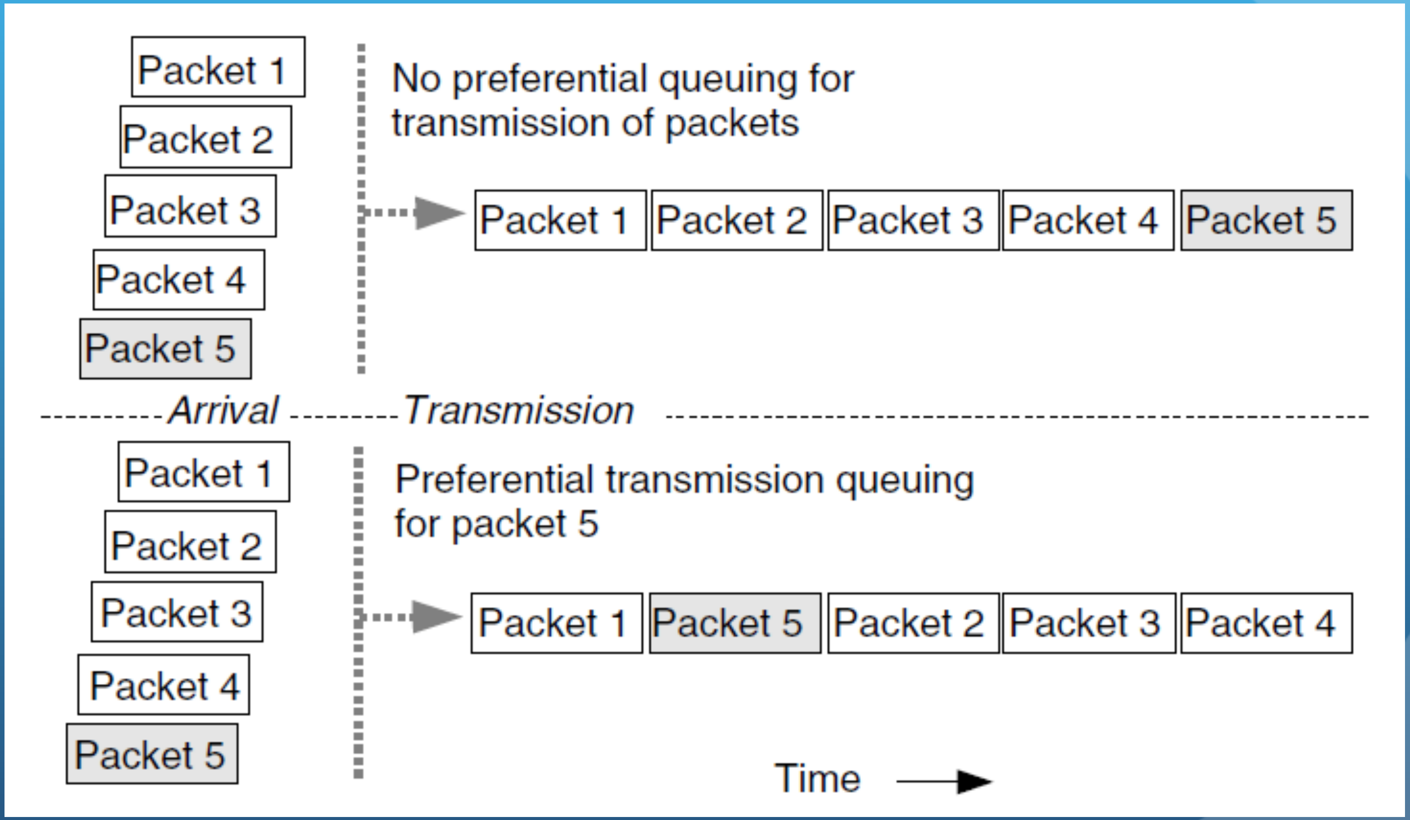

Link Layer Support for Packet Prioritisation

- Classification, queuing and scheduling at the IP packet layer does not solve all problems when facing serialisation delays on low-speed links

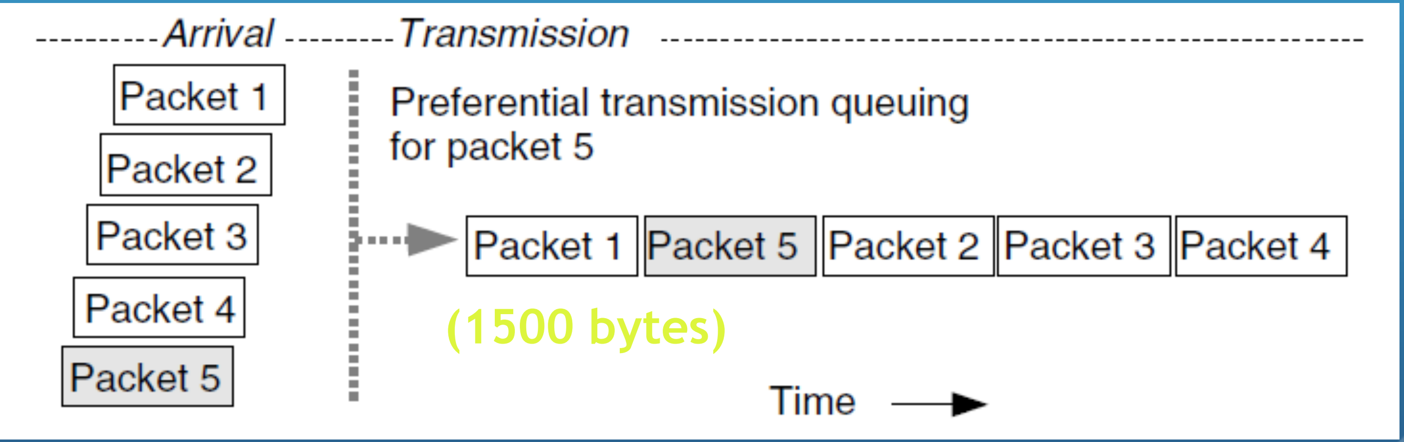

Packet 5 having to wait until packet 1 is transmitted if packet had begun transmission just before packet 5 arrived!

Where to Place and Trust Traffic Classification

- How do ISPs know what packets to give priority to at any given time? (Which packets are game traffic derserving perferential treatment?)
  - Only a game client and game server know what IP packets constitute game traffic
- The game client might use a signaling protocol to inform the ISP of the 5-tuple associated with game traffic when a new flow of game packets begins?

Where to Place and Trust Traffic Classification

- An emerging approach is for ISPs to automatically detect game traffic by looking for particular statistical properties rather than specific 5-tuple values or well-known port numbers
  - Once a flow has been identified as game traffic, the flow's 5-tuple can be passed out to routers along the path who need to provide preferential treatment