Data and Computer Communications

Eighth Edition

by William Stallings

Chapter 13 – Congestion in Data Networks

Lecture slides by Lawrie Brown

Partly modified/translated by J. Porras

What Is Congestion?

- congestion occurs when the no of packets being transmitted through the network approaches the packet handling capacity of the network
- congestion control aims to keep no of packets below a level at which performance falls off dramatically
- a data network is a network of queues
- generally 80% utilization is critical
- finite queues mean data may be lost

Queues at a Node

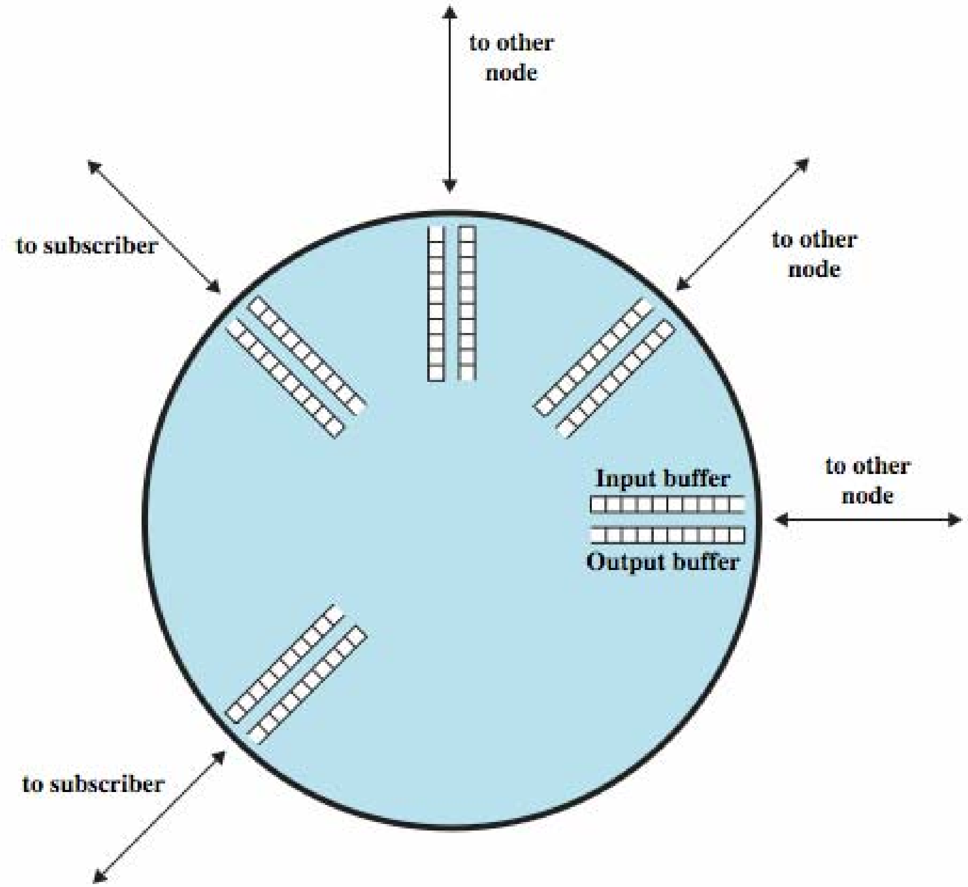

Interaction of Queues

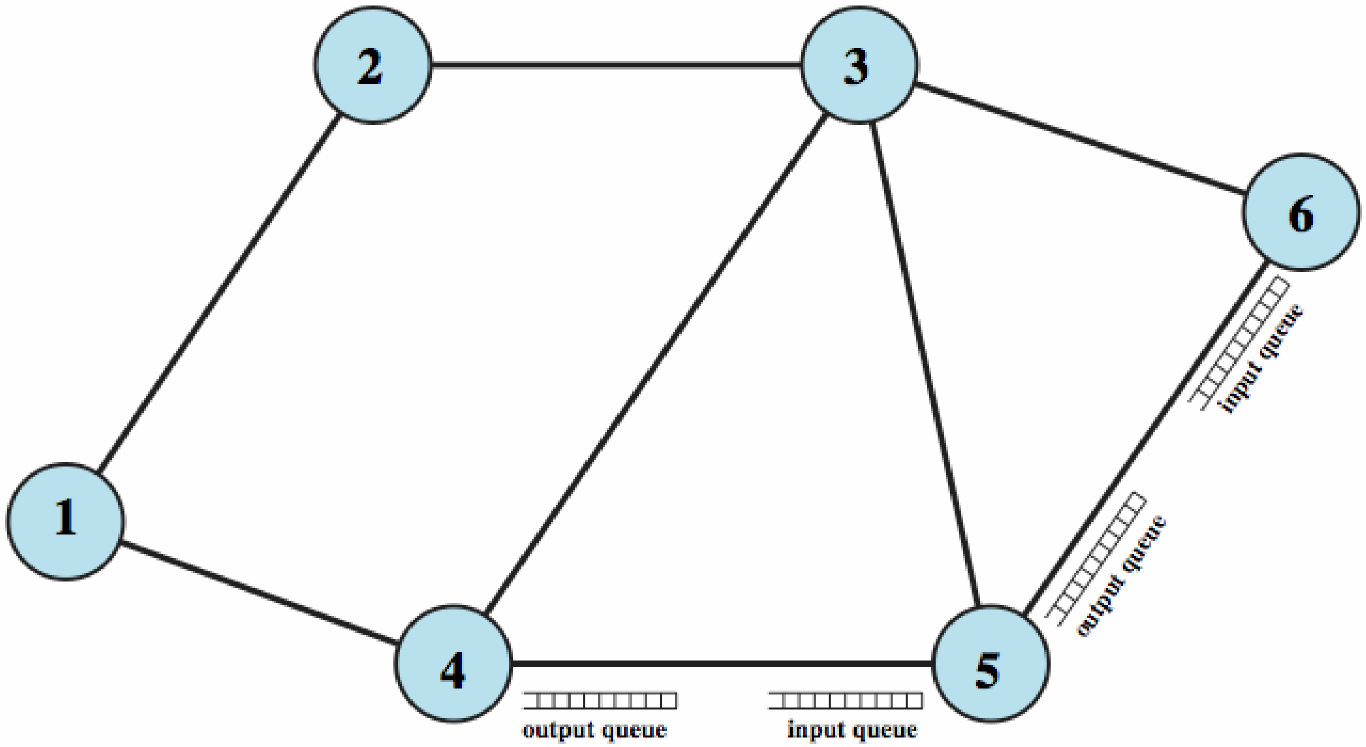

Ideal Network Utilization

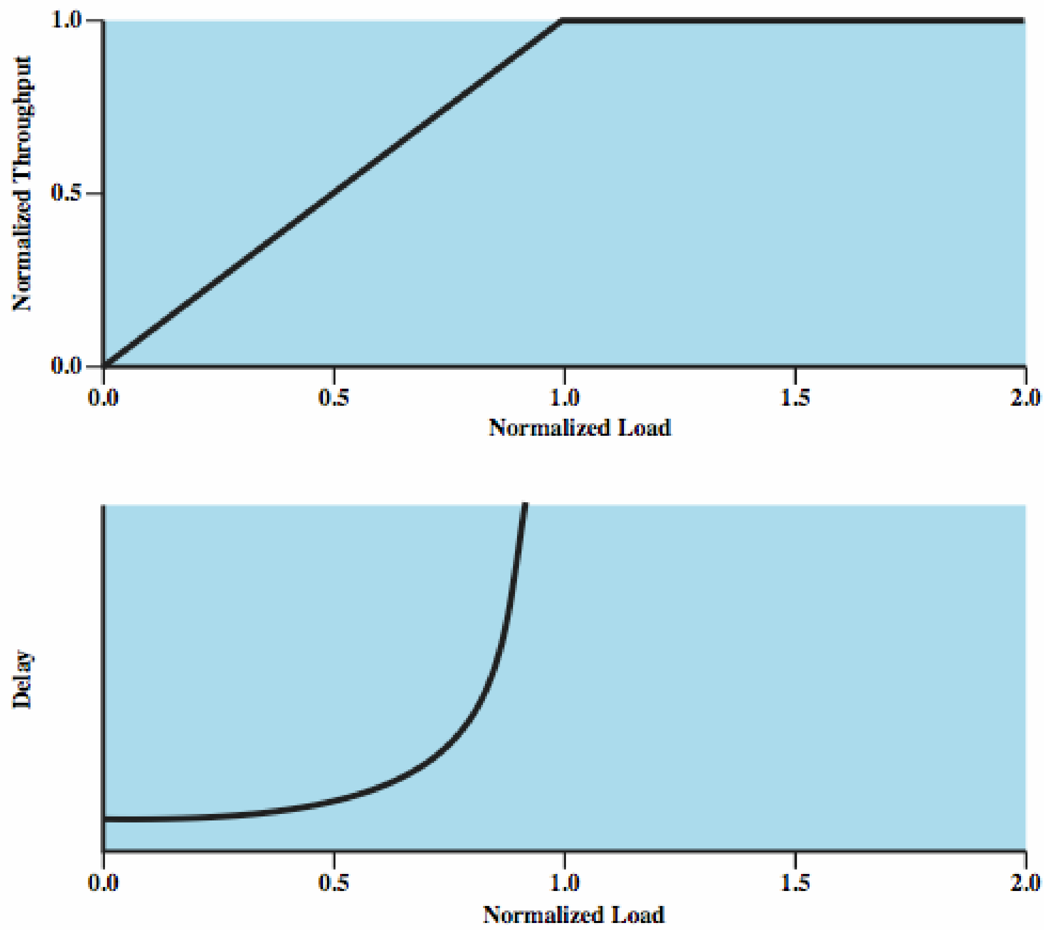

Effects of Congestion - No Control

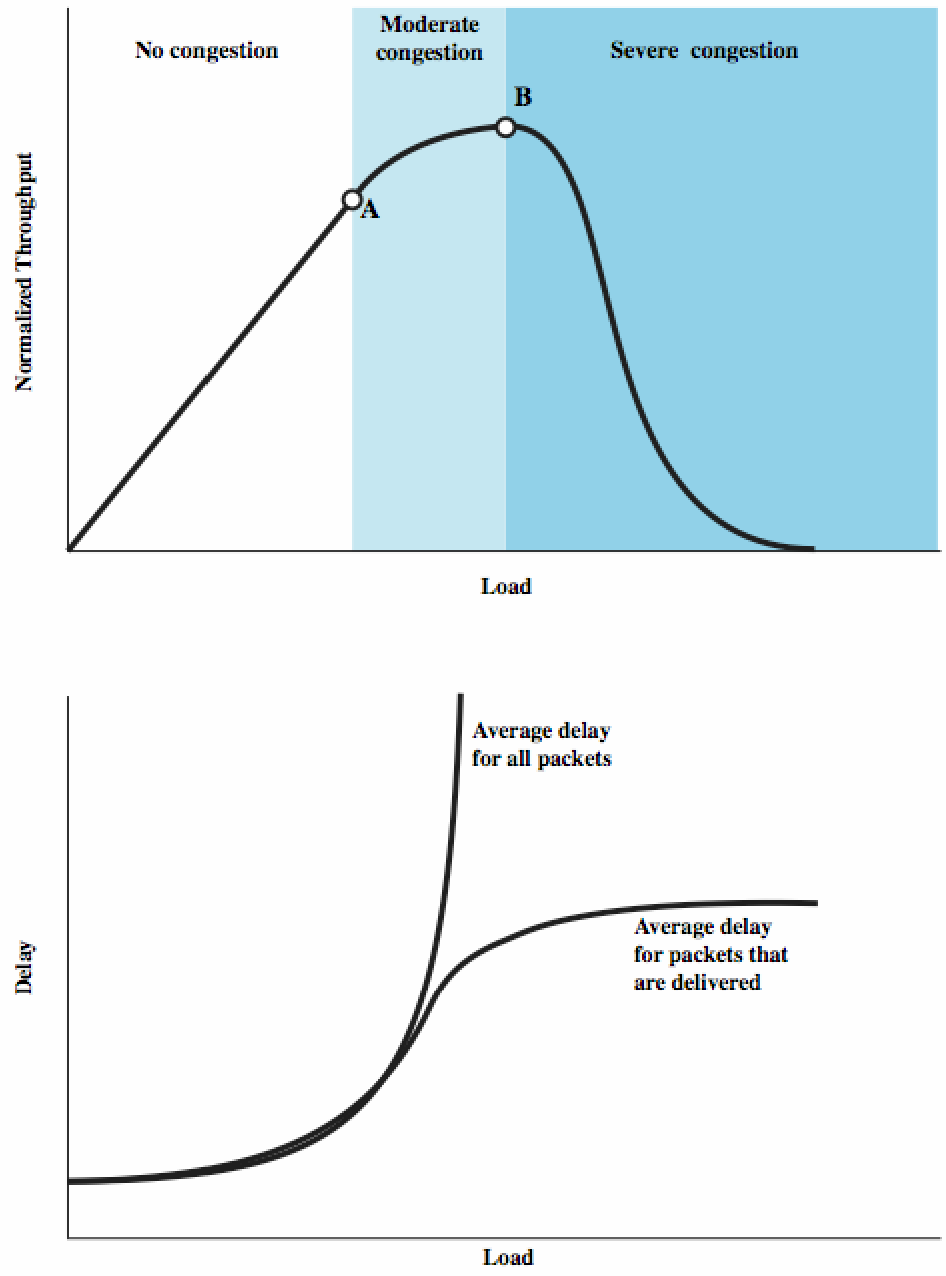

Mechanisms for Congestion Control

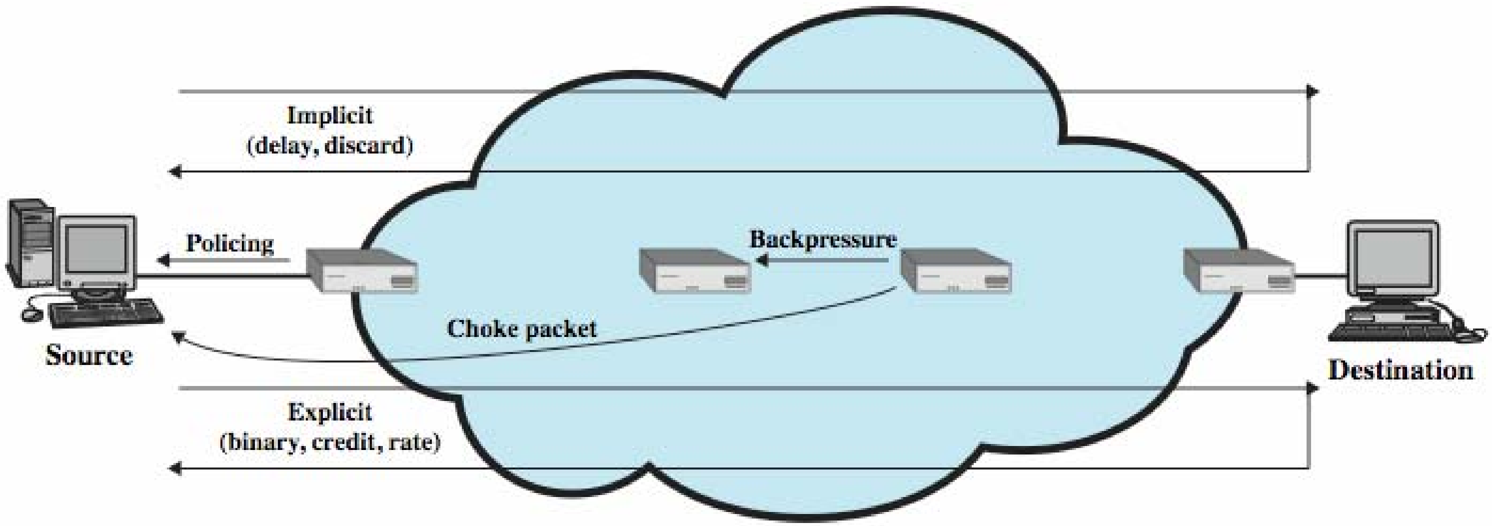

Backpressure

- if node becomes congested it can slow down or halt flow of packets from other nodes
  - cf. backpressure in blocked fluid pipe
  - may mean that other nodes have to apply control on incoming packet rates
  - propagates back to source
- can restrict to high traffic logical connections
- used in connection oriented nets that allow hop by hop congestion control (eg. X.25)
- not used in ATM nor frame relay
- only recently developed for IP

Choke Packet

- a control packet
  - generated at congested node
  - sent to source node
  - eg. ICMP source quench
    - from router or destination
    - source cuts back until no more source quench message
    - sent for every discarded packet, or anticipated
- is a rather crude mechanism

Implicit Congestion Signaling

- transmission delay increases with congestion
- hence a packet may be discarded
- source detects this implicit congestion indication
- useful on connectionless (datagram) networks
  - eg. IP based
    - (TCP includes congestion and flow control - see chapter 17)
- used in frame relay LAPF

Explicit Congestion Signaling

- network alerts end systems of increasing congestion
- end systems take steps to reduce offered load
- Backwards
  - congestion avoidance notification in opposite direction to packet required
- Forwards
  - congestion avoidance notification in same direction as packet required

Explicit Signaling Categories

- Binary
  - a bit set in a packet indicates congestion
- Credit based
  - indicates how many packets source may send
  - common for end to end flow control
- Rate based
  - supply explicit data rate limit
  - nodes along path may request rate reduction
  - eg. ATM

Traffic Management

- fairness
  - provide equal treatment of various flows
- quality of service
  - different treatment for different connections
- reservations
  - traffic contract between user and network
  - carry best-effort or discard excess traffic

Congestion Control in Packet Switched Networks

- send control packet to some or all source nodes
  - requires additional traffic during congestion
- rely on routing information
  - may react too quickly
- end to end probe packets
  - adds to overhead
- add congestion info to packets in transit
  - either backwards or forwards

How do loss and delay occur?

packets *queue* in router buffers

- packet arrival rate to link exceeds output link capacity
- packets queue, wait for turn

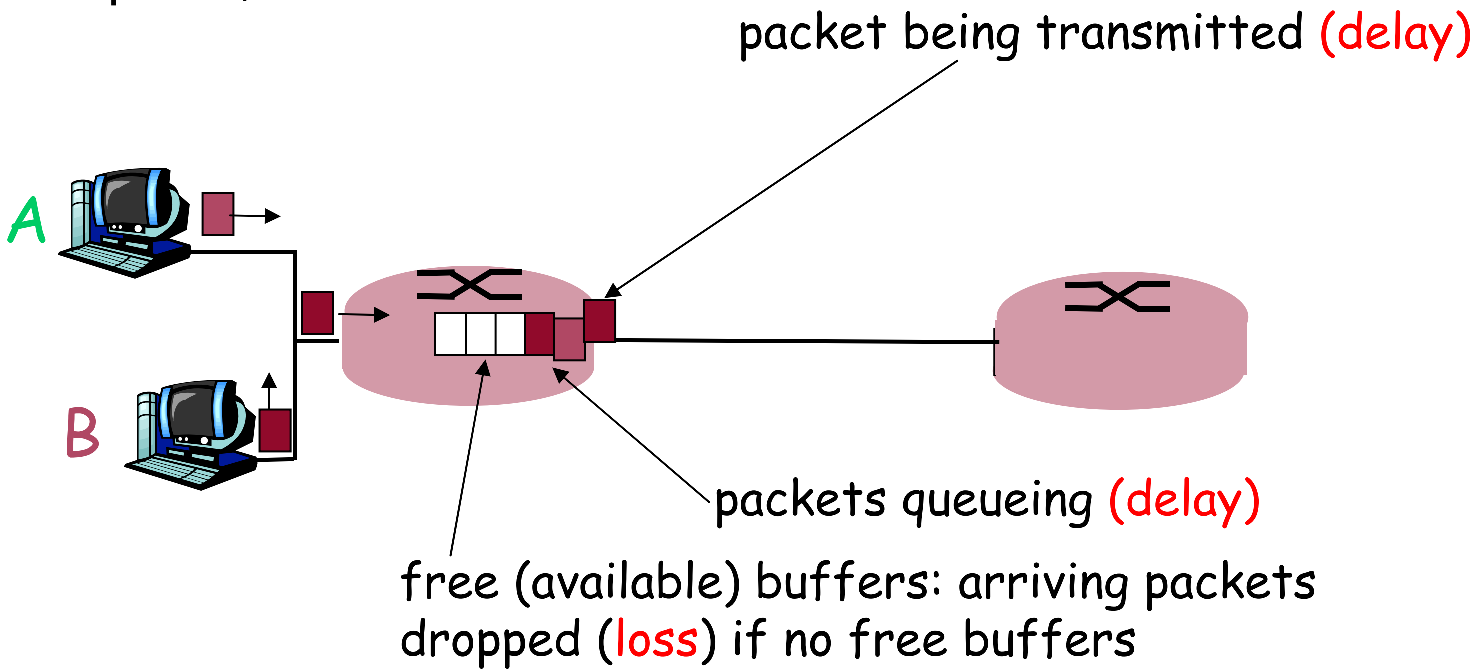

Four sources of packet delay

- 1. nodal processing:
  - check bit errors
  - determine output link
- 2. queueing
  - time waiting at output link for transmission
  - depends on congestion level of router

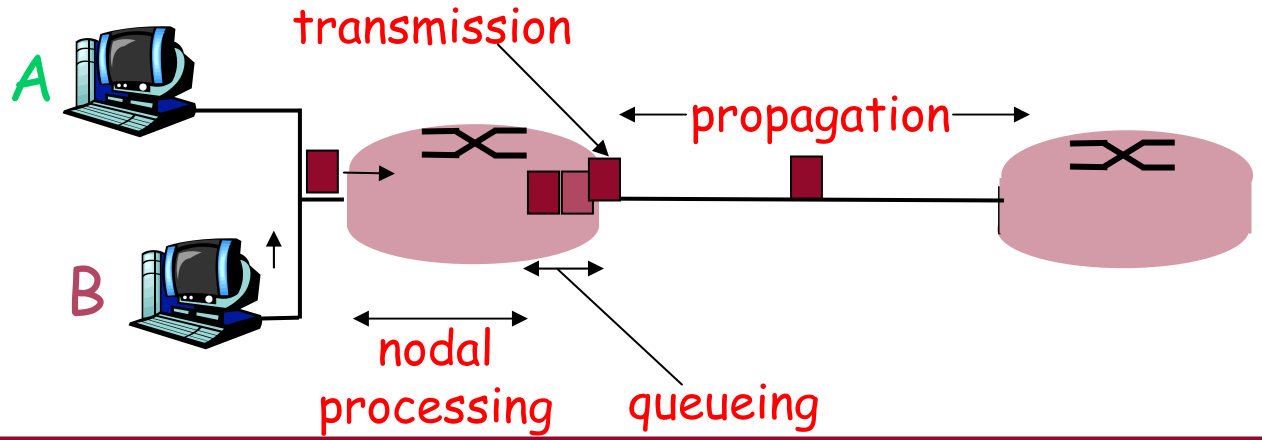

Delay in packet-switched networks

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 (~2×10^8 m/sec)
- propagation delay = d/s

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

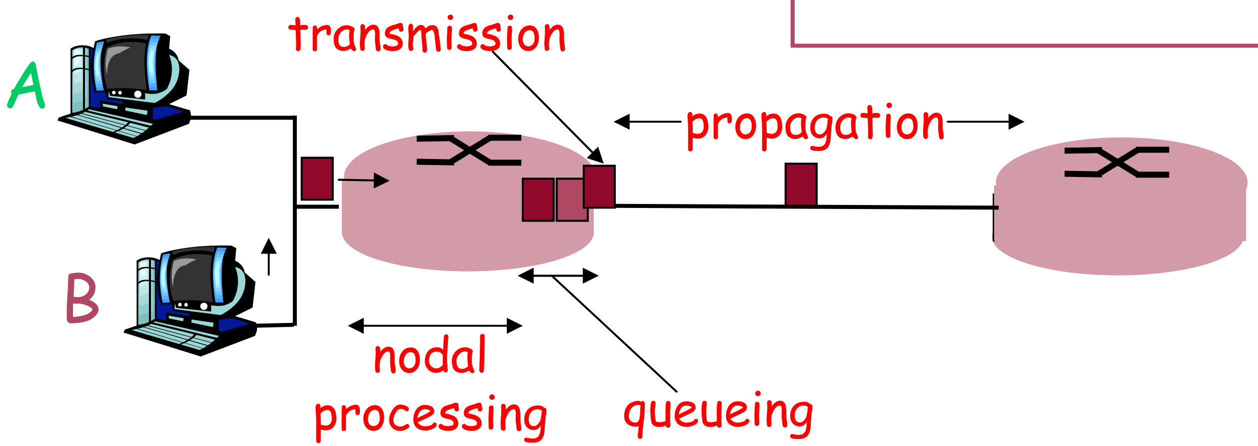

Caravan analogy

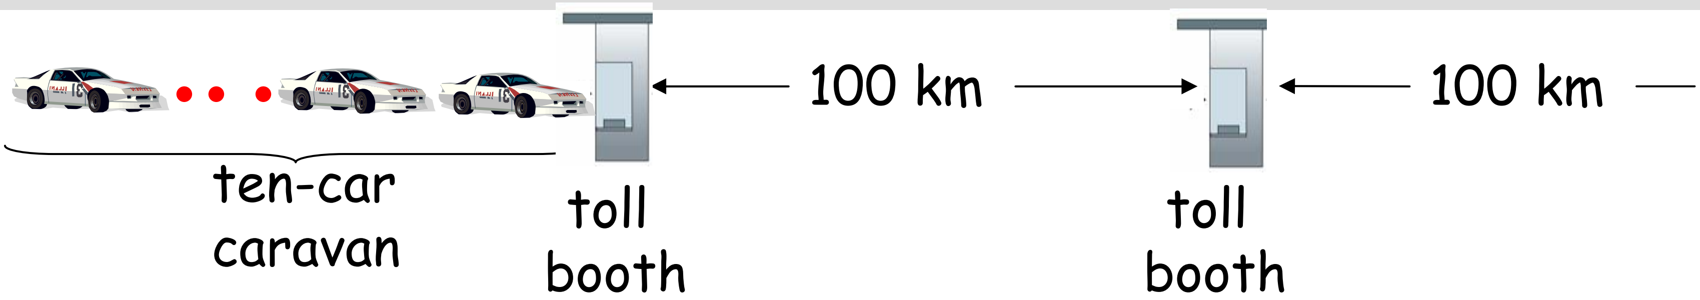

- cars "propagate" at 100 km/hr
- toll booth takes 12 sec to service car (transmission time)
- car~bit; caravan ~ packet
- Q: How long until caravan is lined up before 2nd toll booth?

- 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)

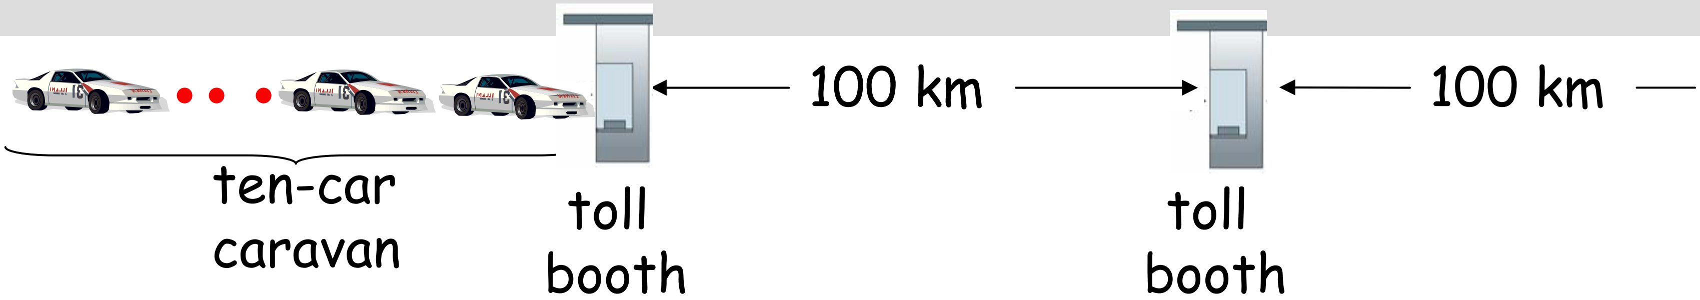

- Cars now "propagate" at 1000 km/hr
- Toll booth now takes 1 min to service a car
- Q: Will cars arrive to 2nd booth before all cars serviced at 1st booth?

- Yes! After 7 min, 1st car at 2nd booth and 3 cars still at 1st booth.
- 1st bit of packet can arrive at 2nd router before packet is fully transmitted at 1st router!
  - See Ethernet applet at AWL Web site

Nodal 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

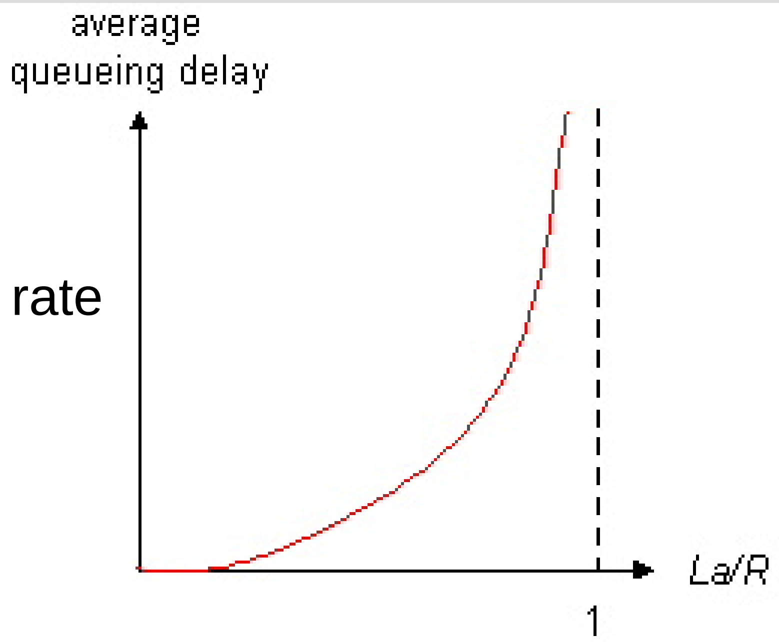

- La/R ~ 0: average queueing delay small
- La/R -> 1: delays become large
- La/R > 1: more "work" arriving than can be serviced, average delay infinite!

“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.

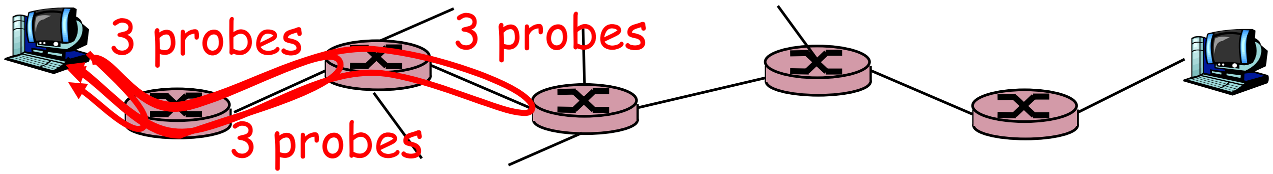

“Real” Internet delays and routes

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

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

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

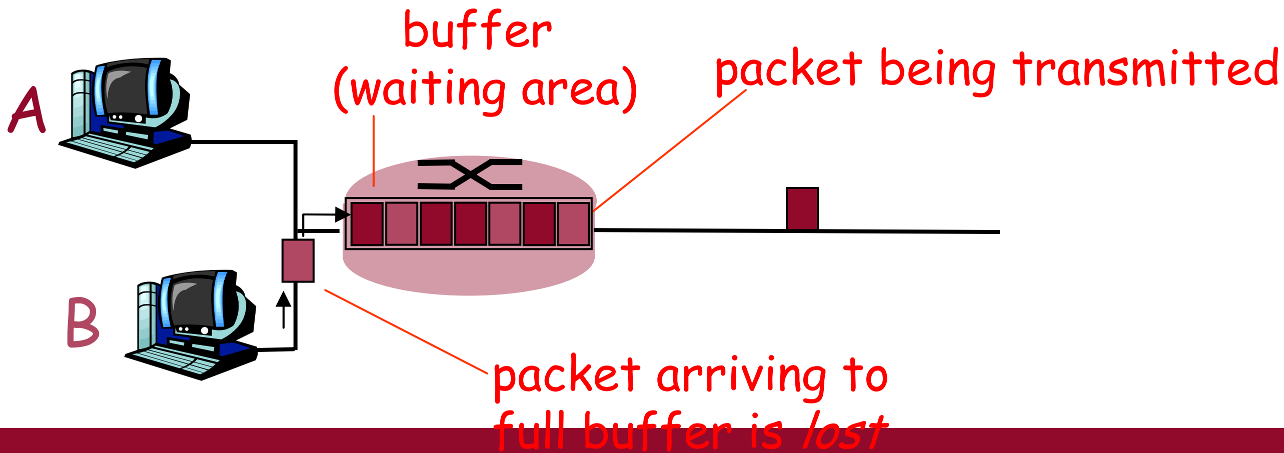

Throughput

- *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

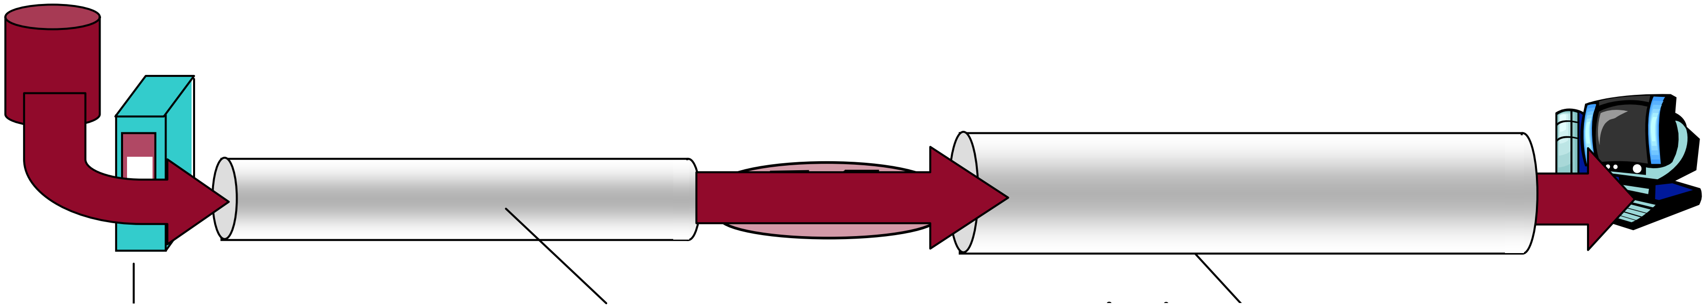

Throughput (more)

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

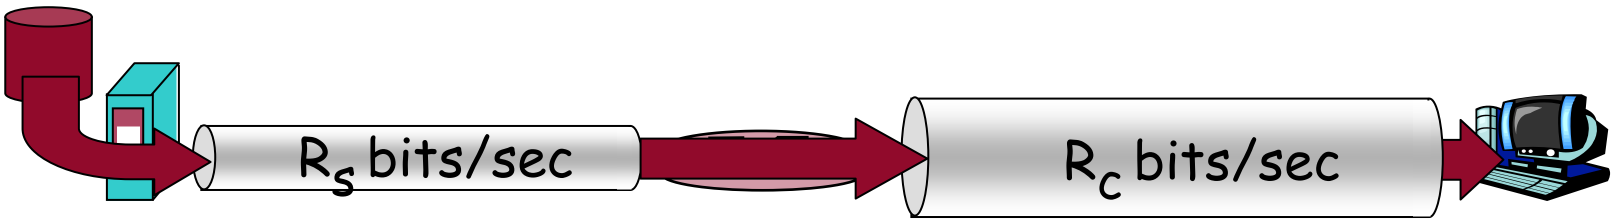

- $R_s > R_c$ What is average end-end throughput?

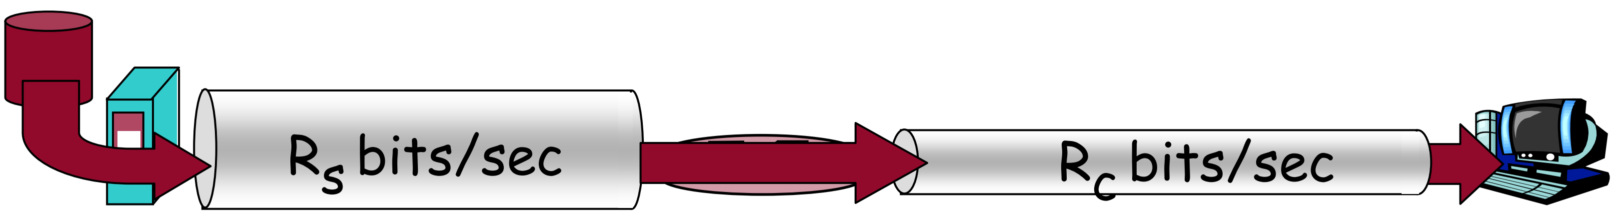

bottleneck link

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

Throughput: Internet scenario

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

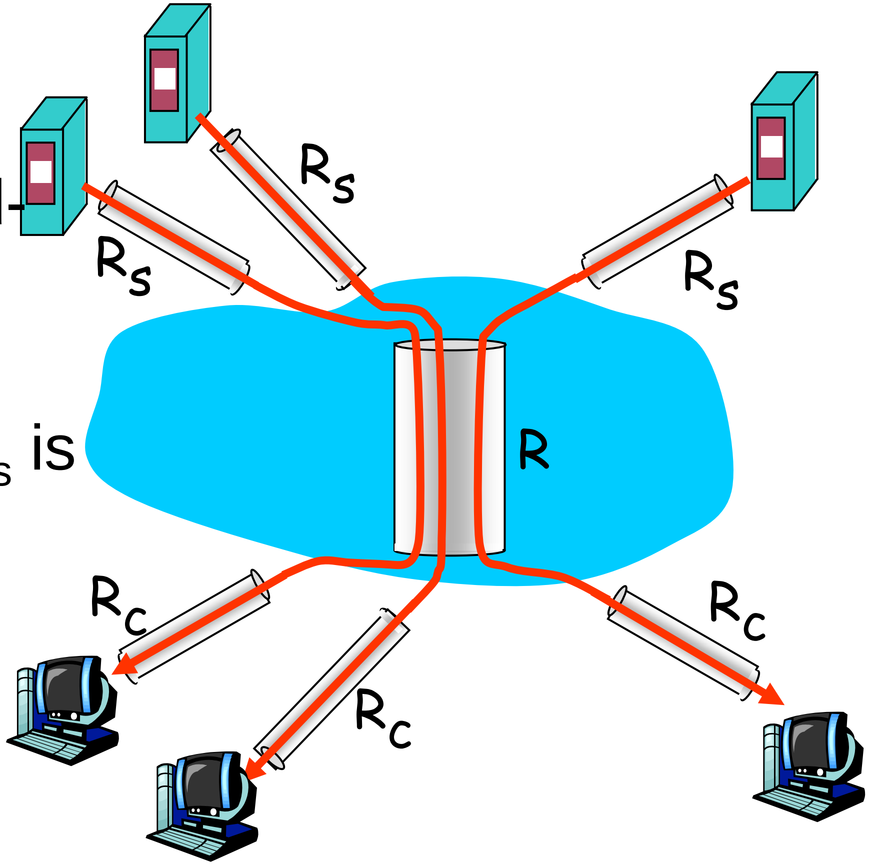

10 connections (fairly) share backbone bottleneck link R bits/sec