

# Computer Networks Performance Metrics



**Advanced Computer Networks**  
**Fall 2014**

# Performance Metrics Outline

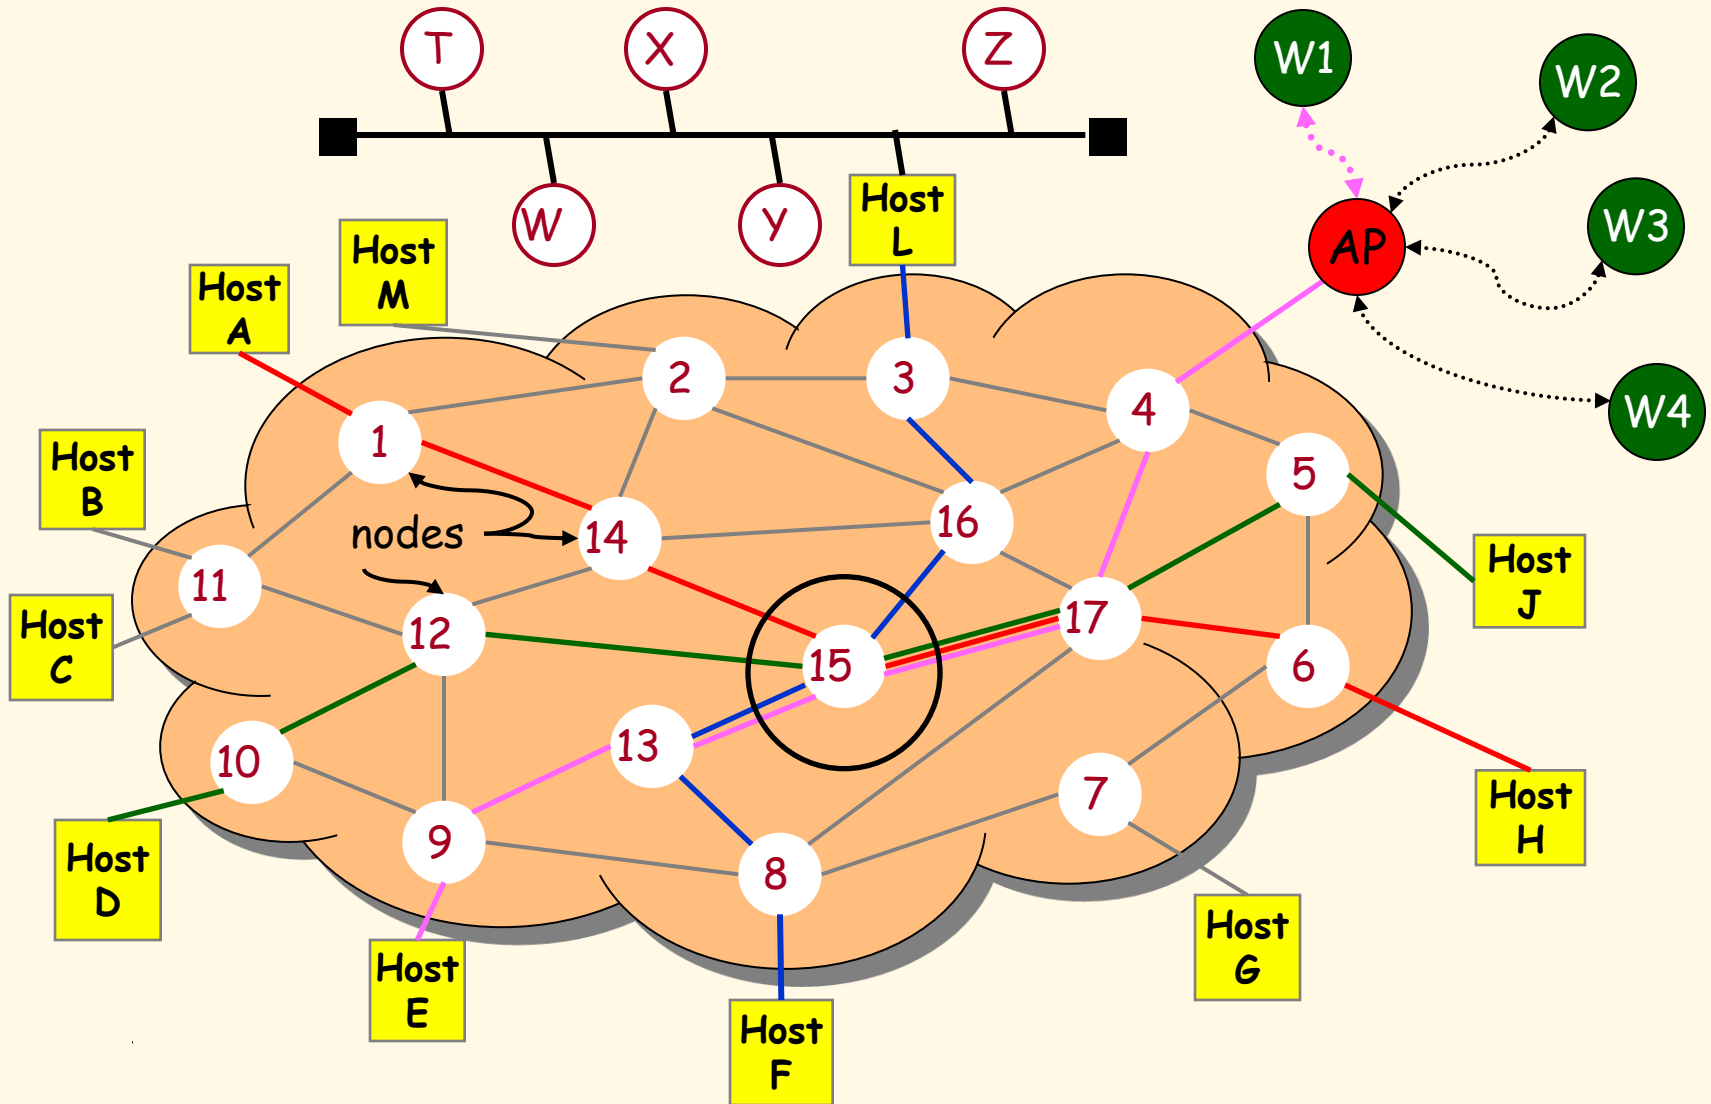
- **Warning: Terminologies can differ!!**
- Introduction to Queueing Model
- Generic Performance Metrics
- Components of Hop and End-to-End Packet Delay
- Traceroute Tool
- Other Performance Measures
- Summary

# Terminology Differences

- Textbook authors, researchers and companies will have different rules and meanings for network performance terms, specifically:
  - Bandwidth versus capacity
  - Link versus channel
  - Latency versus delay versus RTT

***\*\* For this course's exams, use my definitions!!***

# Computer Networks






# How do Loss and Delay occur?


packets *queue* in router buffers

- when the sum of the packet arrival rates at the router exceeds the output link capacity.

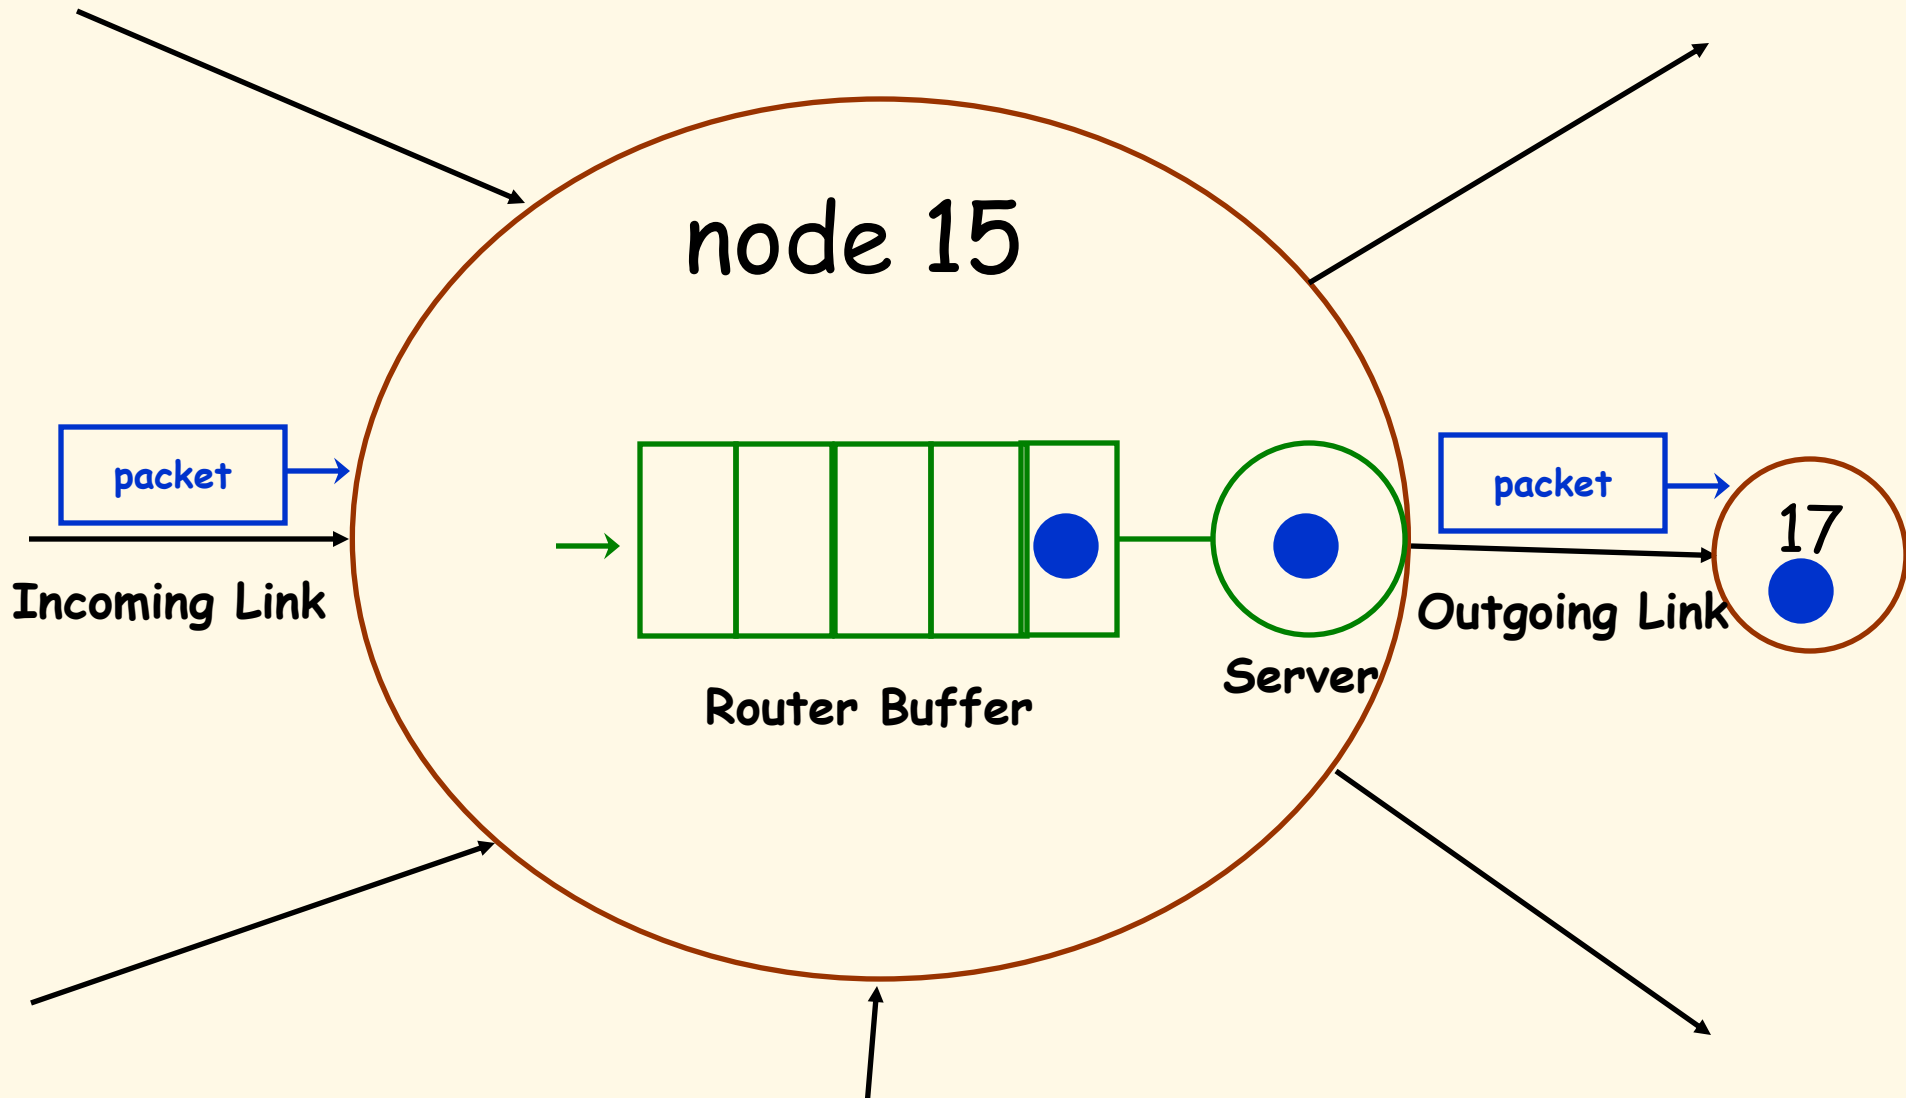


K & R

# Simple Queueing Model



# Router Node



# Generic Performance Metrics

## Utilization ::

- the percentage of time a device is busy servicing a "customer".

## Throughput ::

- the number of jobs processed by the "system" per unit time.

## Response time ::

- the time required to receive a response to a request (round-trip time (RTT) ).

## Delay ::

- the time to traverse from one end to the other of a system.


# Network Performance Measures

- **Channel utilization**:: the average fraction of time a channel is busy  
[e.g.  $Util = 0.8$ ]
  - when overhead is taken into account (i.e., it is excluded from the **useful** bits sent), channel utilization is often referred to as **channel efficiency**.
- **Throughput**:: bits/sec. successfully sent {includes retransmissions}  
[e.g.  $T_{put} = 10 \text{ Mbps}$ ]

# 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



server sends bits  
(fluid) into pipe


pipe that can carry  
fluid at rate  
 $R_s$  bits/sec)

pipe that can carry  
fluid at rate  
 $R_c$  bits/sec)


K & R

# 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

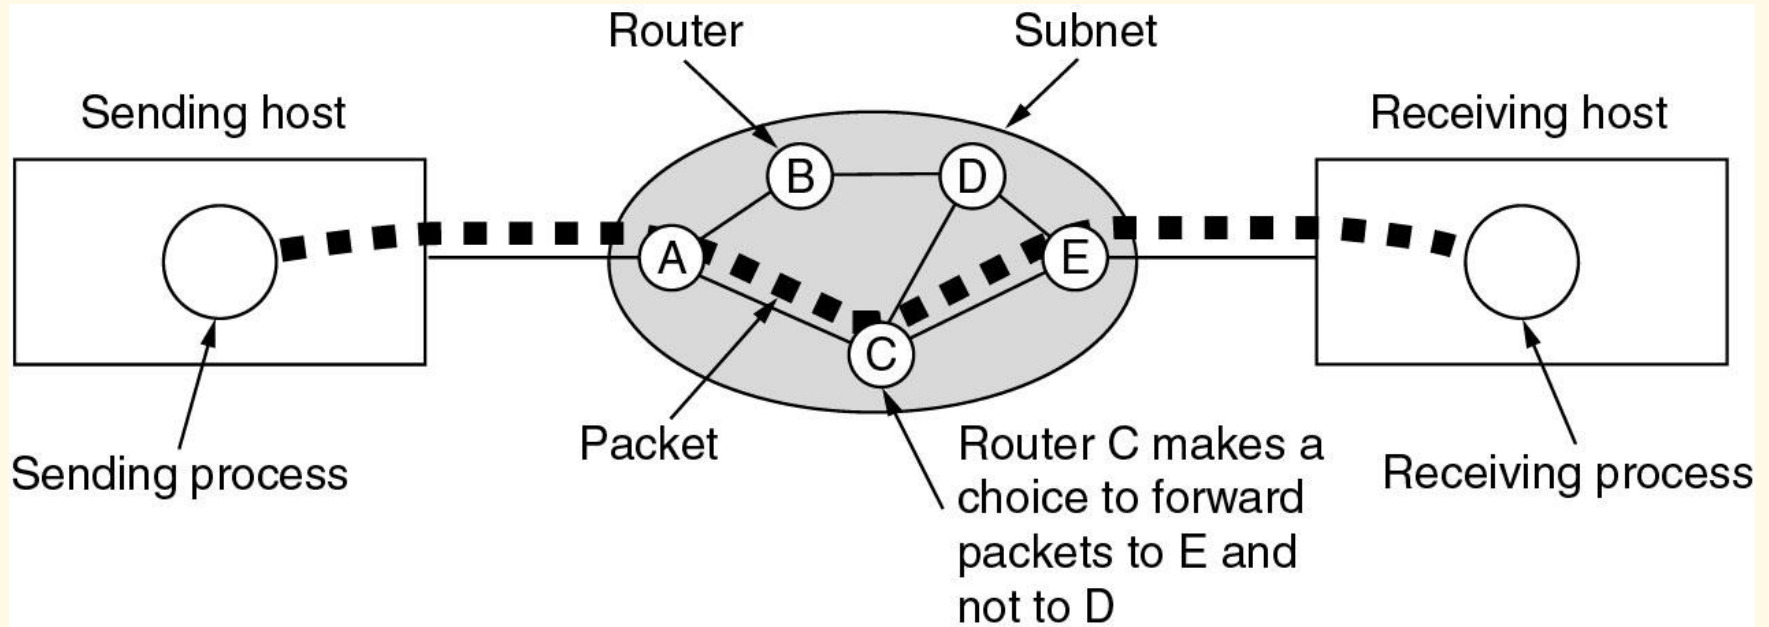
- per-connection end-end throughput:  $\min(R_c, R_s, R/10)$
- in practice:  $R_c$  or  $R_s$  is often the bottleneck
- The **last mile link** has capacity  $R_c$ .



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



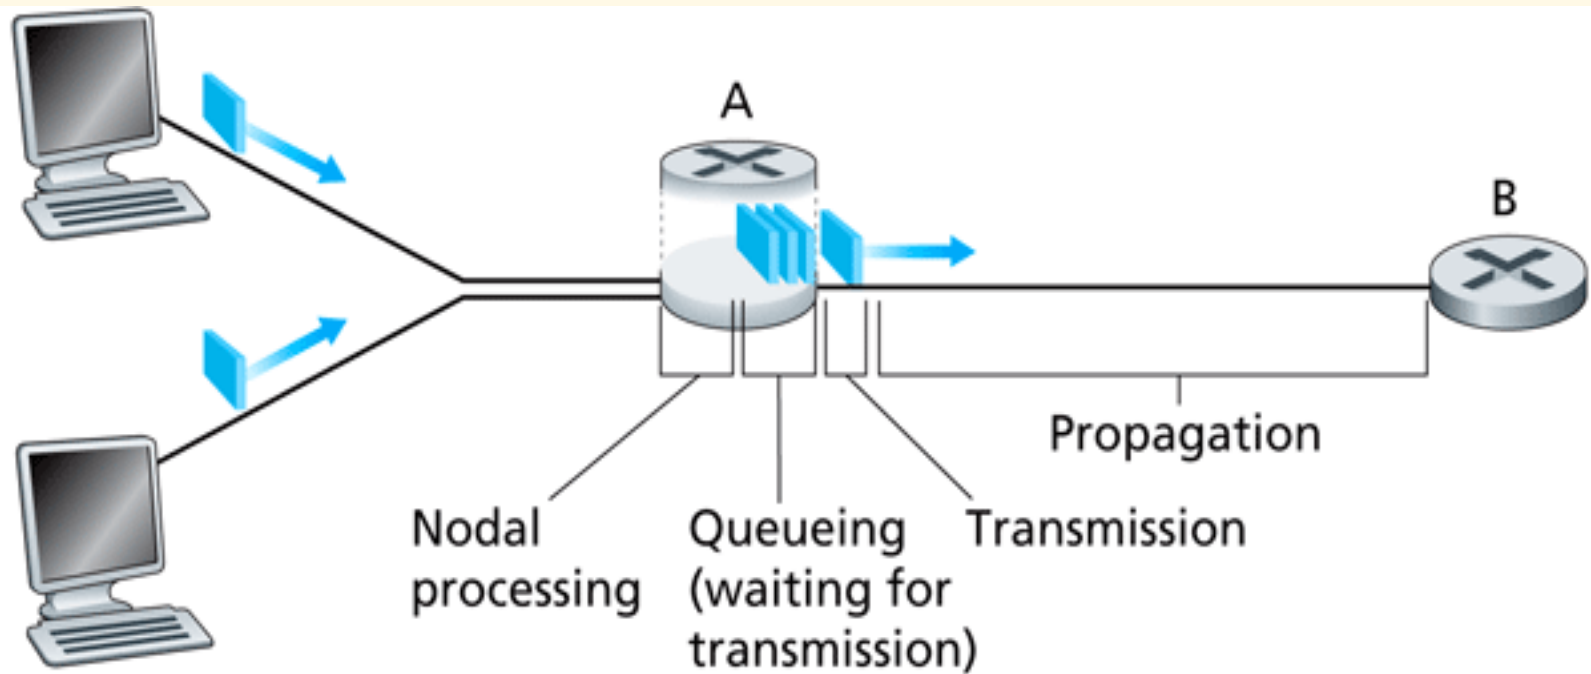
# End-to-End Packet Delay



End-to-end delay includes multiple hop link delays.

Tanenbaum

# Hop Delay Components



**Figure 1.12** ♦ The nodal delay at router A

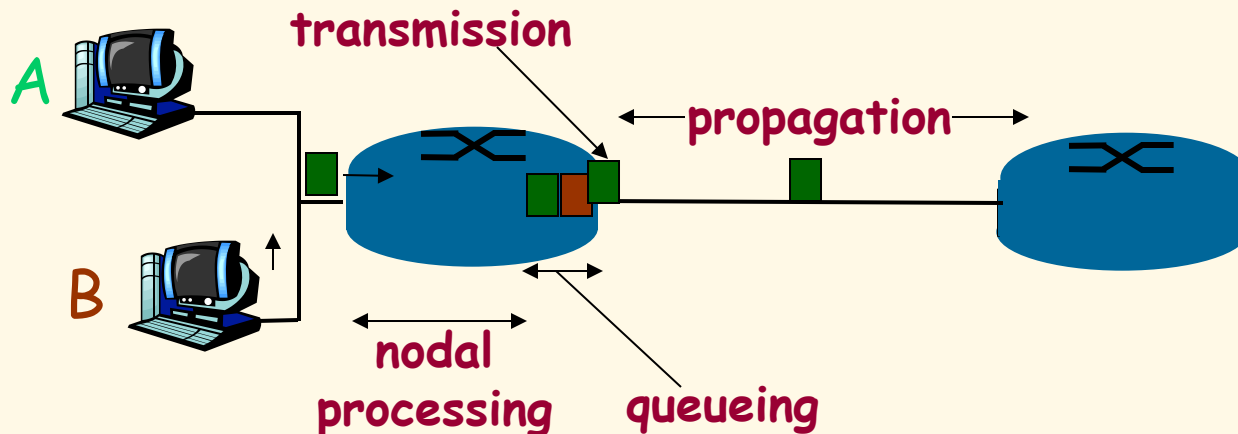
# Four Sources of Packet Delay

## 1. processing at node:

- ❑ Checking for bit errors
- ❑ determine output link
- ❑ Moving packet from input queue to output queue.
  - ❑ Table lookup time (see routing algorithms)

## 2. queueing delay

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



K & R

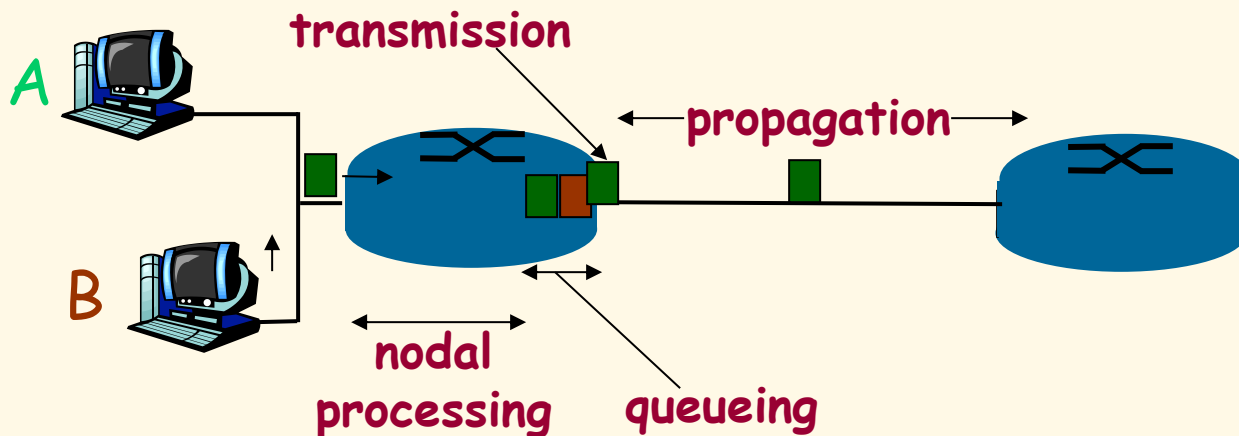
# Delay in packet-switched networks

## 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 ( $\sim 2 \times 10^8$  m/sec)
- propagation delay =  $d/s$



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

# End-to-end Packet Delay

**End-to-end packet delay** :: the time to deliver a packet from source to destination Host (or node).

{Most often, we are interested in the packet delay within the *communications subnet*.} This delay is the sum of the delays on each subnet link traversed by the packet.

Each link delay consists of four components[B&G Bertsekas and Gallager]:

# Link Packet Delay

1. The **processing delay [PROC]** between the time the packet is correctly received at the head node of the incoming link and the time the packet is assigned to an outgoing link queue for transmission.
2. The **queuing delay [QUEUE]** between the time the packet is assigned to a queue for transmission and the time it starts being transmitted. During this time, the packet waits while other packets in the transmission queue are transmitted.

# Link Packet Delay

3. The **transmission delay [TRANS]** between the times that the first and last bits of the packet are transmitted.
4. The **propagation delay [PROP]** is between the time the last bit is transmitted at the head node of the link queue and the time the last bit is received at the next router. This is **proportional** to the physical distance between transmitter and receiver.

# Nodal (Link) Delay

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

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

K & R



# End-to-End Packet Delay

$$\text{Link packet delay} = \text{PROC} + \text{QD} \\ + \text{TRANS} + \text{PROP}$$

**end-to-end packet delay** = sum of  
ALL link packet delays.

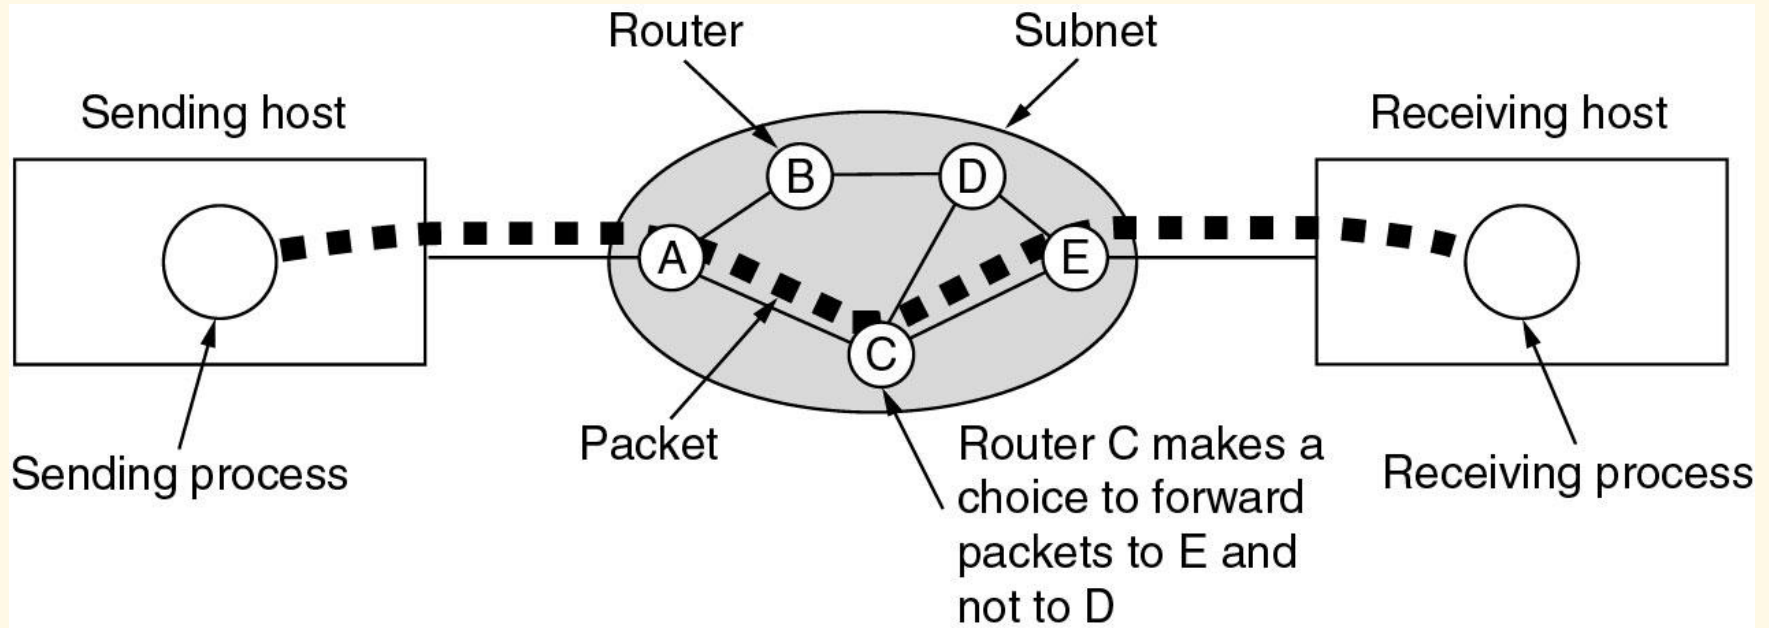
Be Careful !!

**end-to-end** can be defined either:

- from Host-to-Host
- or **only** from end-to-end nodes within the subnetwork.

B & G

# End-to-End Packet Delay




End-to-end delay includes multiple hop link delays.

Tanenbaum

# Bandwidth Delay Product (BDP)

- Amount of data “in flight” or “in the pipe”
- Sometimes relative to RTT (i.e.,  $2 \times$  delay)
- Example:  $100\text{ms} \times 40\text{Mbps} = 4 \text{ M bits}$




- Is this 500 KB?

# Queueing Delay (revisited)

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

traffic intensity =  $\lambda a / R$




- $\lambda a / R \sim 0$ : average queueing delay small
- $\lambda a / R \rightarrow 1$ : delays become large
- $\lambda a / R =$  or  $> 1$ : more "work" arriving than can be serviced, average delay infinite!

K & R

# "Real" Internet Delays and Routes

- What do "real" Internet delay & loss look like?
- *Traceroute* program: provides delay measurement from source to router along end-to-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.



K & R

# "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)

K & R

# Network Performance Measures

- **Latency ::**
  - usually implies the **minimum** possible delay. Latency assumes no queueing and no contention encountered along the path.
- **Goodput ::**
  - *{measured at the receiver}* the rate in bits per second of useful traffic received. Goodput excludes duplicate packets and packets dropped along the path.
- **Fairness ::**
  - either **Jain's fairness** or **max-min fairness** are used to measure fair treatment among competing flows.
- **Quality of Service (QoS) ::**
  - a QoS measure accounts for importance of specific metric to one type of application [e.g. **jitter** and **playable frame rate** for streaming media].

# Wireless Performance Metrics

WLANs and WSNs are concerned with packet loss and employ additional metrics:

## Delivery ratio::

- the ratio of packets received to packets sent {excluding duplicates and retransmissions}.

## Packet loss rate::

- the percentage of packets lost or dropped.


## Link layer retransmission rates::

- the percentage of DL layer frames that are retransmitted.



# Packet Loss

- queue (aka buffer) preceding link in buffer has finite capacity.
- packet arriving to full queue is dropped (aka lost) **[FIFO Drop Tail router]**.
- lost packet may be retransmitted by previous node, by source end system, or not at all.



K & R

# Packet/Frame losses

- BDP often used as 'rule of thumb' in choosing size of the FIFO buffer at a router.
- Other causes of packet/frame loss:
  - Collisions on media (both wired and wireless media {e.g., hidden terminal})
  - Wireless conditions (SNR, multipath fading)
  - Out-of-order arrival.

# Performance Metrics Summary

- The three most general performance measures are : **utilization, throughput and response time.**
- In computer networks, **end-to-end delay** is an important performance metric.
- Queueing models are used to analyze and estimate computer network performance.

# Performance Metrics Summary

- Other useful metrics include: **latency**, **goodput**, **fairness** and QoS metrics such as **jitter** or **playable frame rate**.
- In wireless networks, **delivery ratio**, **packet loss rate** and **link layer retransmission rates** are valuable network measures.