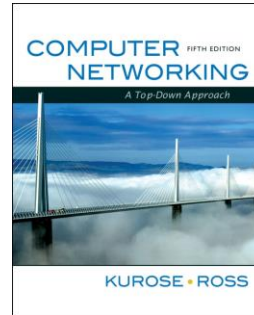# RSC
Part II: Network Layer
1. Basic Concepts

Redes y Servicios de Comunicaciones
Universidad Carlos III de Madrid

These slides are, mainly, part of the companion slides to the book "Computer Networking: A Top Down Approach" generously made available by their authors (see copyright below). The slides have been adapted, where required, to the teaching needs of the subject above.

Computer Networking: A Top Down Approach 5th edition.
Jim Kurose, Keith Ross
Addison-Wesley, April 2009.

RSC Part II: Network Layer

- II. 1 Basic Network layer concepts
  - Network layer
  - Network service model
  - Introduction to router architecture
  - Delay, packet loss, throughput
- II.2 Introduction to IP
  - Datagram format
  - ICMP
- II.3 IP addressing
- II.4 IP in operation
  - ARP

- II.5 Network routing
  - Link state
  - Distance Vector
  - Hierarchical routing
- II.6 Routing in the Internet
  - RIP
  - OSPF
  - BGP
  - Broadcast and multicast
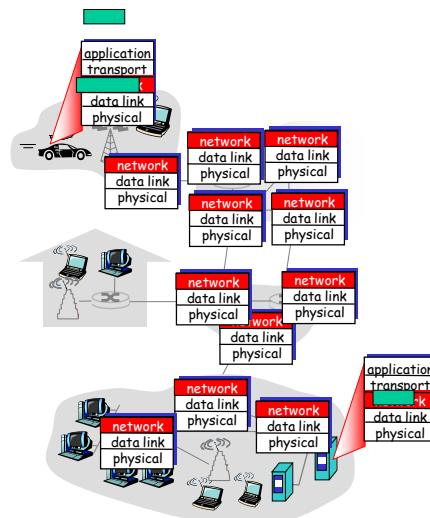
Part II: Network Layer

Chapter goals:
☐ understand principles behind network layer services:
- ○ network layer service models
- ○ forwarding versus routing
- ○ how a router works
- ○ addressing
- ○ routing (path selection)
- ○ dealing with scale

☐ instantiation, implementation in the Internet

Network layer

☐ transport segment from sending to receiving host

☐ on sending side encapsulates segments into datagrams

☐ on receiving side, delivers segments to transport layer

☐ network layer protocols in *every* host, router

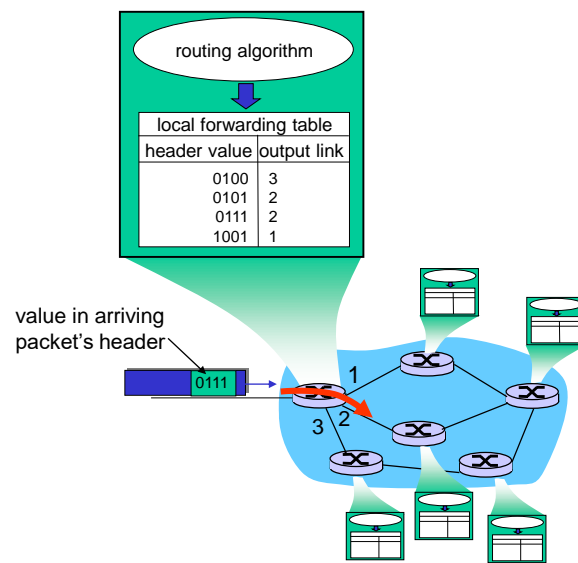☐ router examines header fields in all IP datagrams passing through it

Two Key Network-Layer Functions

❒ *forwarding:* move packets from router's input to appropriate router output

❒ *routing:* determine route taken by packets from source to destination

  ○ *routing algorithms*

<u>analogy:</u>

❒ routing: process of planning trip from source to destination

❒ forwarding: process of getting through single interchange

Interplay between routing and forwarding

routing algorithm

local forwarding table

header value	output link
0100	3
0101	2
0111	2
1001	1

value in arriving packet's header

0111

1

3 2

Connection setup

❒ 3rd important function in *some* network architectures:
  ○ ATM, frame relay, X.25
❒ before datagrams flow, two end hosts *and* intervening routers establish virtual connection
  ○ routers get involved
❒ network vs transport layer connection service:
  ○ network: between two hosts (may also involve intervening routers in case of VCs)
  ○ transport: between two processes

Network service model

Q: What *service model* for "channel" transporting datagrams from sender to receiver?

Example services for individual datagrams:
❒ guaranteed delivery
❒ guaranteed delivery with less than 40 msec delay

Example services for a flow of datagrams:
❒ in-order datagram delivery
❒ guaranteed minimum bandwidth to flow
❒ restrictions on changes in inter-packet spacing

Network layer service models:

| Network Architecture	Service Model	Guarantees ?				Congestion feedback
		Bandwidth	Loss	Order	Timing	
Internet	best effort	none	no	no	no	no (inferred via loss)
ATM	CBR	constant rate	yes	yes	yes	no congestion
ATM	VBR	guaranteed rate	yes	yes	yes	no congestion
ATM	ABR	guaranteed minimum	no	yes	no	yes
ATM	UBR	none	no	yes	no	no

Network layer connection and connection-less service

□ datagram network provides network-layer connectionless service

□ VC network provides network-layer connection service

□ analogous to the transport-layer services, but:
  ○ service: host-to-host
  ○ no choice: network provides one or the other
  ○ implementation: in network core

Router Architecture Overview

Two key router functions:

❑ run routing algorithms/protocol (RIP, OSPF, BGP)

❑ *forwarding* datagrams from incoming to outgoing link
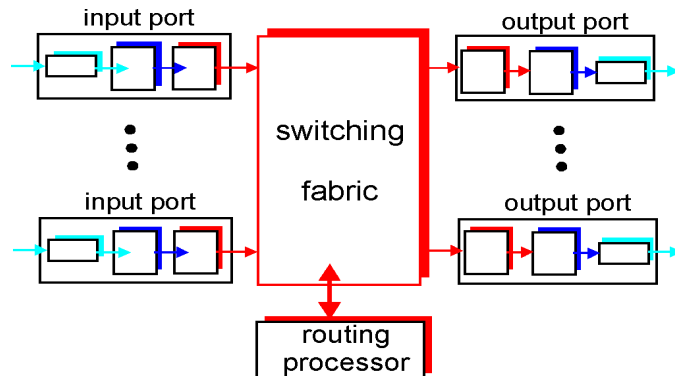
How much buffering?
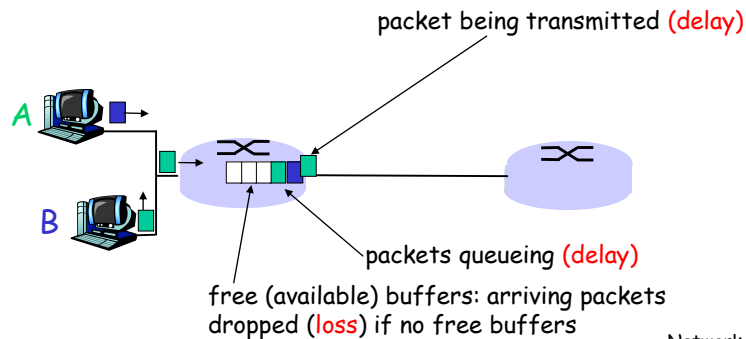
❑ RFC 3439 rule of thumb: average buffering equal to "typical" RTT (say 250 msec) times link capacity C

  ○ e.g., C = 10 Gps link: 2.5 Gbit buffer

❑ Recent recommendation: with *N* flows, buffering equal to $\dfrac{RTT \cdot C}{\sqrt{N}}$

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

packet being transmitted (delay)

A

B

packets queueing (delay)

free (available) buffers: arriving packets dropped (loss) if no free buffers

---

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

A

B

transmission

propagation

nodal processing

queueing
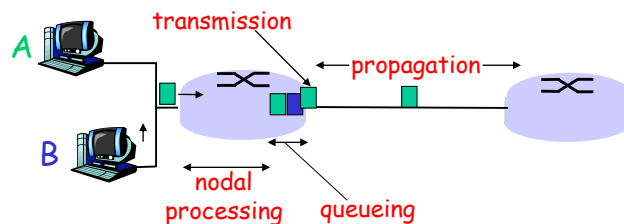
Delay in packet-switched networks

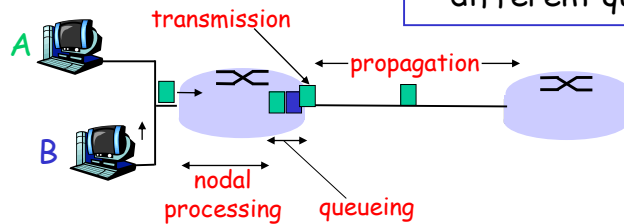**3. Transmission delay:**
- R=link bandwidth (bps)
- L=packet length (bits)
- time to send bits into link = L/R

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!

A

transmission

propagation

B

nodal processing

queueing

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
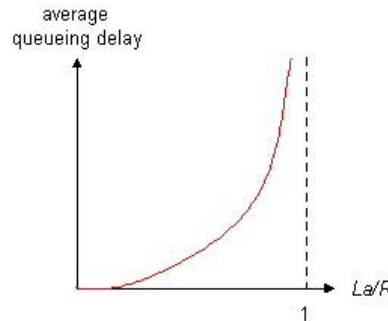 - a few microsecs to hundreds of msecs

Queueing delay (revisited)

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

traffic intensity = La/R

average queueing delay

❑ 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.

3 probes   3 probes   3 probes
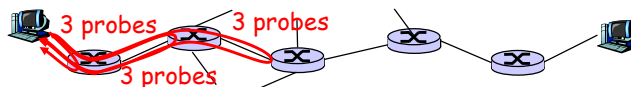
"Real" Internet delays and routes

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
```
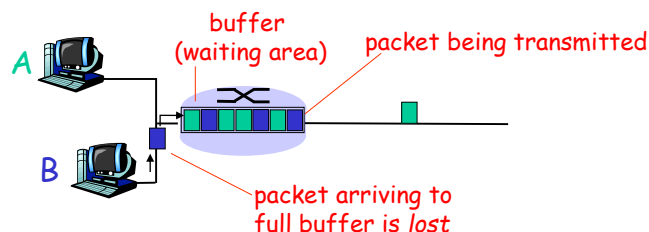
trans-oceanic link

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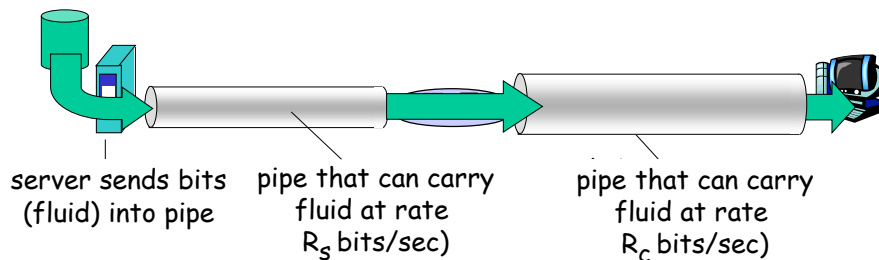

Packet loss

- ☐ 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

buffer (waiting area)

packet being transmitted

A

B

packet arriving to full buffer is *lost*

10

Throughput

- *throughput:* rate (bits/time unit) at which bits transferred between sender/receiver
  - *instantaneous:* rate at given point in time
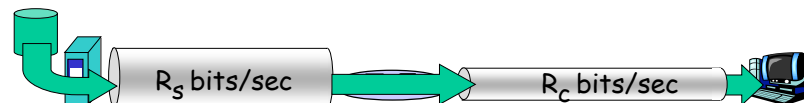 - *average:* rate over longer period of time

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)

Throughput (more)

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

R_s bits/sec R_c bits/sec

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

R_s bits/sec R_c bits/sec
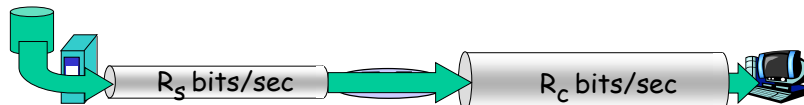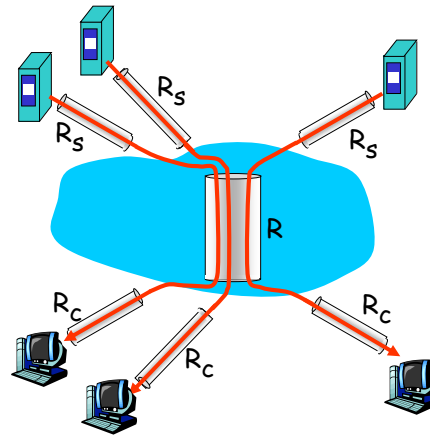
bottleneck link

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