Reti di Calcolatori

Delays and Other Stories

Stefano Basagni
basagni@ece.neu.edu
312 Dana
Dept. of Electrical and Computer Engineering

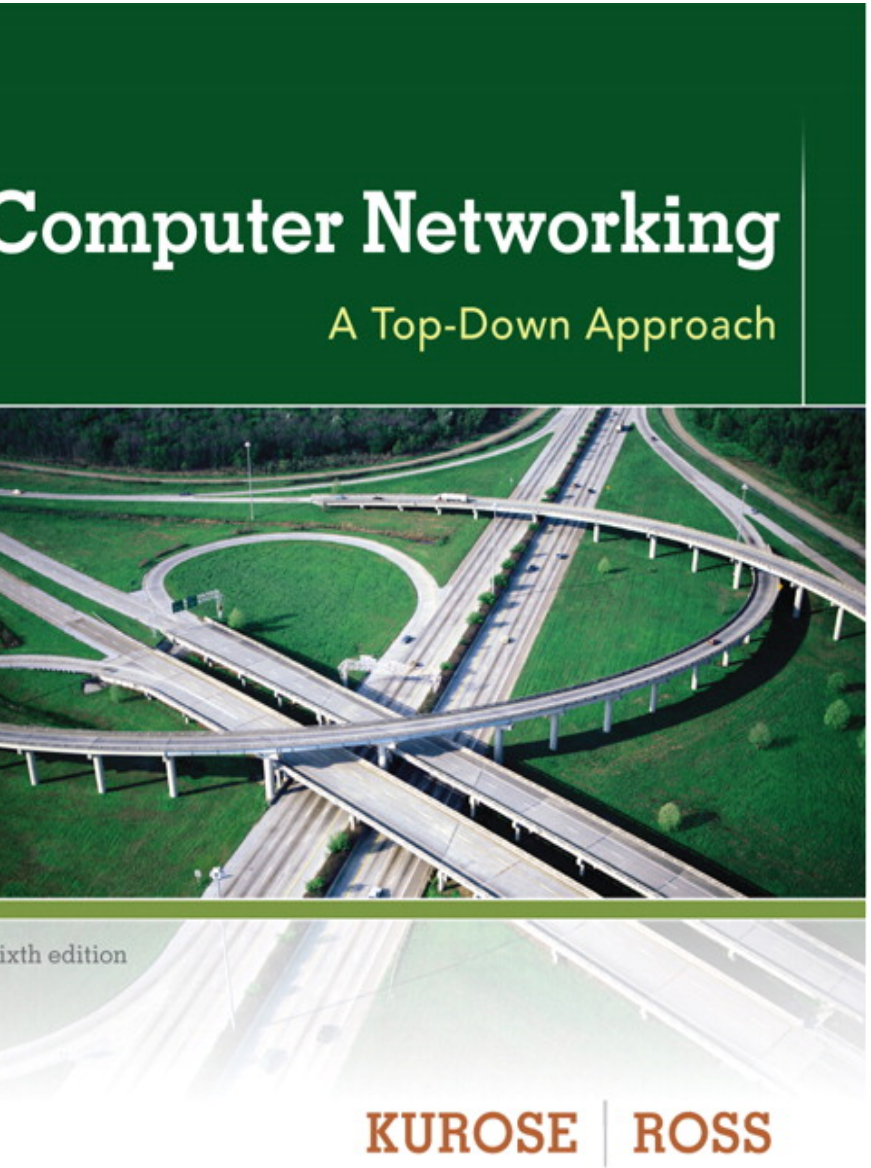

Computer Networking: A Top-Down Approach

6th edition
Jim Kurose, Keith Ross
Addison-Wesley
March 2012

Stefano Basagni

- Associate professor of Computer Engineering
- Office hours: By appointment
- Email: basagni@ece.neu.edu
- Web: http://www.ece.neu.edu/faculty/basagni/
- Come visit! Provide feedback!

How do loss and delay occur?

ackets queue in router buffers

Loss: packet arrival rate to link (temporarily) exceeds output link capacity

Delay: packets queue, wait for turn

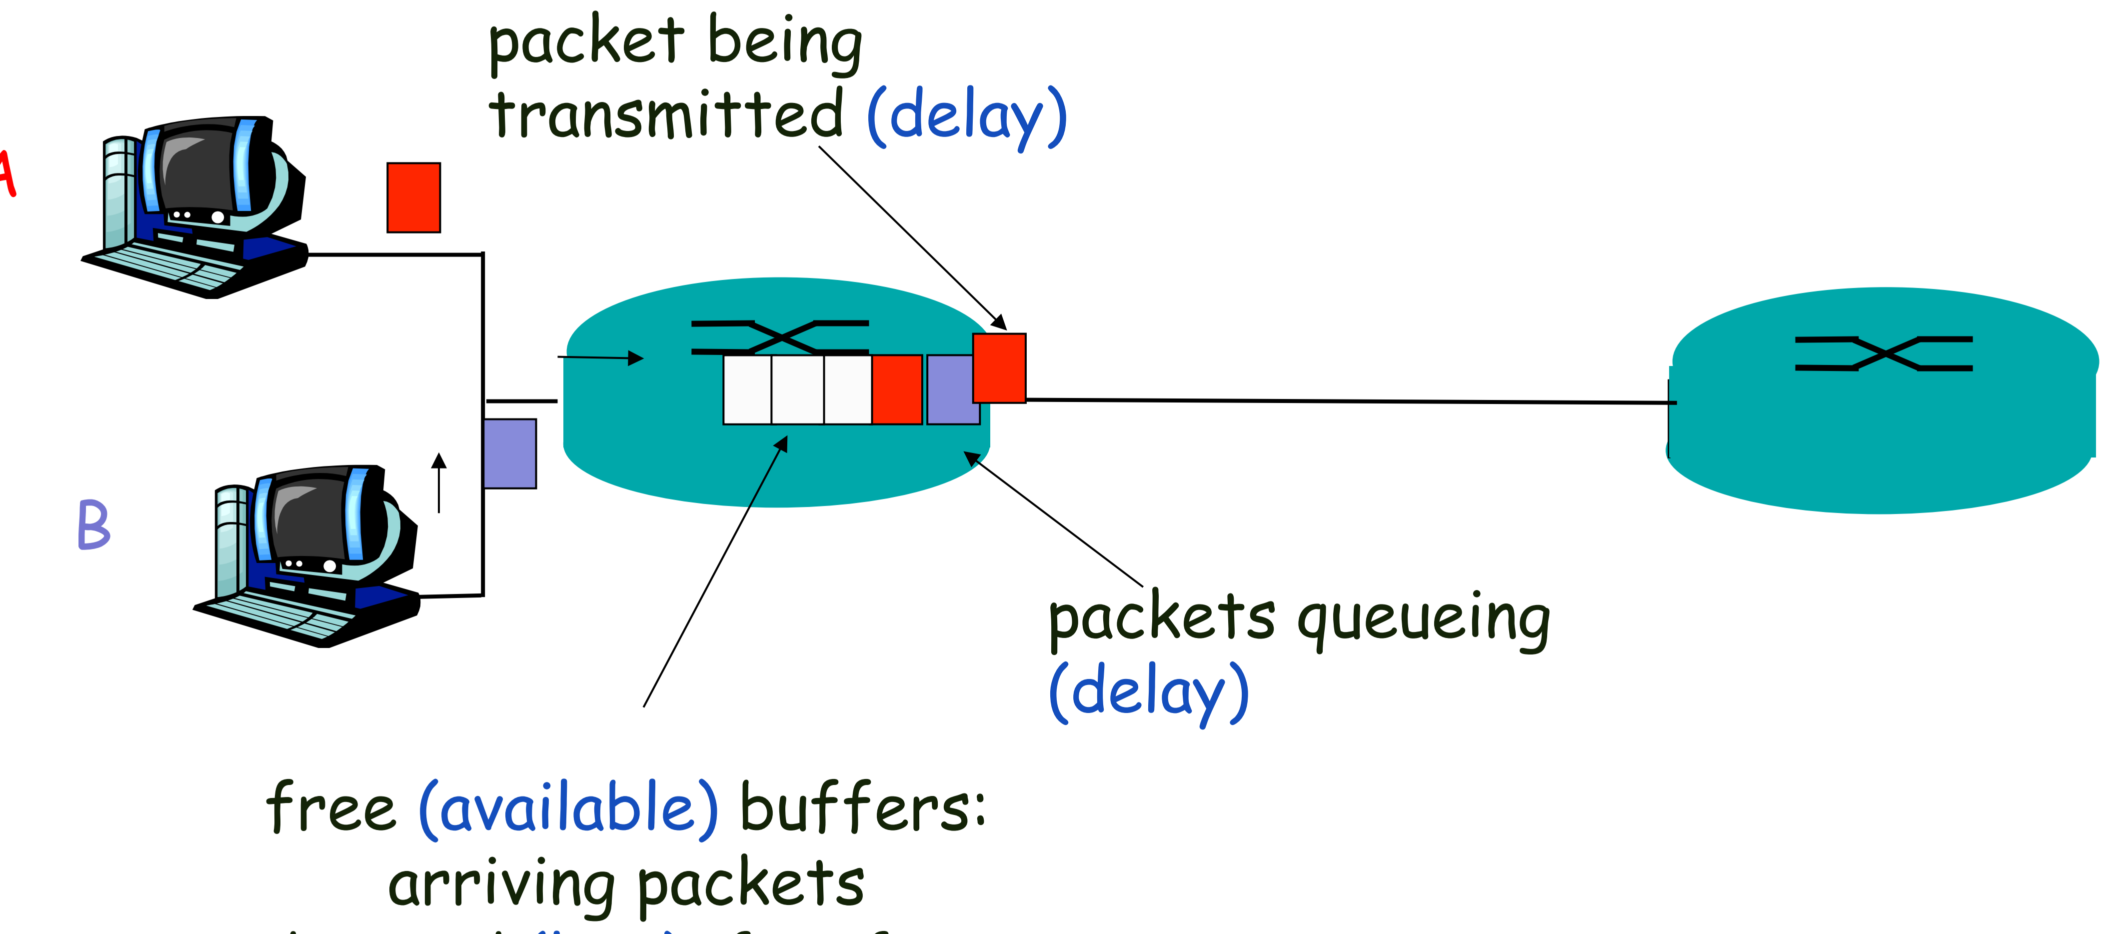

lodal processing:

Check bit errors

Determine output link

Typically < msec

- ## Queueing:

  - time waiting at output link transmission
  - depends on congestion leve router

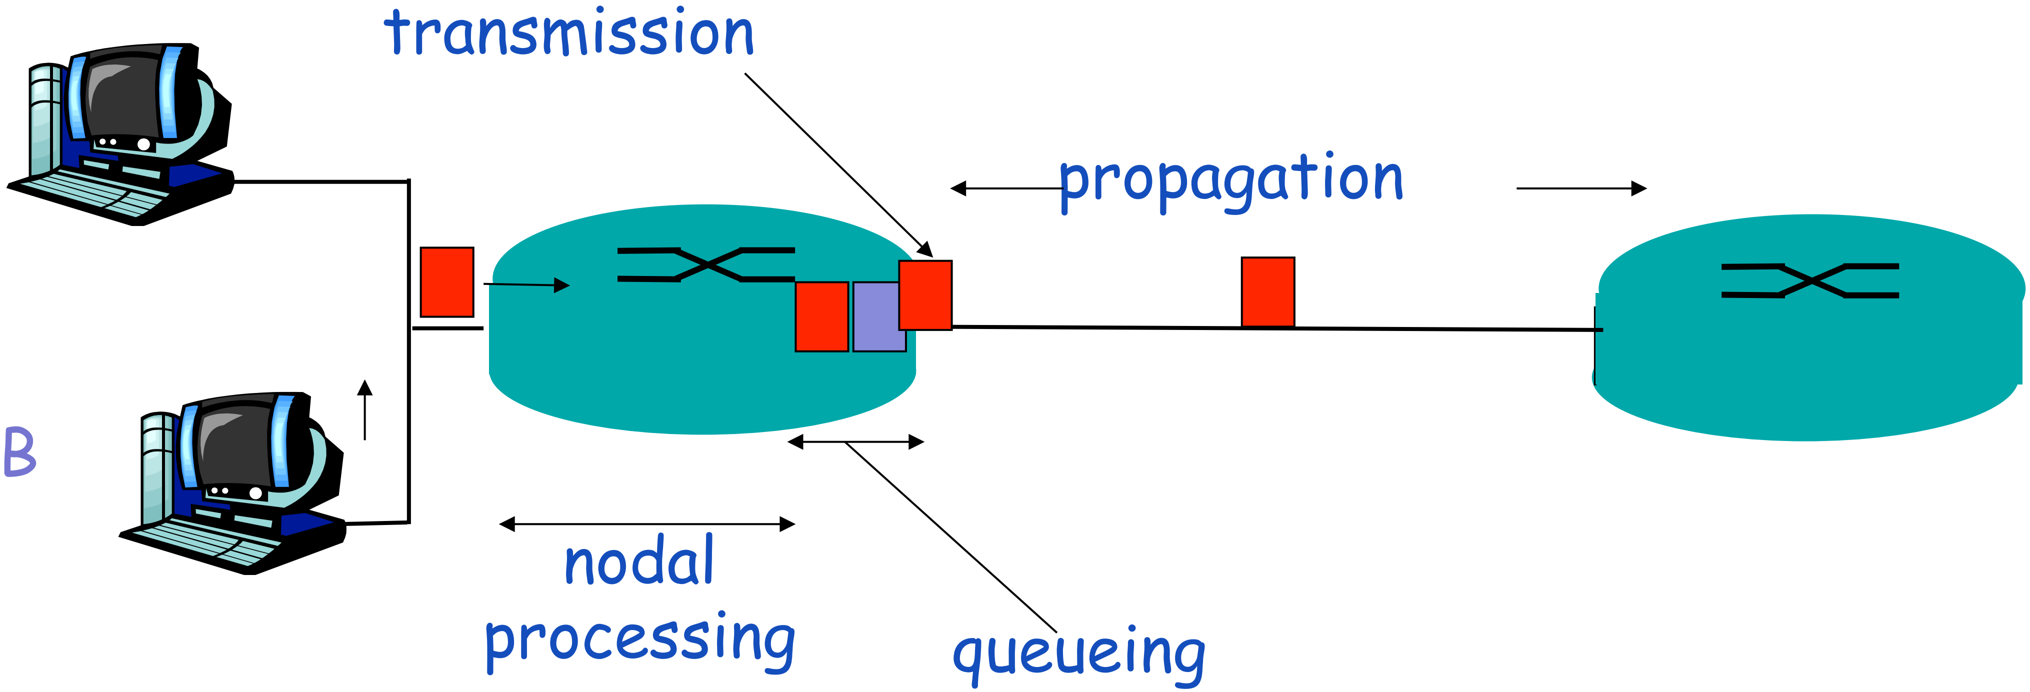

- ransmission delay
  - R=link bandwidth (bps)
  - L=packet length (bits)
  - time to send bits into link = L/R

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

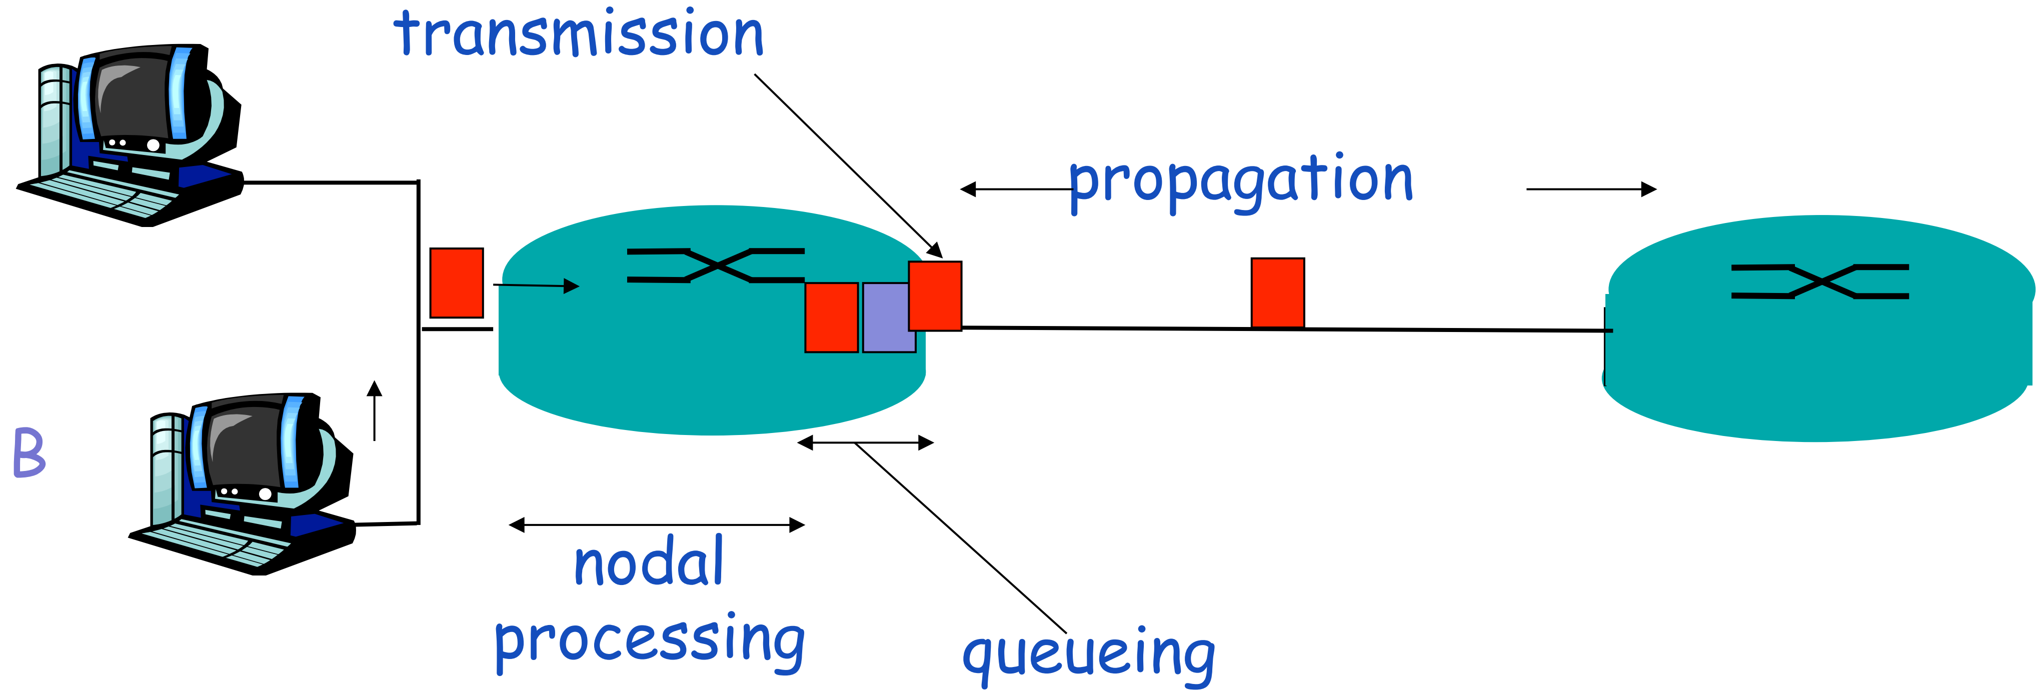

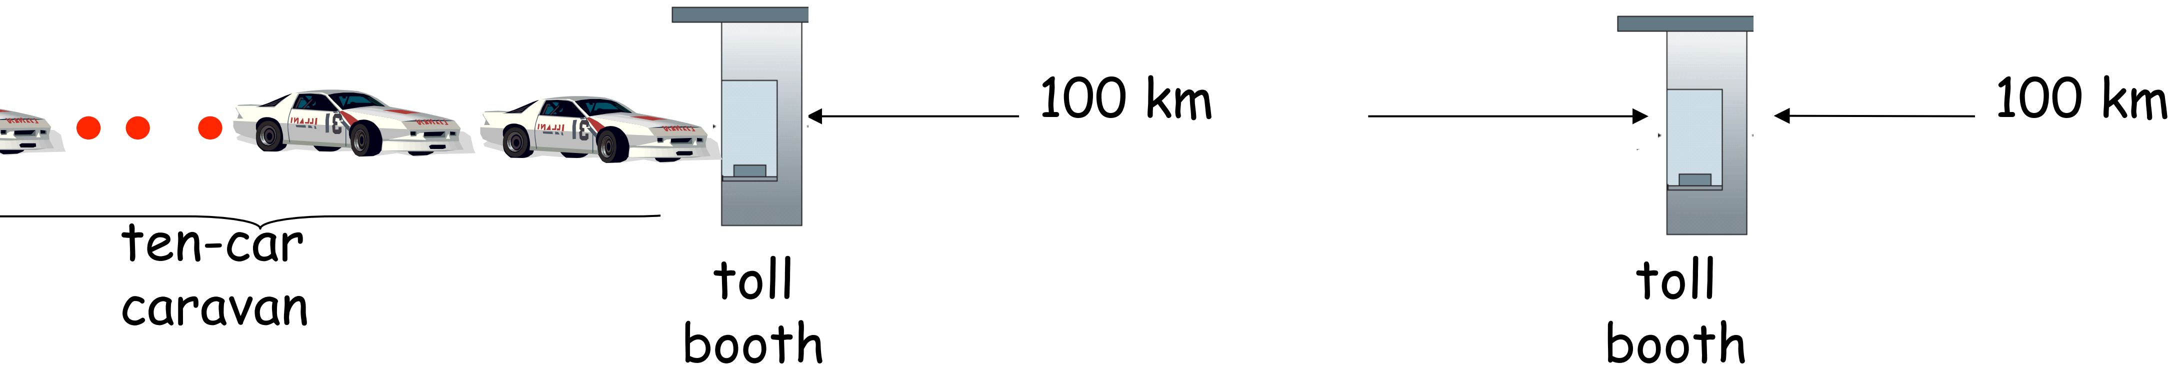

“propagate” at
km/hr

ooth takes 12 sec to service
(car transmission time)

bit; caravan~packet

ow long until caravan is lined up
ore 2^{nd} toll booth?

- time to “push” the entire carav
through the toll booth onto
highway = 12*10 = 120 sec

- time for last car to propagate
from 1^{st} to 2^{nd} toll booth:
  - 100km/(100km/hr) = 1 hr

- A: 62 minutes

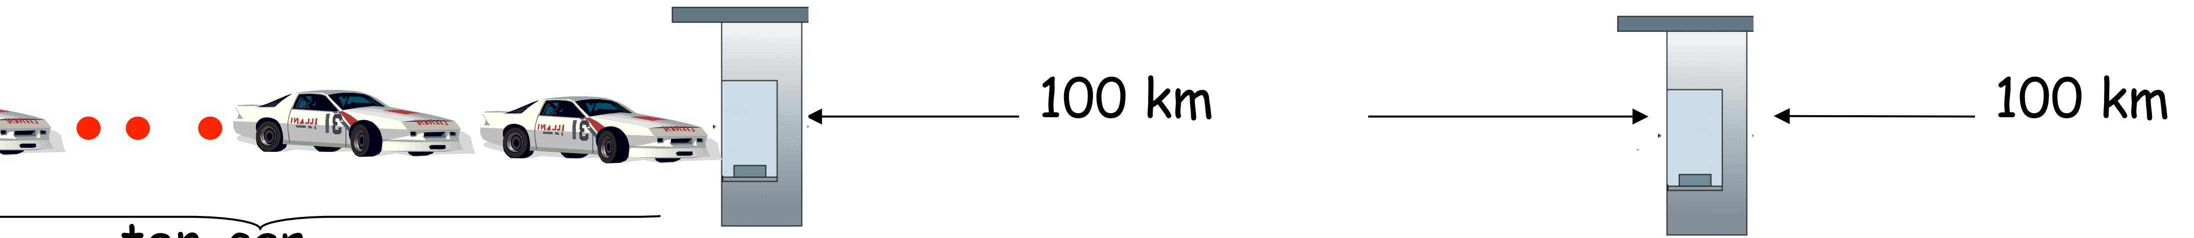

suppose cars now "propagate" at 1000 km/hr

and suppose toll booth now takes one min to service a ca

Q: Will cars arrive to 2nd booth before all cars servicec
at first booth?

- A: Yes! after 7 min, 1st car arrives at second booth; three cars
still at 1st booth

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

- 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

link bandwidth (bps)
packet length (bits)
average packet arrival rate

affic intensity = La/R

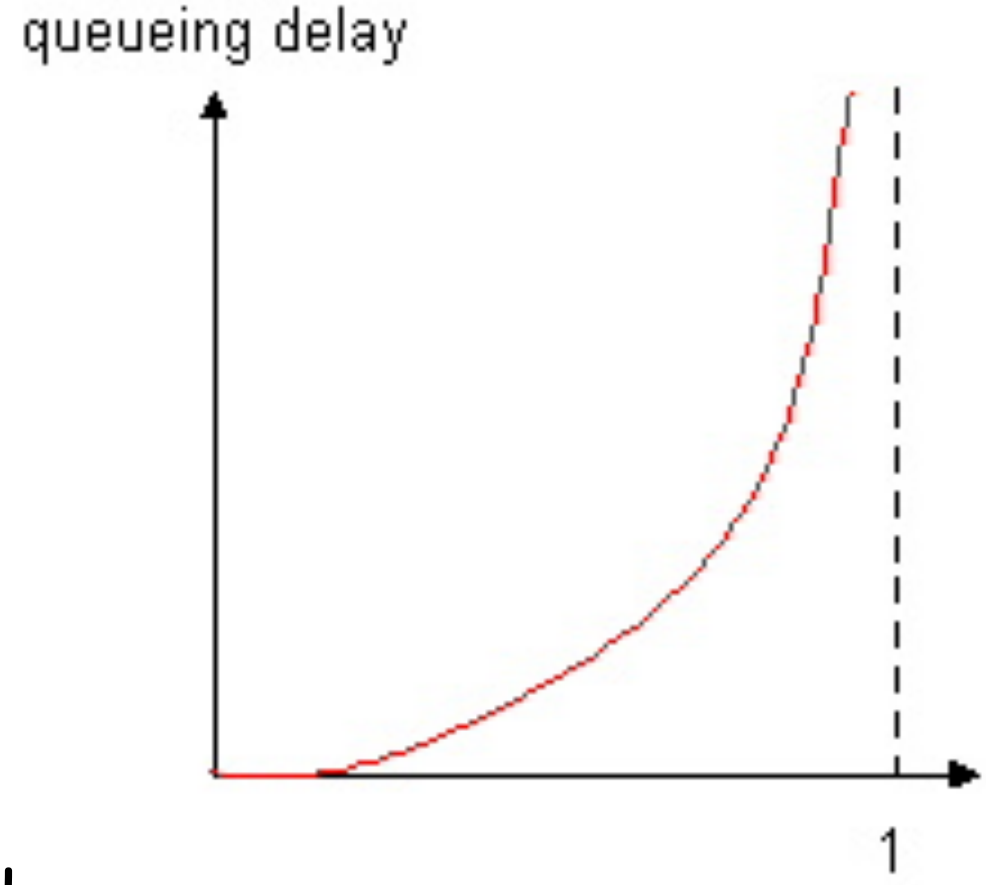

a/R ~ 0: average queueing delay small
a/R -> 1: delays become large
a/R > 1: more "work" arriving than
an be serviced, avg delay infinite

La/R ~ 0

- What do "real" Internet delay & loss look like?
- Traceroute program: provides delay measurements from source to router i along the end-to-end Internet path towards the destination. For all i:
  - sends three packets that will reach router i on path towar
    destination
  - router i will return packets to sender
  - sender times interval between transmission and reply

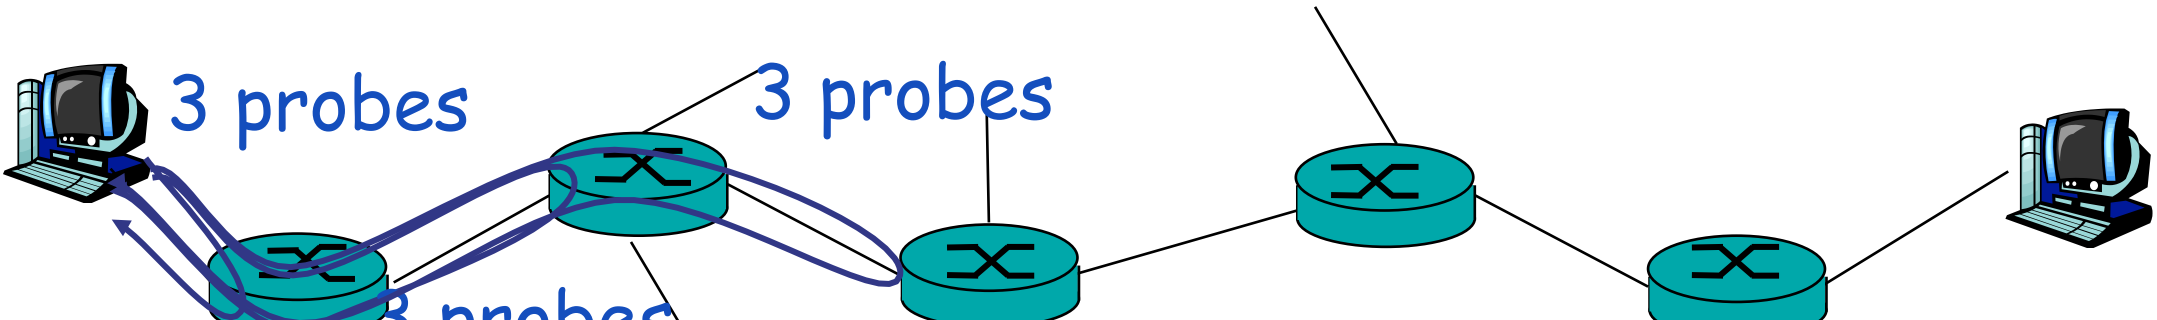

traceroute: gaia.cs.umass.edu to www.eurecom.fr

Three delay measurements from gaia.cs.umass.edu to cs-gw.cs.umass.edu

cs-gw (128.119.240.254) 1 ms 1 ms 2 ms
border1-rt-fa5-1-0.gw.umass.edu (128.119.3.145) 1 ms 1 ms 2 ms
cht-vbns.gw.umass.edu (128.119.3.130) 6 ms 5 ms 5 ms
jn1-at1-0-0-19.wor.vbns.net (204.147.132.129) 16 ms 11 ms 13 ms
jn1-so7-0-0-0.wae.vbns.net (204.147.136.136) 21 ms 18 ms 18 ms
abilene-vbns.abilene.ucaid.edu (198.32.11.9) 22 ms 18 ms 22 ms
nycm-wash.abilene.ucaid.edu (198.32.8.46) 22 ms 22 ms 22 ms
62.40.103.253 (62.40.103.253) 104 ms 109 ms 106 ms
de2-1.de1.de.geant.net (62.40.96.129) 109 ms 102 ms 104 ms
de.fr1.fr.geant.net (62.40.96.50) 113 ms 121 ms 114 ms
renater-gw.fr1.fr.geant.net (62.40.103.54) 112 ms 114 ms 112 ms
nio-n2.cssi.renater.fr (193.51.206.13) 111 ms 114 ms 116 ms
nice.cssi.renater.fr (195.220.98.102) 123 ms 125 ms 124 ms
r3t2-nice.cssi.renater.fr (195.220.98.110) 126 ms 126 ms 124 ms
eurecom-valbonne.r3t2.ft.net (193.48.50.54) 135 ms 128 ms 133 ms
194.214.211.25 (194.214.211.25) 126 ms 128 ms 126 ms
* * *
* * *
fantasia.eurecom.fr (193.55.113.142) 132 ms 128 ms 136 ms

trans-oceanic li

* means no response (probe lost, router not repl

ueue (aka buffer) preceding output link has finite capacity

acket arriving to full queue dropped (aka lost)

st packet may be retransmitted by previous node, by source end ystem, or not at all

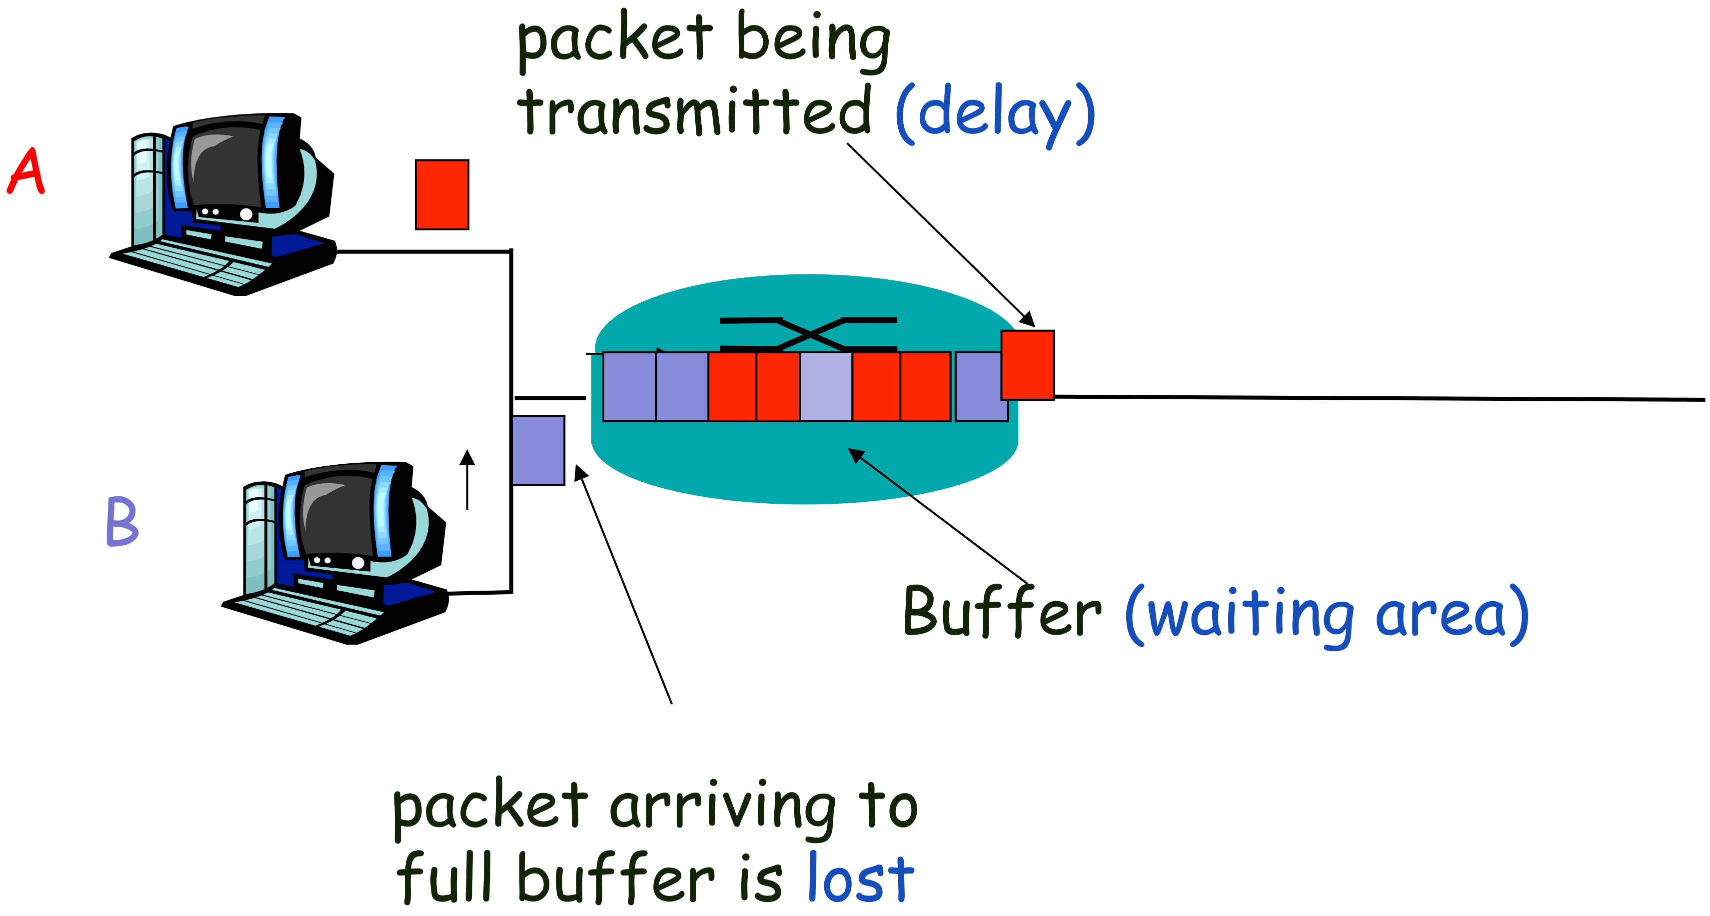

- throughput: rate (bits/time unit) at which bits are transferred between sender/receiver
  - instantaneous: rate at given point in time
  - average: rate over longer period of time

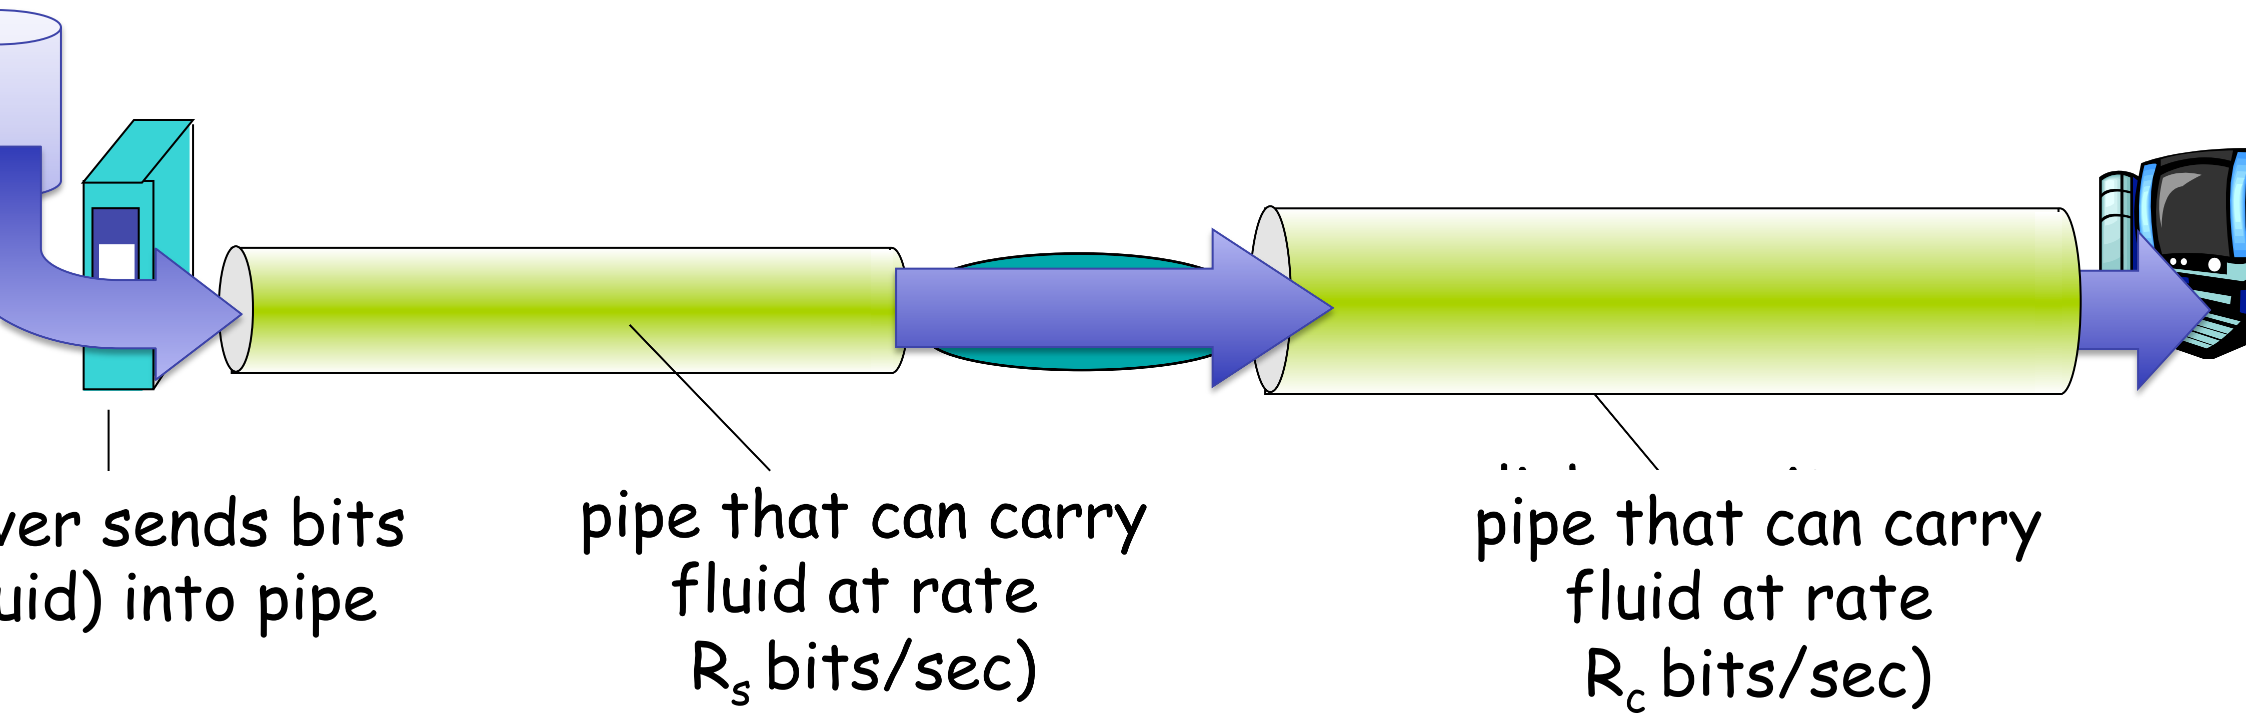

$R_s < R_c$ What is average end-end throughput?

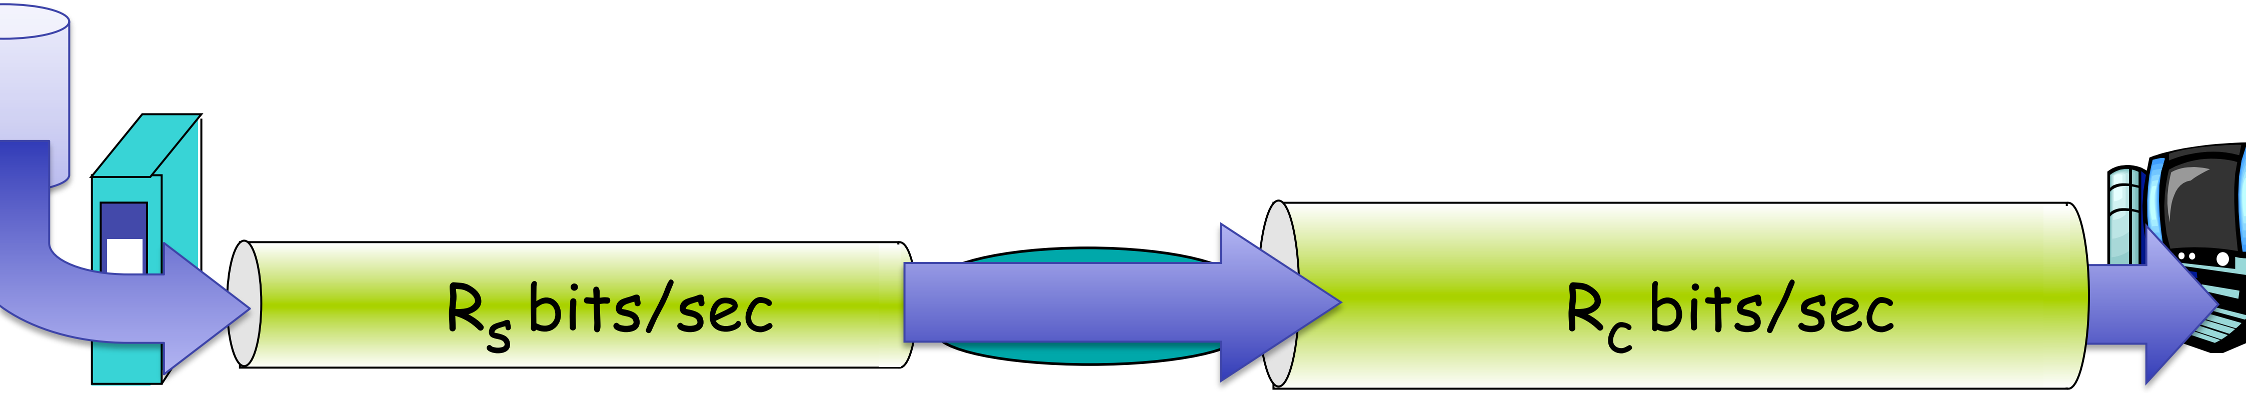

$R_c < R_s$ What is average end-end throughput?

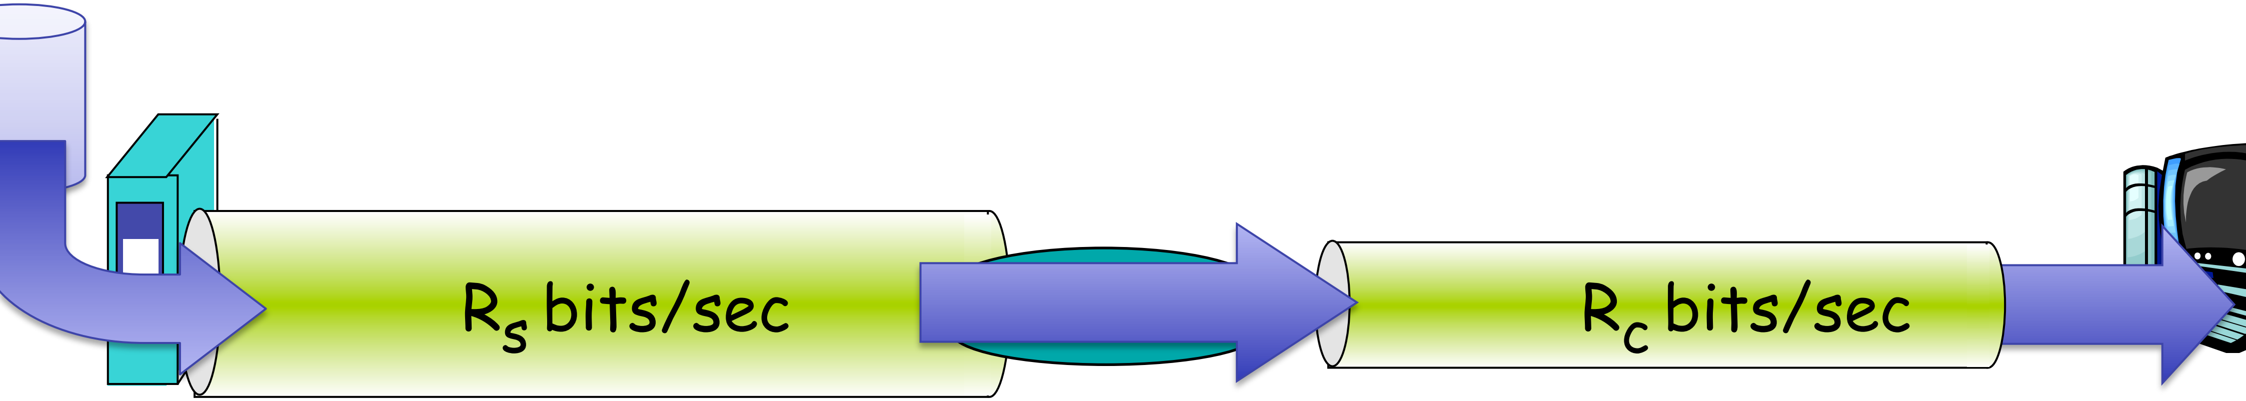

tleneck link on end-end path that constrains end-end
oughput

er-connection end-
nd throughput:
nin(R_c,R_s,R/10)

n practice: R_c or R_s
s often bottleneck

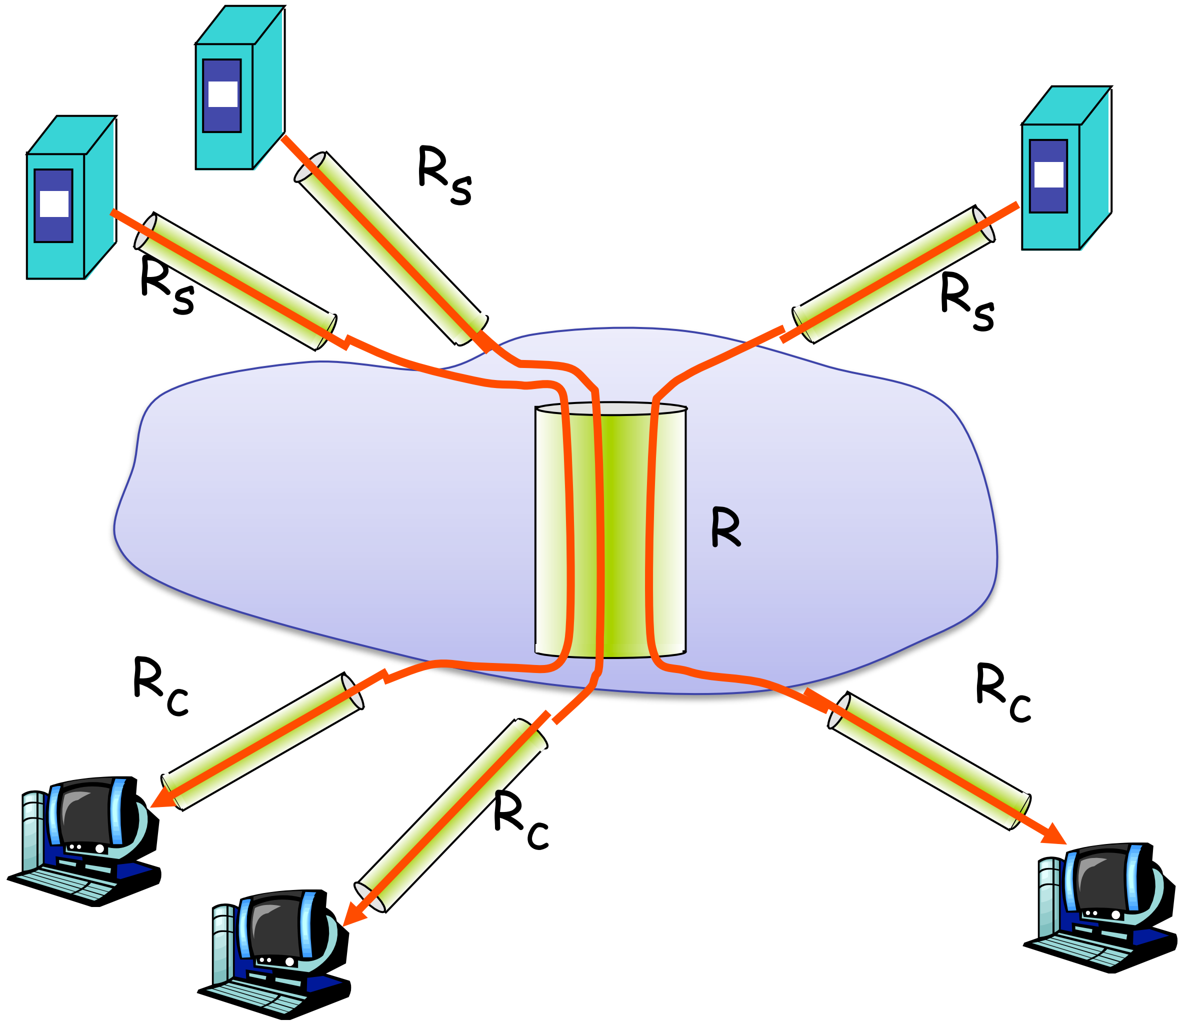

10 connections (fairly) share backbone
bottleneck link R bits/sec

Networks are complex!

- many "pieces":
  - hosts
  - routers
  - links of various media
  - applications
  - protocols
  - hardware, software

Question:

Is there any hope of *organiz*
the structure of network

Or at least our discussion of
networks?

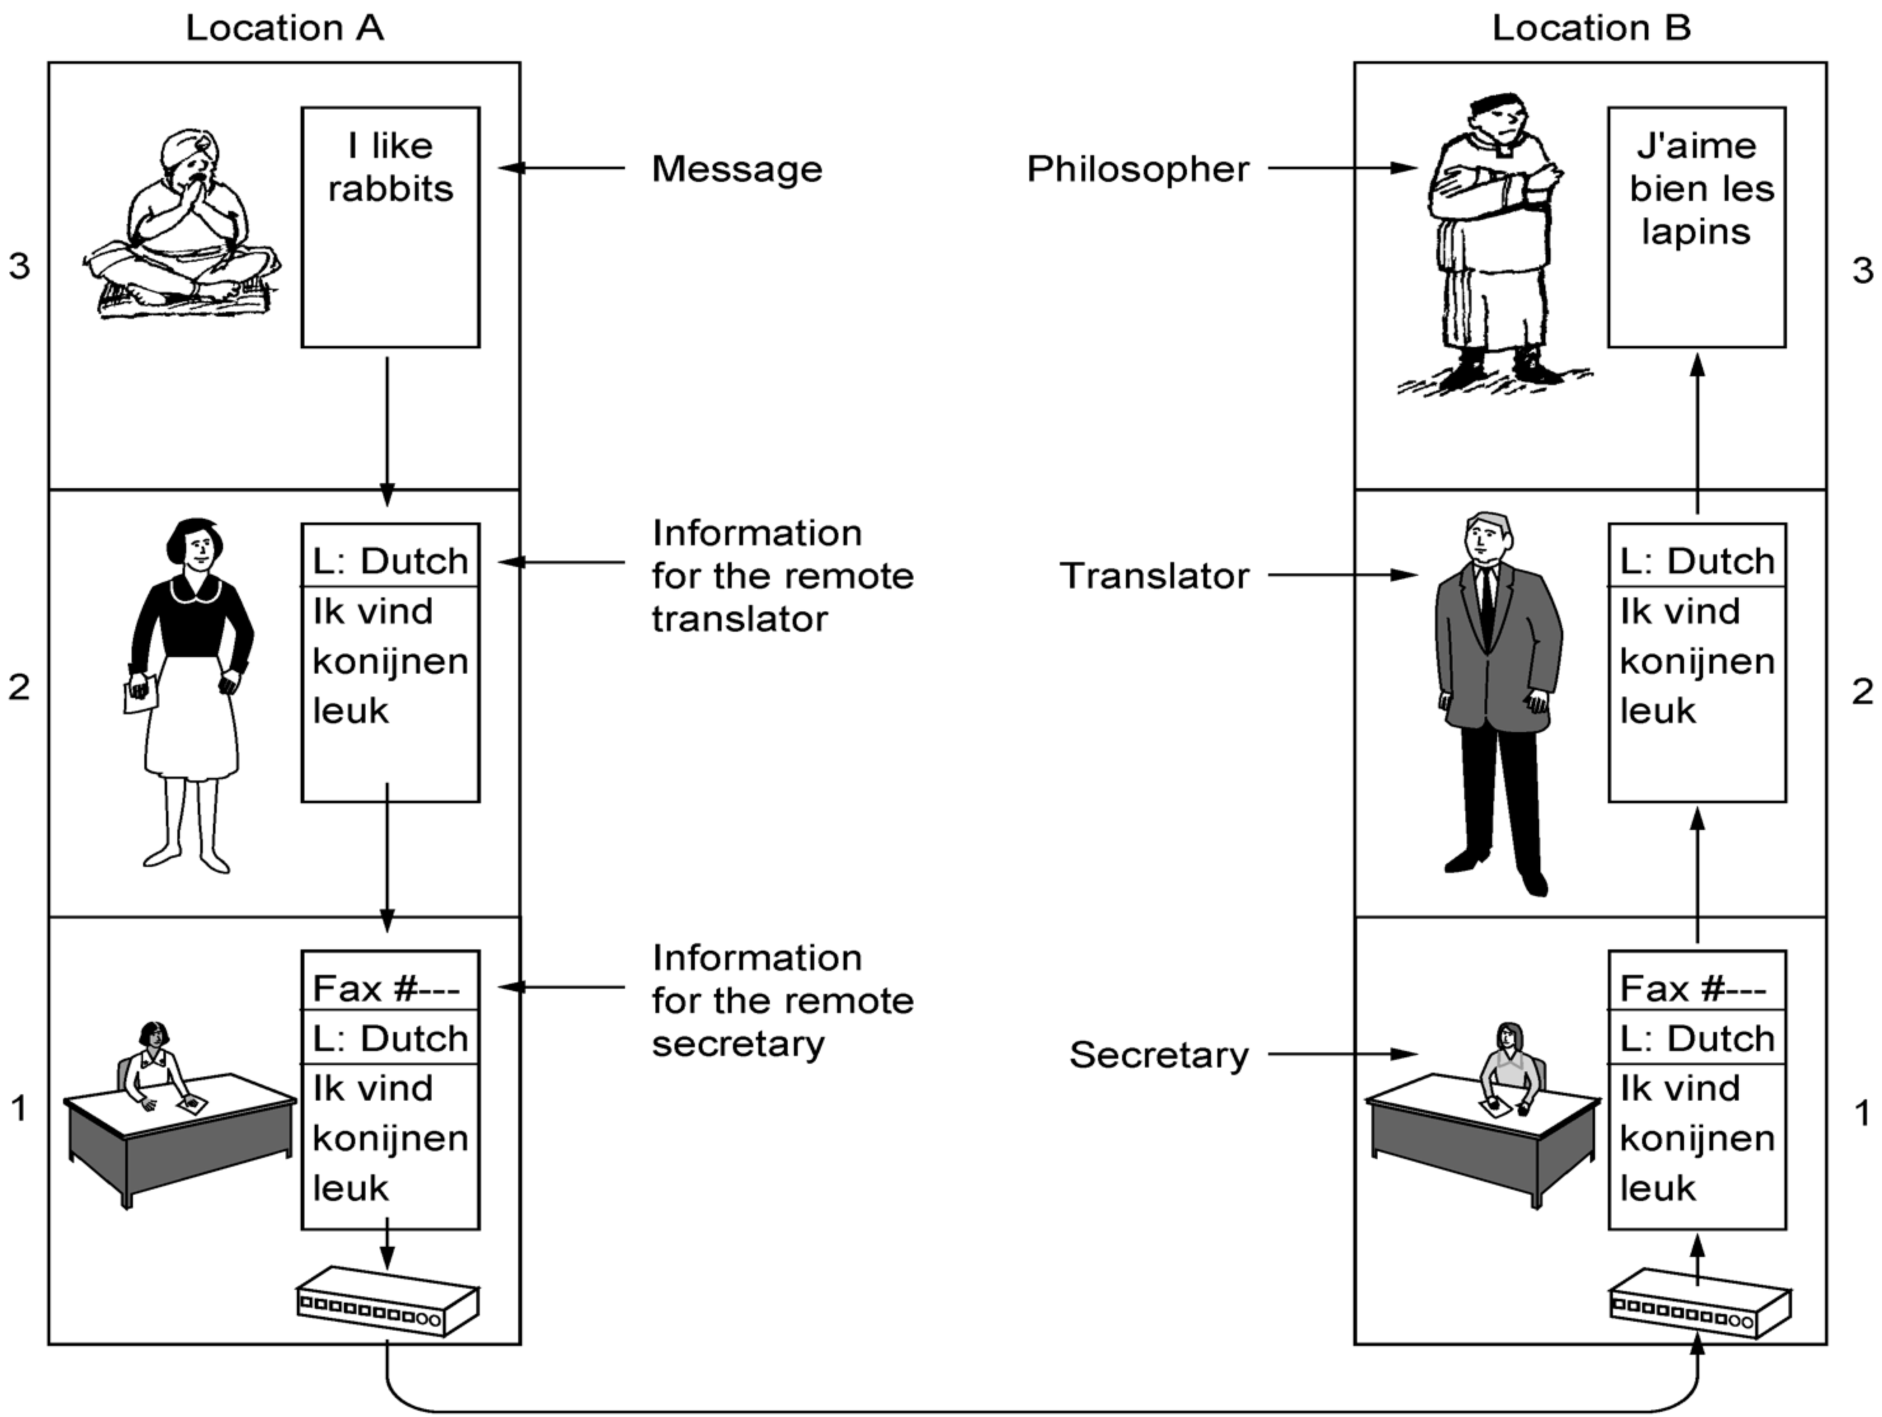
Location A
Location B
3
I like rabbits
Message
Philosopher
J'aime bien les lapins
3
2
L: Dutch
Ik vind konijnen leuk
Information for the remote translator
Translator
L: Dutch
Ik vind konijnen leuk
2
1
Fax #---
L: Dutch
Ik vind konijnen leuk
Information for the remote secretary
Secretary
Fax #---
L: Dutch
Ik vind konijnen leuk
1

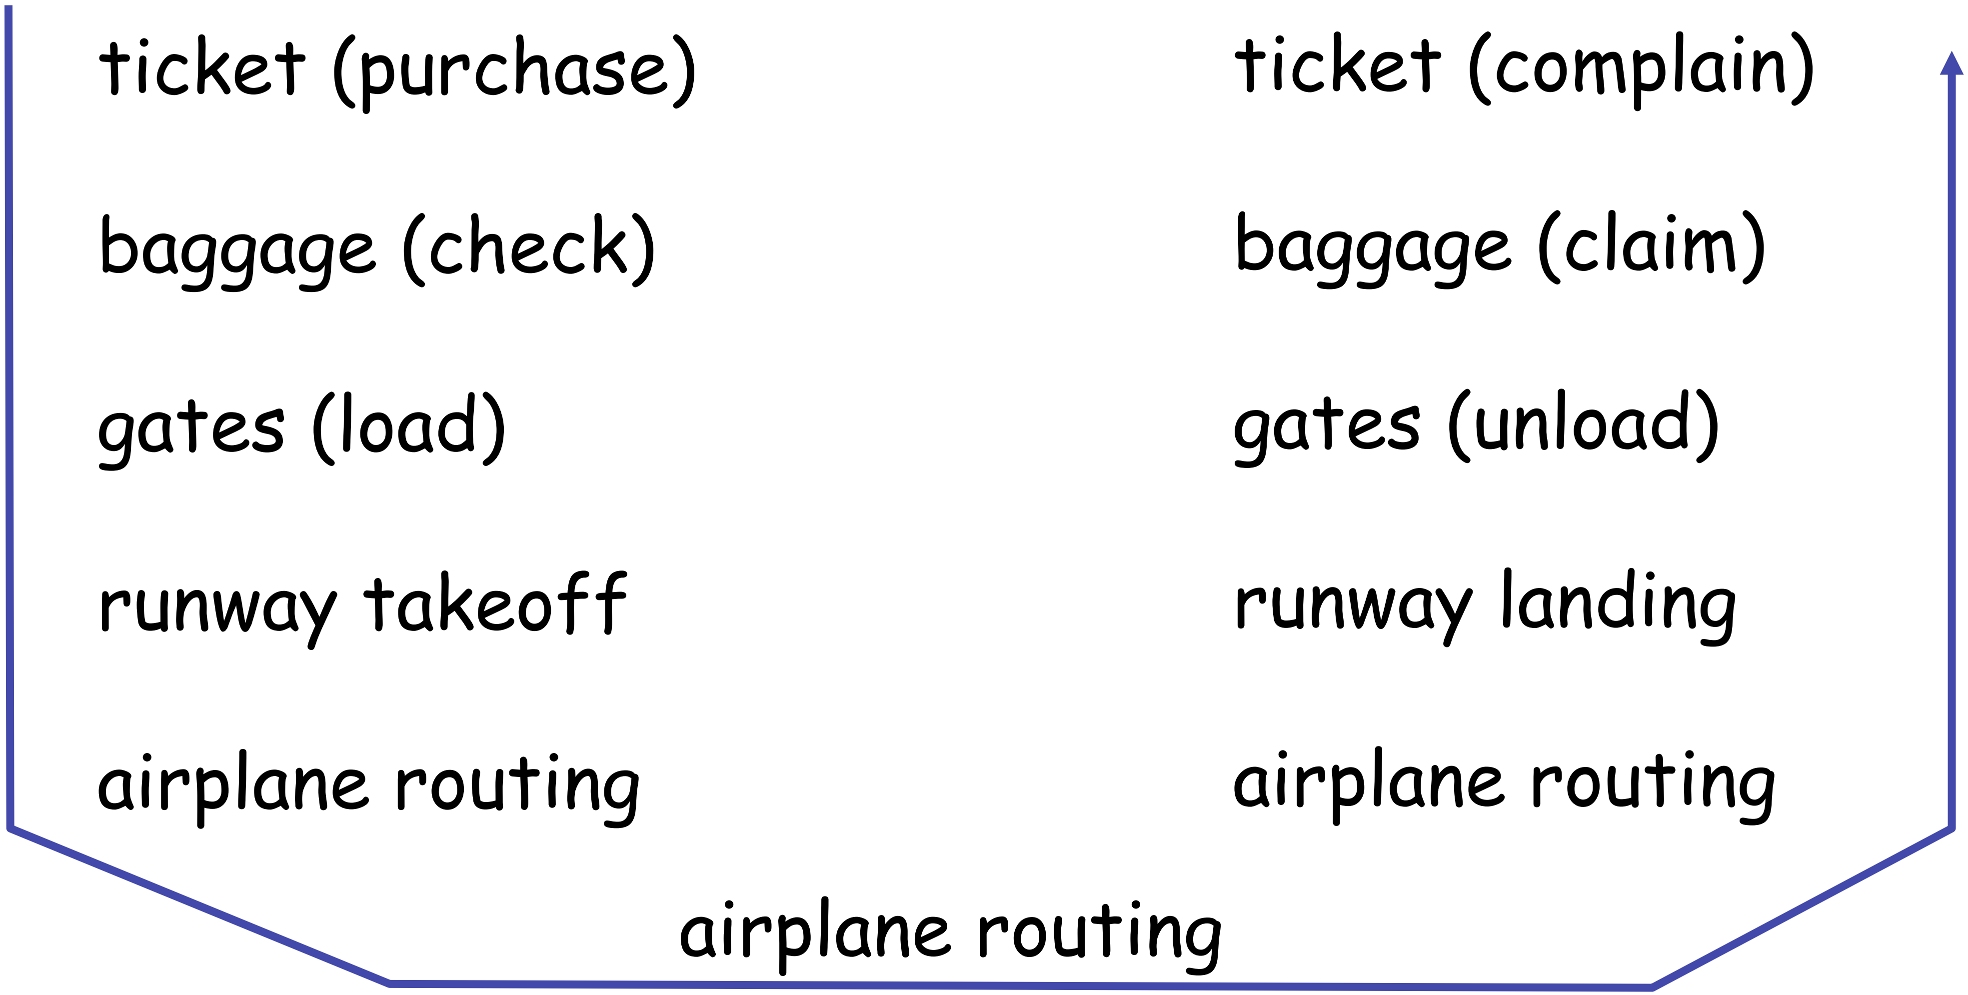

A series of steps!

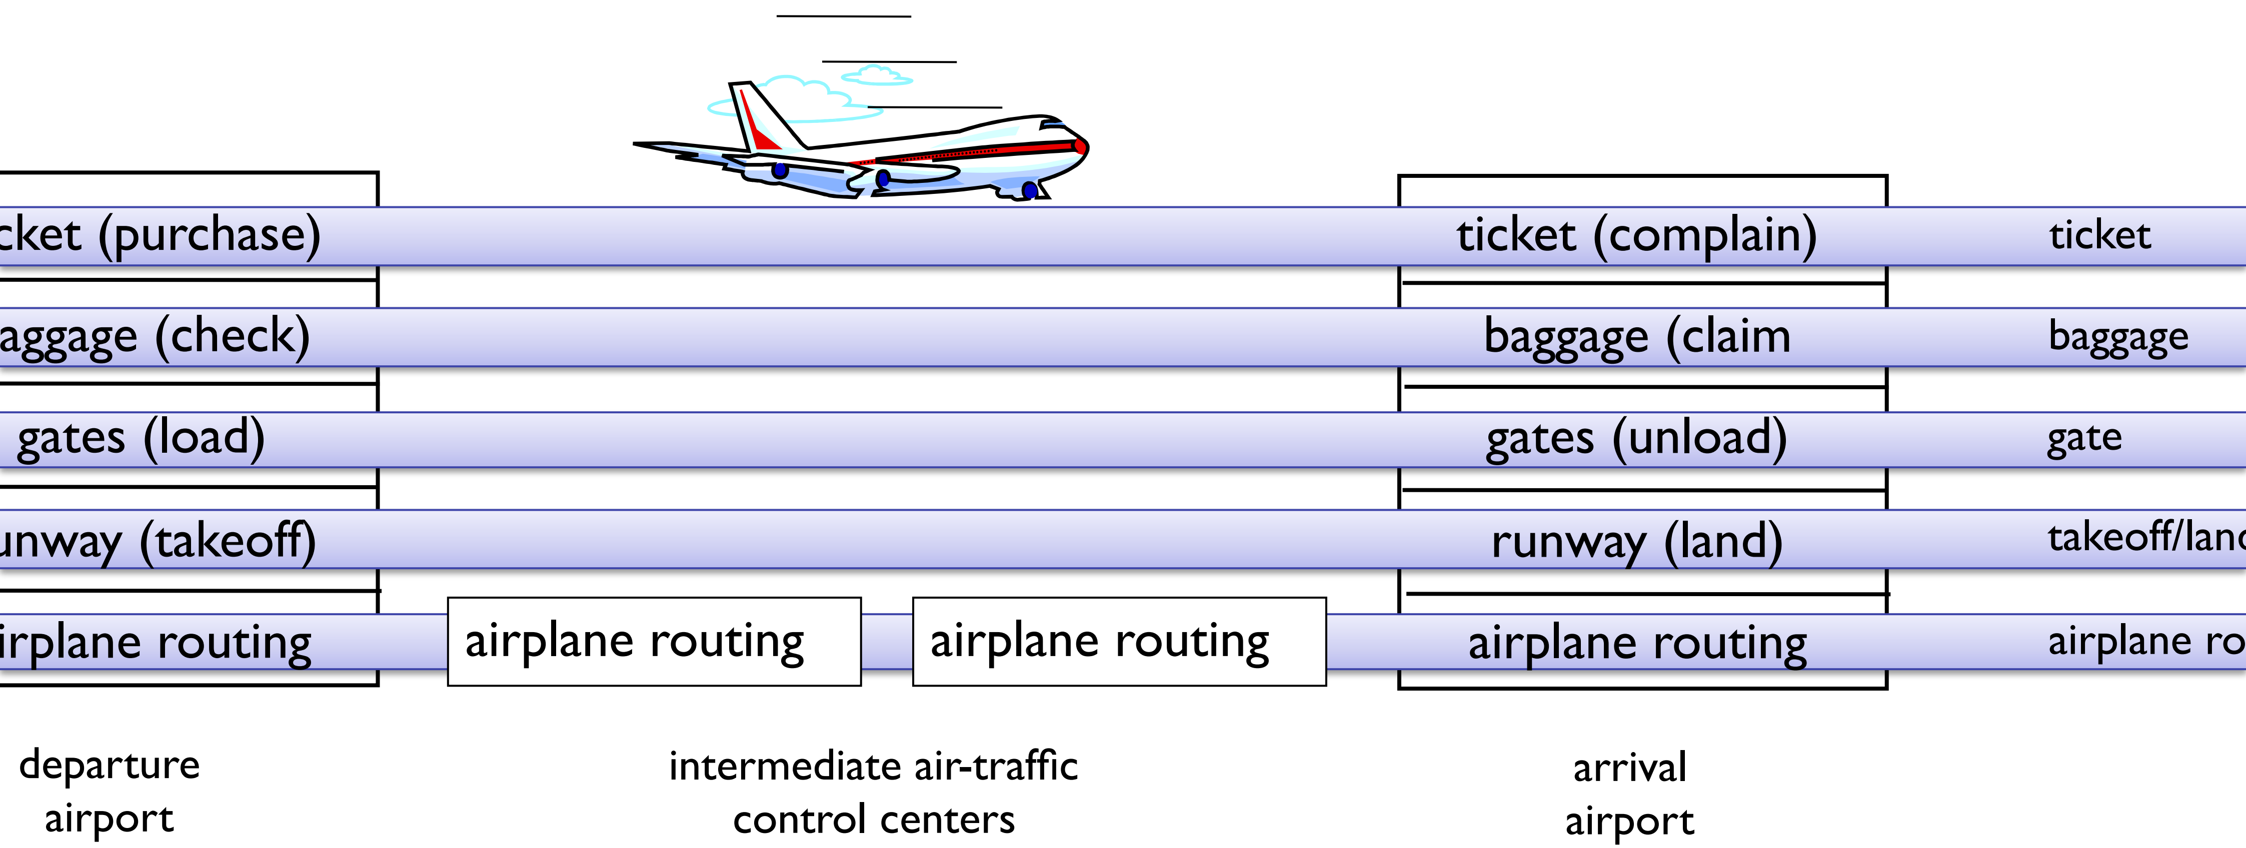

Layers: each layer implements a service

- via its own internal-layer actions
- relying on services provided by layer below

Dealing with complex systems:

- explicit structure allows identification, relationship of complex system pieces
  - layered reference model for discussion
- modularization eases maintenance, updating of system
  - change of implementation of layer service transparent to rest of system
  - e.g., change in gate procedure doesn't affect rest of system
- layering considered harmful?

- application: supporting network applications
  - FTP, SMTP, HTTP
- transport: process-process data transfer
  - TCP, UDP
- network: routing of datagrams from source to destination
  - IP, routing protocols
- link: data transfer between neighboring network elements
  - PPP, Ethernet

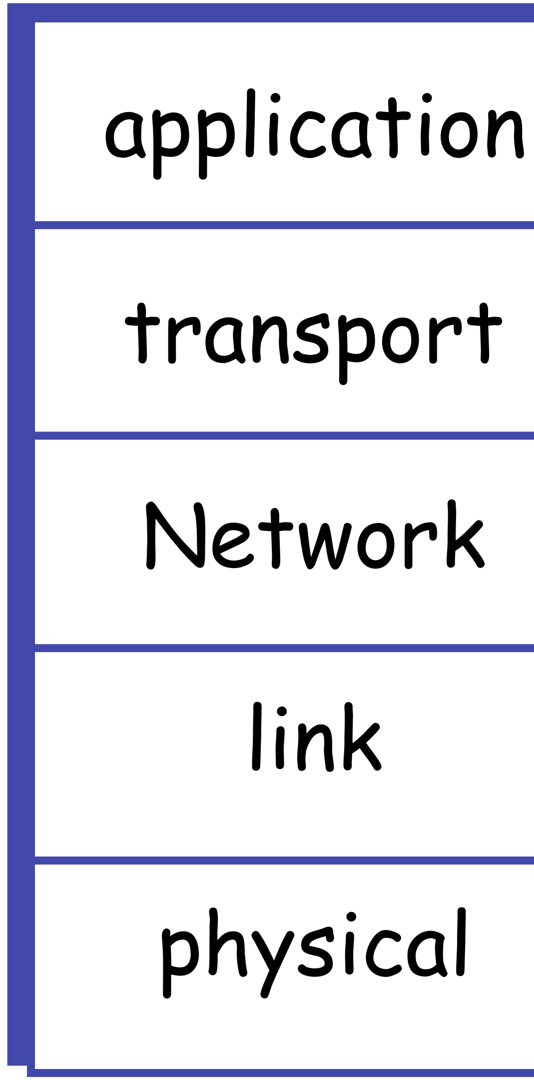

- presentation: allow applications to interpret meaning of data, e.g., encryption, compression, machine-specific conventions
- *session:* synchronization, checkpointing, recovery of data exchange
- Internet stack "missing" these layers!
  - these services, *if needed*, must be implemented in application
  - needed?

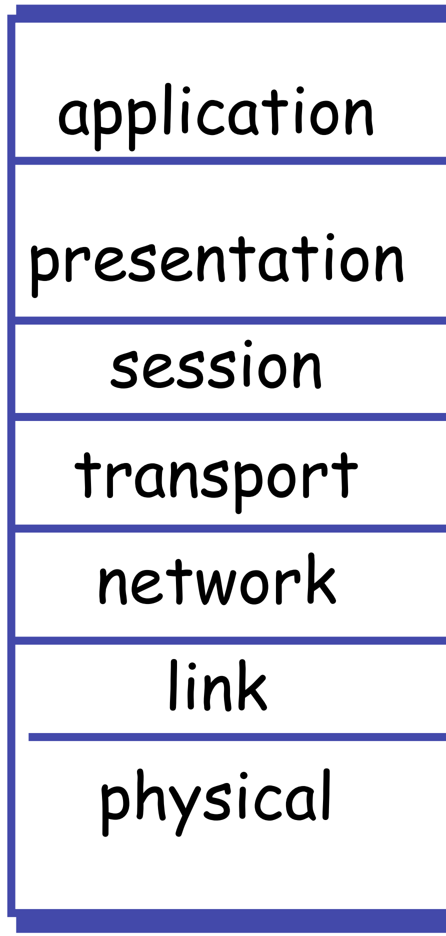

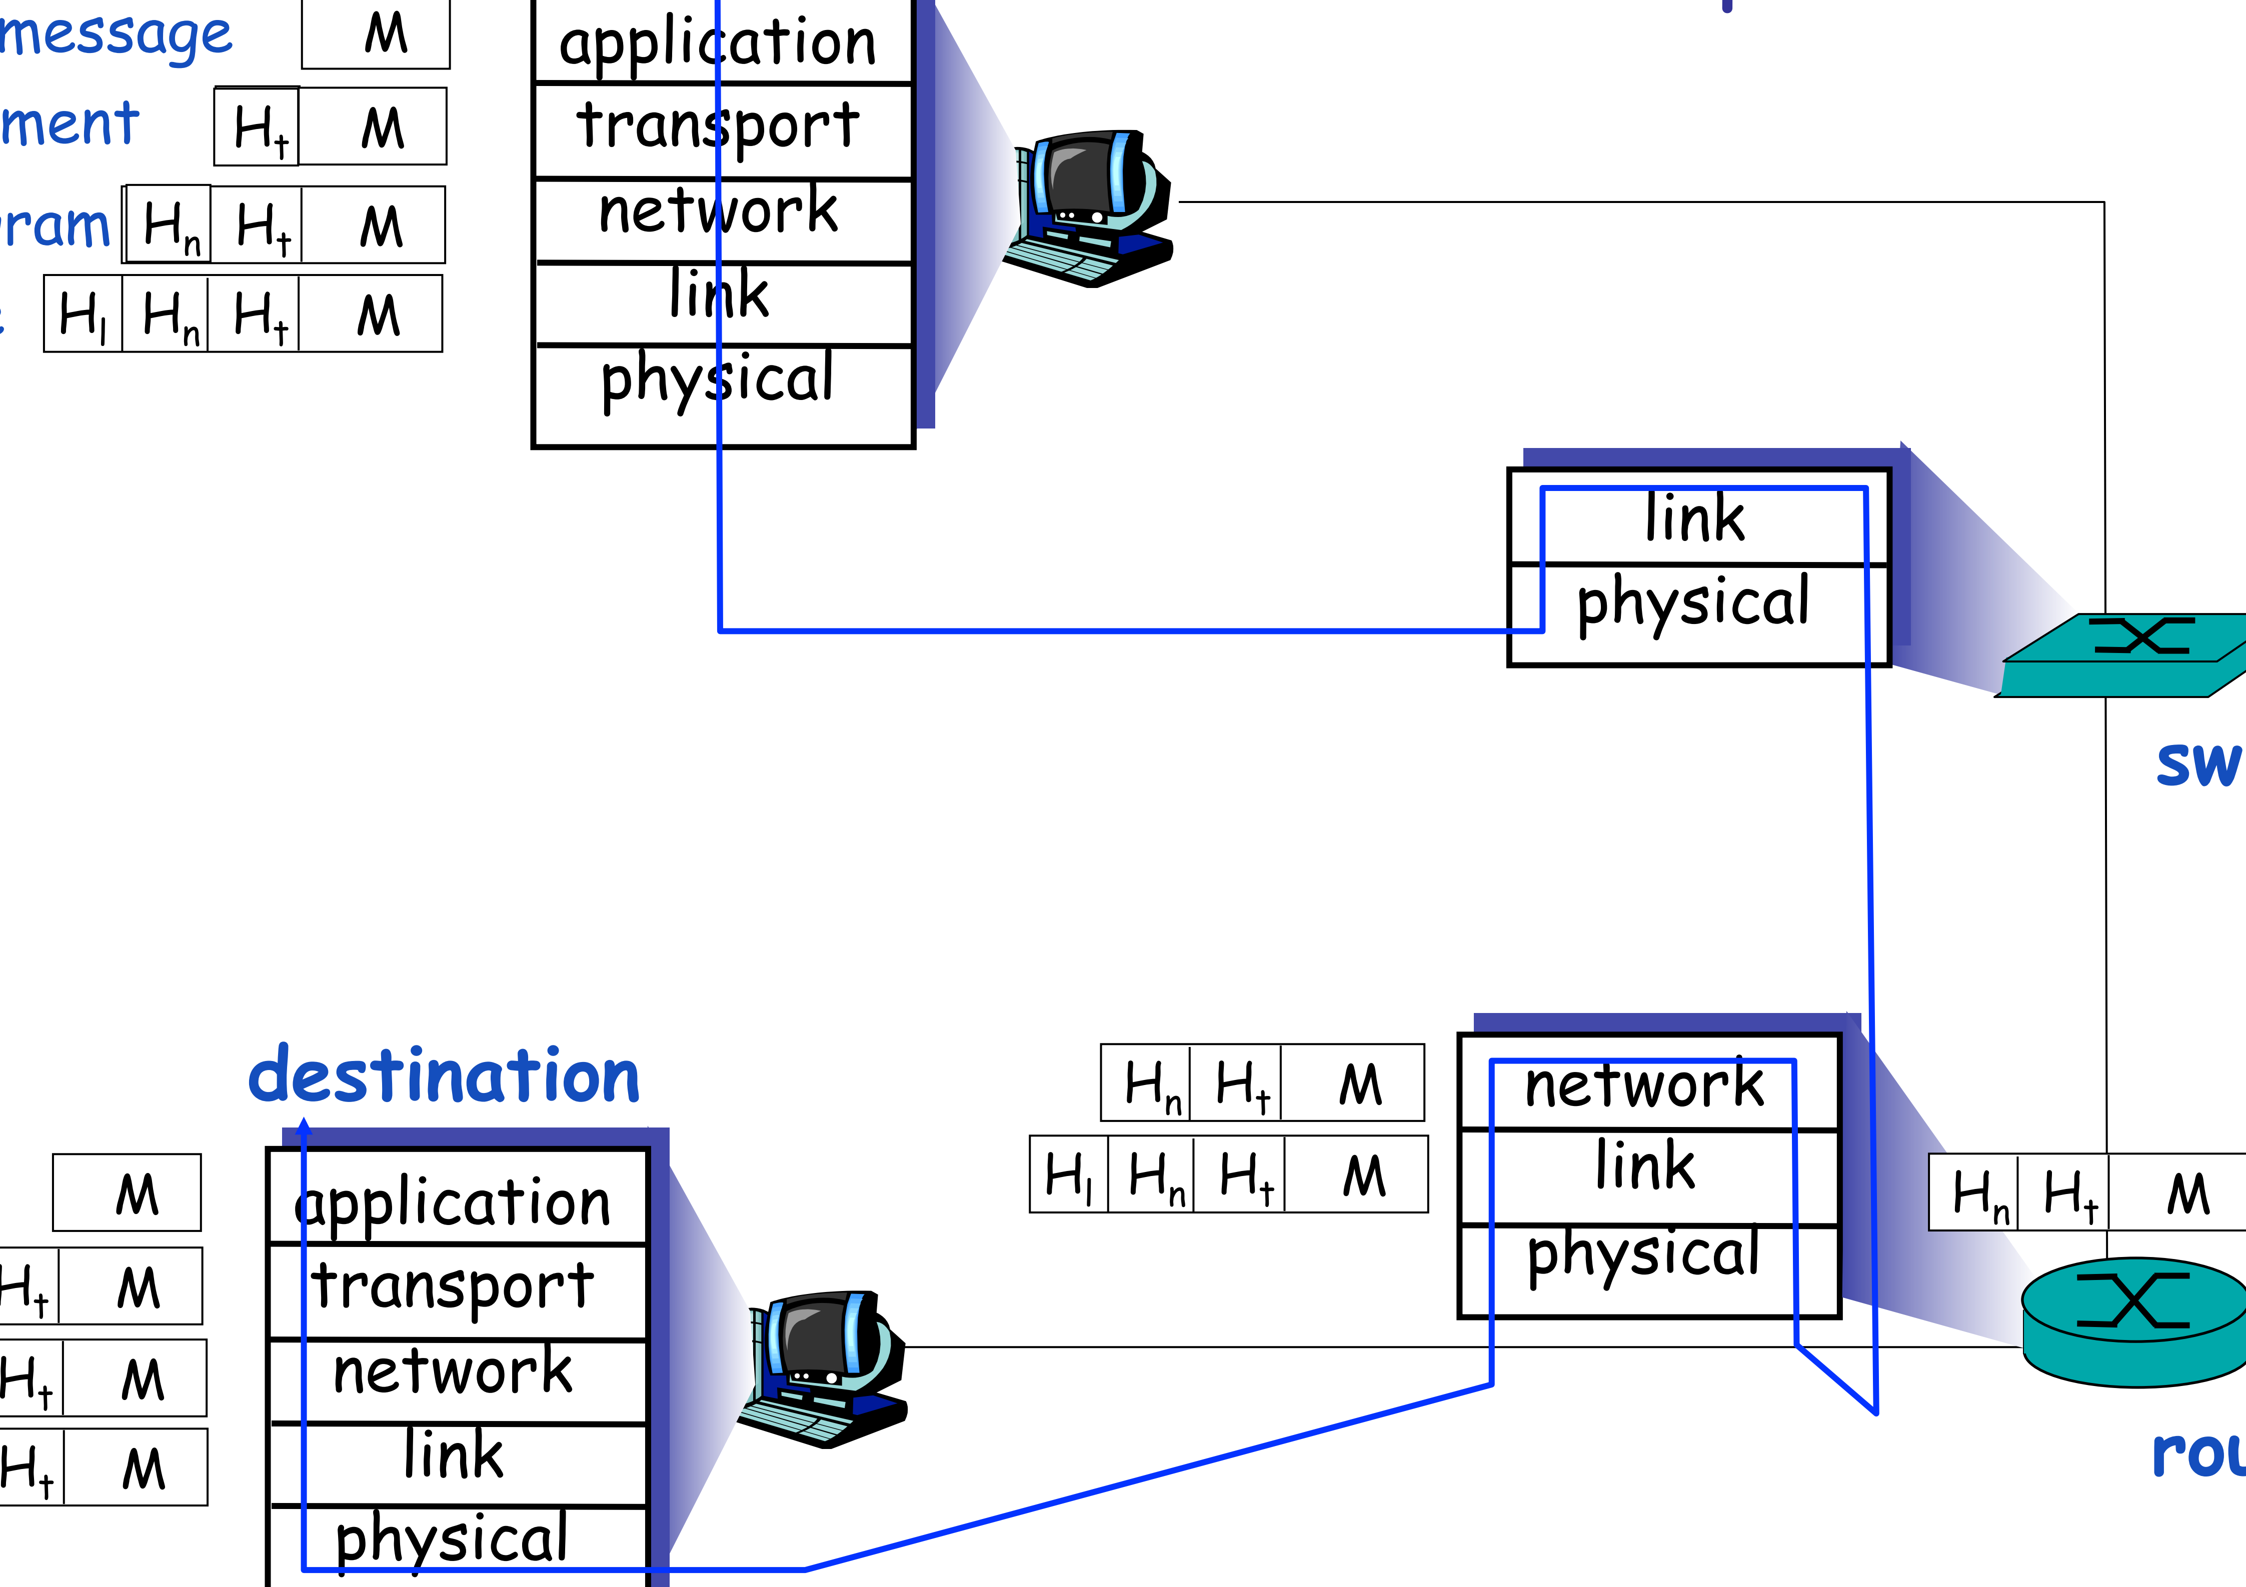

message
M
ment
Ht M
ram
Hn Ht M
Hl Hn Ht M
application
transport
network
link
physical
link
physical
sw
destination
M
Ht M
Ht M
Ht M
application
transport
network
link
physical
Hn Ht M
Hl Hn Ht M
network
link
physical
Hn Ht M
rou