Network Technologies (TCP/IP Suite)

Umar Kalim
Dept. of Communication Systems Engineering

umar.kalim@niit.edu.pk
http://www.niit.edu.pk/~umarkalim

15/03/2007

Outline

- A Detailed FTP Example
- Layering
- Packet Switching and Circuit Switching
- Some terms
  - Data rate, “Bandwidth” and “throughput”
  - Propagation delay
  - Packet, header, address
  - Bandwidth-delay product, RTT

Example: FTP over the Internet Using TCP/IP and Ethernet

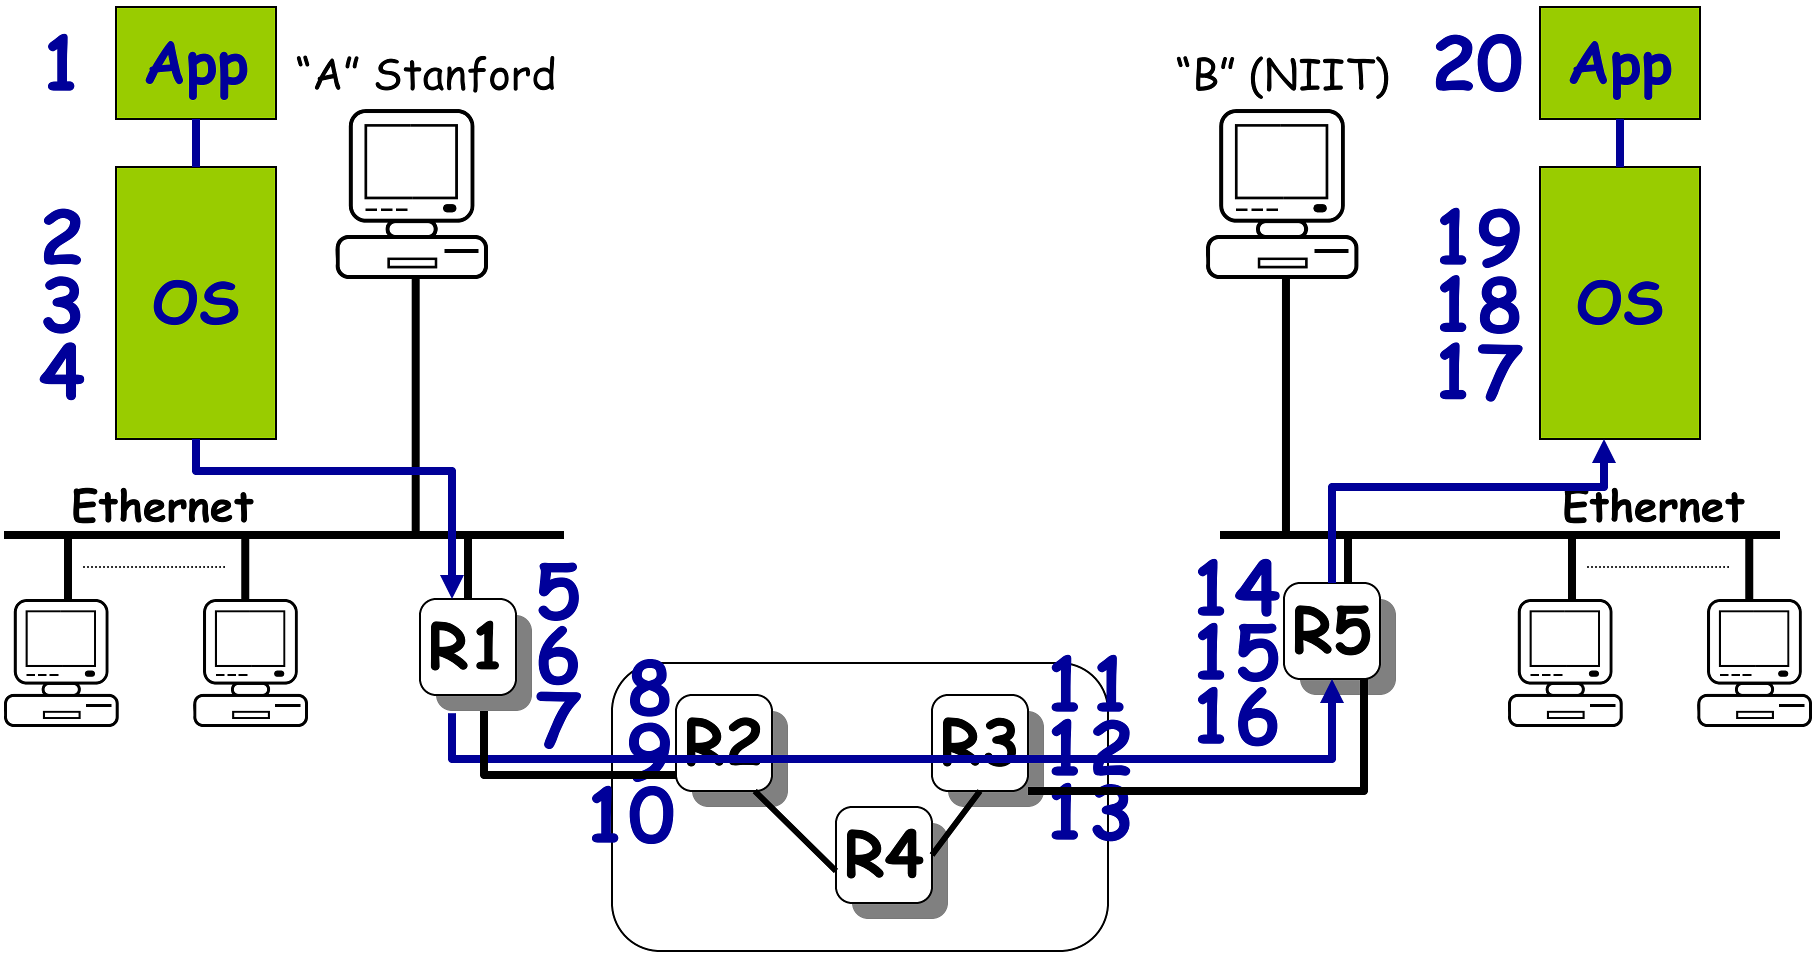

In the sending host

1. **Application-Programming Interface (API)**
   – Application requests TCP connection with "B"
2. **Transmission Control Protocol (TCP)**
   – Creates TCP "Connection setup" packet
   – TCP requests IP packet to be sent to "B"

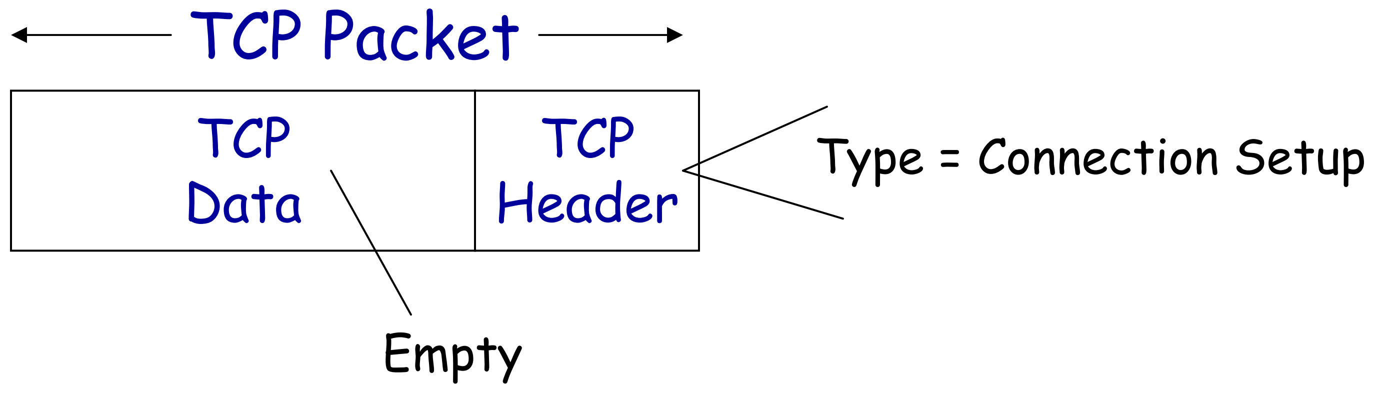

In the sending host (2)

3. Internet Protocol (IP)

- Creates IP packet with correct addresses.
- IP requests packet to be sent to router.

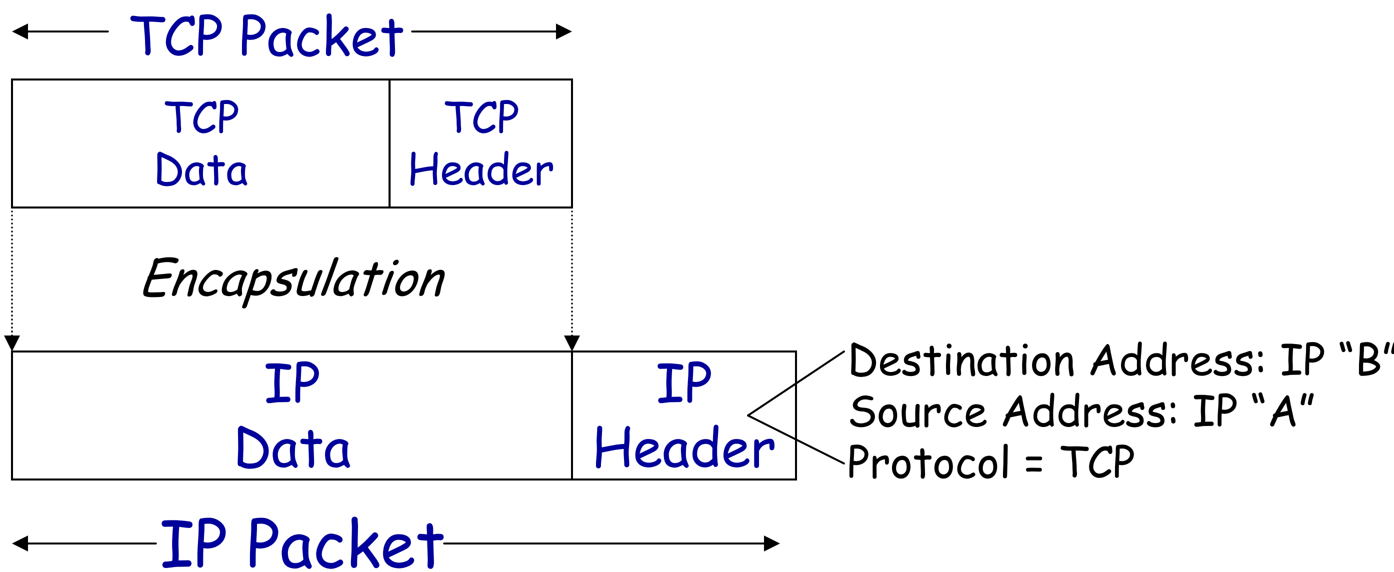

In the sending host (3)

4. Link ("MAC" or Ethernet) Protocol

- Creates MAC frame with Frame Check Sequence (FCS).
- Wait for Access to the line.
- MAC requests PHY to send each bit of the frame.

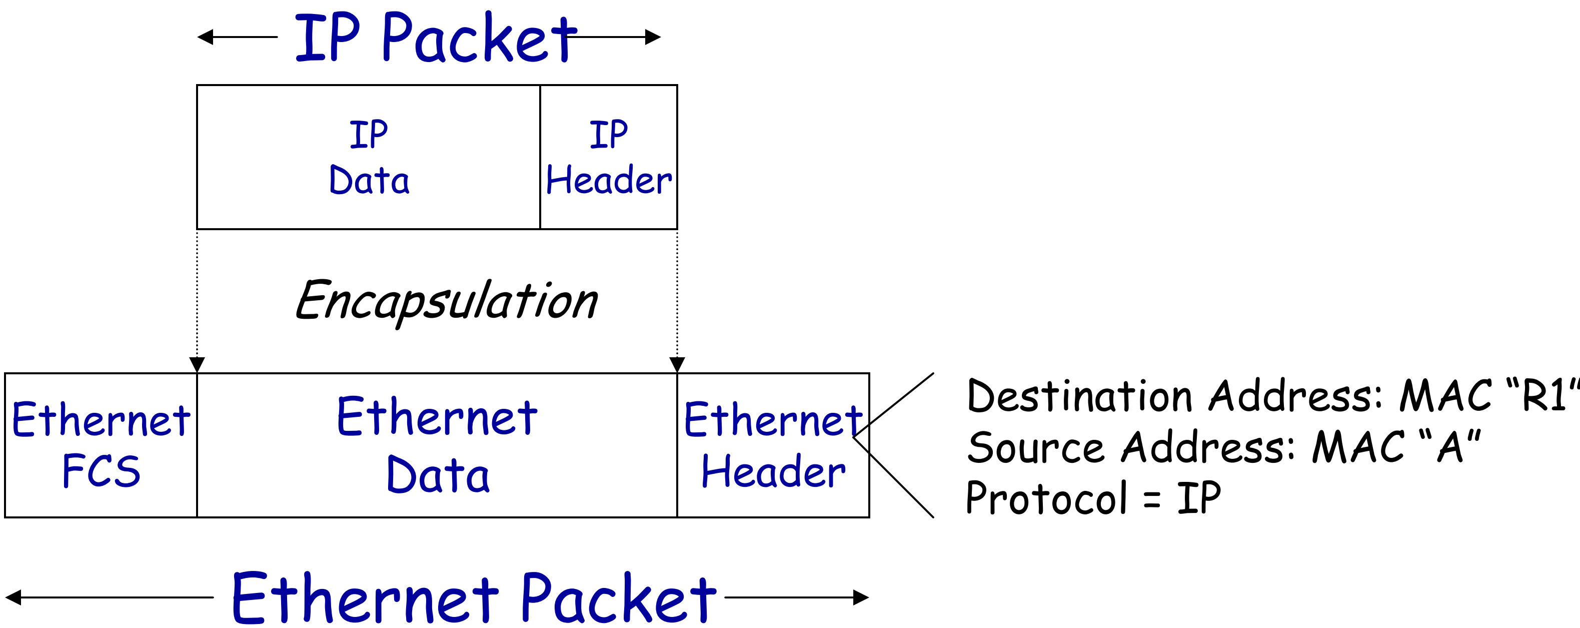

In Router R1

5. Link ("MAC" or Ethernet) Protocol

- Accept MAC frame, check address and Frame Check Sequence (FCS).
- Pass data to IP Protocol.

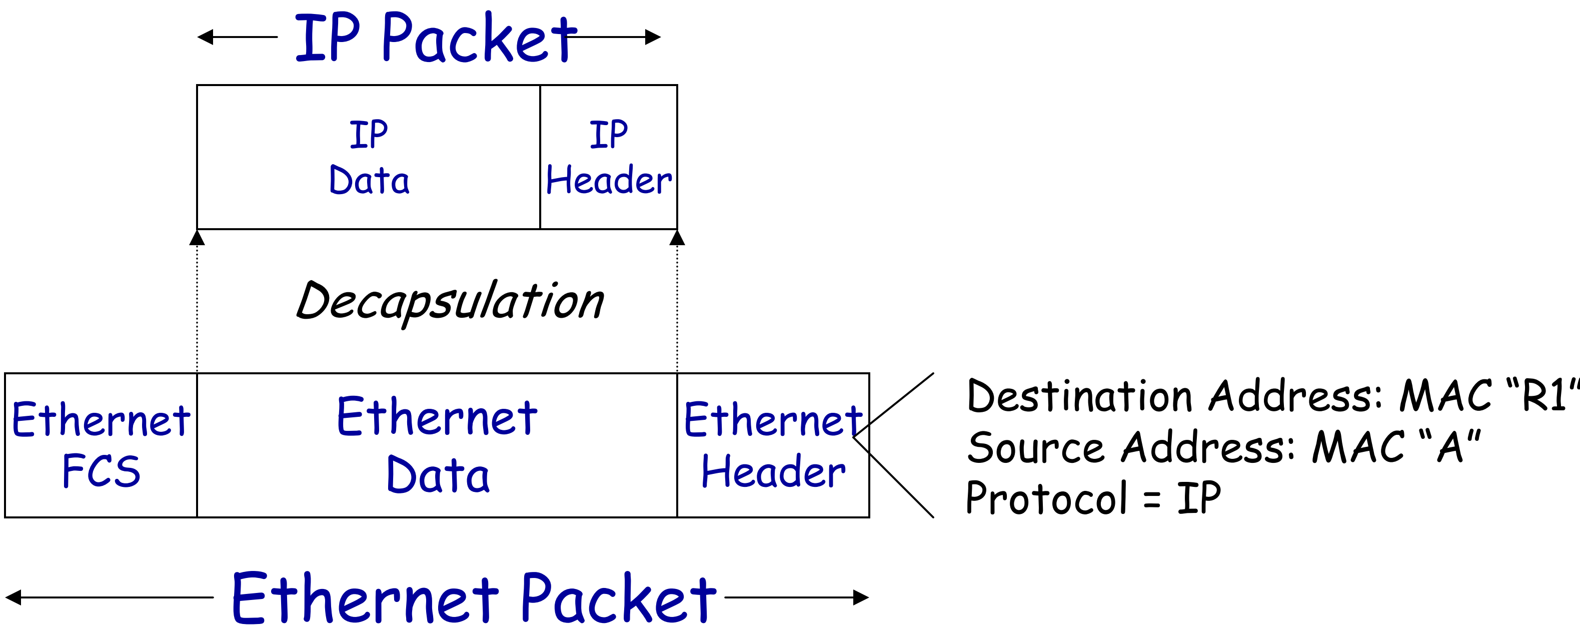

In Router R1

6. Internet Protocol (IP)

- Use IP destination address to decide where to send packet next ("next-hop routing").
- Request Link Protocol to transmit packet.

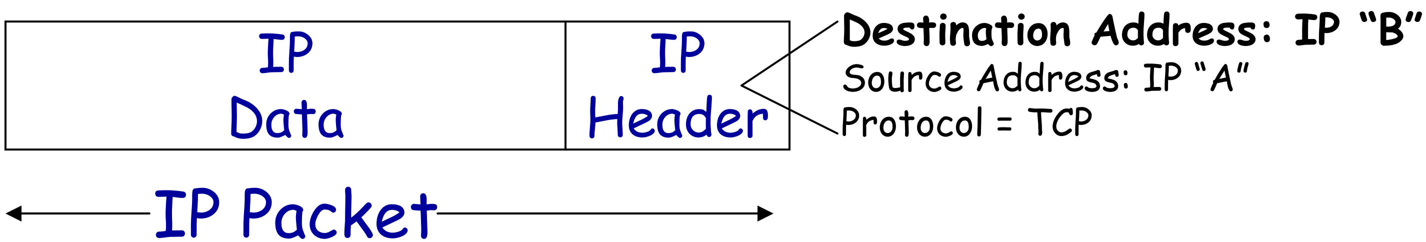

In Router R1

7. Link ("MAC" or Ethernet) Protocol

- Creates MAC frame with Frame Check Sequence (FCS).
- Wait for Access to the line.
- MAC requests PHY to send each bit of the frame.

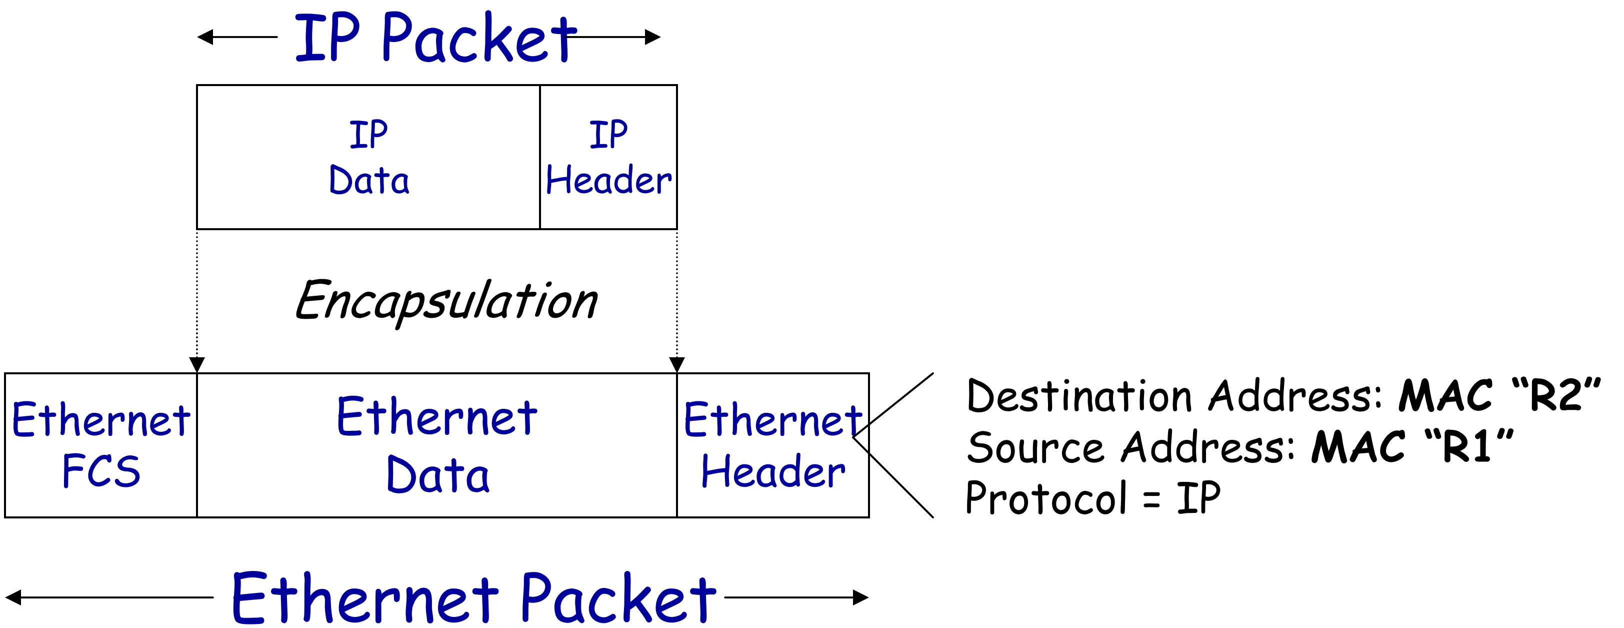

In Router R5

16. Link ("MAC" or Ethernet) Protocol

- Creates MAC frame with Frame Check Sequence (FCS).
- Wait for Access to the line.
- MAC requests PHY to send each bit of the frame.

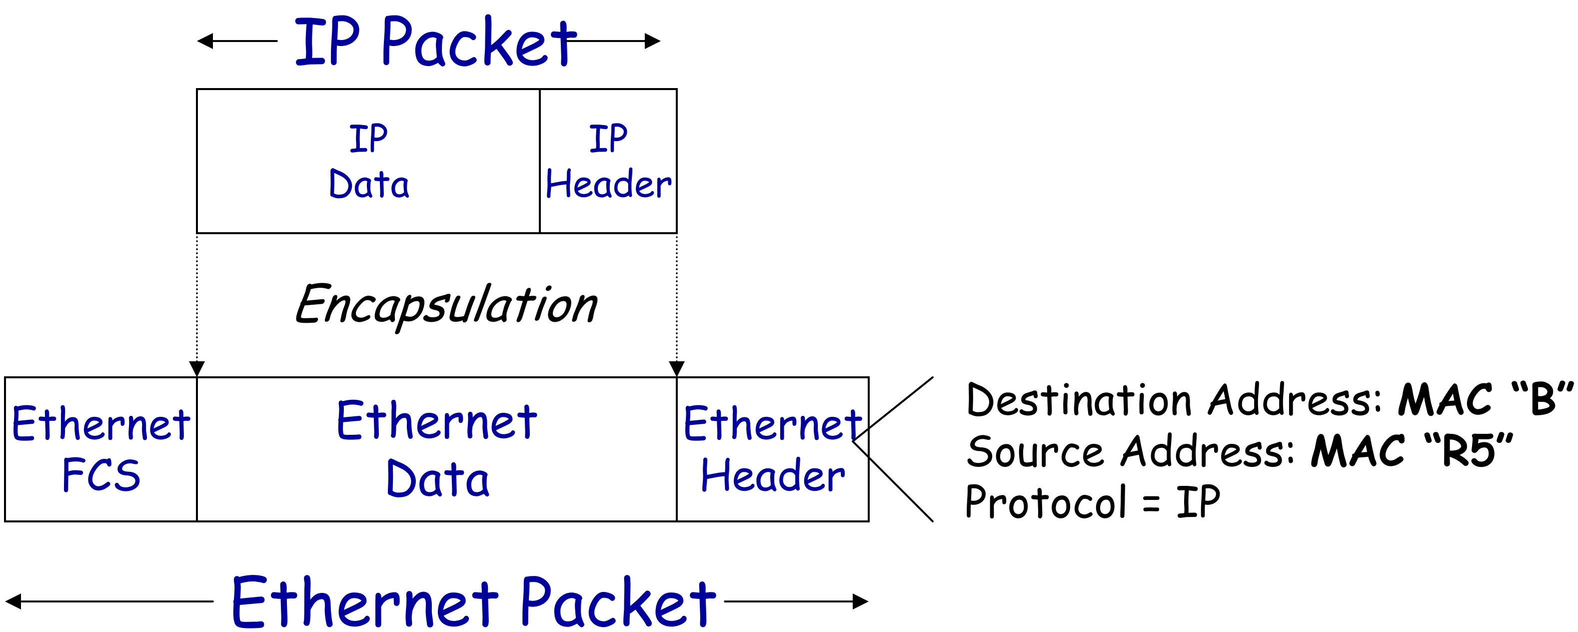

In the receiving host

17. Link ("MAC" or Ethernet) Protocol

- Accept MAC frame, check address and Frame Check Sequence (FCS).
- Pass data to IP Protocol.

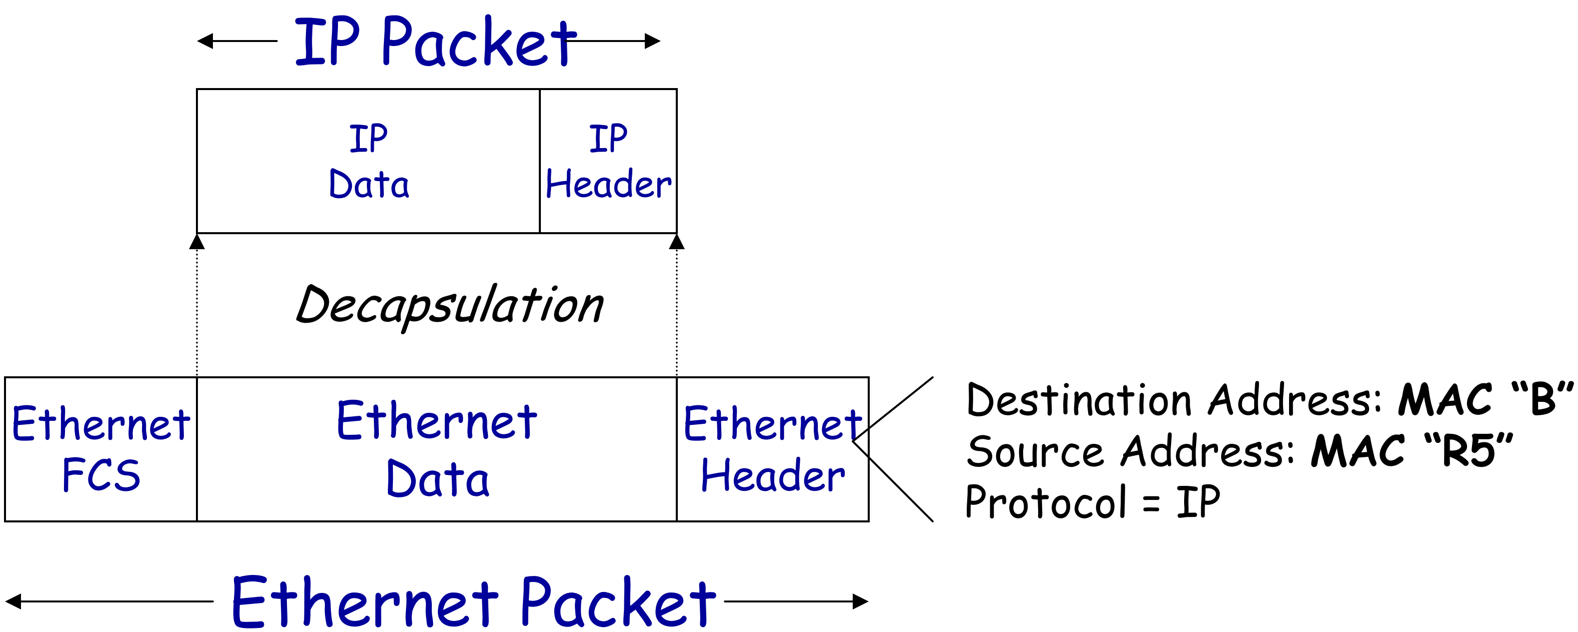

In the receiving host (2)

18. Internet Protocol (IP)

- Verify IP address.
- Extract/decapsulate TCP packet from IP packet.
- Pass TCP packet to TCP Protocol.

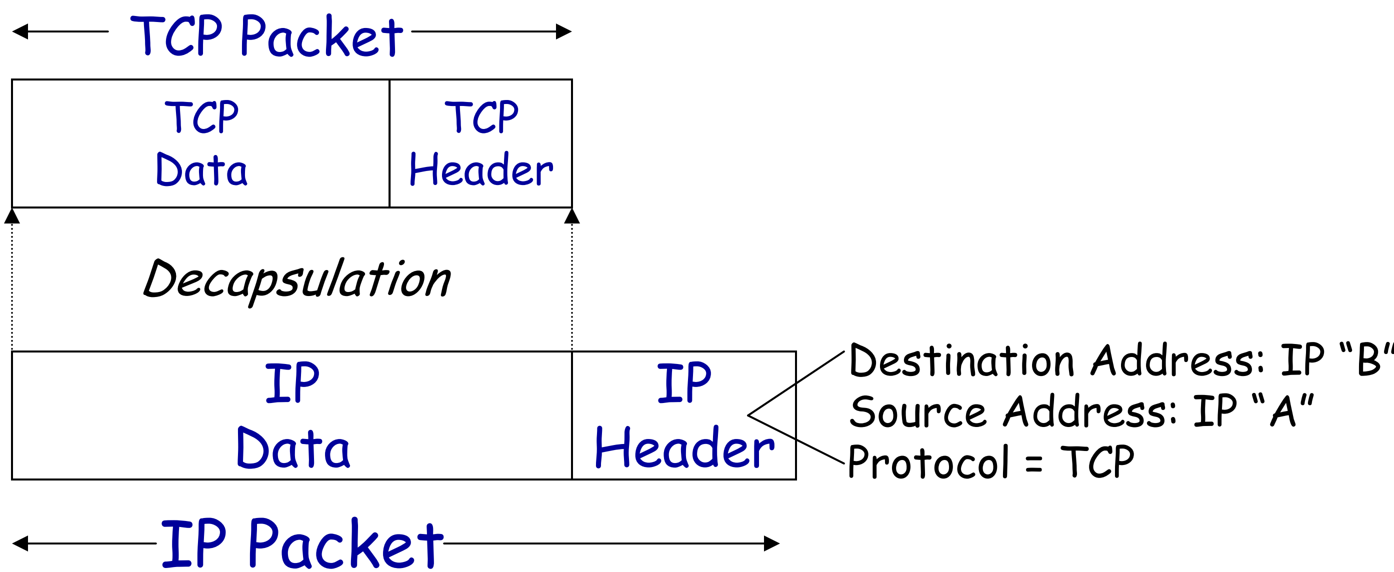

In the receiving host (3)

19. Transmission Control Protocol (TCP)

– Accepts TCP "Connection setup" packet

– Establishes connection by sending "Ack".

20. Application-Programming Interface (API)

– Application receives request for TCP connection with "A".

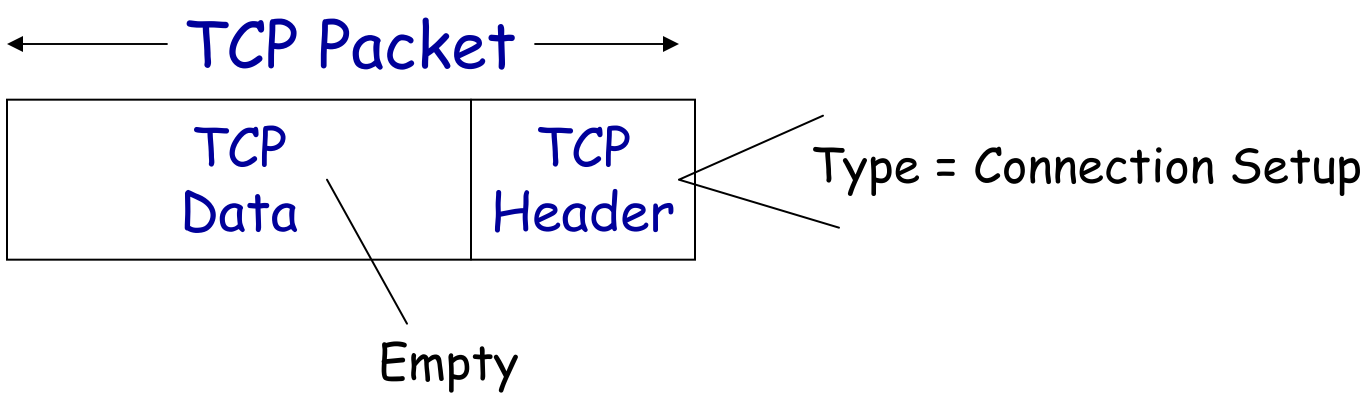

Outline

- A Detailed FTP Example
- Layering
- Packet Switching and Circuit Switching
- Some terms
  - Data rate, “Bandwidth” and “throughput”
  - Propagation delay
  - Packet, header, address
  - Bandwidth-delay product, RTT

Layering: The OSI Model

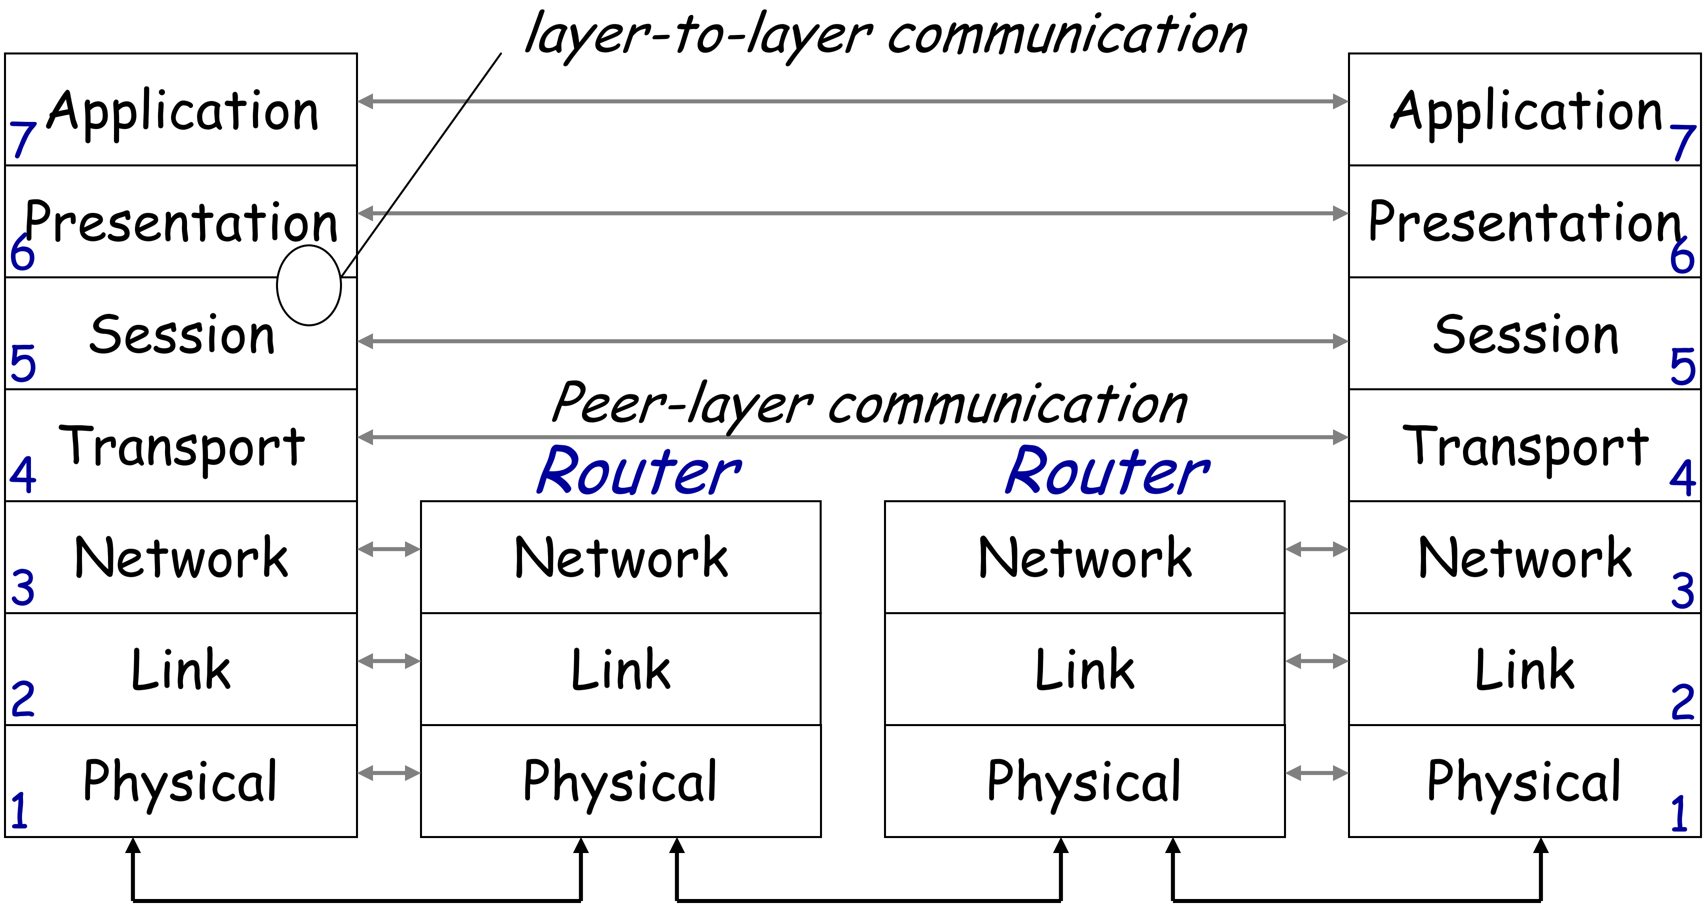

Layering: Our FTP Example

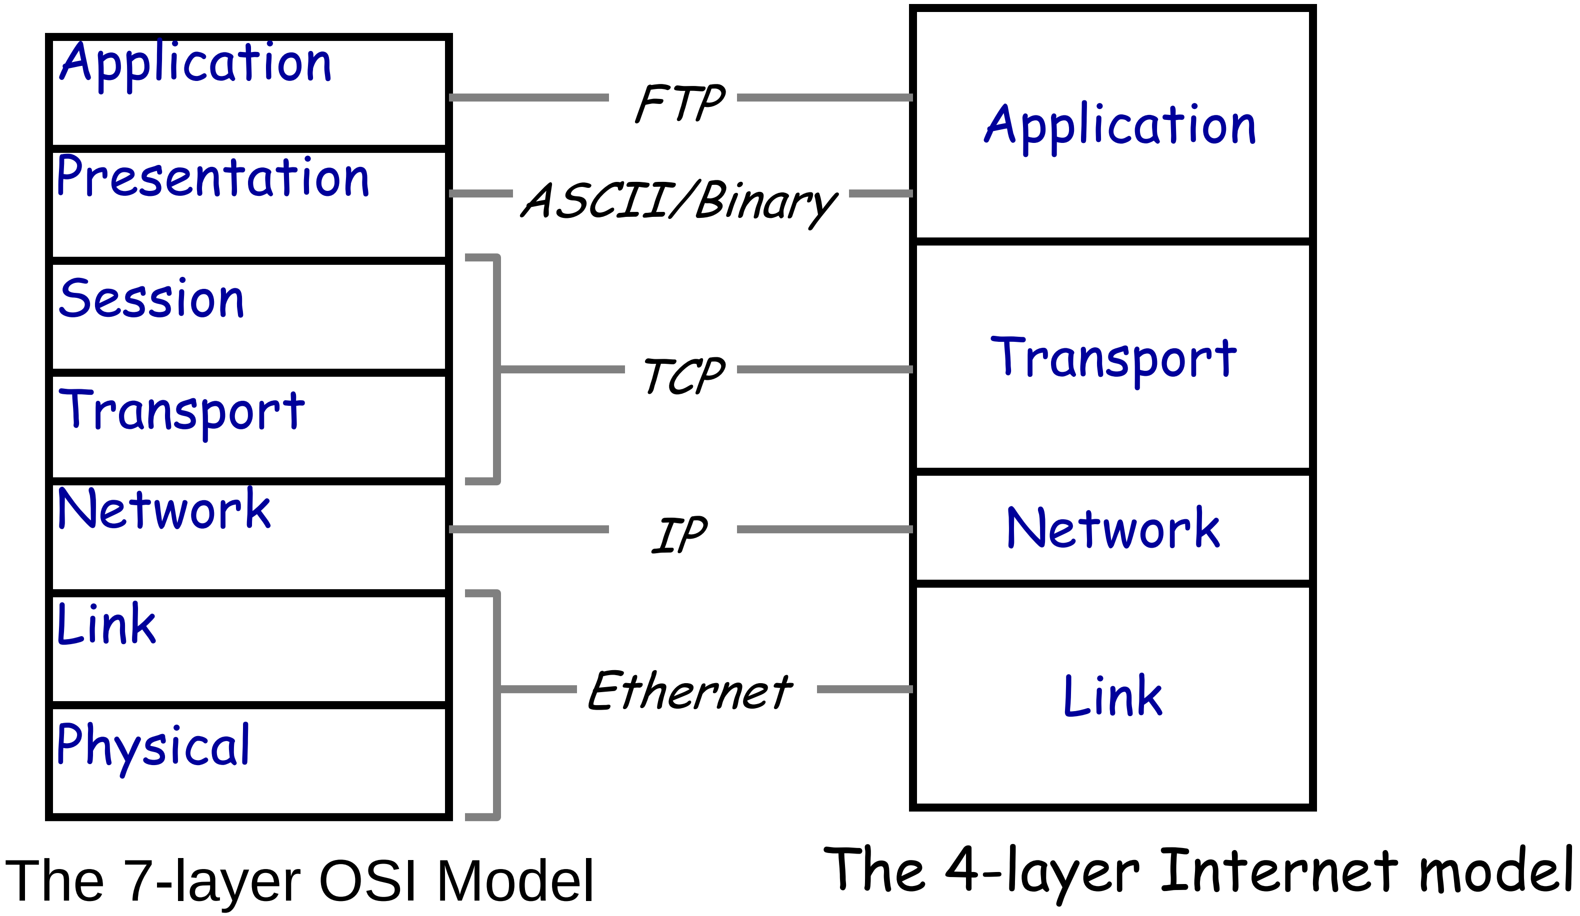

Outline

- A Detailed FTP Example
- Layering
- Packet Switching and Circuit Switching
- Some terms
  - Data rate, “Bandwidth” and “throughput”
  - Propagation delay
  - Packet, header, address
  - Bandwidth-delay product, RTT

Circuit Switching

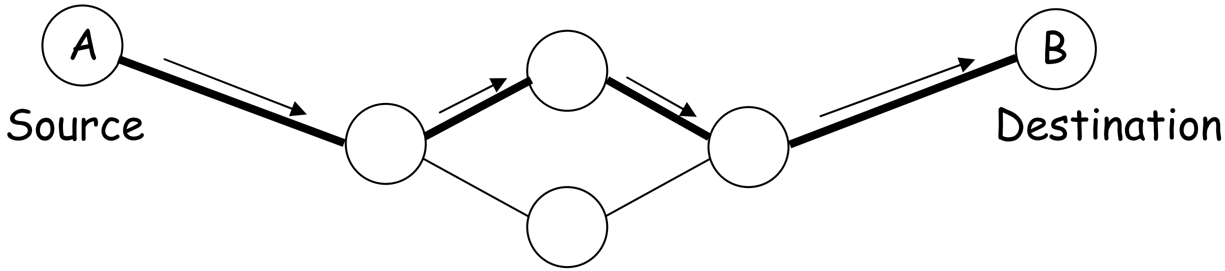

- It's the method used by the telephone network.
- A call has three phases:
  1. Establish circuit from end-to-end ("dialing"),
  2. Communicate,
  3. Close circuit ("tear down").
- Originally, a circuit was an end-to-end physical wire.
- Nowadays, a circuit is like a virtual private wire: each call has its own private, guaranteed data rate from end-to-end.

Circuit Switching

Telephone Network

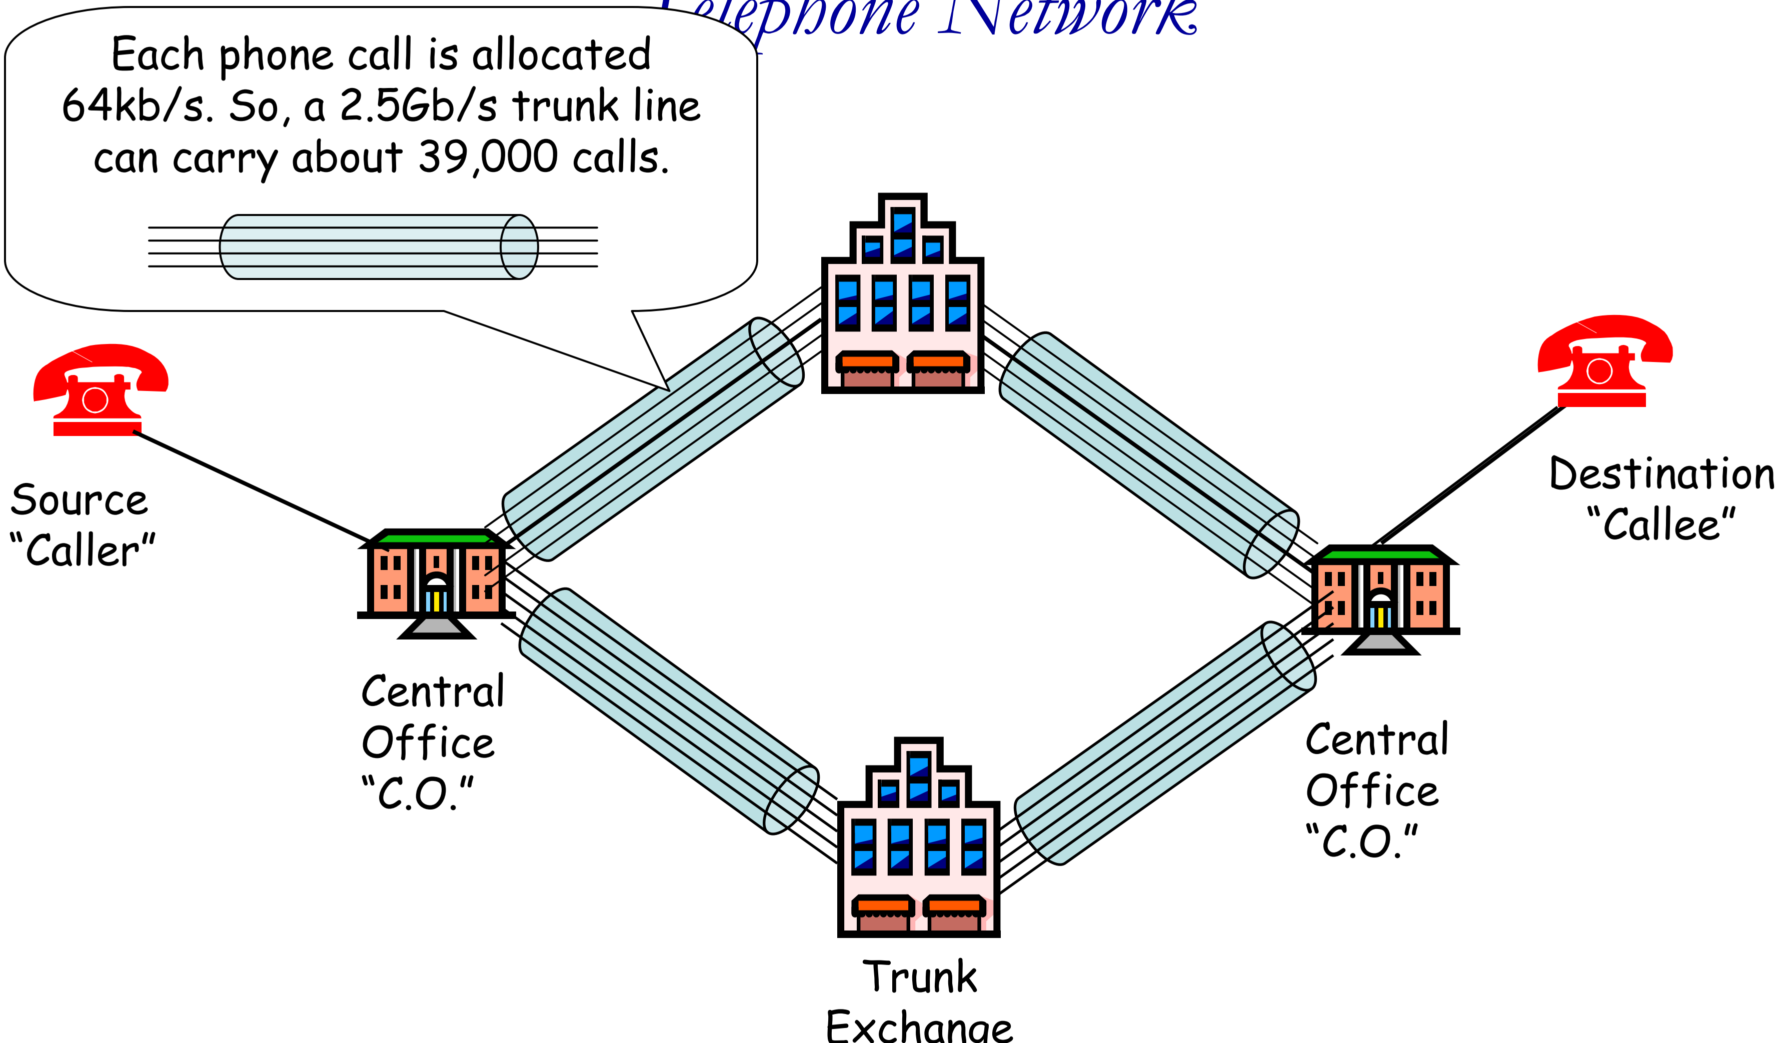

Packet Switching

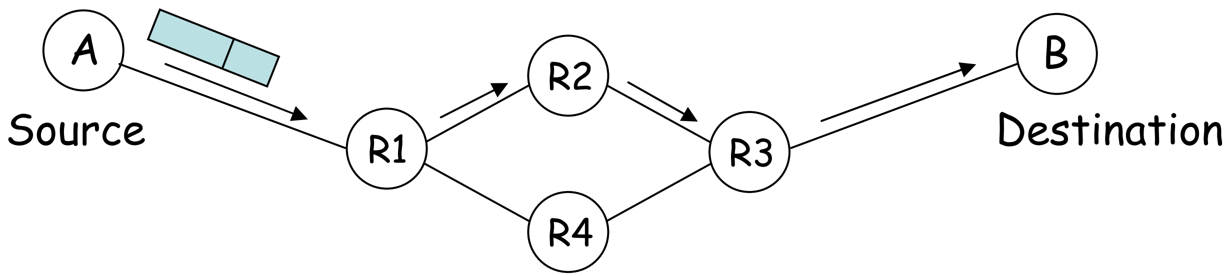

- It's the method used by the Internet.
- Each packet is individually routed packet-by-packet, using the router's local routing table.
- The routers maintain no per-flow state.
- Different packets may take different paths.
- Several packets may arrive for the same output link at the same time, therefore a packet switch has buffers.

Packet Switching

Simple router model

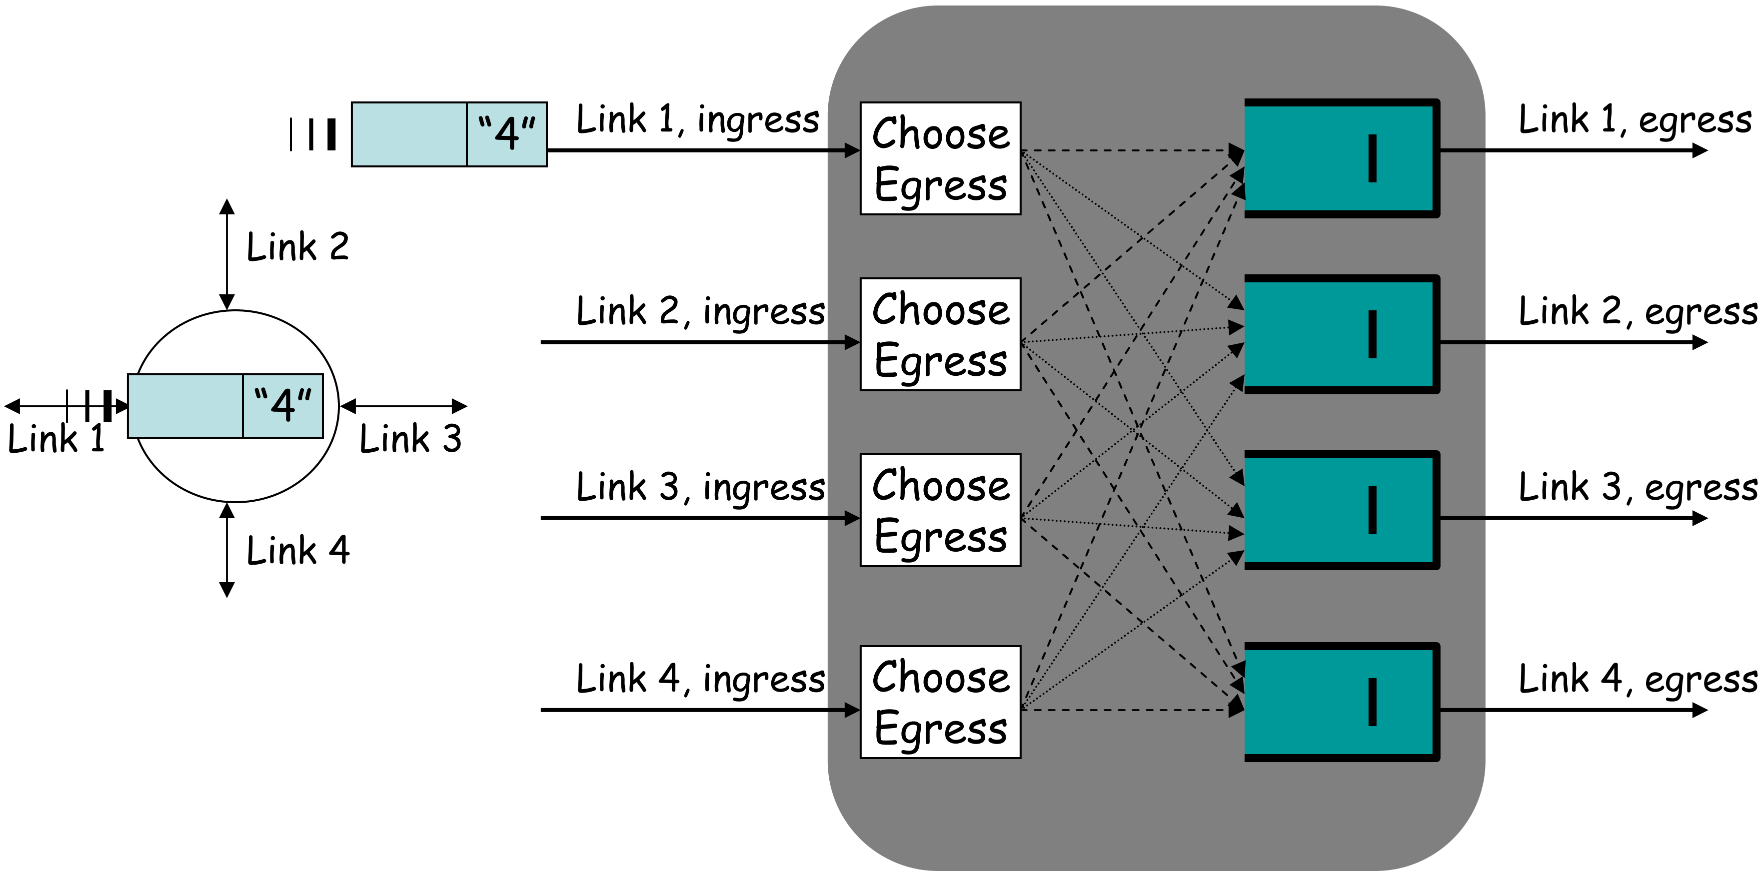

Statistical Multiplexing

Basic idea

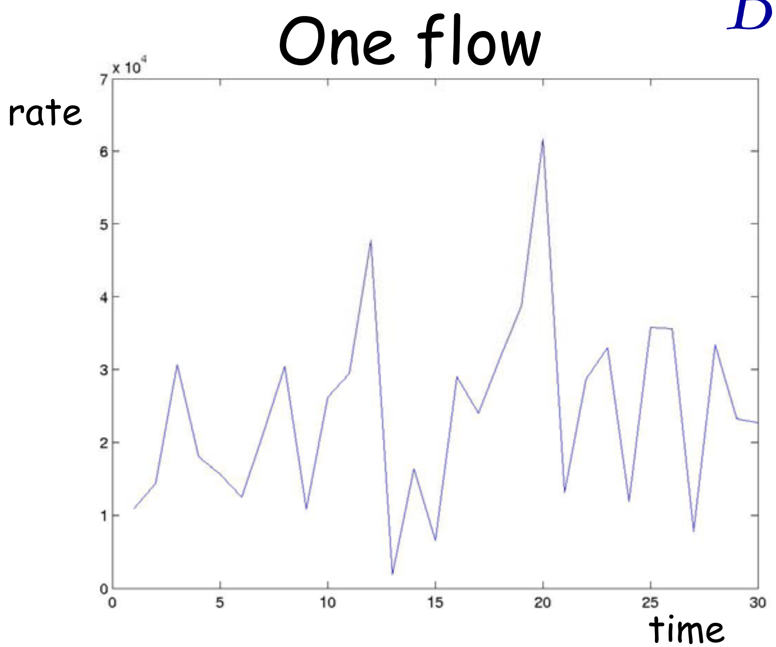

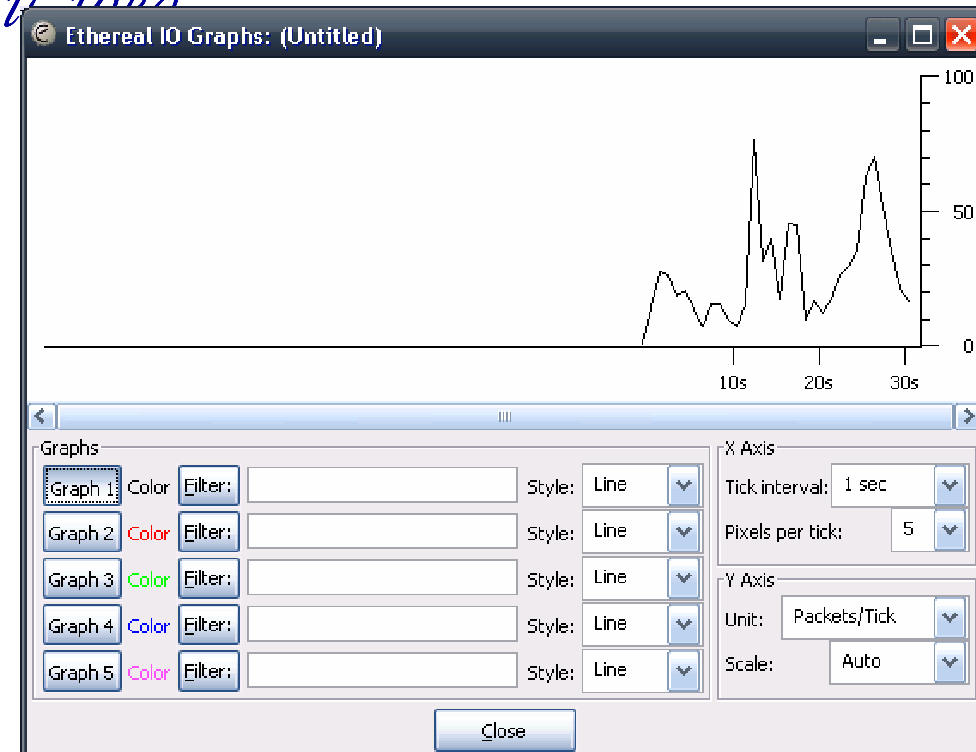

- Network traffic is bursty.
  i.e. the rate changes frequently.
- Peaks from independent flows generally occur at different times.
- Conclusion: The more flows we have, the smoother the traffic.

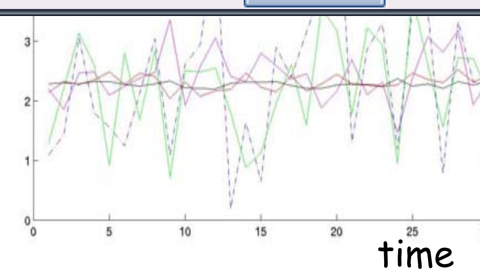

flows.

Assignments

- Study
  - Netflow
    - http://en.wikipedia.org/wiki/Netflow
  - SmokePing
    - http://oss.oetiker.ch/smokeping/
  - RRDTools
    - http://oss.oetiker.ch/rrdtool/doc/rrdtool.en.html

1. Volunteers for 5-min presentations???
2. Problem Set 1 – available online

Packet Switching

Statistical Multiplexing

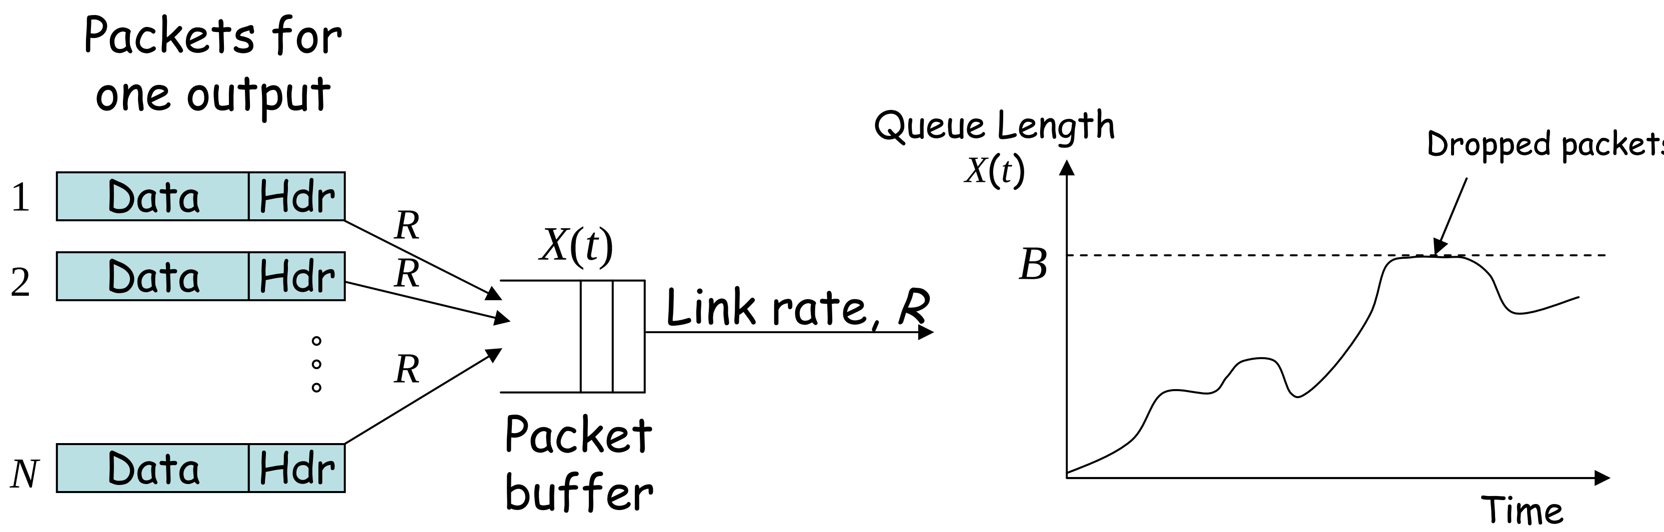

- Because the buffer absorbs temporary bursts, the egress link need not operate at rate $N.R$.
- But the buffer has finite size, B, so losses will occur.

Statistical Multiplexing

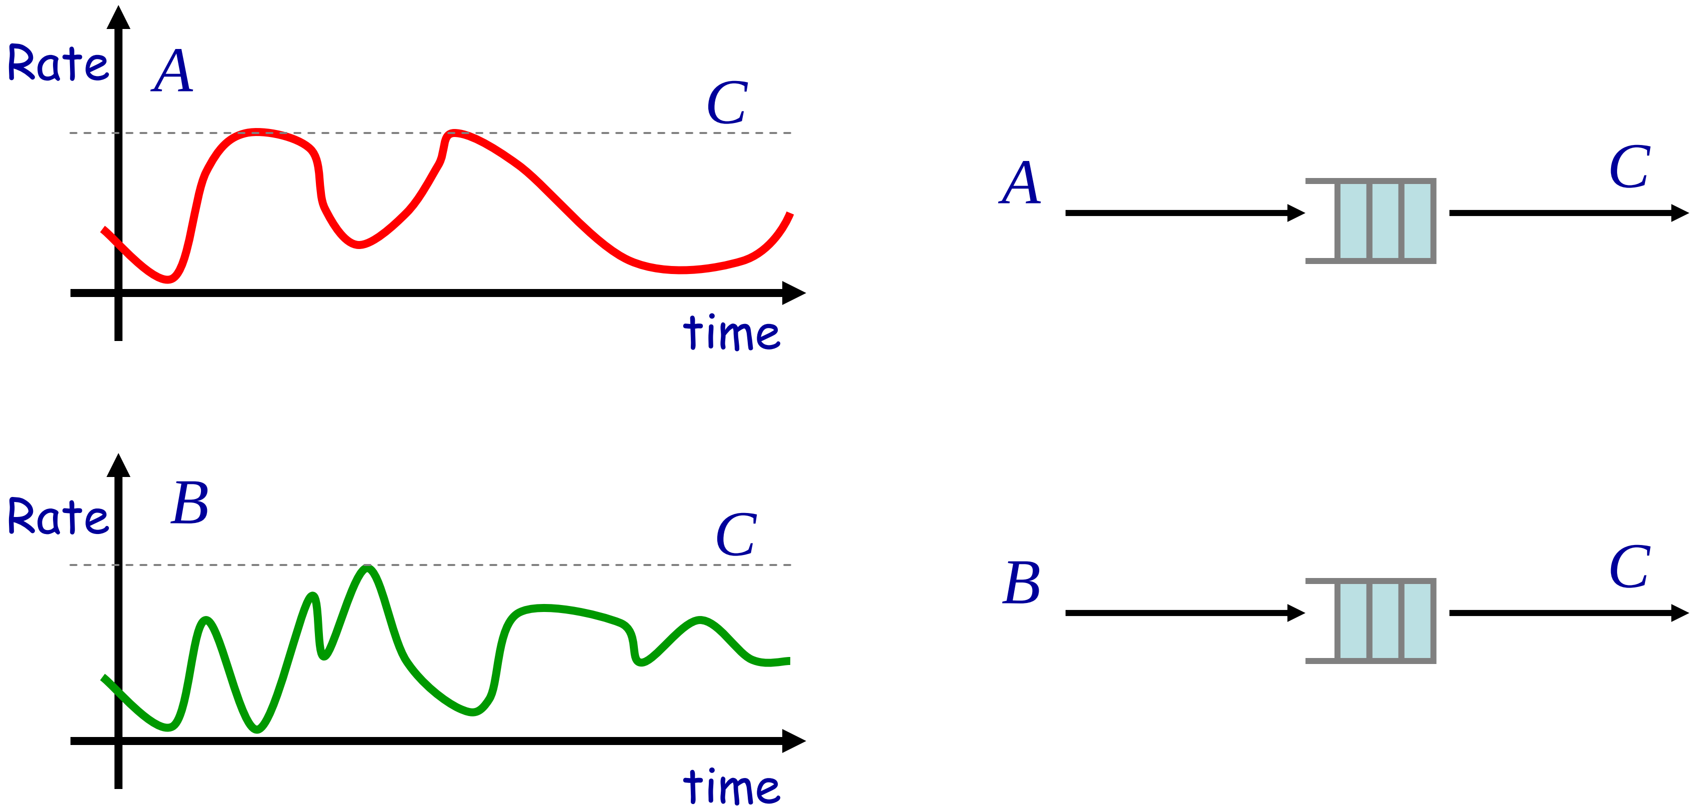

Statistical Multiplexing Gain

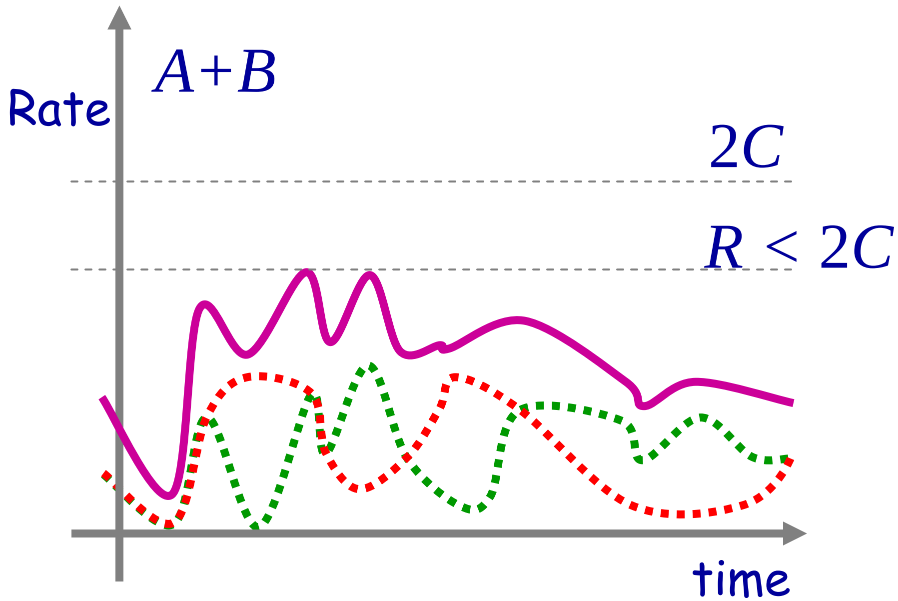

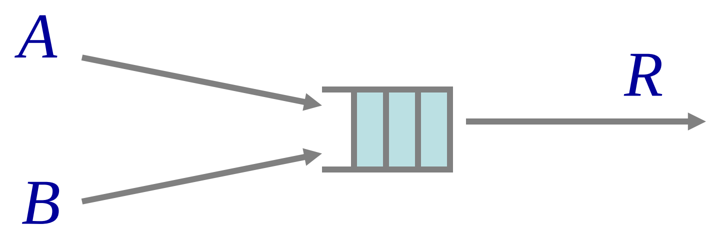

Statistical multiplexing gain = $2C/R$

Other definitions of SMG: The ratio of rates that give rise to a particular queue occupancy, or particular loss probability.

Why does the Internet use packet switching?

1. Efficient use of expensive links:
   - The links are assumed to be expensive and scarce.
   - Packet switching allows many, bursty flows to share the same link efficiently.
   - "Circuit switching is rarely used for data networks, ... because of very inefficient use of the links" - *Gallager*
2. Resilience to failure of links & routers:
   - "For high reliability, ... [the Internet] was to be a datagram subnet, so if some lines and [routers] were destroyed, messages could be ... rerouted" - *Tanenbaum*

Outline

- A Detailed FTP Example
- Layering
- Packet Switching and Circuit Switching
- Some terms
  - Data rate, "Bandwidth" and "throughput"
  - Propagation delay
  - Packet, header, address
  - Bandwidth-delay product, RTT

Some Definitions

- Packet length, *P*, is the length of a packet in bits.
- Link length, *L*, is the length of a link in meters.
- Data rate, *R*, is the rate at which bits can be sent, in bits/second, or b/s.[1]
- Propagation delay, *PROP*, is the time for one bit to travel along a link of length, *L*.

$$PROP = L/c.$$

- Transmission time, *TRANSP*, is the time to transmit a packet of length *P*.

$$TRANSP = P/R.$$

- Latency is the time from when the first bit begins transmission, until the last bit has been received. On a link:

$$Latency = PROP + TRANSP.$$

---

1. Note that a kilobit/second, kb/s, is 1000 bits/second, not 1024 bits/second.

Bandwidth

- In computer networks, Bandwidth → Channel Capacity
  - No. of bits that *can* be transmitted through the system over a period of time

- Throughput
  - The measured performance of a system
  - No. of bits transmitted *actually* transmitted per second

Latency

- The time delay between the moment something is initiated, and the moment one of its effects begins
- Measurement of Latency
  - *one-way*
    - Time taken to send a packet from the source to the destination
  - *round-trip*
    - one-way latency from source to destination + one-way latency from the destination back to the source
      - Doesn't cater for the amount of time that a destination system spends processing the packet

Packet Switching

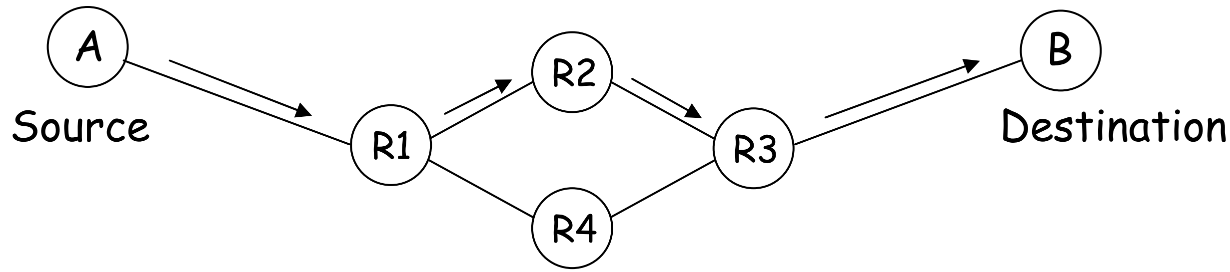

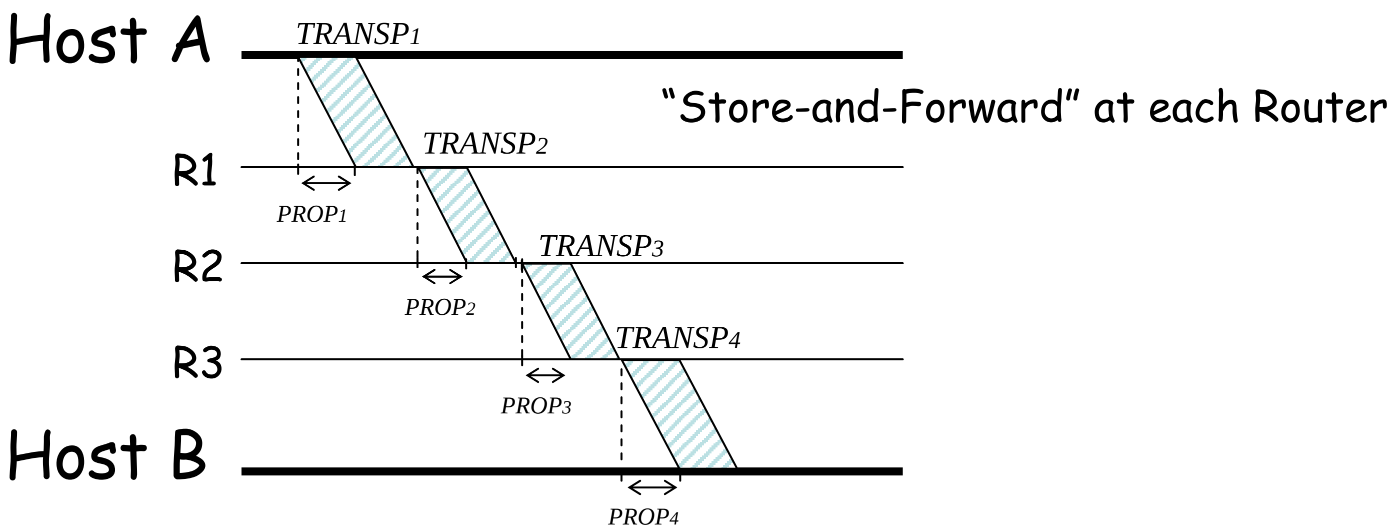

$$\text{Minimum end to end latency} = \sum_i (TRANSP_i + PROP_i)$$

Packet Switching

Why not send the entire message in one packet?

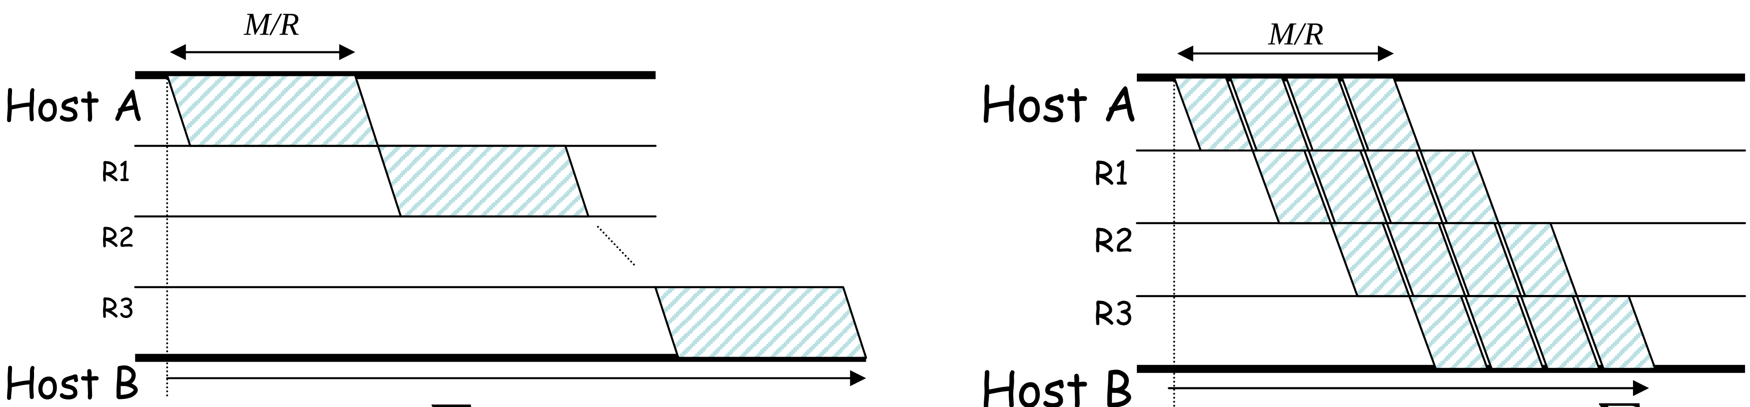

$$\text{Latency} = \sum_i (PROP_i + M / R_i)$$

$$\text{Latency} = M / R_{\min} + \sum_i PROP_i$$

Breaking message into packets allows parallel transmission across all links, reducing end to end latency. It also prevents a link from being "hogged" for a long time by one message.

Packet Switching

Queuing Delay

Because the egress link is not necessarily free when a packet arrives, it may be queued in a buffer. If the network is busy, packets might have to wait a long time.

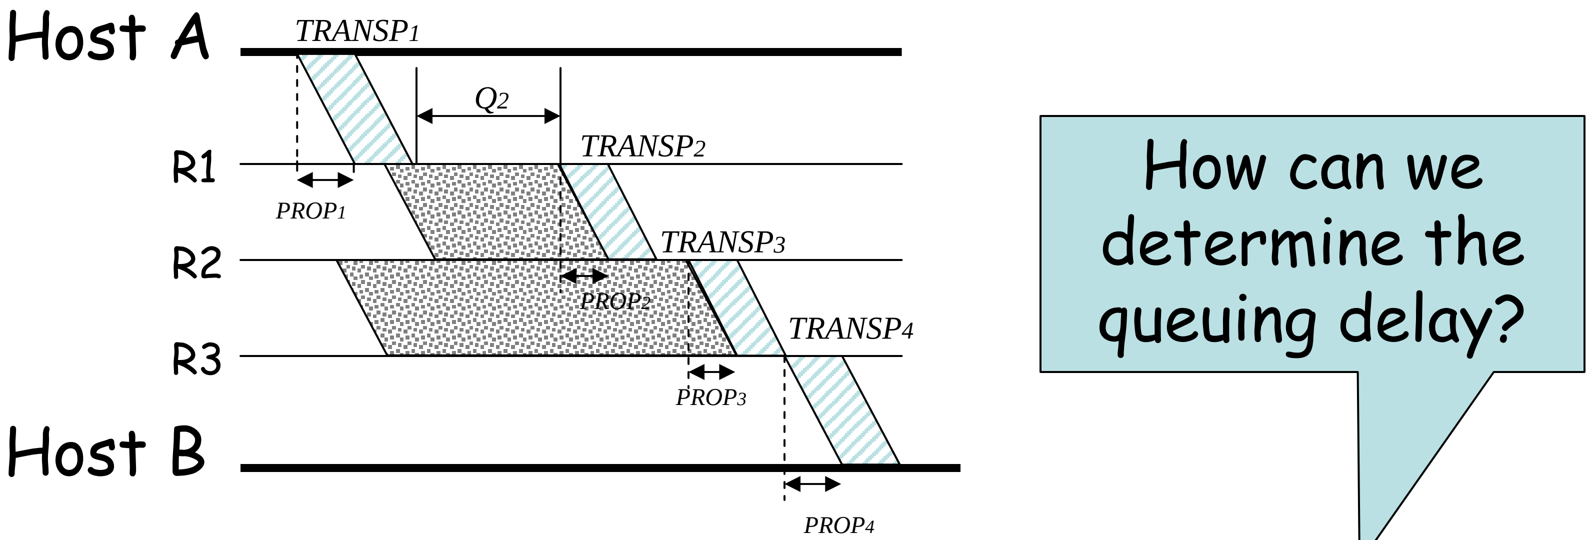

$$\text{Actual end to end latency} = \sum_i (TRANSP_i + PROP_i + Q_i)$$

Components of Latency

1. Propagation Delay
   - Speed of Light limit
   - Dependent on the transmission medium
     - ~ 2.4×10^8 m/s in copper
     - ~ 2×10^8 m/s in fiber

   - Propagation Delay (D_p)= Distance / Speed of Light

Components of Latency

2. Transmission Delay
   - The time from the *start* of packet *transmission* to the *end* of packet *transmission*
   - Transmission Delay (D_t) = Size of Data / Bandwidth

3. Queuing Delay
   - Time a packet waits in a queue until it can be processed
     - Avg. Queuing Delay (D_q) = (D_t / 2) * (N-1)
       - Where N is no. of packets the queue can hold

Latency = Propagation + Transmit + Queue

Bandwidth vs Latency

- Relative importance of each depends upon the application
- The postal mail service can be faster than the Internet?

Bandwidth x Delay

- Considering the network as a pipe, *delay x bandwidth* product gives us its capacity

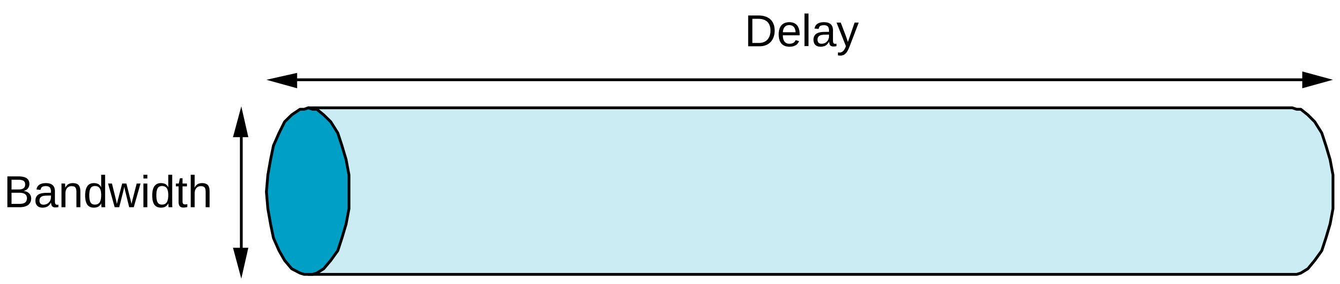

- The pipe should be ideally always full
  - How many bits the sender must transmit before the 1st bit arrives at the receiver?

Queues and Queuing Delay

- To understand the performance of a packet switched network, we can think of it as a series of queues interconnected by links.
- For given link rates and lengths, the only variable is the queuing delay.
- If we can understand the queuing delay, we can understand how the network performs.

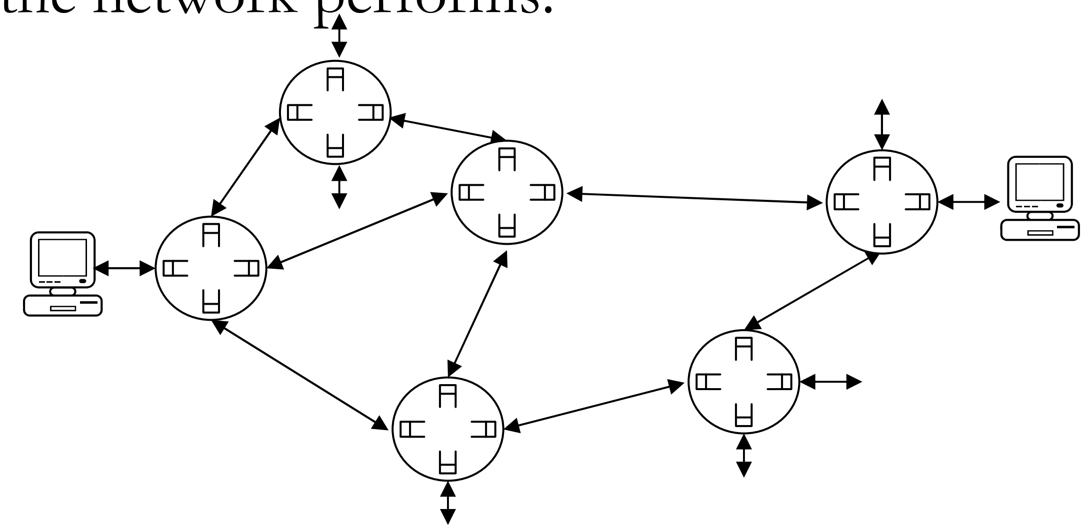

Queues and Queuing Delay

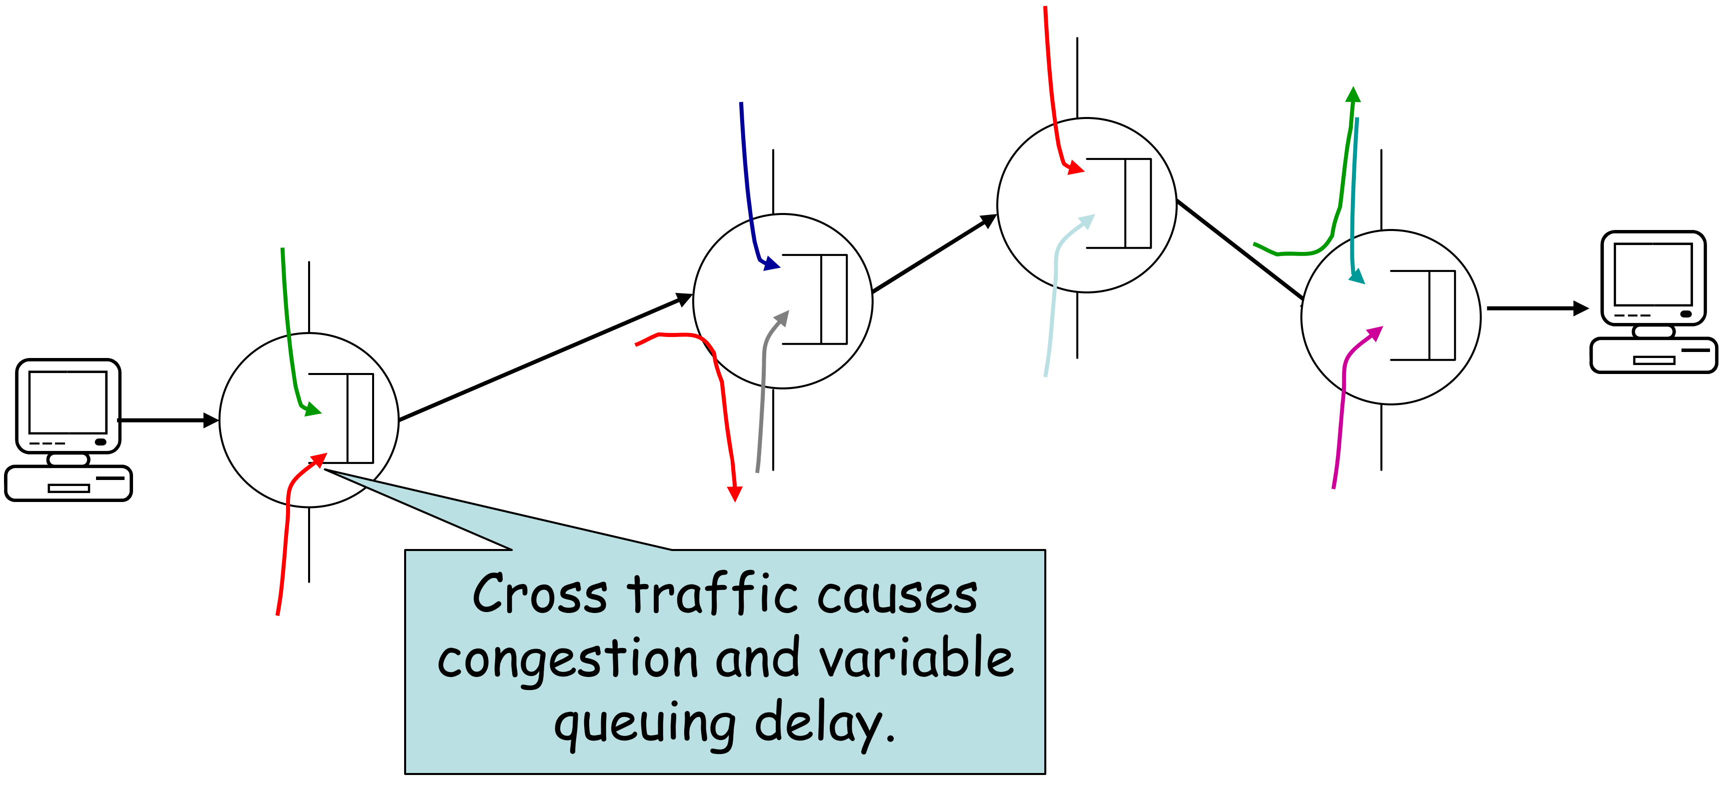

Summary

That's all folks!

A router queue

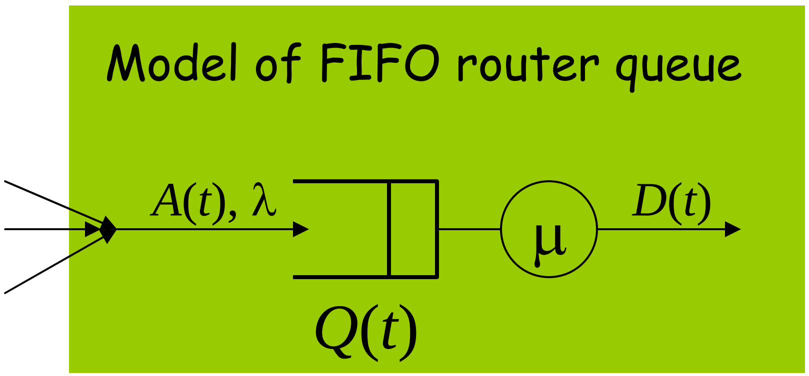

$A(t)$: The arrival process. The number of arrivals in interval $[0,t]$.

λ: The average rate of new arrivals in packets/second.

$D(t)$: The departure process. The number of departures in interval $[0,t]$.

$1/\mu$: The average time to service each packet.

$Q(t)$: The number of packets in the queue at time t.

A simple deterministic model

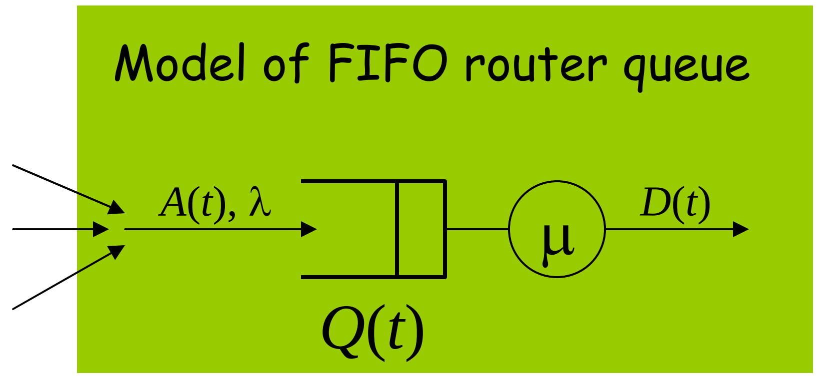

Properties of $A(t)$, $D(t)$:

- $A(t)$, $D(t)$ are non-decreasing
- $A(t) >= D(t)$

A simple deterministic model

bytes or "fluid"

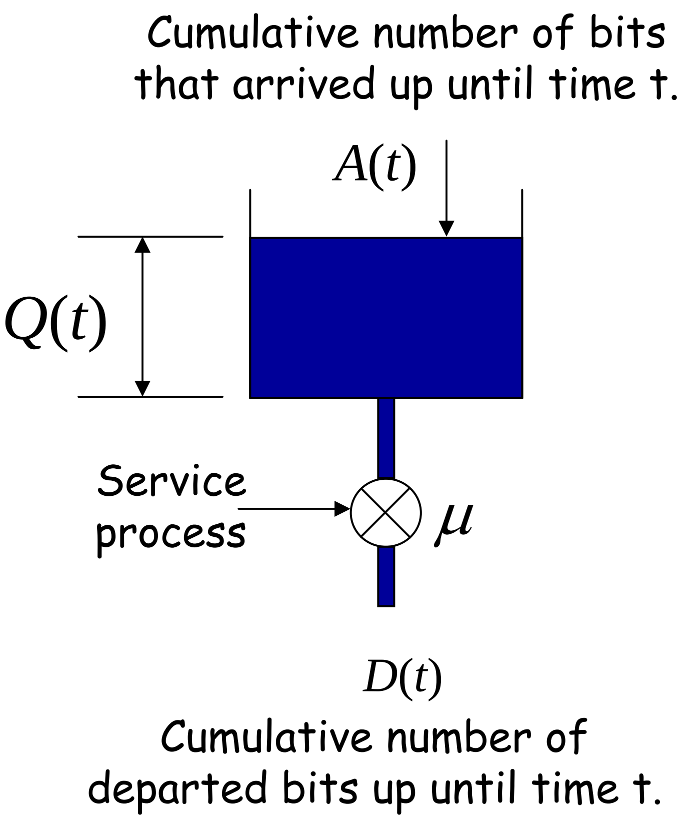

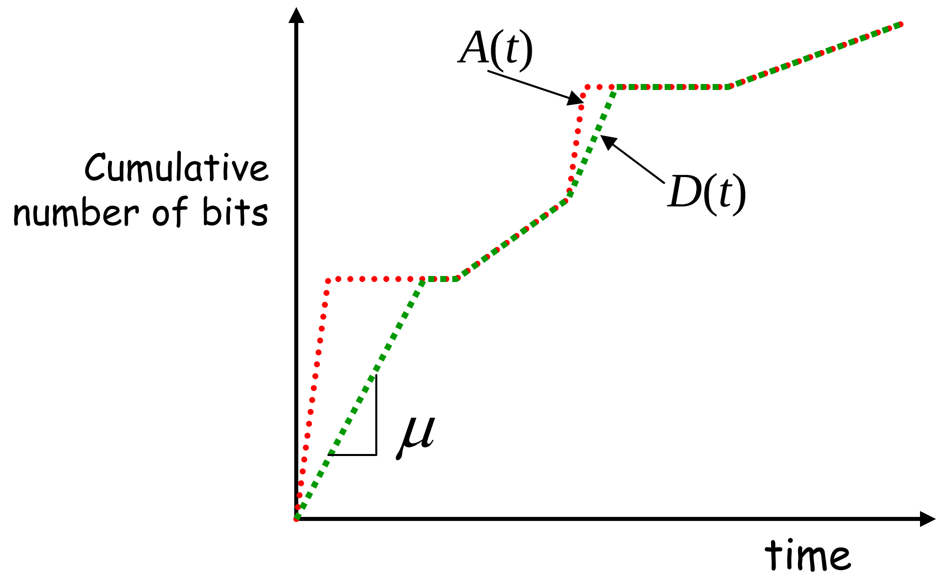

Properties of $A(t)$, $D(t)$:

- $A(t)$, $D(t)$ are non-decreasing
- $A(t) >= D(t)$

Simple deterministic model

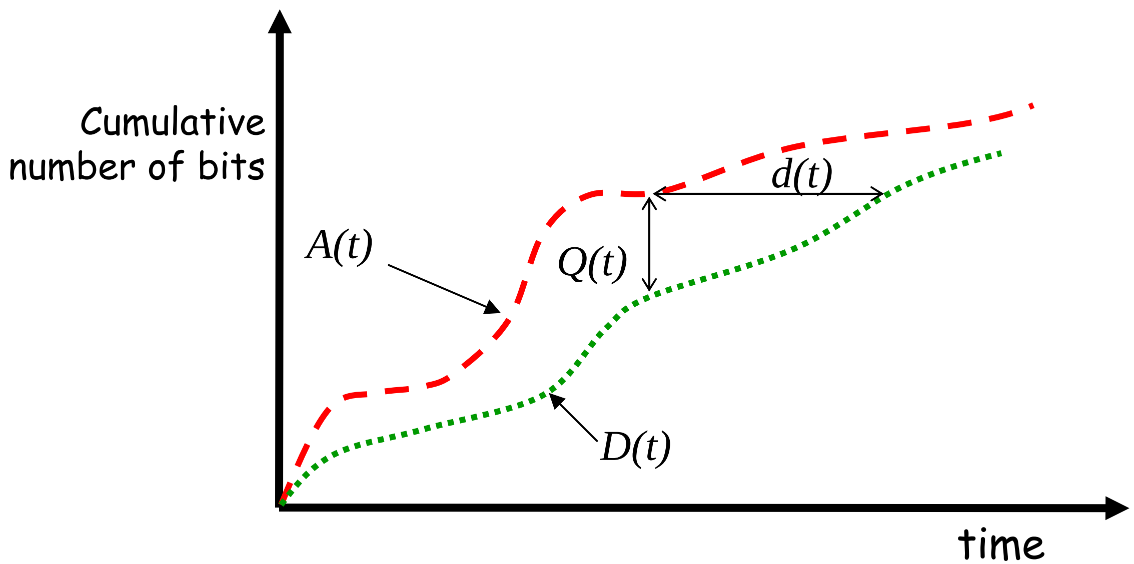

Queue occupancy: $Q(t) = A(t) - D(t)$.

Queuing delay, $d(t)$, is the time spent in the queue by a bit that arrived at time t, and if the queue is served first-come-first-served (FCFS or FIFO)

Example

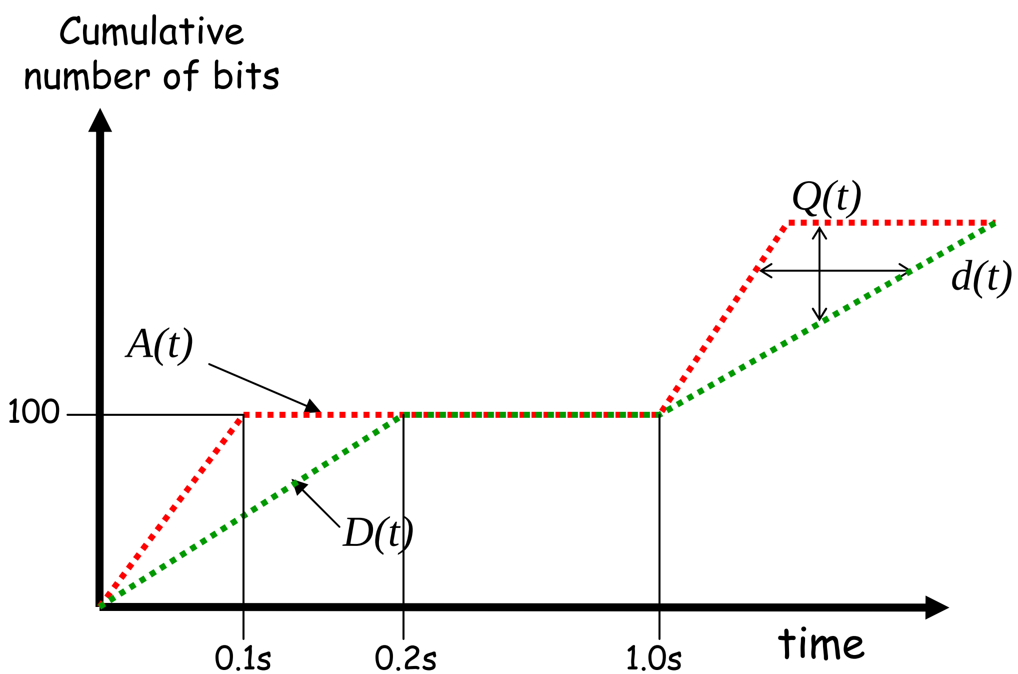

Example: Every second, a train of 100 bits arrive at rate 1000b/s. The maximum departure rate is 500b/s. What is the average queue occupancy?

Solution: During each cycle, the queue fills at rate 500b/s for 0.1s, then drains at rate 500b/s for 0.1s. The average queue occupancy when the queue is non-empty is therefore: $(\bar{Q}(t) | Q(t) > 0) = 0.5 \times (0.1 \times 500) = 25$ bits. The queue is empty for 0.8s each cycle, and so: $\bar{Q}(t) = (0.2 \times 25) + (0.8 \times 0) = 5$ bits. (You'll probably have to think about this for a while...).