We were doing very well, up to the kind of sum when a bath is filling at the rate of so many gallons and two holes are letting the water out, and please to say how long it will take to fill the bath, when my mother put down the socks she was darning and clicked her tongue in impatience.

"Filling up an old bath with holes in it, indeed. Who would be such a fool?"

"A sum it is, girl," my father said. "A sum. A problem for the mind."

"Filling the boy with old nonsense," Mama said.

"Not nonsense, Beth," my father said. "A sum, it is. The water pours in and takes so long. It pours out and takes so long. How long to fill? That is all."

"But who would pour water into an old bath with holes?" my mother said. "Who would think to do it, but a lunatic?"

—How Green Was My Valley,

Richard Llewellyn

ETSF05/ETSF10 – Internet Protocols

Performance & QoS
Congestion Control

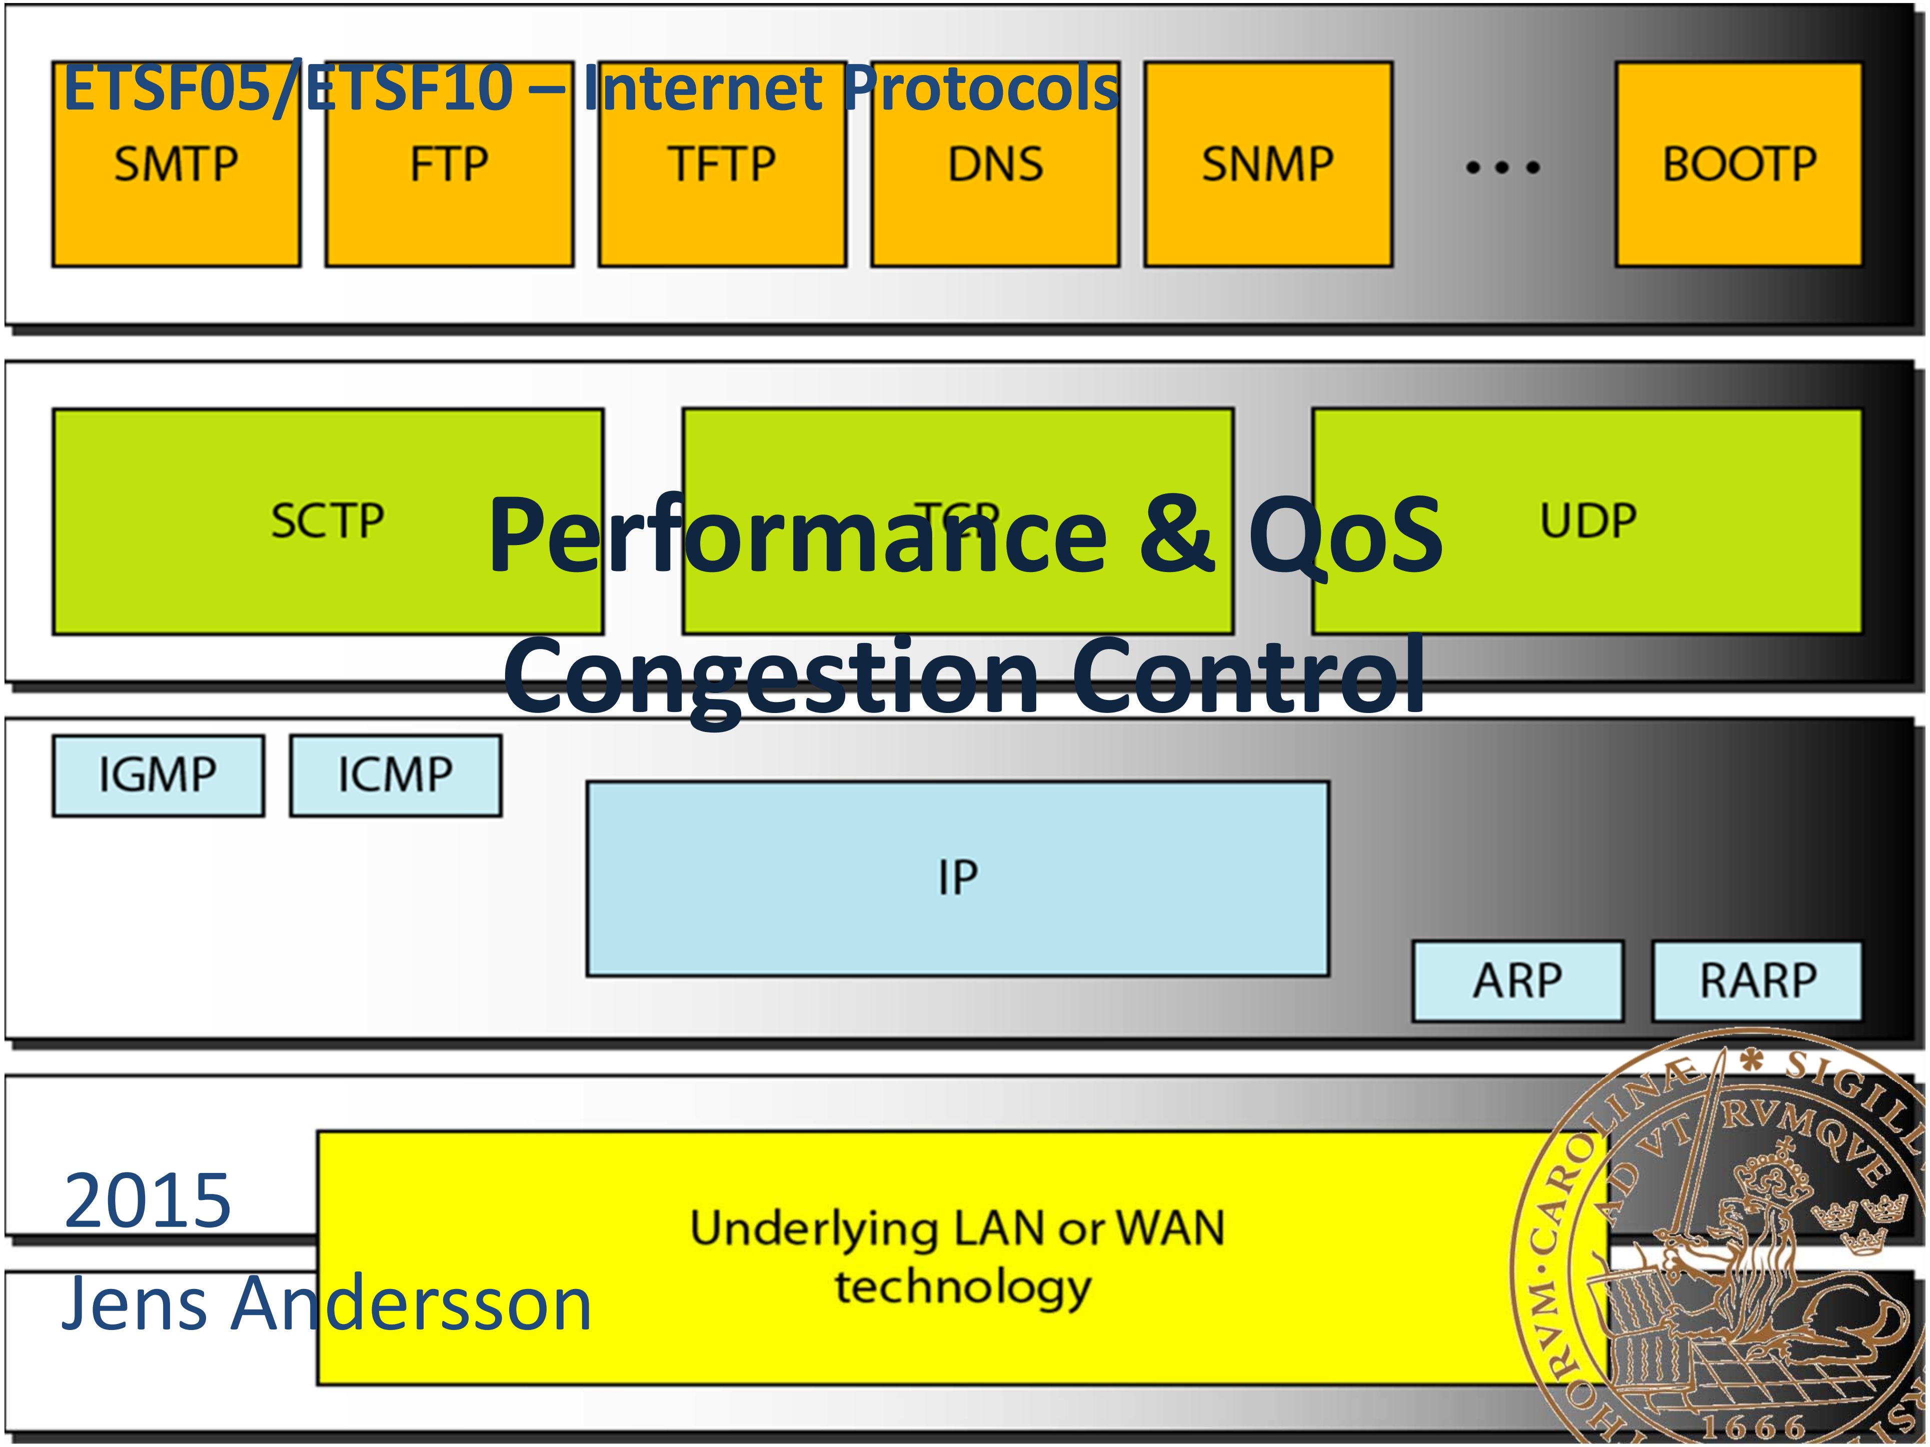

2015

Jens Andersson

Quality of Service (QoS)

- Maintaining a functioning network
  - Meeting applications' demands
    - User's demands = QoE
  - Dealing with flow characteristics

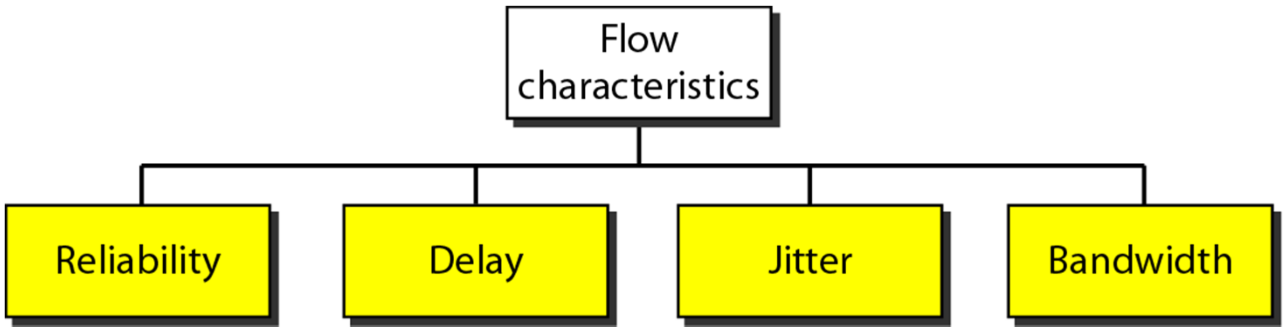

Jitter = Packet Delay Variations

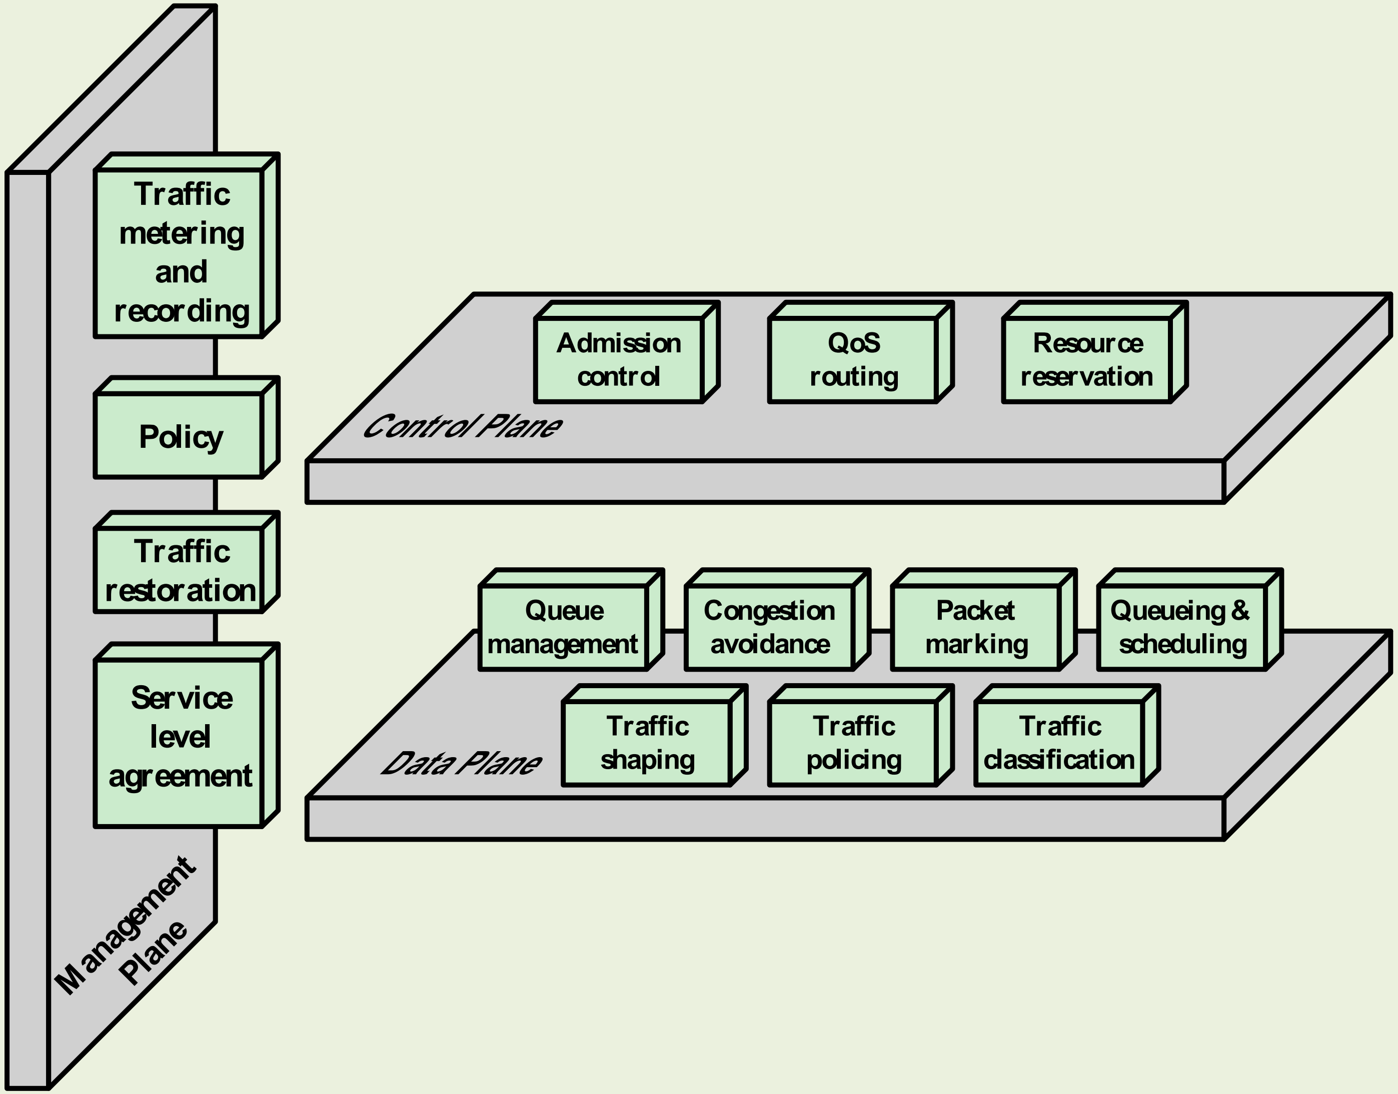

Figure 22.1 Architectural Framework for QoS Support

Data Plane

- Includes those mechanisms that operate directly on flows of data

- Queue management algorithms
- Queueing and scheduling
- Congestion avoidance
- Packet marking
- Traffic classification
- Traffic policing
- Traffic shaping

Control Plane

- Concerned with creating and managing the pathways through which user data flows
- It includes:
  - Admission control
  - QoS routing
  - Resource reservation

Management Plane

- Contains mechanisms that affect both control plane and data plane mechanisms
- Includes:
  - Service level agreement (SLA)
  - Traffic metering and recording
  - Traffic restoration
  - Policy

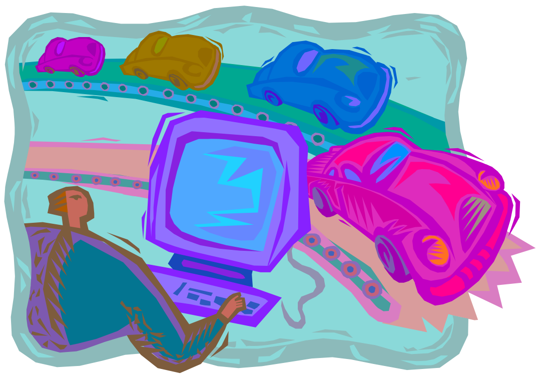

Network performance

- Bandwidth
  - Bits per second (capacity)
- Throughput
  - Efficiency, always less than capacity (<1)
- Latency (Delay)
  - Transmission, propagation, processing, queueing
- PDV = Packet Delay Variation (Jitter) → real-time data!

Other parameters

- Bit Error Rate
  – L1 parameter that heavily impacts on L3
  – Frame/Packet Loss on higher layers
- Inter Packet Gap variations
  – "Jitter"
  – Could be non-zero already at sender
- Ratio of packets out of order
  – Impact on delay in TCP

Circuit Switching	Datagram Packet Switching	Virtual Circuit Packet Switching
Dedicated transmission path	No dedicated path	No dedicated path
Continuous transmission of data	Transmission of packets	Transmission of packets
Fast enough for interactive	Fast enough for interactive	Fast enough for interactive
Messages are not stored	Packets may be stored until delivered	Packets stored until delivered
The path is established for entire conversation	Route established for each packet	Route established for entire conversation
Call setup delay; negligible transmission delay	Packet transmission delay	Call setup delay; packet transmission delay
Busy signal if called party busy	Sender may be notified if packet not delivered	Sender notified of connection denial
Overload may block call setup; no delay for established calls	Overload increases packet delay	Overload may block call setup; increases packet delay
Electromechanical or computerized switching nodes	Small switching nodes	Small switching nodes
User responsible for message loss protection	Network may be responsible for individual packets	Network may be responsible for packet sequences
Usually no speed or code conversion	Speed and code conversion	Speed and code conversion
Fixed bandwidth	Dynamic use of bandwidth	Dynamic use of bandwidth
No overhead bits after call setup	Overhead bits in each packet	Overhead bits in each packet

Table 9.1

Comparison of Communication Switching Techniques

(Table can be found on page 315 in textbook)

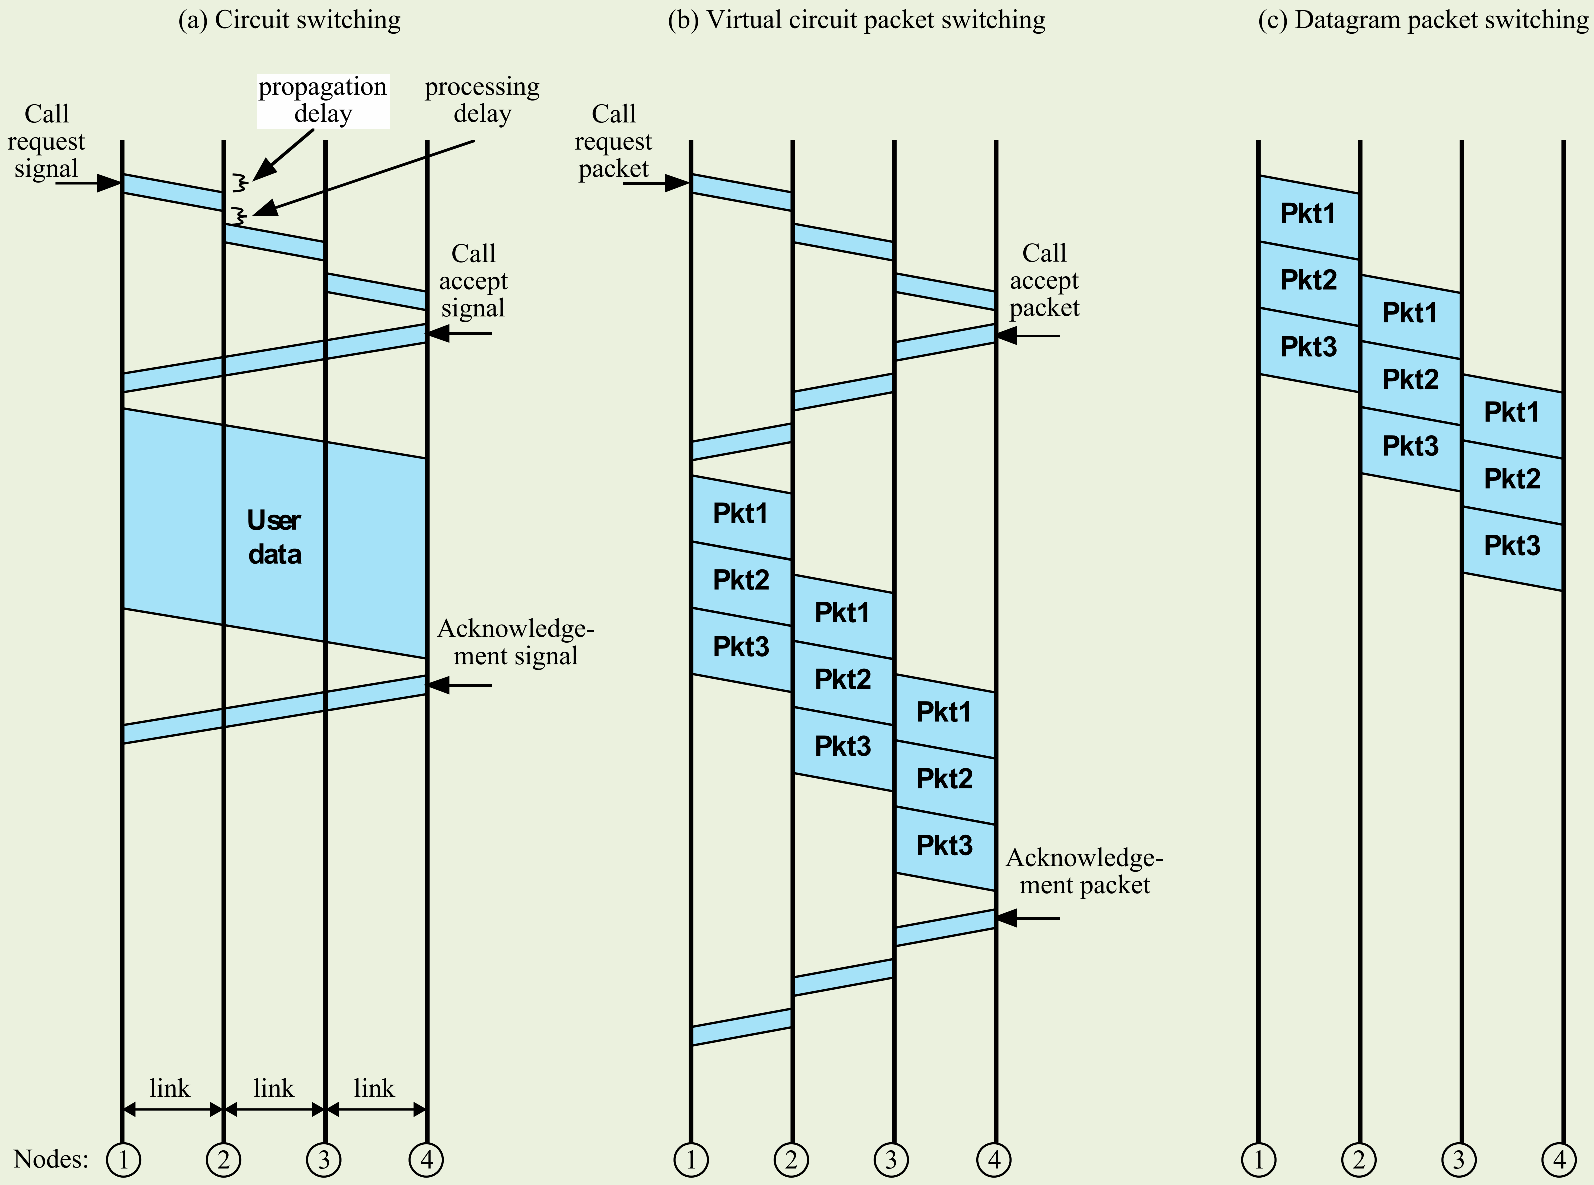

Figure 9.15 Event Timing for Circuit Switching and Packet Switching

Virtual Circuits vs. Datagram

- Virtual circuits
  - Network can provide sequencing and error control
  - Packets are forwarded more quickly
  - Less reliable (compare Circuit Switching)
- Datagram (Best Effort)
  - No call setup phase
  - Individual packet handling
  - More flexible
  - More reliable

IP Performance Metrics (IPPM wg)

- Chartered by IETF to develop standard metrics that relate to the quality, performance, and reliability of Internet data delivery
- Need for standardization:
  - Internet has grown and continues to grow at a dramatic rate
  - Internet serves a large and growing number of commercial and personal users across an expanding spectrum of applications

Table 22.3
IP Performance Metrics

Metric Name	Singleton Definition	Statistical Definitions
One-Way Delay	Delay = dT, where Src transmits first bit of packet at T and Dst received last bit of packet at T + dT	Percentile, median, minimum, inverse percentile
Round-Trip Delay	Delay = dT, where Src transmits first bit of packet at T and Src received last bit of packet immediately returned by Dst at T + dT	Percentile, median, minimum, inverse percentile
One-Way Loss	Packet loss = 0 (signifying successful transmission and reception of packet); = 1 (signifying packet loss)	Average
One-Way Loss Pattern	Loss distance: Pattern showing the distance between successive packet losses in terms of the sequence of packets Loss period: Pattern showing the number of bursty losses (losses involving consecutive packets)	Number or rate of loss distances below a defined threshold, number of loss periods, pattern of period lengths, pattern of inter-loss period lengths.
Packet Delay Variation	Packet delay variation (pdv) for a pair of packets with a stream of packets = difference between the one-way-delay of the selected packets	Percentile, inverse percentile, jitter, peak-to-peak pdv

Src = IP address of a host
Dst = IP address of a host

(a) Sampled metrics

Table 22.3
IP Performance Metrics

Metric Name	General Definition	Metrics
Connectivity	Ability to deliver a packet over a transport connection.	One-way instantaneous connectivity, Two-way instantaneous connectivity, one-way interval connectivity, two-way interval connectivity, two-way temporal connectivity
Bulk Transfer Capacity	Long-term average data rate (bps) over a single congestion-aware transport connection.	BTC = (data sent)/(elapsed time)

(b) Other metrics

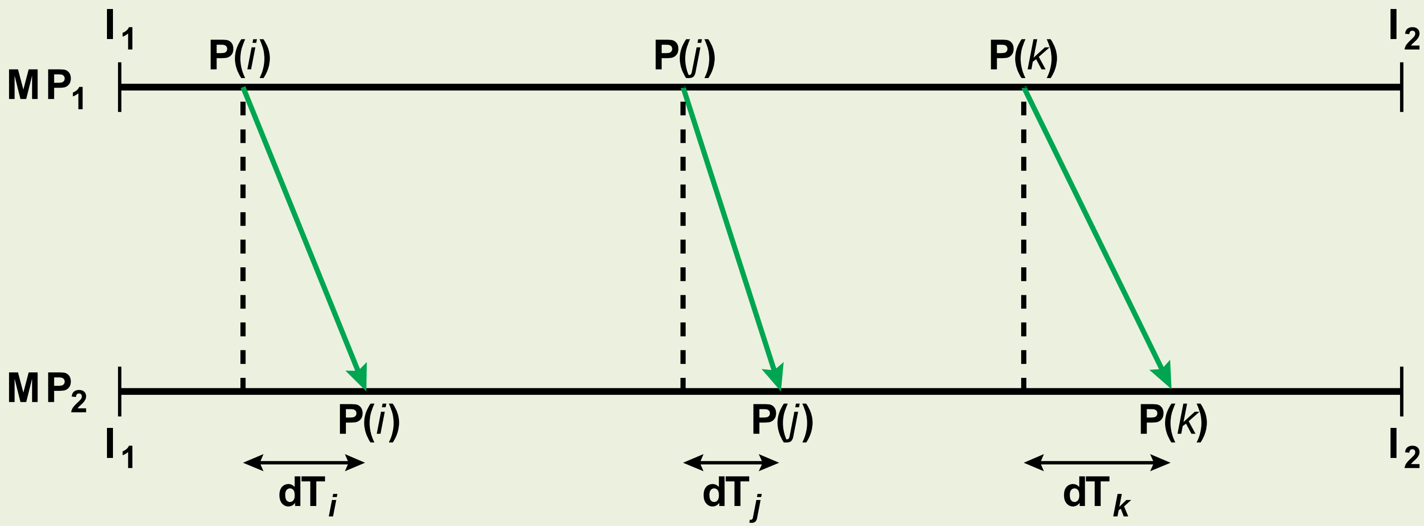

$$PDV_i = dT_i - dT_{i-1}$$

I_1, I_2 = times that mark that beginning and ending of the interval in which the packet stream from which the singleton measurement is taken occurs.
MP_1, MP_2 = source and destination measurement points
P(*i*) = *i*th measured packet in a stream of packets
dT_i = one-way delay for P(*i*) → Time synch!

Figure 22.12 Model for Defining Packet Delay Variation

Bandwidth-delay product

- How much data fills the link
  - (or full path?)

- One Way Delay
- Two Way Delay = Round Trip Time (RTT)
  Time for data + time for ACK

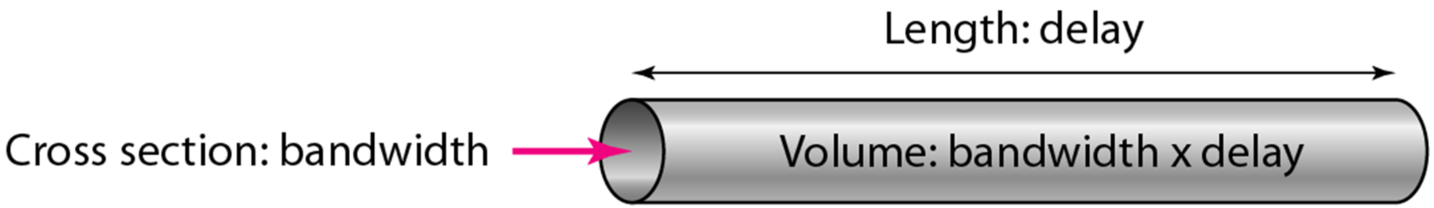

Bandwidth-delay product

bandwidth: 5 bps, delay: 5s
bandwidth x delay = 25 bits

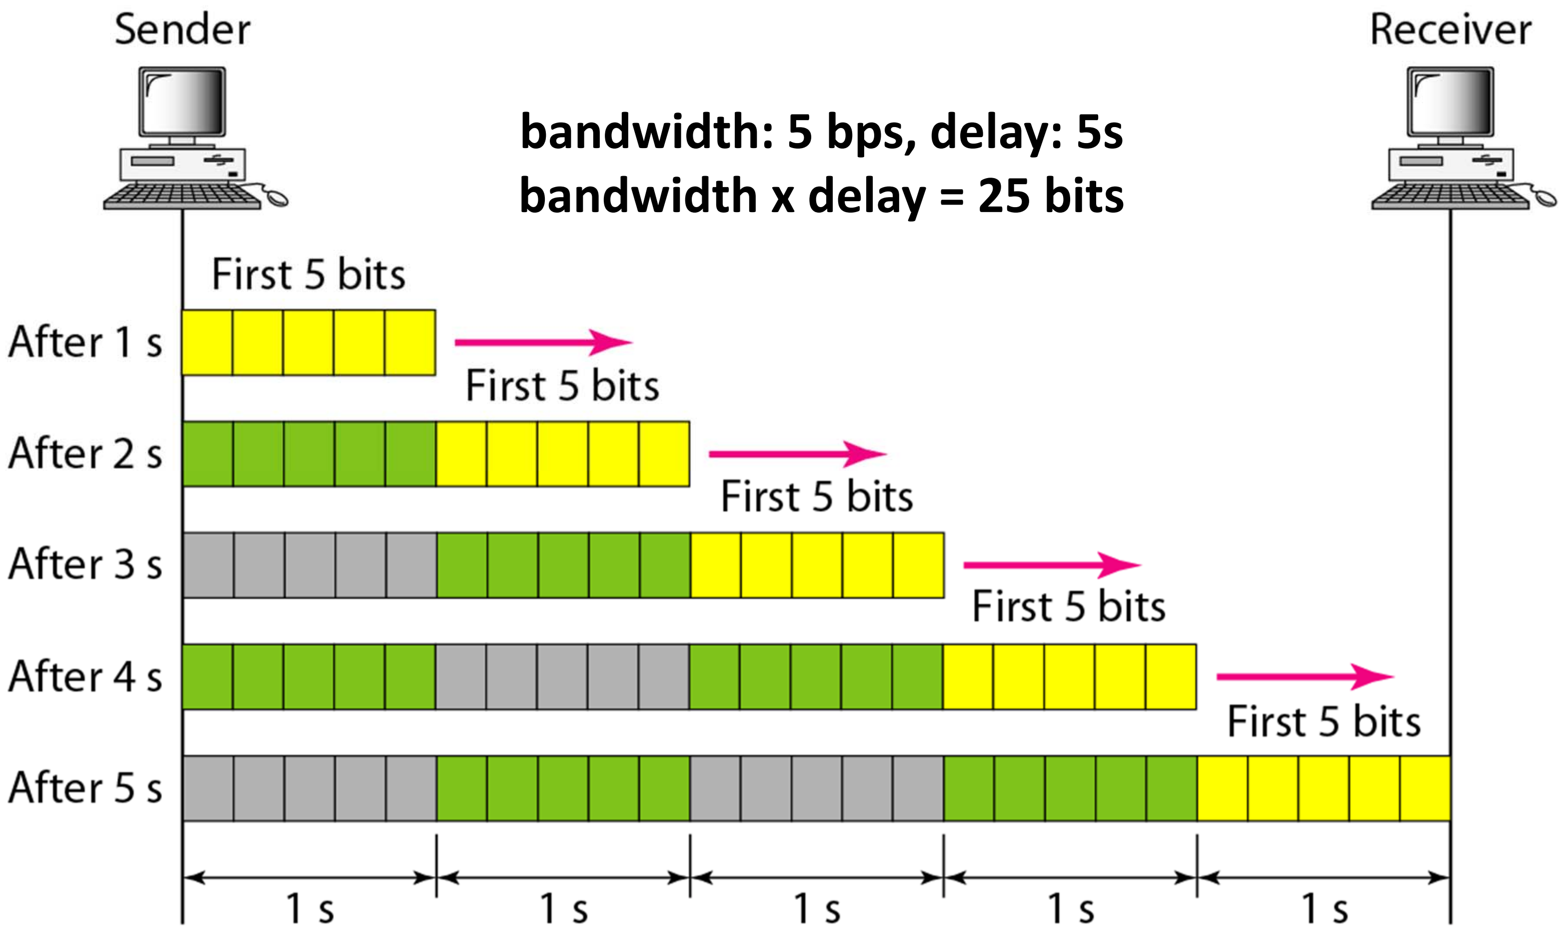

Bandwidth-delay product

- Important for congestion avoidance
  - Don’t overfill the link
- Important for efficiency
  - Keep the link filled at all times
  - For max efficiency
    Data chunks > 2 * bandwidth * delay

Bandwidth-delay product

- Important for tuning (TCP)
- Long Fat Network (LFN, ”elephant”)
  BDP >> 10^5 bits

- Very long (high delay) links:
  -> Bandwidth = BDP/delay
  But it takes long time before ACK arrives …

Performance vs ARQ

- Method
  - Stop-&-Wait
  - Go-Back-N
  - Selective-Repeate
- Utilisation = f(window size)

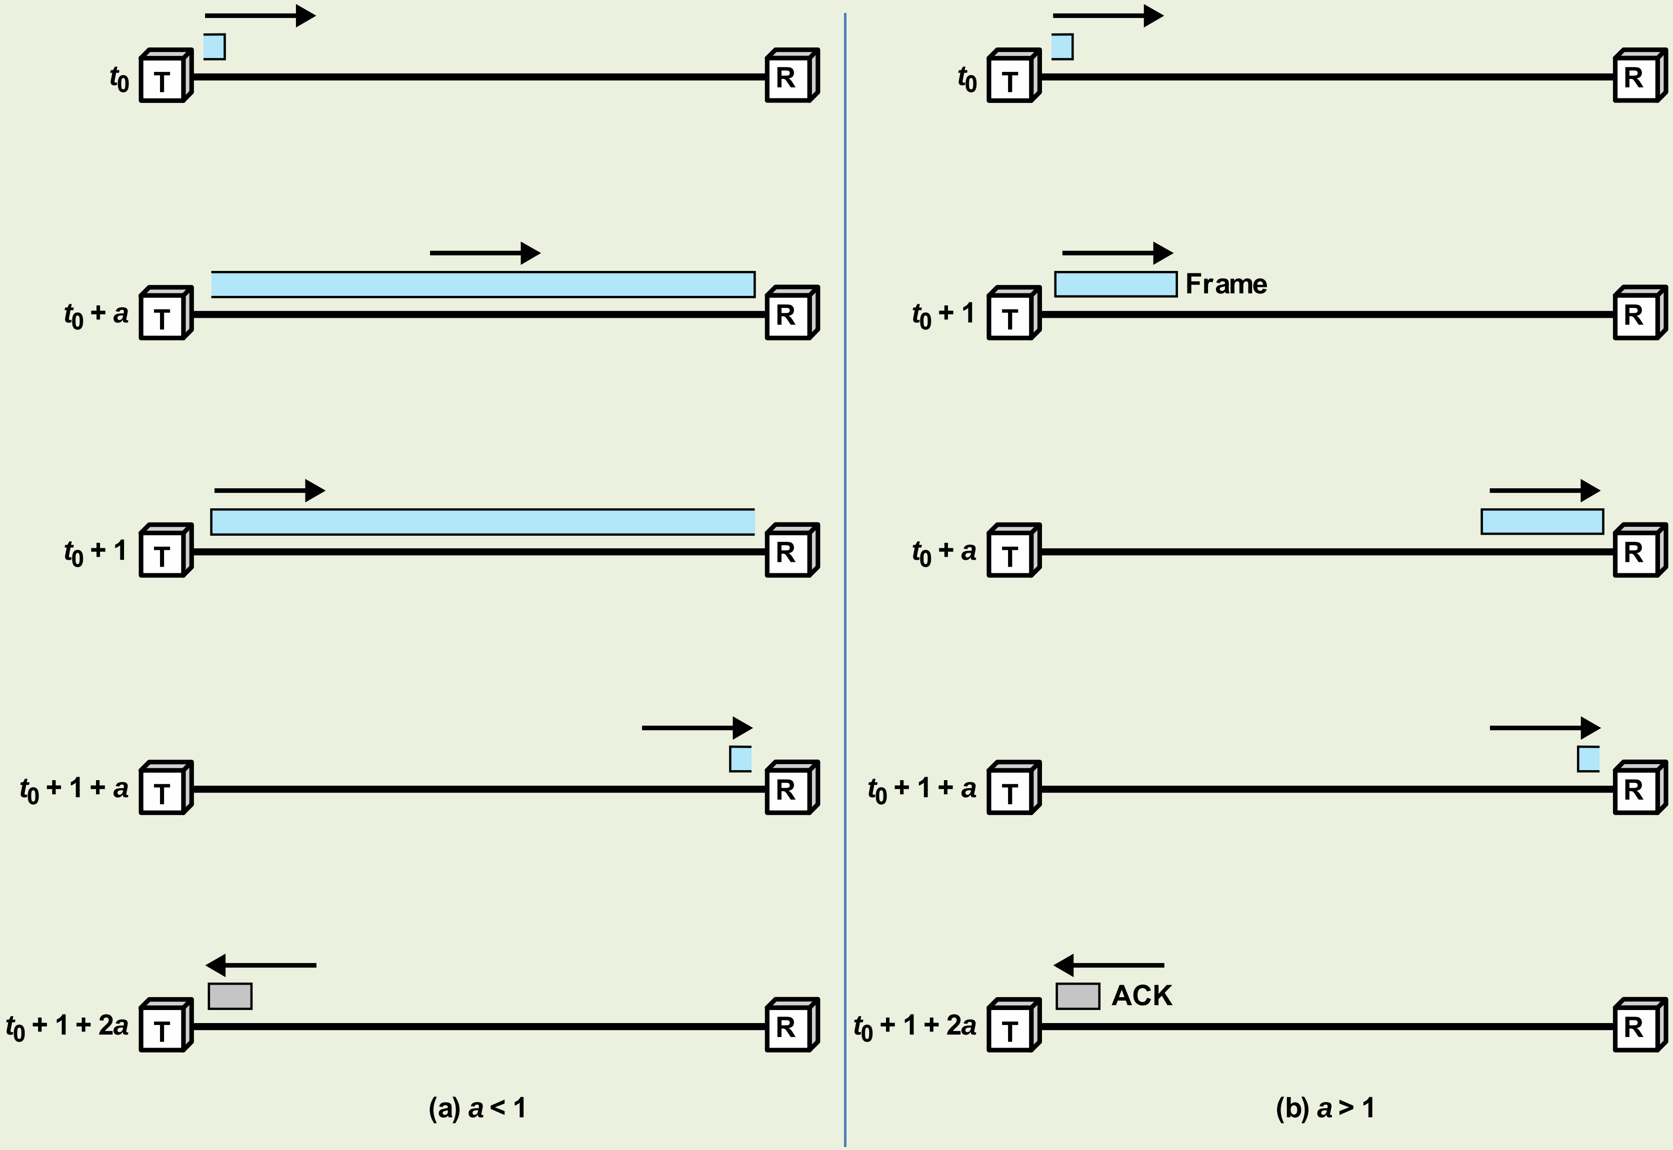

Figure 16.8 Stop-and-Wait Link Utilization (transmission time = 1; propagation time = a)

Sliding Windows based

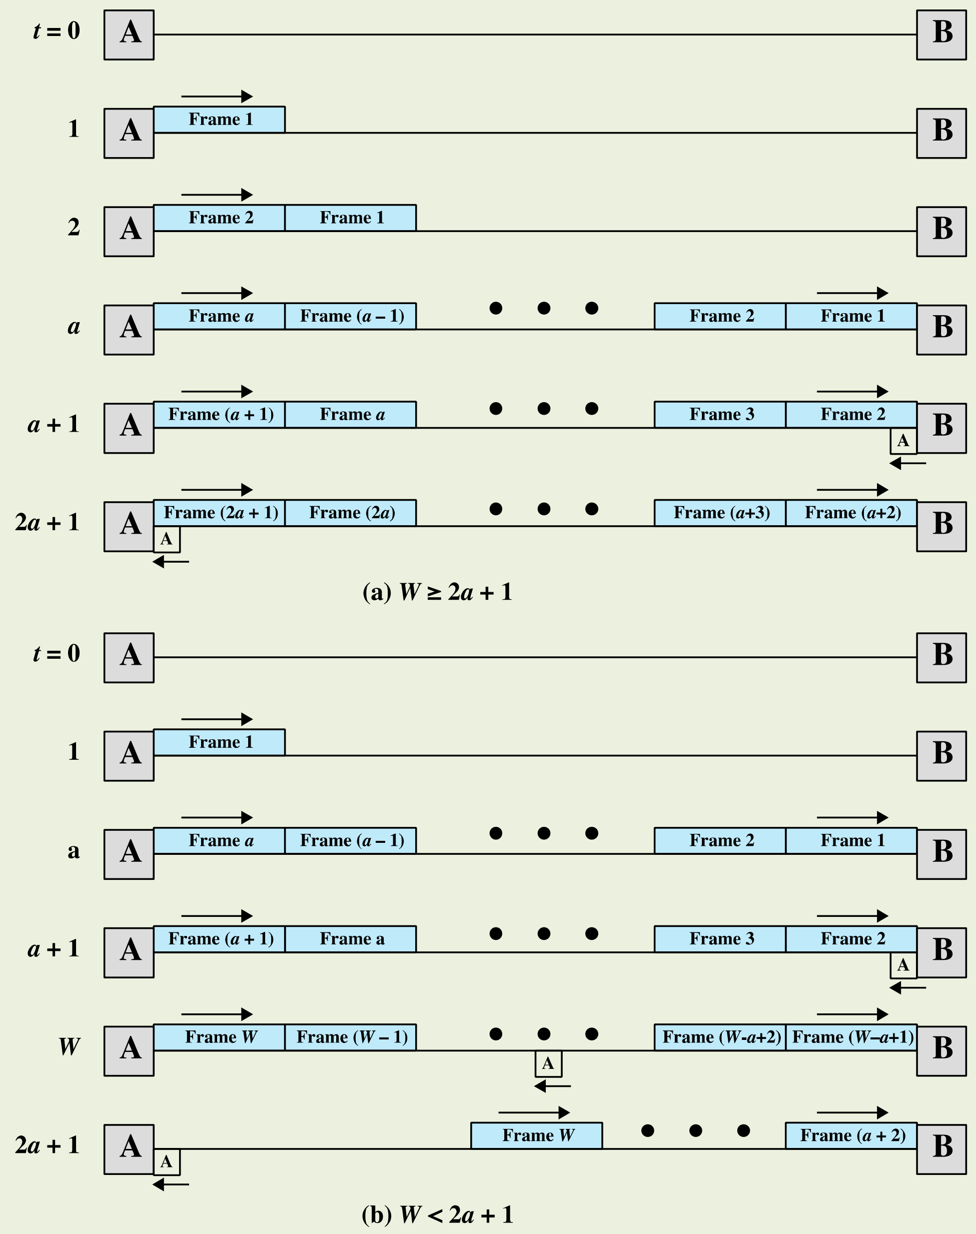

Figure 16.9 Timing of Sliding-Window Protocol

- a = propagation time
- w = window size
- Compare with Bandwitdh-Delay Product

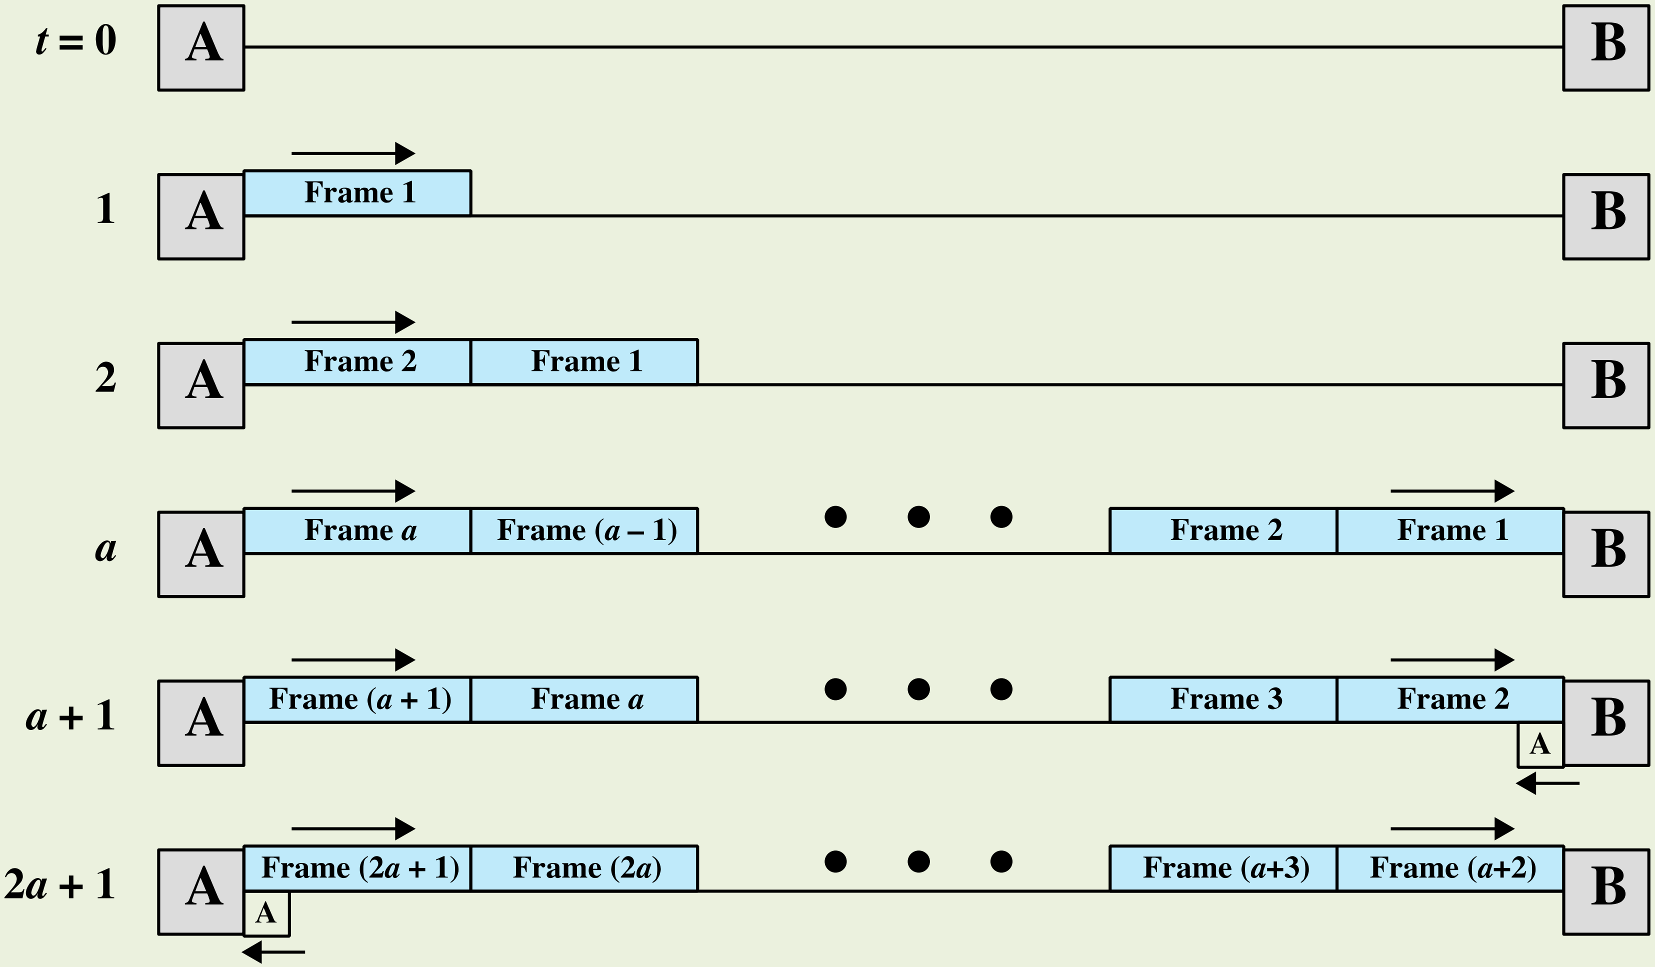

t = 0
A
B
1
Frame 1
2
Frame 2
Frame 1
a
Frame a
Frame (a – 1)
Frame 2
Frame 1
a + 1
Frame (a + 1)
Frame a
Frame 3
Frame 2
2a + 1
Frame (2a + 1)
Frame (2a)
Frame (a+3)
Frame (a+2)

(a) $W \geq 2a + 1$

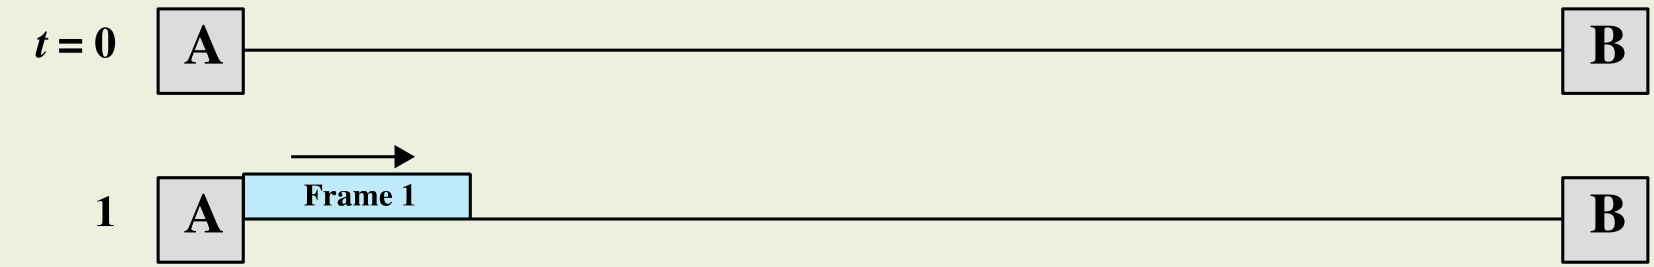

t = 0
A
B
1
Frame 1

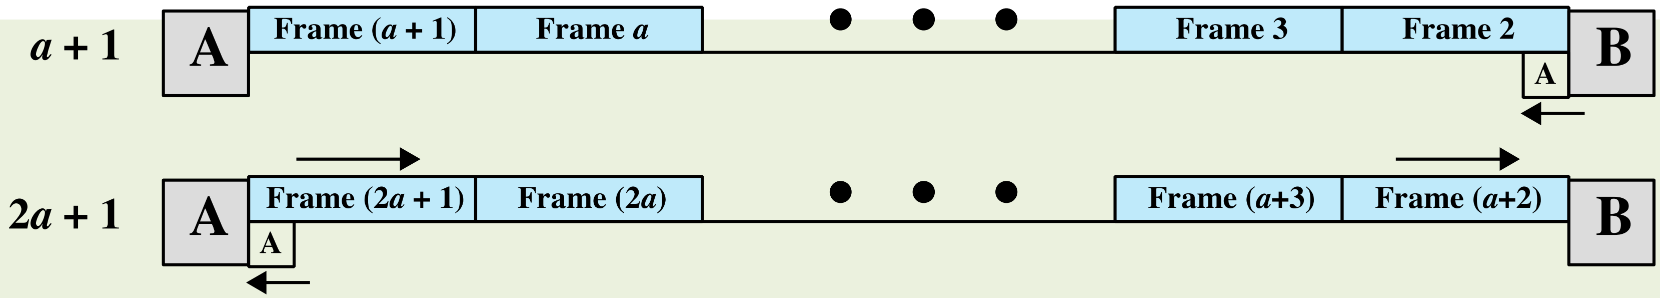

(a) $W \geq 2a + 1$

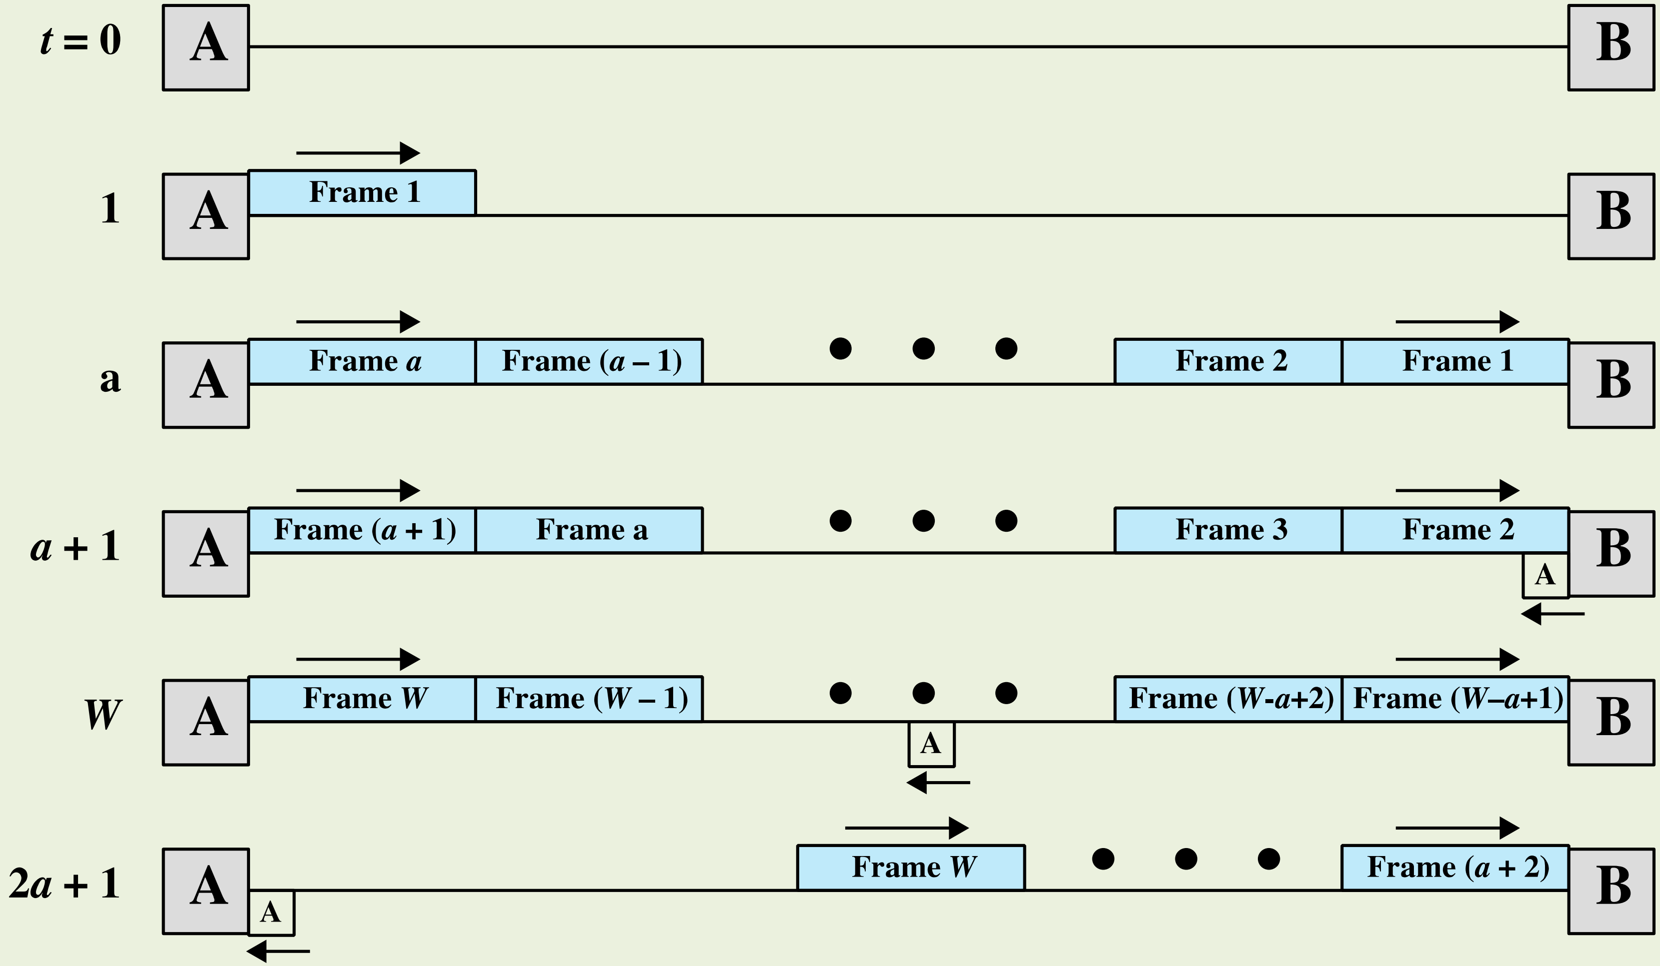

(b) $W < 2a + 1$

Packet Size vs Transmission Time

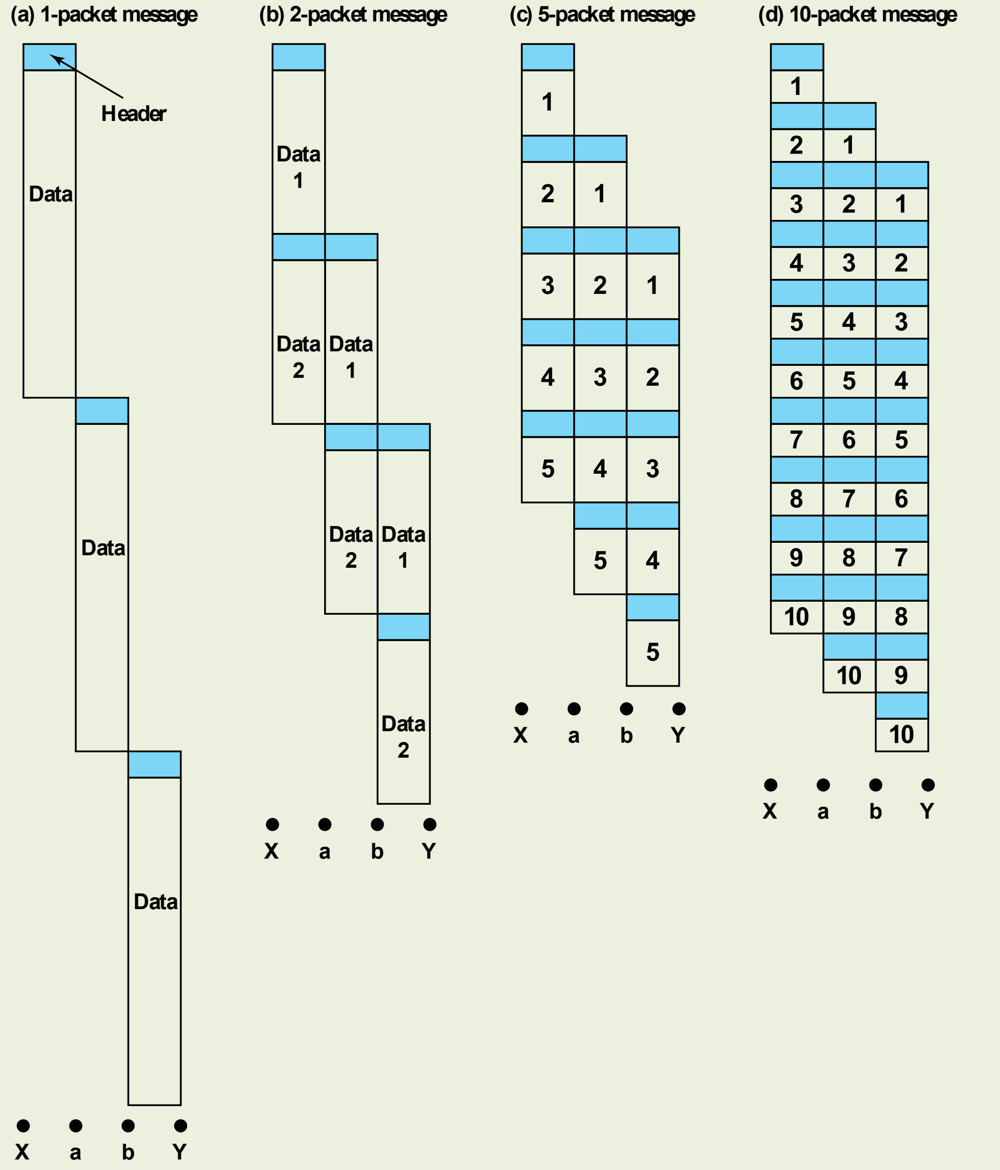

Figure 9.13 Effect of Packet Size on Transmission Time

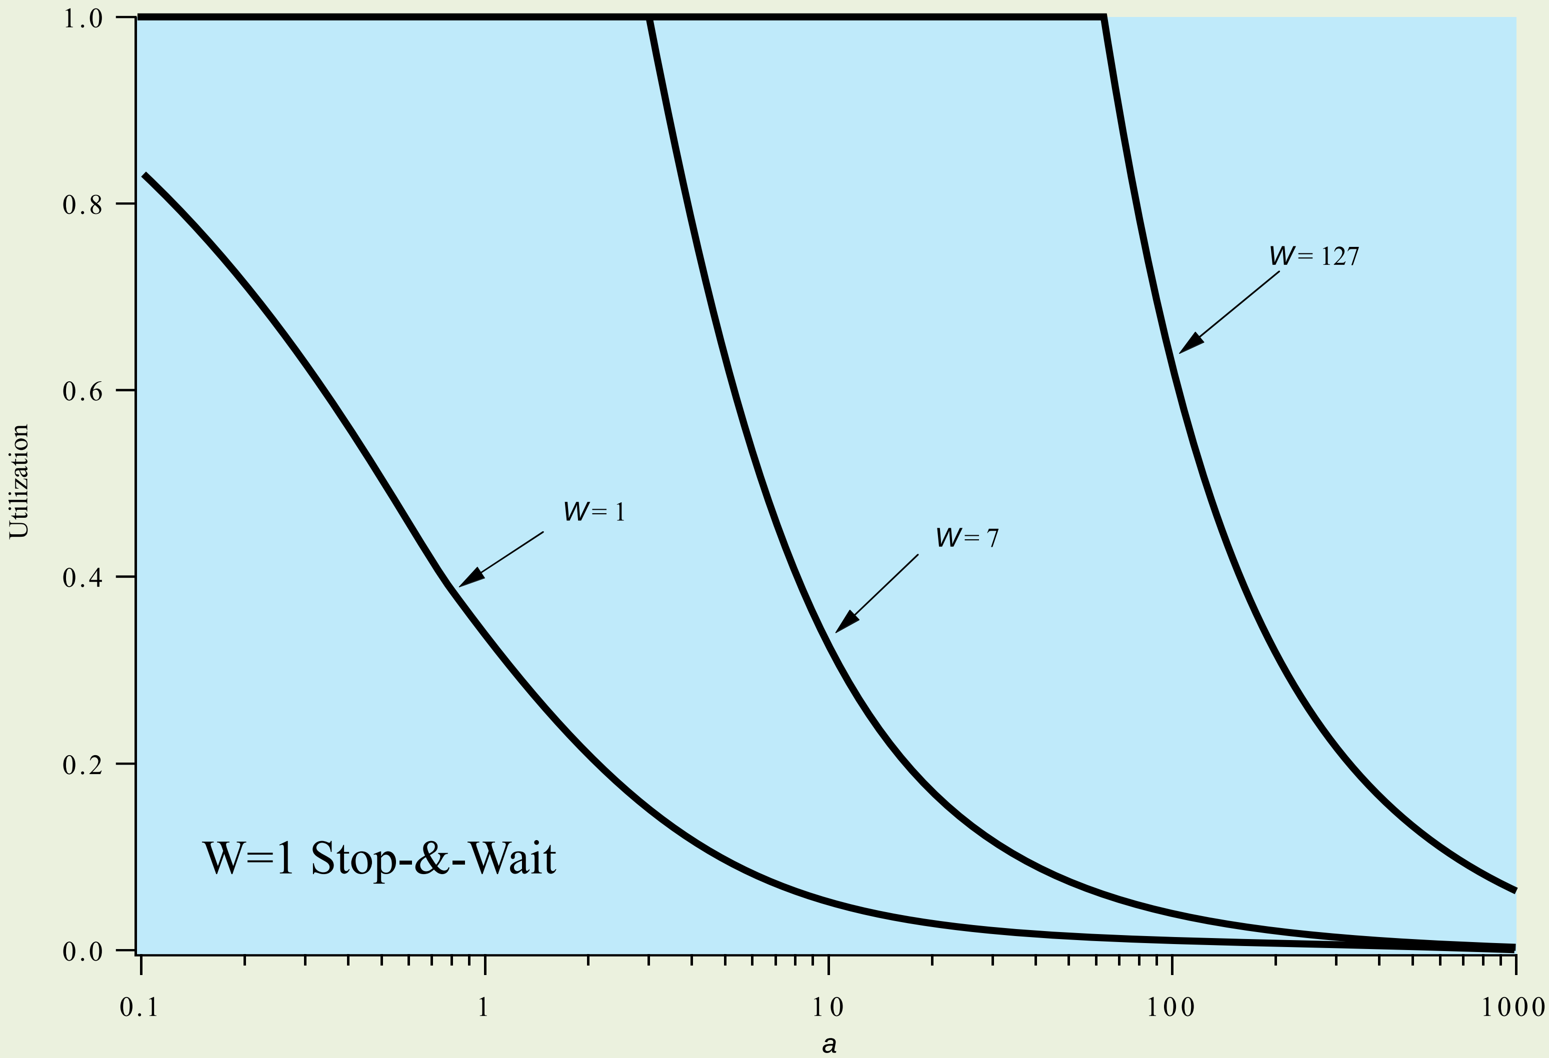

Figure 16.10 Sliding-Window Utilization as a function of *a*

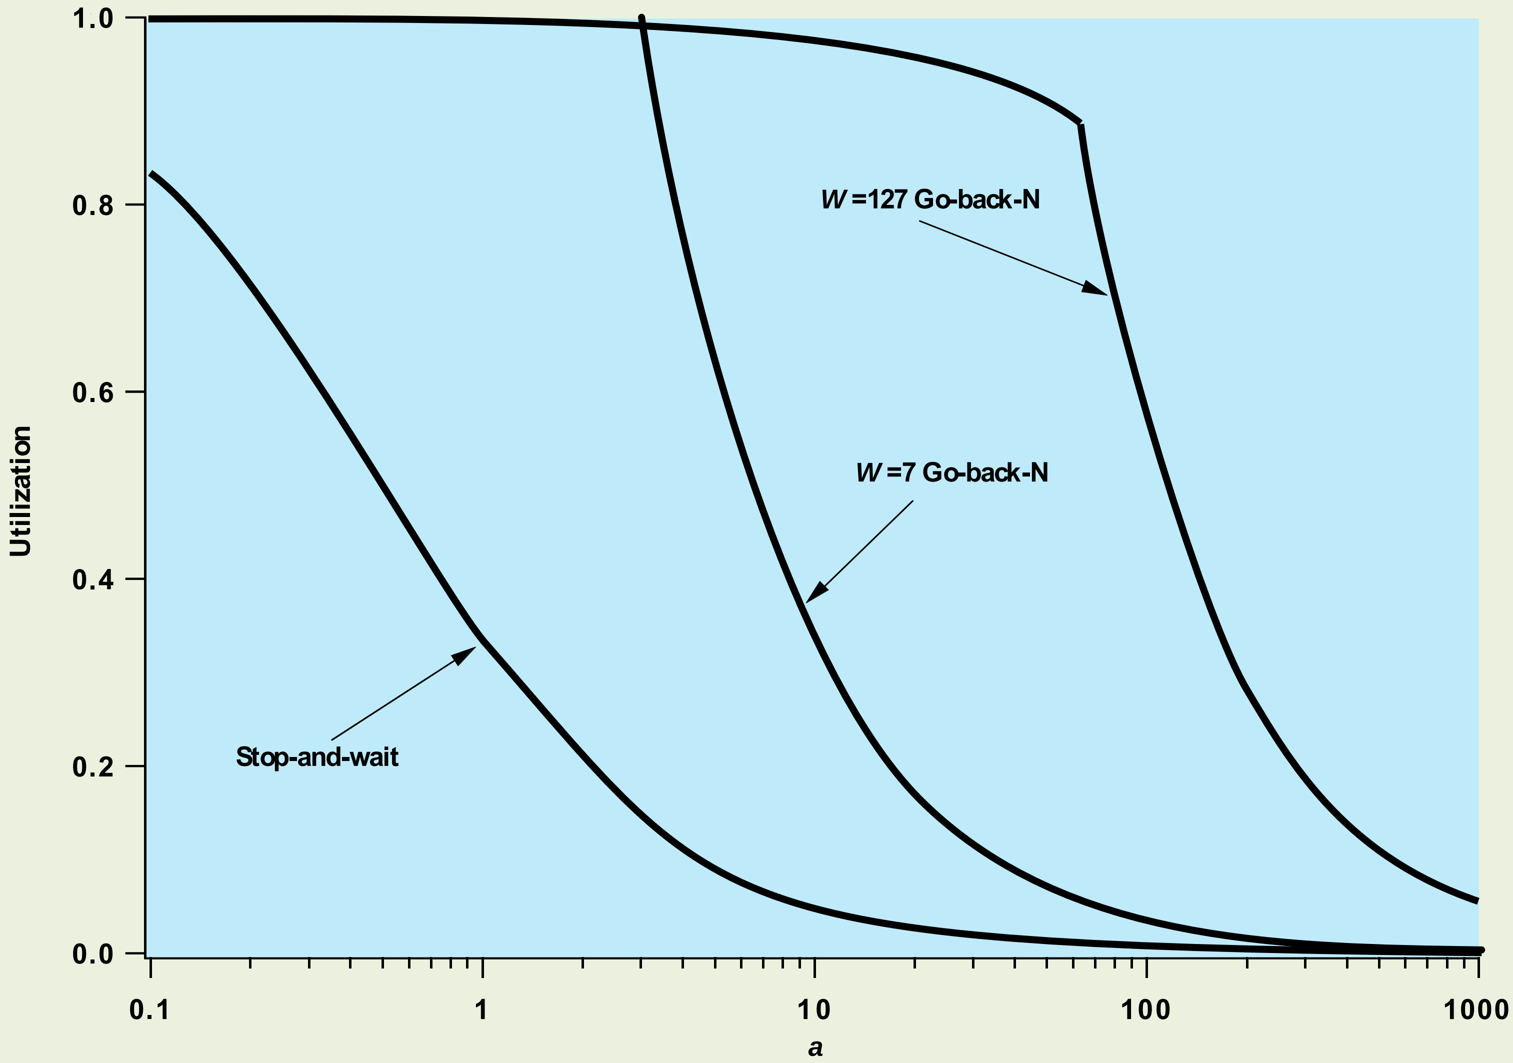

Figure 16.11 ARQ Utilization as a Function of *a* ($P = 10^{-3}$)

P = Propability of frame error

Delay and throughput: Finite buffers No congestion control

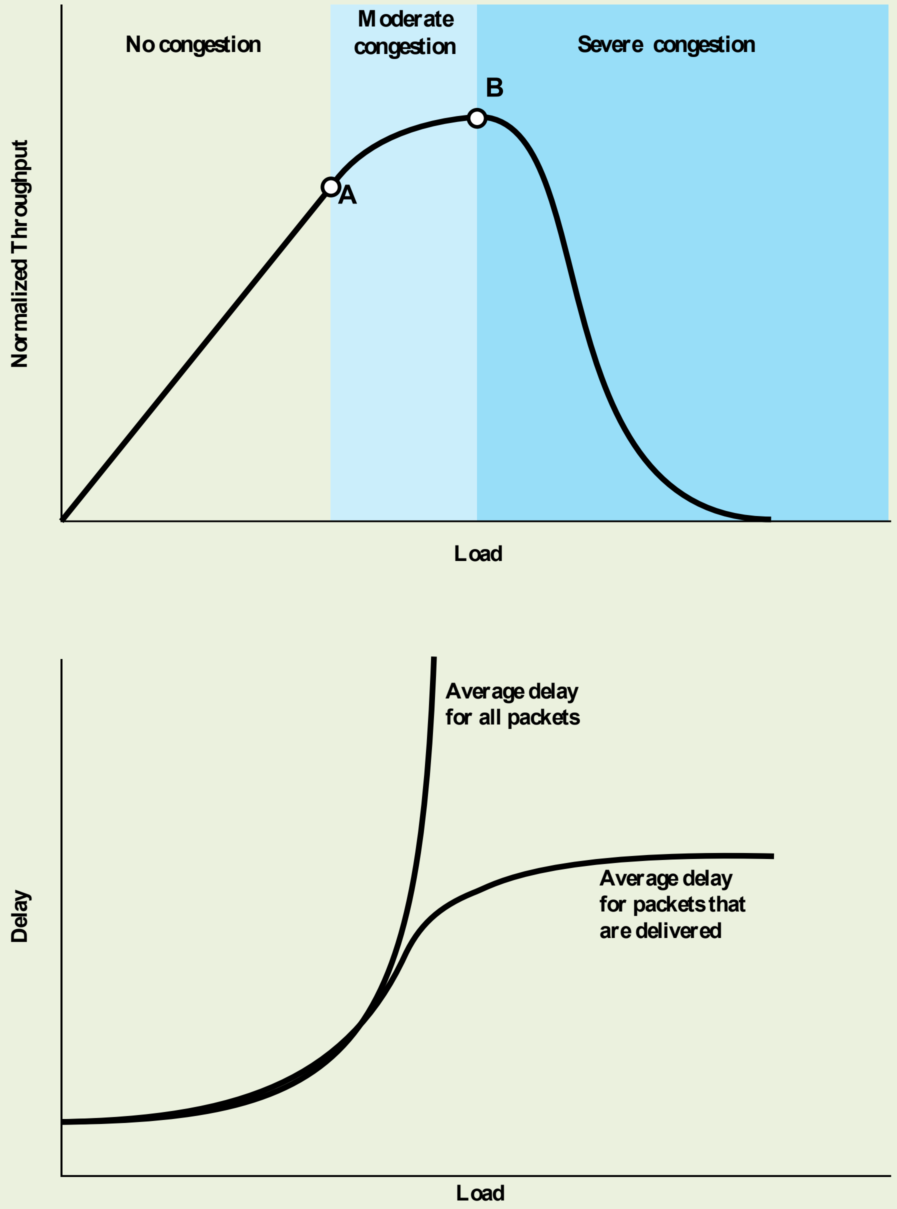

Figure 20.4 The Effects of Congestion

Packet loss

- Due to
  - Bit error in packet
    - Discard erronous packet
    - Link or Physical Layer?
  - Queue overflow
    - Discard packets
    - Node problems
- In real time multimedia late packets considered lost
- Packet loss ratio (%)
- Note TCP's sensitivity to packet loss

Congestion Control in Packet-Switching 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 information to packets in transit

- Either backwards or forwards

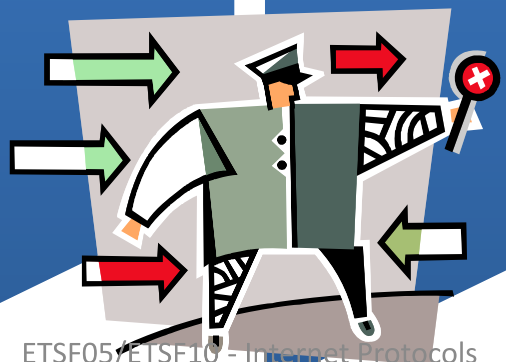

Explicit Congestion Signaling

- Backward
  - Congestion avoidance notification in opposite direction to packet required
- Forward
  - Congestion avoidance notification in same direction as packet required

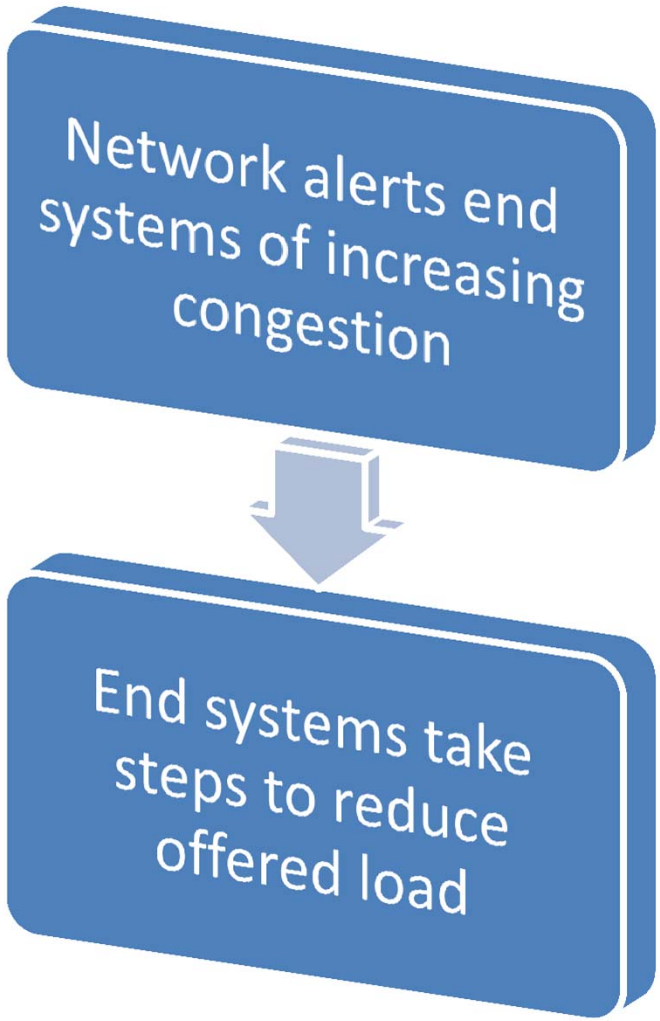

Congestion control

- Avoiding and eliminating congestion
  - Open-loop = proactive
  - Closed-loop = reactive

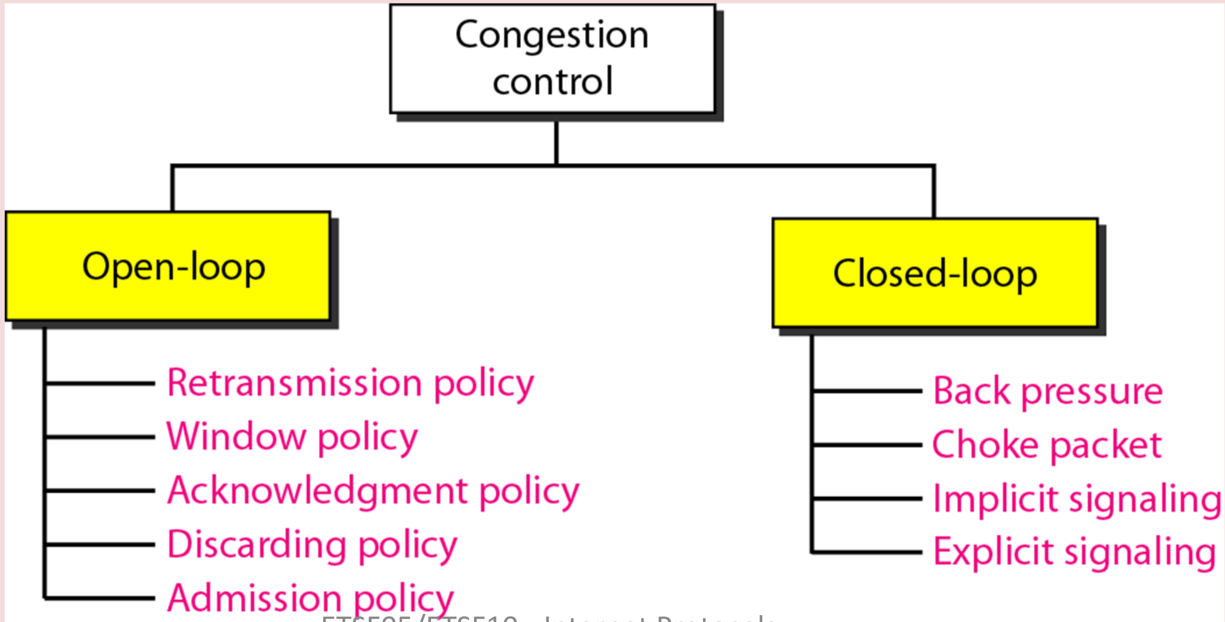

Closed-loop congestion control (1)

- Backpressure

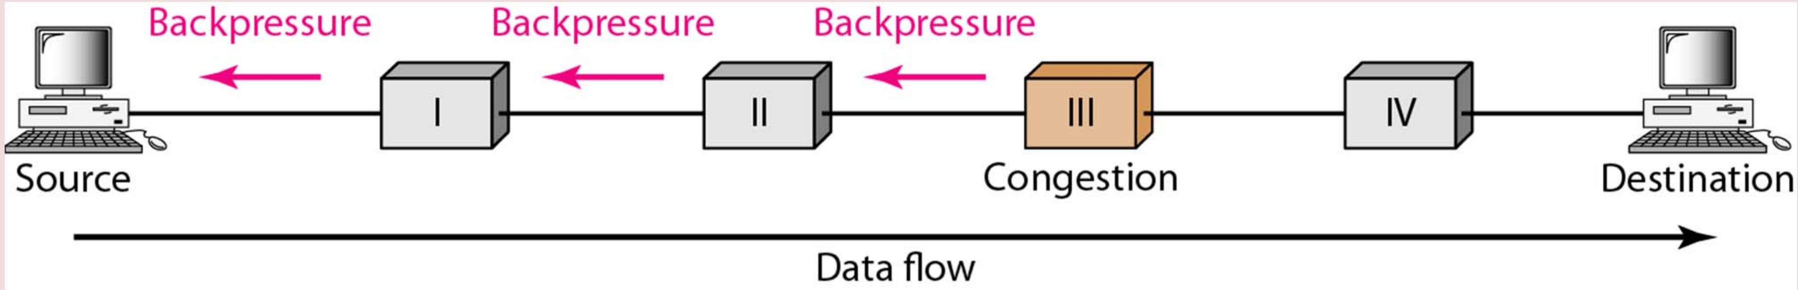

Closed-loop congestion control (2)

- Choke packet

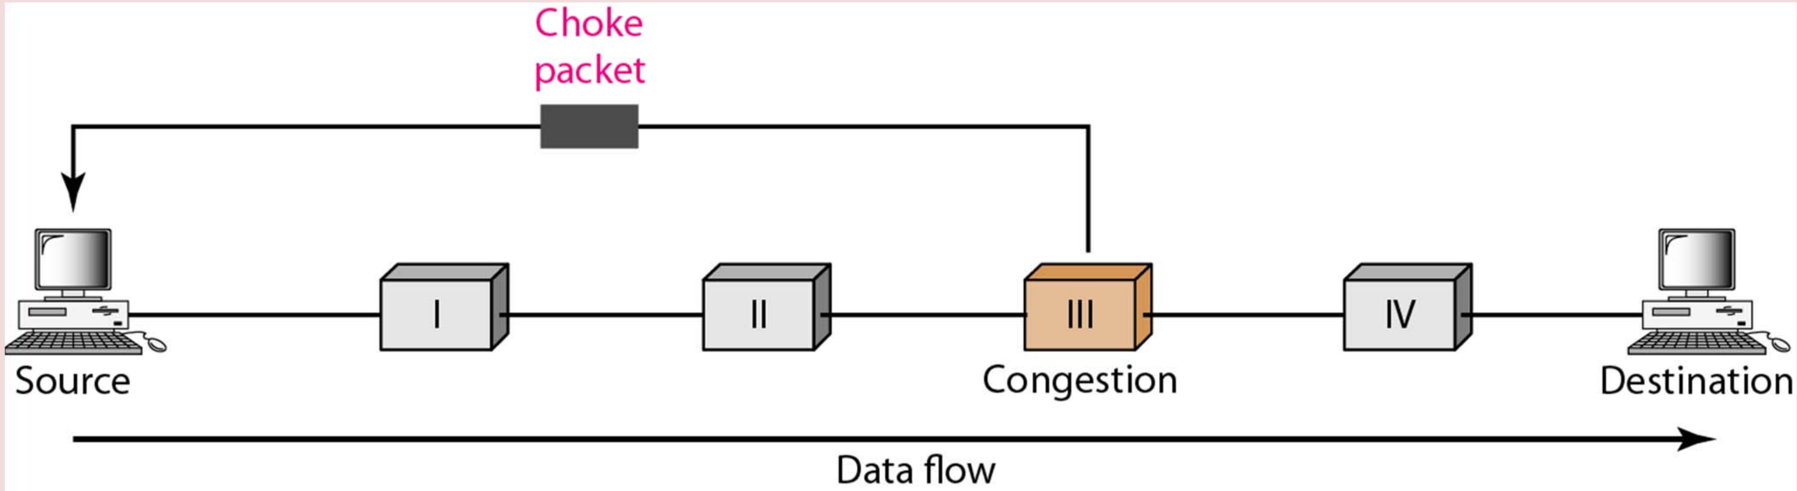

Implicit Congestion Signaling

- With network congestion:
  - Transmission delay increases
  - Packets may be discarded
- Source can detect congestion and reduce flow
- Responsibility of end systems
- Effective on connectionless (datagram) networks
- Also used in connection-oriented networks

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

How to improve QoS

- Admission control
- Resource reservation
- **Scheduling**
- **Traffic shaping**

- **Routing?**

Where to improve QoS?

- Admission control
  - INTSERV, DIFFSERV
- Resource reservation
  - RSVP

- **Scheduling**
- **Traffic shaping**

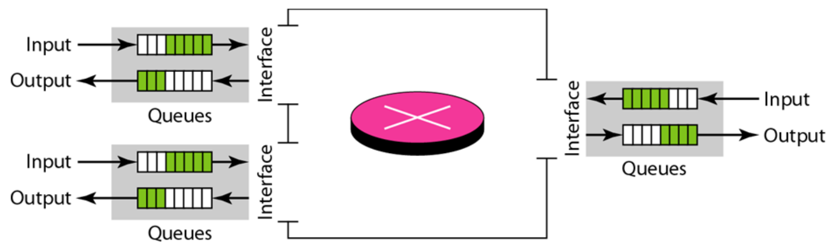

ANYWHERE YOU FIND QUEUES!

Scheduling: *FIFO queuing*

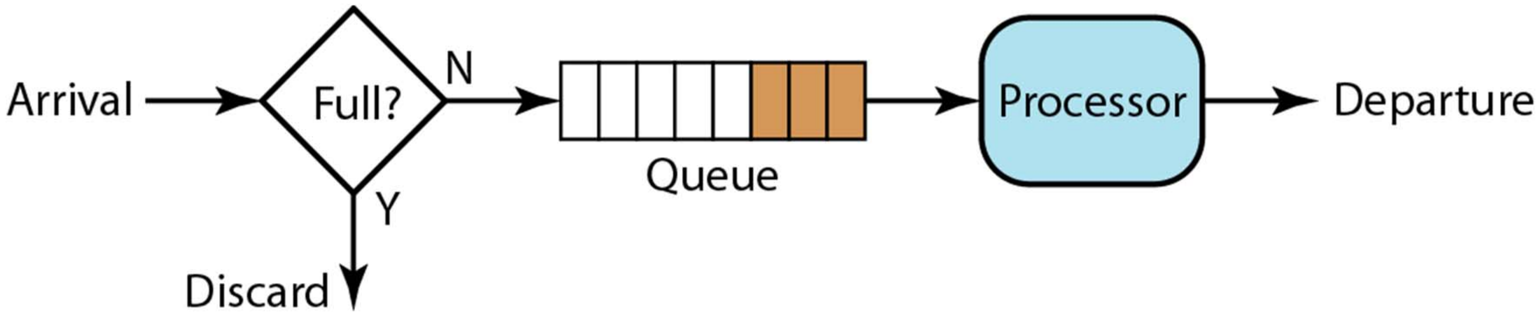

Scheduling: *Priority queuing*

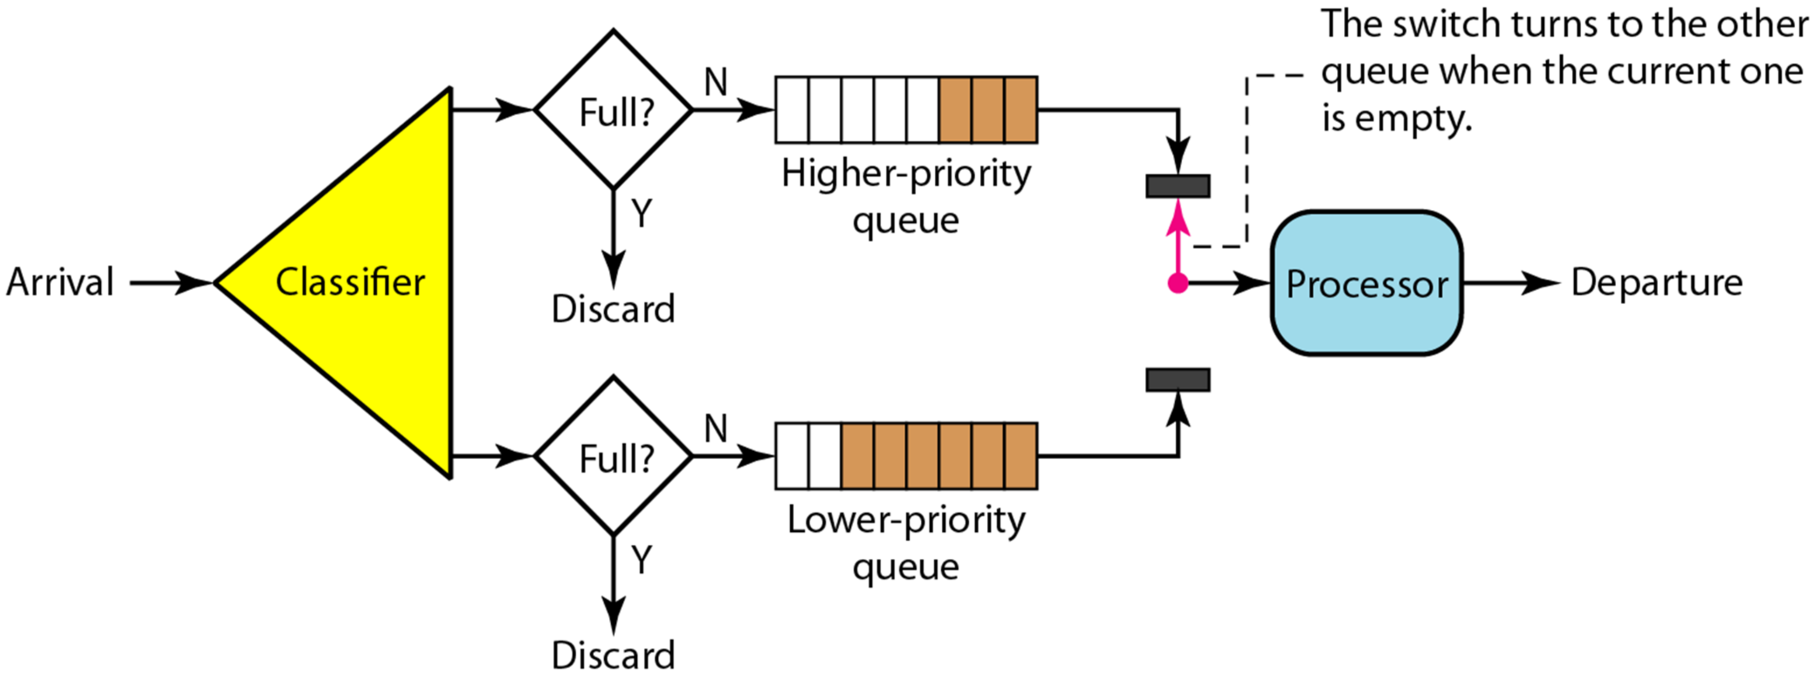

Scheduling: *Weighted fair queuing*

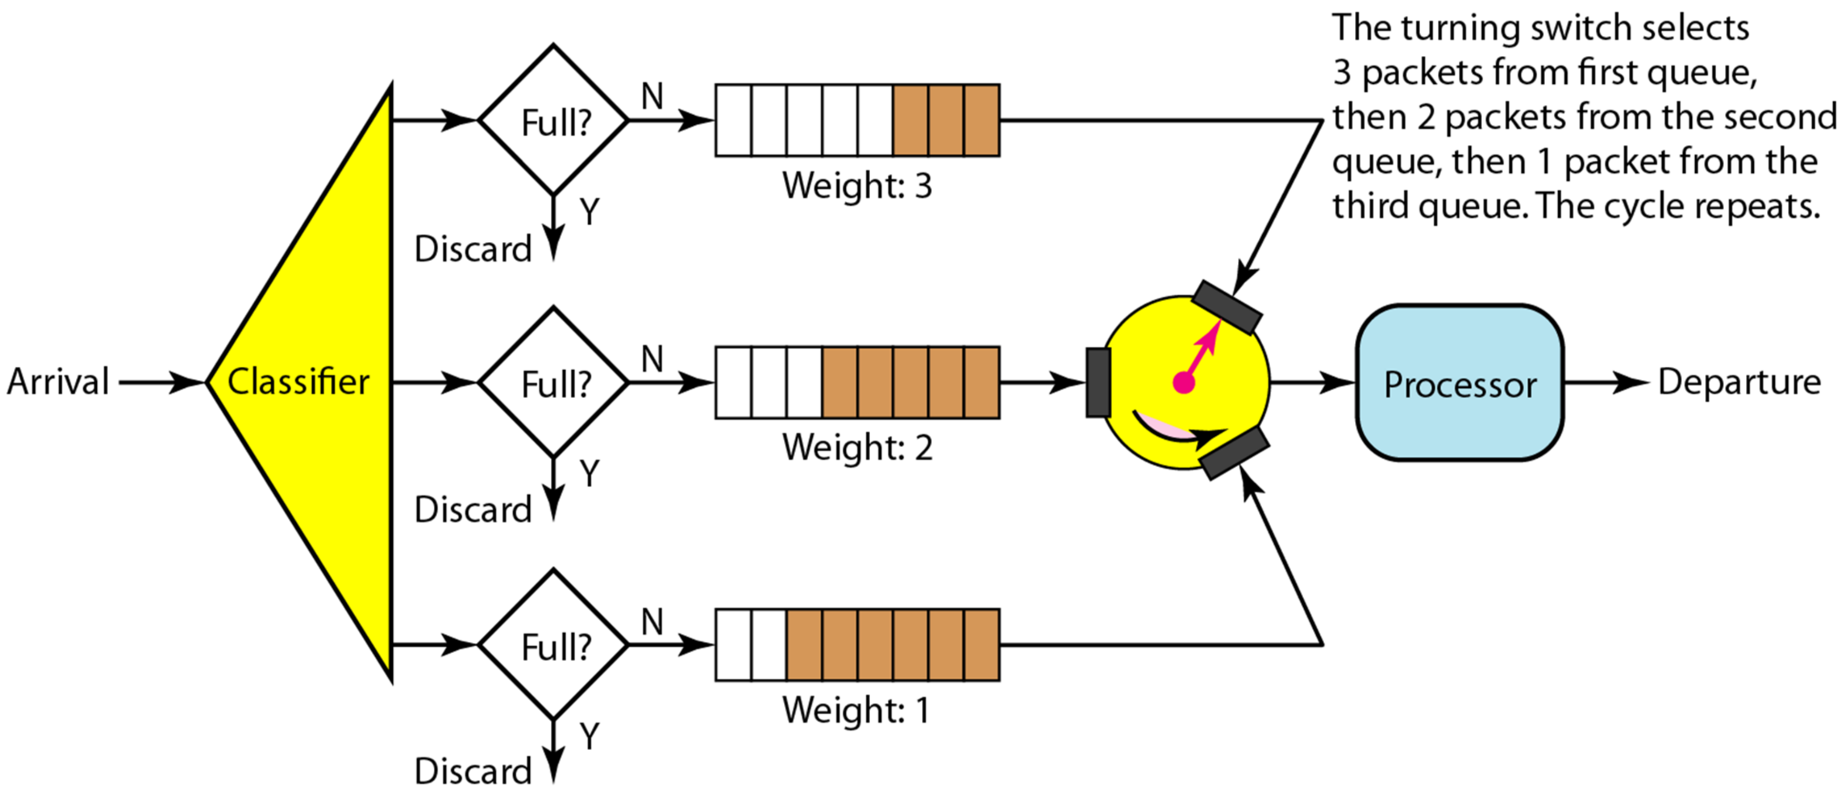

Traffic Shaping/Traffic Policing

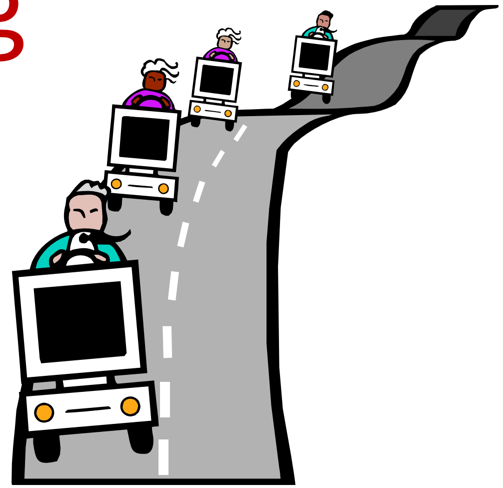

- Two important tools in network management:
  - **Traffic shaping**
    - Concerned with traffic leaving the switch
    - Reduces packet clumping
    - Produces an output packet stream that is less burst and with a more regular flow of packets
  - Traffic policing
    - Concerned with traffic entering the switch
    - Packets that don’t conform may be treated in one of the following ways:
      - Give the packet lower priority compared to packets in other output queues
      - Label the packet as nonconforming by setting the appropriate bits in a header
      - Discard the packet

Traffic Management

Fairness

- Provide equal treatment of various flows

Quality of service

- Different treatment for different connections

Reservations

- Traffic contract between user and network
- Excess traffic discarded or handled on a best-effort basis

Support from routing protocols?

Yes!

- Optimal path
  - Single metric
  - Multiple metrics?
- Multiple Equal Cost Paths
  - Load sharing
  - Load balancing
- QoS routing
  - Cross-layer approaches
  - OSPF extensions (RFC2676)

Well, sort of!

- Applies to all traffic
  - No differentiation between flow types
- What about inter-domain?
  - No control over network resources
- More sophisticated mechanisms needed
  - Multiprotocol Label Switching (MPLS)
  - Resource Reservation

Layer3 congestion avoidance

- Congestion = data load > network capacity
  - Arrival rate > processing rate
  - Processing rate > departure rate
- A simple method
  - Random early discard (RED)
  - UDP?

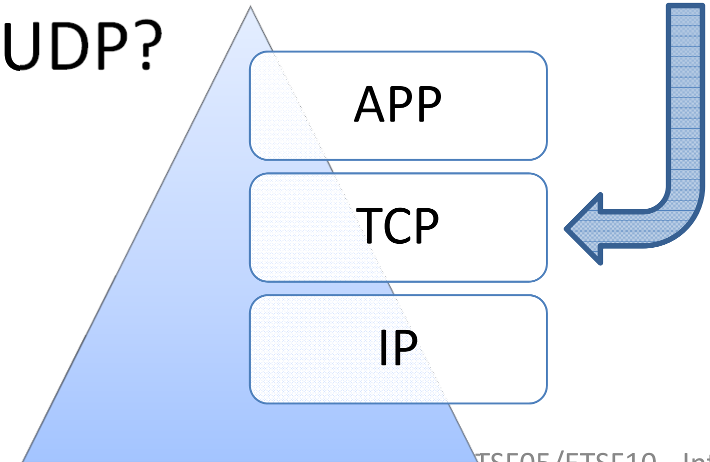

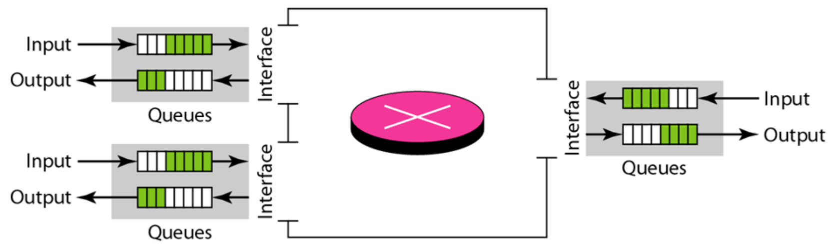

Traffic descriptors

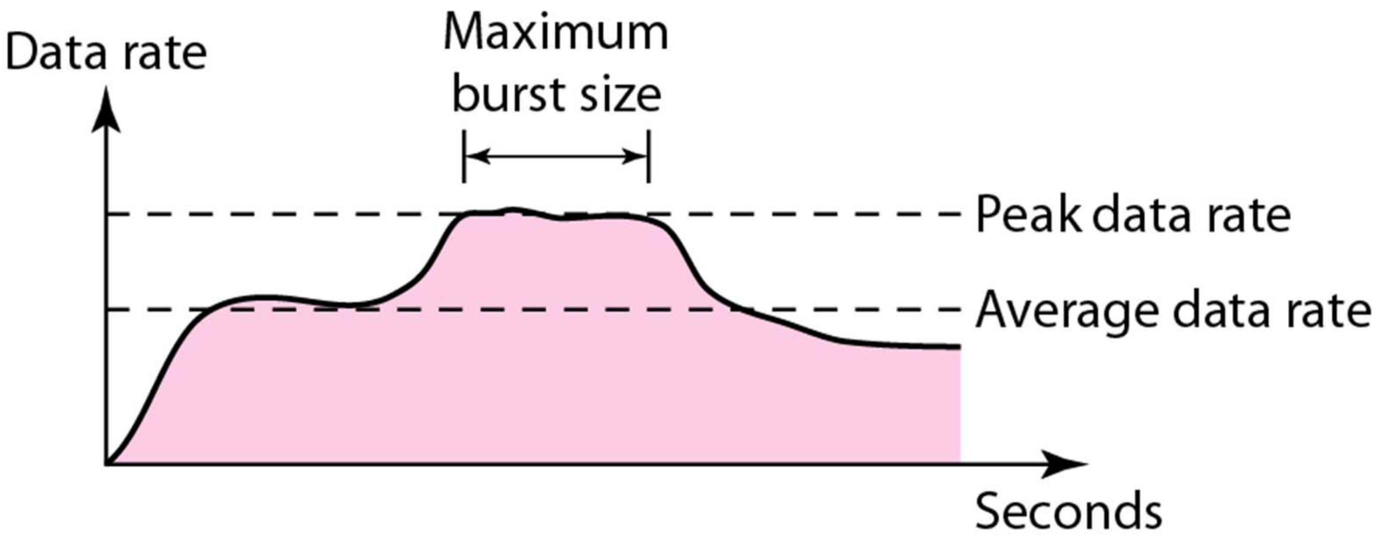

Effective bandwidth requirement

Traffic profiles

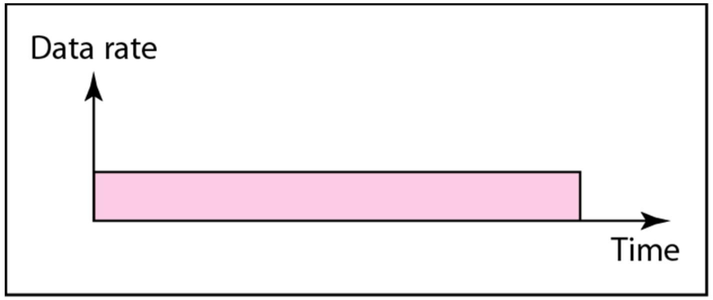

a. Constant bit rate

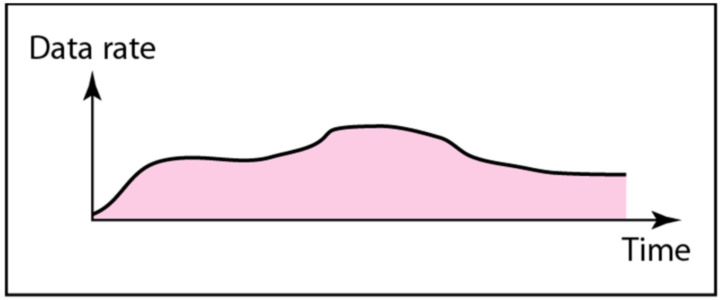

b. Variable bit rate

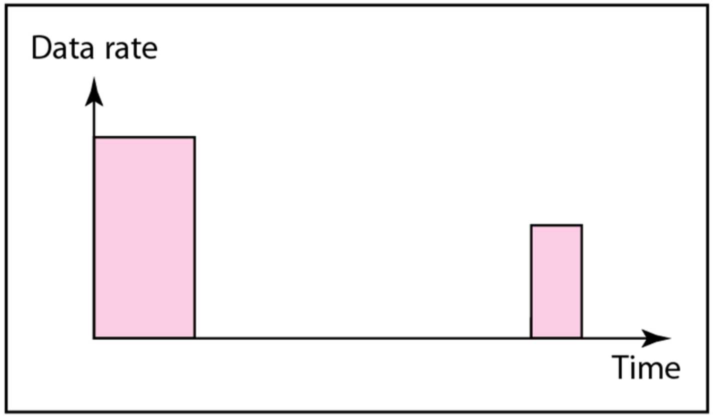

c. Bursty

Traffic shaping: *Two approaches*

Leaky bucket

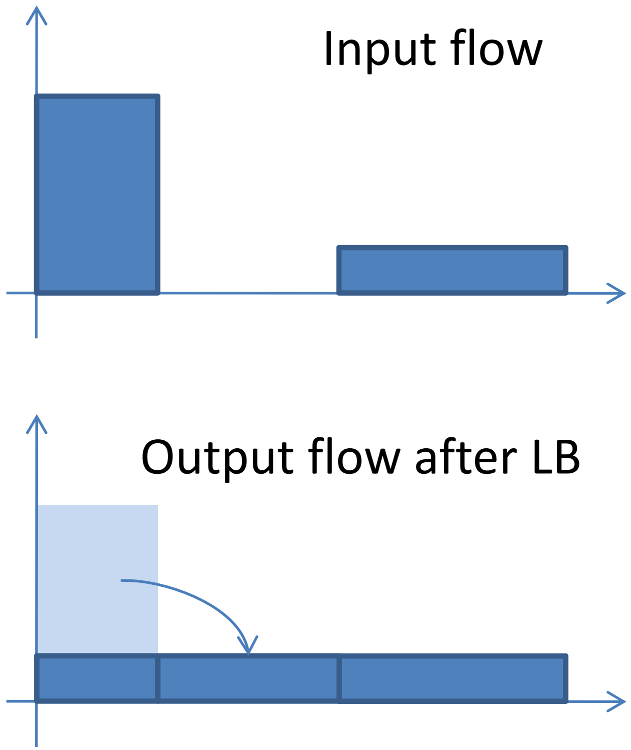

Token bucket

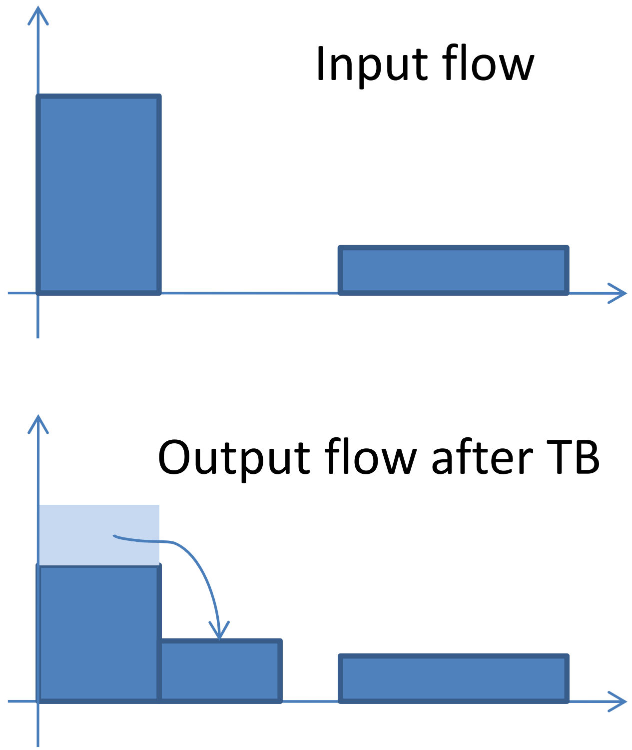

Traffic shaping: *Leaky bucket*

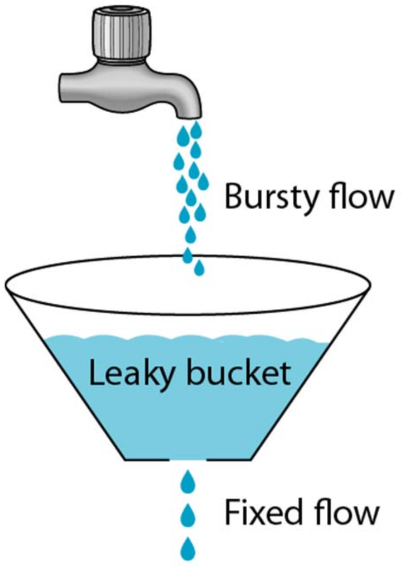

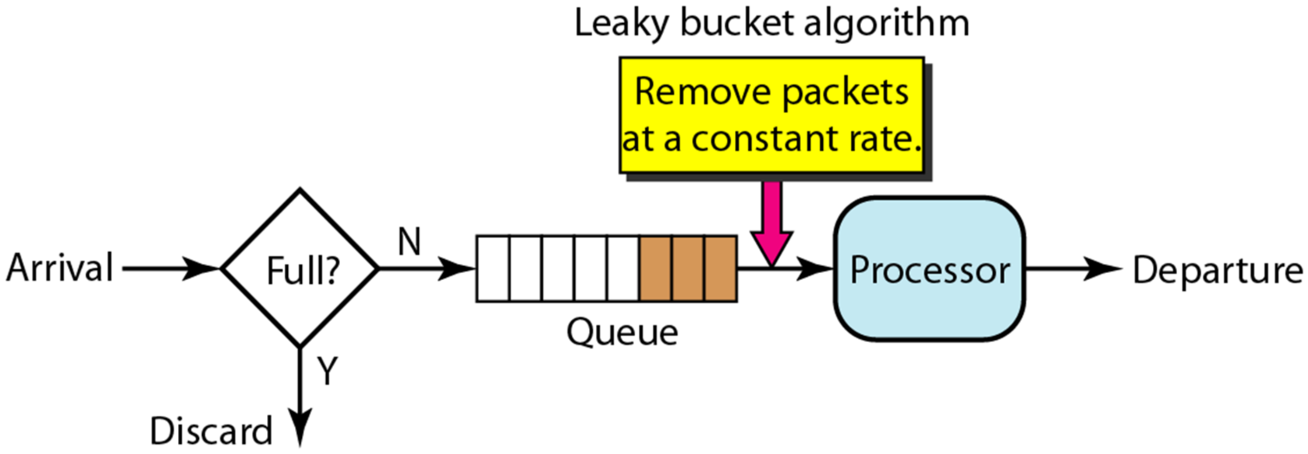

See also Figure 20.7

Token Bucket

- Widely used traffic management tool
- Advantages:
  - Many traffic sources can be defined easily and accurately
  - Provides a concise description of the load to be imposed by a flow, enabling the service to determine easily the resource requirement
  - Provides the input parameters to a policing function

Traffic shaping: *Token bucket*

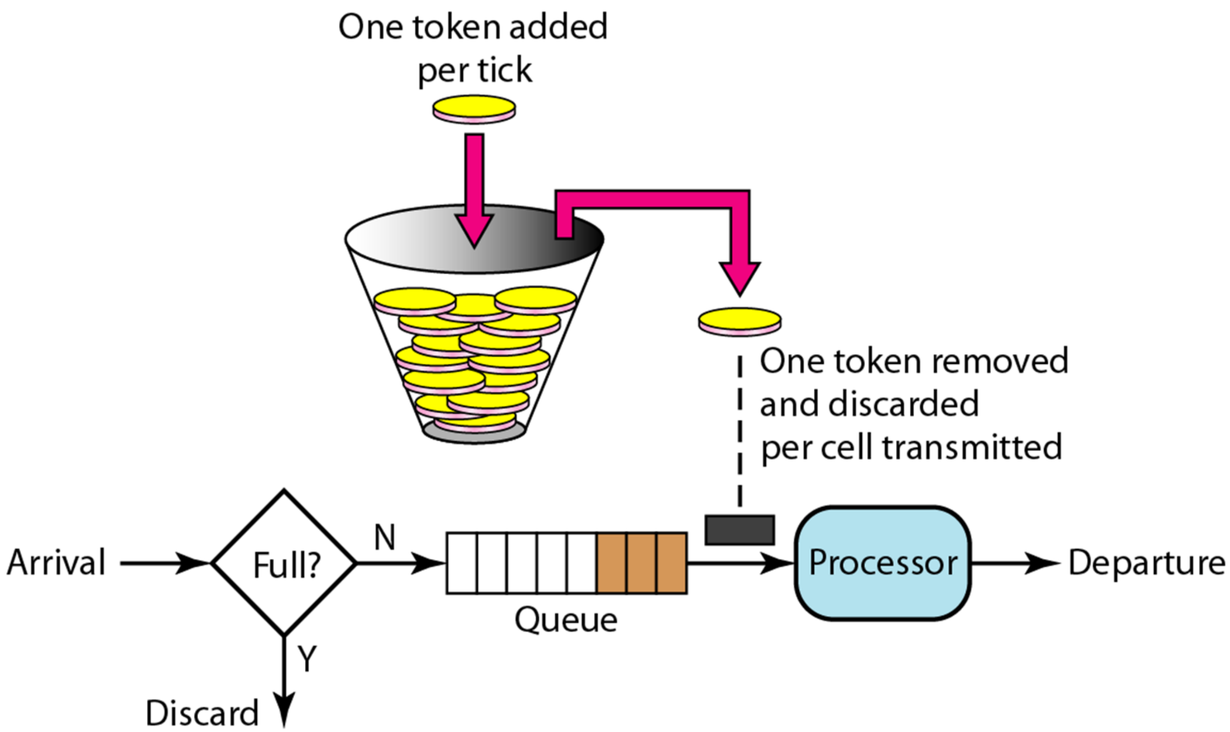

See also Figure 20.6

Bonus: QoE, Quality of Experience

- The user’s subjective perception of the presentation of the content
- Mean Opinion Score
- Research for to find objective measures
  - Full reference
  - No reference
  - Hybrid
- What QoS give what QoE