LANs
Local Area Networks
focused on the
Media Access Control (MAC) Sub Layer

Advanced Computer Networks
C13

LANs Outline

- Channel Allocation Problem
- Relative Propagation Time
- LAN Utilization Upper Bound
- Multiple Access Protocols
  - TDMA, FDMA
  - Aloha, Slotted Aloha
  - CSMA (non-persistent, 1-persistent, p-persistent), CSMA/CD
  - Performance Results

Local Area Networks

- Aloha
- Slotted Aloha
- CSMA
  - non-persistent
  - 1-persistent
  - p-persistent
- CSMA/CD
- Ethernet
- Token Ring

Data Link Sub Layers

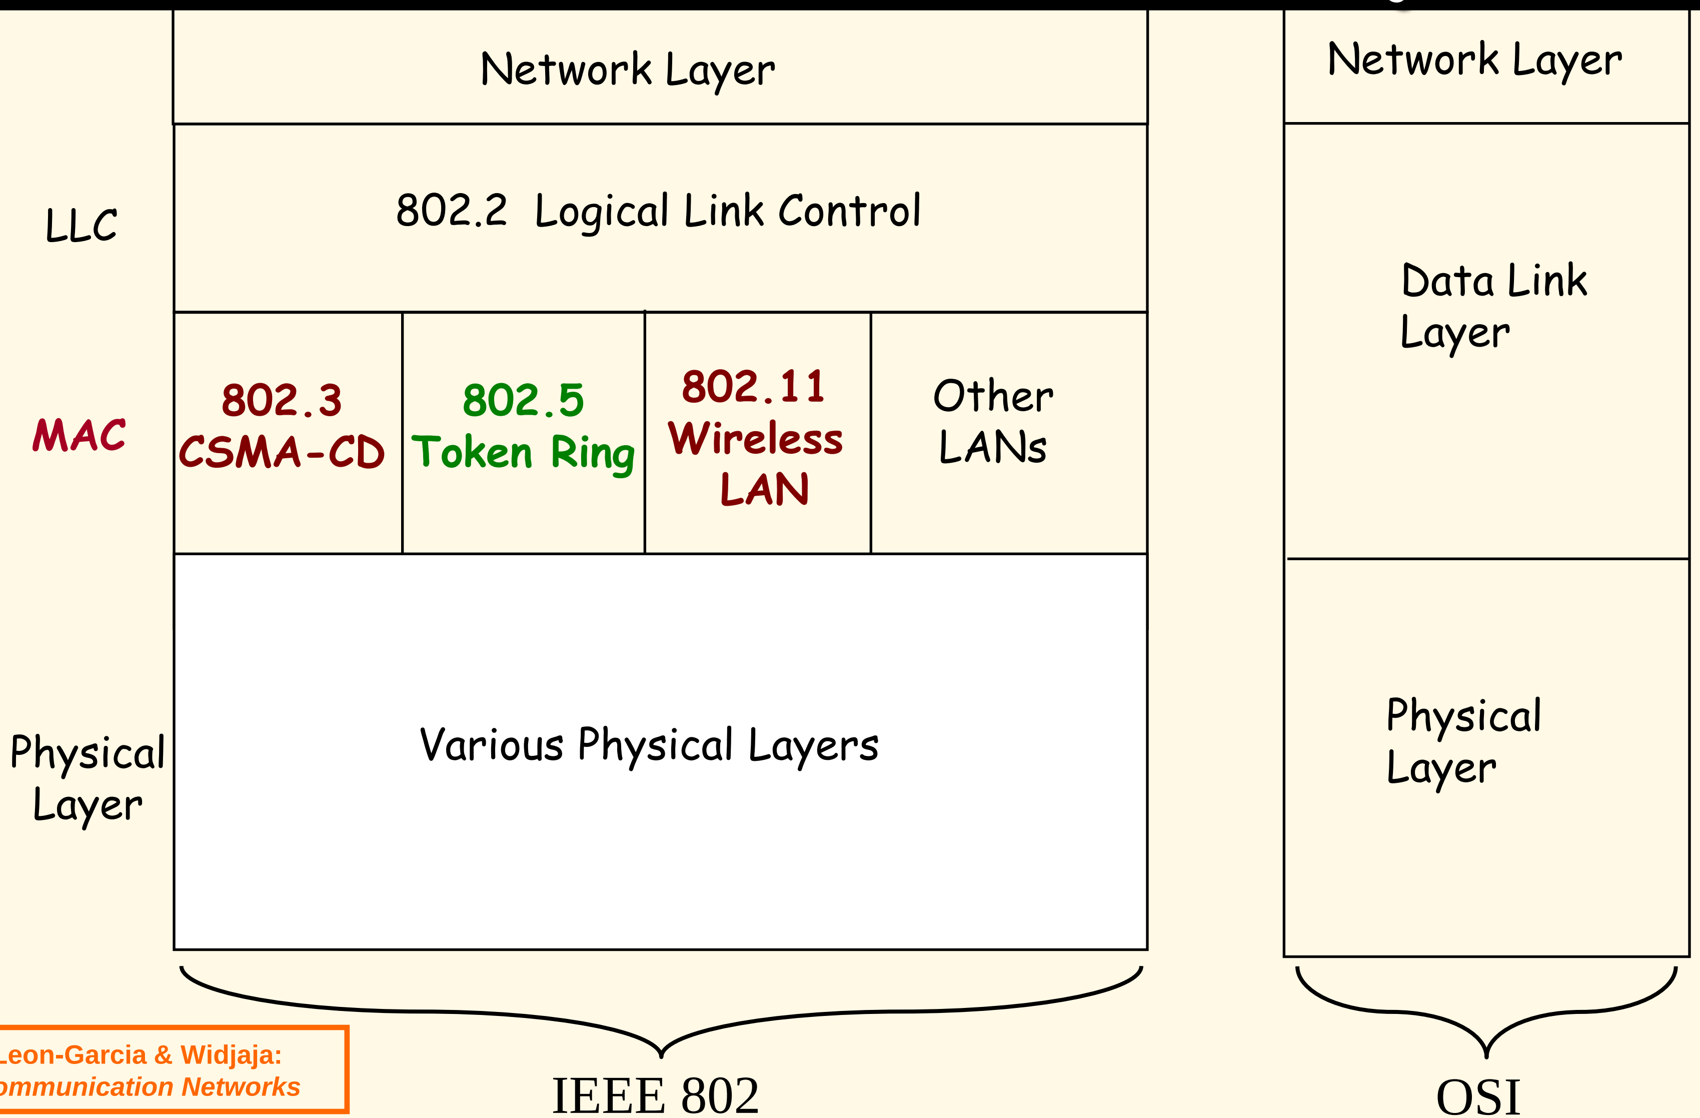

Leon-Garcia & Widjaja:
Communication Networks

Media Access Control (MAC)

- Can divide networks into two categories:
  - Point-to-Point
  - Using broadcast channels*

*deal here with broadcast channels and their protocols

- MAC sub-layer is important in LANs, MANs and wireless networks (both WiFi and cellular)!!

Channel Access Abstraction

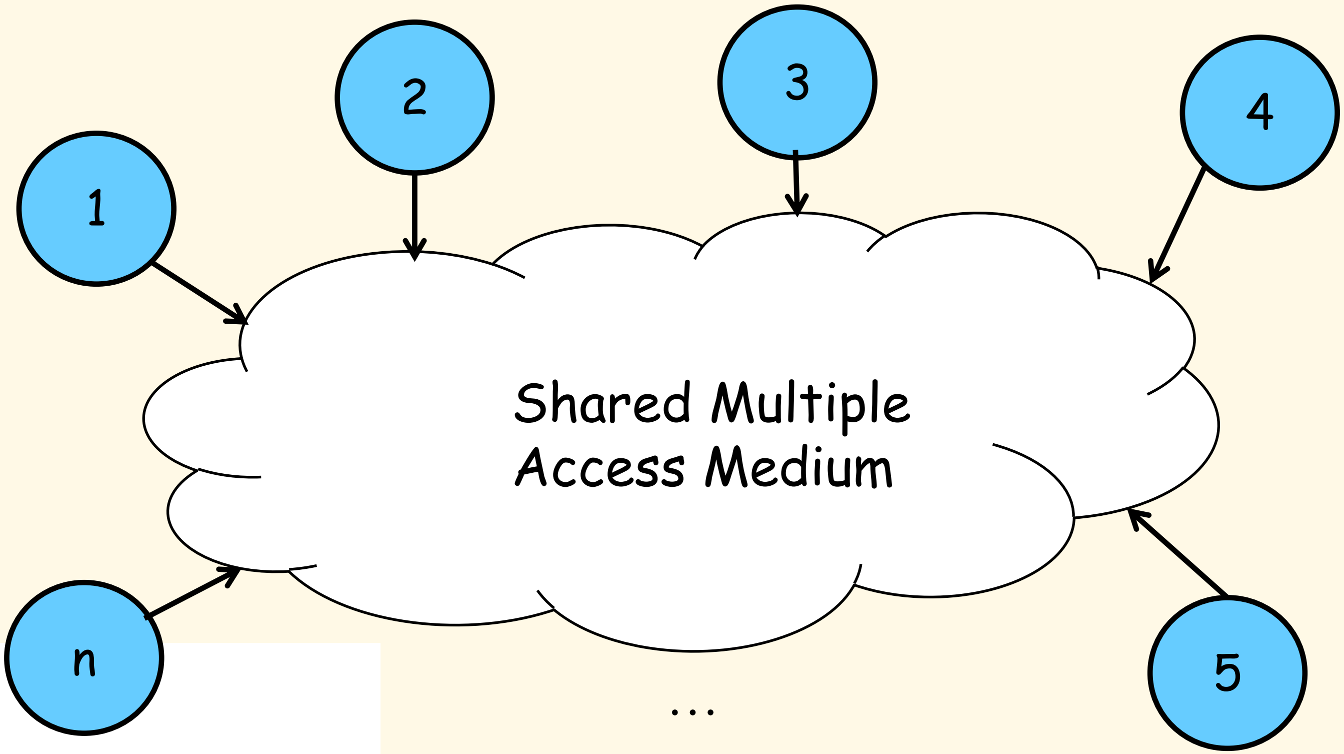

Leon-Garcia & Widjaja:
Communication Networks

Static Channel Allocation Problem

The history of broadcast networks includes satellite and packet radio networks.

Let us view a satellite as a repeater amplifying and rebroadcasting everything that comes in.

To generalize this problem, consider networks where every frame sent is *automatically* received by every site (node).

Satellite Channel

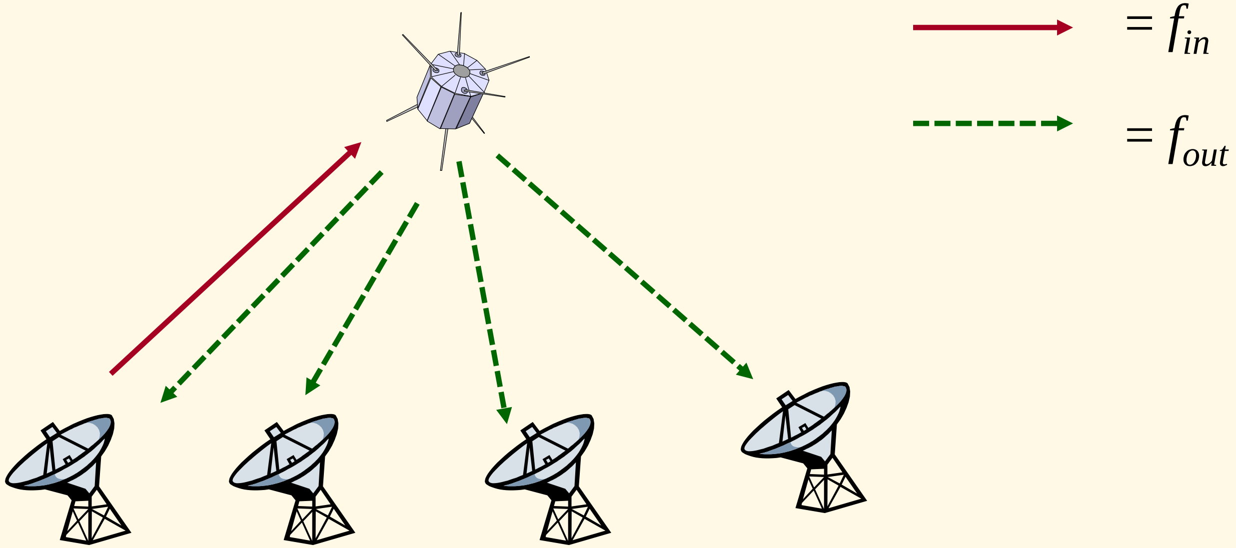

Leon-Garcia & Widjaja: *Communication Networks*

Static Channel Allocation Problem

We model this situation as **n** independent users (one per node), each wanting to communicate with another user and they have no other form of communication.

The Channel Allocation Problem

To manage a single broadcast channel which must be shared **efficiently** and **fairly** among **n** uncoordinated users.

Specific LAN Topologies

Ring networks

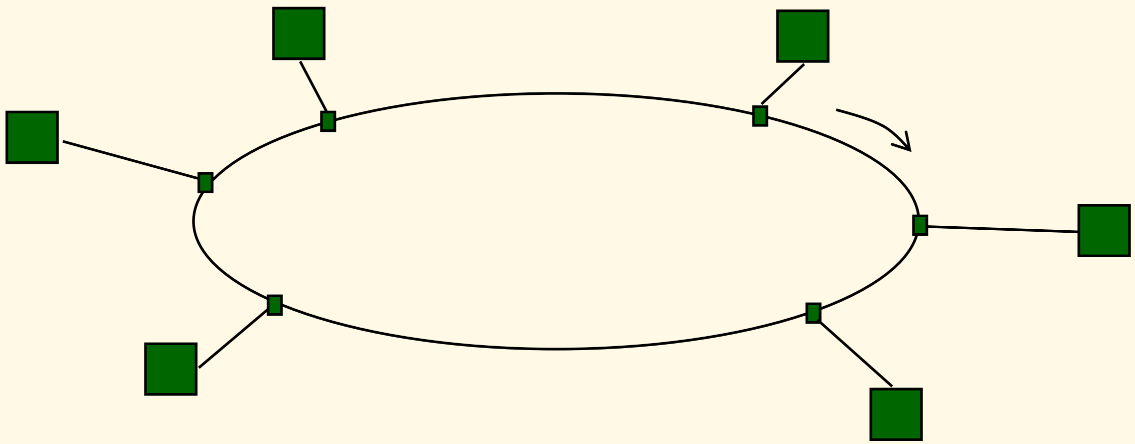

Multitapped Bus Networks

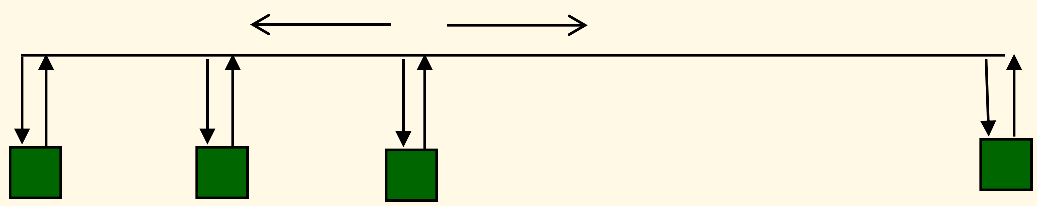

Leon-Garcia & Widjaja:
Communication Networks

Possible Model Assumptions

0. Listen property :: (applies to satellites)

   The sender is able to listen to sent frame one round-trip after sending it.

   → no need for explicit ACKs.

1. The model consists of n independent stations.

2. A single channel is available for communications.

Possible Model Assumptions

3. Collision Assumption :: If two frames are transmitted simultaneously, they overlap in time and the resulting signal is garbled. This event is a collision.

4a. Continuous Time Assumption :: frame transmissions can begin at any time instant.

4b. Slotted Time Assumption :: time is divided into discrete intervals (slots). Frame transmissions always begin at the start of a time slot.

Possible Model Assumptions

5a. **Carrier Sense Assumption (CS) ::**
Stations can tell if the channel is busy (in use) before trying to use it. If the channel is busy, no station will attempt to use the channel until it is idle.

5b. **No Carrier Sense Assumption ::**
Stations are unable to sense channel before attempting to send a frame. They just go ahead and transmit a frame.

a :: Relative Propagation Time

$$a = \frac{\text{length of the data path (in bits)}}{\text{length of a standard frame (in bits)}}$$

-OR-

$$a = \frac{\text{propagation time (in seconds)}}{\text{transmission time (in seconds)}}$$

-OR-

$$a = \frac{\text{bandwidth-delay product}^{*}}{\text{average frame size}}$$

[LG&W def p.346]

* bandwidth-delay product :: the product of the capacity (bit rate) and the delay.

Effect of a on Utilization

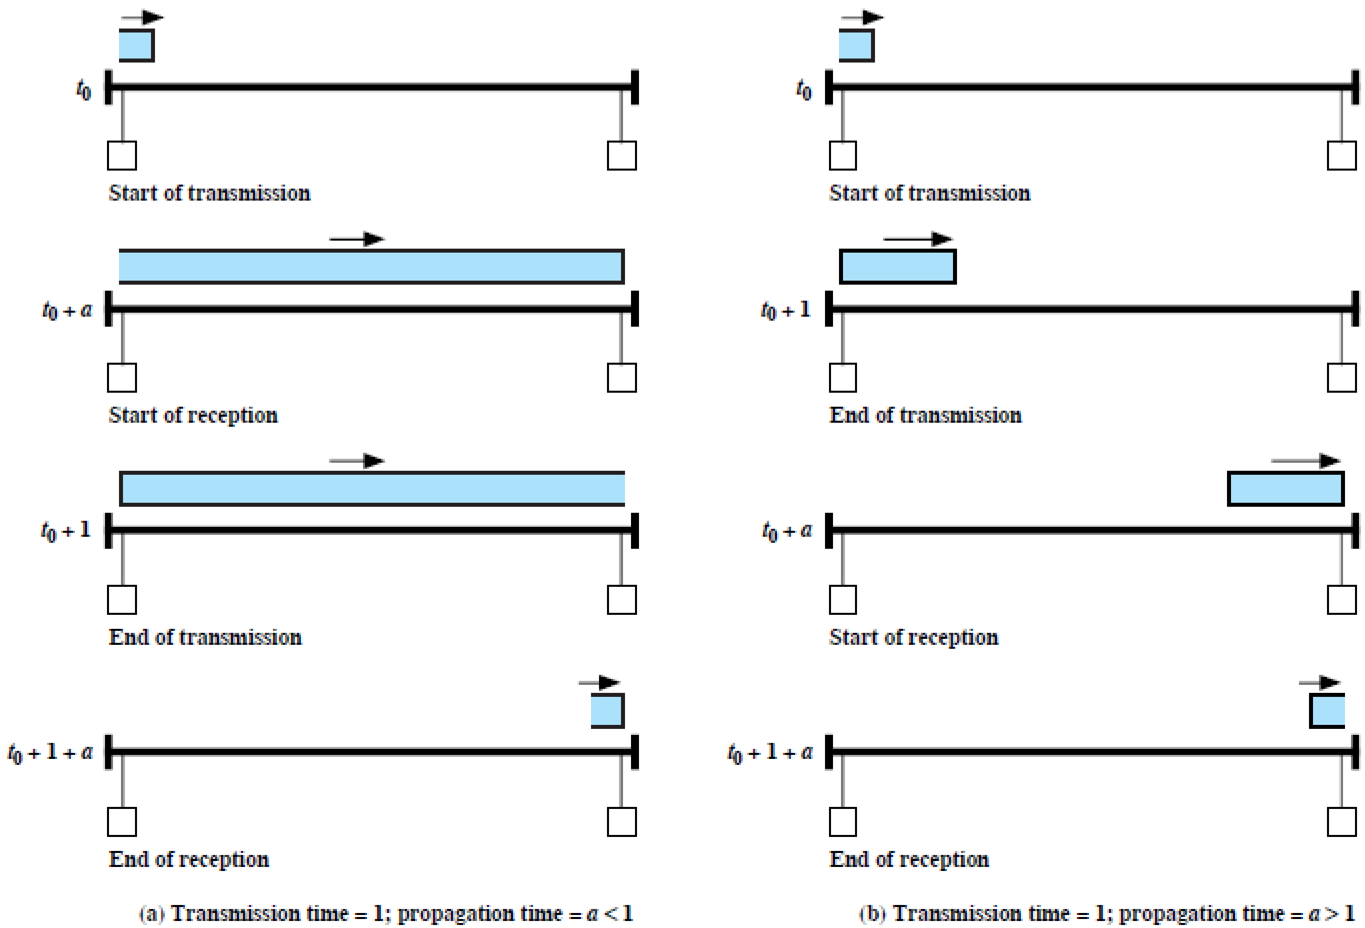

Figure 16.14 The Effect of a on Utilization for Baseband Bus

DCC 8th Ed.
Stallings

Relative Propagation Time

R = capacity (data rate)

d = maximum distance of communications path

v = propagation velocity (Assume v = 2/3 speed of light
2×10^8 meters/second)

L = frame length

$$a = \frac{d / v}{L/R} = \frac{Rd}{vL}$$

Upper Bound on Utilization for Shared Media LAN

Assume a perfect, efficient access that allows one transmission at a time where there are no collisions, no retransmissions, no delays between transmissions and no bits wasted on overhead. {These are best-case assumptions!!}

$$\text{Util} = \frac{\text{Tput}}{\text{Capacity}} = \frac{L}{\textit{propagation time} + \textit{transmission time}} \Big/ R$$

Best Case LAN Utilization

$$\text{max. Util} = \frac{\dfrac{L}{R}}{\dfrac{d}{v} + \dfrac{L}{R}} = \frac{L}{\dfrac{Rd}{v} + L} = \frac{1}{a + 1}$$

$$\text{max. Util} = \frac{1}{1 + a}$$

Worst Case Collision Scenario

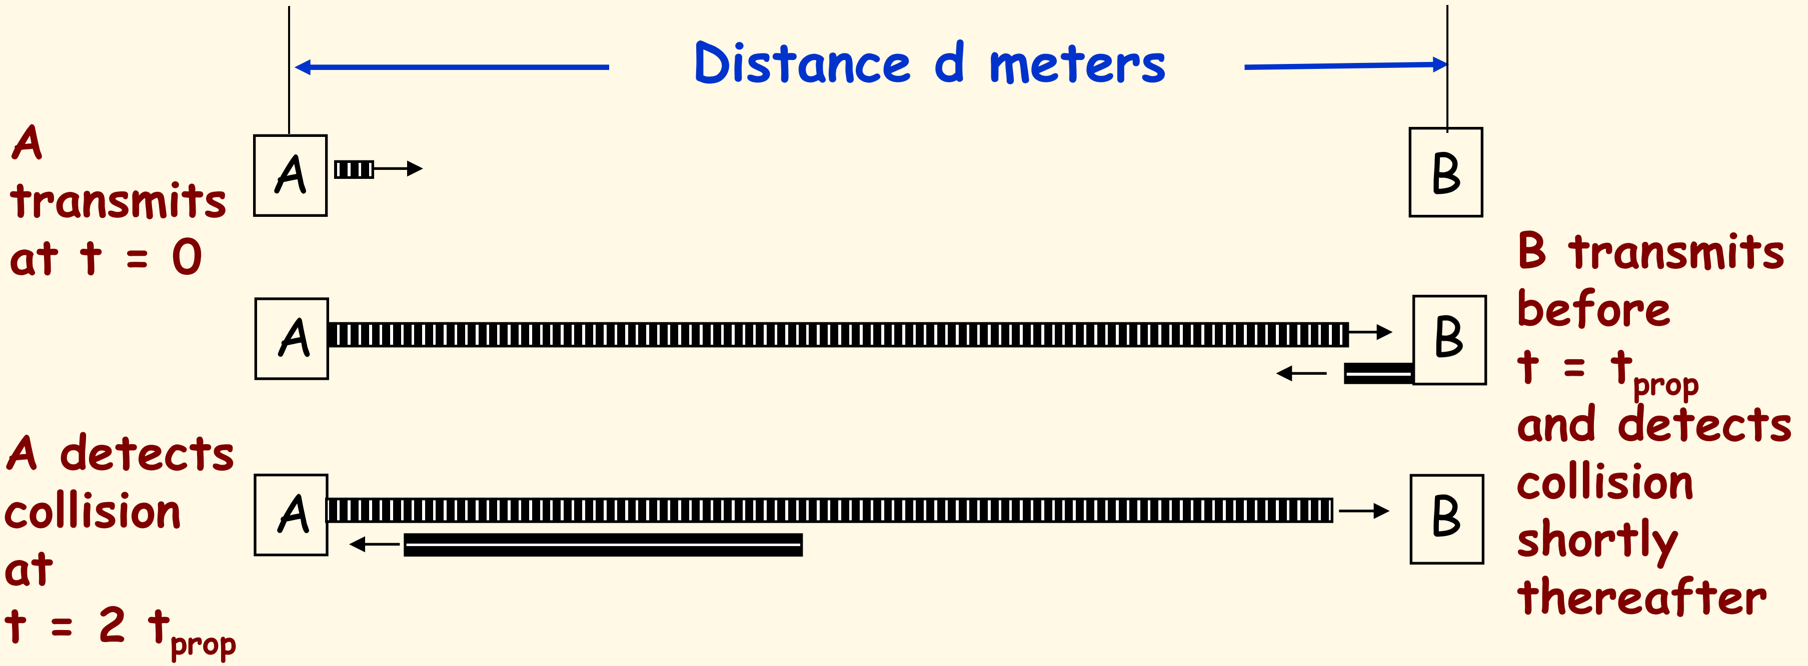

$t_{prop} = d / \nu$ seconds

Leon-Garcia & Widjaja: *Communication Networks*

LAN Design Performance Issues

For broadcast LANs what factors under the designer's control affect LAN performance?

- Capacity {function of media}
- Propagation delay {function of media, distance}
- Bits /frame (frame size)
- MAC protocol
- Offered load {depends on retransmission scheme}
- Number of stations
- Bit error rate {function of media}

Typical Frame Delay versus Throughput Performance

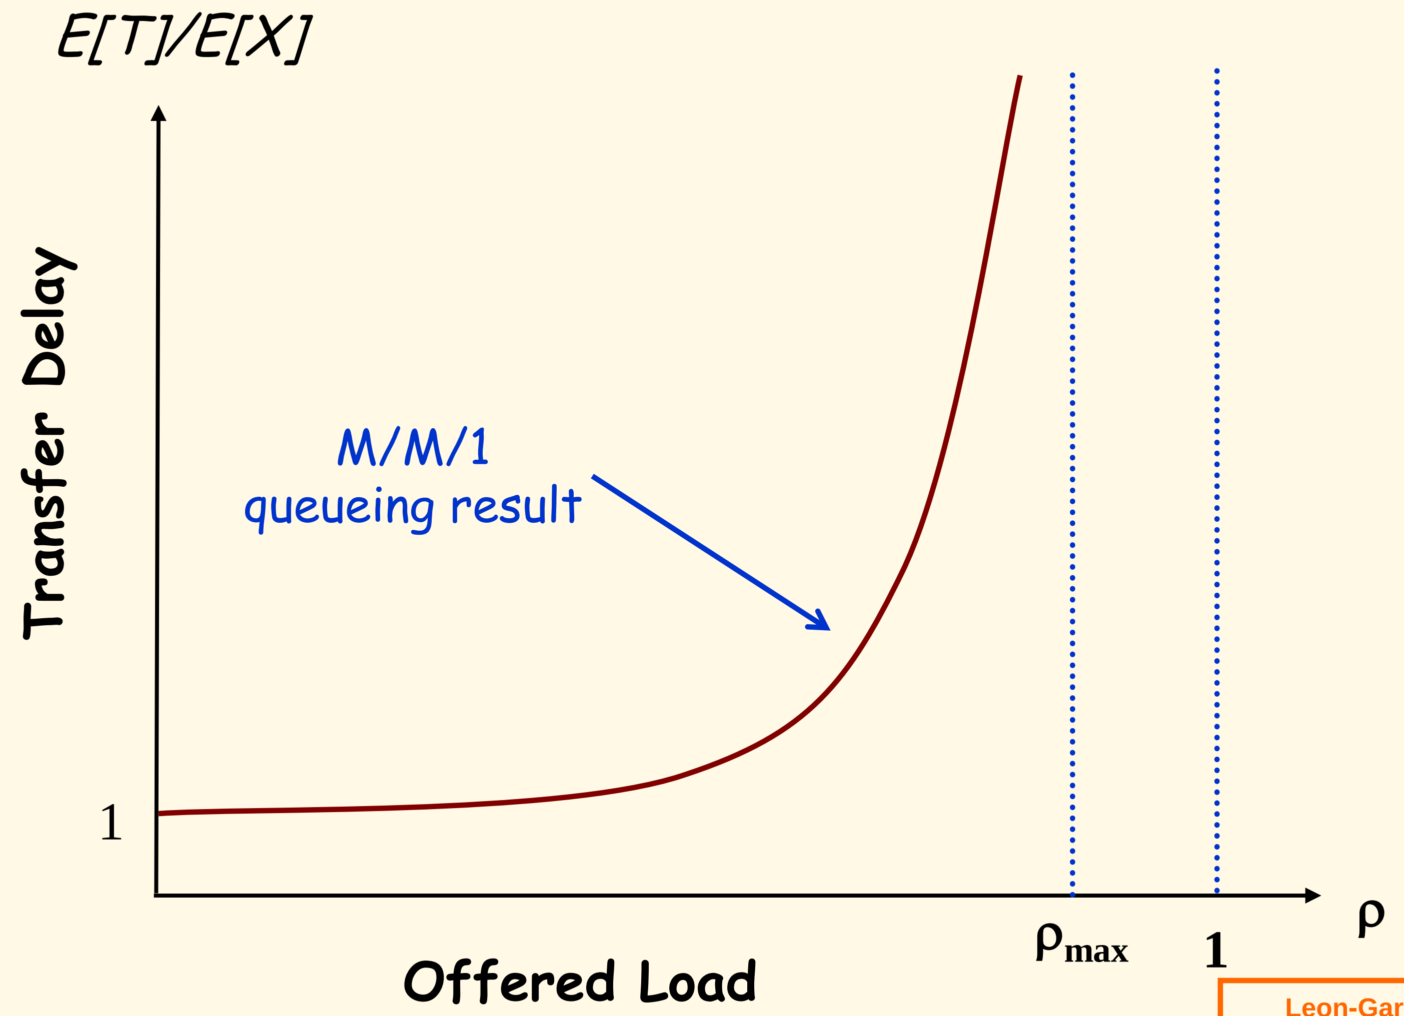

Leon-Garcia & Widjaja: *Communication Networks*

Delay-Throughput Performance Dependence on a

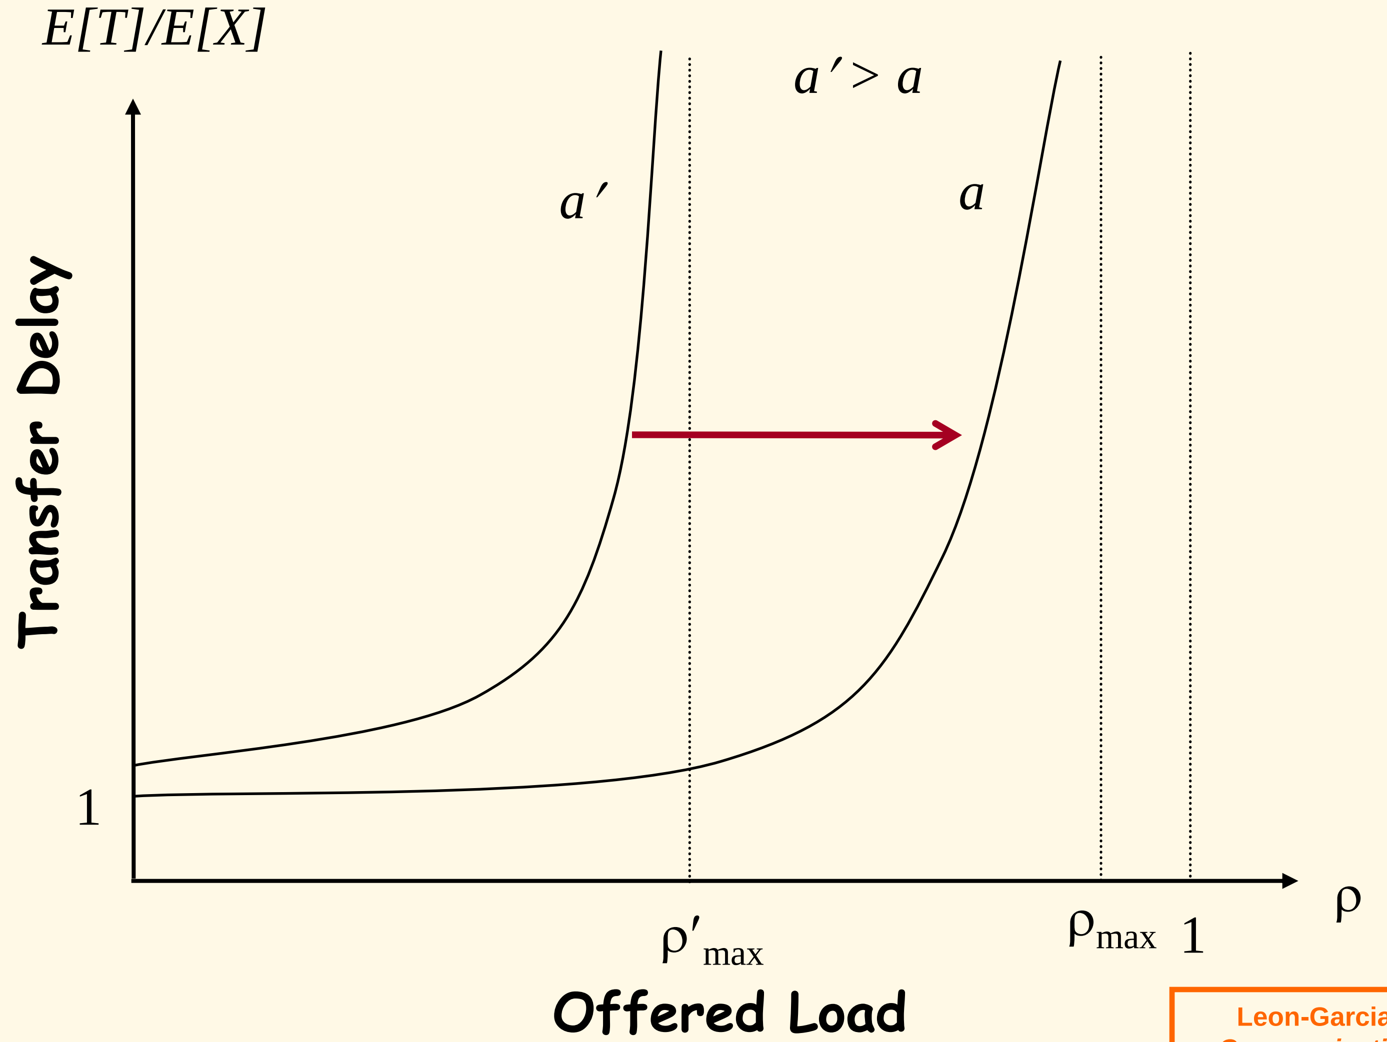

Leon-Garcia & Widjaja: *Communication Networks*

Multiple Access Protocols

Multiple Access Links and Protocols

K & R

Two types of "links":

- point-to-point
  - PPP for dial-up access
  - point-to-point link between Ethernet switch and host
- broadcast (shared wire or medium)
  - old-fashioned Ethernet
  - upstream HFC
  - 802.11 wireless LAN

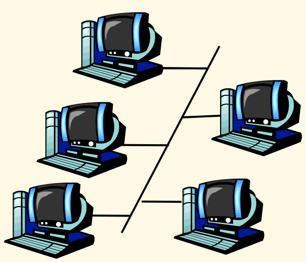

shared wire (e.g., cabled Ethernet)

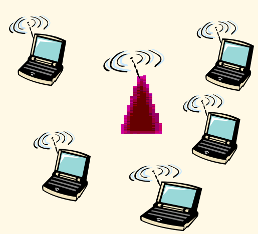

shared RF (e.g., 802.11 WiFi)

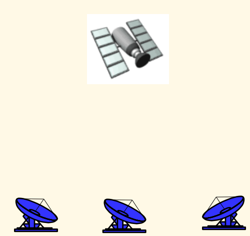

shared RF (satellite)

humans at a cocktail party (shared air, acoustical)

Multiple Access Protocols

- single shared broadcast channel
- two or more simultaneous transmissions by nodes → interference, namely
  - a collision if any node receives two or more signals at the same time.

Multiple Access Protocol (MA)

- distributed algorithm that determines how nodes share channel, i.e., determine when a node can transmit.
- communication about channel sharing must use channel itself!
  - Assumes no out-of-band channel for coordination.

MAC Protocols Taxonomy

Three broad classes:

- Channel Partitioning
  - divide channel into smaller "pieces" (time slots, frequency, code).
  - allocate piece to node for exclusive use.
- Random Access
  - channel not divided, allow collisions.
  - "recover" from collisions.
- "Taking Turns"
  - nodes take turns, but nodes with more to send can take longer turns.

K & R

Channel Partitioning MAC Protocols: TDMA

TDMA: Time Division Multiple Access

- access to channel in "rounds".
- each station gets fixed length slot (length = pkt transmission time) in each round.
- unused slots go idle.
- example: 6-station LAN, 1,3,4 have pkt, slots 2,5,6 wasted (idle).

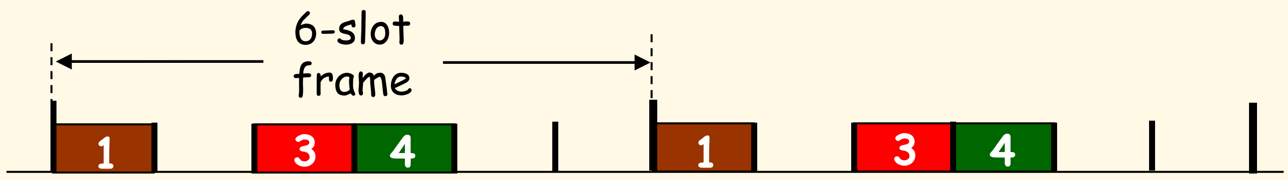

K & R

Channel Partitioning MAC Protocols: FDMA

FDMA: Frequency Division Multiple Access

- channel spectrum divided into frequency bands.
- each station assigned fixed frequency band.
- unused transmission time in frequency bands go idle.
- example: 6-station LAN, 1,3,4 have pkt, frequency bands 2,5,6 idle.

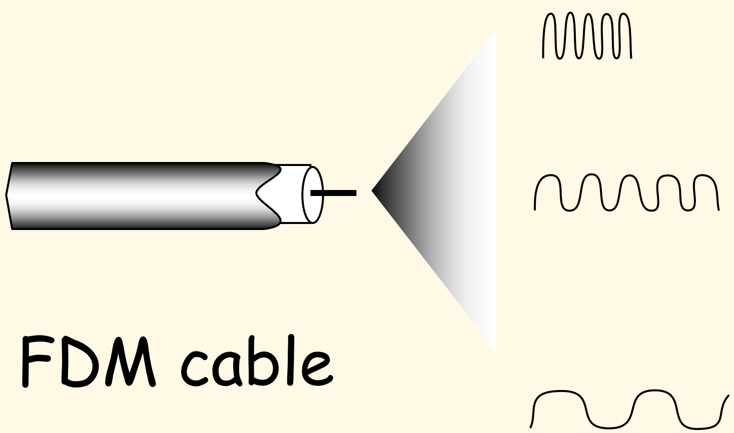

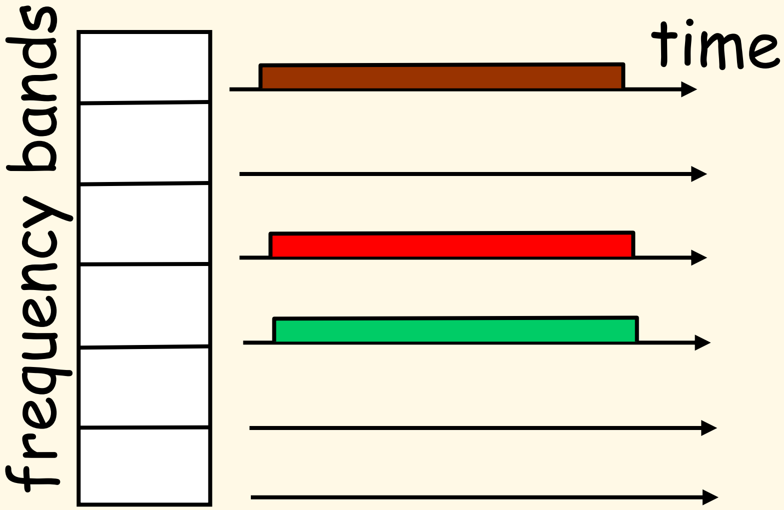

K & R

Random Access Protocols

- When node has packet to send
  - transmit at full channel data rate R.
  - no *a priori* coordination among nodes.
- two or more transmitting nodes ➜ "collision",
- random access MAC protocol specifies:
  - how to detect collisions?
  - how to recover from collisions (e.g., via delayed retransmissions)?
- Examples of random access MAC protocols:
  - ALOHA
  - slotted ALOHA
  - CSMA, CSMA/CD, CSMA/CA

K & R

Historic LAN Performance Notation

I :: input load - the total (normalized) rate of data generated by all n stations.

G :: offered load - the total (normalized) data rate presented to the network including retransmissions.

S :: LAN throughput - the total (normalized) data rate transferred between stations.

D :: average frame delay - the time from when a frame is ready for transmission until completion of a successful transmission.

Normalizing Throughput (S)

[*assuming one packet = one frame*]

S is normalized using packets/packet time where

packet time :: the time to transmit a *standard* fixed-length packet.

i.e.,

$$\text{packet time} = \frac{\text{packet length}}{\text{bit rate}}$$

NOTE: Since the channel capacity is one packet /packet time, S* can be viewed as throughput as a fraction of capacity.

*Represented in LG&W by ρ in later graphs.

Historic LAN Performance Notation

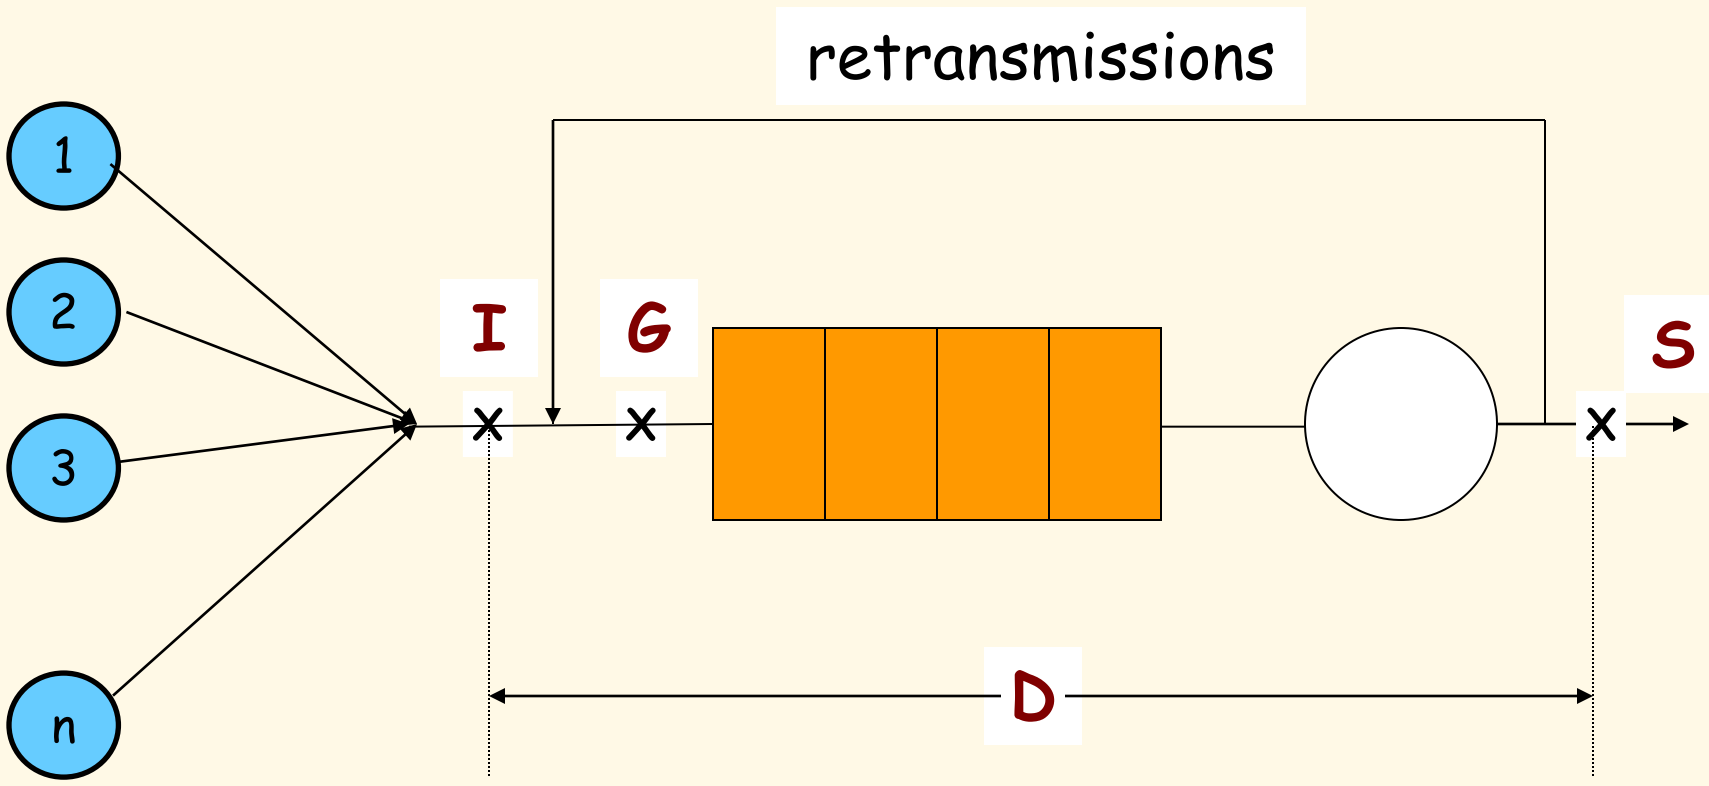

ALOHA

- **Abramson solved the channel allocation problem for ground radio at University of Hawaii in 1970's.**

Aloha Transmission Strategy

Stations transmit whenever they have data to send.

- Collisions will occur and colliding frames are destroyed.

Aloha Retransmission Strategy

Station waits a **random amount of time** before sending again.

ALOHA

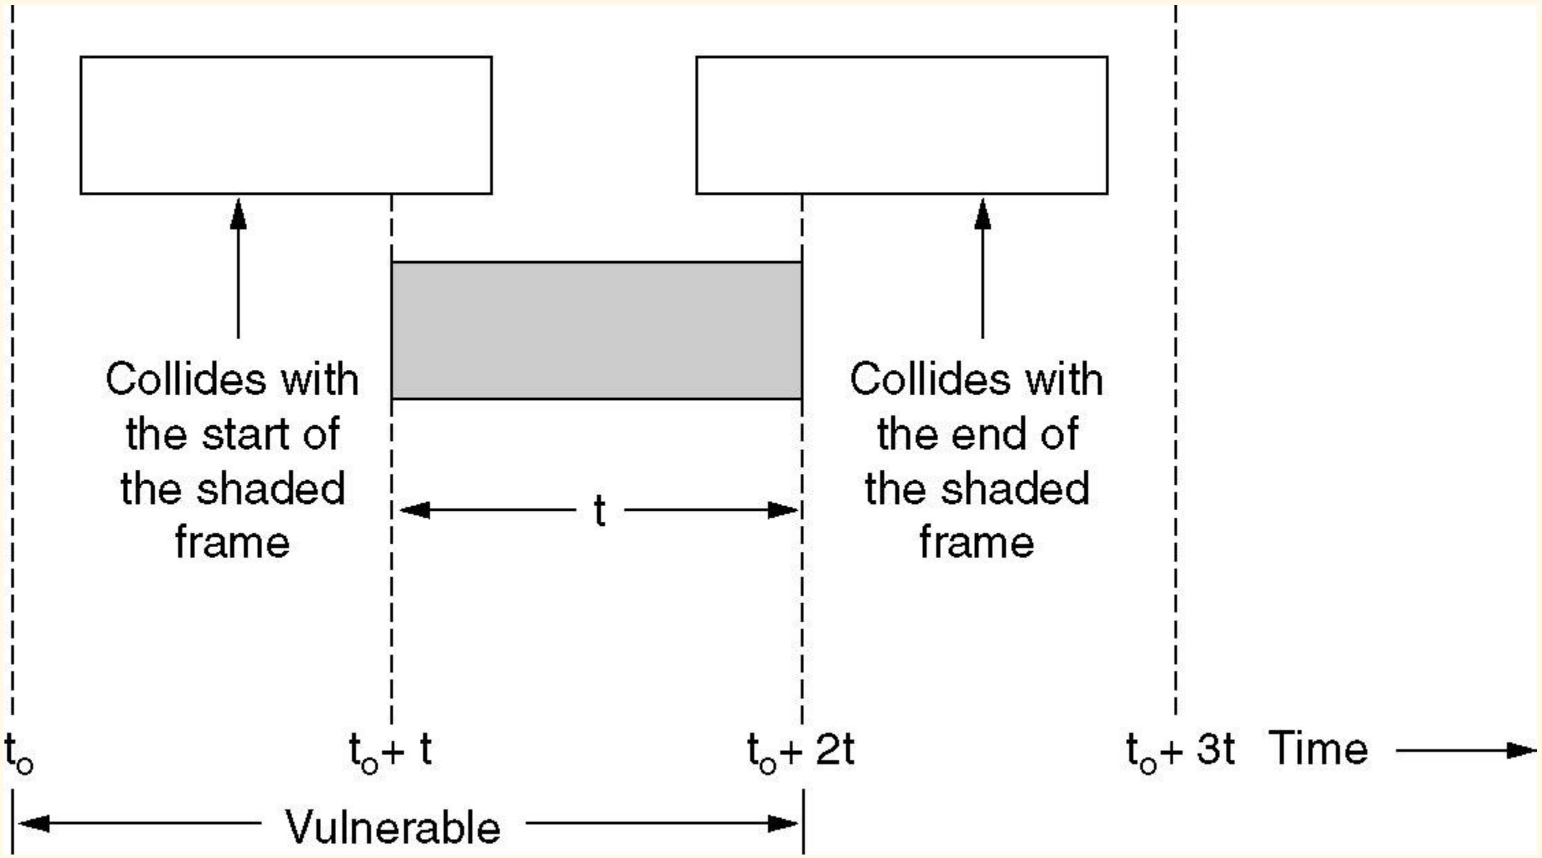

Figure 4-2. Vulnerable period for the shaded frame.

Tanenbaum

ALOHA Retransmissions

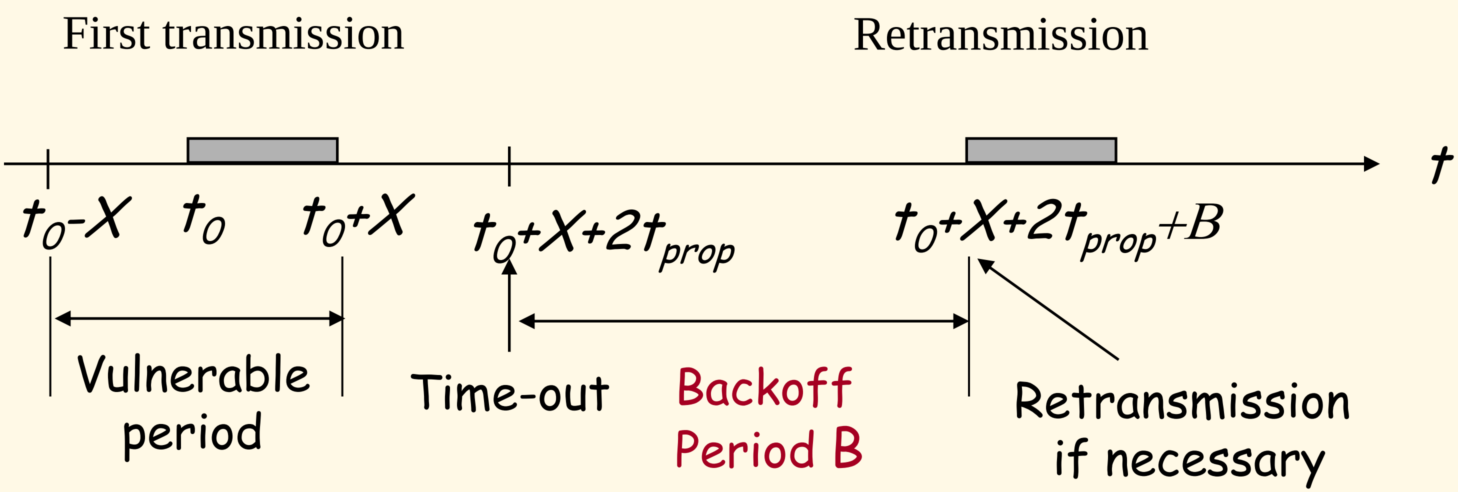

random backoff period B

Leon-Garcia & Widjaja: *Communication Networks*

ALOHA

Vulnerable period :: $t_0 - X$ to $t_0 + X$ two frame transmission times
Assume: Poisson Arrivals with average number of arrivals of
2G arrivals/ 2 X

$$S = G e^{-2(1+a)G}$$

Slotted ALOHA (Roberts 1972)

- uses discrete time intervals as slots (i.e., slot = one packet transmission time) and synchronize the send time (e.g., use "pip" from a satellite).

Slotted Aloha Strategy

Station transmits ONLY at the beginning of a time slot.

- Collisions will occur and colliding frames are destroyed.

Slotted Aloha Retransmission Strategy

Station waits a random amount of time before sending again.

Slotted ALOHA

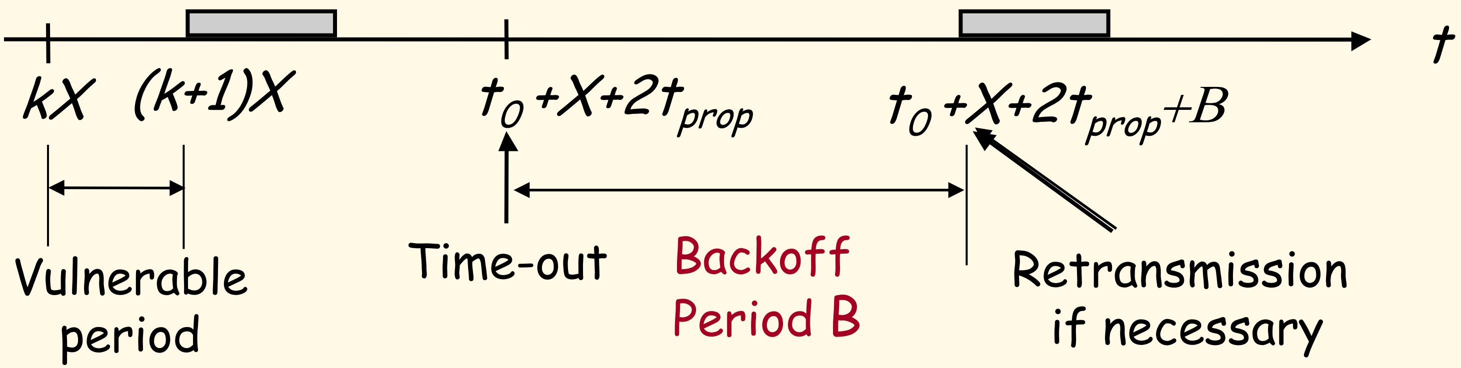

random backoff period B slots

Leon-Garcia & Widjaja: *Communication Networks*

Slotted ALOHA

Vulnerable period :: $t_0 - X$ to t_0 **one frame transmission time**
Assume: Poisson Arrivals with average number of arrivals of G arrivals/ X

$$P_0 = P[k=0, t=1] = e^{-G}$$
$$S = G P_0$$
$$S = G e^{-G}$$

and an adjustment for **a** yields:

$$S = G e^{-(1+a) G}$$

ALOHA and Slotted AlOHA Throughput versus Load

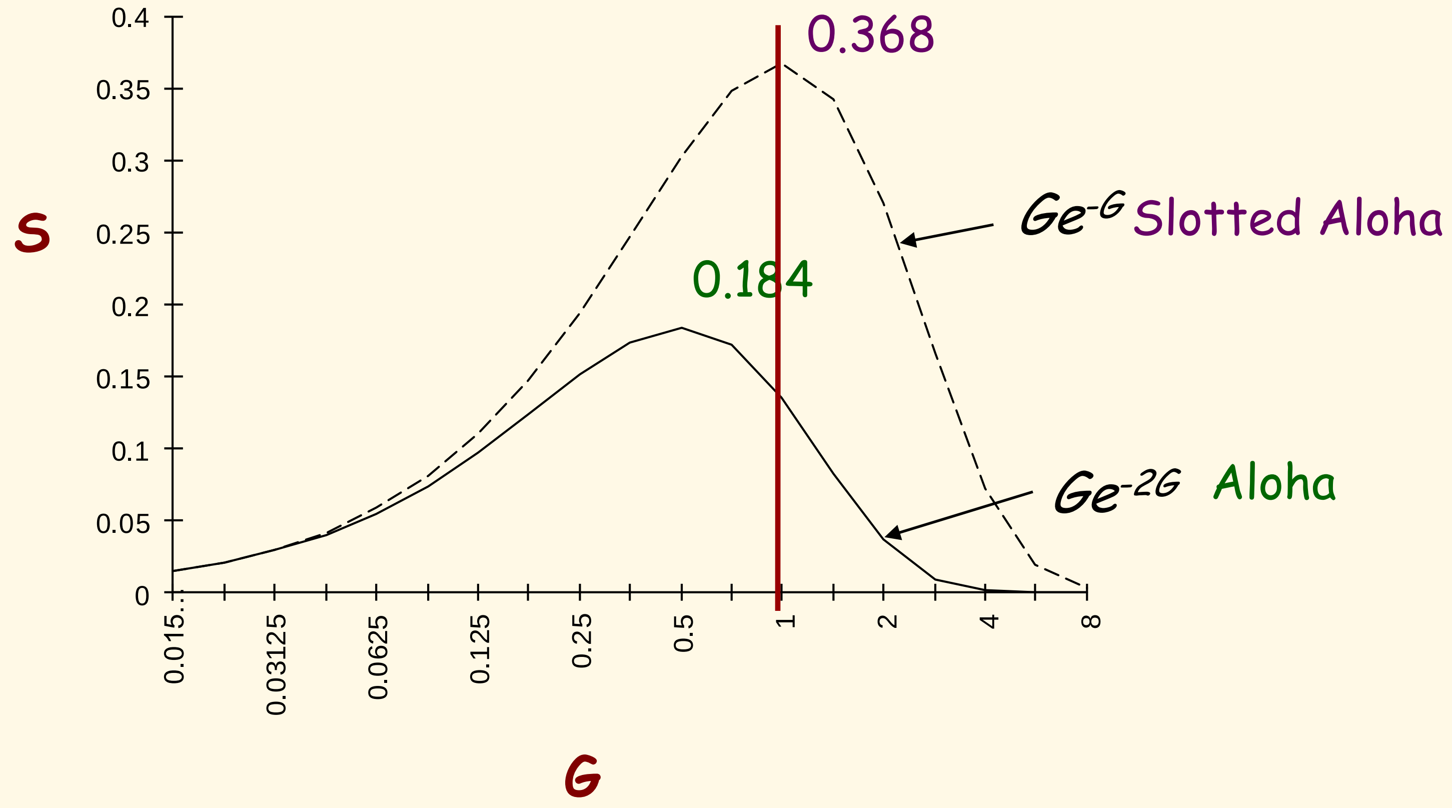

Leon-Garcia & Widjaja: *Communication Networks*

Carrier Sense with Multiple Access (CSMA)

1-persistent CSMA Transmission Strategy

‘the greedy algorithm’

1. Sense the channel.
2. IF the channel is *idle*, THEN transmit.
3. IF the channel is *busy*, THEN continue to listen until channel is *idle* and transmit immediately.

nonpersistent CSMA Transmission Strategy

'the less-greedy algorithm'

1. Sense the channel.
2. IF the channel is *idle*, THEN transmit.
3. IF the channel is *busy*, THEN wait a **random amount of time** and repeat the algorithm.

p - persistent CSMA Transmission Strategy

'a slotted approximation'

1. Sense the channel.

2. IF the channel is *idle*, THEN with probability p transmit and with probability (1-p) delay *one time slot* and repeat the algorithm.

3. IF the channel is *busy*, THEN delay *one time slot* and repeat the algorithm.

P – Persistent CSMA details

- the time slot is usually set to the **maximum propagation delay**.
- as p decreases, stations wait longer to transmit but the number of collisions decreases.
- Consideration for the choice of p
  - ($n \times p$) must be < 1 for stability, where n is maximum number of stations, i.e.,

$$p < 1/n$$

CSMA Collisions

- In all three strategies, a collision is possible.
- CSMA determines collisions by the lack of an ACK which results in a TIMEOUT. {This is extremely expensive with respect to performance!!}
- If a collision occurs, THEN wait a random amount of time and retransmit.

CSMA Persistence Summary

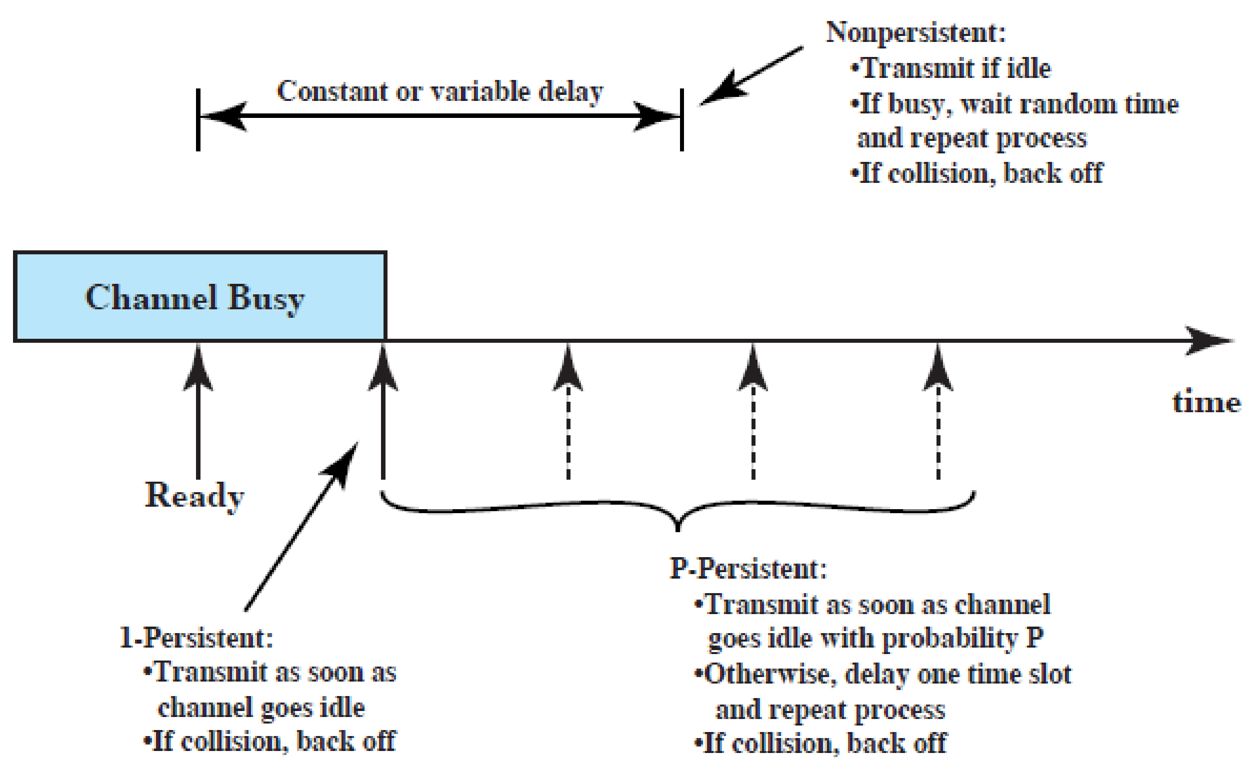

Figure 16.1 CSMA Persistence and Backoff

CSMA Collisions

Collisions can still occur:
propagation delay means two nodes may not hear each other's transmission.

Collision:
entire packet transmission time wasted.

Note:
The role of distance & propagation delay in determining collision probability

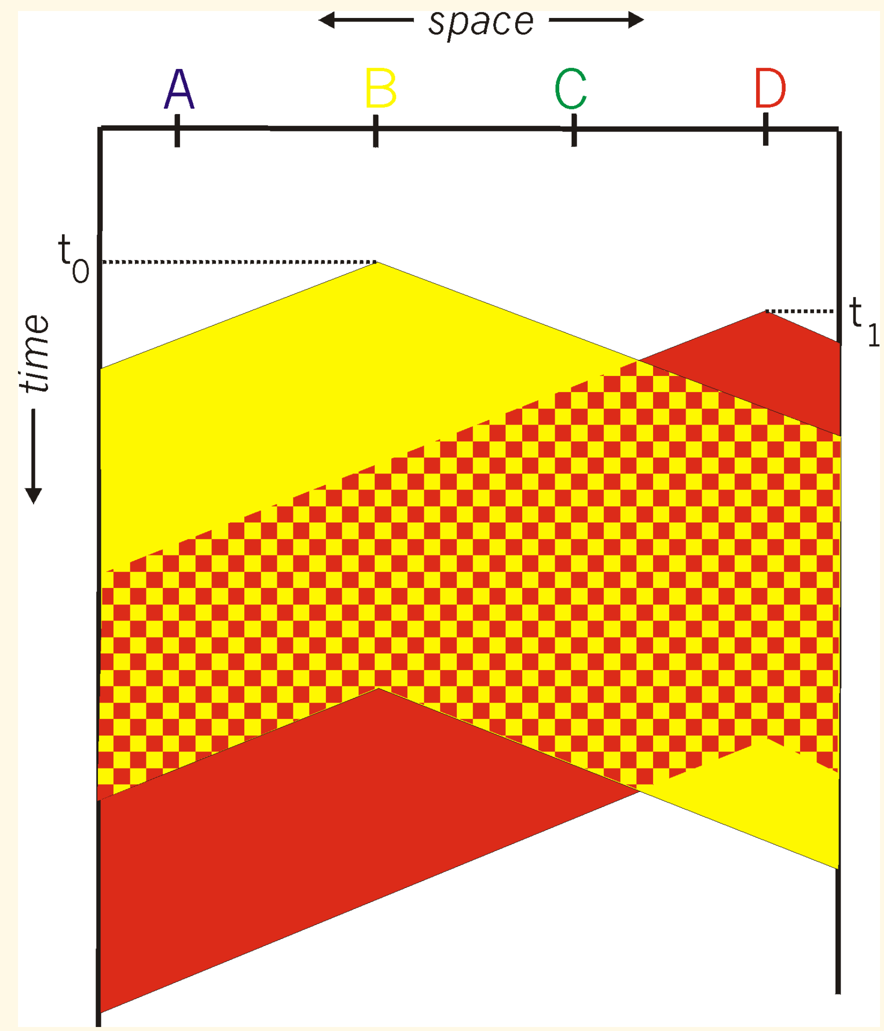

K & R

Persistent and Non-persistent CSMA

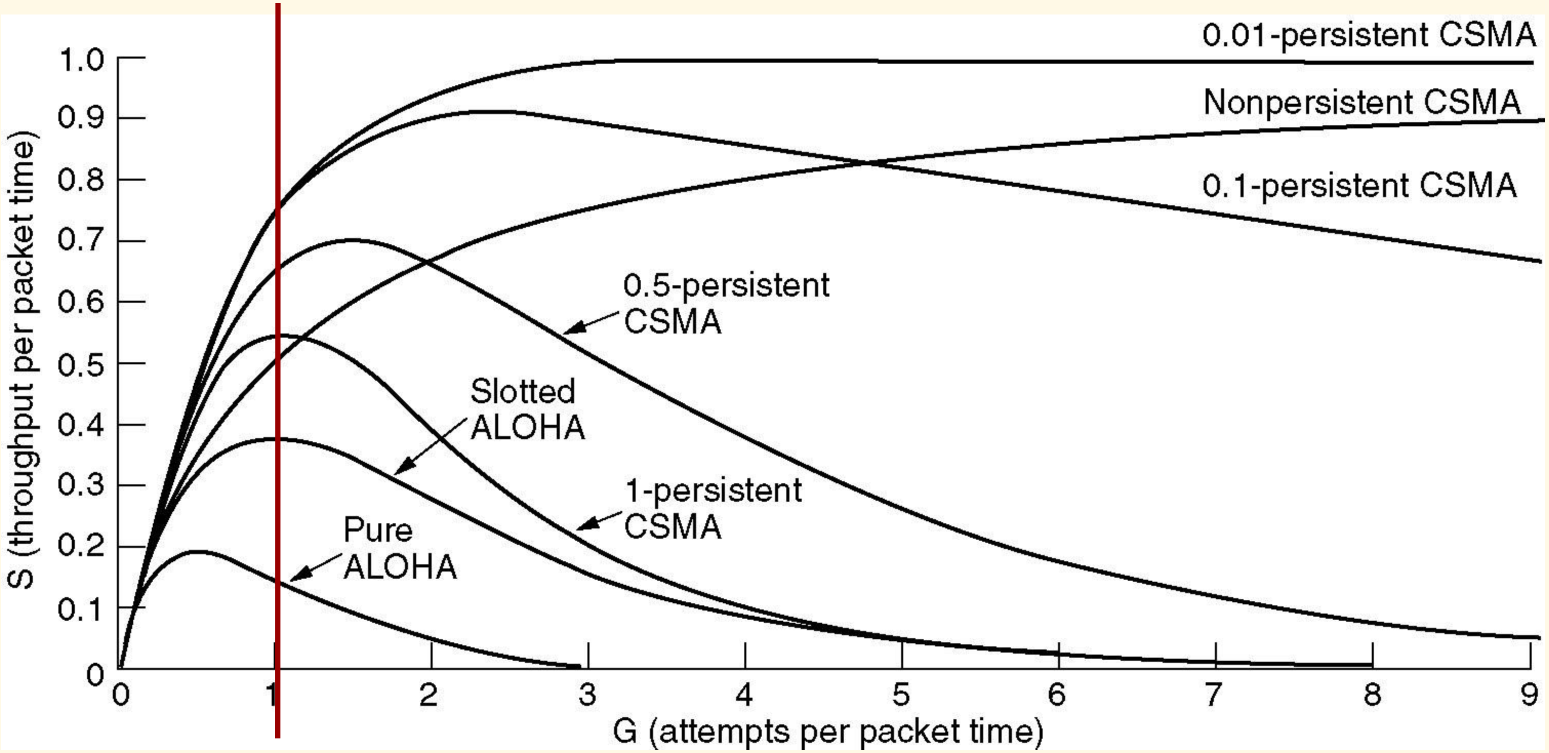

Figure 4-4. Comparison of the channel utilization versus load for various random access protocols.

Tanenbaum

Throughput versus Load with varying a

1-Persistent CSMA

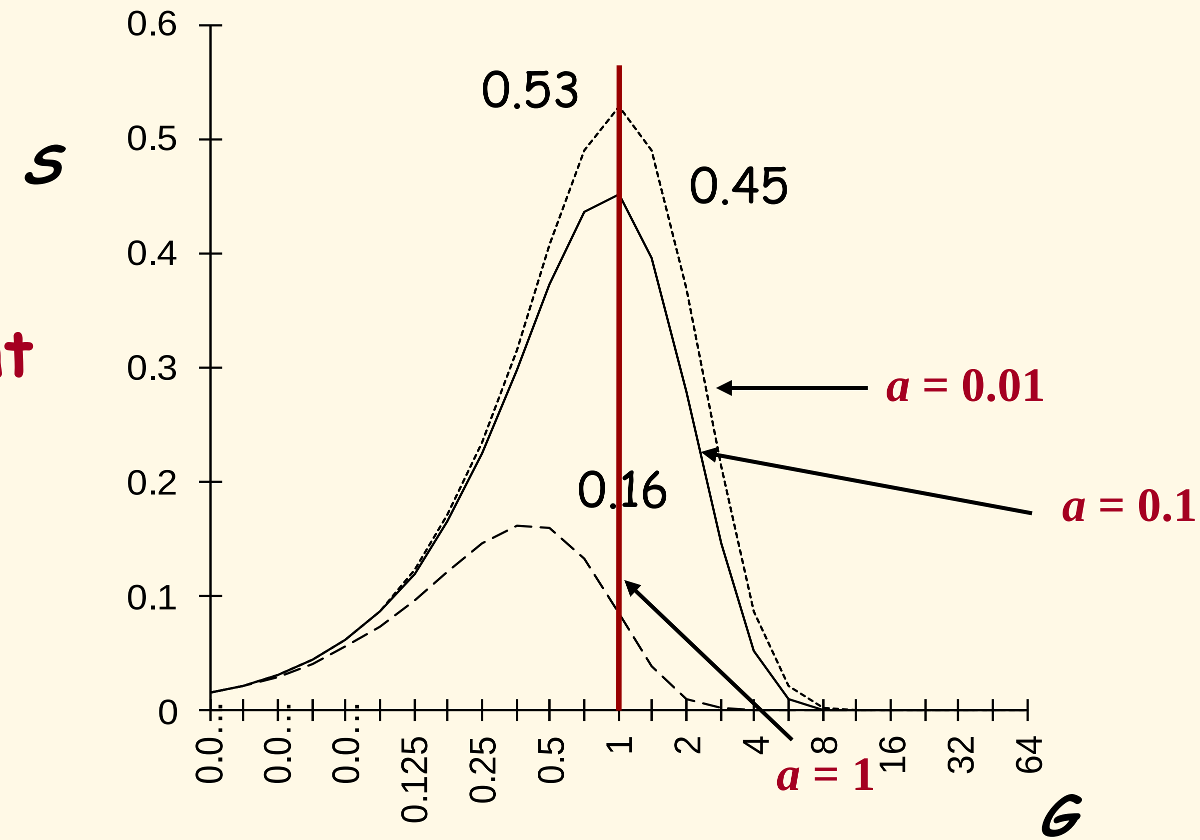

Leon-Garcia & Widjaja: *Communication Networks*

Throughput versus Load with varying a

Non-Persistent CSMA

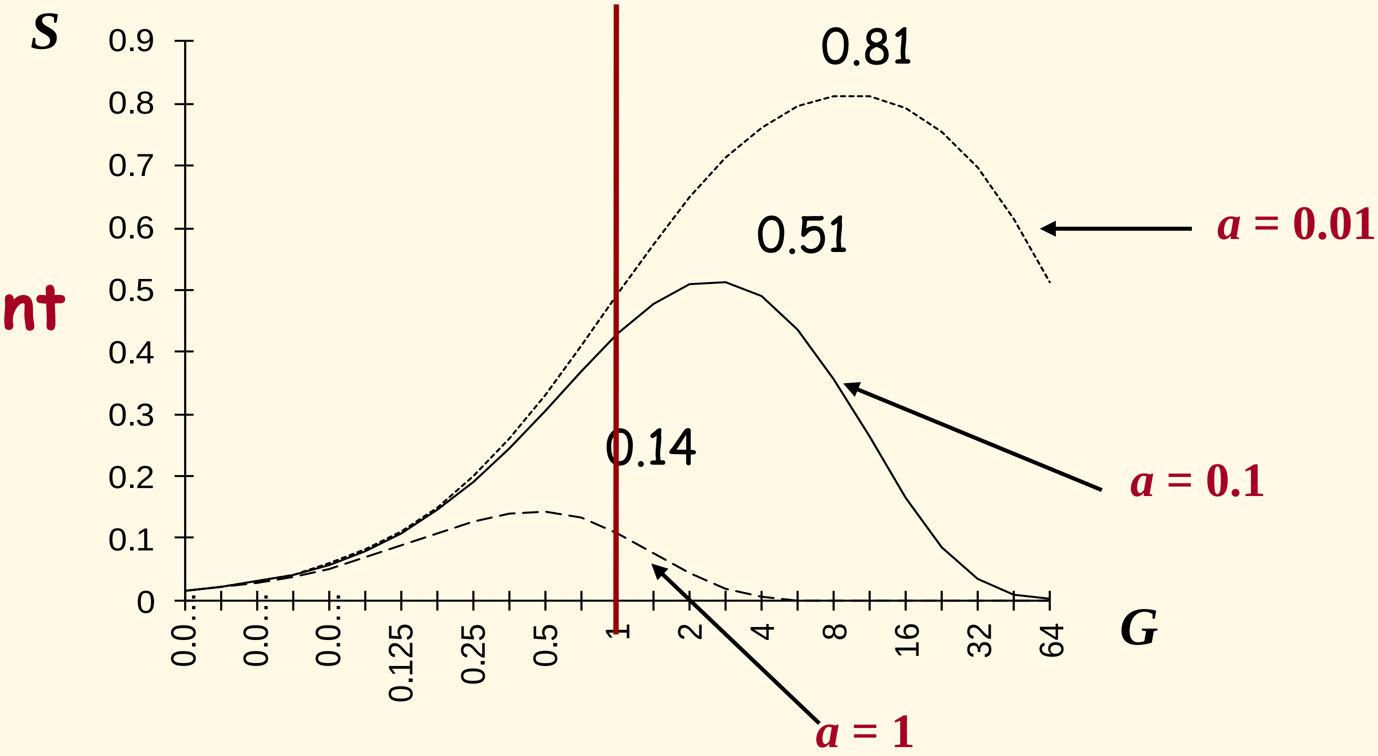

Leon-Garcia & Widjaja: *Communication Networks*

CSMA/CD (Collision Detection)

CSMA/CD:

- collisions detected within short time.
- colliding transmissions aborted, reducing channel wastage.

Collision Detection:

- easy in wired LANs: measure signal strengths, compare transmitted, received signals
- difficult in wireless LANs: received signal strength overwhelmed by local transmission strength.

CSMA/Collision Detection

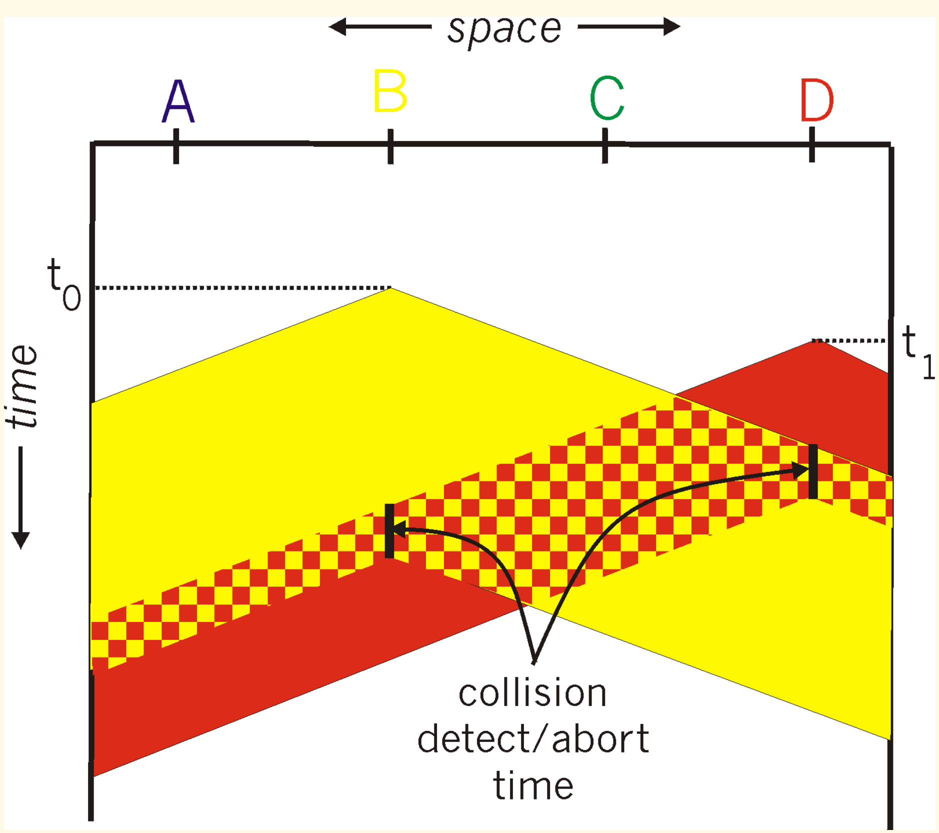

K & R

CSMA/CD
CSMA *with* Collision Detection

- If a collision is detected during transmission, then immediately cease transmitting the frame.
- The first station to detect a collision sends a *jam signal* to all stations to indicate that there has been a collision.
- After receiving a *jam signal*, a station that was attempting to transmit waits a random amount of time before attempting to retransmit.
- The maximum time needed to detect a collision is 2 x propagation delay.

CSMA vs CSMA/CD

- CSMA is essentially a historic technology until we include Wireless LANs.
- If propagation time is short compared to transmission time, station can be *listening before sending* with CSMA.
- Collision detection (CD) is accomplished by detecting voltage levels outside acceptable range. Thus attenuation limits distance without a repeater.
- If the collision time is short compared to packet time (i.e., small *a*), performance will increase due to CD.

CSMA/CD

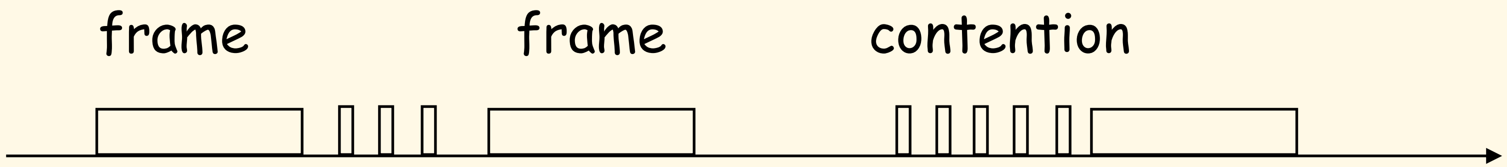

Probability of 1 successful transmission:

$$P_{success} = np(1-p)^{n-1}$$

$P_{success}$ is maximized at $p = 1/n$:

$$P_{success}^{max} = n(1-\frac{1}{n})^{n-1} \to \frac{1}{e}$$

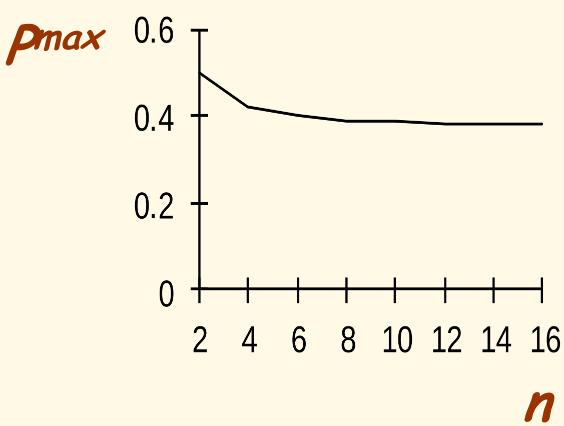

Tanenbaum

Maximum Achievable Throughputs

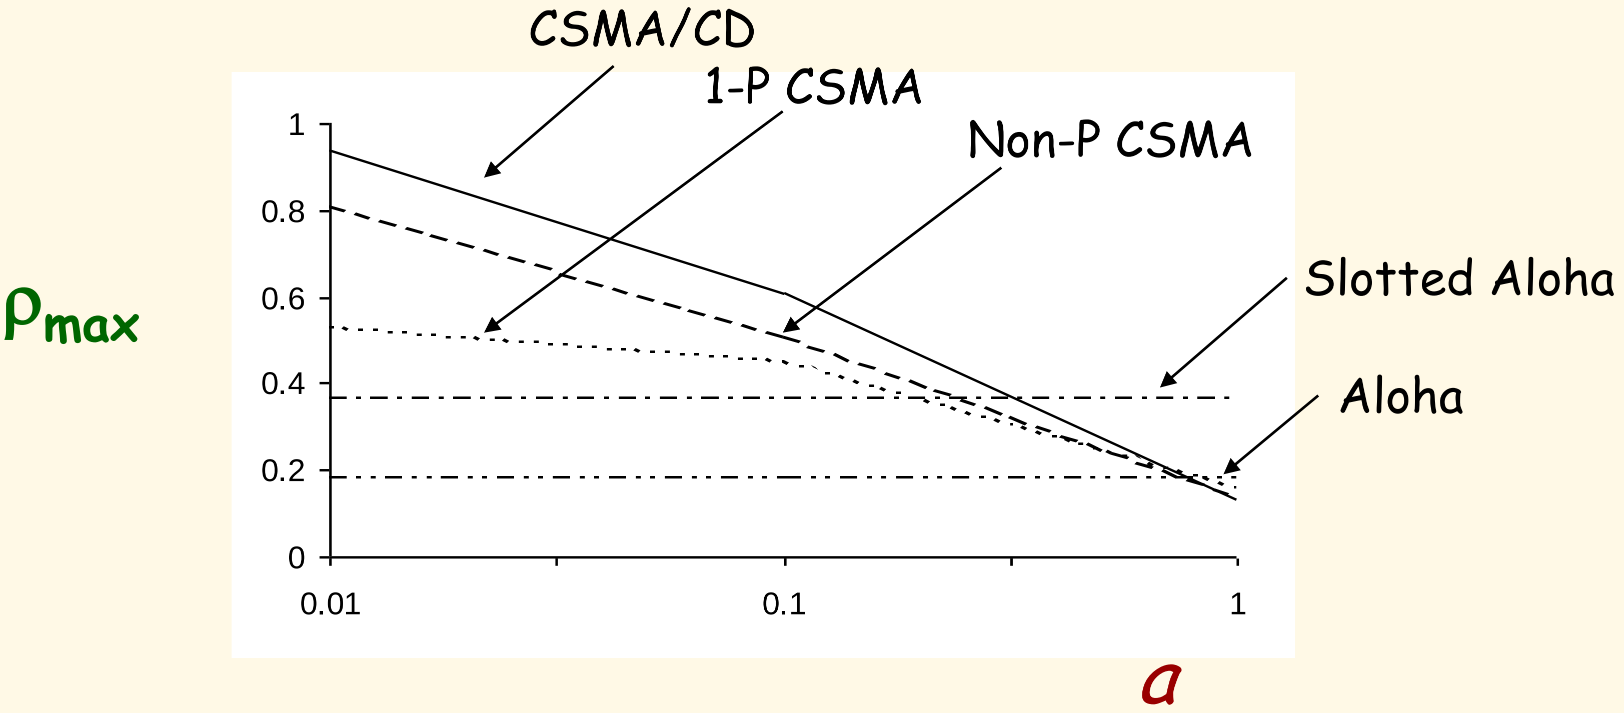

Leon-Garcia & Widjaja: *Communication Networks*

Frame Delay with varying a

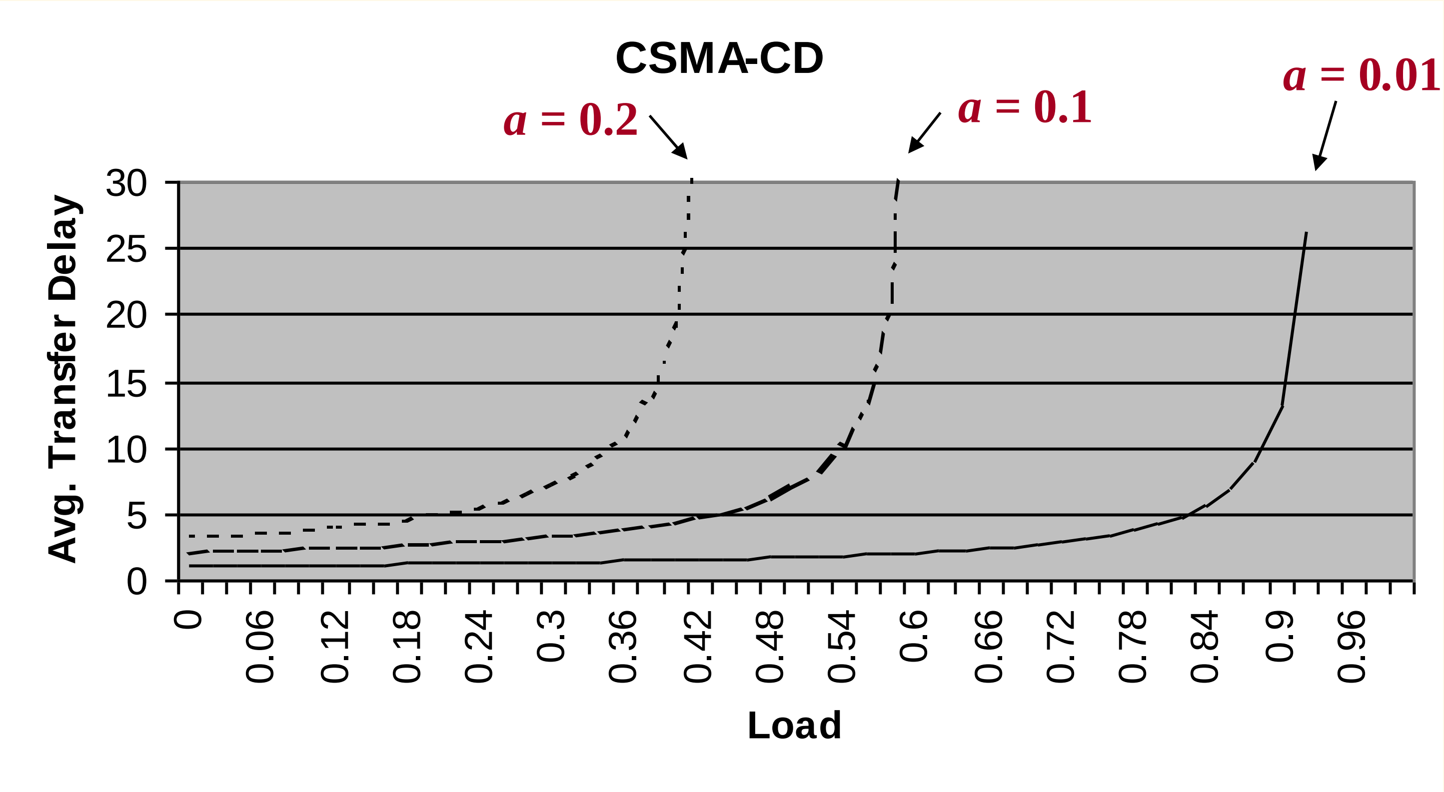

Leon-Garcia & Widjaja:
Communication Networks

"Taking Turns" MAC protocols

Channel Partitioning MAC protocols:

- share channel *efficiently* and *fairly* at high load.
- inefficient at low load: delay in channel access, 1/N bandwidth allocated even if only 1 active node!

Random Access MAC protocols:

- efficient at low load: single node can fully utilize channel.
- high load: collision overhead

"Taking Turns" protocols:

look for best of both worlds!

K & R

"Taking Turns" MAC protocols

Polling:

- master node "invites" slave nodes to transmit in turn
- typically used with "dumb" slave devices
- concerns:
  - polling overhead
  - latency
  - single point of failure (master)

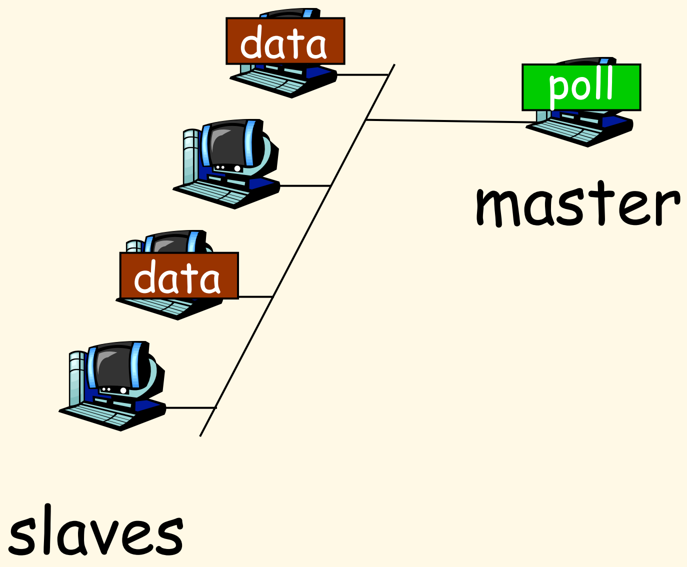

K & R

"Taking Turns" MAC protocols

Token passing:

- control **token** passed from one node to next sequentially.
- token message
- concerns:
  - token overhead
  - latency
  - single point of failure (token)

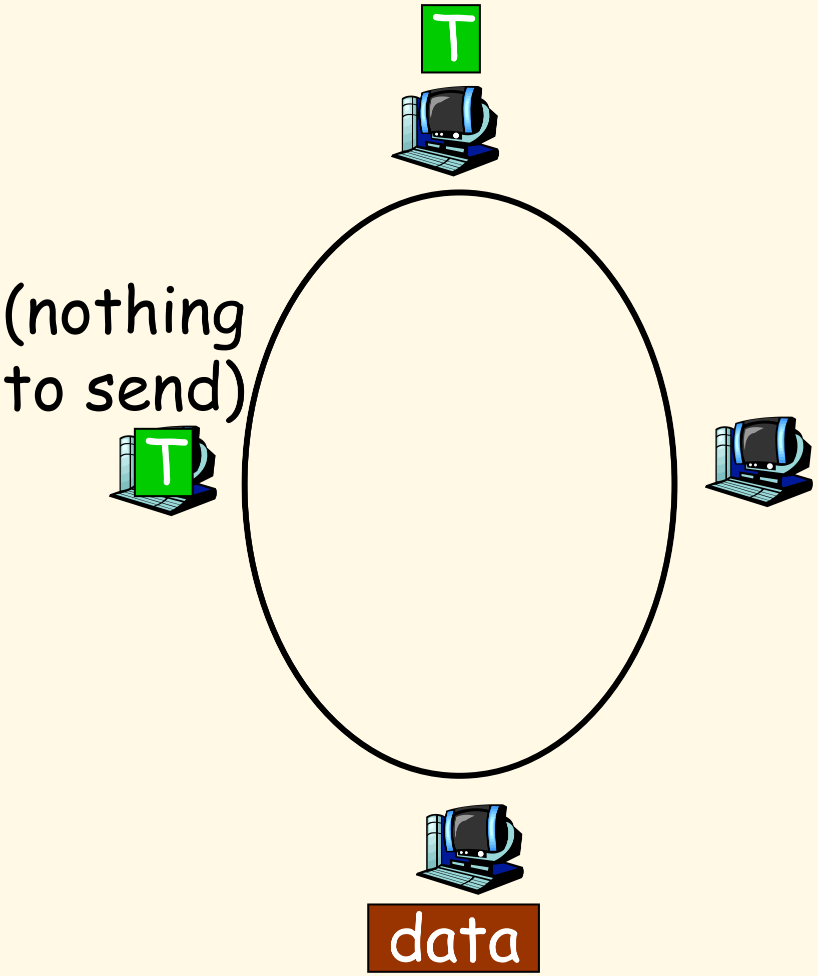

K & R

LANS Summary

- Channel Allocation Problem
- Relative Propagation Time
- LAN Utilization Upper Bound
- Multiple Access Protocols
  - TDMA, FDMA
  - Aloha, Slotted Aloha
  - CSMA (non-persistent, 1-persistent, p-persistent), CSMA/CD
  - Token Passing
  - Performance Results