CS4700/CS5700
Fundamentals of Computer Networks

Lecture 7: Ethernet / Wi-Fi media access control

Slides used with permissions from Edward W. Knightly, T. S. Eugene Ng, Ion Stoica, Hui Zhang

Overview

- Ethernet and Wi-Fi are both "multi-access" technologies
  - Broadcast medium, shared by many hosts
  - Simultaneous transmissions will result in collisions
- Media Access Control (MAC) protocol required
  - Rules on how to share medium

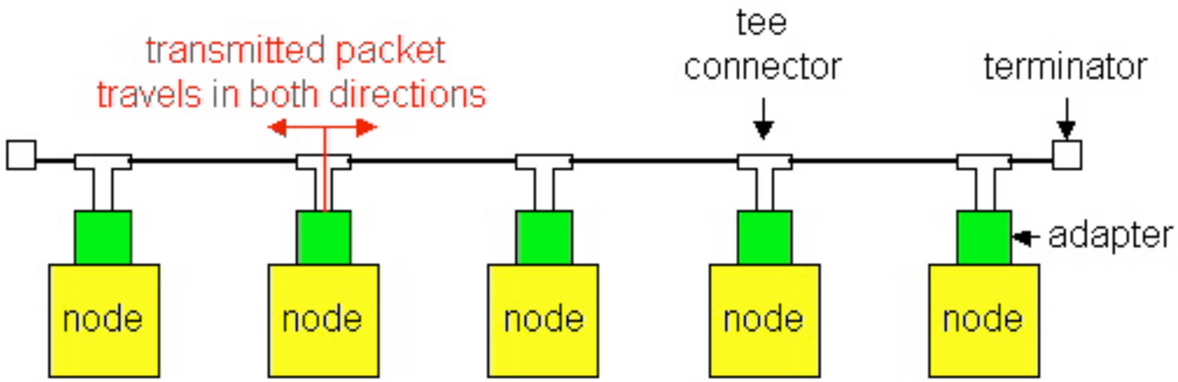

Media Access Control Protocols

- Channel partitioning
  - Divide channel into smaller "pieces" (e.g., time slots, frequency)
  - Allocate a piece to node for exclusive use
  - E.g. Time-Division-Multi-Access (TDMA) cellular network
- Taking-turns
  - Tightly coordinate shared access to avoid collisions
  - E.g. Token ring network
- Contention
  - Allow collisions
  - "recover" from collisions
  - E.g. Ethernet, Wi-Fi

Contention Media Access Control Goals

- Share medium
  - If two users send at the same time, *collision* results in no packet being received (interference)
  - If no users send, channel goes idle
  - Thus, want to have only one user send at a time
- Want high network utilization
  - TDMA doesn't give high utilization
- Want simple distributed algorithm
  - no fancy token-passing schemes that avoid collisions

Evolution of Contention Protocols

Evolution of Contention Protocols

Aloha

Evolution of Contention Protocols

Aloha

Developed in the 1970s for a packet radio network

Evolution of Contention Protocols

Aloha

Developed in the 1970s for a packet radio network

Slotted Aloha

Evolution of Contention Protocols

Aloha

Developed in the 1970s for a packet radio network

Slotted Aloha

Improvement: Start transmission only at fixed times (slots)

Evolution of Contention Protocols

Aloha

Developed in the 1970s for a packet radio network

Slotted Aloha

Improvement: Start transmission only at fixed times (slots)

CSMA

Evolution of Contention Protocols

Aloha — Developed in the 1970s for a packet radio network

Slotted Aloha — Improvement: Start transmission only at fixed times (slots)

CSMA — CSMA = Carrier Sense Multiple Access
Improvement: Start transmission only if no transmission is ongoing

Evolution of Contention Protocols

Aloha

Developed in the 1970s for a packet radio network

Slotted Aloha

Improvement: Start transmission only at fixed times (slots)

CSMA

CSMA = Carrier Sense Multiple Access
Improvement: Start transmission only if no transmission is ongoing

CSMA/CD

Evolution of Contention Protocols

Aloha — Developed in the 1970s for a packet radio network

Slotted Aloha — Improvement: Start transmission only at fixed times (slots)

CSMA — CSMA = Carrier Sense Multiple Access
Improvement: Start transmission only if no transmission is ongoing

CSMA/CD — CD = Collision Detection
Improvement: Stop ongoing transmission if a collision is detected (e.g. Ethernet)

(Pure) ALOHA

- **Topology:** Broadcast medium with multiple stations
- **Aloha Protocol:**
  - Whenever a station has data, it transmits immediately
  - Receivers ACK all packets
  - No ACK = collision. Wait a random time and retransmit

Simple, but Radical

- Previous attempts all partitioned channel
  - TDMA, FDMA, etc.

- Aloha optimized the common case (few senders) and dealt with collisions through retries

Trade-off Compared to TDMA

- In TDMA, you always have to wait your turn
  - delay proportional to number of sites
- In Aloha, can send immediately
- Aloha gives much lower delays, at the price of lower utilization (as we will see)

Collisions in (Pure) ALOHA

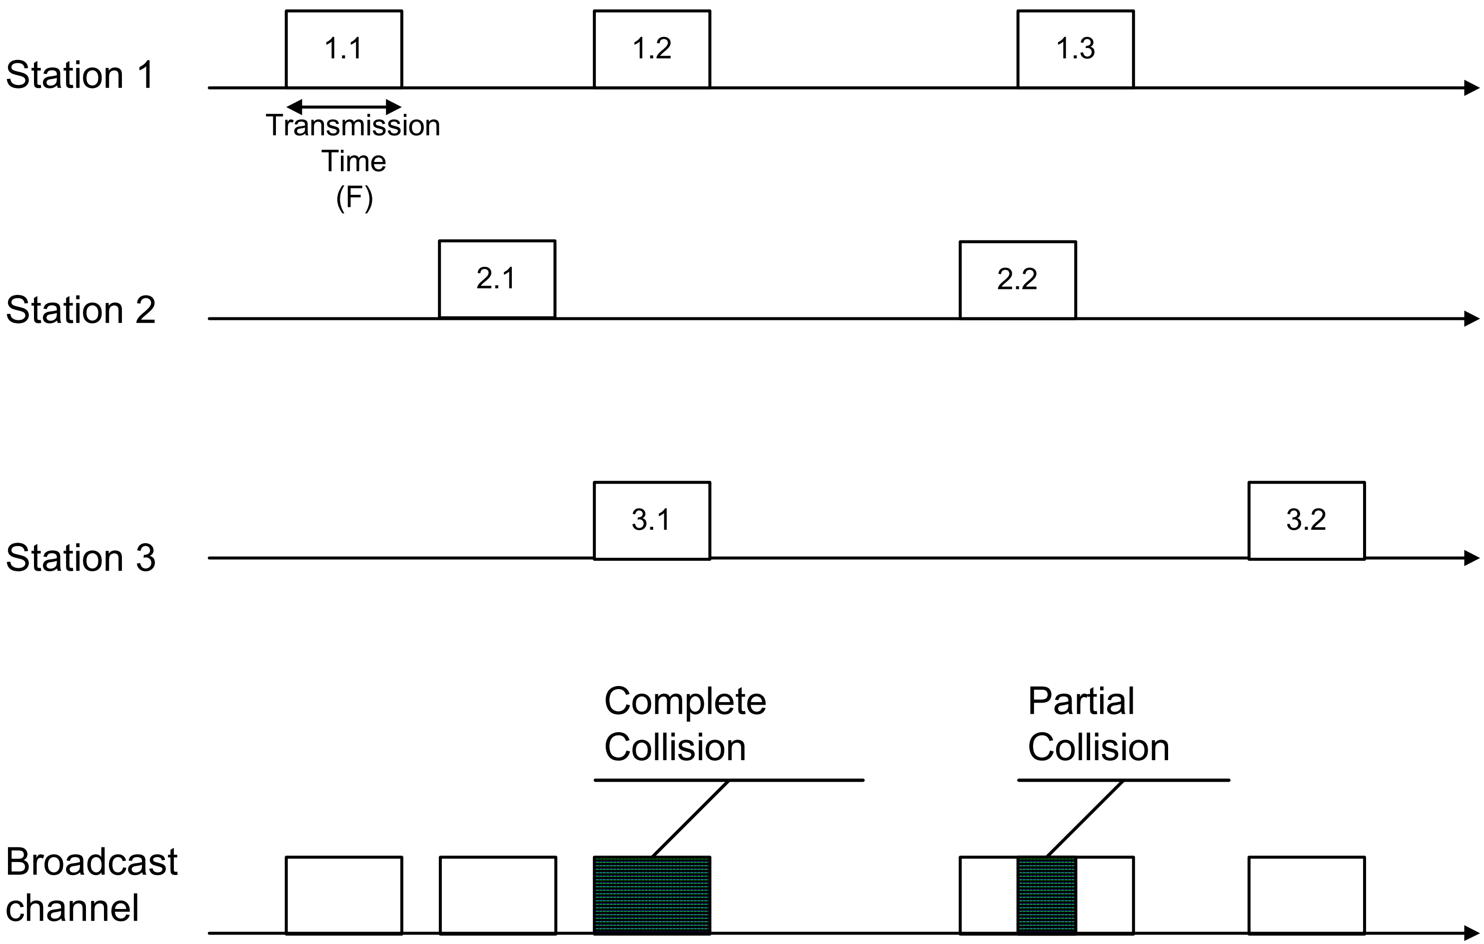

Performance of (Pure) ALOHA

- **Performance questions:**
  - What is the collision probability?
  - What is the maximum throughput?
- **Notation:**

  C: link capacity in (bits/sec)

  s: packet size (bits)

  F: packet transmission time (sec)

$$F = \frac{s}{C}$$

Collisions and Vulnerable Period

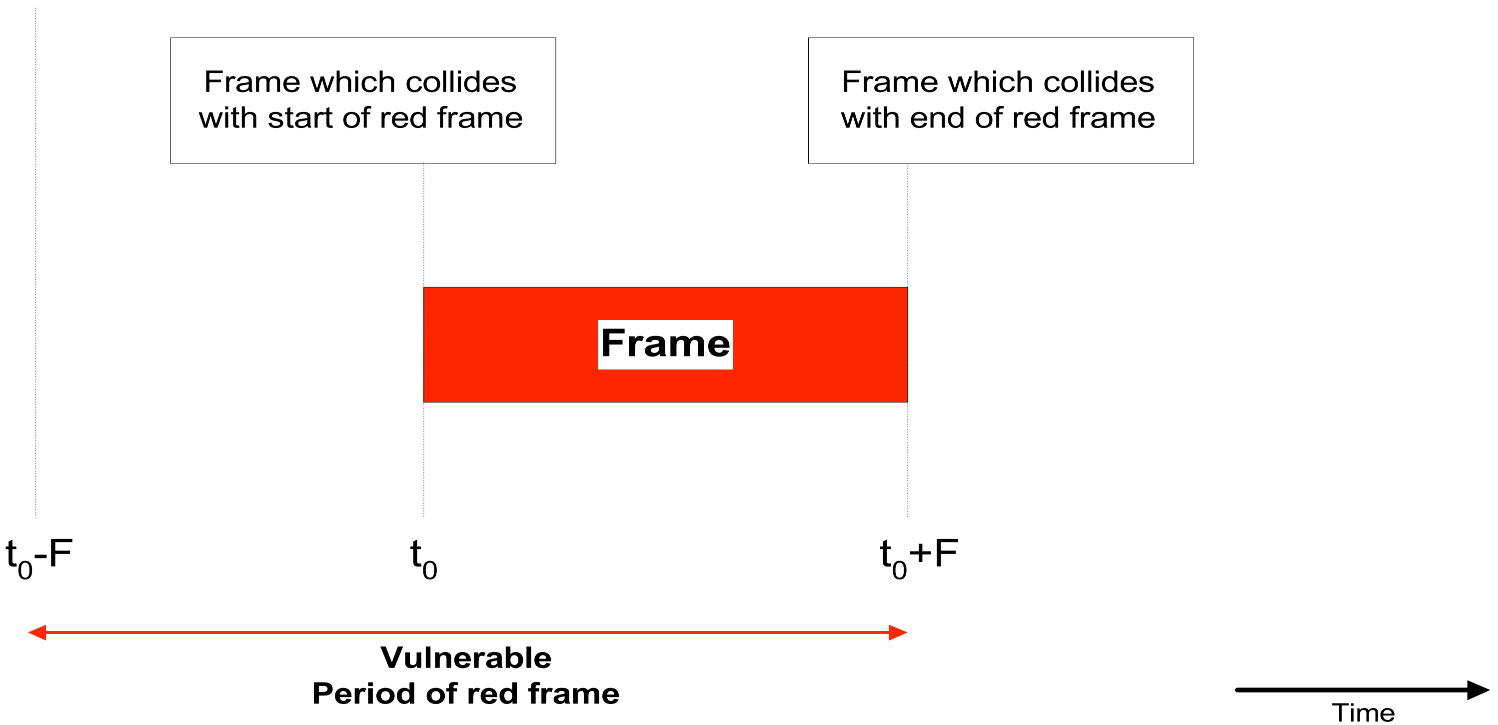

- A frame (red frame) will be in a collision if and only if another transmission begins in the vulnerable period of the frame
- Vulnerable period has the length of 2 frame times

Traffic Model

$\lambda \equiv$ Poisson rate from all stations (Poisson Process's are additive)

$$P[no\_event\_in\_t] = e^{-\lambda t}$$

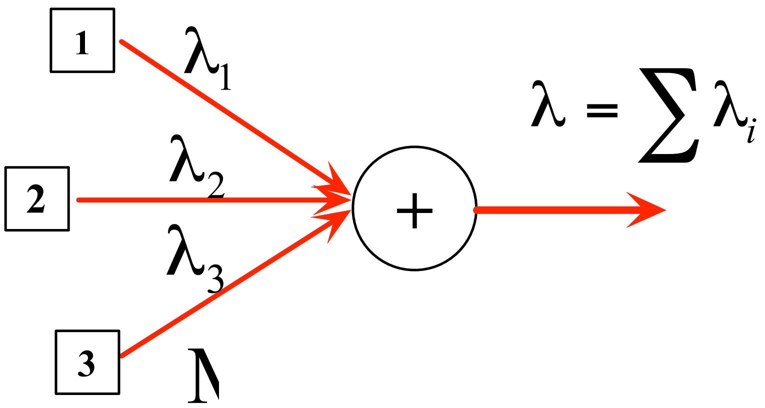

$$\rho \equiv \text{offered load} = \frac{\text{total arrival rate}}{\text{total service rate}} = \frac{\lambda}{C/s} = F\lambda$$

System Model

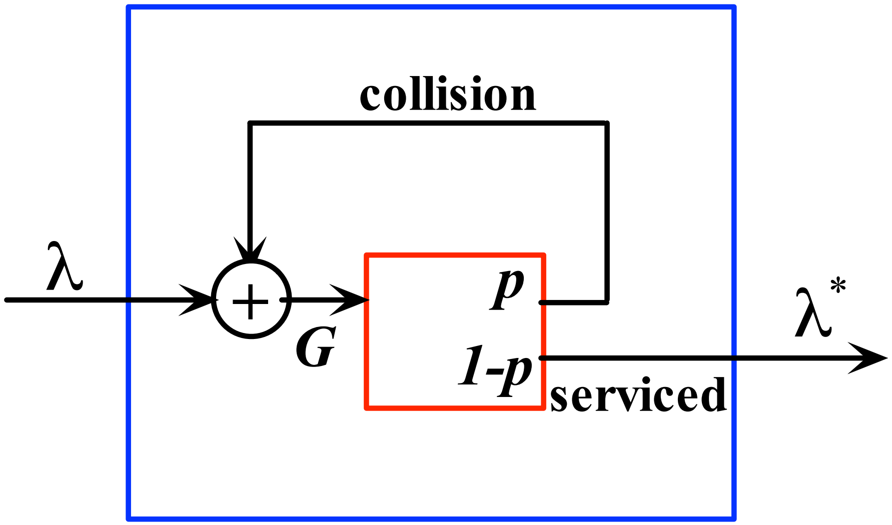

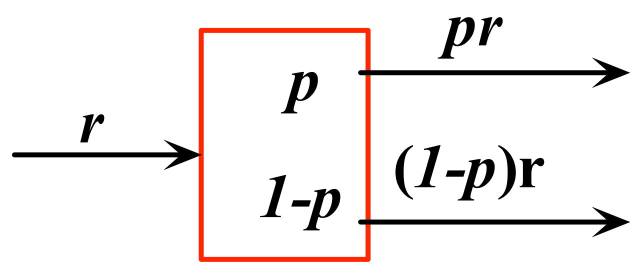

Fact: Bernoulli sampling of a PP is also a PP

$p \equiv$ Probability of collision

$G \equiv$ Total carried load (incl. Retransmissions)

$\Rightarrow$

System Model

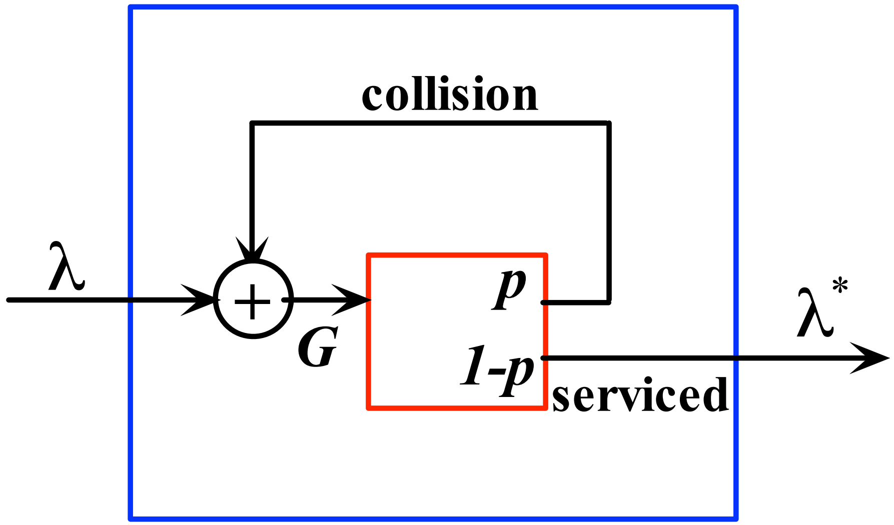

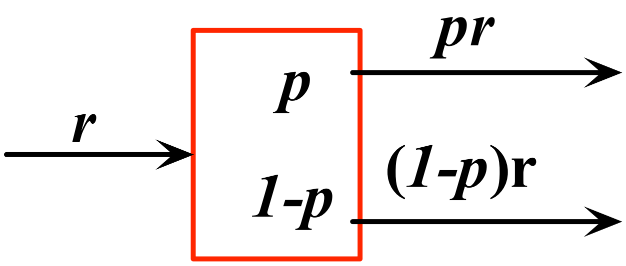

Fact: Bernoulli sampling of a PP is also a PP

$p \equiv$ Probability of collision

$G \equiv$ Total carried load (incl. Retransmissions)

$$\Rightarrow G = \lambda + pG$$

Probability of Collision

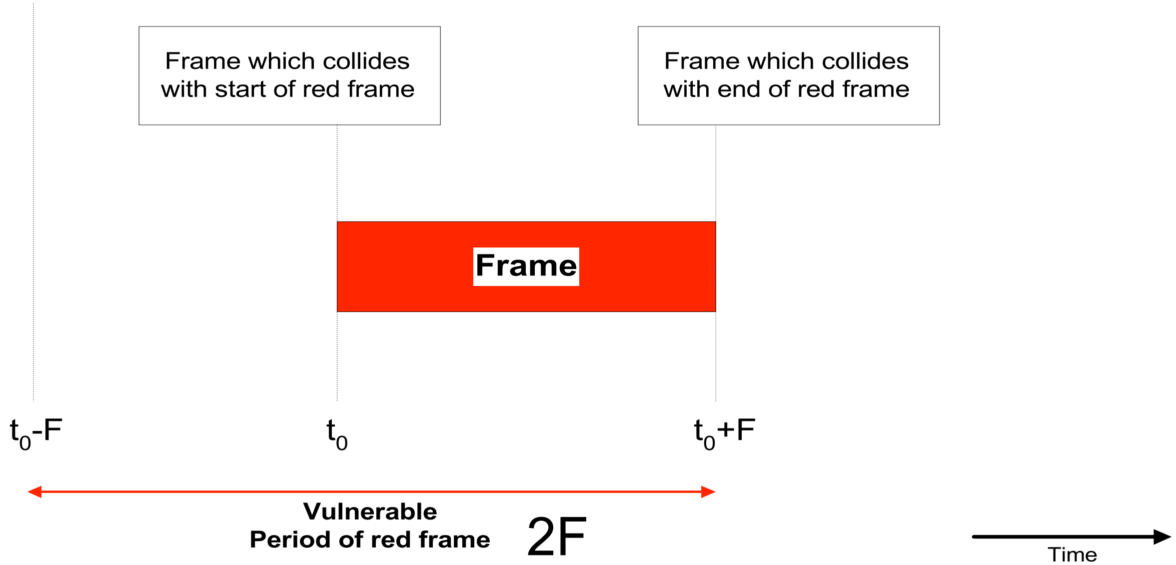

If Poisson events occur at rate λ,

$$P(\text{no event in T seconds}) = e^{-\lambda T}$$

Probability of Collision

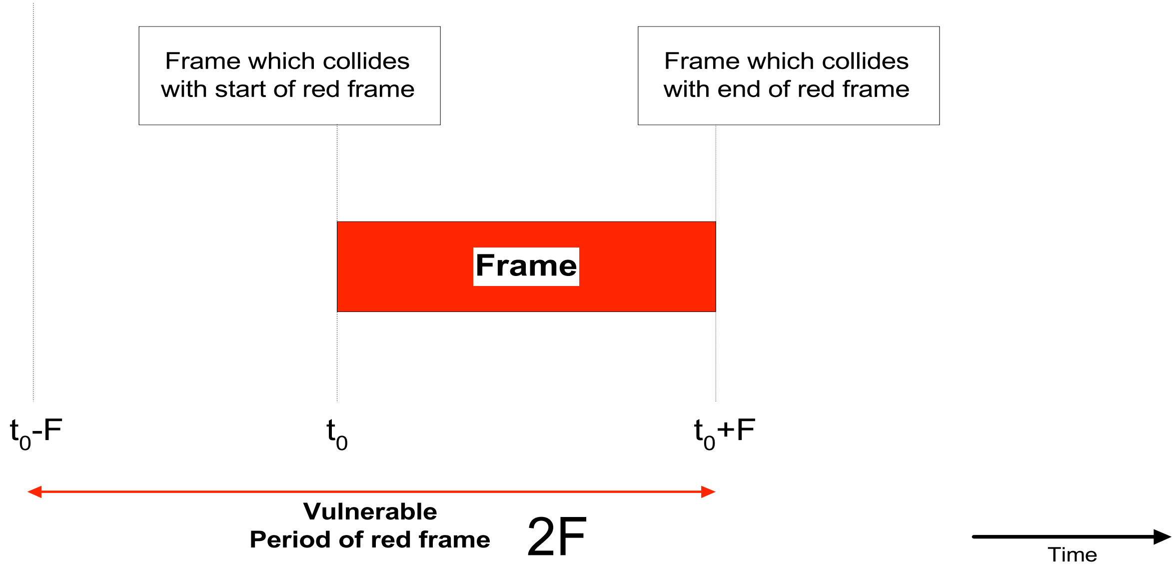

If Poisson events occur at rate λ,

$$P(\text{no event in T seconds}) = e^{-\lambda T}$$

$$p = P(some\_event\_in < 2F) = 1 - e^{-G2F}$$

Throughput and Total Carried Load

From $G = \lambda + pG \quad \Rightarrow \quad \lambda = Ge^{-2GF}$

$R \equiv \text{normalized total carried load} = \dfrac{G}{C/s} = FG$

$$\Rightarrow \quad \rho = F\lambda = \mathrm{R}\mathrm{e}^{-2R}$$

If stable, all offered traffic is serviced, and ρ is also the throughput.

$\text{stable} \equiv E(Delay) < \infty$

Expression characterizes throughput vs. total carried load including retransmissions. What is ALOHA's maximum throughput?

Maximum Throughput

Maximum achievable throughput:

$$\frac{d\rho}{dR} = 0 \Rightarrow R(-2e^{-2R}) + e^{-2R} = 0 \Rightarrow R = \frac{1}{2}$$

$$\rho_{max} = \frac{1}{2}e^{-1} \approx 0.18$$

Observe: if offered load > .18*C, unstable

Performance of ALOHA

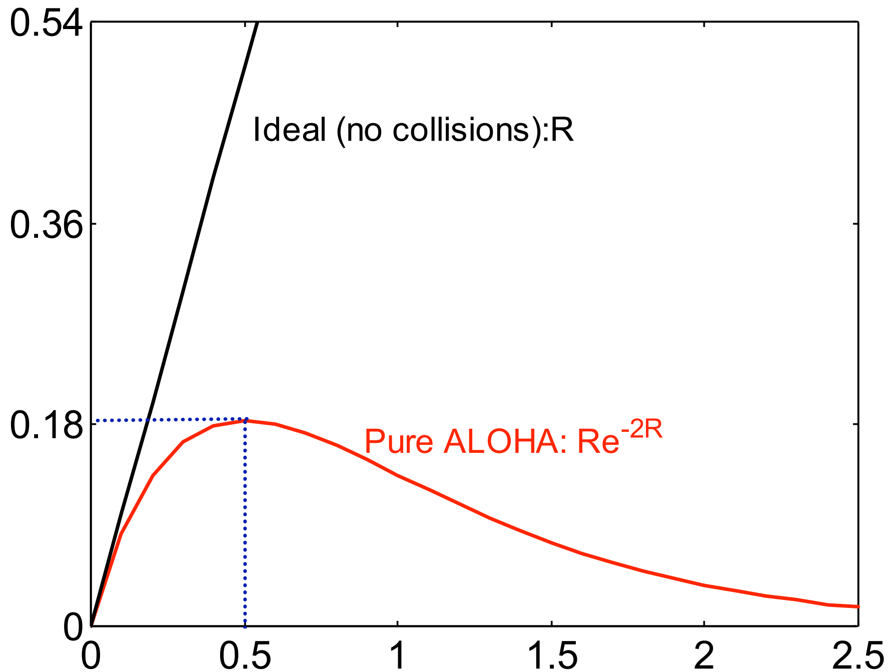

- Maximum throughput approximately 18% of the capacity
- Can do better with improved control
- However, ALOHA is still used for its simplicity
    - Ex. Cell phone call establishment

Slotted ALOHA (S-ALOHA)

- **The Slotted Aloha Protocol**
  - Slotted Aloha - Aloha with an additional constraint
  - Time is divided into discrete time intervals **(=slot)**
  - A station can transmit only at the beginning of a frame

- As a consequence:
  - Frames either collide completely or do not collide at all
  - Vulnerable period = ?

Collisions in S-ALOHA

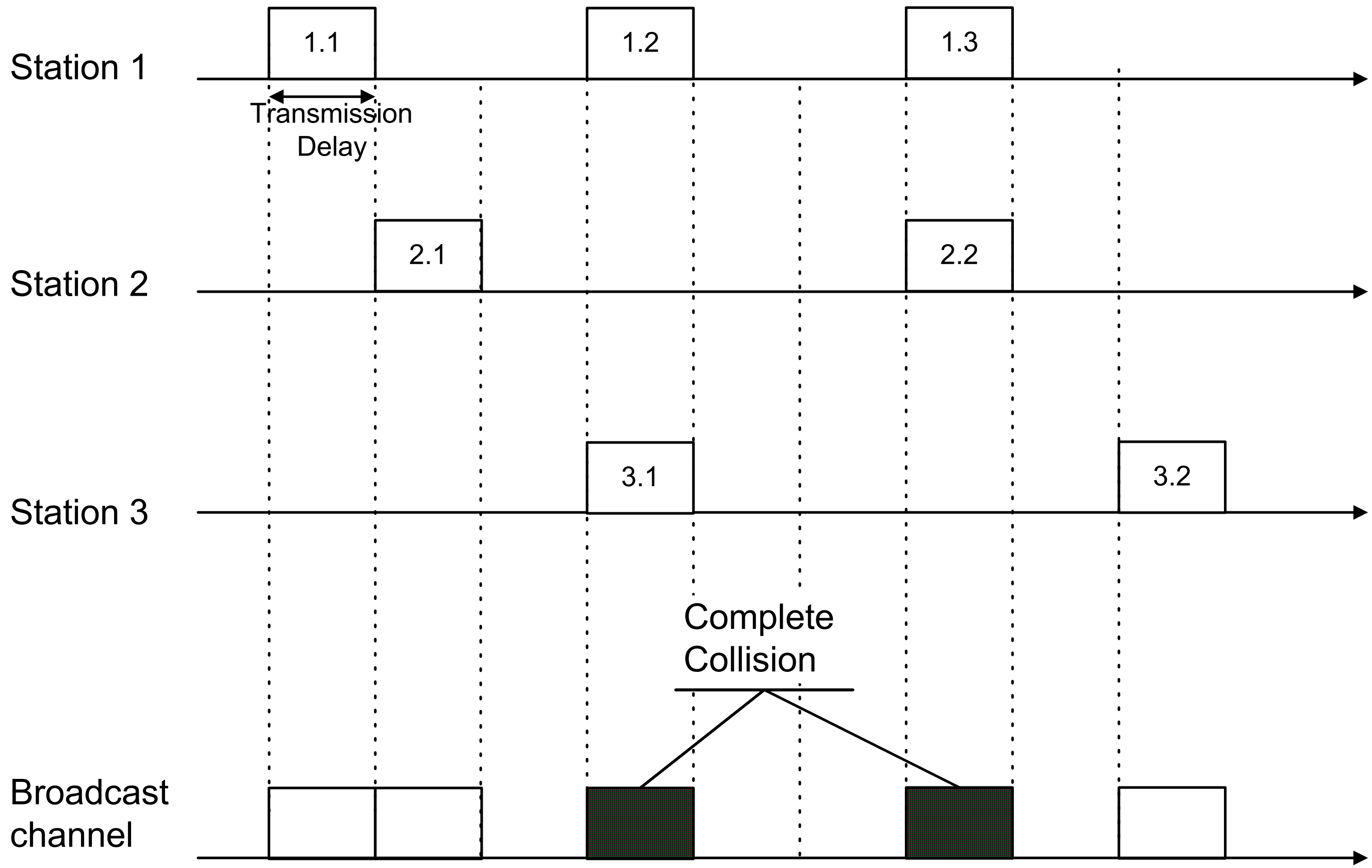

Collisions in S-ALOHA

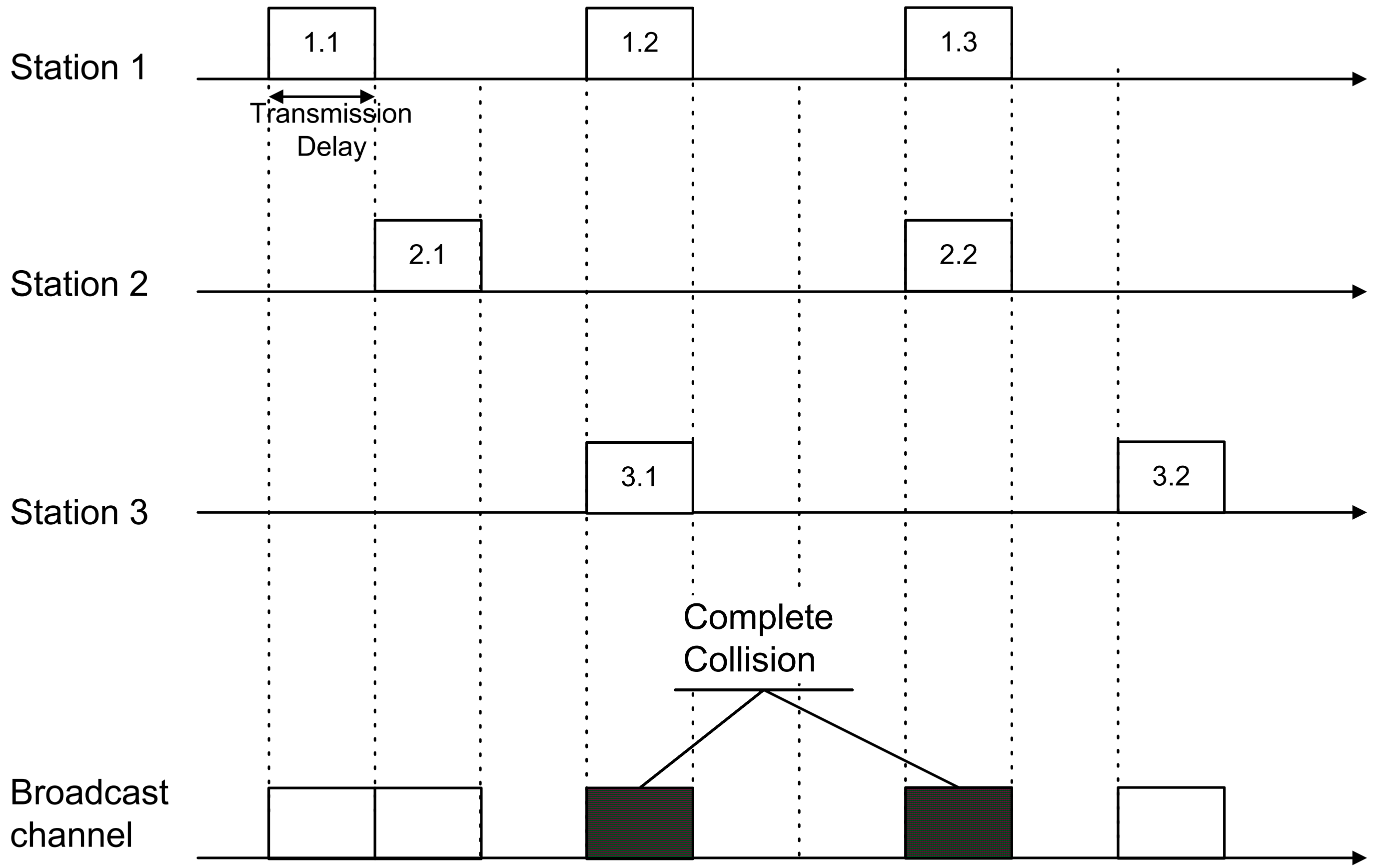

$$p = P(Y_i < F) = 1 - e^{-GF}$$

Performance of S-ALOHA

- Total Throughput in S-ALOHA:

$$\rho = Re^{-R}$$

- Maximum achievable throughput:

$$\rho_{\max} = e^{-1} = 0.37$$

- Performance gain but requires nodes to have synchronized frame boundaries

Comparison of ALOHA and S-ALOHA

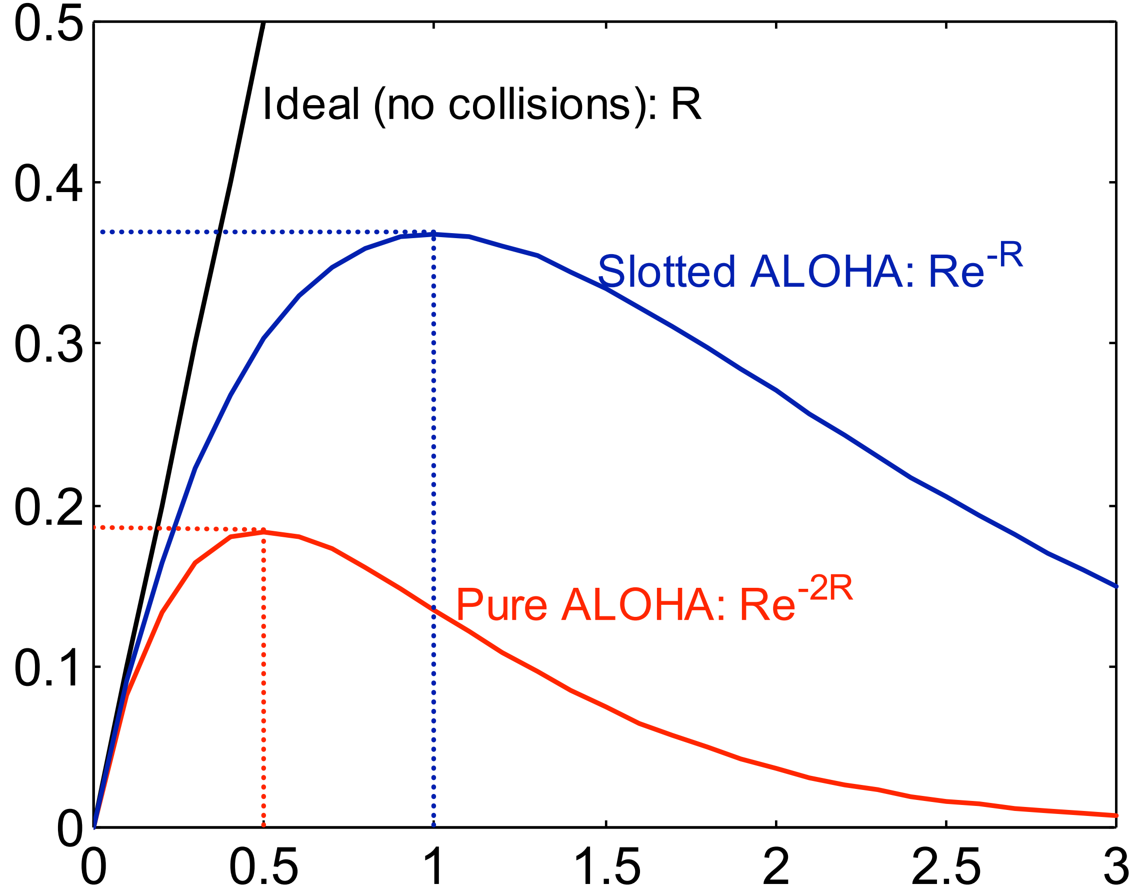

802.3 Ethernet

Broadcast technology

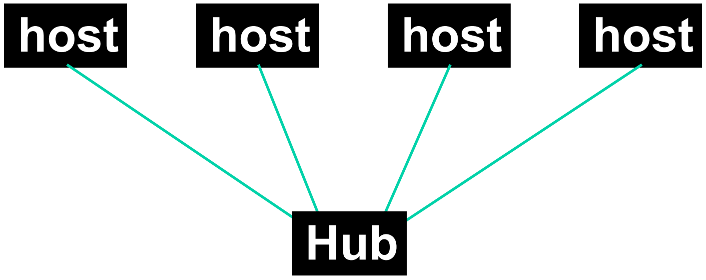

- Carrier-sense multiple access with collision detection (CSMA/CD).
  - MA = multiple access
  - CS = carrier sense
  - CD = collision detection
- Base Ethernet standard is 10 Mbps.
  - Original design was ~2 Mbps
  - 100Mbps, 1Gbps, 10Gbps

CSMA/CD Algorithm

CSMA/CD Algorithm

- Sense for carrier.

CSMA/CD Algorithm

- Sense for carrier.
- If carrier present, wait until carrier ends.
  - Sending would force a collision and waste time

CSMA/CD Algorithm

- Sense for carrier.
- If carrier present, wait until carrier ends.
  - Sending would force a collision and waste time
- Send packet and sense for collision.

CSMA/CD Algorithm

- Sense for carrier.
- If carrier present, wait until carrier ends.
  - Sending would force a collision and waste time
- Send packet and sense for collision.
- If no collision detected, consider packet delivered.

CSMA/CD Algorithm

- Sense for carrier.
- If carrier present, wait until carrier ends.
  - Sending would force a collision and waste time
- Send packet and sense for collision.
- If no collision detected, consider packet delivered.
- Otherwise, abort immediately, perform “exponential back off” and send packet again.
  - Start to send at a random time picked from an interval
  - Length of the interval increases with every retransmission

CSMA/CD: Some Details

- When a sender detects a collision, it sends a “jam signal”.
  - Make sure that all nodes are aware of the collision
  - Length of the jam signal 48 bits
- Exponential backoff operates in multiples of 512 bit time.

CSMA collisions

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

Collision:
entire packet transmission time wasted

Note:
role of distance and propagation delay in determining collision prob.

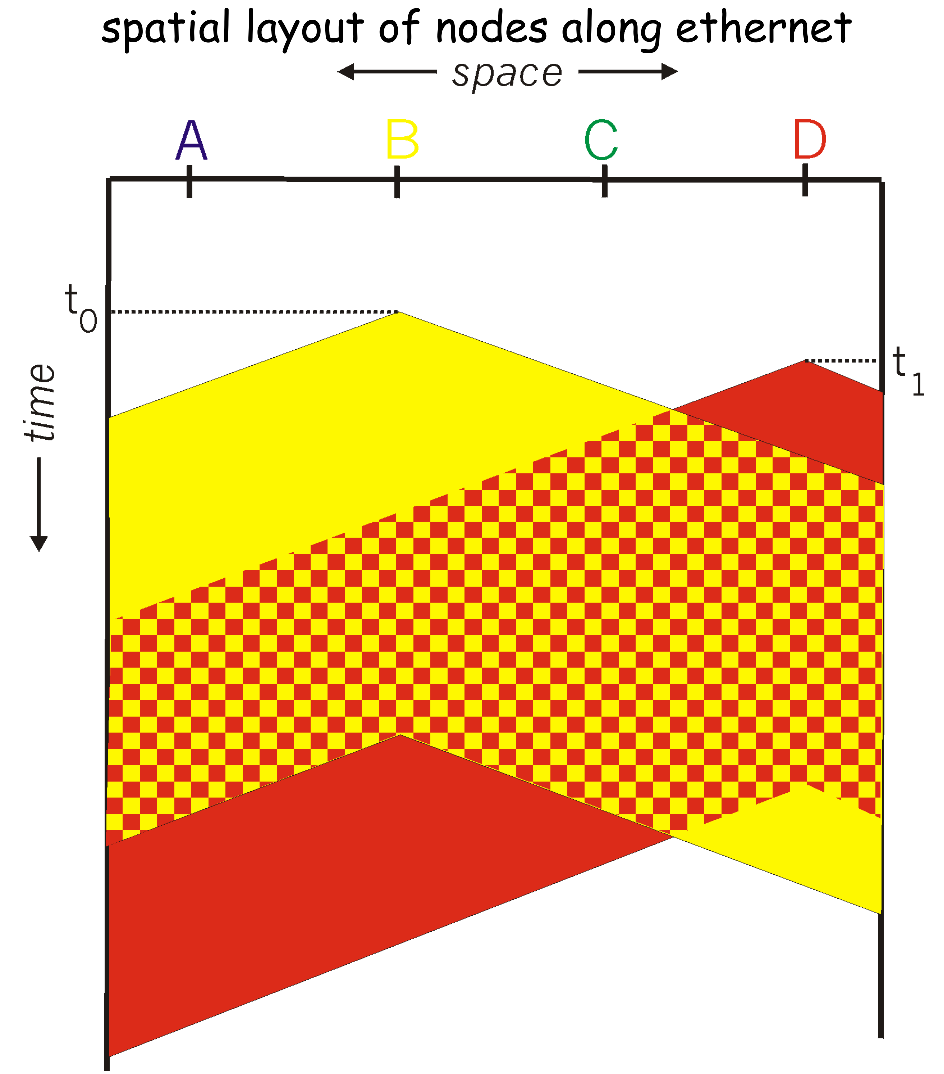

CSMA/CD (Collision Detection)

- Collisions *detected* within short time
- Colliding transmissions aborted, reducing channel wastage
- Easy in wired LANs:
  - measure signal strengths,
  - compare transmitted, received signals
- Difficult in wireless LANs

CSMA/CD collision detection

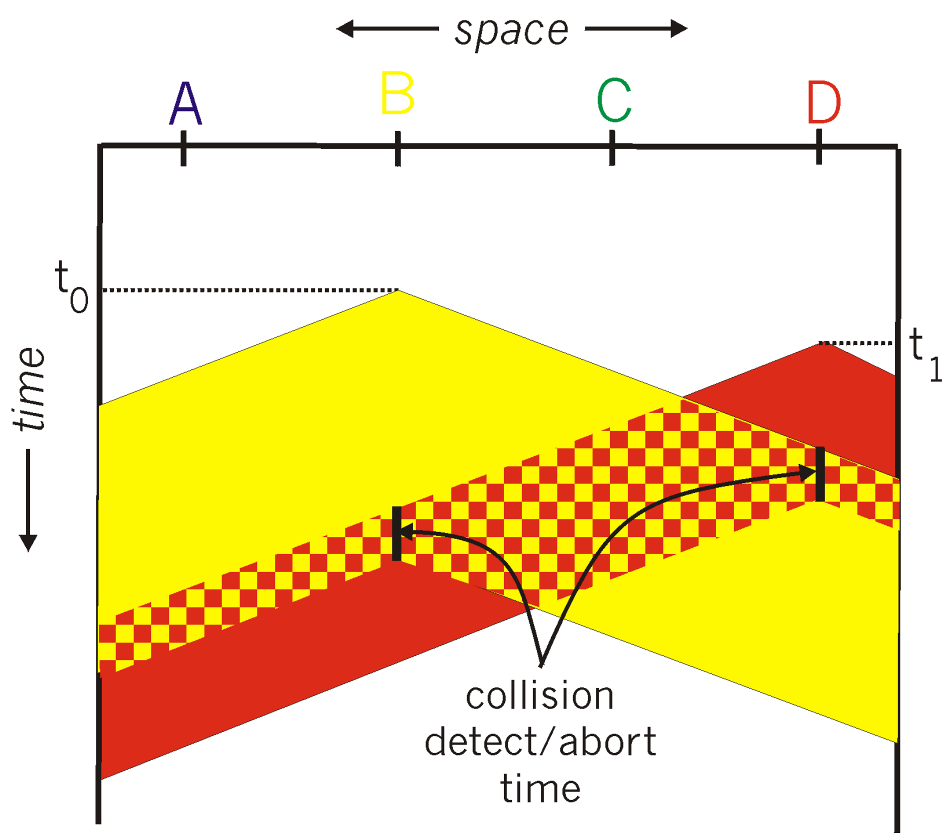

Minimum Packet Size

- Why put a minimum packet size?
- Give a host enough time to detect collisions
- In Ethernet, minimum packet size = 64 bytes (two 6-byte addresses, 2-byte type, 4-byte CRC, and 46 bytes of data)
- If host has less than 46 bytes to send, the adaptor pads (adds) bytes to make it 46 bytes
- What is the relationship between minimum packet size and the length of the LAN?

Minimum Packet Size (more)

a) Time = t; Host 1 starts to send frame

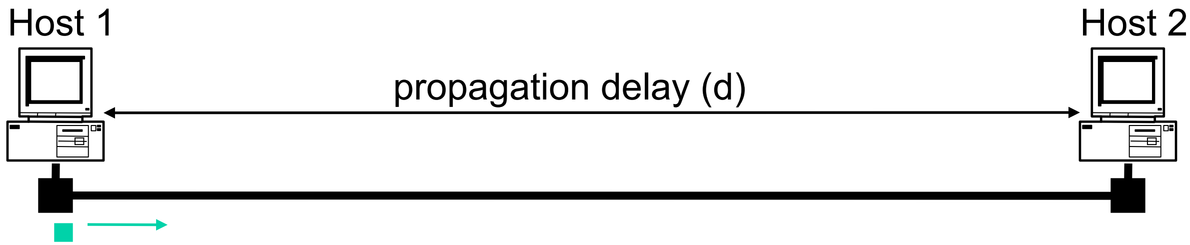

Minimum Packet Size (more)

a) Time = t; Host 1 starts to send frame

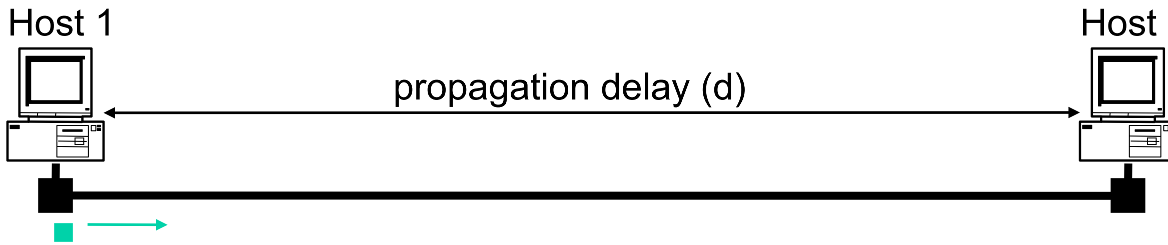

b) Time = t + d; Host 2 starts to send a frame just before it hears from host 1's frame

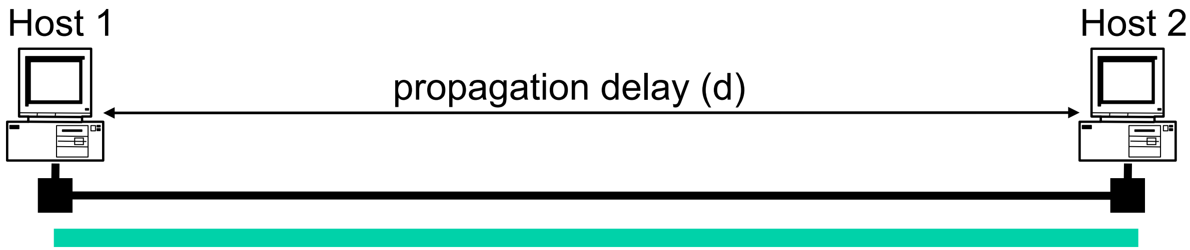

Minimum Packet Size (more)

a) Time = t; Host 1 starts to send frame

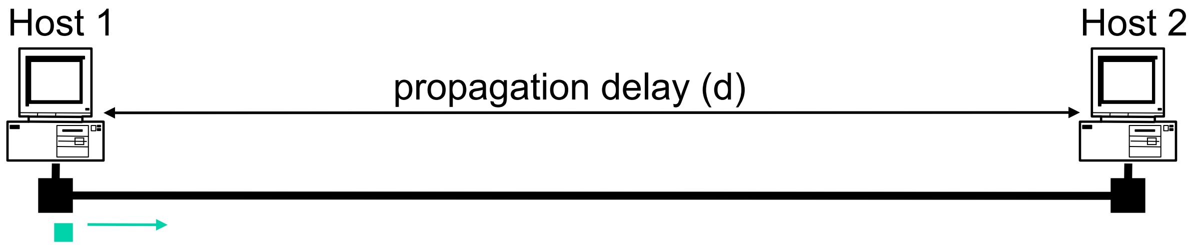

b) Time = t + d; Host 2 starts to send a frame just before it hears from host 1's frame

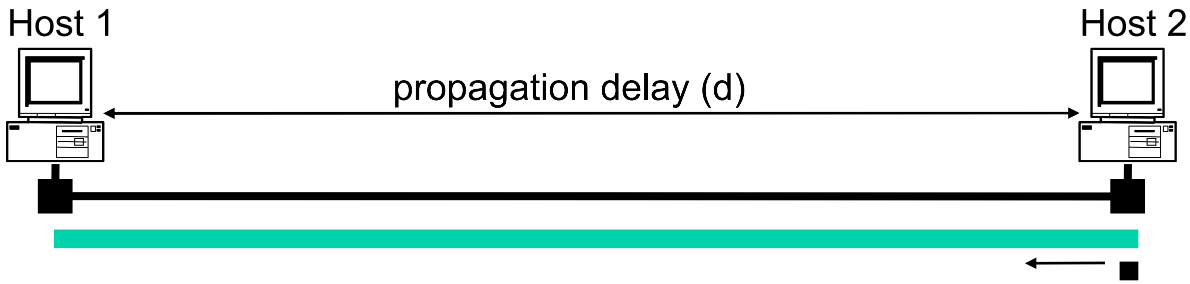

Minimum Packet Size (more)

a) Time = t; Host 1 starts to send frame

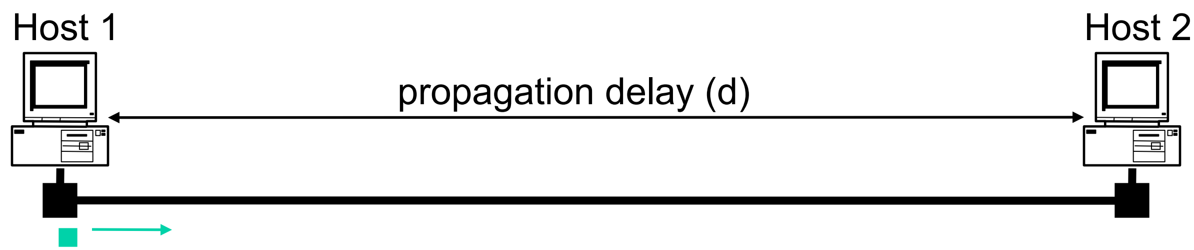

b) Time = t + d; Host 2 starts to send a frame just before it hears from host 1's frame

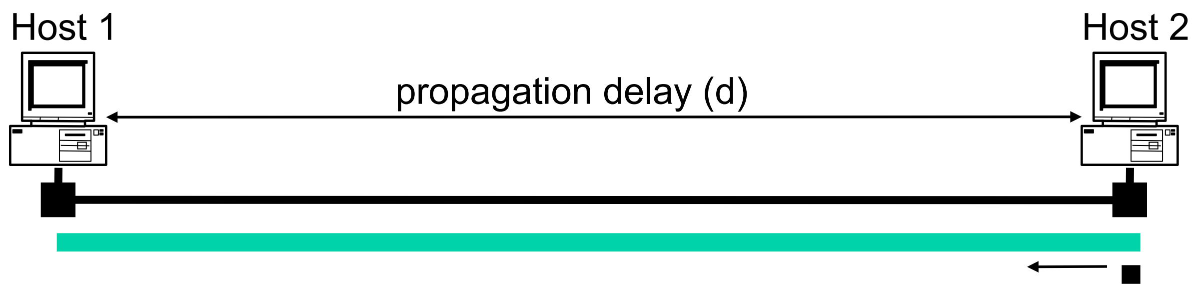

c) Time = t + 2*d; Host 1 hears Host 2's frame → detects collision

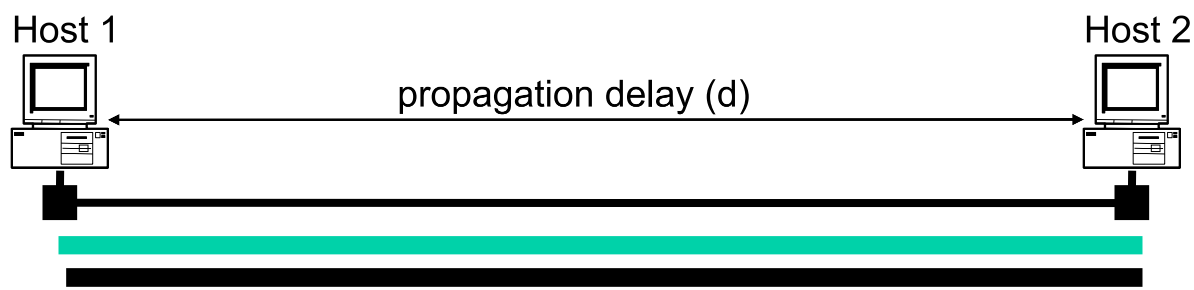

Minimum Packet Size (more)

a) Time = t; Host 1 starts to send frame

b) Time = t + d; Host 2 starts to send a frame just before it hears from host 1's frame

c) Time = t + 2*d; Host 1 hears Host 2's frame → detects collision

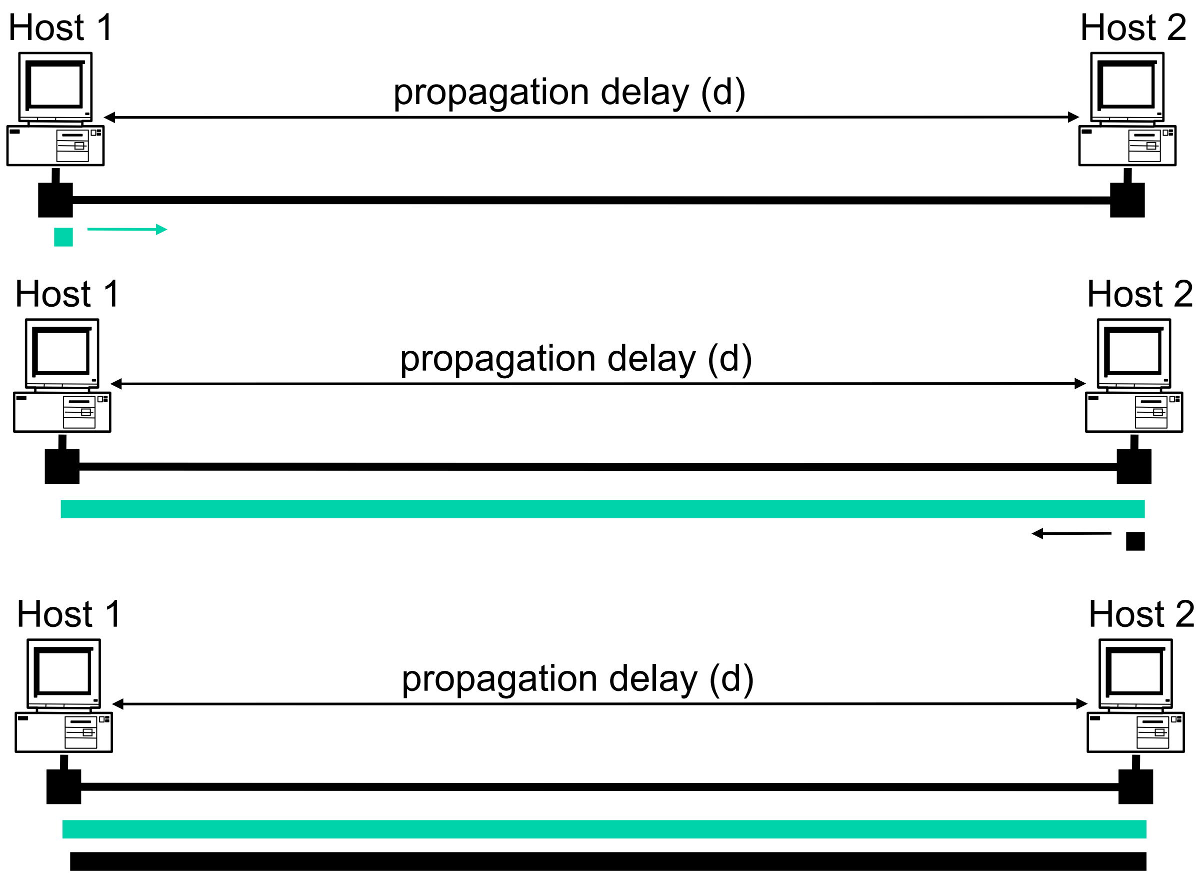

$$\text{LAN length} = (\text{min\_frame\_size})*(\text{light\_speed})/(2*\text{bandwidth}) =$$
$$= (8*64b)*(2.5*10^8 \text{mps})/(2*10^7 \text{ bps}) = 6400\text{m approx}$$

Exponential Backoff Algorithm

- Ethernet uses the exponential backoff algorithms to determine when a station can retransmit after a collision

Algorithm:

- Set "slot time" equal to 512bit time
- After first collision wait 0 or 1 slot times
- After i-th collision, wait a random number between 0 and 2^i-1 time slots
- Do not increase random number range, if i=10
- Give up after 16 collisions

CSMA/CD Contention Interval

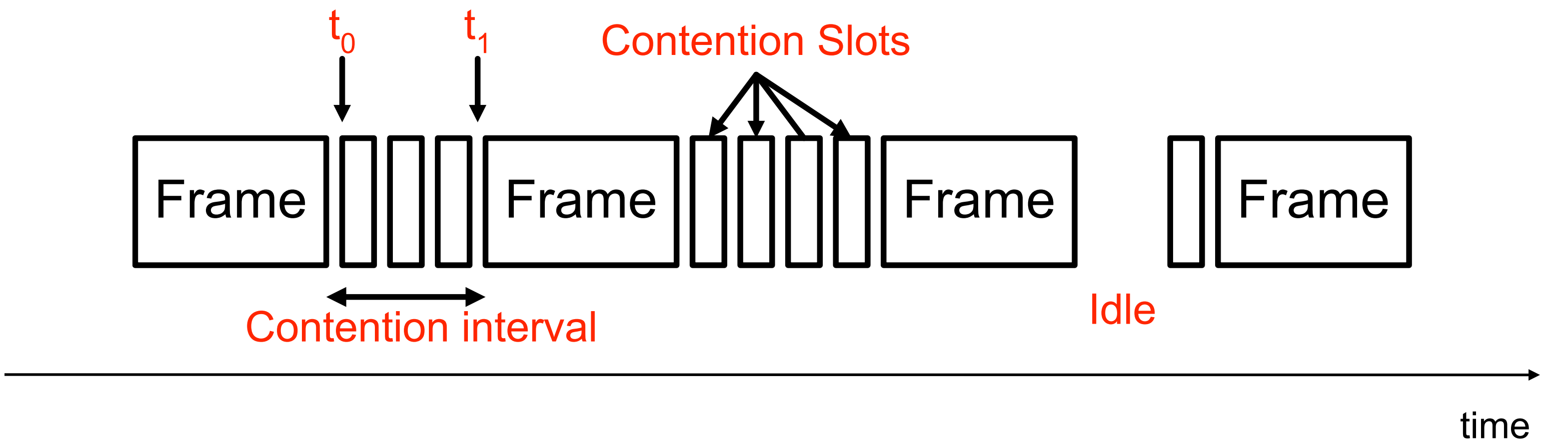

- Contention slots end in a collision
- Contention interval is a sequence of contention slots
- Length of a slot in contention interval is 512 bit time

Min packet size & slot time

- Min packet size is 512 bits
- Slot time is the transmission of 512 bits
- Coincident?
- If slot time is the transmission of 256 bits, then two stations picking 0 and 1 slot to wait respectively can still collide

Ethernet Frame Structure

- Sending adapter encapsulates IP datagram

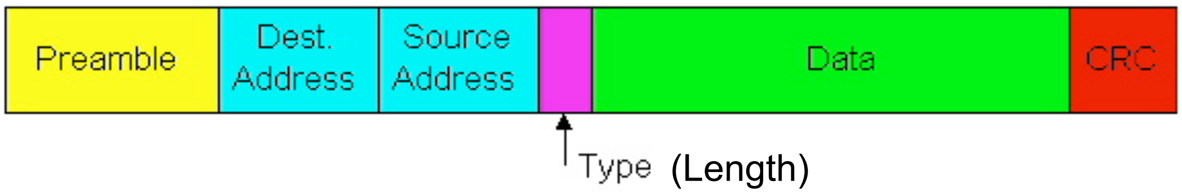

- Preamble:
  - 7 bytes with pattern 10101010 followed by one byte with pattern 10101011
  - Used to synchronize receiver, sender clock rates

Ethernet Frame Structure (more)

- Addresses: 6 bytes, frame is received by all adapters on a LAN and dropped if address does not match
- Type: 2 bytes, is actually a length field in 802.3
- CRC: 4 bytes, checked at receiver, if error is detected, the frame is simply dropped
- Data payload: maximum 1500 bytes, minimum 46 bytes
  - If data is less than 46 bytes, pad with zeros to 46 bytes

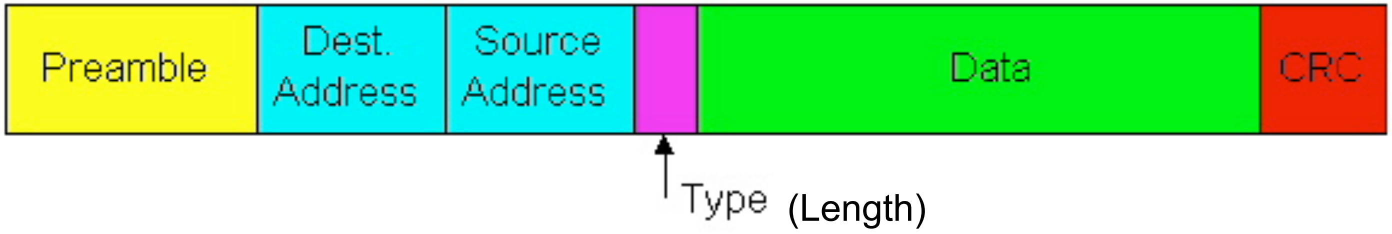

Ethernet Technologies: 10Base2

- 10: 10Mbps; 2: under 200 meters max cable length
- Thin coaxial cable in a bus topology

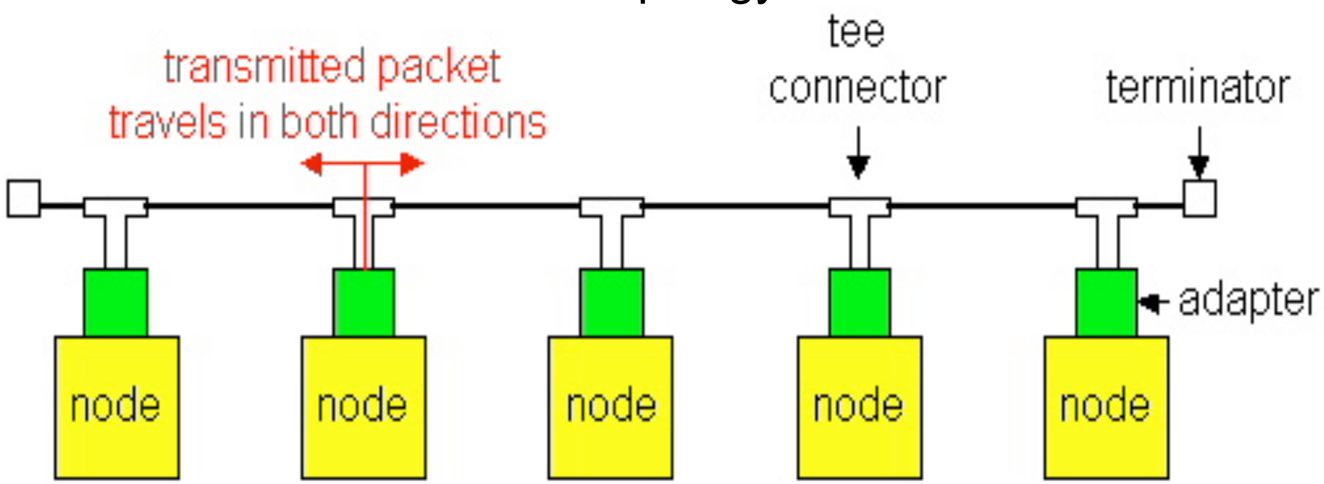

- Repeaters used to connect up to multiple segments
- Repeater repeats bits it hears on one interface to its other interfaces: physical layer device only!

10BaseT and 100BaseT

- 10/100 Mbps rate; latter called “fast ethernet”
- T stands for Twisted Pair
- Hub to which nodes are connected by twisted pair, thus “star topology”

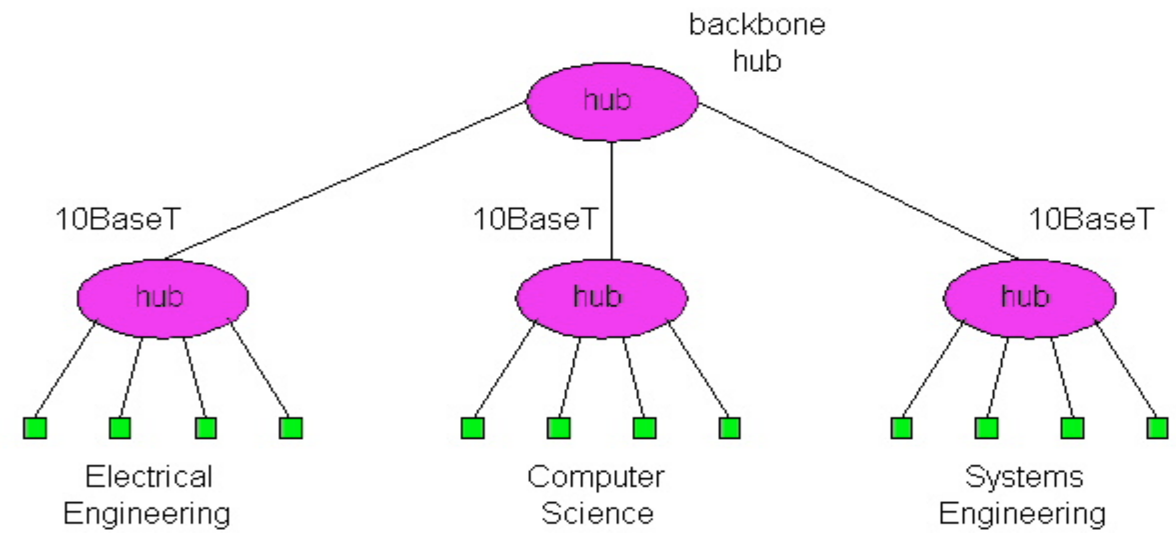

10BaseT and 100BaseT (more)

- Max distance from node to Hub is 100 meters
- Hub can gather monitoring information, statistics for display to LAN administrators

- Hubs still preserve one collision domain
  - Every packet is forwarded to all hosts
- Use bridges to address this problem
  - Bridges forward a packet only to the destination leading to the destination
  - Next lecture

Gbit Ethernet

- Use standard Ethernet frame format
- Allows for point-to-point links and shared broadcast channels
- In shared mode, CSMA/CD is used; short distances between nodes to be efficient
- Full-Duplex at 1 Gbps for point-to-point links

802.3 Ethernet vs 802.11 Wi-Fi

- Ethernet: one shared “collision” domain
- 802.11: radios have small range compared to overall system: collisions are local
  - collisions are at receiver, not sender
  - carrier-sense plays different role
- CSMA/CA not CSMA/CD
  - collision avoidance, not collision detection

Collision Avoidance: The Problems

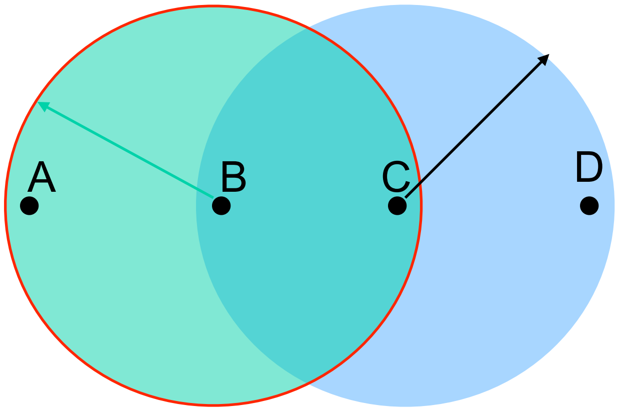

Collision Avoidance: The Problems

- Reachability is not transitive: if A can reach B, and B can reach C, it doesn't necessary mean that A can reach C

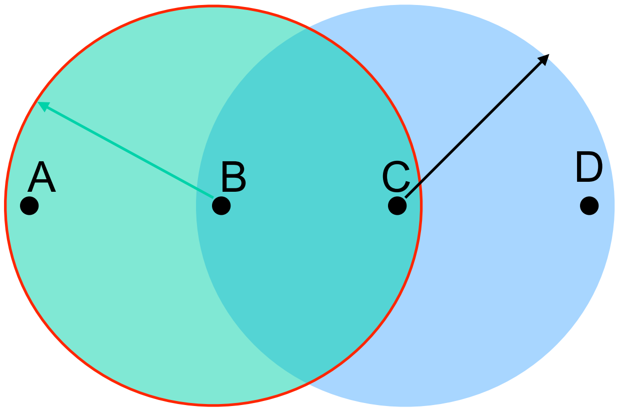

Collision Avoidance: The Problems

- Reachability is not transitive: if A can reach B, and B can reach C, it doesn't necessary mean that A can reach C

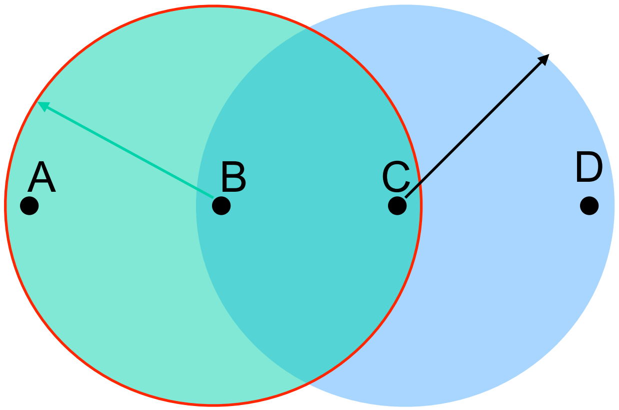

Collision Avoidance: The Problems

- Reachability is not transitive: if A can reach B, and B can reach C, it doesn't necessary mean that A can reach C

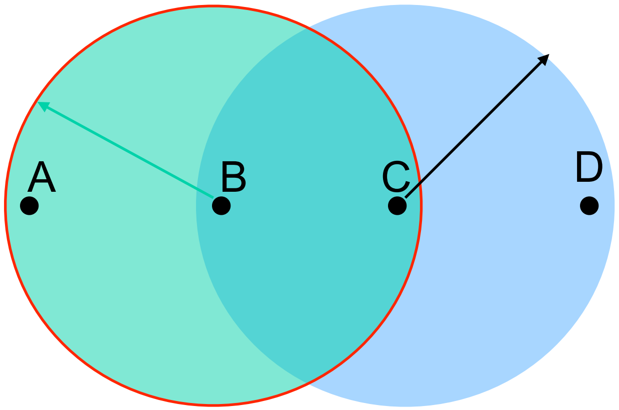

Collision Avoidance: The Problems

- Reachability is not transitive: if A can reach B, and B can reach C, it doesn't necessary mean that A can reach C

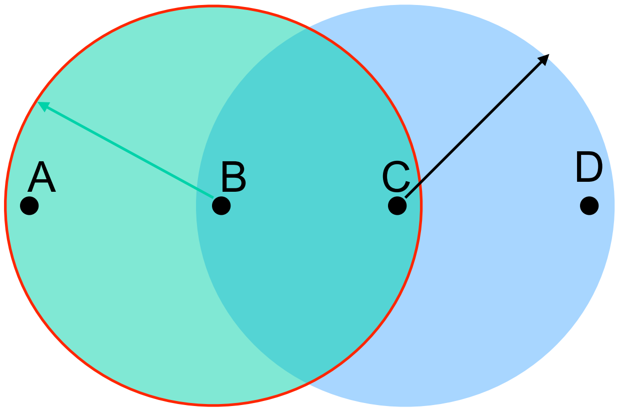

Collision Avoidance: The Problems

- Reachability is not transitive: if A can reach B, and B can reach C, it doesn't necessary mean that A can reach C

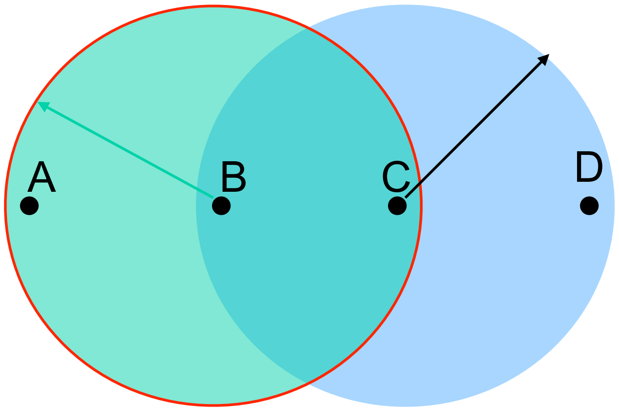

Collision Avoidance: The Problems

- Reachability is not transitive: if A can reach B, and B can reach C, it doesn’t necessary mean that A can reach C

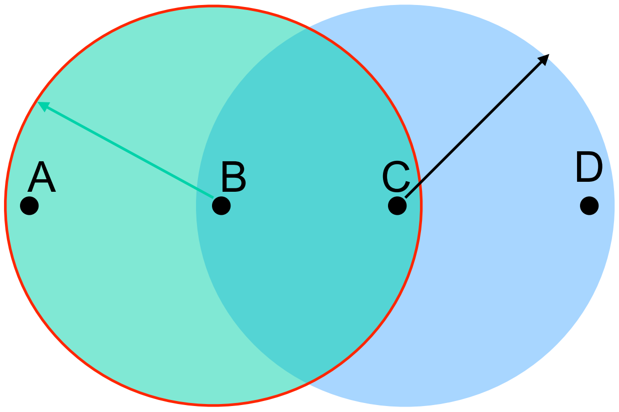

Collision Avoidance: The Problems

- Reachability is not transitive: if A can reach B, and B can reach C, it doesn't necessary mean that A can reach C

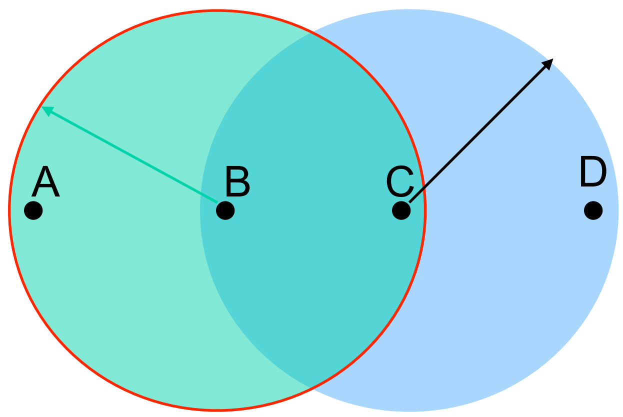

- Hidden nodes: A and C send a packet to B; neither A nor C will detect the collision!

Collision Avoidance: The Problems

- Reachability is not transitive: if A can reach B, and B can reach C, it doesn't necessary mean that A can reach C

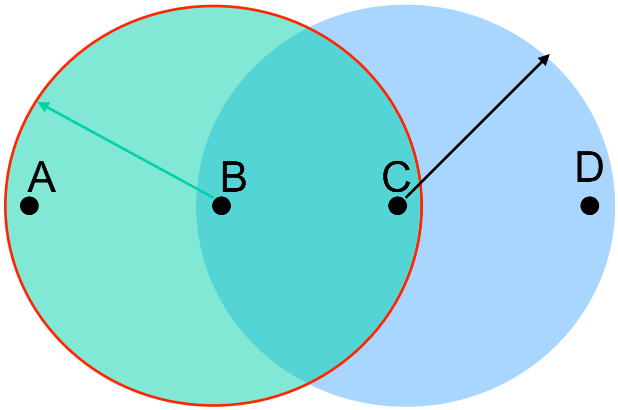

- Hidden nodes: A and C send a packet to B; neither A nor C will detect the collision!
- Exposed node: B sends a packet to A; C hears this and decides not to send a packet to D (despite the fact that this will not cause interference)!

Multiple Access with Collision Avoidance (MACA)

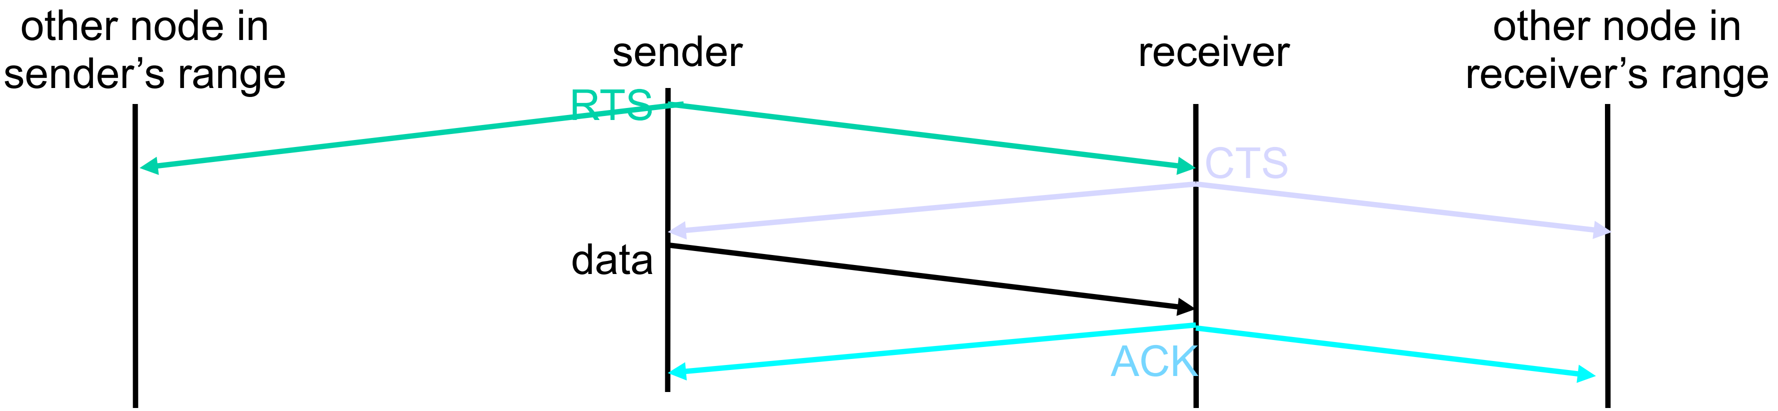

- Before every data transmission
  - Sender sends a Request to Send (RTS) frame containing the length of the transmission
  - Receiver respond with a Clear to Send (CTS) frame
  - Sender sends data
  - Receiver sends an ACK; now another sender can send data
- When sender doesn't get a CTS back, it assumes collision

Other Nodes

- When you hear a CTS, you keep quiet until scheduled transmission is over (hear ACK)
- If you hear RTS, but not CTS, you can send
  - interfering at source but not at receiver is ok
  - can cause problems when a CTS is interfered with