Media Access Control (MAC)

Roman Dunaytsev

The Bonch-Bruevich Saint-Petersburg
State University of Telecommunications

roman.dunaytsev@spbgut.ru

Lecture № 5

Outline

Outline

Types of MAC

- Why MAC?
- Generally, all networks can be grouped into 2 categories:
  - Non-broadcast
  - Broadcast
- **Non-broadcast networks** (aka **switched**) provide interconnection between users by means of transmission lines, multiplexers, and switches
- A key characteristic of non-broadcast networks is **adjacency** – nodes can only communicate with nodes they are next to
  - The best way of describing a non-broadcast network is to imagine a mail correspondence – at any time, only the sender and the receiver of the mail know what the conversation is about, the rest of the people do not
  - Examples of non-broadcast networks are Frame Relay and ATM

Types of MAC (cont'd)

- In **broadcast networks** (aka **shared-medium**), a single transmission medium is shared by a number of users
- The information from a user is broadcast into the medium and all stations attached to the medium listen to all transmissions
  - Broadcast networks are much simpler than non-broadcast ones: because all information is received by all users, dedicated wiring (point-to-point) and switching are not necessary
  - To provide low cost and simplicity, **LANs have been traditionally based on the broadcast approach**
- When 2 or more stations transmit simultaneously, their signals will collide and interfere with each other
  - So a technique must be provided for to prevent or minimize such interference
- Thus, broadcast networks need MAC for the same reason streets need traffic lights and rules of the road – to prevent traffic accidents

Types of MAC (cont'd)

Traffic code	MAC
Regulates the use of the road	Regulates the use of the transmission medium
Determines who can drive and at what time	Determines who can send data and at what time
If 2 or more cars try to use the same place at the same time, the cars collide, causing car crash	If 2 or more stations try to send data at the same time, the signals interfere with each other, corrupting the data being transmitted
Prevents or minimizes automobile accidents	Prevents or minimizes data collisions

Types of MAC (cont'd)

- In **broadcast networks**, the physical transmission medium is a shared resource
- Thus, access to this shared resource needs to be coordinated either centrally or in a decentralized fashion
- The goal of controlled access is to avoid or minimize simultaneous transmission attempts (that will result in data collisions)
- **Collisions** occur if multiple stations transmit at the same time and the destination cannot resolve the composite signal created due to this uncontrolled superposition
- **Media Access Control (MAC)** refers to the algorithms used by devices to control access to the shared transmission medium
  - MAC is especially important in broadcast LANs
  - WANs, in contrast, use point-to-point links, except for satellite networks

Types of MAC (cont'd)

- MAC techniques can be classified based on their functionality as:
  - Centralized vs. distributed control
  - Deterministic vs. random access
  - Fixed vs. dynamic resource allocation
  - etc.
- Some of MAC techniques are variants of others
- **There exists no technique that performs better than all others over the entire range of performance criteria**
  - Each class has its own benefits and shortcomings
- MAC techniques can also be classified based on types of users:
  - **Passive users** may access the transmission medium only when specifically polled by the central controller
  - **Active users** actively seek access to the transmission medium instead of waiting to be polled

Types of MAC (cont'd)

- **Centralized control** – a **central controller** polls stations to specify when access and transmission by each station is allowed to occur
- Stations transmit data (if any) when requested to do so, or when a station transmission request is acknowledged and granted
- **Benefits** of centralized MAC:
  - It provides great flexibility, including priorities and guaranteed capacity
  - It enables the use of relatively simple access logic at each station
- **Shortcomings** of centralized MAC:
  - The process of polling requires the passing of control frames, adding a significant overhead under light traffic load
  - The central controller may act as a bottleneck, reducing the overall performance
  - The failure of the central controller will disrupt the entire network
  - But in practice this rarely occurs because one of the controllers is used as an active controller and the other is used as a standby controller

Types of MAC (cont'd)

- **Distributed control** – each station is responsible for controlling its access to the transmission medium
- Additionally, each station is charged with responsibility for detecting and resolving any data collisions which might occur due to transmission overlap with other stations
- Thus, all stations **collectively** perform MAC
- **Benefits** of distributed MAC:
  - Distributed control schemes have no single point of overall failure
- **Shortcomings** of distributed MAC:
  - It does not inherently provide Quality of Service (QoS) guarantees
  - Quite complex access logic should be implemented at each station
  - Performance tends to collapse under heavy traffic load
- The benefits and shortcomings of distributed MAC are mirror images of the pros and cons of centralized MAC

Types of MAC (cont'd)

- **Deterministic access** – MAC allows to determine the maximum medium access delay
  - I.e., each station can be guaranteed the right to communicate within a certain time interval
- Deterministic access is also known as **contention-free**, because the stations do not contend for access
- **Benefits** of contention-free MAC:
  - Contention-free techniques are especially efficient under heavy traffic load, since they completely eliminate collisions through pre-allocation of the resources
- **Shortcomings** of contention-free MAC:
  - Contention-free techniques lead to inefficient resource utilization, adding a significant overhead under light traffic load

Types of MAC (cont'd)

- **Random access** – MAC does not allow to determine exactly (only as a probabilistic estimate) the maximum medium access delay
  - I.e., each station can not be guaranteed the right to communicate within a certain time interval
- Random access is also known as **contention-oriented**
  - No control is exercised to determine whose turn to transmit
  - Rather, all stations contend for access
- **Benefits** of contention-oriented MAC:
  - They are simple to implement and efficient under light traffic load
- **Disadvantage** of contention-oriented MAC:
  - Performance tends to collapse under heavy traffic load
- Again, the benefits and shortcomings of contention-oriented MAC are mirror images of the pros and cons of contention-free MAC

Types of MAC (cont'd)

- Taxonomy of MAC techniques

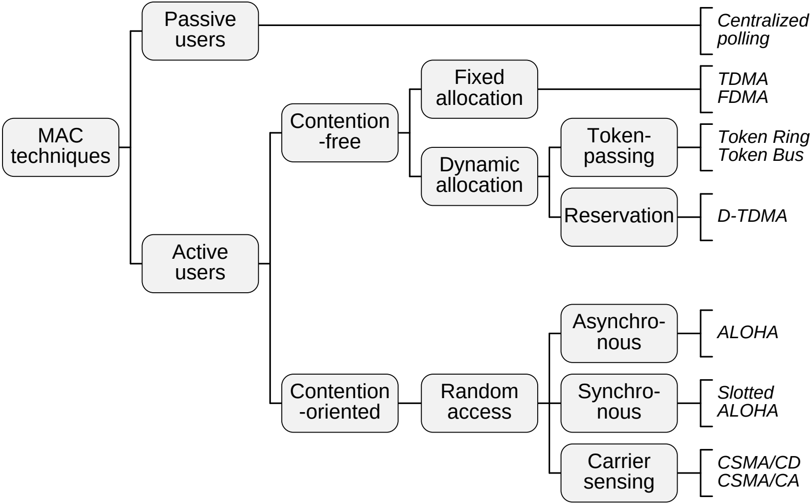

Outline

1. Types of MAC
2. Contention-free MAC
   - Centralized polling
   - Resource allocation
   - Token-passing
   - Reservation
3. Contention-oriented MAC
   - ALOHA
   - Slotted ALOHA
   - CSMA
   - CSMA/CD
   - CSMA/CA

Centralized Polling

- In MAC with passive users, stations can access the transmission medium only when specifically polled by the central controller
- A central controller transmits **polling messages** to stations according to some predefined order
  - Usually, based on the round-robin scheduling
- With **round-robin**, each station in turn is given the opportunity to transmit data
  - During that opportunity, the station may decline to transmit (nothing to transmit) or may transmit subject to a specified upper bound

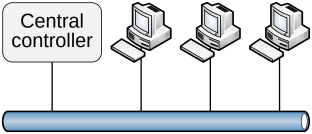

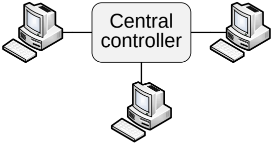

Centralized Polling (cont'd)

- How much is a station allowed to transmit per polling message?
- **Exhaustive service** – until the send buffer is empty (including new frame arrivals)
- **Gated service** – all data in the send buffer when a polling message arrives
  - In both cases, the station indicates the completion of its transmission through a **go-ahead message**
- **Frame-limited service** – 1 frame per polling message
- **Time-limited service** – transmission time is limited to some predefined maximum duration

Centralized Polling (cont'd)

- **The polling cycle** starts by polling the first station
- A certain time, called **the walk time**, elapses while the polling message propagates and is received and until the station begins its transmission
- The next period is occupied by the transmission from this station
- This period is followed by the walk time that elapses while the next station is polled and then by the transmission from it
- This process continues until the last station is polled and has completed its transmission

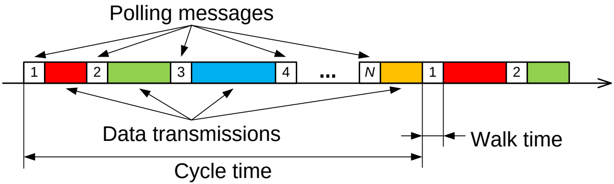

Centralized Polling (cont'd)

- **Exhaustive/gated service**
  - As the load increases, the cycle time can increase without bound
  - Medium access delays increase as well
  - But **the total walk time** (the sum of the walk times in the cycle) remains constant
  - The walk time can be viewed as a form of overhead
  - As a result, the overhead due to polling becomes negligible
  - Thus, a polling LAN can achieve almost 100% network utilization when stations are allowed to send all of the frames in their buffers
- **Frame-/time-limited service**
  - In situations where stations carry delay-sensitive traffic, it is desirable to have a cycle time that has a strict upper bound
  - The stations will receive polling messages at some minimum frequency
  - However, the maximum network utilization will now be less than 100%

$$\text{Network utilization} = \left(\frac{\text{Cycle time} - \text{Total walk time}}{\text{Cycle time}}\right) * 100\%$$

Resource Allocation

- MAC can be made contention-free either by implementing fixed or dynamic resource allocation
  - I.e., a certain resource is assigned to a station either on a continuing or temporarily basis
- **Fixed allocation** – generally not optimal in LANs because the needs of stations are usually unpredictable
  - E.g., TDMA, FDMA
- **Dynamic allocation** – particularly useful for meeting the varying requirements on channel capacity, QoS, etc.
  - E.g., Token Ring, Token Bus, D-TDMA

Token-Passing

- Token-passing can be considered as **decentralized polling**
  - I.e., receiving a free token corresponds to receiving a polling message
- Token-passing MAC is based on the use of a special bit pattern, called **a token**, that circulates between stations
  - Even when all stations are idle
- A token is used to assign the right to transmit data
- A token has a special **flag**, which indicates the **token status**
  - If it is '0', the token has not been taken and a station that has sufficient priority may capture it
  - If it is '1', the token has been captured by another station and the frame is in use (contains data)

Token-Passing (cont'd)

- Data frame and token formats

1 byte	1 byte	1 byte	2 or 6 bytes	2 or 6 bytes		4 bytes	1 byte
Preamble	Start delimiter	Frame control	Destination address	Source address	Data	FCS	End delimiter

Token Bus data frame

1 byte	1 byte	1 byte	2 or 6 bytes	2 or 6 bytes	4 bytes	1 byte
Preamble	Start delimiter	Frame control	Destination address	Source address	FCS	End delimiter

Token Bus token

1 byte	1 byte	1 byte	2 or 6 bytes	2 or 6 bytes		4 bytes	1 byte	1 byte
Start delimiter	Access control	Frame control	Destination address	Source address	Data	FCS	End delimiter	Frame status

Token Ring data frame

1 byte	1 byte	1 byte
Start delimiter	Access control	End delimiter

Token Ring token

2 or more bytes	1 byte	1 byte	2 or 6 bytes	2 or 6 bytes		4 bytes	1 byte	1 byte
Preamble	Start delimiter	Frame control	Destination address	Source address	Data	FCS	End delimiter	Frame status

FDDI data frame

2 or more bytes	1 byte	1 byte	1 byte
Preamble	Start delimiter	Frame control	End delimiter

FDDI token

Token-Passing (cont'd)

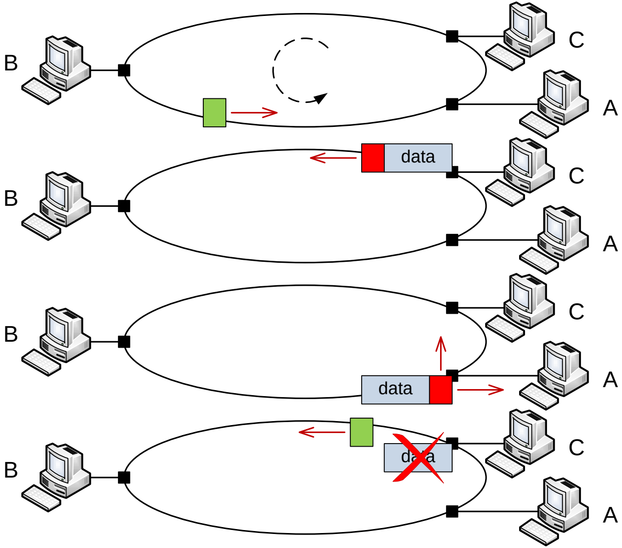

Token-Passing (cont'd)

- Let us suppose that there are 3 stations: A, B, and C:
  - Station C wishes to transmit a frame of data to station A
  - First, station C captures the token and transmits the frame, which circulates through the ring
  - The frame header includes the address of station A
  - As the frame propagates along the ring, station B receives it
  - Since the destination is station A, station B examines the frame header and then passes the frame to the next station
  - In turn, station A sees that the frame is addressed to itself and, therefore, copies the data from the frame into the local buffer
  - Station A also confirms the reception to station C by changing value of a special bit in the 'Frame Status' field, which provides a form of acknowledgement
  - The frame continues its traveling until it reaches station C
  - Station C removes the returning frame and releases a new token

Token-Passing (cont'd)

- How much is a station allowed to transmit per token?
- **Exhaustive/gated service** – allow a station to transmit an unlimited number of frames each time a token is received
  - This approach minimizes the delay experienced by frames from the transmitting station
  - But allows the time between consecutive arrivals of a free token to a station to be unbounded
- **Frame-/time-limited service** – place a certain limit either on the number of frames that can be transmitted each time a token is received, or on the total time that a station may transmit data into the network
  - These limits have the effect of placing an upper bound on the time between consecutive arrivals of a free token to a station

Token-Passing (cont'd)

- **Benefits** of token-passing MAC:
  - When many stations have data to transmit over an extended period of time, token-passing with the round-robin scheduling can be very effective
- **Shortcomings** of token-passing MAC:
  - If only a few stations have data to transmit over an extended period of time, then there is a considerable overhead in passing the turn from station to station, because most of the stations will not transmit but simply pass their turns
- MAC efficiency largely depends on the data traffic characteristics:
  - **Bursty traffic** – short, sporadic transmissions (e.g., Web-surfing)
  - **Stream traffic** – lengthy, continuous transmissions (bulk transfers)

Reservation

- **For stream traffic, reservation techniques are well suited**
- For reservation techniques, time on the transmission medium is divided into slots, much as with **Time Division Multiple Access (TDMA)**
- A station wishing to transmit reserves future slots for an extended or even indefinite period
- An example of reservation with dynamic allocation is **Dynamic Time Division Multiple Access (D-TDMA)**
- D-TDMA improves TDMA by taking slots that would have gone idle and making them available to other stations
  - D-TDMA is used in IEEE 802.16 WiMAX and ETSI HIPERLAN/2 (HIgh PErformance Radio LAN, version 2)

Reservation (cont'd)

- Let us suppose that there are 4 stations

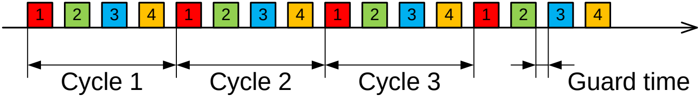

Time Division Multiple Access (TDMA)

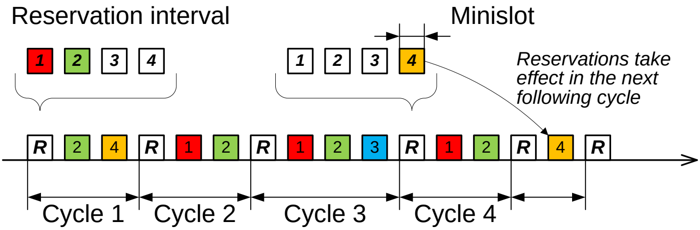

Dynamic Time Division Multiple Access (D-TDMA)

Reservation (cont'd)

- D-TDMA works as follows:
  1. The stations take turns transmitting a single frame and the transmissions from the stations are organized into cycles that can be variable in length
  2. Each cycle begins with **the reservation interval**
  3. In the simplest case, the reservation interval consists of N minislots, 1 minislot per station
  4. The stations announce their intention to transmit a frame by broadcasting their reservation bit during the appropriate minislot
  5. By listening to the reservation interval, the stations can determine the order of frame transmissions
  6. The length of the cycle will then correspond to the number of stations that have a frame to transmit
  7. Note that variable-length frames can be handled if the reservation message includes frame-length information
  8. Usually, reservations take effect in the next following cycle

Reservation (cont'd)

- Reservation intervals can be viewed as a form of overhead
- If the number of stations becomes very large, this overhead can become very significant
  - This situation becomes a serious problem when a very large number of stations transmit frames infrequently
  - The reservation minislots appear in every cycle, even though most stations do not transmit data
- The problem can be addressed by not allocating a minislot to each station and instead making **stations contend for a reservation minislot** by using a random access technique such as ALOHA or slotted ALOHA

Outline

Random Access

- Taxonomy of MAC techniques

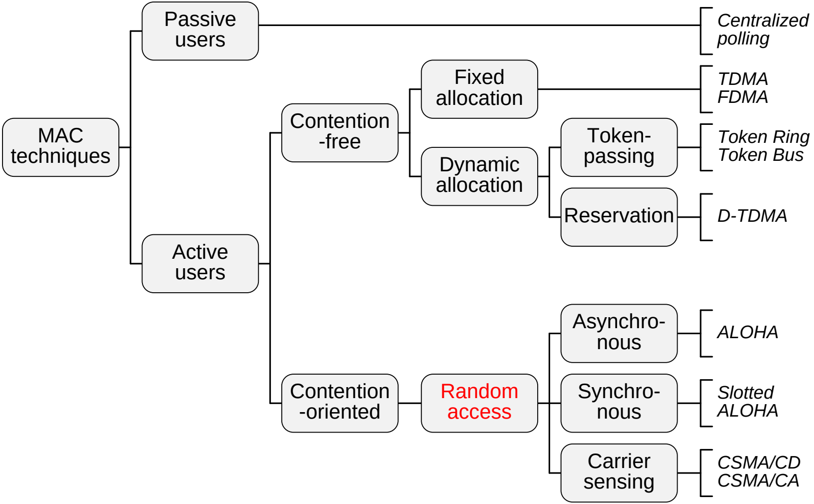

Random Access (cont'd)

- **2 types of random access**:
  - Noncarrier-sensing
  - Carrier-sensing

- **Noncarrier-sensing random access** (aka **'impolite'**)
  - A station transmits frames across the transmission medium at a random time, without coordinating its access with other stations
  - E.g., ALOHA, Slotted ALOHA

- **Carrier-sensing random access** (aka **'listen-before-talk'**)
  - A station first senses the carrier to gain certain channel state information and then uses this information to schedule its frames for transmission, again without coordinating its access with other stations
  - CSMA, CSMA/CD, CSMA/CA

ALOHA

- **ALOHA** (aka **Pure ALOHA**) is the first and the simplest random access technique
  - It derives its name from **ALOHAnet**, a communication network developed at the University of Hawaii to enable wireless communication among the campuses located at different Hawaiian islands and first put into operation in 1971
- ALOHA works as follows:
  1. Stations do not monitor the status of the shared transmission medium
  2. Whenever a station has a frame to send, it is transmitted right away
  3. If the transmitting station is the only transmitter in the medium, then the frame transmission is successful
  4. A transmitting station knows that its data transmission is successful by the reception of an ACK transmitted in response to the data frame
  5. When a station has determined that its transmitted frame has collided, it then schedules this frame for retransmission after a random retransmission delay

ALOHA (cont'd)

- Let us suppose that there are 3 stations: A, B, and C
  - Frame A1 transmitted by station A and frame B1 transmitted by station B are successfully received by station C
  - Due to the random nature of the medium access, data collisions are unavoidable, like the collision of frames A2 and B2
  - Since neither station A nor station B receives an ACK from station C, they know that their frame transmissions are not successful and they retransmit after waiting a random amount of time to avoid successive collisions

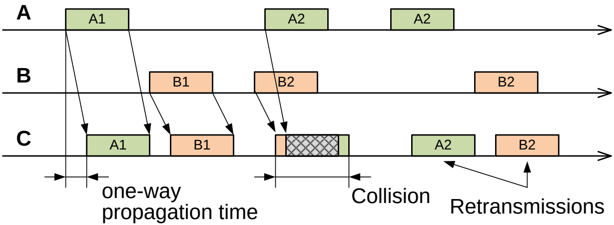

ALOHA (cont'd)

- When traffic load is very light, the probability of collision is very small, and so retransmissions need to be carried infrequently

- There is a big difference between transmission errors caused by transmission impairments and those due to collisions:
  - Transmission errors due to transmission impairments affect only a single station at a time
  - However, in any collision more than 1 station is involved, and, therefore, more than 1 retransmission is necessary
  - These retransmissions can trigger additional collisions

- To prevent a snowball effect with multiple collisions, stations use **a backoff algorithm**, which chooses a random number in a certain retransmission time interval
  - This **randomization** is intended to spread out the retransmissions and reduce the likelihood of additional collisions

ALOHA (cont'd)

- Frame transmission using ALOHA:
  - The first transmission is done without any scheduling
  - Information about the outcome of the transmission is obtained at the earliest after $2T_{\text{pr}}$ seconds, where T_{pr} is the one-way propagation time between the communicating stations

$$\boxed{\text{Propagation time} = \frac{\text{Distance (in meters)}}{\text{Signal velocity (in m/s)}}}$$

- If no ACK is received after $2T_{\text{pr}}$ seconds, a backoff algorithm is used to select a random retransmission time

- Assuming the use of fixed-size frames, the transmission time T_{tr} can be found as

$$\boxed{\text{Transmission time} = \frac{\text{Frame size (in bits)}}{\text{Data rate (in bits/s)}}}$$

ALOHA (cont'd)

- The transmission time is due to a finite data rate of the channel
- The propagation time is due to a finite velocity of signal propagation

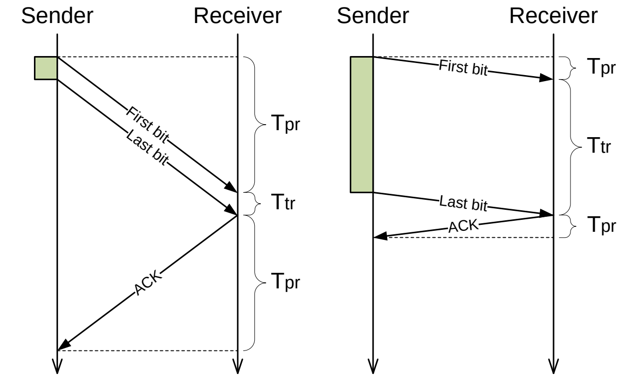

Transmission time < propagation time (high data rate and/or long distance)

Transmission time > propagation time (small data rate and/or short distance)

ALOHA (cont'd)

- Consider a reference frame that is transmitted starting at time T_0 and completed at time $(T_0 + T_{tr})$
  - It will be successfully transmitted if no other frame collides with it
  - Any frame that begins its transmission in the interval T_0 to $(T_0 + T_{tr})$ will collide with the reference frame
  - Moreover, any frame that begins its transmission in the prior T_{tr} seconds will also collide with the reference frame
  - Thus, the probability of a successful transmission is the probability that there are no additional frame transmissions in **the vulnerable period** $(T_0 - T_{tr})$ to $(T_0 + T_{tr})$

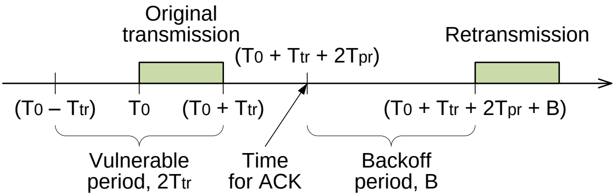

ALOHA (cont'd)

- **The vulnerable period for a frame transmission** is a period, where any transmission attempt by any other station will result in a collision
- **The vulnerable period in ALOHA** $= 2T_{\text{tr}}$
  - The vulnerable period is a factor that limits the maximum throughput achievable by ALOHA
- **The throughput of a shared-medium LAN** = the fraction of time the channel carries useful information, namely noncolliding frames
- **The maximum throughput of ALOHA** $= 18.4\%$
  - This corresponds to the frame arrival rate (aka **offered load**) of exactly 1 frame per vulnerable period (i.e., 1 frame per $2T_{\text{tr}}$ seconds)
  - Thus, 81.6% of the total available bandwidth is wasted due to losses from collisions

Slotted ALOHA

- The performance of ALOHA can be improved by reducing the probability of collisions
- **Slotted ALOHA** (aka **S-ALOHA**) reduces collisions by constraining the stations to transmit in a synchronized fashion
  - I.e., time is divided into time slots and stations can start their frame transmissions only in the beginning of each slot
- Frames are fixed-size and occupy exactly 1 slot

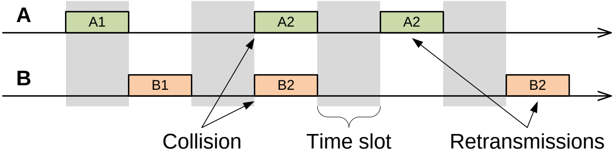

Slotted ALOHA (cont'd)

- **A centralized clock** sends out small clock tick frames to stations
  - Stations are only allowed to send their frames immediately after receiving a clock tick
  - If there is only 1 station with a frame to send, the frame transmission is successful
  - On the other hand, if there are 2 or more stations with frames to send, there is a collision and the whole slot up to the next clock tick is wasted
- **Slotted ALOHA reduces the vulnerable period to 1 slot**

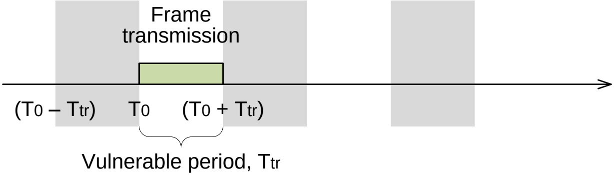

Slotted ALOHA (cont'd)

- **The vulnerable period in Slotted ALOHA** $= T_{\text{tr}}$
  - As opposed to $2T_{\text{tr}}$ in ALOHA
  - However, Slotted ALOHA requires network-wide synchronization, which is difficult to achieve in practice
- **The maximum throughput of Slotted ALOHA** = is 36.8%
  - This is because the probability of collision is reduced by a factor of 2
- Both ALOHA and Slotted ALOHA exhibit **a bistable system behavior**
  - Under light traffic load, the system state resides at the first region, yielding an acceptable delay-throughput performance
  - Under heavy traffic load, the system can experience a high level of collisions, which in turn lead to further retransmissions and collisions, causing the system to produce very few successful transmissions

Slotted ALOHA (cont'd)

- Stability is an important problem in ALOHA and Slotted ALOHA, which may degrade the system performance significantly
  - If the offered load exceeds the optimal operating point, throughput starts to decrease, eventually reaching 0 due to excessive collisions

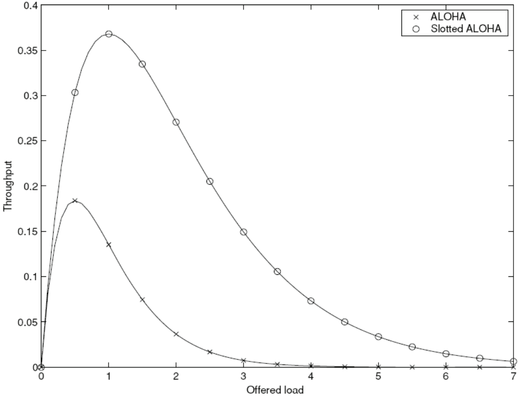

Slotted ALOHA (cont'd)

- **Benefits** of ALOHA and Slotted ALOHA:
  - They are efficient in light-load environments with varying numbers of bursty stations
  - They enable the use of relatively simple access logic at each station
  - ALOHA has no single point of overall failure
- **Shortcomings** of ALOHA and Slotted ALOHA:
  - Inefficient resource utilization due to collisions, which limits the maximum achievable throughput to low values
  - Performance tends to collapse under heavy traffic load
  - In Slotted ALOHA, the failure of the centralized clock will disrupt the entire network

CSMA

- Both ALOHA and Slotted ALOHA exhibit poor performance which can be attributed to the **'impolite'** behavior of stations
- Consider more polite, the **'listen-before-talk'** approach
  - For $T_{\text{pr}} < T_{\text{tr}}$, when a station launches a frame, all the other stations know it almost immediately
  - So, if the other stations would not try transmitting until the the first station was done, then collisions would occur only when 2 or more stations began to transmit almost simultaneously
  - The process of listening to the medium is not demanding because every station is equipped with a receiver anyway
  - Moreover, to detect another station's transmission does not require receiving the transmitted data – it is enough just to sense the carrier that is present when signals are transmitted
- **Carrier-sensing MAC** is characterized by sensing the carrier and deciding according to it whether another transmission is ongoing

CSMA (cont'd)

- **Carrier Sense Multiple Access (CSMA)** works as follows:
  1. A station wishing to transmit first listens to the medium to determine if another transmission is in progress
  2. If the medium is in use, the station must wait
  3. Otherwise, the station may transmit
  4. If 2 or more stations will attempt to transmit at about the same time, there will be a collision
  5. To account for this, a station waits a reasonable amount of time after transmitting for an ACK, taking into account the **round-trip time (RTT)** and the fact that the acknowledging station must also contend for the medium to respond
  6. If there is no ACK, the station assumes that a collision has occurred and retransmits the frame

CSMA (cont'd)

- CSMA is effective only for networks in which $T_{pr} < T_{tr}$
  - Collisions can occur only if 2 or more stations will attempt to transmit within the period of propagation time (T_{pr})
  - If a station begins to transmit a frame and there are no collisions during the time it takes for the leading edge of the frame to propagate to the most remote station, then there will be no collision for this frame because all other stations are now aware of the transmission
- **The vulnerable period in CSMA** $= T_{pr}$

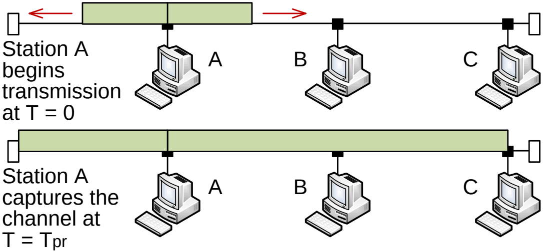

CSMA (cont'd)

- CSMA techniques differ according to the behavior of stations that have a frame to transmit **when the transmission medium is busy**
- **1-Persistent CSMA** works as follows:
  1. When a station has data to send, it senses the medium
  2. If the medium is busy, the station senses it continuously, waiting until the medium becomes idle
  3. As soon as the medium is sensed idle, the station transmits its frame
  4. If more than one station is waiting, a collision will occur
  5. After sending the frame, the station waits for an ACK, and if none is received in a specified amount of time, the station will wait a random amount of time and then resume listening to the medium
  6. When the medium is again sensed to be idle, the frame is retransmitted immediately
- In 1-Persistent CSMA, stations attempt to access the medium as soon as possible, acting in a greedy fashion
  - As a result, 1-Persistent CSMA has a relatively high collision rate

CSMA (cont'd)

- **Nonpersistent CSMA** works as follows:
  1. When a station has data to send, it senses the medium
  2. If the medium is busy, the station does not sense it continuously
  3. Instead, after sensing the busy condition, it waits a randomly selected interval of time before sensing again
  4. As with 1-Persistent CSMA, the station begins transmitting immediately when the medium is sensed to be idle
- In Nonpersistent CSMA, the randomized waiting times between medium sensings eliminate most of the collisions that would result from multiple stations transmitting simultaneously upon sensing the transition from busy to idle condition
  - By immediate rescheduling a resensing time and not persisting, the collision rate is reduced relative to 1-Persistent CSMA
  - This leads to throughput values much higher than 1-Persistent CSMA at high traffic loads
  - However, waiting times between medium sensings result in longer access delays than in 1-Persistent CSMA at light traffic loads

CSMA (cont'd)

- **p-Persistent CSMA** works as follows:
  1. When a station has data to send, it senses the medium
  2. If the medium is sensed idle, the station transmits with probability p
  3. With probability $(1 - p)$, the station decides to wait some time before again sensing the medium
  4. If the medium is still idle, the station transmits with probability p or defers again with probability $(1 - p)$
  5. This procedure is repeated until either the frame has been transmitted or the medium is sensed to be busy
  6. When the medium is busy, the station then senses it continuously
  7. When the medium is again sensed to be idle, the station restarts this procedure
- Thus, p-Persistent CSMA is a combination of 1-Persistent CSMA and Nonpersistent CSMA
  - 1-Persistent CSMA is so called due to its strategy, which is to transmit with probability $p = 1$ as soon as the medium is available

CSMA (cont'd)

- The basic idea behind p-Persistent CSMA is to spread out the transmission attempts by the stations that have been waiting for a transmission to be completed
- For small values of p, p-Persistent CSMA has a much smaller collision rate than 1-Persistent CSMA or Nonpersistent CSMA
- But this is achieved at the expense of much longer access delays than in 1-Persistent CSMA or Nonpersistent CSMA
- 1-Persistent CSMA, Nonpersistent CSMA, and p-Persistent CSMA also have **slotted forms**, in which stations are synchronized and all transmissions, whether initial transmissions or retransmissions, are synchronized to the time slots

CSMA (cont'd)

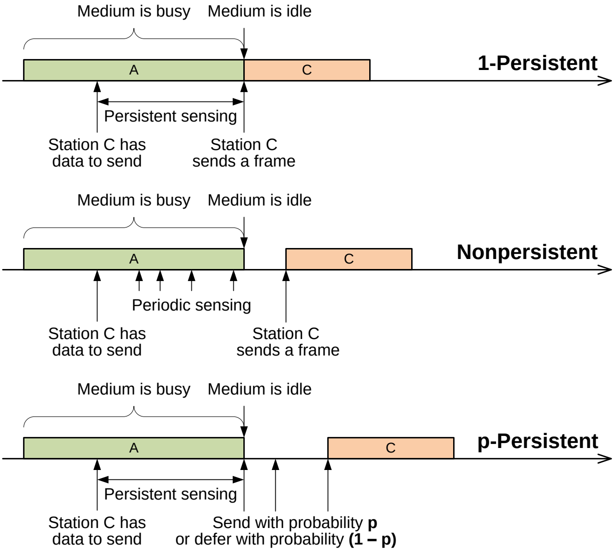

Medium is busy
Medium is idle
A
C
1-Persistent
Persistent sensing
Station C has data to send
Station C sends a frame
Medium is busy
Medium is idle
A
C
Nonpersistent
Periodic sensing
Station C has data to send
Station C sends a frame
Medium is busy
Medium is idle
A
C
p-Persistent
Persistent sensing
Station C has data to send
Send with probability **p**
or defer with probability **(1 – p)**

CSMA (cont'd)

- Throughput vs. offered load (packet time = T_{tr}; $T_{pr} < T_{tr}$)
  - CSMA outperforms ALOHA and Slotted ALOHA when $T_{pr} < T_{tr}$
  - However, for longer propagation delays CSMA becomes relatively inefficient due to the inability of the stations to accurately sense the busy medium

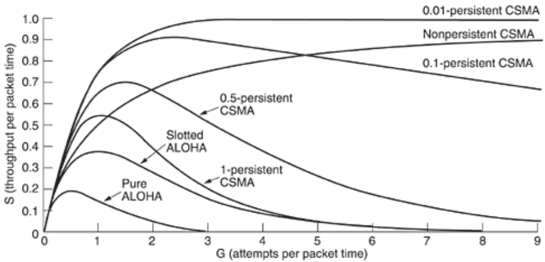

CSMA/CD

- In both ALOHA and CSMA, when frames collide, the medium remains unusable for the duration of transmission of the damaged frames
  - For long transmissions, compared to T_{pr}, the amount of wasted bandwidth can be considerable
- But if a station can determine whether a collision is taking place, then the amount of wasted bandwidth can be reduced by **aborting the transmission when a collision is detected**
- **Carrier Sense Multiple Access with Collision Detection (CSMA/CD)** uses this approach
  - I.e., a station continues to listen to the medium while transmitting
  - Sometimes referred to as the **'listen-while-talk'** approach

CSMA/CD (cont'd)

- CSMA/CD can be used in 1-Persistent, Nonpersistent, or p-Persistent variations of CSMA, each with a slotted or unslotted version
- **Unslotted 1-Persistent CSMA/CD** works as follows:
  1. When a station has data to send, it senses the medium
  2. If the medium is idle, the station transmits its frame
  3. Otherwise, the station continues to listen until the medium is idle, then transmits immediately
  4. If a collision detected during transmission, the station transmits a brief **jam signal** to assure that all stations know that there has been a collision and then stops its transmission
  5. After transmitting a jam signal, the station waits a random amount of time, then attempts to transmit again

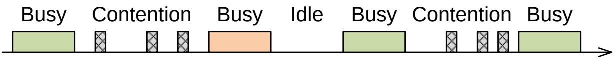

CSMA/CD (cont'd)

- With CSMA/CD, the maximum time it takes to detect a collision is about $2T_{\text{pr}}$ seconds
  - Also referred to as the **'worst-case scenario'**
- Thus, a station cannot be sure that it has captured the channel until it has transmitted for $2T_{\text{pr}}$ seconds without hearing a collision
- As a result, an important rule followed in most CSMA/CD implementations is that **frames should be long enough to allow collision detection prior to the end of transmission**
  - I.e., $2T_{\text{pr}} < T_{\text{tr}}$
- If shorter frames are used, then collision detection does not occur
  - In that case, CSMA/CD exhibits the same performance as the less efficient CSMA

CSMA/CD (cont'd)

- Let us compare the maximum collision durations of CSMA and CSMA/CD in the worst-case scenario
  - **The duration of a collision in the network** = the time from the beginning of the first transmission in the collision until the earliest time at which a fresh transmission can begin

$$\frac{\text{Maximum collision duration of CSMA}}{\text{Maximum collision duration of CSMA/CD}} = \frac{2T_{\text{pr}} + T_{\text{tr}}}{3T_{\text{pr}}}$$

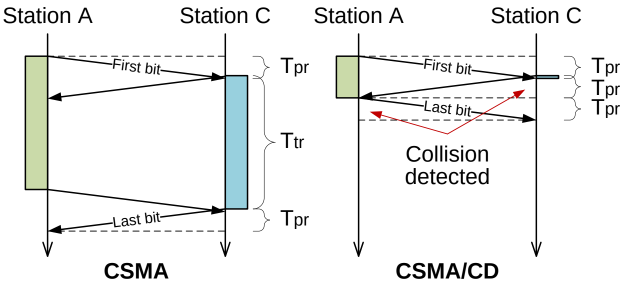

CSMA/CD (cont'd)

- Cable length = 500 m
- Signal velocity $= 0.7 * 3 * 10^8 = 2.1 * 10^8$ m/s
- Frame size = 1500 bytes $= 12 * 10^3$ bits
- Data rate = 10 Mbit/s $= 10 * 10^6$ bit/s

$$T_{\text{pr}} = \frac{\text{Cable length}}{\text{Signal velocity}} = \frac{500}{2.1 * 10^8} \approx 2.4 * 10^{-6} \text{ s}$$

$$T_{\text{tr}} = \frac{\text{Frame size}}{\text{Data rate}} = \frac{12 * 10^3}{10 * 10^6} = 1.2 * 10^{-3} \text{ s}$$

- With CSMA, the medium will be released after approximately $1.2 * 10^{-3}$ s
- With CSMA/CD, it will be released after approximately $7.2 * 10^{-6}$ s

$$\frac{2T_{\text{pr}} + T_{\text{tr}}}{3T_{\text{pr}}} \approx 167$$

CSMA/CD (cont'd)

- **Collisions are absolutely normal events** and are simply an indication that the CSMA/CD algorithm is functioning as designed
- Collisions resolve quickly
  - E.g., the design of the CSMA/CD algorithm ensures that the majority of collisions on a 10 Mbit/s Ethernet LAN will be resolved in microseconds
- A normal collision does not result in lost data
  - In the event of a collision, the station backs off (waits) for some time, and then tries to retransmit the frame
- As more stations are added to a given network, there will be more traffic, resulting in more collisions
  - Repeated collisions for a given frame indicate that the network is congested

CSMA/CD (cont'd)

- After transmitting the jam signal, the station waits a random period of time before attempting to transmit again, a process known as backoff
- **The backoff process** is the process by which a transmitting station determines how long to wait after a collision before attempting to retransmit
- If all stations waited the same amount of time, secondary collisions would inevitably occur
- Stations avoid this by generating a random number that determines the time they must wait before restarting the CSMA/CD process
- This period is called **the backoff delay**
- The algorithm used to manage this process in Ethernet is known as **truncated binary exponential backoff (TBEB)**

CSMA/CD (cont'd)

- When a collision occurs, each station generates a random number r that always falls within a specific range of values
- The random number determines the number of slot times it must wait before listening to determine whether the medium is idle
- **A slot time** represents the minimum period of time that a station needs access to the medium to send the smallest legal frame size (thus, 64 bytes on 10 Mbit/s Ethernet results in 51.2 microseconds)
- The number of slot times to delay before the n-th attempt is chosen as a uniformly distributed random integer r in the range

$$0 \leq r \leq 2^k - 1, \qquad k = \min(n, 10)$$

CSMA/CD (cont'd)

- TBEB works as follows:
  1. In the first attempt, the range of slot times the station waits before listening to determine whether the medium is idle lies between 0 and $2^1 - 1 = 1$
  2. On the second attempt, it is between 0 and $2^2 - 1 = 3$
  3. For the third, it falls between 0 and $2^3 - 1 = 7$
  4. The process continues, expanding each time to a larger range
  5. If multiple collisions occur, the range expands until it reaches 10 attempts, at which point the range is between 0 and 1023
  6. From that point on, the range of values remains fixed between 0 and 1023
  7. Thus, the term **'truncated'** simply means that after a certain number of increases, the exponentiation stops
  8. If a station fails to transmit following 16 attempts, the MAC protocol issues **an excessive collision error**
  9. The frame being transmitted is dropped, requiring that the application software reinitiate the transmission process

CSMA/CD (cont'd)

- CSMA/CD collision backoff ranges for 10 Mbit/s Ethernet LANs
  - Note that it is not required always to pick a larger integer and develop a larger backoff time after each collision on a given frame
  - Instead, the range of integers to choose from just gets larger and the probability that the station will generate a longer backoff increases

Attempt	**Range**	**Max number**	**Max backoff time**
1st	0 – 1	$2^1 - 1$	51.2 μs
2nd	0 – 3	$2^2 - 1$	153.6 μs
3rd	0 – 7	$2^3 - 1$	358.4 μs
4th	0 – 15	$2^4 - 1$	768.0 μs
5th	0 – 31	$2^5 - 1$	1587.2 μs
6th	0 – 63	$2^6 - 1$	3225.6 μs
7th	0 – 127	$2^7 - 1$	6502.4 μs
8th	0 – 255	$2^8 - 1$	13056.0 μs
9th	0 – 511	$2^9 - 1$	26163.2 μs
10th – 15th	0 – 1023	$2^{10} - 1$	52377.6 μs

CSMA/CA

- CSMA/CD is widely used in wired LANs
  - E.g., IEEE 802.3 Ethernet
- However, CSMA/CD is difficult to implement in wireless LANs:
- First, the radios used in the wireless LANs are half-duplex and, therefore, collision detection is not possible
  - Implementing the collision detection scheme would require the implementation of a full-duplex radio, capable of transmitting and receiving at once, an approach that would increase the price significantly
- Second, on a wireless environment we cannot assume that all stations hear each other (which is the basic assumption of the collision detection scheme)
  - I.e., the fact that a station willing to transmit senses the medium idle, does not necessarily mean that the medium is idle around the receiver

CSMA/CA (cont'd)

- In order to overcome these problems, wireless LANs use **Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)** together with positive acknowledgements
- The idea behind CSMA/CA is to prevent collisions at the moment they are most likely to occur:
  - Carrier sensing involves monitoring the medium to determine whether the medium is idle or not
  - If the medium is busy, the station should wait until the medium becomes idle
  - But other stations may have also been waiting for the medium to become idle
  - So if the station transmits immediately after the medium becomes idle, then collisions are likely to occur, and because collision detection is not possible, the channel will be wasted for an entire frame duration
  - A solution to this problem is to randomize the times at which the contending stations attempt to capture the medium

CSMA/CA (cont'd)

- Thus, to minimize collisions, wireless LANs use CSMA/CA
  - E.g., IEEE 802.11 Wi-Fi
- Collision avoidance is implemented by mandating that, when the channel is sensed idle, the station has to wait for a randomly chosen duration before attempting to transmit
- This mechanism drastically decreases the probability that more than 1 station attempts to transmit at the same time
- However, CSMA/CA minimizes, but does not eliminate, the potential for collisions
- In these cases, transmissions are corrupted and the corresponding stations retry later on

CSMA/CA (cont'd)

- Distributed InterFrame Space (DIFS); Short InterFrame Space (SIFS)