ECE 333: Introduction to Communication Networks
Fall 2002

Lecture 17: Medium Access Control V

- Ethernet
- Contention approaches for wireless networks

In the past several lectures we have considered ***contention-based*** approaches for medium access control (these are also called ***random access*** approaches). The basic approach here is Aloha; we also looked at CSMA and CSMA/CD - these approaches take into account more knowledge about the other users than was assumed in Aloha. For both CSMA and CSMA/CD a key performance parameter was β where this is defined as

$$\beta = \frac{\textit{max. propagation \& detection delay}}{\textit{time to send 1 packet}}$$

Better performance is achieved for smaller values of β.

We also observed that "stability" is a concern in these protocols; this is more of a concern in Aloha, but still arises in CSMA and CSMA/CD.

Today we will look at the Ethernet protocol. Ethernet is the most widely used LAN technology; the original version of Ethernet was based on a CSMA/CD protocol.

We also will look at contention-based protocols for wireless LAN's - wireless channels introduce some new difficulties, that preclude the direct application of protocols like CSMA/CD. We will look at some basic protocols for this setting and briefly look at IEEE 802.11 (Wi-Fi) wireless LANS.

Ethernet & IEEE 802.3

Two researchers at the Xerox Palo Alto Research Center, Bob Metcalfe[1] and David Boggs, invented the Ethernet protocol in the 1970's. Researchers at Xerox implemented the protocol, where it was used for research and internal networking. These early versions ran at approximately 3 Mbps. Xerox formed an alliance with Intel and Digital (DEC) to establish an Ethernet standard in 1978 (referred to as the DIX Ethernet standard), but little interest was shown at Xerox to commercialize the product. In 1979, Metcalfe left Xerox to found 3Com, with a business case built on Ethernet.

In 1982, the IEEE 802 group published the 802.3 standard; this standard was based on and very similar to the DIX standard. The IEEE 802 group has produced a variety of standards for LANs - many of these have been adopted by other standard agencies including ANSI, NIST, and ISO. All Ethernet products since the mid-1980's have been based on the IEEE 802.3 standards.

IEEE 802.3 standard describes the MAC sub-layer as well as the physical layer for an Ethernet LAN. The IEEE 802.2 standard described in Lecture 10 is used for the Logical Link Control sub-layer in these LANs.

10base5

[1] Before starting at Xerox and after graduating with a PhD from Harvard, Metcalfe spent time at the University of Hawaii studying the Aloha protocol.

The original 802.3 standard specified the use of a shared bus with thick coaxial cable as the physical medium. Since then many alternative physical configurations for Ethernet have been standardized. This original configuration is referred to as **10base5.** This notation means that the data rate is 10 Mbps, base-band signaling is used (with Manchester encoding), and the maximum cable segment is 500 meters.

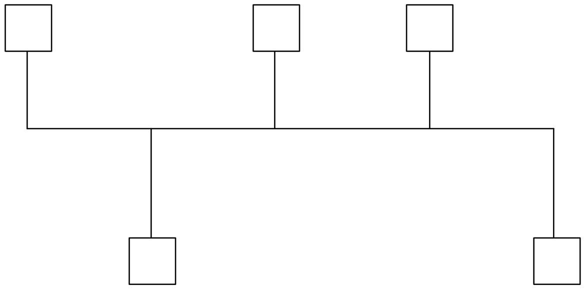

Medium Access

For medium access, 802.3 uses ***1-persistent, non-slotted CSMA/CD***, this operates as follows:

1) Stations listen to hear if channel is in use (carrier sense)

2a) If channel is idle, station transmits.

2b) If channel is busy, station waits for channel to become idle, then transmits (immediately).

3) When collisions occur, stations stop transmitting and retransmissions are controlled by ***binary exponential back-off procedure***.

Exponential Binary Back-off

The back-off and retransmission procedure used in Ethernet is slightly different from what we have discussed before. After a collision each station retransmits after a random number of time slots, where the length of a time slot is approximately 2β. The number of time slots is chosen as follows:

After 1st collision, stations randomly transmit after either 0 or 1 time-slots with equal probability (if no carrier is sensed).

After 2nd collision, stations randomly transmit after either 0,1,2 or 3 time-slots, with equal probability.

After i^{th} collision, stations wait a random number of slots between 0 and 2^i-1.

In the standard, i stops increasing when it gets to 10. After 16 retries the transmission is considered to be a failure.

Minimum Message Size

In 802.3, if a station does not detect a collision before it finishes transmitting a packet, it assumes that the packet was correctly received. Recall, in Lecture 16, we argued that it could take up to 2β seconds to detect a collision from when a station begins transmitting. Thus the transmission time of a packet must be longer than 2β, i.e. $2\beta < 1$, or else the transmitter may assume that the packet arrived correctly when it actually did not. (Also, as shown in Leon-Garcia, as β increases the average delay for a given load increases which is also undesirable.) Because of these concerns, Ethernet standards limit the minimum packet size and the maximum propagation delay. In the original 802.3 standard, up to 5 segments of 500 meter cable can be used (segments are connected together by repeaters). In this case,

$$\text{Maximum propagation delay } = \frac{2500\text{m}}{2.31\times 10^{8}\text{ m/sec}} = 10.82\mu\text{sec}$$

(The 802.3 standard assumes a propagation speed of 0.77 * c = 2.31 x 10^8 m/s)

So, the maximum round trip propagation delay = 21.65 μsec

Additionally there are several other sources of delay including the delay incurred in each repeater and other processing delays, Adding all of these delays gives a worst-case round-trip delay of approximately 49.47 μsec.

At 10 Mbps with a frame of size F bits, we have:

$$2\beta = \frac{49.47\times 10^{-6}}{F/10\times 10^{6}}$$

Thus the requirement that $2\beta < 1$, means that we must have $F > 494.7$ bits.

Based on these considerations, the minimum message size for 10base5 is set at 512 bits, or 64 bytes.

The resulting time to send a minimum size message, 51.2 μsec, is the slot-time used in the back-off algorithm described above.

Minimum Message Size

Note that given a minimum message size, this places an **absolute** limit on the maximum length of an 802.3 network.

The original standard allowed for segments up to 500 meters and up to 4 repeaters giving the maximum length of 2.5 km. (This was reduced to 2 km in later versions of the standard)

What happens in this analysis if data rate increases to 100 Mbps?

⇒ Either the LAN must be 10 times shorter <u>or</u> minimum message size must be 10 times larger.

802.3 Frame Format

The 802.3 frame format is shown below. The number of bytes in each field is indicated.

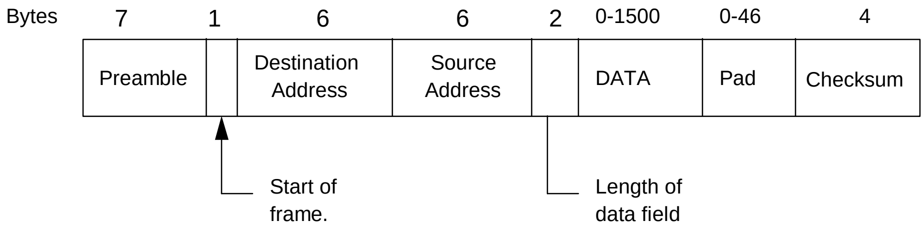

Bytes	7	1	6	6	2	0-1500	0-46	4
	Preamble		Destination Address	Source Address		DATA	Pad	Checksum

Start of frame.

Length of data field

The **preamble** contains a 10101010 pattern; this is used for the receiver to synchronize to the transmitted signal. The **Start of Frame** is the pattern 10101011. The Preamble and Start of Frame are not counted in minimum frame length. The **Length field** contains the length of the data field. The data field contains a packet from the LLC sub-layer. The **Pad field** is used to generate the minimum 64 byte message size if the data field is too short. The **Checksum** is a 32 bit Cycle Redundancy Code (CRC), used for error detection.

Addresses

The Ethernet frame contains fields for both the senders and receiver's addresses. Standard Ethernet addresses are 48 bits long (6 bytes), and several special addresses are reserved, for example the all 1's address is used to indicate a message that is broadcast to all stations on the network.

Ethernet addresses are assigned and hard-coded into 802.3 Network Interface Cards (NIC) when they are manufactured. Thus every NIC card must have a globally unique Ethernet address. These addresses are administered by IEEE to ensure that they are all unique. Ethernet addresses are often written in hexadecimal notation, with each byte written as a pair of hexadecimal numbers. For example, the broadcast address mentioned above is written as ff:ff:ff:ff:ff:ff.

802.3 Extensions

Since its introduction, the 802.3 standard has been expanded to allow for different kinds of physical media, with different restrictions. Some examples are given in the table below

10Base5	Thick Coax	500 meters/segment – the original standard
10Base2	Thin Coax	200 meters/segment – “thin ethernet”
10Base-T	Twisted Pair	100 meters in a star configuration with a hub
10Base-FP	Optical fiber	1km segments/ start topology.

10Base-T:

The 10Base-T configuration is one of the most widely used for connecting desktop LANs[2]. In this configuration each station is connected to a central device called a hub using 2 (category 5) twisted pairs; both pairs are twisted together in one cable. Each station receives on one pair and transmits on another. The hub takes frames received from one station and forwards it to all other stations. The hub blindly forwards received packets; if more than one packet arrive at the hub to be forwarded, then a collision occurs, and the hub broadcasts a signal indicating this. Note that while this network physically has a star topology, from the point-of-view of the stations, it is no different than if they were connected to a shared bus. This is sometimes described by saying that the **logical topology** of the network is a shared bus.

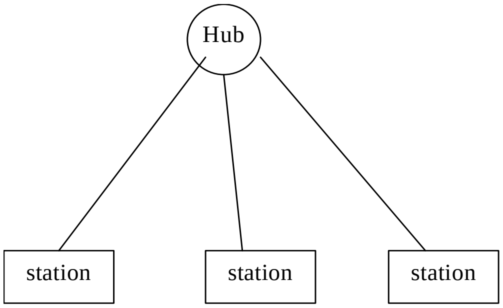

[2] The primary reason for this is that the 10Base-T configuration uses the same wiring that is used for telephone wiring in most office buildings; many buildings have extra telephone wiring already installed.

Fast and Faster Ethernet

Over time, Ethernet has evolved to provide higher transmission rates. Several methods have been used to accomplish this.

The first method for improving the transmission rate of Ethernet is by replacing the hub with a ***Ethernet switch***. A switch can buffer incoming frames and avoid collisions. In this case, it is also possible to allow nodes to use full-duplex communication (in traditional Ethernet, half-duplex communication is used). With such a configuration, collisions do not occur, and the CSMA/CD algorithm is no longer needed. In some cases, a single input line to a switch is shared by several stations using a hub; these stations are said to be in the same ***collision domain***.

In 1995, a 100 Mbps 802.3 standard was approved for several different configurations including twisted pair/hub (100baseT) and a fiber optic hub (100baseFX). This standard is also called ***fast Ethernet.*** This standard keeps exactly the same frame sizes and formats as in 10 Mbps Ethernet[3]. Only hub configurations where used and the maximum distance from a user to a hub is 100 meters (in the twisted pair versions).

[3] This ensures compatibility between 10Mbps and 100Mbps Ethernet standards. Indeed, Ethernet cards are available that can operate over either type network.

In 1998, ***gigabit Ethernet*** (1000base-*x*) was introduced for several different transmission media. Again the basic frame structure and procedures of 10Mbps Ethernet were retained. However now, to ensure that the CSMA/CD algorithm works properly, the minimum packet size was extended to 512 bytes and several other modifications where made. Most implementations of Gigabit Ethernet use a switched configuration, so the CSMA/CD protocol is not used.

Currently a 10 Gigabit Ethernet standard is being developed; it is not yet approved, but some products have already been announced.

Contention protocols for wireless networks

Contention-based protocols are also used for wireless networks; indeed Aloha was originally developed in a wireless setting.

Wireless networks can be classified as having either a ***peer-to-peer***, ***centralized*** or ***hierarchical*** architecture.

- In a peer-to-peer architecture, wireless nodes transmit directly to each other (and possibly multi-hop). This is sometimes referred to as an ***ad hoc wireless network.***

- In a centralized architecture, all communication goes through a centralized hub node; one example of this is a satellite network.

- In a hierarchical architecture nodes are organized in a hierarchy; lower nodes in the hierarchy only communicate with higher nodes, e.g. a cellular network.

In the following we look at some issues that arise when trying to use a contention-based medium access protocol in a wireless setting, for example a wireless LAN.

In the following, we ignore issues of routing and location management and assume that these are addressed at a different layer.

Several properties of wireless channels can add a new twist to the design of a MAC protocol.

1st issue - in a wireless setting nodes cannot transmit and receive at the same time and in the same frequency band. Thus, the only way for Collision Detection to occur is for some other node to broadcast this on a separate channel.

This can work in centralized architecture but is difficult in a peer-to-peer setting.

2nd issue - Not all nodes may be in range of each other. Thus a node may only be able to detect if a subset of the other nodes are transmitting. This can result in problems for a CSMA approach.

Hidden terminal problem **-** Suppose A is transmitting to B and C wants to transmit to D. Then C may not hear A, resulting in a collision.

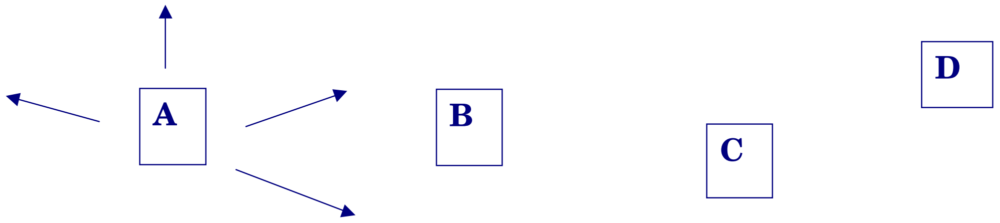

Exposed terminal problem **-** Assume B is transmitting to A and C wants to transmit to D. Here C hears B, but could transmit anyway because this would not interfere at A.

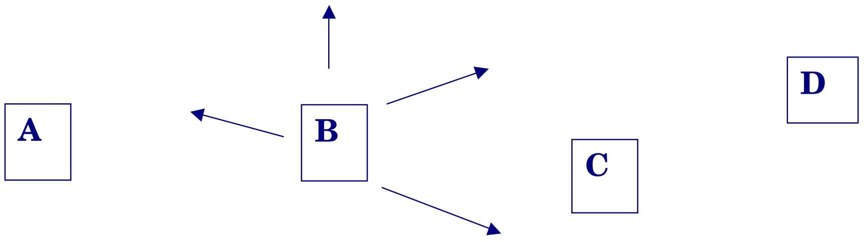

Basic Problem- when transmitting a node really cares about the activity around the receiver not itself, but it can only detect the activity around itself.

One solution- when a node is receiving have it broadcast a ***busy signal*** on a control channel, so other nodes in range of receiver can detect this.

(Common assumption is that all nodes have the same transmission range)

A variation of this idea is used in a protocol called ***Multiple Access with Collision Avoidance (MACA)***.

MACA protocol:

1. Sending node sends **Request to Send (RTS)** packet to destination containing the length of packet to be sent.
2. Receiving node replies to a successfully received RTS with a **Clear to Send (CTS)** packet, also containing the length of the packet to be sent.
3. When the sending node receives the CTS, it begins transmitting the packet.
4. Any other node that hears a CTS packet must not transmit for the duration of the packet.

Note: Multiple nodes can transmit RTS at same time resulting in collisions. When collisions occur, nodes retransmit their RTS packets using Aloha with binary exponential back-off.

Can other collisions occur? (Yes, for example the CTS, may not be heard by one node due to other traffic around that node.)

MACAW (MACA for wireless LANs):

MACAW is a modification of MACA protocol including the following changes:

1) Receiver sends an ACK to sender after successful reception. Other nodes cannot transmit until they hear ACK.

2) Nodes use CSMA when sending RTS packets.

This results in protocol with improved performance.

Issues -

In most MAC protocols, ACKs and retransmission are done at LLC sub-layer. What are advantages of doing it at the MAC layer?

If node only hears RTS but not CTS it is an ***exposed terminal*** and should be allowed to transmit - but if it is transmitting then the node that sent the RTS will not be able to hear the ACK for it packet, and may retransmit.

This can be avoided if nodes are not allowed to transmit if they hear either a RTS or CTS. However this results in lower throughput.

IEEE 802.11 Wireless LAN Standard

The IEEE 802.11 is a wireless LAN standard, first standardized in 1997. IEEE 802.11 networks are also called Wi-Fi networks, and in some cases "wireless ethernet." The original 802.11 standard supports rates of 1 and 2 Mbps using one of 3 different physical layer transmission techniques:

- Infrared between 850 and 950 nm.
- Direct-sequence spread spectrum in the 2.4 GHz band.
- Frequency-hopping spread spectrum in the 2.4 GHz band.

This has been extended to 54 Mbps in the 5 GHz band (802.11a), and to 5.5 and 11 Mbps in the 2.4 GHz band (802.11b). The 2.4GHz and 5 GHz bands are un-licensed frequency bands - meaning that users do not have to get permission from the FCC to operate in these bands.

802.11 is designed to operate in both a peer-to-peer (ad-hoc) architecture and a hierarchical architecture where nodes communicate with ***access points*** that are connected to a wired distribution system (often an Ethernet network -which 802.11 was designed to be compatible with.)

802.11 supports several different MAC techniques - these are called ***coordination functions***.

The ***Point Coordination Function*** is designed for the case where there is a centralized access point that can control each user's transmissions.

The ***Distributed Coordination Function*** does not use any central control and used a contention based MAC technique called CSMA/CA (CSMA with collision avoidance). One mode is similar to CSMA with a binary exponential back-off. The other mode is similar to the MACAW protocol discussed earlier. One interesting variation here is that each RTS and CTS frame contains a duration field; this indicates the length of time the channel will be utilized. Any node within range of the sender/receiver will receive these packets and uses this information to determine when the channel will become idle - this technique is called ***virtual carrier sensing***. This is combined with ***physical carrier sensing*** at the physical layer.

802.11 Frame Structure

16	16	48	48	48	16	48	0-18,496	32
Control	Duration	Addr1	Addr2	Addr3	SeqCtrl	Addr4	Payload	CRC

Control field - contains 6 bit **type** sub-field used to indicate if frame is data, RTS, CTS or other types of control frames.

Duration field - indicates time in microseconds that channel will be allocated for transmission.

The multiple address fields are provided for multi-hop situations.

The standard also provides provisions for number of other issues including location management and security issues.