CT2108 – Nets and Comms 1

Medium Access Control Sub Layer

Contents

- Channel allocation
- Multiple Access Protocols
  - Carrier Sense Multiple Access Protocols
  - Collision-Free Protocols
  - Wireless LAN Protocols
- Ethernet
  - Classical Ethernet
  - Fast Ethernet
  - Gigabit Ethernet

This presentation deals with broadcast networks and their protocols.

The Channel Allocation Problem

- In broadcast networks the key issue is how to determine how gets to use the channel when there is competition for it
  - Static Channel Allocation in LANs and MANs
    - FDM or TDM allocation
      - Problems when there is a large number of users, since spectrum will be wasted
  - Dynamic Channel Allocation in LANs and MANs
    - A number of assumptions are in place

CSMA Protocols

- Are protocols in which stations listen for a carrier (i.e. transmission) and act accordingly
- Networks based on these protocols can achieve better channel utilization than 1/e
- Protocols
  - 1 persistent CSMA
  - Non persistent CSMA
  - p persistent CSMA
  - CSMA CD

1 Persistent CSMA

- 1 persistent CSMA
  - When a station has data to send, it first listens to the channel
  - If channel is busy, the station waits until the channel is free. When detects an idle channel, it transmits the frame
  - If collision occurs, it will wait an random amount of time and starts again
  - The protocol is called 1 persistent, because the station sends with probability of 1 when finds the channel idle, meaning that is continuously listening

The propagation delay has an important effect on the performance of the protocol. There is a small chance that just after a station begins sending, another station will sense the channel and start sending (before the signal from the first one reached it). In this situation a collision can occur. The longer the propagation delay, the more important this effect becomes, and the worse the performance of the protocol.

Even if the propagation delay is zero, there will still be collisions – imagine two station wanting to transmit data at the same time. But a third one is buy transmitting. Both station will wait until the third will finish is transmission, will sense idle channel and will start sending.

Non Persistent CSMA

- Before sending a station senses the channel. If no activity, it sends its frame
- If channel is busy, then will not continue to sense the channel until it becomes idle, but it will retry at a latter time (waiting a random period of time and repeating the algorithm)
- With this algorithm, fewer collisions will happen; thus better channel utilization but with longer delays than 1 persistent CSMA algorithm

p Persistent CSMA

- It applies to slotted channels
- When a station becomes ready to send, it senses the channel. If it is idle will transmit with a probability of p. With a probability of q it defers to the next slot.
- If next slot is also idle, it transmits or it defers again with probabilities of p and q
- This process is repeated until the frame gets either transmitted or another station it began transmission
- For latter case, the unlucky station acts the same as it would have been a collision (waits a random time and starts again)

Persistent and Non-persistent CSMA

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

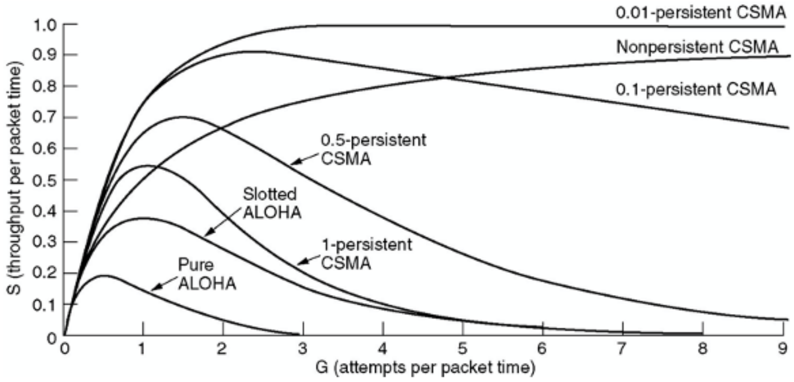

CSMA with Collision Detection

- An improvement over CSMA protocols is for a station to abort its transmission when it senses a collision.
- If two stations sense the channel idle and begin transmission at the same time, they will both detect the collision immediately; there is no point in continuing to send their frames, since they will be garbled.
- Rather than finishing the transmission, they will stop as soon as the collision is detected
  - Saves time and bandwidth

This protocol is widely used CSMA/CD. In particular is the basis for Ethernet LAN.

CSMA with Collision Detection

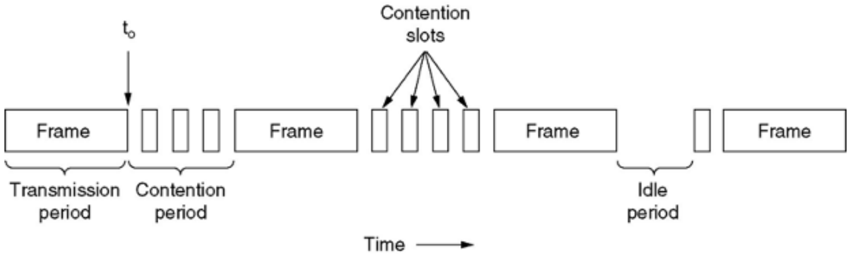

CSMA/CD can be in one of three states: contention, transmission, or idle.

At the point marked t0, a station has finished transmitting its frame. Any other station having a frame to send may now attempt to do so. If two or more stations decide to transmit simultaneously, there will be a collision. Collision can be detected by looking at the power of the line or at the pulse width of the received signal and comparing it with the transmitted signal.

After a station detects a collision, aborts its transmission, waits a random period of time and tries again latter. Therefore the CSMA/CD will consist of alternating contention and transmission periods, with idle periods occurring when all stations are quiet.

The minimum time to detect a collision is two times the time it takes the signal to propagate from one station to another. In worst case scenario, the two stations can be at the ends of the cable … therefore, the minimum time to detect a collision is the round trip propagation delay for the whole cable segment. The sending station has to monitor the channel for collisions during transmission. Therefore the CSMA/CD is a half duplex system.

It is important to realize that the collision detection is an analog process. The station's hardware must listen to the cable while is transmitting. If what it reads back is different than what it sends, then a collision must have had happened. The implication is that the signal encoding must allow collisions to be detected (i.e. a collision of two 0 V signals will never be detected). For this reason, special encoding is used.

Collision Free Protocols

- Collisions adversely affect the system performance, especially if the cable is long and the frames are short
- The collision free protocols solve the contention for the transmission channel without an collisions at all
- N stations are assumed to be connected to the same transmission channel
- Protocols
  - Bit-Map Protocol
  - Binary Countdown

Collision-Free Protocols (1)

The basic bit-map protocol.

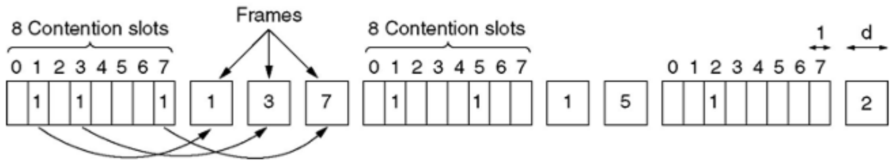

If station j has a frame to send, it will transmit a 1 in j-th contention slot

Each contention period consists of exact N slots. If station 0 has a frame to send, then it transmits a slot during zero-th contention slot. No other station is allowed to transmit during this slot. Regardless of what station 0 does, station 1 gets the opportunity to transmit a 1 bit during the contention slot 1, but only if it wants to send a frame (has a queued frame).

After all N slots have passed, each station has complete knowledge of which stations whish to transmit. Al that point they begin transmitting in numerical order. Since everyone agrees who goes next, there is no collisions.

Protocols like this are called RESERVATION PROTOCOLS.

Collision-Free Protocols (2)

The binary countdown protocol. A dash indicates silence.

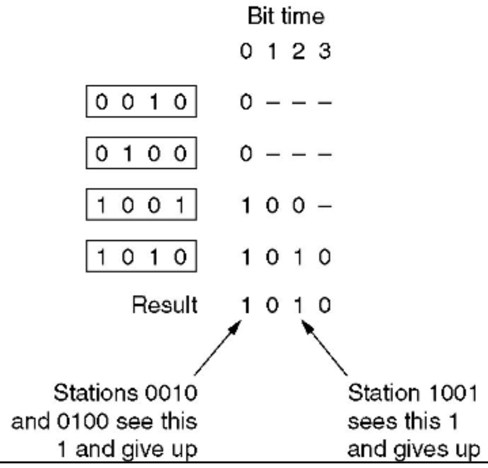

A problem with bit map protocols is that the overhead per station is fixed (on bit), therefore it is not scaling well with large number of stations.

In binary countdown protocol, all stations will have same length addresses. A station wanting to use the channel broadcasts its address as a binary bit string, starting with the high order bit. The bits in each address position, from each station are BOOLEAN OR-ed together. This is called binary countdown. All the stations will see the result of the OR operation instantaneously. So an station that placed an 0 on the channel sees a 1, than it will stop trying. A higher priority station wants to transmit therefore it has to stop trying this time.

In other words, as soon as a station has seen that a high order bit in its address with value 0 has been overwritten with a 1, it gives up.

Consider the stations 0010, 0100, 1001 and 1010 are all trying to get the channel. Station having address 1010 gets the channel (due to the fact that this protocol is giving priority to the stations with higher address).

Wireless LAN Protocols (1)

A wireless LAN. (a) A transmitting. (b) B transmitting.

CSMA doesn’t work because of the hidden station problem.

It is important to realize that in a wireless LAN system not all stations are within the range of the others. Due to this, HIDDEN STATION PROBLEM can occur (when A transmits, C can’t sense it and it can transmit too … the frame from A to B gets garbled if C decides to transmit at the same time). Therefore a simple CSMA approach will not work for wireless LANs.

EXPOSED STATION PROBLEM can also occur (the reverse situation – B transmits to A. If C senses the medium, it will hear an ongoing transmission and falsely conclude that may not send to D, when in fact such transmission would cause bad reception only in regions between B and C.

In conclusion, before starting a transmission, a station wants to know if there is activity around the receiver.

Wireless LAN Protocols (2)

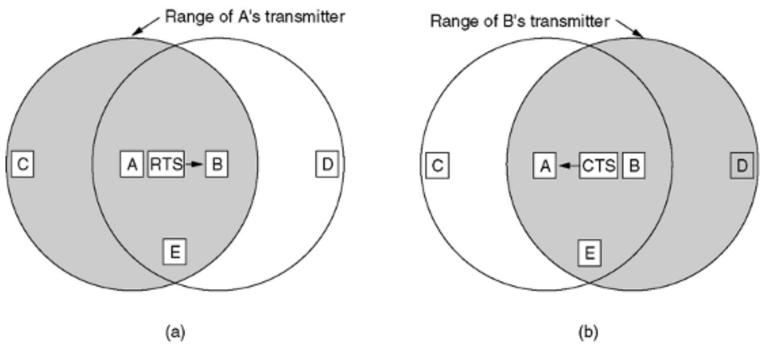

The MACA protocol. (a) A sending an RTS to B. (b) B responding with a CTS to A.

MACA (Multiple Access with Collision Avoidance) – the sender stimulates the receiver to output a short packet, so the stations around it will hear it and understand to avoid transmission during the upcoming (large) data frame.

In this slide we have station A wanting to send data to station B. A short exchange of packets RTS (Request to Send) and CTS (Clear To Send) takes place. RTS (short frame, less than 30 bytes) contains the length of the data that is about to be transferred (data frame). This data is copied also into the CTS frame by the receiver.

When A receives the CTS frame, it begins the transmission of the data frame.

Ethernet

- Historical Ethernet Cabling
- The Ethernet MAC Sublayer Protocol
- The Binary Exponential Back-off Algorithm
- Modern Switched Ethernet
- Fast Ethernet
- Gigabit Ethernet

Historical Ethernet Cabling

Three kinds of Ethernet cabling.

(a) 10Base5, (b) 10Base2, (c) 10Base-T.

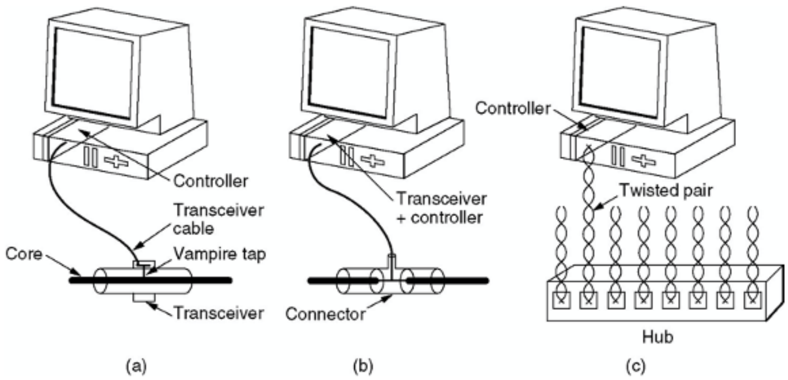

In 10Base5 the transceiver and the interface board are separated. The interface board goes inside the computer, while the transceiver goes on the cable (close to the vampire tap). The transceiver contains the electronics that handle the collision detection.

In 10Base2, the connection to the cable is just a passive BNC T junction. The transceiver electronics are situated onto the controller board (inside the PC).

In 10BaseT there is no shared cable at all, just the hub. Each station is connected to the hub, using a dedicated, non shared cable.

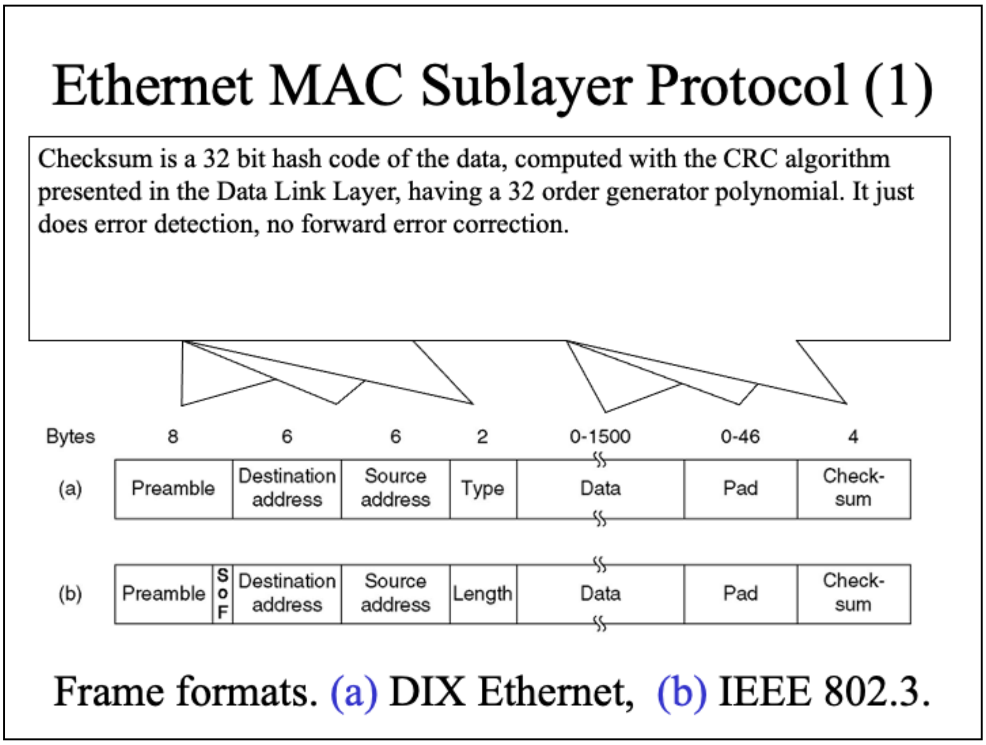

There are two variants of Ethernet. The initial one (DIX, initiated by Xerox, Intel and DEC) and the IEEE 802.3 standard. The differences are minor and they are interoperable.

Each frame starts with Preamble of 8 bytes, each containing the bit pattern 10101010. The Manchester encoding on this pattern produces a 10MHz square wave for 6.4 us to allow the receiver's clock to synchronize with the sender. The clock will stay in sync for the rest of the frame, using the transitions in the middle of the bit boundaries for adjustment.

Ethernet addresses: High order bit distinguishes between normal and group addresses (1 for group, 0 for normal) (allow multiple stations to listen to one address, when a frame is sent to a group address, all the stations in that group receive it). Second high order bit distinguishes between local and global addresses (1 for local, 0 for global). All 1s (ff:ff:ff:ff:ff:ff) is reserved for the broadcast address.

Type field in DIX: - tells to the receiver what to do with the frame. Multiple network protocols could be supported. So this field is unique for each network protocol. It is used to dispatch the incoming frames. IEEE 802.3 changed the type field in Length field – the type is handled as part of the data itself, in a small header. The length field contains the length in bytes of the data, up to 1500 bytes.

There is also a minimum frame length, related to the collision detection. The explanation is given in the next slide. The minimum Ethernet frame length is 64 bytes (without the preamble), and if the data portion is less than 46 bytes (64 – headers + checksum) then the PAD field is used to fill the frame to the minimum size.

Checksum is a 32 bit hash code of the data, computed with the CRC algorithm presented in the Data Link Layer, having a 32 order generator polynomial. It just does error detection, no forward error correction.

Ethernet MAC Sublayer Protocol (2)

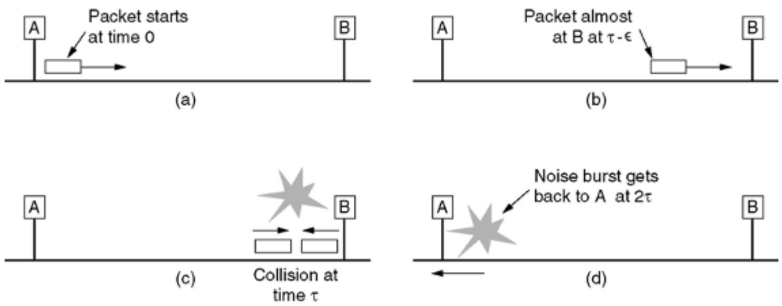

Collision detection can take as long as 2τ.

The strongest reason why there is a minimum length for the Ethernet frame is to allow the CSMA/CD protocol to work properly. In other words, the collisions should be detected by the transmitting station, or preventing a transmitting station to finish the transmission of a frame until the first bit reached the far end of the cable, where it may collide with another frame.

If we have a collision an time t, the station B (that detected the collision first) issues a 48 bit noise burst (known as jam signal) to warn the other stations. At about 2t time, the sender (station A) will see the noise burst and aborts its transmission too. It waits a random time until it tries again.

For a 2500 meters LAN with 4 repeaters, the round trip time (the propagation time for a signal to travel from one end of the LAN to the other end of the LAN and back) is about 50us. That is equivalent to about 500 bit times (one bit time for 10Mbps network is 100ns). Therefore, a station should be in transmit mode for at least 500 bit times, which gives us the minimum Ethernet frame of about 500 bits. The standard's choice was 512 bits or 64 bytes minimum frame.

So frames with fewer than 64 bytes will be padded out to 64 bytes.

As the network speed goes up, the minimum frame length goes up too.

Binary Exponential Backoff Algorithm

- After a collision the time is divided in discrete slots (equal to worst round trip propagation, which is 512 bits time or 51.2 us)
- After the first collision, each station waits 0 or 1 slot time before tries again
  - If two station collide and they pick same number, they will collide again
- After a second collision, each station waits 0, 1, 2 or 3 at random and waits that number of slot times.
- After a third collision will happen, the next number to pick is between 0 and 2^3 -1 and that number of slots is skipped.
- After 10 collisions have been reached, the number interval is frozen at 0 – 1023.
- After 16 collisions, the station gives up to send the frame and reports the failure. Further recovery it is up to the higher layers.

Modern Switched Ethernet

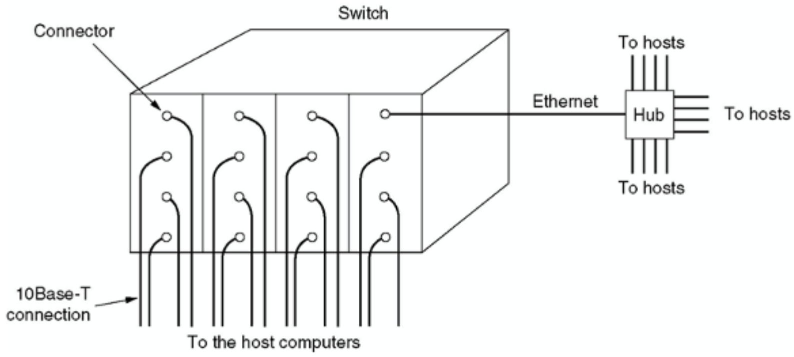

- Way to deal with saturated Ethernet LANs
- Switch – contains a high speed backplane
  - switches frames from incoming ports to destination ports
  - Avoids collisions

More stations added on an Ethernet, the traffic will go up. Eventually, at one point the LAN will saturate.

Switches are typically devices that have 4 to 32 plug-in cards, each with one to 8 Ethernet connectors (RJ45). When a station wants to send an Ethernet frame, the plug-in card getting the frame will check to see if the destination address is of one of the cards connected onto one of its ports. If yes, then the packet is copied onto that port. If not, then, using a multi-gigabit per second back plane, the frame is copied to the card containing the destination port.

Fast Ethernet

- Approved as IEEE 802.3 u standard in 1995
- Keeps all the old frame formats, interfaces and procedural rules
- Reduces the bit time from 100ns to 10ns
- It is based only on the 10Base-T wiring
  - It is using only hubs and switches; drop cables with vampire taps and BNC connectors are not possible
- It supports both UTP Cat 3 (for backwards compatibility with preinstalled infrastructure) and UTP Cat 5 cables

Fast Ethernet - Cabling

Name	Cable	Max. segment	Advantages
100Base-T4	Twisted pair	100 m	Uses category 3 UTP
100Base-TX	Twisted pair	100 m	Full duplex at 100 Mbps
100Base-FX	Fiber optics	2000 m	Full duplex at 100 Mbps; long runs

The original fast Ethernet (802.3u) cabling.

Anyone who understands classic Ethernet already understands much about Fast Ethernet. Fast Ethernet uses the same cabling and access method as 10Base-T. With certain exceptions, Fast Ethernet is simply regular Ethernet, just ten times faster.

Probably the most popular form of Fast Ethernet is **100BASE-TX**. 100BASE-TX runs on UTP Category 5 unshielded twisted pair, sometimes called UTP-5. It uses the same pair and pin configurations as 10Base-T, and is topologically similar in running from a number of stations to a central hub.

As an upgrade to 10Mbps Ethernet over Multimode fiber (10Base-F), **100BASE-FX** is Fast Ethernet over fiber. Distances up to 2km are supported.

Fast Ethernet is possible on Category 3 UTP with **100BASE-T4**. There is a popular misconception that Fast Ethernet will only run on Category 5 cable. That is true only for 100BASE-TX. If you have Category 3 cable with all four pairs (8 wires) connected between station and hub, you can still use it for Fast Ethernet by running 100BASE-T4.

100BASE-T4 sends 100Mbps over the relatively slow UTP-3 wire by fanning out the signal to three pairs of wire. This "demultiplexing" slows down each byte enough that the signal won't overrun the cable. Category 3 cable has four pairs of wire, eight wires total, running from point to point. 10Base-T only uses four wires, two pairs. Some cables only have these two pairs connected in the RJ-45 plug. If the category 3 cabling at your site has all four pairs between hub and workstation, you can use Fast Ethernet by running 100BASE-T4.

Fast Ethernet – 100Base-TX (1)

- Is using only two pairs out of the 4 available in the UTP cable (one for transmit and one for receive)
- For 100 Mbps, the waveform frequency would peak at 50MHz, while with Manchester encoding would pick at 100MHz
  - Category 5 UTP is only rated at 100MHz, so Fast Ethernet would be difficult to implement using Manchester encoding
- 100BASE-TX uses two encoding techniques:
  - 4B/5B coding schema is used to avoid loss of synchronization
  - To decrease the frequency on UTP cable, MLT-3 (Multiple Level Transition - 3 levels) encoding is used

The primary difficulty with 100 Mbps transmission of data is that high-frequency signals don't propagate well over either twisted pair or fiber. 10 Mbps Ethernet uses Manchester encoding to include a clock signal with every data bit. However, the clocking almost doubles the rate of transmission, so a worst-case scenario would transmit 10Mbps of data with a 10MHz waveform. For 100 Mbps, the waveform frequency would peak at 100MHz. Category 5 UTP is only rated at 100MHz, so Fast Ethernet would be difficult to implement.

100BASE-FX uses NRZI (Non-Return-to-Zero, Invert-on-one). To decrease the frequency even further on UTP, 100BASE-TX adds a variation of NRZI at the PMD sub-layer called either MLT-3 (Multiple Level Transition - 3 levels) or NRZI-3.

Fast Ethernet –100Base-TX (2)

- In order to send information using 4B5B encoding, the data byte to be sent is first broken into two nibbles.
  - If the byte is 0E, the first nibble is 0 and the second nibble is E.
- Next each nibble is remapped according to the 4B5B table
  - Hex 0 is remapped to the 4B5B code 11110.
  - Hex E is remapped to the 4B5B code 11100.
- The 4B5B replacement happens at the Physical layer, followed by MLT-3 encoding
- **4B5B Encoding Table:**

Data	Binary	4B/5B code
0	0000	11110
1	0001	01001
2	0010	10100
…		
E	1110	11100
F	1111	11101

Fast Ethernet –100Base-TX (3)

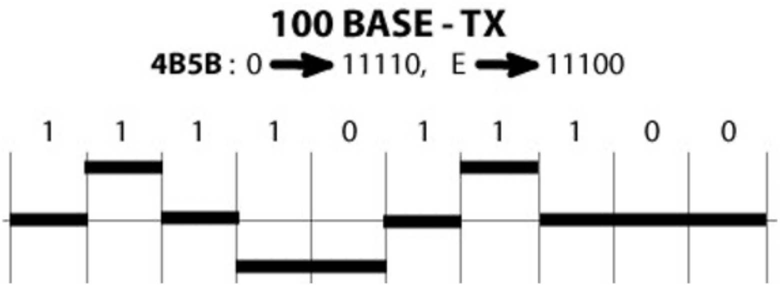

- MLT-3 (Multiple Level Transition) encoding
- It is using three voltage levels: +1V, 0V and -1V
- Encodes a bit as a presence or lack of transition
  - 1 encoded as a presence of transition
  - 0 as a lack of transition

MLT-3 encodes a bit as presence or lack of transition, exactly as in NRZI. What makes MLT-3 different is that the base waveform is a 3-state alternating wave. Rather than alternating between 0 and 1 as in Manchester encoding and NRZI, MLT-3 alternates from -1 to 0 to +1, back to 0, then back to -1, repeating indefinitely. A zero is encoded as a halt in the back-and-forth progression. It may be useful to think of MLT-3 as a stop-and-go sine wave, encoding 0 as stop and 1 as go. Using MLT-3, it is therefore possible to represent four or more bits with every complete waveform, at 0, +1, 0, and -1

With 4B/5B encoding, the data rate of the cable goes up to 125Mb. After the MLT-3 encoding the base frequency of the encoded signal goes down by 4 (due to the way the encoding works), therefore the signal that is transmitted over the cable will have a maximum frequency component of about 31.25Mhz, which is handled very well by the UTP5 cable.

The data being transmitted in this example is the hexadecimal byte "0E". First, the byte is broken down into four bit "nibbles", giving a nibble of "0" and a nibble of "E". Then, each nibble is looked up in a table to find the byte code associated with that hexadecimal number. The code for "0" hex is "11110" and the code for "E" hex is "11100". Finally, the byte codes are transmitted onto the wire using MLT-3 encoding.

Fast Ethernet – 100BaseFX (1)

- Is using optical fiber
- 100BASE-FX uses two encoding techniques:
  - 4B/5B coding schema is used to avoid loss of synchronization
  - NRZI (Non-Return-To-Zero, Invert-on-one) encoding is used
- The distance between a station and a hub can be up to 2 KM.

Fast Ethernet – 100BaseFX (2)

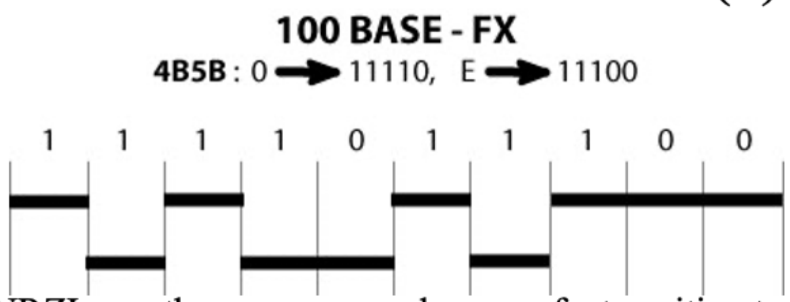

- NRZI uses the presence or absence of a transition to signify a bit
  - indicating a logical 1 by inverting the state and a 0 by maintaining the state.
- NRZI uses a half-wave to encode each bit (having a better usage of the bandwidth)
  - Rather than transition through ground at each bit like Manchester encoding,

The data being transmitted in this example is the hexadecimal byte "0E". First, the byte is broken down into four bit "nibbles", giving a nibble of "0" and a nibble of "E". Then, each nibble is looked up in a table to find the byte code associated with that hexadecimal number. The code for "0" hex is "11110" and the code for "E" hex is "11100". Finally, the byte codes are transmitted onto the wire using NRZI encoding.

Gigabit Ethernet (1)

- IEEE 802.3z approved in 1998
- All configurations are point to point rather than multidrop as in classical Ethernet
- Supports two modes of operation:
  - Full Duplex mode (the normal mode of operation)
    - All the stations are connected through a switch that allows traffic in both directions at the same time
    - CSMA/CD is not employed anymore
    - Max cable length is governed by propagation issues
  - Half duplex mode
    - All stations are connected through a hub, that internally electrically connects all the lines, simulating a multidrop environment as with classic Ethernet electrically connects all the lines
    - CSMA/CD protocol is required
    - Max cable length is still governed by CSMA/CD protocol

It supports two modes of operation:

Full Duplex mode

All the stations are connected through a switch that allows traffic in both directions at the same time

Since all possible connections are point to point, the CSMA/CD protocol is not employed anymore

Since on the line between the computer and the switch, no other devices are present

The max length of the cable is now governed by signal propagation issues rather than the round trip propagation

Half duplex mode

All stations are connected through a hub, that internally electrically connects all the lines, simulating a multidrop environment as with classic Ethernet electrically connects all the lines

Gigabit Ethernet (2)

- Half Duplex Mode Max Cable length extended using:
  - Carrier extension
    - Hardware adds its own padding to extend the frame to 512 bytes; since this padding is added by sending hardware and removed by the receiving hardware, the software is not aware of it
  - Frame Bursting
    - Allows the sender to transmit a concatenated sequence of multiple frames in one single transmission; if the total length of burst is less then 512 bytes, the hardware pads it again.
- Those techniques extend the radius of the network to 200 meters

Since 25 meters is not an acceptable radius (to maintain the minimum 64 bytes frame), two methods have been adopted by the standards committee:

1. Carrier extension
2. Frame bursting

Gigabit Ethernet (3)

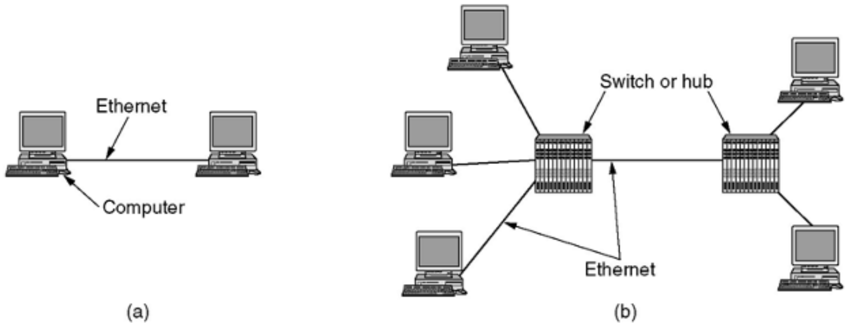

(a) A two-station Ethernet. (b) A multi-station Ethernet

Gigabit Ethernet (4)

Gigabit Ethernet cabling.

Name	Cable	Max. segment	Advantages
1000Base-SX	Fiber optics	550 m	Multimode fiber (50, 62.5 microns)
1000Base-LX	Fiber optics	5000 m	Single (10 μ) or multimode (50, 62.5 μ)
1000Base-CX	2 Pairs of STP	25 m	Shielded twisted pair
1000Base-T	4 Pairs of UTP	100 m	Standard category 5 UTP