20IS709

Communication Systems For Industrial Networking

Ethernet

Ethernet

- Invented in 1973 @ Xerox.
- Complex networking technology that defines wiring and signaling standards for the *physical layer* through the means of the *data link layer* and a common *addressing format*.
- Standardized as **IEEE 802.3**
- Originally a LAN technology – extended to MAN / WAN

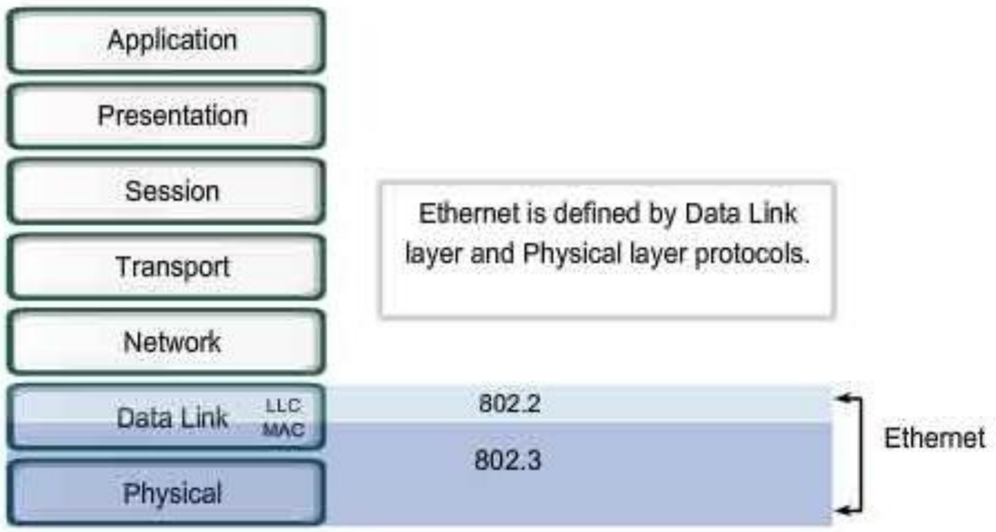

Ethernet

Advantages

- Versatile
  - Best known use is in linking desktop computers in offices
  - Can transfer any kind of data, from short messages to huge files
  - Uses higher-level protocols such as TCP and IP, or it can use an application-specific protocol.
- Easy to Use
  - All of the computers in the network follow standard Ethernet specifications for interconnecting, managing network traffic, and exchanging data.
  - Allow choices - a network may use twisted pair, fiber-optic, or coaxial cable.
- A Wide Selection of Products is Available - Hardware, software, and debugging tools
- The Hardware Controls Network Access - hardware manages the network traffic, so there's no need for software to control network access
- Fast - supports speeds from 10 Megabits per second (Mb/s) to 10 Gigabits per second (Gb/s)
- Can Span Long Distances – from 100 m to 5 km; extended using repeater hubs or switches
- Cost Is Reasonable

Ethernet Physical Layer

- IEEE 802.3 standard defines a range of cable types - coaxial cable, twisted pair cable and fiber optic cable

Table : Ethernet allows a choice of three types of cable.

Cable Type	Twisted Pair	Fiber Optic	Coaxial
Maximum data rate (Mb/sec.)	1000	10,000	10
Maximum length per segment (meters)	100+	2000 (half duplex, 10 Mb/s), 5000 (full duplex, 10 Mb/s)	500 (thick coax), 185 (thin coax)
Cost	low	high	moderate
Noise immunity	good	excellent	good
Ease of Installation	excellent	good with prefabricated cables, fair/poor if attaching connectors on raw cable	fair/poor

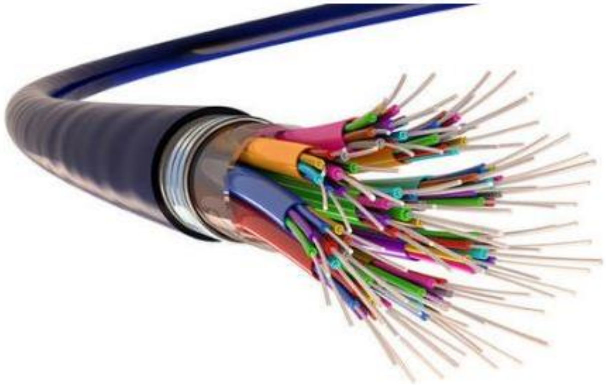

Ethernet Physical Layer

- 10Base2
  - Thin wire coaxial cable (0.25 inch diameter), 10 Mbps, single cable bus
- 10Base5
  - Thick wire coaxial cable (0.5 inch diameter), 10 Mbps, single cable bus
- 10BaseT
  - Unscreened twisted pair cable (0.4 to 0.6 mm conductor diameter), 10 Mbps, twin cable bus
- 10BaseF
  - Optical fiber cables, 10 Mbps, twin fiber bus
- 1Base5
  - Unscreened twisted pair cables, 1 Mbps, twin cable bus
- 10Broad36
  - Cable television (CATV) type cable, 10 Mbps, broadband

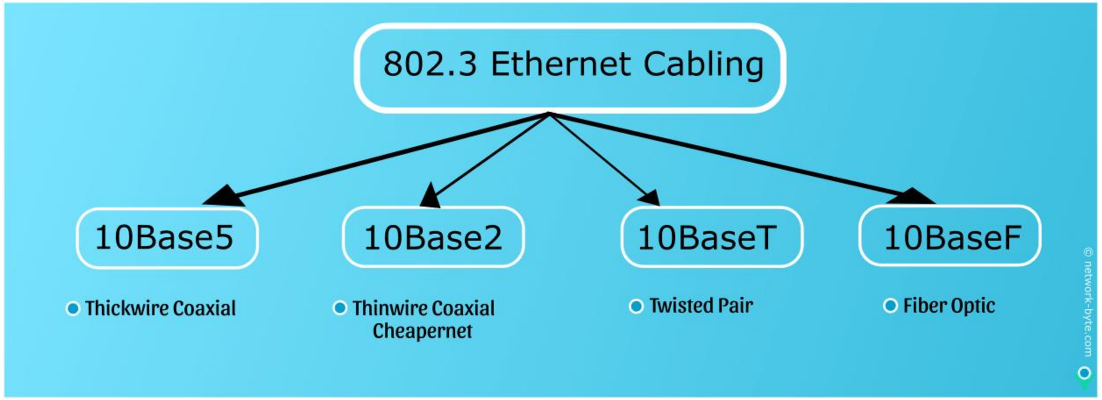

Ethernet Physical Layer

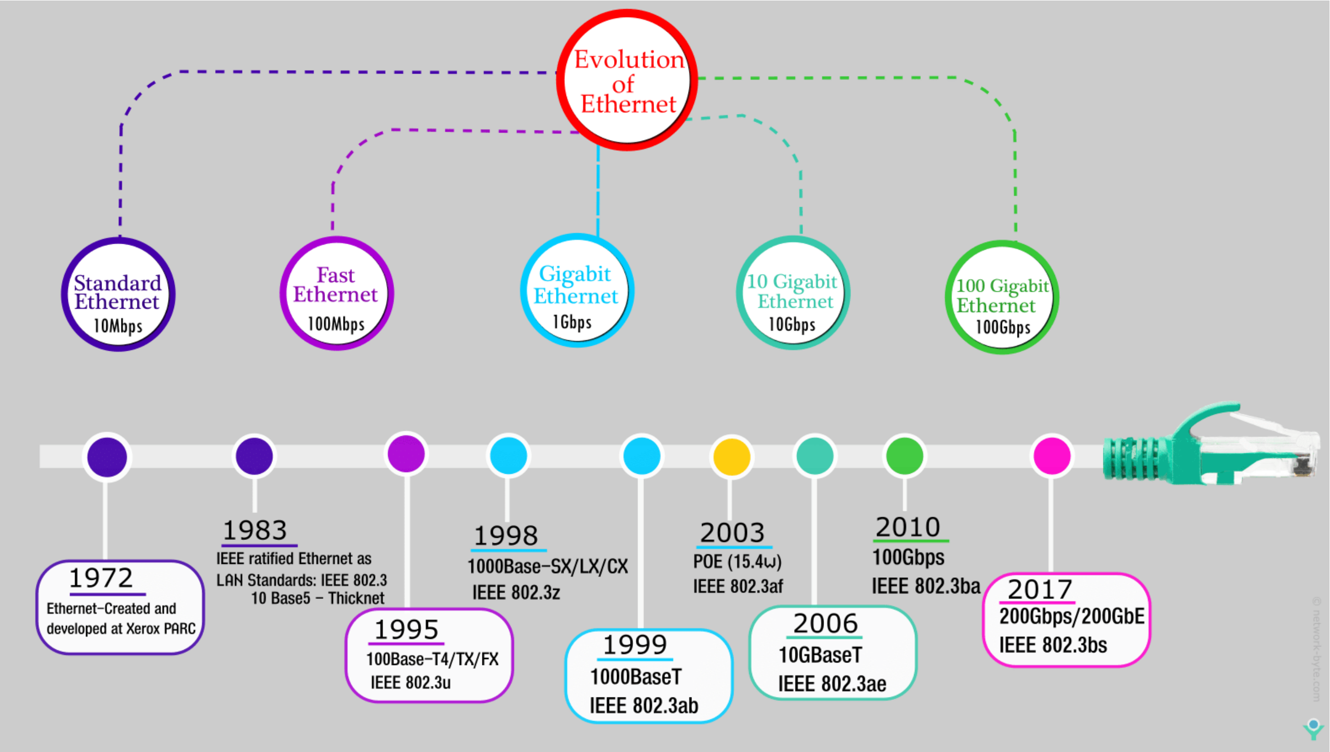

Ethernet Type	Bandwidth	Cable Type	Duplex	Maximum Distance
10Base-5	10 Mbps	Thicknet Coaxial	Half	500 m
10Base-2	10 Mbps	Thinnet Coaxial	Half	185 m
100Base-TX	10 Mbps	Cat3/Cat5 UTP	Half	100 m
100Base-TX	100 Mbps	Cat5 UTP	Half	100 m
100Base-FX	200 Mbps	Cat5 UTP	Full	100 m
100Base-FX	100 Mbps	Multimode Fiber	Half	400 m
1000Base-T	200 Mbps	Multimode Fiber	Full	2 km
1000Base-TX	1 Gbps	Cat5e UTP	Full	100 m
1000Base-SX	1 Gbps	Cat6 UTP	Full	100 m
1000Base-LX	1 Gbps	Multimode Fiber	Full	550 m
10GBase-CX4	1 Gbps	Single-Mode Fiber	Full	2 km
10GBase-T	10 Gbps	Twin-axial	Full	100 m
10GBase-LX4	10 Gbps	Cat6a/Cat7 UTP	Full	100 m
10GBase-LX4	10 Gbps	Multimode Fiber	Full	300 m
10 Mbps	10 Gbps	Single-Mode Fiber	Full	10 km

Ethernet Frames

- All data in an Ethernet network travels in structures called frames.
- An Ethernet frame has defined fields for data and other information to help the data get to its destination and to help the destination computer determine whether the data has arrived intact.
- The Ethernet controller’s hardware places information to be sent in frames for transmitting, and extracts and stores the information in received frames

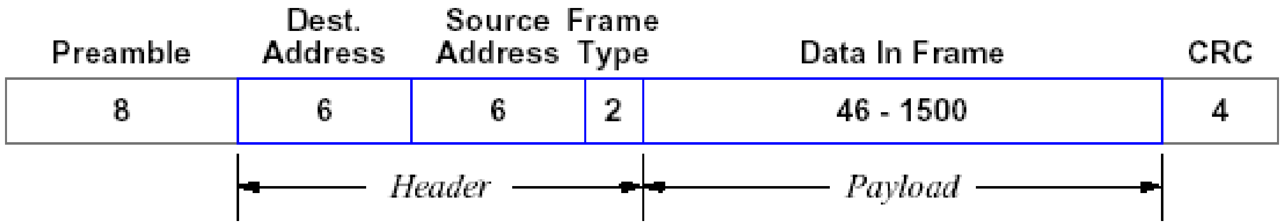

- Synchronizing bits, addressing information, an error-checking sequence, and additional identifying information are added to the data being sent

Ethernet Frames

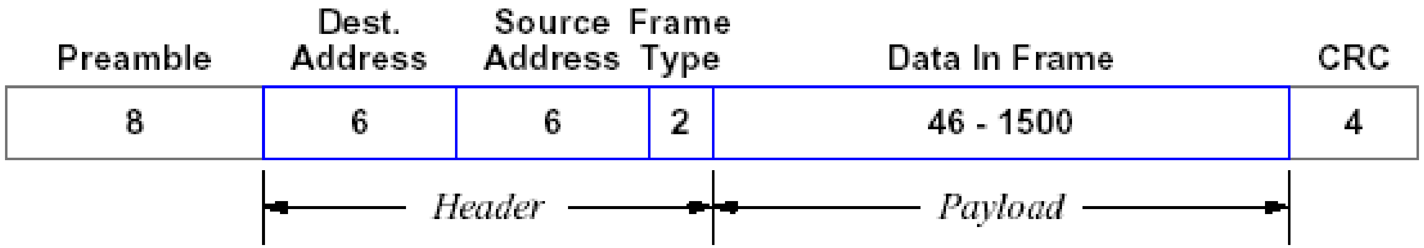

Preamble and Start Frame Delimiter

- A predictable bit pattern that enables the interfaces on a 10-Mb/s network to synchronize to a new frame being transmitted.
- Each transmitted word begins with a Start bit.
- The receiver uses the leading edge of the Start bit as a timing reference to predict when to read each of the bits that will follow.
- The Preamble consists of seven identical bytes, each with the value 10101010.
- The Start Frame Delimiter follows the Preamble, and consists of the byte 10101011
- Final two bits in the Start Frame Delimiter indicate the end of the Preamble.

Ethernet Frames

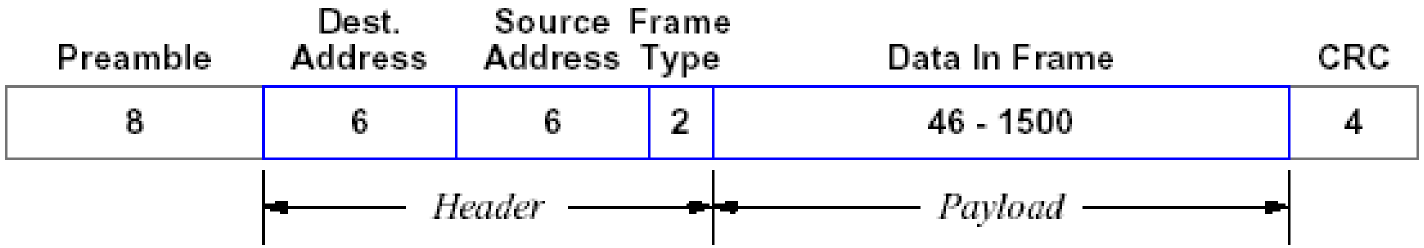

Destination Address

- Every Ethernet interface has a 48-bit physical, or hardware, address that identifies the interface on the network.
- The Destination Address field contains the physical address of the intended receiver of the frame.
- The first bit is 0 if the address is for a single interface, and 1 if the address is a multicast or broadcast address.
- A broadcast address is all 1s (FFFFFFFFFFFF) and is directed to every interface in the network.
- The second bit of the destination address is zero if the address was assigned by the manufacturer of the interface and is 1 if the address is administered locally.

Source Address

- The Source Address field contains the 48-bit physical address of the transmitting interface.

Ethernet Frames

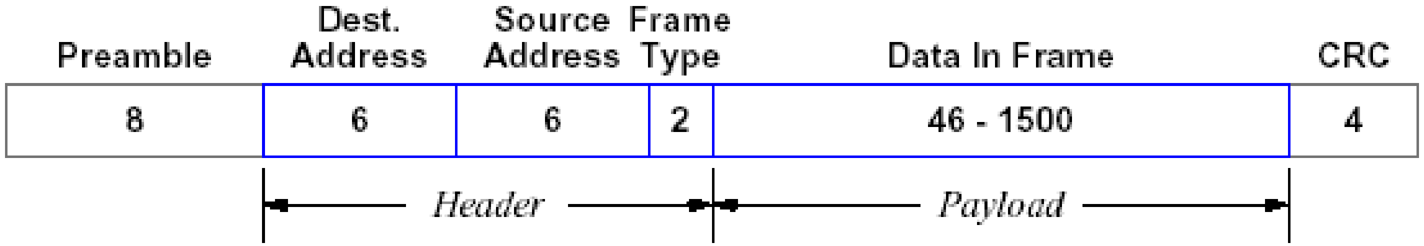

Length/Type

- 16 bits - indicate the number of bytes of valid data in the data field or the protocol used by the data in the field that follows.
- If the value is less than or equal to 1500 decimal (5DCh), the value indicates length - data field must contain between 46 and 1500 bytes.
- If there are less than 46 bytes of valid, or usable, data, the length field can indicate how many of the bytes are valid data
- If the value is greater than or equal to 1536 decimal (600h), the Length/Type field indicates the protocol that the contents of the data field use. Eg: The value for the Internet Protocol (IP) is 800h

Ethernet Frames

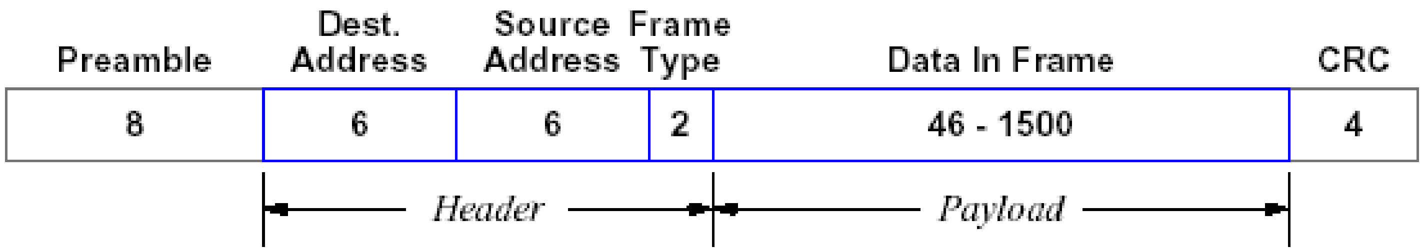

Data

- Between 46 and 1500 bytes
- If there are fewer than 46 bytes of data, the field must include pad bytes to increase the size to 46 bytes.
- If the transmitting interface has more than 1500 bytes to send, it uses multiple frames.
- Frames with a full 1500 data bytes are the most efficient because they have just 26 bytes, or less than 2 percent, of overhead.
- An Ethernet frame must be at least 512 bits (64 bytes) not including the Preamble and Start-of-Frame bits - size of a frame with the minimum 46 data bytes - receiving interfaces ignore frames that are shorter than this minimum size.

Ethernet Frames

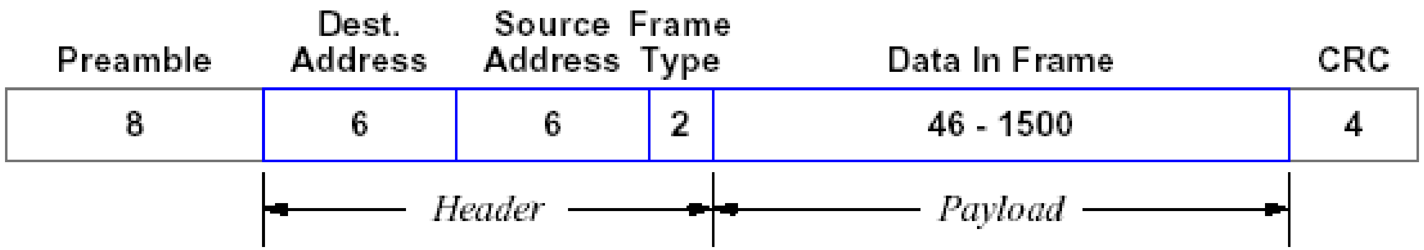

Frame Check Sequence

- A receiving interface can detect corrupted data by using the 32-bit cyclic redundancy check (CRC) value in the frame check sequence field.
- The transmitting interface performs a calculation, called the cyclic redundancy check, on the bytes to be sent and places the result in the frame check sequence field.
- The receiving interface performs the same calculation on the received bytes.
- The Ethernet controller’s hardware typically performs the CRC calculations on both ends.
- On detecting an error in a received frame, the controller typically sets a bit in a status register.

Media Access Control

- Method of deciding who gets to transmit – Deciding when to transmit
- determines the efficiency of the network
- One computer is the master, and the other computers transmit only after receiving permission from the master - USB interface
- Token-passing network - the computers take turns - a register bit or sequence of bits as token is set to indicate possession
- Only the computer holding the token can transmit and when finishes transmitting, it passes the token to another computer.
- *Carrier sense multiple access with collision detection*, or CSMA/CD - allows any interface to attempt to transmit any time the network is idle
- Ethernet uses CSMA/CD.

Media Access Control

CSMA/CD

- *Carrier sense* - an interface that wants to transmit must monitor the network and sense, or detect, when the network is idle, indicated by the absence of a carrier.
- *Multiple access* - no single interface controls the network traffic - any interface can attempt to transmit on a network that has been idle for at least the amount of time defined as the interframe gap.
- In a 10Mb/s network, the interframe gap equals 9.6 microseconds.
- Ethernet controller's hardware handles the sending and receiving of frames, including detecting collisions and deciding when to try again after a collision.
- The CPU writes the data to send into memory that the controller can access, and the controller stores received data in memory that the CPU can access.
- CPU uses interrupts or polling to learn of the success or failure of a transmission and the arrival of received.

Collisions

- When two interfaces in the same *collision domain* try to transmit at the same time.

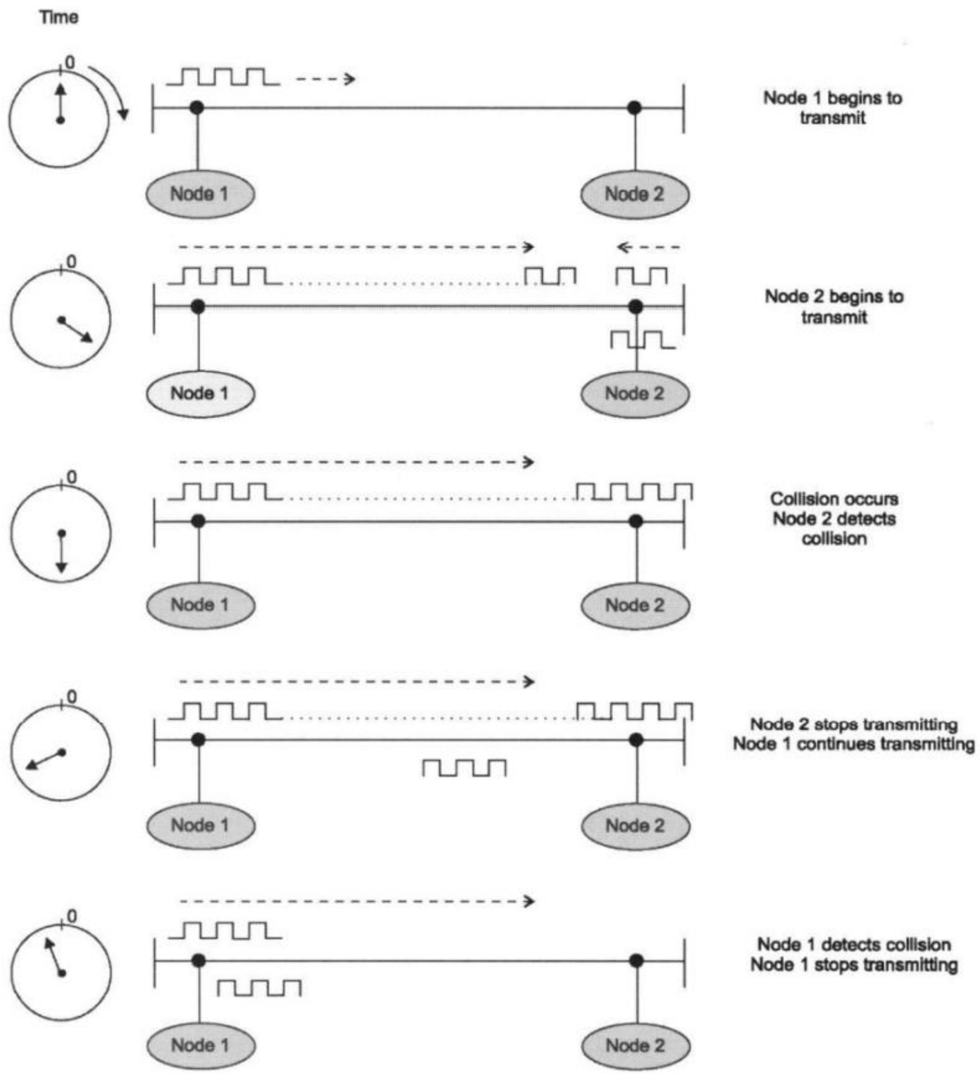

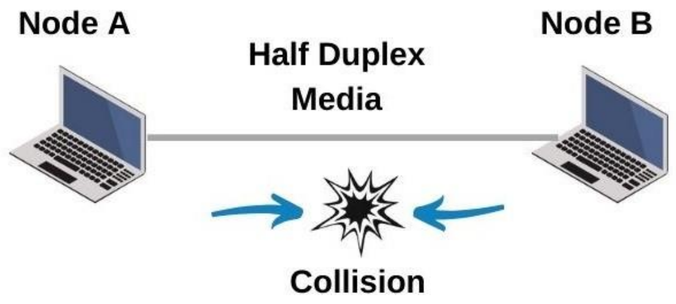

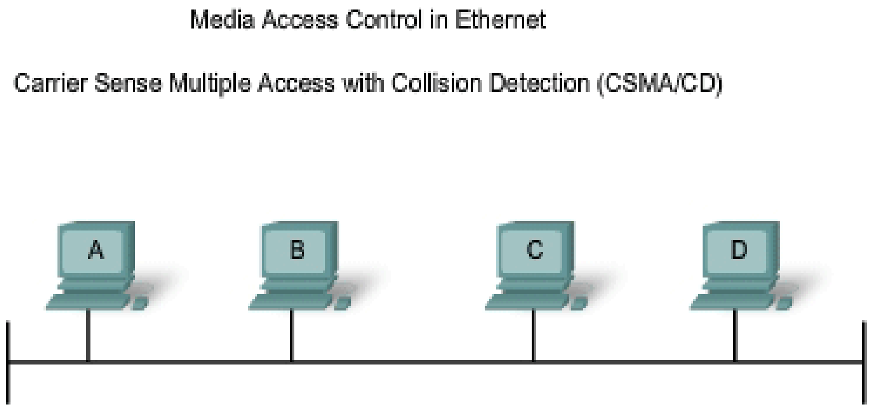

Collisions

- When two interfaces in the same *collision domain* try to transmit at the same time.
- On detecting a collision, the transmitting interface doesn’t stop transmitting immediately - continues long enough to be sure that the other transmitting interface(s) have time to detect the collision.
- A transmitting interface that has detected a collision always finishes sending the 64 bits of the Preamble and Start of Frame Delimiter
- Following these, the interface sends an additional 32 bits called the *jam signal*, then stops transmitting. The jam signal can be any arbitrary data except the previous frame’s CRC value – random bits.
- After an interface stops transmitting due to a collision, the next task is deciding when to try again – wait for the same amount of time and then retry - another collision will occur.
- Ethernet standard defines a backoff process where each interface selects a randomly chosen delay time before attempting to retransmit.
- The delays are multiples of the interface’s slot time - specified in units of bit times.
- For 10-Mb/s, the slot time is 512 bit times, works out to 51.2 microseconds.

Reducing Collisions

- The main reasons for collision rates on an Ethernet network are:
  - The number of packets per second
  - The signal propagation delay between transmitting nodes
  - The number of stations initiating packets
  - The bandwidth utilization

- Suggestions on reducing collisions in an Ethernet network are:
  - Keep all cables as short as possible
  - Keep all high activity sources and their destinations as close as possible.
  - Possibly isolate these nodes from the main network backbone with bridges/routers to reduce backbone traffic
  - Use buffered repeaters rather than bit repeaters
  - Check for unnecessary broadcast packets that are aimed at non existent nodes
  - Remember that the monitoring equipment to check out network traffic can contribute to the traffic (and the collision rate)

Physical Address or MAC address

- To send an Ethernet frame on the network, a computer places its physical address in the Source Address field and the places destination's physical address in the Destination Address field.
- Two parts - a 24-bit *Organizationally Unique Identifier* (OUI) that identifies the interface's manufacturer and an additional 24 bits that are unique to the piece of hardware.
- The IEEE grants the rights to use an OUI.
- An interface card purchased for a PC or an embedded-system module with an Ethernet interface typically has a physical address programmed into the hardware.
- Expressed as a series of six hexadecimal bytes:

00-90-C2-C0-D3-EA

  - 00-90-C2 is the OUI and C0-D3-EA is the unique value assigned by the owner of the OUI to a specific piece of hardware.

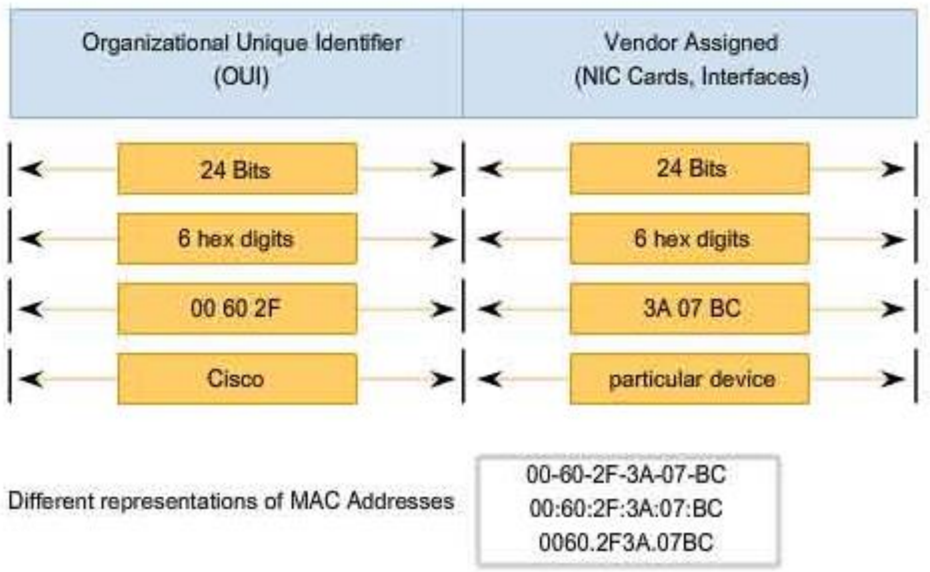

10Base5 systems

- Coaxial cable system and uses the original cable for Ethernet systems - generically called 'Thicknet'
- 10 Mbps base band signaling on a cable that will support 500-meter segment lengths.
- Laid in a cabling tray and the transceiver electronics (medium attachment unit, MAU) is installed directly on the cable.
- From there an intermediate cable, known as an attachment unit interface (AUI) cable is used to connect to the network interface card (NIC).
- This cable can be a maximum of 50 meters long, compensating for the lack of flexibility of placement of the segment cable.
- The AUI cable consists of 5 individually shielded pairs - two each (control and data) for both transmit and receive; plus one for power.

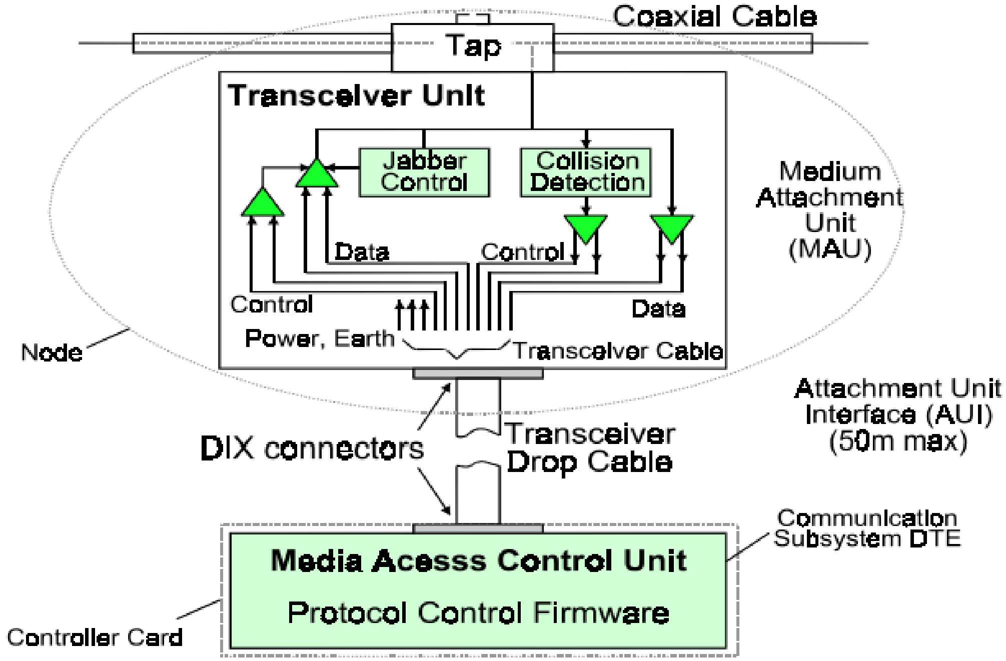

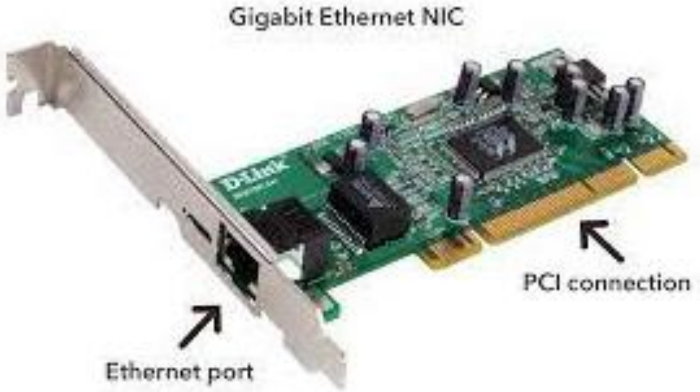

10Base5 systems

- The location of the connection is important to avoid multiple electrical reflections on the cable, and the Thicknet cable is marked every 2.5 meters with a black or brown ring to indicate where a tap should be placed.
- Fan out boxes can be used if there are a number of nodes for connection.
- The connection at either end of the AUI cable is made through a 25-pin D-connector, with a slide latch - called a DIX connector.

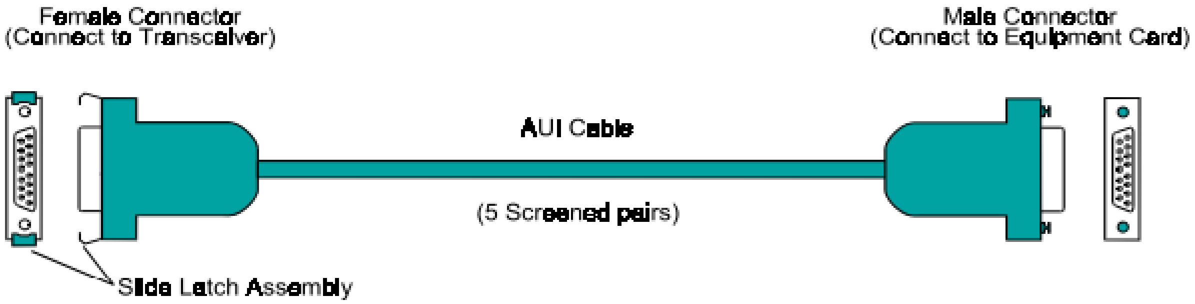

- There are certain requirements if this cable architecture is used in a network
  - Segments must be less than 500 meters in length to avoid signal attenuation problems
  - No more than 100 taps on each segment i.e. not every potential connection point can support a tap
  - Taps must be placed at integer multiples of 2.5 meters
  - The cable must be terminated with a 50-ohm terminator at each end
  - One end of the cable shield must be earthed

10Base5 systems

- Transceiver is remote from the NIC, the node needs to be aware that the termination can detect collisions if they occur.
- This confirmation is performed by a signal quality error (SQE), or heartbeat, test function in the MAU.
- The SQE signal is sent from the MAU to the node on detecting a collision on the bus.

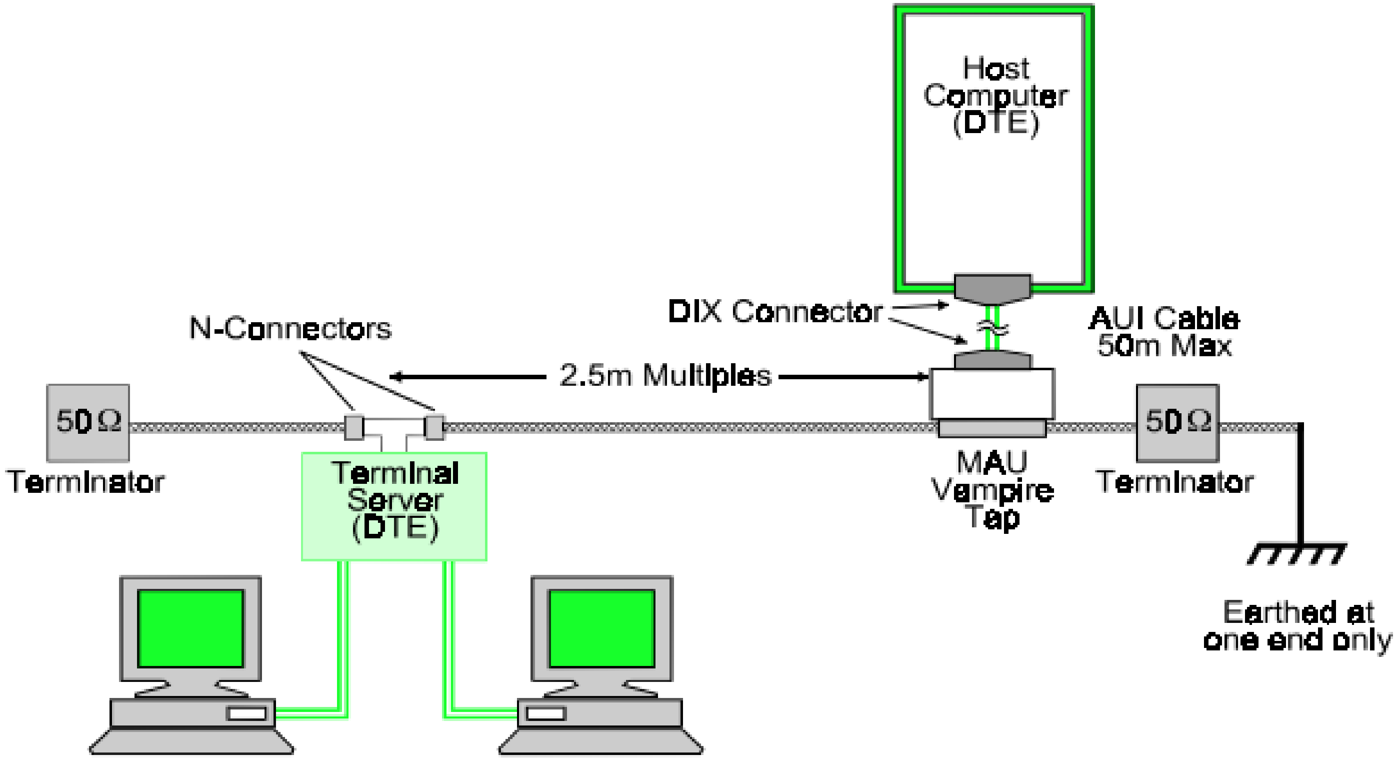

10Base2 systems

- Coaxial cable Ethernet referred to as 'Thinnet' or sometimes 'thinwire Ethernet'
- The cable is normally connected to the NICs in the nodes by means of a BNC T-piece connector.
- Connectivity requirements include:
  - It must be terminated at each end with a 50-ohm terminator
  - The maximum length of a cable segment is 185 meters and NOT 200 meters
  - No more than 30 transceivers can be connected to any one segment
  - There must be a minimum spacing of 0.5 meters between nodes.
  - It may not be used as a link segment between two 'Thicknet' segments
  - The minimum bend radius is 5 cm

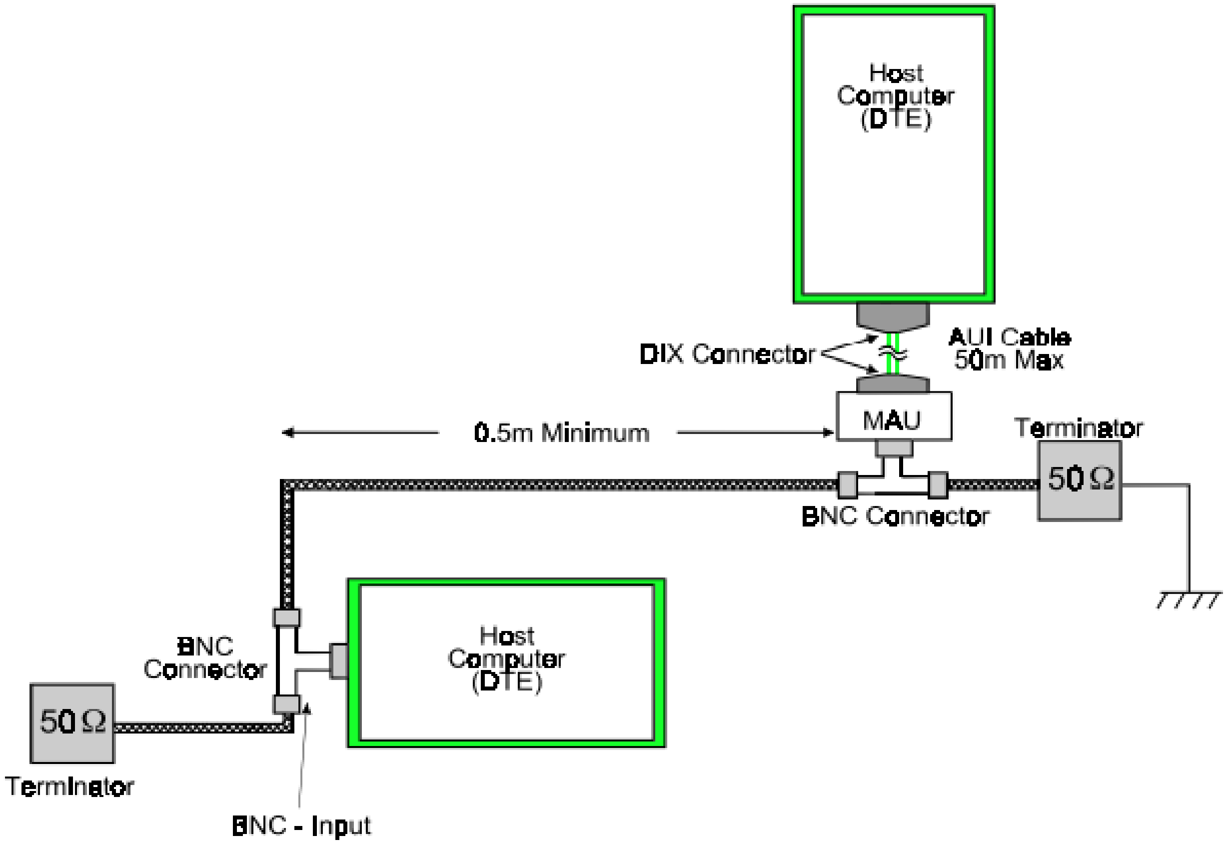

10BaseT systems

- Uses AWG24 unshielded twisted pair (UTP) cable for connection to the node.
- Physical topology of the standard is a star, with nodes connected to a wiring hub, or concentrator.
- The node cable can be category 3 or category 4 cable.
- Advised to consider category 5 for all new installations - allow an upgrade path as higher speed networks become more common.
- The node cable has a maximum length of 100 meters; consists of two pairs for receive and transmit and is connected via RJ45 plugs

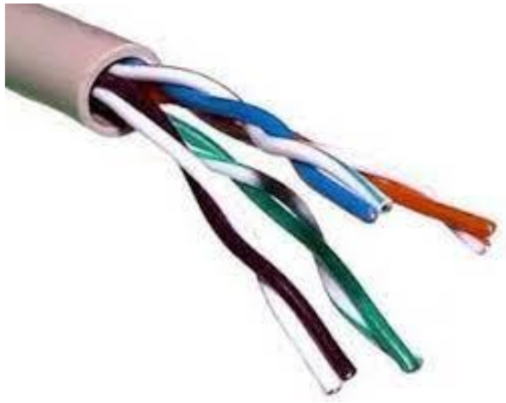

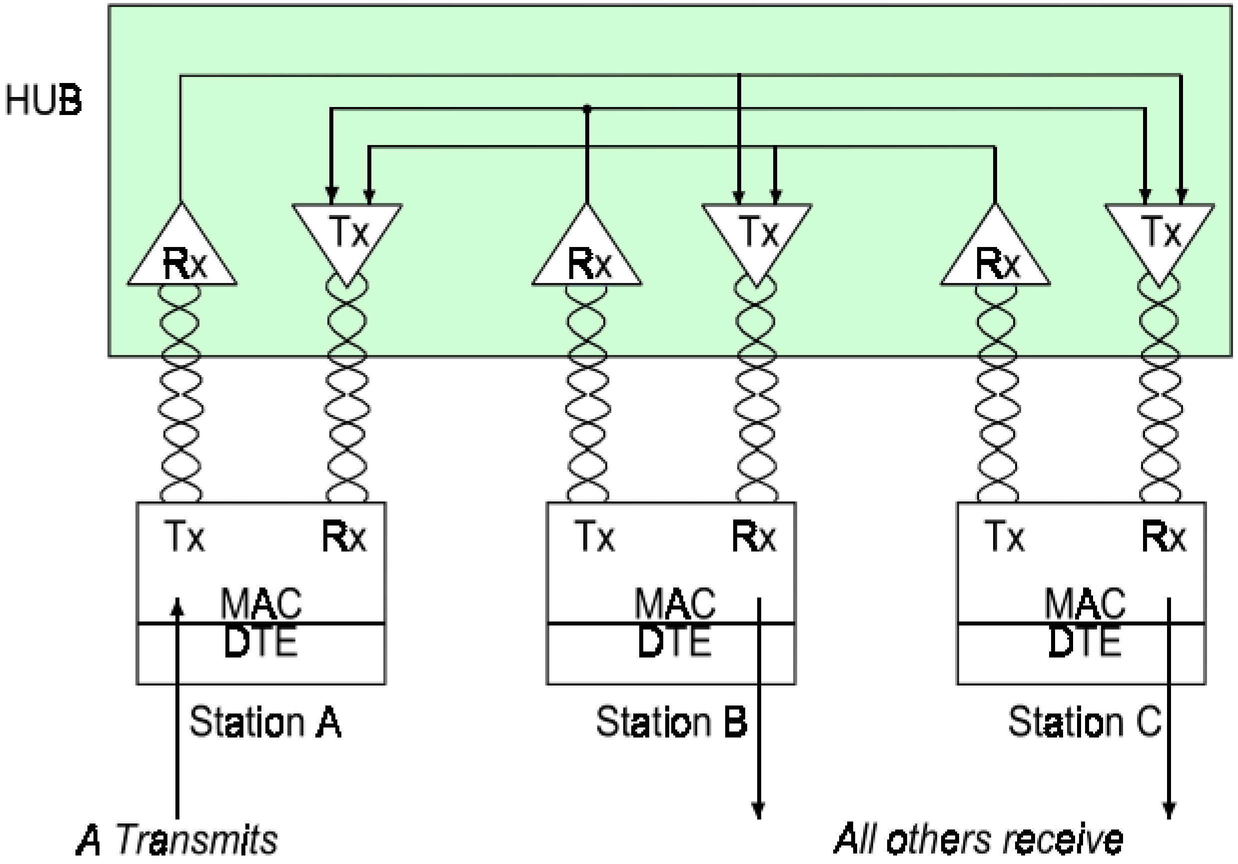

10BaseF systems

10BaseFL

- The fiber link segment standard of 2 km is upgraded to the existing fiber optic inter repeater link (FOIRL) standard.
- The original FOIRL as specified in the 802.3 standard was limited to a 1 km fiber link between two repeaters, with a maximum length of 2.5 km if there are 5 segments in the link.

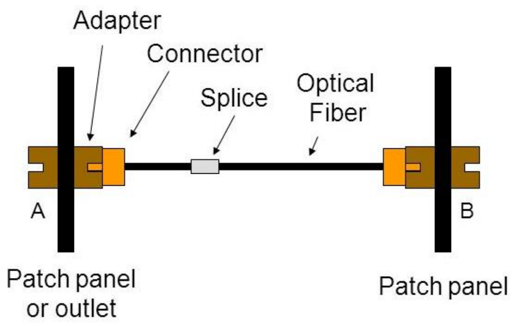

10BaseFP

- A star topology network based on the use of a passive fiber optic star coupler
- Up to 33 ports are available per star, and each segment has a maximum length of 500 m.
- The passive hub is completely immune to external noise and is an excellent choice for noisy industrial environments

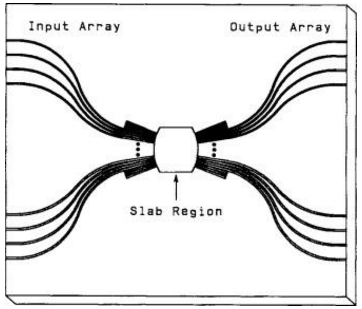

10BaseFB

- A fiber backbone link segment in which data is transmitted synchronously.
- Designed only for connecting repeaters, and for repeaters to use this standard, they must include a built in transceiver.
- This reduces the time taken to transfer a frame across the repeater hub.

10Broad36 systems

- Broadband version of Ethernet, and uses a 75-ohm coaxial cable for transmission.
- Each transceiver transmits on one frequency and receives on a separate one.
- The Tx/Rx streams require a 14 MHz bandwidth and an additional 4 MHz is required for collision detection and reporting - total bandwidth requirement is thus 36 MHz.

Ethernet design rules

Length of the cable segment

- Each segment has a particular maximum length allowable.
- For example, 10Base2 allows 200 m maximum length. The recommended maximum length is 80% of this.
- Multiple lengths joined by coaxial connectors
- Thicknet (10Base5) and Thinnet (10Base2) cables have the same nominal 50-ohm impedance, but they can only be mixed within the same 10Base2 cable segment to achieve greater segment length

System	Maximum	Recommended
10Base5	500m	400m
10Base2	185m	150m
10BaseT	100m	80m
1Base5	500m	400m

Ethernet design rules

Node placement rules

- Connection of the transceiver media access units (MAU) to the cable causes signal reflections due to their bridging impedance.
- In 10Base5 systems the MAUs are spaced at 2.5 m multiples, coinciding with the cable markings.
- In 10Base2 systems the minimum MAU spacing is 0.5 m.

Maximum transmission path

- Made of five segments connected by four repeaters.
- The total number of segments can be made up of a maximum of three coax segments containing station nodes and two link segments, having no intermediate nodes : 5-4-3-2 rule.

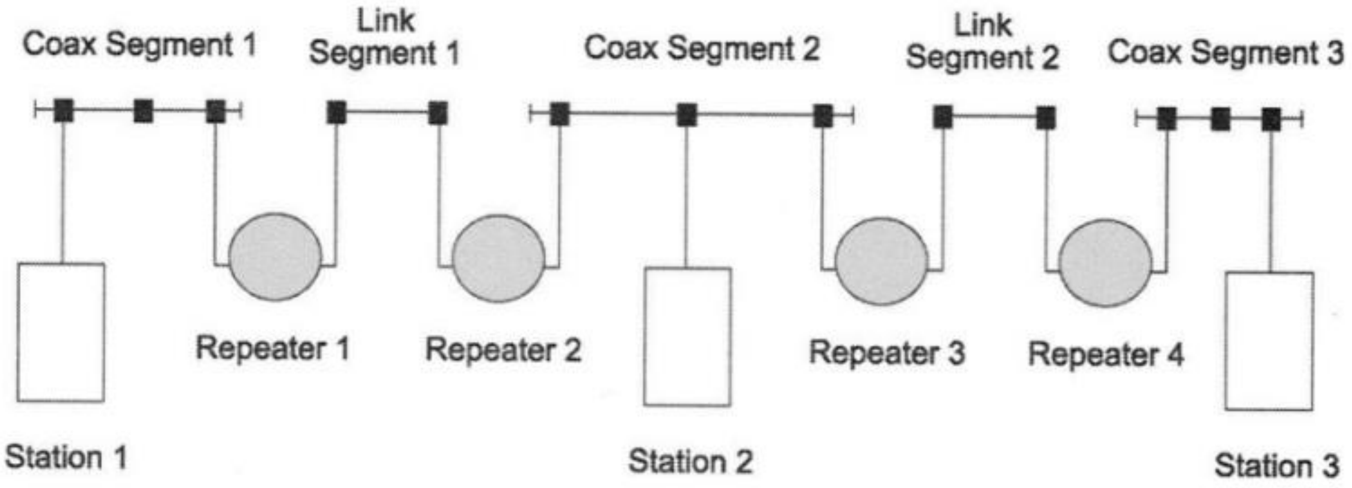

5 segments	OR	5 segments
4 repeaters		4 repeaters
3 coax segments		3 link segments
2 link segments		2 coax segments

Ethernet design rules

Maximum network size

- 10Base5 = 2800 m node to node

  (5 X 500 m segments + 4 repeater cables + 2 AUI)
- 10Base2 = 925 m node to node (5 x 185 m segments)
- 10BaseT = 100 m node to hub

Repeater rules

- Repeaters are connected to transceivers that count as one node on the segments.
- Special transceivers are used to connect repeaters and these do not implement the signal quality error test (SQE).

Cable system grounding

- Safety and noise connotations.
- Single point earth reference is usually located at one of the terminators.

Fast Ethernet systems

- 10 Mbps speed is too slow for very data intensive or real-time applications.

- Ways to increase speed on a network
  - To increase the bandwidth and allow faster changes of the data signal - requires a high bandwidth medium and generates a considerable amount of high frequency electrical noise on copper cables, which is difficult to suppress
  - To move away from the serial transmission of data on one circuit to a parallel method of transmitting over multiple circuits at each instant.
  - To use data compression techniques to enable more than one bit to be transferred for each electrical transition.

  - With 1000 Mbps gigabit Ethernet - to operate circuits full-duplex, enabling simultaneous transmission in both directions

100 Mbps
fast Ethernet

1000 Mbps
gigabit Ethernet

100Base-T

- 100 Mbps transmission - uses the existing Ethernet MAC layer with various enhanced physical media dependent (PMD) layers to improve the speed.
- IEEE 802.3u defines three different versions based on the physical media
  - 100Base-TX uses two pairs of category 5 UTP or STP
  - 100Base-T4 uses four pairs of wires of category 3, 4 or 5 UTP
  - 100Base-FX uses multimode or single-mode fiber optic cable
- IEEE 802.3y - 100Base-T2 which uses two pairs of wires of category 3, 4 or 5 UTP

- The 10 Mbps Ethernet interframe gap is defined as an absolute time interval of 9.60 microseconds, equivalent to 96 bit times; while the 100 Mbps system reduces this by ten times to 960 nanoseconds.

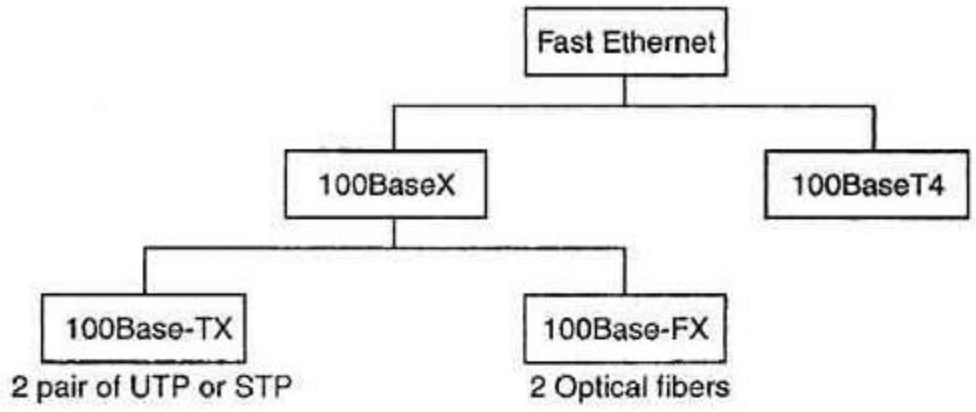

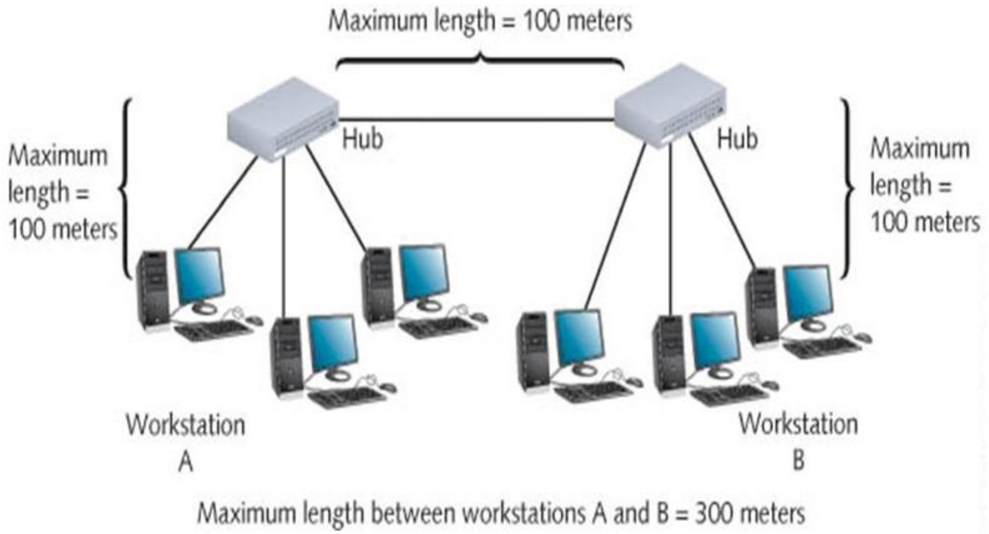

100Base-T

100Base-TX and FX

- Operates on two pairs of category 5 twisted pair or two multimode fibers.
- Uses stream cipher scrambling for data security and multilevel threshold-3 (MLT-3) bit encoding.
- The MLT-3 bit coding uses three voltage levels: +1 volts, 0 volts and –1 volts.
- The level remains the same for consecutive sequences of the same bit, i.e. continuous ‘1s’. When a bit changes, the voltage level changes to the next state in the circular sequence 0V, +1V, 0V, –1V, 0V etc - results in a coded signal resembling a smooth sine wave of much lower frequency than the incoming bitstream.
- For a 31.25 MHz baseband signal this allows for a 125 Mbps signaling bit stream providing a 100 Mbps throughput (4B/5B encoder).
- Two pair wire, RJ-45 connectors and a hub are requirements for 100BASE-TX

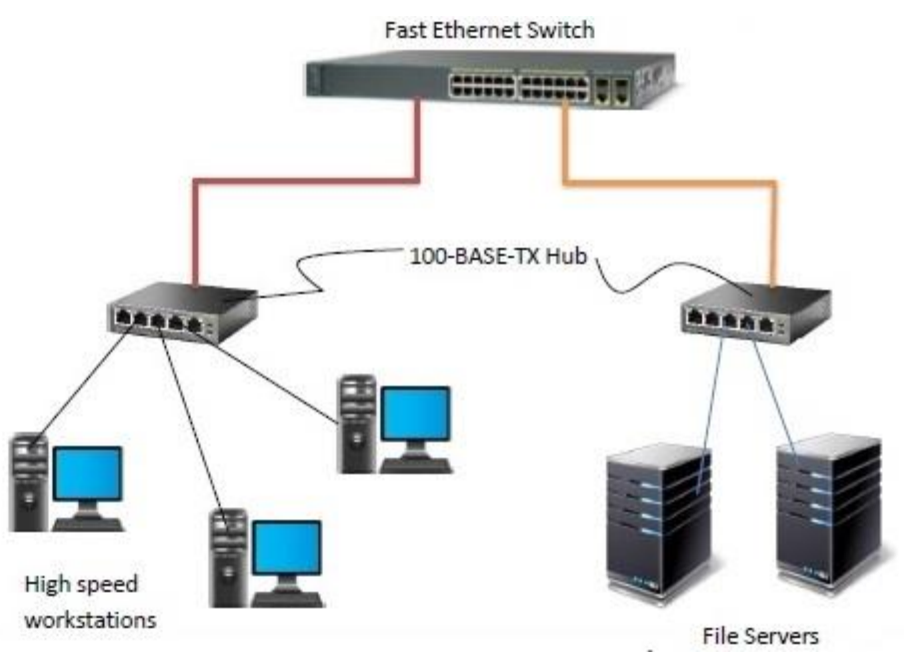

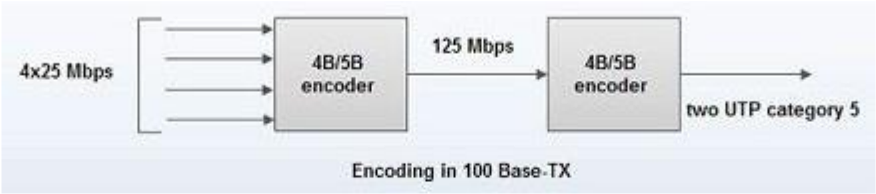

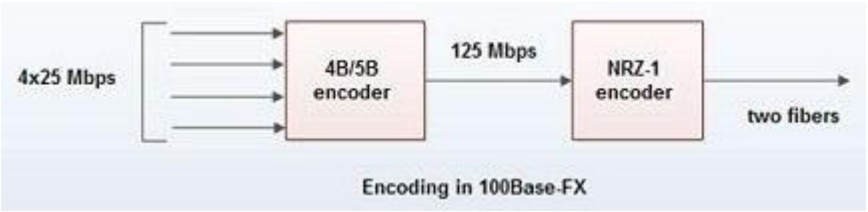

100Base-T

100Base-T4

- Uses four pairs of category 3 unshielded twisted pair (UTP) - two of the four pairs are bi-directional, the other two are unidirectional
- In each direction, three pairs are used at the same time to carry data, one pair switches between sending and receiving
- Uses data encoded in an eight binary six ternary (8B/6T) coding scheme similar to the MLT-3 code.
- The data is encoded using three voltage levels per bit time of +V, 0 volts and –V, these are usually written as simply +, 0 and –.
- This coding scheme allows the eight bits of binary data to be coded into six ternary symbols and reduces the required bandwidth to 25 MHz.

  (100 x (6/8) /3 = 25)
- The 256 codewords are chosen so the line has a mean line signal of zero - helps the receiver to discriminate the positive and negative signals relative to the average zero level.

Using four wires in 100Base-T4

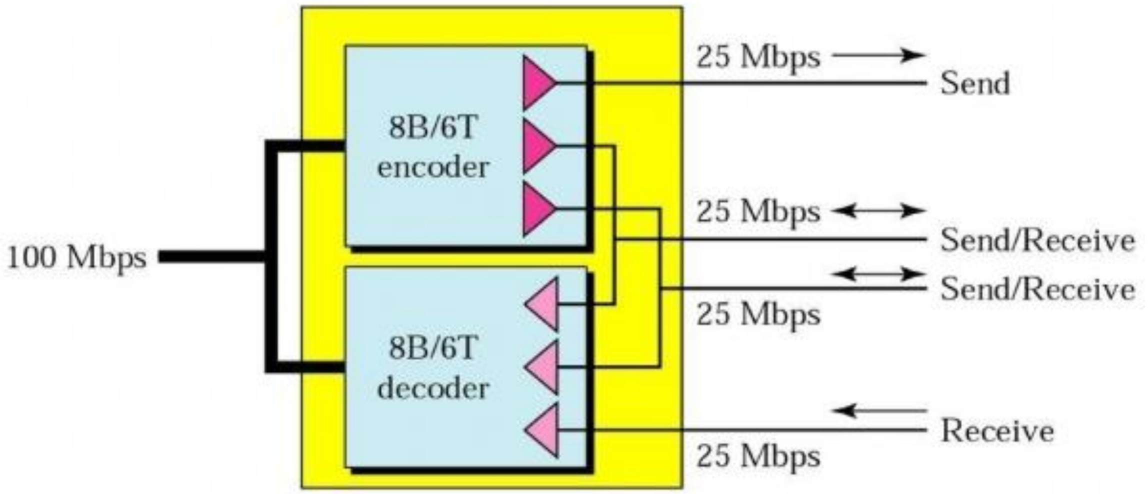

100BASE-TX Repeaters

Class I repeater

- Allowed to have larger timing delays, and operates by translating line signals on an incoming port to digital form, and then retranslating them to line signals when sending them out on the other ports.
- This makes it possible to repeat signals between media segments that use different signaling techniques, such as 100BASE-TX/FX segments and 100BASE-T4 segments, allowing these segment types to be mixed within a single repeater hub.
- The translation process in Class I repeaters uses up a number of bit times, so that only one Class I repeater can be used in a given collision domain when maximum cable lengths are used.

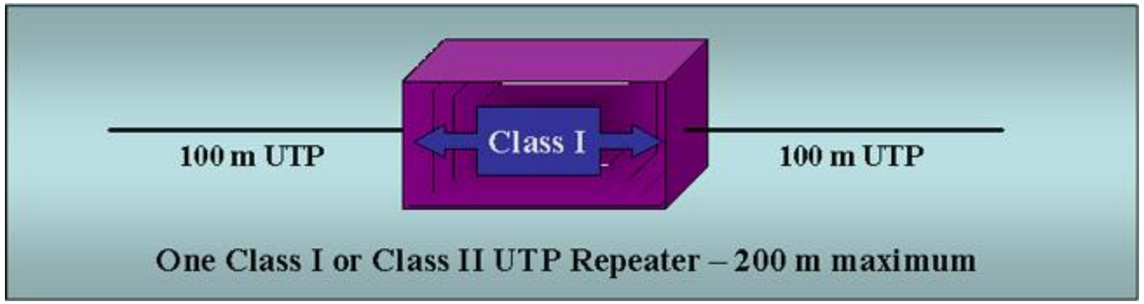

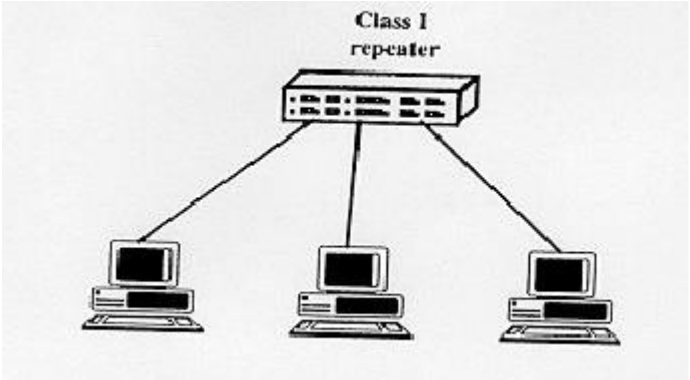

100BASE-TX Repeaters

Class II repeater

- Restricted to smaller timing delays, and immediately repeats the incoming signal to all other ports without a translation process.
- To achieve the smaller timing delay, Class II repeaters connect only to segment types that use the same signaling technique, such as 100BASE-TX and 100BASE-FX segments.
- A maximum of two Class II repeaters can be used within a given collision domain when maximum cable lengths are used.
- Segment types with different signaling techniques (e.g. 100BASE-TX/FX and 100BASE-T4) typically cannot be mixed together in a Class II repeater hub.

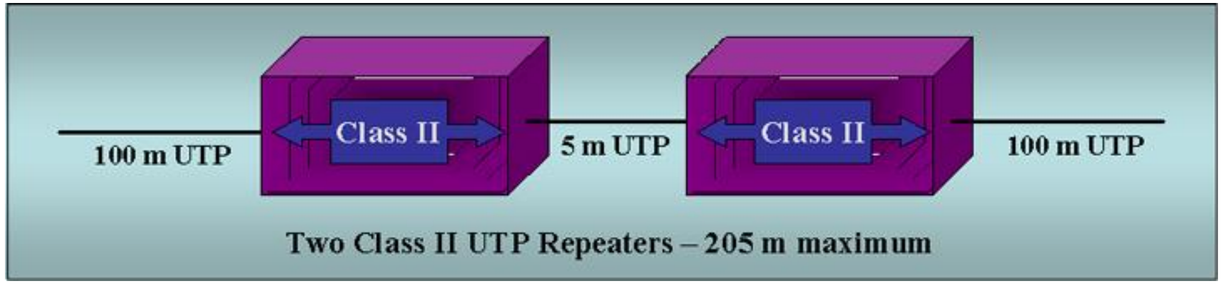

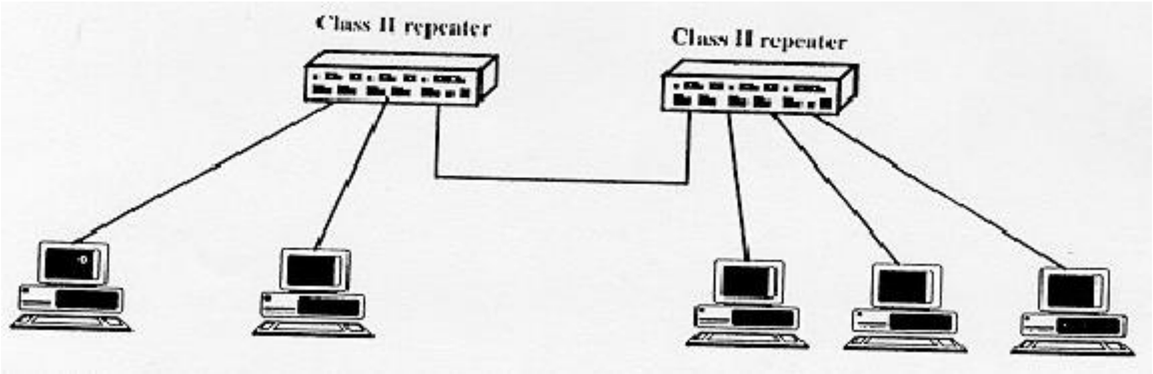

Fast Ethernet design considerations

UTP cabling distances 100Base-TX/T4

- Maximum distance between UTP hub and desktop NIC is 100 meters, made up as follows:
  - 5 meters from hub to patch panel
  - 90 meters horizontal cabling from patch panel to office punch-down block
  - 5 meters from punch-down block to desktop NIC

Fiber optic cable distances 100Base-FX

- Node to hub: maximum distance of multimode cable (62.5/125) is 160 meters (for connections using a single Class II repeater).
- Node to switch: maximum multimode cable distance is 210 meters.
- Switch-to-switch: maximum distance of multimode cable for a backbone connection between two 100Base-FX switch ports is 412 meters.
- Switch-to-switch, full-duplex: maximum distance of multimode cable for a full-duplex connection between two 100Base-FX switch ports is 2000 meters

Fast Ethernet design considerations

100Base-T repeater rules

- The cable distance and the number of repeaters, used in a 100Base-T collision domain, depends on the delay in the cable and the time delay in the repeaters and NIC delays.
- The maximum round-trip delay for 100Base-T systems is the time to transmit 64 bytes or 512 bits and equals 5.12 μs.
- A frame has to go from the transmitter to the most remote node then back to the transmitter for collision detection within this round trip time - the one-way time delay will be half this.
- The maximum sized collision domain can then be determined as:

Repeater delays + Cable delays + NIC delays + Safety factor (5 bits minimum)< 2.56 μs

Component	Maximum Delay (μs)
Fast Ethernet NIC	0.25
Fast Ethernet Switch Port	0.25
Class I Repeater	0.7 max
Class II Repeater	0.46 max
UTP Cable (per 100 meters)	0.55
Multimode Fiber (per 100 meters)	0.50

Fast Ethernet design considerations

Sample Calculation

- Can two fast Ethernet nodes be connected together using two class II repeaters connected by 50 m fiber? One node is connected to the first repeater with 50 m UTP while the other has a 100 m fiber connection.

Component	Maximum Delay (μs)
Fast Ethernet NIC	0.25
Fast Ethernet Switch Port	0.25
Class I Repeater	0.7 max
Class II Repeater	0.46 max
UTP Cable (per 100 meters)	0.55
Multimode Fiber (per 100 meters)	0.50

NIC	0.25μs
50 m UTP	0.275μs
Repeater Class II	0.46μs
50 m fiber	0.25μs
Repeater Class II	0.46μs
100 m fiber	0.50μs
NIC	0.25μs
TOTAL DELAY	2.445μs

- The total one-way delay of 2.445 μs is within the required interval (2.56 μs) and allows at least 5 bits safety factor, so this connection is permissible.

Gigabit Ethernet 1000Base-T

- Uses the same 802.3 frame format as 10 Mbps and 100 Mbps Ethernet systems
- Operates at ten times the clock speed of fast Ethernet at 1 Gbps.
- Defined by the IEEE 802.3z standard.
- Three different physical layers: 1000Base-LX and 1000Base-SX using fiber and 1000Base-CX using copper.
- Used 8B/10B encoding to reduce the bandwidth required to send high speed signals.

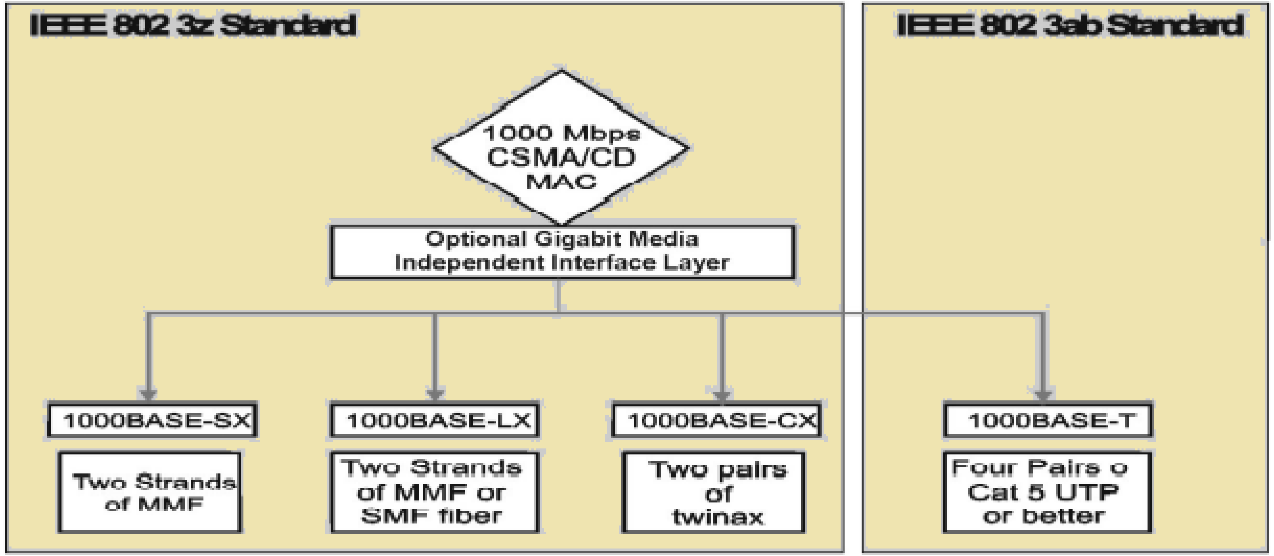

Gigabit Ethernet MAC layer

- Retains the standard 802.3 frame format, CSMA/CD algorithm is modified to enable it to function effectively at 1 Gbps.
- The slot time of 64 bytes, used with both 10 Mbps and 100 Mbps systems, has been increased to 512 bytes.
- Slot time defines the time during which the transmitting node retains control of the medium, and in particular is responsible for collision detection.
- With gigabit Ethernet it was necessary to increase this time by a factor of eight to 4.096 µs to compensate for the tenfold speed increase - gives a collision domain of about 200 m.
- If the transmitted frame is less than 512 bytes the transmitter continues transmitting to fill the 512 byte window.
- A carrier extension symbol is used to mark frames, which are shorter than 512 bytes and to fill the remainder of the frame.

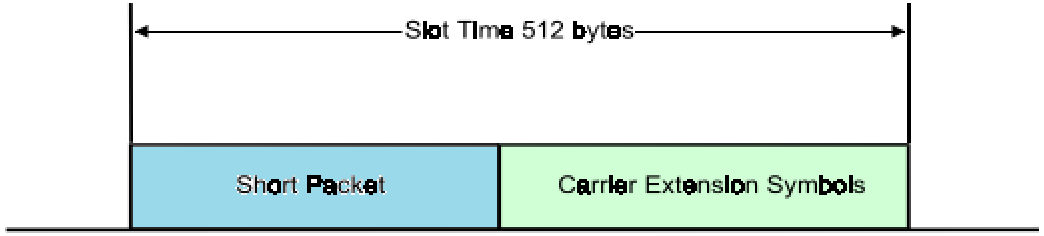

Gigabit Ethernet MAC layer

- Very low utilization if many short frames are sent.
- A 64 byte frame would have 448 carrier extension symbols attached and result in a utilization of less than 10%.
- This effect can be minimized if many small frames are sent by a technique called packet bursting.
- Once the first frame in a burst has successfully passed through the 512 byte collision window, using carrier extension if necessary, transmission continues with additional frames being added to the burst until the burst limit of 1500 bytes is reached.
- This process averages the time wasted sending carrier extension symbols over a number of frames.
- Frames are added to the burst in real-time with carrier extension symbols filling the interpacket gap.
- The total number of bytes sent in the burst is totaled after each frame and transmission continues until at least 1500 bytes have been transmitted.

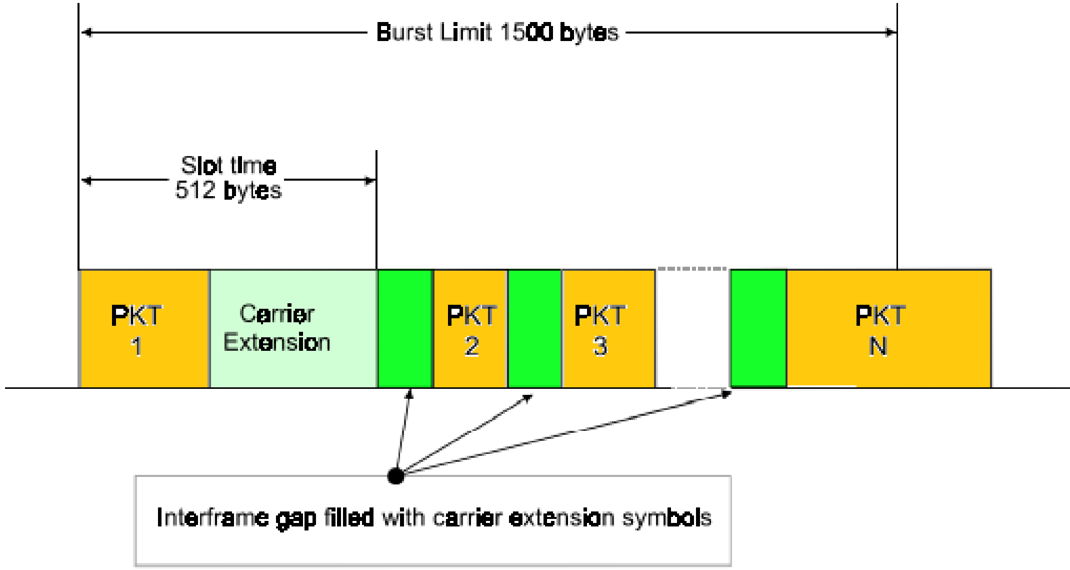

Gigabit Ethernet MAC layer

- Retains the standard 802.3 frame format, CSMA/CD algorithm is modified to enable it to function effectively at 1 Gbps.
- The slot time of 64 bytes, used with both 10 Mbps and 100 Mbps systems, has been increased to 512 bytes.
- Slot time defines the time during which the transmitting node retains control of the medium, and in particular is responsible for collision detection.
- With gigabit Ethernet it was necessary to increase this time by a factor of eight to 4.096 µs to compensate for the tenfold speed increase - gives a collision domain of about 200 m.
- If the transmitted frame is less than 512 bytes the transmitter continues transmitting to fill the 512 byte window.
- A carrier extension symbol is used to mark frames, which are shorter than 512 bytes and to fill the remainder of the frame.

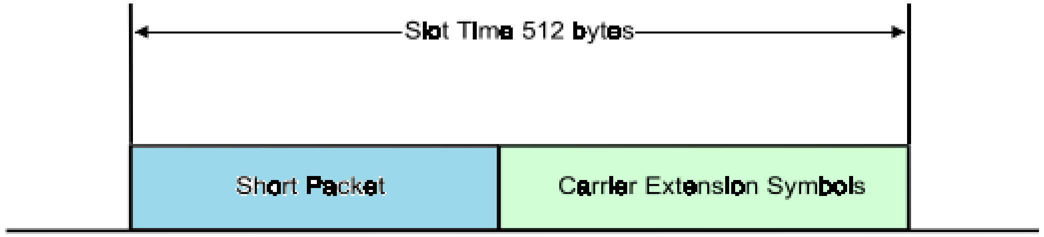

Gigabit Ethernet Physical layer

1000Base-SX for horizontal fiber (Short wavelength)

- Developed for the short backbone connections of the horizontal network wiring.
- The SX systems operate full duplex with multimode fiber only, using the cheaper 850 nm wavelength laser diodes.
- The maximum distance supported varies between 200 and 550 meters depending on the bandwidth and attenuation of the fiber optic cable used.

1000Base-LX for vertical backbone cabling (Long wavelength)

- Developed for use in the longer backbone connections of the vertical network wiring.
- The LX systems can use single mode or multimode fiber with the more expensive 1300 nm laser diodes.
- The maximum distances recommended by the IEEE for these systems operating in full-duplex is 5 kilometers for single mode cable and 550 meters for multimode fiber cable.
- Many 1000Base-LX vendors guarantee their products over much greater distances, typically 10 km.
- Fiber extenders are available to give service over as much as 80 km

Gigabit Ethernet Physical layer

1000Base-CX for copper cabling

- Developed for the short interconnection of switches, hubs or routers within a wiring closet.
- It is designed for 150 ohm twinax cable similar to that used for IBM token ring systems.
- The IEEE specified two types of connectors:
  - High-speed serial data connector (HSSDC) known as the fiber channel style 2 connector
  - Nine pin D-subminiature connector
- The maximum cable length is 25 meters for both full- and half-duplex systems.

1000Base-T for category 5 UTP

- Developed under the IEEE 802.3ab standard for transmission over four pairs of category 5 or better cable.
- This is achieved by simultaneously sending and receiving over each of the four pairs (existing 100Base-TX system has individual pairs for transmitting and receiving.
- Has less noise immunity - digital signal processors (DSP) associated with each pair overcomes any problems

Gigabit Ethernet full-duplex repeaters

- Gigabit Ethernet nodes are connected to full-duplex repeaters also known as non-buffered switches or buffered distributors.
- These devices have a basic MAC function in each port, which enables them to verify that a complete frame is received and compute its frame check sequence (CRC) to verify the frame validity.
- Then the frame is buffered in the internal memory of the port before being forwarded to the other ports of the repeater - combining the functions of a repeater with some features of a switch.
- All ports on the repeater operate at the same speed of 1 Gbps, and operate in full duplex so it can simultaneously send and receive from any port.
- The repeater uses 802.3x flow control to ensure the small internal buffers associated with each port do not overflow - When the buffers are filled to a critical level, the repeater tells the transmitting node to stop sending until the buffers have been sufficiently emptied.

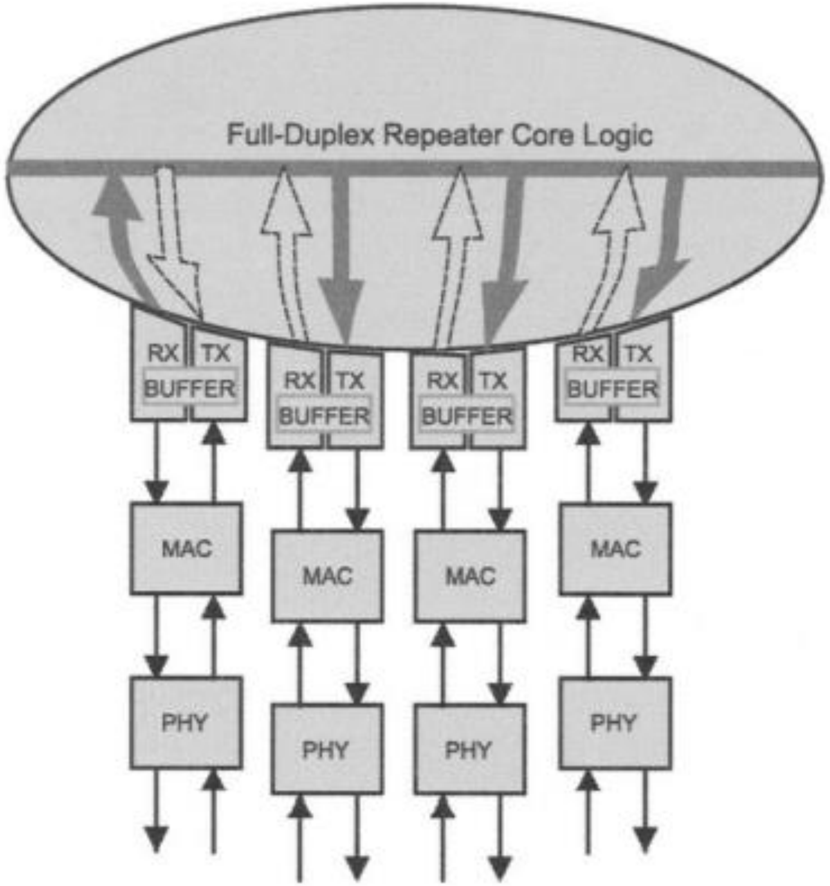

PHY - Physical medium independent

Gigabit Ethernet design considerations

Fiber optic cable distances

- Maximum cable distances which can be used between the node and a full-duplex repeater depend mainly on the chosen wavelength, the type of cable, and its bandwidth.
- Maximum distances for full-duplex 1000Base-X repeaters

Wavelength (nm)	Cable Type	Bandwidth (MHz. km)	Attenuation (dB/km)	Maximum distance (m)
850	50/125 Multimode	400	3.25	500
850	50/125 Multimode	500	3.43	550
850	62.5/125 Multimode	160	160	220
850	62.5/125 Multimode	200	200	275
1300	50/125 Multimode	500	2.32	550
1300	62.5/125 Multimode	500	1.0	550
1300	9/125 Single mode	Infinite	0.4	5000

Gigabit Ethernet design considerations

Gigabit repeater rules

- The cable distance and the number of repeaters, used in a half-duplex 1000Base-T collision domain, depend on the delay in the cable and the time delay in the repeaters and NIC delays.
- The maximum round-trip delay for 1000Base-T systems is the time to transmit 512 bytes or 4096 bits and equals 4.096 μs.
- A frame has to go from the transmitter to the most remote node then back to the transmitter for collision detection within this round trip time - one-way time delay will be half this.
- The maximum sized collision domain can then be determined by the following calculation:

Repeater delays + Cable delays + NIC delays + Safety factor (5 bits minimum) <2.048 μs

System	Maximum collision diameter point-to-point Half-duplex	Maximum collision diameter One repeater segment
1000Base-CX	25 m	50m
1000Base-T	100 m	200m
1000Basc-SX or LX	316 m	220 m

Summary

Format	Data rate	Max segment length *	Max nodes per segment	Topology	Media	Connectors	Encoding	Notes
10BASE-T	10 Mbps half duplex 20 Mbps full duplex	100 m Max network length = 100 m node to hub	2	Star	Category 3, 4, or 5 UTP cable with two pairs of voice-grade/ telephone twisted pair, 100 ohms	8-pin RJ-45 style modular jack; industrial variants include M18, M12, and DB9	Manchester	Most popular 10-m format
10BASE2 "Thinnet" or "Cheapernet"	10 Mbps half duplex only	185 m Max network length = 925 m = 5 × 185 m	30	Bus with drops. Minimum spacing between nodes = 0.5 m, max drop length = 4 cm	5-mm "thin" coax, e.g., RG58A/U or RG58C/U, Belden 9907 (PVC), and 89907 (ple-num); 50 ohms	BNC "T" coax connectors, bar-rel connectors, and terminators	Manchester	5-cm min bend radius; may not be used as link between 10BASE5 sys-tems
10BASE5 "Thicknet"	10 Mbps half duplex only	500-m (50-m max AUI length) Max network length = 2800 m = 5 × 500 m segments + 4 repeater cables + 2 AUI cables	100	Bus with drops	10-mm ("thick") coax, e.g., Belden 9880 (PVC) and 89880 (ple-num); bend radius min 25 cm; 50 ohm media and ter-mination	N-type coaxial connectors, bar-rel-style insula-tion displacement connectors and terminators	Manchester	MAU links trunk to NIC via AUI cable; taps must be spaced at 2.5-m intervals; ground at one end of cable**

Summary

Format	Data rate	Max segment length *	Max nodes per segment	Topology	Media	Connectors	Encoding	Notes
10BROAD36	10 Mbps half duplex only	1800-m single segment; 3600 m total for multiple segments			75-ohm CATV broadband cable		Modulated RF	Dead
10BASE-FL	10 Mbps half duplex; 20 Mbps full duplex	2000 m	2	Star	2 MMF cables, RX and TX, typically 62.5/125 fiber, 850-nm wavelength	BFOC/2.5, also called “ST”	Manchester	Uncommon
100BASE-TX	100 Mbps half duplex; 200 Mbps full duplex	100 m	2	Star	2 pairs of Category 5 UTP cabling; 100-ohm impedance (optionally supports 150-ohm STP)	RJ-45 style modular jack (8 pins) for UTP cabling (optionally supports 9-pin D-shell connector for STP cabling)	4B/5B	Most popular 100-m format IEEE 802.3u
100BASE-FX	100 Mbps half duplex; 200 Mbps full duplex	Half duplex: 412 m; full duplex: 2000 m	2	Star	2 MMF optical channels, one for TX, one for RX. Typ. 62.5/125 MMF, 1300-nm wavelength	Duplex SC, ST, or FDDI MIC connectors	4B/5B	IEEE 802.3u

Summary

Format	Data rate	Max segment length *	Max nodes per segment	Topology	Media	Connectors	Encoding	Notes
100BASE-T4	100 Mbps half duplex only	100 m	2	Star	Category 3, 4, or 5 UTP (uses 4 pairs or wires); 100 ohm	RJ-45 style modular jack (8 pins)	8B/6T	Uncommon IEEE 802.3u Useful where existing CAT3 telecom cables are available
1000BASE-LX	1000 Mbps half duplex; 2000 Mbps full duplex	Half-duplex MMF & SMF: 316 m; full-duplex MMF: 550 m; full-duplex SMF: 5000-m; 10-micron SMF: 3000-m max segment length	2	Star	2 62.5/125 or 50/125 multimode optical fibers (MMF), or 2 10-micron single-mode optical fibers (SMF), 1270- to 1355-nm light wavelength	duplex SC connector	8B/10B	803.z Multimode: longer-building backbones; Single mode: campus-wide backbones
1000BASE-SX	1000 Mbps half duplex; 2000 Mbps full duplex	Half-duplex 62.5/125: 275 m; half-duplex 50/125: 316 m; full-duplex 62.5/125: 275 m; full-duplex 50/125: 550 m	2	Star	2 62.5/125 or 50/125 MMF, 770 to 860 nm	duplex SC connector	8B/10B	
1000BASE SX	1000 Mbps half duplex; 2000 Mbps full duplex	Half duplex: 25 m; full duplex: 25 m 62.5/125 MMF full duplex: 260 m	2	Star	Specialty shielded balanced copper jumper cable ("twinax" or "short haul copper")	9-pin shielded D-subminiature connector, or 8-pin ANSI Fibre Channel Type 2 (HSSC) connector	8B/10B	802.3z Intended for short backbones

Summary

Format	Data rate	Max segment length *	Max nodes per segment	Topology	Media	Connectors	Encoding	Notes
1000BASE-T	1000 Mbps half duplex; 2000 Mbps full duplex	100 m (328 ft)	2	Star	4 pairs of CAT5 or better cabling, 100 ohms	8-pin RJ-45 con-nector	PAM5	802.3ab Replace exist-ing 10/ 100BASE-T runs in floors of buildings
1000BASE-CX		25 m	2	Star	STP copper Twinax, 150 ohms	DB9 or HSSC	8B/10B	802.3z Short jumper connection in computer rooms or switching clos-ets Common ground required on devices at both ends of the cable

AUI = Attachment Unit Interface; MAU= Medium Attachment Unit; MMF = Multimode fiber; NIC = Network interface card; SMF = Single-mode fiber; STP = Shielded Twisted Pair; UTP = Unshielded Twisted Pair; 62.5/125 means 62.5-micron fiber core with 125-micron outer cladding.

Reference

1. David Bailey Edwin Wright, “Practical SCADA for Industry”, 1st Edition, Elsevier, 2003.
2. Deon Reynders and Edwin Wright, “Practical TCP/IP and Ethernet Networking for Industry”, Elsevier, 2003.
3. Jan Axelson, “Embedded Ethernet and Internet Complete: Designing and Programming Small Devices for Networking”, Lakeview Research, 2003.
4. Perry S. Marshall, John S. Rinaldi, “Industrial Ethernet - How to Plan, Install, and Maintain TCP/IP Ethernet Networks: The Basic Reference Guide for Automation and Process Control Engineers”, The Instrumentation, Systems, and Automation Society (ISA), Second Edition, 2004.