Introduction to Ethernet

1. Introduction

Ethernet was originally developed by Digital, Intel and Xerox (DIX) in the early 1970's and has been designed as a 'broadcast' system, i.e. stations on the network can send messages whenever and wherever it wants. All stations may receive the messages, however only the specific station to which the message is directed will respond.

The original format for Ethernet was developed in Xerox Palo Alto Research Centre (PARC), California in 1972. Using Carrier Sense Multiple Access with Collision Detection (CSMA/CD) it had a transmission rate of 2.94Mb/s and could support 256 devices over cable stretching for 1km. The two inventors were Robert Metcalf and David Boggs.

Ethernet versions 1.0 and 2.0 followed until the IEEE 802.3 committee re-jigged the Ethernet II packet to form the Ethernet 802.3 packet. (IEEE's Project 802 was named after the time it was set up, February 1980. It includes 12 committees 802.1 to 802.12, 802.2 is the LLC, 802.4 Token Bus, 802.11 Wireless, 802.12 100VG-AnyLAN etc.) Nowadays you will see either Ethernet II (DIX) (invented by **D**igital, **I**ntel and **X**erox) format or Ethernet 802.3 format being used.

The 'Ether' part of Ethernet denotes that the system is not meant to be restricted for use on only one medium type, copper cables, fibre cables and even radio waves can be used.

802.3 Ethernet uses **Manchester Phase Encoding (MPE)** for coding the data bits on the outgoing signal. The next few sections describe how Ethernet works and how Ethernet is structured.

2. CSMA/CD

As mentioned earlier, Ethernet uses Carrier Sense Multiple Access with Collision Detection (CSMA/CD). When an Ethernet station is ready to transmit, it checks for the presence of a signal on the cable i.e. a voltage indicating that another station is transmitting. If no signal is present then the station begins transmission, however if a signal is already present then the station delays transmission until the cable is not in use. If two stations detect an idle cable and at the same time transmit data, then a collision occurs. On a star-wired UTP network, if the transceiver of the sending station detects activity on both its receive and transmit pairs before it has completed transmitting, then it decides that a collision has occurred. On a coaxial system, a collision is detected when the DC signal level on the cable is the same or greater than the combined signal level of the two transmitters, i.e.. significantly greater than +/- 0.85v. Line voltage drops dramatically if two stations transmit at the same and the first station to notice this sends a high voltage jamming signal around the network as a signal. The two stations involved with the collision lay off transmitting again for a time interval which is randomly selected. This is determined using **Binary Exponential Backoff**. If the collision occurs again then the time interval is doubled, if it happens more than 16 times then an error is reported.

A **Collision Domain** is that part of the network where each station can 'see' other stations' traffic both unicast and broadcasts. The Collision Domain is made up of one segment of Ethernet coax (with or without repeaters) or a number of UTP shared hubs. A network is segmented with bridges (or microsegmented when using switches) that create two segments, or two Collision Domains where a station on one segment can not see traffic between stations on the other segment unless the packets are destined for itself. It can however still see all broadcasts as a segmented network, no matter the number of segments, is still one **Broadcast Domain**. Separate Broadcast Domains are created by VLANs on switches so that one physical network can behave as a number of entirely separate LANs such that the only way to allow stations on different VLANs to communicate is at a layer 3 level using a router, just as if the networks were entirely physically separate.

3. Ethernet Frame

3.1 Frame Formats

The diagrams below describe the structure of the older DIX (Ethernet II) and the now standard 802.3 Ethernet frames. The numbers above each field represent the number of bytes.

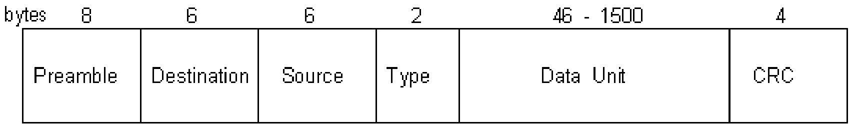

DIX Ethernet Packet

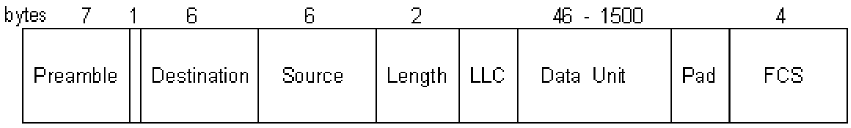

IEEE 802.3 Frame

- **Preamble field:** Establishes bit synchronisation and transceiver conditions so that the PLS circuitry synchs in with the received frame timing. The DIX frame has 8 bytes for the preamble rather than 7, as it does not have a Start Frame Delimiter (or Start of Frame).
- **Start Frame Delimiter:** Sequence 10101011 in a separate field, only in the 802.3 frame.
- **Destination address:** Hardware address (MAC address) of the destination station (usually 48 bits i.e. 6 bytes).
- **Source address:** Hardware address of the source station (must be of the same length as the destination address, the 802.3 standard allows for 2 or 6 byte addresses, although 2 byte addresses are never used, N.B. Ethernet II can *only* uses 6 byte addresses).
- **Type:** Specifies the protocol sending the packet such as IP or IPX (only applies to DIX frame).
- **Length:** Specifies the length of the data segment, actually the number of LLC data bytes, (only applies to 802.3 frame and replaces the Type field).
- **Pad:** Zeros added to the data field to 'Pad out' a short data field to 46 bytes (only applies to 802.3 frame).
- **Data:** Actual data which is allowed anywhere between 46 to 1500 bytes within one frame.
- **CRC:** Cyclic Redundancy Check to detect errors that occur during transmission (DIX version of FCS).
- **FCS:** Frame Check Sequence to detect errors that occur during transmission (802.3 version of CRC). This 32 bit code has an algorithm applied to it which will give the same result as the other end of the link, provided that the frame was transmitted successfully.

From the above we can deduce that the maximum 802.3 frame size is 1518 bytes and the minimum size is 64 bytes. Packets that have correct CRC's (or FCS's) but are smaller than 64 bytes, are known as 'Runts'.

The hardware address, or MAC address is transmitted and stored in Ethernet network devices in **Canonical** format i.e. Least significant Bit (LSB) first. You may hear the expression **Little-Endian** to describe the LSB format in which Ethernet is transmitted. Token Ring and FDDI, on the other hand, transmit the MAC address with the Most Significant Bit (MSB) first, or **Big-Endian**, This is known as **Non-Canonical** format. Note that this applies on a byte by byte basis i.e. the bytes are transmitted in the same order it is just the bits in each of those bytes that are reversed! The storage of the MAC addresses in Token Ring and FDDI devices however, may sometimes still be in Canonical format so this can sometimes cause confusion. The reference to, the distribution of MAC addresses and the OUI desinations are always carried out in Canonical format.

Some discussion is warranted on the LLC field. The 802.2 committee developed the **Logical Link Control (LLC)** to operate with 802.3 Ethernet as seen in the above diagram. LLC is based on the **HDLC** format and more detail can be found by following the link. Whereas Ethernet II (2.0) combines the MAC and the Data link layers restricting itself to connectionless service in the process, IEEE 802.3 separates out the MAC and Data Link layers. 802.2 (LLC) is also required by Token Ring and FDDI but cannot be used with the Novell 'Raw' format. There are three types of LLC, Type 1 which is connectionless, Type 2 which is connection-oriented and Type 3 for Acknowledged Connections.

8	8	8 or 16		bits
DSAP Address	SSAP Address	Control	Information	

The **Service Access Point (SAP)** is used to distinguish between different data exchanges on the same end station and basically replaces the Type field for the older Ethernet II frame. The **Source Service Access Point (SSAP)** indicates the service from which the LLC data unit is sent, and the **Destination Service Access Point (DSAP)** indicates the service to which the LLC data unit is being sent. As examples, NetBIOS uses the SAP address of **F0** whilst IP uses the SAP address of **06**. The following lists common SAPs:

- **00** - Null LSAP
- **02** - Individual LLC Sublayer Management Function
- **03** - Group LLC Sublayer Management Function
- **04** - IBM SNA Path Control (individual)
- **05** - IBM SNA Path Control (group)
- **06** - ARPANET Internet Protocol (IP)
- **08** - SNA
- **0C** - SNA
- **0E** - PROWAY (IEC955) Network Management & Initialization
- **14** - ICL OSLAN SSAP (TP4 over 802.3)
- **18** - Texas Instruments
- **42** - IEEE 802.1 Bridge Spanning Tree Protocol
- **4E** - EIA RS-511 Manufacturing Message Service
- **54** - ICL OSLAN DSAP (TP4 over 802.3)
- **7E** - ISO 8208 (X.25 over IEEE 802.2 Type 2 LLC)
- **80** - Xerox Network Systems (XNS)
- **86** - Nestar
- **8E** - PROWAY (IEC 955) Active Station List Maintenance
- **98** - ARPANET Address Resolution Protocol (ARP)
- **BC** - Banyan VINES
- **AA** - SubNetwork Access Protocol (SNAP)
- **E0** - Novell NetWare
- **F0** - IBM NetBIOS
- **F4** - IBM LAN Management (individual)
- **F5** - IBM LAN Management (group)

- **F8** - IBM Remote Program Load (RPL)
- **FA** - Ungermann-Bass
- **FE** - ISO Network Layer Protocol
- **FF** - Global LSAP

The Control Field identifies the type of LLC, of which there are three:

- **Type 1** - uses **Unsequenced Information (UI)** (Indicated by a Control Field value of **03**) frames to set up unacknowledged connectionless sessions.
- **Type 2** - uses **Information (I)** frames and maintains the sequence numbers during an acknowledged connection-oriented transmission.
- **Type 3** - uses **Acknowledged Connection (AC)** frames in an acknowledged connectionless service.

3.2 I/G and U/L within the MAC address

With an Ethernet MAC address, the first octet uses the lowest significant bit as the I/G bit (Individual/Group address) only and does not have such a thing as the U/L bit (Universally/Locally administered). The U/L bit is used in Token Ring A destination Ethernet MAC address starting with the octet '05' is a group or multicast address since the first bit (LSB) to be transmitted is on the right hand side of the octet and is a binary '1'. Conversely, '04' as the first octet indicates that the destination address is an individual address. Of course, in Ethernet, all source address will have a binary '0' since they are always individual.

The first 3 octets of the MAC address form the Organisational Unique Identifier (OUI) assigned to organisations that requires their own group of MAC addresses. A list of OUIs can be found at **OUI Index**.

3.3 Subnetwork Access Protocol (SNAP)

The SNAP protocol was introduced to allow an easy transition to the new LLC frame format for vendors. SNAP allows older frames and protocols to be encapsulated in a Type 1 LLC header so making any protocol 'pseudo-IEEE compliant'. SNAP is described in **RFC 1042**. The following diagram shows how it looks:

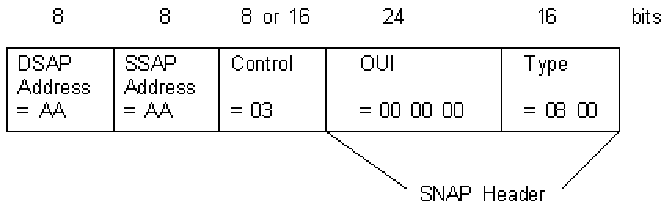

As you can see, it is an LLC data unit (sometimes called a **Logical Protocol Data Unit (LPDU)**) of Type 1 (indicated by 03). The DSAP and SSAP are set to **AA** to indicate that this is a SNAP header coming up. The SNAP header then indicates the vender via the **Organisational Unique Identifier (OUI)** and the protocol type via the Ethertype field. In the example above we have the OUI as 00-00-00 which means that there is an Ethernet frame, and the Ethertype of 08-00 which indicates IP as the protocol. The official list of types can be found at **Ethertypes**. More and more vendors are moving to LLC1 on the LAN but SNAP still remains and crops up time and time again.

Have a look at the document **IPX** for further discussion of 802.3 and 802.5 headers (SNAP etc.) in an IPX environment.

4. Media

4.1 10Base5

Traditionally, Ethernet is used over 'thick' coaxial cable (Normally yellow in colour) called 10Base5 (the '10' denotes 10Mbps, base means that the signal is baseband i.e. takes the whole bandwidth of the cable (so that only one device can transmit at one time on the same cable), and the '5' denotes 500m maximum length). The minimum length between stations is 2.5m.

The cable is run in one long length forming a 'Bus Topology'. Stations attach to it by way of inline N-type connections or a transceiver which is literally screwed into the cable (by way of a 'Vampire Tap') providing a 15-pin AUI (Attachment Unit Interface) connection (also known as a DIX connector or a DB-15 connector) for a drop lead connection (maximum of 50m length) to the station. The segments are terminated with 50 ohm resistors and the shield should be grounded at one end only.

5-4-3 Rule

The segment could be appended with up to a maximum of 4 repeaters, therefore 5 segments (total length of 2,460m) can be connected together. Of the 5 segments only 3 can have devices attached (100 per segment). A total of 300 devices can be attached on a Thicknet broadcast domain.

4.2 10Base2

It was common to see the Thick coax used in Risers to connect Repeaters which in turn provide 'Thin Ethernet' coaxial connections for runs around the floors to up to 30 workstations. Thin ethernet (Thinnet) uses RG-58 cable and is called 10Base2 (The '2' now denoting 200m maximum length, strictly speaking this is 185m). The minimum length between stations is 0.5m. Following is a table detailing various types of coaxial cable:

- RG-58 /U - solid copper core (0.66mm or 0.695mm), 53.5 ohms.
- RG-58 A/U - stranded copper core (0.66mm or 0.78mm), 50 ohms.
- RG-58 C/U - military version of RG58 A/U (0.66mm), 50 ohms.
- RG-59 - broadband transmissions e.g. cable TV.
- RG-6 - higher frequency broadband transmissions. A larger diameter than RG-59.
- RG-62 - ArcNet.
- RG-8 - Thicknet, 50 ohms.

Each station connects to the thinnet by way of a Network Interface Card (NIC) which provides a BNC (British Naval Connector). At each station the thinnet terminates at a T-piece and at each end of the thinnet run (or 'Segment') a 50-ohm terminator is required to absorb stray signals, thereby preventing signal bounce. The shield should be grounded at one end only.

A segment can be appended with other segments using up to 4 repeaters, i.e. 5 segments in total. 2 of these segments however, cannot be tapped, they can only be used for extending the length of the broadcast domain (to 925m). What this means is that 3 segments with a maximum of 30 stations on each can give you 90 devices on a Thinnet broadast domain.

(There is also a little used 10Broad36 standard where 10 Mbps Ethernet runs over broadband up to 3.6km. With broadband, a number of devices can transmit at the same time using multiple basebands e.g. multiple TV stations each with its own baseband signal frequency on one wire).

4.3 10BaseT

Nowadays, it is becoming increasingly important to use Ethernet across Unshielded Twisted Pair (UTP) or Shielded Twisted Pair (STP), this being called 10BaseT (the 'T' denoting twisted pair). For instance, Category 5 UTP is installed in a 'Star-wired' format, with runs recommended at no greater than 100m (including patch leads, cable run and flyleads) and Ethernet Hubs with UTP ports (RJ45) centrally located. It has been found though that runs of up to 150m are feasible, the limitations being signal strength. Also, there should be no more than a 11.5dB signal loss and the minimum distance between devices is 2.5m. The maximum delay for the signal in a 10Mbps network is 51.2 microseconds. This comes from the fact that the bit time (time to transmit one bit) is 0.1 microseconds and that the slot time for a frame is 512 bit times.

The wires used in the RJ45 are 1 and 2 for transmit, 3 and 6 for receive.

In order to connect to ethernet in this 'Star Topology', each station again has a NIC which, this time, contains an RJ45 socket which is used by a 4-pair RJ45 plug-ended droplead to connect to a nearby RJ45 floor or wall socket.

Each port on the hub sends a 'Link Beat Signal' which checks the integrity of the cable and devices attached, a flickering LED on the front of the port of the hub tells you that the link is running fine. The maximum number of hubs (or, more strictly speaking, repeater counts) that you can have in one segment is 4 and the maximum number of stations on one broadcast domain is 1024.

The advantages of the UTP/STP technology are gained from the flexibility of the system, with respect to moves, changes, fault finding, reliablity and security.

The following table shows the RJ45 pinouts for 10BaseT:

RJ45 Pin	Function	Colour
1	Transmit	White/Orange
2	Transmit	Orange/White
3	Receive	White/Green
4		Blue/White
5		White/Blue
6	Receive	Green/White
7		White/Brown
8		Brown/White

If you wish to connect hub to hub, or a NIC directly to another NIC, then the following 10BaseT cross-over cable should be used:

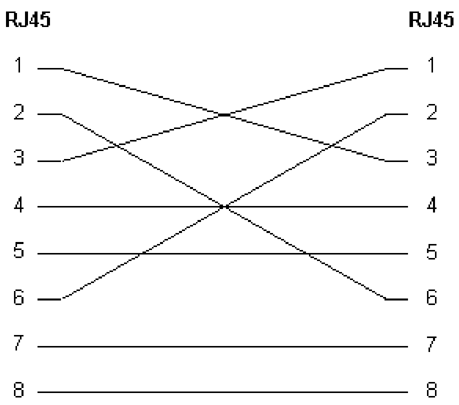

10BaseT Crossover

The 4 repeater limit manifests itself in 10/100BaseT environments where the active hub/switch port is in fact a repeater, hence the name multi-port repeater. Generally, the hub would only have one station per port but you can cascade hubs from one another up to the 4 repeater limit. The danger here of course, is that you will have all the traffic from a particular hub being fed into one port so care would need to be taken on noting the applications being used by the stations involved, and the likely bandwidth that the applications will use.

There is a semi-standard called Lattisnet (developed by Synoptics) which runs 10MHz Ethernet over twisted pair but instead of bit synchronisation occurring at the sending (as in 10BaseT) the synchronisation occurs at the receiving end.

4.4 10BaseF

The 10BaseF standard developed by the IEEE 802.3 committee defines the use of fibre for ethernet. 10BaseFB allows up to 2km per segment (on multi-mode fibre) and is designed for backbone applications such as cascading repeaters. 10BaseFL describes the standards for the fibre optic links between stations and repeaters, again allowing up to 2km per segment on multi-mode fibre. In addition, there is the 10BaseFP (Passive components) standard and the FOIRL (Fibre Optic Inter-Repeater Link) which provides the specification for a fibre optic MAU (Media Attachment Unit) and other interconnecting components.

The 10BaseF standard allows for 1024 devices per network.

4.5 Fast Ethernet (802.3u) 100BaseTx

Fast Ethernet uses the same frame formats and CSMA/CD technology as normal 10Mbps Ethernet. The difference is that the maximum delay for the signal across the segment is now 5.12 microseconds instead of 51.2 microseconds. This comes from the fact that the bit time (time to transmit one bit) is 0.01 microseconds and that the slot time for a frame is 512 bit times. The Inter-Packet Gap (IPG) for 802.3u is 0.96 microseconds as opposed to 9.6 microseconds for 10Mbps Ethernet.

Fast Ethernet is the most popular of the newer standards and is an extension to 10BaseT, using CSMA/CD. The '100' denotes 100Mbps data speed and it uses the same two pairs as 10BaseT (1 and 2 for transmit, 3 and 6 for receive) and must only be used on Category 5 UTP cable installations with provision for it to be used on Type 1 STP. The Copper physical layer being based on the **Twisted Pair-Physical Medium Dependent (TP-PMD)** developed by ANSI X3T9.5 committee. The actual data throughput increases by between 3 to 4 times that of 10BaseT.

Whereas 10BaseT uses **Normal Link Pulses (NLP)** for testing the integrity of the connection, 100BaseT uses **Fast Link Pulses (FLP)** which are backwardly compatible with NLPs but contain more information. FLPs are used to detect the speed of the network (e.g. in 10/100 switchable cards and ports).

The ten-fold increase in speed is achieved by reducing the time it takes to transmit a bit to a tenth that of 10BaseT. The **slot-time** is the time it takes to transmit 512 bits on 10Mbps Ethernet (i.e. 5.12 microseconds) and listen for a collision (see earlier). This remains the same for 100BaseT, but the network distance between nodes, or span, is reduced. The encoding used is **4B/5B** with **MLT-3** wave shaping plus **FSR**. This wave-shaping takes the clock frequency of 125MHz and reduces it to 31.25MHz which is the frequency of the carrier on the wire.

The round trip signal timing is the critical factor when it comes to the distance that the signal can run on copper UTP. The cable has to be Category 5 and the distance must not exceed 100m.

The IEEE use the term **100BaseX** to refer to both 100BaseTx and 100BaseFx and the **Media-Independent Interface (MII)** allows a generic connector for transceivers to connect to 100BaseTx, 100BaseFx and 100BaseT4 LANs.

There is no such thing as the 5-4-3 rule in Fast Ethernet. All 10Base-T repeaters are considered to be functionally identical. Fast Ethernet repeaters are divided into two classes of repeater, **Class I** and **Class II**. A Class I repeater has a repeater propagation delay value of 140 bit times, whilst a Class II repeater is 92 bit times. The Class I repeater (or **Translational Repeater**) can support different signalling types such as 100BaseTx and 100BaseT4. A Class I repeater transmits or repeats the incoming line signals on one port to the other ports by first translating them to digital signals and then retranslating them to line signals. The translations are necessary when connecting different physical media (media conforming to more than one physical layer specification) to the same collision domain. Any repeater with an MII port would be a Class I device. Only one Class I repeater can exist within a single collision domain, so this type of repeater cannot be cascaded. There is only allowed one Class I repeater hop in any one segment.

A Class II repeater immediately transmits or repeats the incoming line signals on one port to the other ports: it does not perform any translations. This repeater type connects identical media to the same collision domain (for example, TX to TX). At most, two Class II repeaters can exist within a single collision domain. The cable used to cascade the two devices is called and unpopulated segment or IRL (Inter-Repeater Link). The Class II repeater (or **Transparent Repeater**) can only support one type of physical signalling, however you can have two Class II repeater hops in any one segment (Collision Domain).

4.6 100BaseT4

100BaseT4 uses all four pairs and is designed to be used on Category 3 cable installations. Transmit is on pairs 1 and 2, receive is on pairs 3 and 6, whilst data is bidirectional on 4 and 5 and on 7 and 8. The signaling is on three pairs at 25MHz each using **8B/6T** encoding. The fourth pair is used for collision detection. Half-Duplex is supported on 100BaseT4.

4.7 100BaseFx

100BaseFx uses two cores of fibre (multi-mode 50/125um, 60/125um or single-mode) and 1300nm wavelength optics. The connectors are SC, Straight Tip (ST) or Media Independent Connector (MIC). The 100BaseT MAC mates with the ANSI X3T9.5 FDDI Physical Medium Dependent (PMD) specification. At half-duplex you can have distances up to 412m, whereas Full-duplex will give 2km.

There is also a proposed **100BaseSx** which uses 850nm wavelength optics giving 300m on multi-mode fibre.

The encoding used is **4B/5B** with **NRZ-I** wave shaping with a clock frequency of 125MHz.

4.8 100BaseT2

This little known version of Fast Ethernet is for use over two pairs of Category 3 cable and uses PAM-5 for encoding. There is simultaneous transmission and reception of data in both pairs and the electronics uses DSP technology to handle alien signals in adjacent pairs.

100BaseT2 can run up to 100m on Category 3 UTP.

4.9 100VG-AnyLAN

Based on 802.12 (Hewlett Packard), 100VG-AnyLAN uses an access method called **Demand Priority**. The 'VG' stands for 'Voice Grade' as it is designed to be used with Category 3 cable. This is where the repeaters (hubs) carry out continuous searches round all of the nodes for those that wish to send data. If two devices cause a 'contention' by wanting to send at the same time, the highest priority request is dealt with first, unless the priorities are the same, in which case both requests are dealt with at the same time (by alternating frames). The hub only knows about connected devices and other repeaters so communication is only directed at them rather than broadcast to every device in the broadcast domain (which could mean 100's of devices!). This is a more efficient use of the bandwidth. This is the reason why a new standard was developed called 802.12 as it is not strictly Ethernet. In fact 802.12 is designed to better support both Ethernet and Token Ring.

The encoding techniques used are **5B/6B** and **NRZ**.

All four pairs of UTP are used. On Cat3 the longest cable run is 100m but this increases to 200m on Cat5.

The clock rate on each wire is 30MHz, therefore 30Mbits per second are transmitted on each pair giving a total data rate of 120Mbits/sec. Since each 6-bits of data on the line represents 5 bits of real data due to the 5B/6B encoding, the rate of real data being transmitted is 25Mbits/sec on each pair, giving a total rate of real data of 100Mbits/sec. For 2-pair STP and fiber, the data rate is 120Mbits/sec on the transmitting pair, for a real data transmission rate of 100Mbits/sec.

4.10 Gigabit Ethernet

Although the functional principles of Gigabit Ethernet are the same as Ethernet and Fast Ethernet i.e. CSMA/CD and the Framing format, the physical outworking is very different. One difference is the slot time. The standard Ethernet slot time required in CSMA/CD half-duplex mode is not long enough for running over 100m of copper, so **Carrier Extension** is used to guarantee a 512-bit slot time.

1000BaseX (802.3z)

802.3z is the committee responsible for formalising the standard for **Gigabit Ethernet**. The 1000 refers to 1Gb/s data speed. The existing Fibre Channel interface standard (ANSI X3T11) is used and allows up to 4.268Gbps speeds. The Fibre Channel encoding scheme is **8B/10B**.

Gigabit Ethernet can operate in half or full duplex modes and there is also a standard 802.3x which manages XON/XOFF flow control in full duplex mode. With 802.3x, a receiving station can send a packet to a sending station to stop it sending data until a specified time interval has passed.

There are three media types for 1000BaseX. 1000BaseLX, 1000BaseSX and 1000BaseCX.

With 1000BaseSX, 'S' is for Short Haul, and this uses short-wavelength laser (850nm) over multi-mode fibre. 1000BaseSX can run up to 300m on 62.5/125um multimode fibre and up to 550m on 50/125um multimode fibre.

Using 1300nm wavelength, Gigabit Ethernet (1000BaseLX where the 'L' is for Long wavelength laser, or Long Haul) can run up to 550m on 62.5/125um multi-mode fibre or 50/125um multi-mode fibre. In addition, 1000BaseLX can run up to 5km (originally 3km) on single-mode fibre using 1310nm wavelength laser.

1000BaseCX is a standard for STP copper cable and allows Gigabit Ethernet to run up to

25m over STP cable.

There is currently an issue as many multimode fibre installations using 62.5/125um fibre and so 220m is often the limit for the backbone when it should be 500m to satisfy ISO 11801 and EIA/TIA 568A.

1000BaseT (802.3ab)

Many cable manufacturers are enhancing their cable systems to 'enhanced Category 5' standards in order to allow Gigabit Ethernet to run at up to 100m on copper. The Category 6 standard has yet to be ratified, and is not likely to be due for a while.

In order to obtain the 1000Mbps data bit rate across the UTP cable without breaking the FCC rules for emission, all 4 pairs of the cable are used. Hybrid circuits at each end of each pair are used to allow simultaneous transmission and reception of data (full-duplex) by separating the transmission signal from the receiving signal. Because some transmission signal still manages to couple itself to the receiving side there is an additional echo canceller built in, this is called a NEXT canceller. This system minimises the symbol rate.

Encoding is carried out with **PAM-5**.

Go next to ***Ethernet Terminology and Errors***.

Home