Gigabit Ethernet (1000BaseT)

Gigabit Ethernet Summary

Gigabit Ethernet uses the same 802.3 frame format as 10 Mbps and 100 Mbps Ethernet systems. This operates at ten times the clock speed of Fast Ethernet at 1 Gbps. By retaining the same frame format as the earlier versions of Ethernet, backward compatibility is assured with earlier versions, increasing its attractiveness by offering a high bandwidth connectivity system to the Ethernet family of devices.

Gigabit Ethernet is defined by the IEEE 802.3z standard. This defines the Gigabit Ethernet Media Access Control (MAC) layer functionality as well as three different physical layers: 1000BaseLX and 1000BaseSX using fiber and 1000BaseCX using copper. These physical layers were originally developed by IBM for the ANSI Fiber Channel systems and used 8B/10B encoding to reduce the bandwidth required to send high speed signals.

The IEEE merged the Fibre Channel to the Ethernet MAC using a Gigabit Media Independent Interface (GMII), which defines an electrical interface enabling existing Fiber Channel PHY chips to be used, and enabling future physical layers to be easily added.

1000BaseT is being developed to provide service over four pairs of Category 5 or better copper cable. As discussed earlier this uses the same technology as 100BaseT2. This development is defined by the IEEE 802.3ab standard.

These Gigabit Ethernet versions are summarized in **Figure 1**.

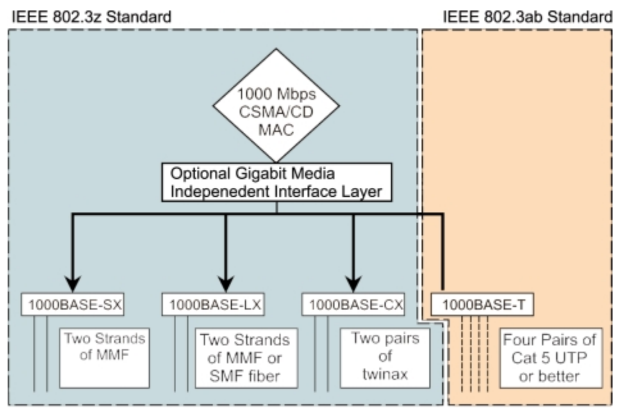

Figure 1. Gigabit Ethernet Versions

Gigabit Ethernet MAC Layer

Gigabit Ethernet retains the standard 802.3 frame format, however the CSMA/CD algorithm has had to undergo a small change 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. Without this increased slot time the network would have been impractically small at one tenth of the size of Fast Ethernet - only 20 meters!

The 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 6.096 ms to compensate for the tenfold speed increase. This then gives a collision domain of about 200 m (660 feet).

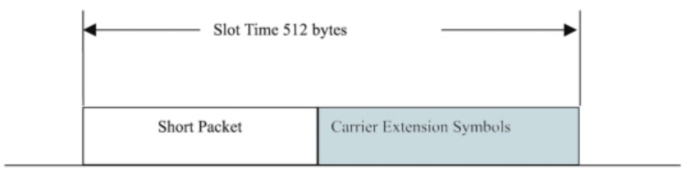

Figure 2. Carrier Extension

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 that are shorter than 512 bytes and to fill the remainder of the frame.

This is shown in **Figure 2**.

While this is a simple technique to overcome the network size problem, it could cause problems with very low utilization if we send a lot of short frames, typical of some industrial control systems.

For example, a 64 byte frame would have 448 carrier extension symbols attached and result in a utilization of less than 10%. This is unavoidable, but its effect can be minimized if we are sending a lot of small frames 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.

The size of the burst varies depending on how many frames are being sent and their size. 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.

This is shown in **Figure 3**.

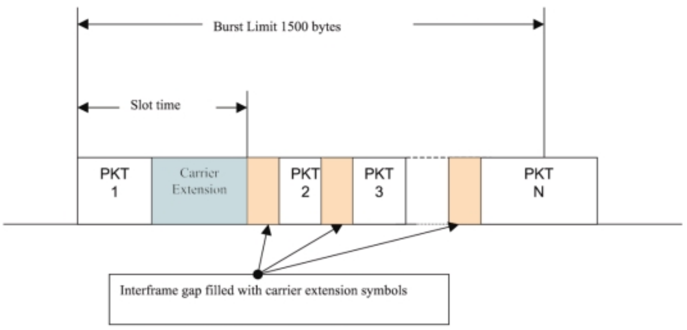

Figure 3. Packet Bursting

Physical Medium Independent (PHY) Sublayer

The 802.3z Gigabit Ethernet standard used the three PHY sublayers from the ANSI X3T11 Fiber Channel standard for the 1000BaseSX and 1000BaseLX versions using fiber optic cable and 1000BaseCX using shielded 150 Ohm twinax copper cable.

The Fiber Channel PMD sub layer ran at 1 GBaud and specifies the 8B/10B coding of the data, data scrambling and the Non-Return to Zero Inverted (NRZI) data coding together with the clocking, data and clock extraction processes. This translated to a data rate of 800 Mbps. The IEEE then had to increase the speed of the Fiber Channel PHY layer to 1250 MBaud to obtain the required throughput of 1 Gbps.

The 8B/10B technique selectively codes each group of eight bits into a ten bit symbol. Each symbol is chosen so that there are at least two transitions from "1" to "0" in each symbol. This ensures there will be sufficient signal transitions to allow the decoding device to maintain clock synchronization from the incoming data stream. The coding scheme allows unique symbols to be defined for control purposes, such as denoting the start and end of packets and frames as well as instructions to devices.

The coding also balances the number of 1s and 0s in each symbol, called DC balancing. This is done so that the voltage swings in the data stream would always average to zero, and not develop any residual DC charge, which could result in any AC-coupled devices distorting the signal. This phenomenon is called "baseline wander".

1000BaseSX for Horizontal Fiber

This Gigabit Ethernet version was 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 (660 and 1800 feet) depending on the bandwidth and attenuation of the fiber optic cable used. The standard 1000BaseSX NICs available today are full-duplex and incorporate SC fiber connectors.

1000BaseLX for Vertical Backbone Cabling

This version was 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 km (3 miles) for single mode cable and 550 meters (1800 feet) for multimode fiber cable. Many 1000BaseLX vendors guarantee their products over much greater distances, typically 10 km (6 miles). Fiber extenders are available to give service over as much as 80 km (50 miles). The standard 1000BaseLX NICs available today are full-duplex and incorporate SC fiber connectors.

1000BaseCX for Copper Cabling

This version of Gigabit Ethernet was 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: The High-Speed Serial Data Connector (HSSDC) known as the Fiber Channel Style 2 connector and also the nine pin D-subminiature connector from the IBM Token Ring systems. The maximum cable length is 25 meters (80 feet) for both full- and half-duplex systems.

These systems are not currently available in the marketplace for connecting different switches. The preferred connection arrangements are to connect chassis-based products via the common backplane and stackable hubs via a regular fiber port.

1000BaseT for Category 5 UTP

This version of the Gigabit Ethernet is 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. Compare this to the existing 100BaseTX system, which has individual pairs for transmitting and receiving. This is shown in Figure 4.

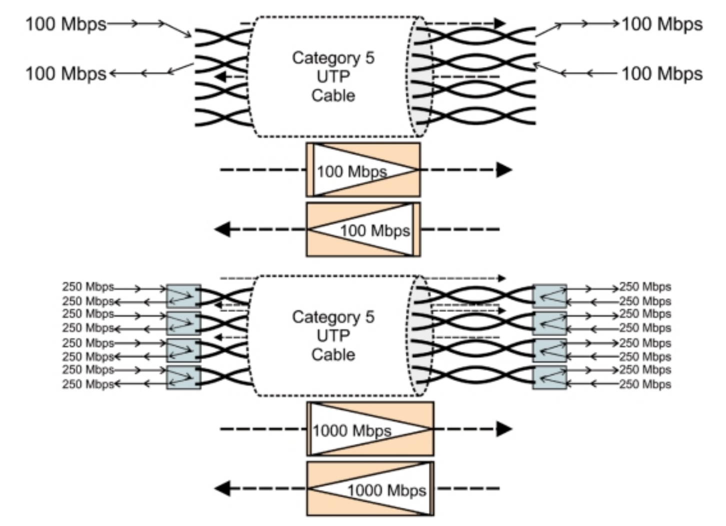

Figure 4. Comparison of 100BaseTX and 1000BaseT

This system uses the same data encoding scheme developed for 100BaseT2, which is PAM5. This utilizes five voltage levels so has less noise immunity, however the Digital signal Processors (DSP) associated with each pair overcomes any problems in this area. The system achieves its tenfold speed improvement over 100BaseT2 by transmitting on twice as many pairs (4) and operating at five times the clock frequency (125 MHz). Active noise cancellation circuitry is needed to enable this system to function correctly, as shown in Figure 5.

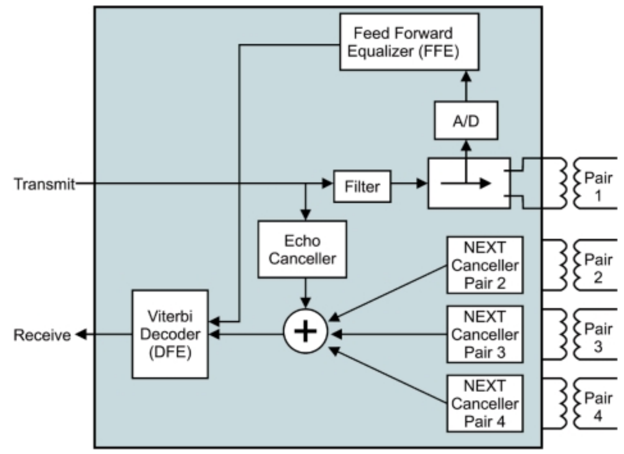

Figure 5. 1000BaseT Receiver uses DSP Technology

Gigabit Ethernet Full-Duplex Repeaters

Gigabit Ethernet nodes are connected to full-duplex repeaters also known as non-buffered switches or buffered distributors.

As shown in Figure 6 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. It is therefore combining the functions of a repeater with some features of a switch (see Figure 6).

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.

The repeater does not analyze the packet address fields to determine where to send the packet, like a switch does, but simply sends out all valid packets to all the other ports on the repeater.

The IEEE does allow for half-duplex Gigabit repeaters - however none exist at this time.

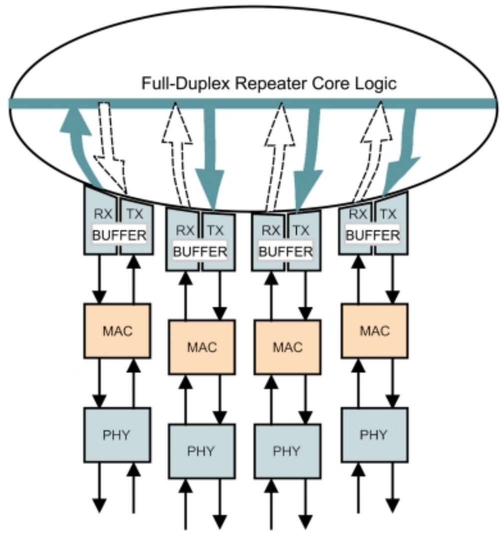

Figure 6. Gigabit Ethernet Full-Duplex Repeaters

Gigabit Ethernet Design Considerations

Fiber Optic Cable Distances

The maximum cable distances, which can be used between the node and a full-duplex 1000BaseSX and LX repeater, depend mainly on the chosen wavelength, the type of cable, and its bandwidth. The maximum transmission distances on multimode cable are limited by the Differential Mode Delay (DMD).

The very narrow beam of laser light injected into the multimode fiber results in a relatively small number of rays going through the fiber core. These rays each have different propagation times because they are going through differing lengths of glass by zigzagging through the core to a greater or lesser extent. These pulses of light can cause jitter and interference at the receiver. This is overcome by using a conditioned launch of the laser into the multimode fiber. This spreads the laser light evenly over the core of the multimode fiber so the laser source looks more like a Light Emitting Diode (LED) source. This spreads the light in a large number of rays across the fiber resulting in smoother spreading of the pulses, so less interference. This conditioned launch is done in the 1000BaseSX transceivers.

Table 1. gives the maximum distances for full-duplex 1000BaseX 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	huge	0.4	5000

Table 1. Maximum Fiber Distances for 1000BaseX (Full-duplex)

Gigabit Repeater Rules

The cable distance and the number of repeaters, which can be used in a half-duplex 1000BaseT 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 1000BaseT systems is the time to transmit 512 bytes or 4096 bits and equals 6.096 ms. 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. Therefore the 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 ms

Table 2. gives typical maximum one-way delays for various components. Repeater and NIC delays for your specific components can be obtained from the manufacturer.

These calculations give the maximum collision diameter for IEEE 802.3z half-duplex Gigabit Ethernet systems. The maximum Gigabit Ethernet network diameters specified by the IEEE are shown in Table 3.

Component	Maximum Delay (µs)
Gigabit NIC	0.432
Gigabit Repeater	0.488
UTP Cable (per 100 m)	0.550
Multimode Fiber (per 100 m)	0.500

Table 2. Maximum One-Way Gigabit Ethernet Component Delays

System	Maximum Collision Diameter (point-to point, half-duplex)	Maximum Collision Diameter (one repeater segment)
1000BaseCX	25 m (82 ft)	50 m (164 ft)
1000BaseT	100 m (330 ft)	200 m (656 ft)
1000BaseSX or LX	316 m (1036 ft)	220 m (721 ft)

Table 3. Maximum Half-Duplex Gigabit Ethernet Network Diameters

Note: Half-duplex Gigabit Ethernet Repeaters are not available for sale. Use full duplex repeaters with the point-to-point cable distances between node and repeater or node and switch as listed in Table 1.