the **EXTENSION**

A Technical Supplement to control NETWORK

INTRODUCTION TO FAST ETHERNET

By George Thomas, Contemporary Controls

With this issue of the Extension, we will return to our discussion on Ethernet by introducing the topic of Fast Ethernet. Fast Ethernet increases the data rate by 10 fold while maintaining interoperability with legacy Ethernet.

Since its introduction in the early 70's, Ethernet has continued to evolve evidenced by the fact the IEEE 802.3-2000 standard, which describes the technology, is now over 1500 pages! Ethernet was created with a multi-drop physical layer using a thick backbone eventually operating at 10 Mbps. A thin coaxial bus topology was introduced to reduce cabling expense, but it was the twisted-pair version that received the greatest acceptance. Throughout these changes, Ethernet maintained its 10 Mbps data rate. What was going to be Ethernet's encore?

Fiber Distributed Data Interface (FDDI) operating a 100 Mbps did exist but was considered expensive. The question was, "Could Ethernet operate at 100 Mbps?" The answer was "yes," but there were competing approaches to Fast Ethernet. One approach was to not make any changes at all, just scale the protocol to 100 Mbps. This would provide backward compatibility and was favored by most vendors. The second approach favored a redesign of the medium access control (MAC) in order to gain a feature called demand-priority at the sake of backward compatibility. In the end the "technically elegant" demand-priority solution, which utilized a token-passing protocol, lost out both in the IEEE 802.3 committee and in the marketplace. Introduced as 100VG-AnyLAN, the technology never gained widespread support outside its two proponents, Hewlett Packard and AT&T. Although the words "Fast Ethernet" are not used, the IEEE 802.3u was adopted as the Fast Ethernet standard in 1995.

How fast is fast?

It is generally accepted that Fast Ethernet implies Ethernet at 100 Mbps. Compared to 10 Mbps, 100 Mbps seems quite fast. In fact, to most microcontrollers it is exceptionally fast and a difficult speed to maintain. However, there are still much faster technologies. Can Ethernet scale to 1000 Mbps? Well not unless you are content with a maximum network diameter of 20m. Still many attributes of 10 Mbps Ethernet are retained in the 1000BASE-T standard. This technology is called Gigabit Ethernet and its proposed successor will operate at 10 Gigabits/second. Although 100 Mbps is not the fastest, we will still concentrate on this technology since it has the most application on the plant floor.

SCALING ETHERNET

In Table 1 you will see a summary of attributes for both Ethernet and Fast Ethernet. You will notice values of the various Ethernet attributes, in terms of bits or bit-times, are maintained even as Ethernet is scaled from 10 Mbps to

	Ethernet/802.3	Fast Ethernet 802.3u
Slot Time	512 bit-times	512 bit-times
Interframe Gap	96 bit-times	96 bit-times
Attempt Limit	16 tries	16 tries
Backoff Limit	10 (exponent)	10 (exponent)
Jam Size	32 bits	32 bits
Max Frame Size	1518 bytes	1518 bytes
Min Frame Size	64 bytes	64 bytes
Address Size	48 bits	48 bits

Table 1—Scaling Ethernet by 10 fold yields Fast Ethernet without any change in attribute parameters.

100 Mbps. However, one bit-time at 10 Mbps is 100ns and at 100 Mbps it is 10ns. Look at Table 2 where the actual times replace the bit and bit-time figures. The good news is that an equivalent amount of data sent at 10 Mbps will be received in a tenth of the time at 100 Mbps. Also notice that the minimum frame size and slot time have been reduced by a factor of ten as well. These attributes are directly linked to the collision domain of Ethernet and determine the maximum network diameter of the network based upon the maximum round-trip delay of a signal propagating between the two furthest points on the network. While the Ethernet protocol scales nicely from 10 Mbps to 100 Mbps, the actual time it takes for signals to pass down wires does not. The result is that the maximum network diameter must be reduced from about 2800m for 10 Mbps Ethernet to a low of 205m for Fast Ethernet. The only way to increase distance while retaining 100 Mbps speed is to use full-duplex links which eliminate collisions altogether; therefore, the link segments are not limited by timing.

	Ethernet/802.3	Fast Ethernet 802.3u
Slot Time	51.2 µs	5.12 µs
Interframe Gap	9.6 µs	0.96 µs
Attempt Limit	16 tries	16 tries
Backoff Limit	10 (exponent)	10 (exponent)
Jam Size	3.2 µs	0.32 µs
Max Frame Size	1214.4 µs	121.44 µs
Min Frame Size	6.4 µs	0.64 µs
Address Size	4.8 µs	0.48 µs

Table 2—In terms of actual time, there are significant differences between Ethernet and Fast Ethernet.

HALF-DUPLEX or FULL-DUPLEX

Full-duplex links are the key to extending the maximum network diameter of Fast Ethernet. Full-duplex requires separate receive and transmit paths and link segments consisting of no more than two devices. These devices can be Ethernet adapters or switching hub ports. Notice that I did not mention repeating hub ports. A repeating hub is part of the collision domain and reinforces collisions received on any of its other ports. A repeating hub is not capable of full-duplex operation. Although it is possible to have just two Ethernet adapters configured for full-duplex, expansion beyond two adapters requires a switching hub capable of supporting full-duplex operation.

Half-duplex means transmitting and receiving over the same medium but not at the same time. Full-duplex allows for simultaneous sending and receiving. Coaxial-based transceivers such as 10BASE5 and 10BASE2 were never able to invoke full-duplex since they do not have separate receive and transmit paths. However, 10BASE-T and 10BASE-FL do have separate receive or transmit paths and are capable of full-duplex operation depending upon the sophistication of the Ethernet adapter or switching hub. If these interfaces are configured for half-duplex, then the simultaneous detection of receive and transmit activity will trigger collision detection. These same interfaces, configured for full-duplex, would disable this collision detection logic since full-duplex does not follow the CSMA/CD rules of shared Ethernet.

It is very important that a full-duplex link be configured properly. A station or switching hub port will send out frames at will, ignoring the CSMA/CD protocol of shared Ethernet, if it is configured for full-duplex. If the other end is configured for half-duplex, it will incorrectly detect collisions and take actions that could cause late collisions (which are not automatically resent) and CRC errors. The result is a general slowdown of the network negating the benefits of migrating to Fast Ethernet.

200 Mbps?

One virtue of full-duplex is the claim that throughput is doubled to 200 Mbps by simply invoking full-duplex over half-duplex. Although mathematically correct that two simultaneous data streams can travel on the medium—one going 100 Mbps in one direction and another one going 100 Mbps in the other direction—this occurrence is not common when using industrial automation protocols. Most protocols operate as either master/slave or command/response. With these types of protocols, a station initiates a command to another station which must decode and execute the command before providing the response. This type of communication could function quite well with a half-duplex link since only a command stream or a response stream is active at one time. In backbone applications, full-duplex operation can be helpful, but do not assume that a system will double in performance by simply invoking full-duplex. The benefits of full-duplex are mostly to gain a larger network diameter and eliminate the complexity and uncertainty of the CSMA/CD protocol.

ETHERNET ACCORDING TO 802.3

If you look at Figure 1, you will see a very complex diagram of the various flavors of Ethernet using the definitions of 802.3. Notice that the possible data rates can range from 1 Mbps to 1000 Mbps. The 10000 Mbps rate is yet to be added to the standard and 1 Mbps is not popular. We will concentrate on the 10 Mbps and 100 Mbps options. In terms of the OSI Reference Model, all the data rate variations are considered physical layer attributes, while the data link layer is common to all data rate variations. The 802.3 model does not even address the upper layers of the OSI model.
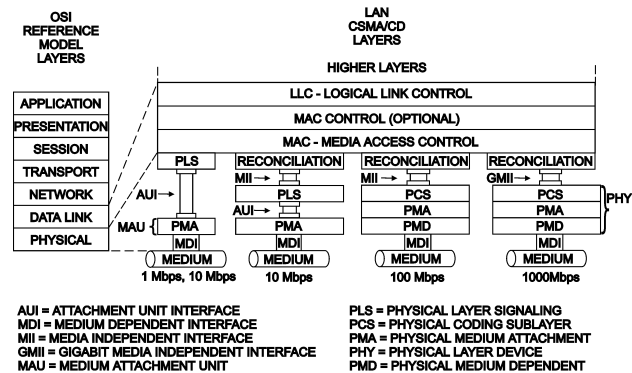

Figure 1—The complex 802.3 model only addresses two layers of the OSI Reference Model.

Of course, the 10 Mbps versions are the simplest. However, notice that two are shown in the figure. The one on the left is the traditional model showing the Attachment Unit Interface (AUI). At this point any physical medium can be attached by way of the DB-15 AUI connector. This was the connector originally used to attach to a 10BASE5 (thick trunk) transceiver which was usually mounted in the ceiling. In 802.3 terms, this transceiver would be called a Physical Medium Attachment (PMA) or simply a Medium Attachment Unit (MAU). In order to get down to the Physical Layer Signaling (PLS) of the Ethernet interface, an AUI cable, or drop cable, is used acting like an extension cord. On one end is a male DB-15 and the other end is a female DB-15. The AUI port is medium independent. If a fiber optic medium is desired, the appropriate MAU would be connected to the AUI port. On the fiber optic side, the fiber optic cables would be connected to the Medium Dependent Interface (MDI) which would include ST-style fiber optic connectors. This completes the connection to the medium.

With the introduction of Fast Ethernet, the model changed somewhat. A new connection was defined and was called the Media Independent Interface (MII). It provided more bandwidth at 100 Mbps than could be provided by the AUI interface. Notice that at 10 Mbps, the MII is inserted between the PLS and the Reconciliation interface. This is shown to the immediate right of the traditional model. Further to the right is the 100 Mbps variation which introduces additional terms. The Physical Medium Dependent now attaches to the MDI and the Physical Coding Sublayer (PCS) attaches to the MII. In the middle is the PMA. Collectively, the PCS, PMA and PMD are termed the PHY for Physical Layer Device. The MII connector is defined as 50 mm wide dual row, 40 pin. Like the AUI, it is only an option. Most Ethernet interfaces only provide the necessary

MDI connector and nothing else. To understand the standard these terms and connections must be defined. You might ask why a MII interface was defined for 10 Mbps Ethernet when an AUI seems to have served the industry quite well? The MII interface supports both 10 Mbps and 100 Mbps data rates which is key to providing dual speed support between legacy 10 Mbps and newer 100 Mbps systems. On the other hand, the AUI only supports 1 Mbps and 10 Mbps while the new Gigabit Medium Independent Interface (GMII) only supports 1000 Mbps.

100BASE-T

Fast Ethernet was introduced with the 100BASE-T standard. In this standard, new concepts such as Reconciliation Sublayer (RS), Media Independent Interface, Auto-Negotiation and Management are introduced along with several different physical layers. Therefore, unlike 10BASE-T, specifying 100BASE-T is not sufficient when designating a physical layer since three are mentioned in the standard—100BASE-T4, 100BASE-TX and 100BASE-FX. What 100BASE-T basically describes is the new MII interface which is much more sophisticated than the AUI. Instead of a serial bit stream for data, four-bit nibbles are transferred between the PHY and RS over the MII for both reception and transmission. For 10 Mbps operation, clocking is reduced from that of 100 Mbps otherwise the interface is the same. It is not a requirement that both speeds be supported.

Management Interface

Another unique feature of Fast Ethernet is the ability to manage the PHY itself through a two-wire interface in the MII. As a minimum, one control and one status register must be provided. The PHY must report its abilities and respond to commands. The PHY reports its type of physical layer, its ability to auto-negotiate and if it is capable of handling full-duplex operation. The station management might command the PHY to a fixed speed or to auto-negotiate. The PHY reporting its abilities ensures that it will not be commanded to perform what it is incapable of doing.

	10BASE-T	100BASE-TX	100BASE-FX
Data Rate	10 Mbps	100 Mbps	100 Mbps
Encoding	Manchester	4B/5B	4B/5B
Signaling	5V Differential	MLT-3	NRZI
Wires	4	4	2 (Fibers)
Cable	Cat. 3 UTP	Cat. 5 UTP	62.5/125 µm
Connector	RJ-45	RJ-45	SC or ST
Max Segment	100m	100m	2km
Max Transceivers	2	2	2

Table 3—Ethernet physical layers differ in terms of data rate, encoding and signaling.

100BASE-T2, 100BASE-T4

Two Fast Ethernet physical layers that are mentioned are 100BASE-T4 and 100BASE-T2. Both were developed to utilize lower speed cabling that may exist in an installation. 100BASE-T4 requires four pairs and is incapable of full-duplex operation. 100BASE-T2 only needs two pairs but uses a sophisticated signaling scheme that is difficult and expensive to implement. The industry seems to have rewired for the higher speed Category 5 cable so these two interfaces are not popular.

100BASE-X

Like 100BASE-T, 100BASE-X is not a unique physical layer but details the encoding for the two most popular physical layers—100BASE-TX and 100BASE-FX. One physical layer is for copper and the other for fiber optics, yet the standard applies to both. Much of the 100BASE-X standard comes from the FDDI standard including the 4B/5B encoding.

4B/5B

Data transfers over the MII are done with 4-bit nibbles that represent actual data. With 10BASE-T, Manchester encoding is used which guarantees a transition within every bit cell regardless of logic state. This effectively creates a 20 Mbaud signal for a 10 Mbps data rate. If the same encoding were used for Fast Ethernet, a 200 Mbaud signal would result making it difficult to maintain the same 100m maximum segment length due to high frequency attenuation. If we could use a code such as non-return to zero (NRZ), we could match the data rate with the baud rate. The problem with NRZ is that it has a DC component, which transformers do not like, and it provides little information about clocking. A compromise is the 4B/5B code where the 4-bit nibbles being transferred over the MII are actually encoded as five-bit symbols sent over the medium. The encoding efficiency is 80% and the baud rate increases to 125 Mbaud. This is still fast but not as fast as 200 Mbaud. The actual codes used are chosen so that sufficient transitions occur in the resulting bit pattern so as not to lose clocking. With twice as many codes as necessary, there are many that are left unused and some that are defined for control purposes. The 4B/5B scheme is used for both the 100BASE-TX and 100BASE-FX physical layers.

100BASE-TX

The 100BASE-TX twisted-pair physical layer retains the same MDI (RJ-45) connector and pinout as 10BASE-T so that auto-negotiation is possible. Two twisted-pairs with separate receive and transmit paths allow for full-duplex operation. Cabling requires a higher performance Category 5. This cable type is an unshielded twisted-pair characterized as 100 ohms. The standard mentions a 150 ohm shielded twisted-pair (STP) option with a DB-9 connector that can be used as well. A unique feature of 100BASE-X encoding is that the link is always active even without data since a unique symbol called IDLE is sent continuously during inactivity. Reception of IDLE serves as a link integrity function. Maximum segment length for 100BASE-TX is 100m just like 10BASE-T. Signaling on the twisted-pair incorporates a three level multi-level technique called MLT-3. The benefit of this code is to further reduce the electromagnetic emissions (EMI) and the bandwidth requirements of the medium.

100BASE-FX

The 100BASE-FX fiber optic physical layer is very similar in performance to 10BASE-FL. Maximum segment length is

2km for both technologies; however, for 100BASE-FX this is only achieved on full-duplex links. On half-duplex links the segment length cannot exceed 412m. Either SC or ST fiber optic connectors can be used, but SC is recommended. Multi-mode fiber optic cable (62.5/125µm) is what is normally used; however, it is possible to use single mode fiber optics for greater distances on full-duplex links. The signaling on fiber optics is NRZI (non-return to zero inverted) since there is no concern for EMI on fiber optic links. With NRZI, the state of signaling inverts on logic 1s, thereby providing some clocking information unlike NRZ.
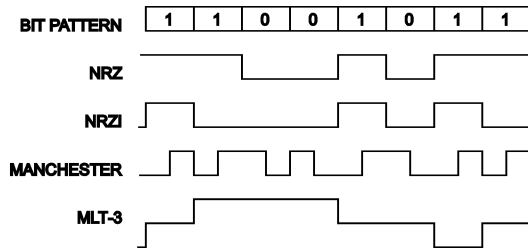

Figure 2—Various encodings for the same bit pattern. Notice that MLT-3 approximates a sine wave with a much lower fundamental frequency than the data rate.

Table 3 summarizes the three most important Ethernet physical layers and Figure 2 shows the various encoding schemes.

Auto-Negotiation

With the proliferation of Fast Ethernet and the similarity of the cabling components to conventional Ethernet, a means was proposed in IEEE 802.3u to automatically configure Fast Ethernet ports to work with either legacy Ethernet ports or other Fast Ethernet ports. This configuration protocol was based upon National Semiconductor's NWay standard. There is a way for twisted-pair links to automatically configure compatible formats in order for links to begin communi-cating. This scheme is only suitable for twisted-pair links and not coaxial or fiber optic links. Coaxial cable is a legacy 10 Mbps standard that is not in the plans for evolving Ethernet. Fiber optics is a different story. Although fiber optics is very much in the plans for evolving Ethernet, there is no simple way for two fiber optic devices to auto-negotiate data rates since a 10BASE-FL device operates at 850nm while a 100BASE-FX device operates at 1300nm. These devices will not interoperate. The Auto-Negotiation protocol is only intended for twisted-pair links where there are only two devices on a segment. Notice that twisted-pair bus segments, which are favored in industrial automation, are nowhere supported by Ethernet standards. Industrial customers must be content wiring in a star topology and using either repeating hubs or switching hubs.

The benefit of auto-negotiation is to provide hands-free configuration of the two devices attached to the link segment. At connection time, each of the two devices will advertise all their technical abilities. These abilities have been ranked by the standard as shown in Table 4. The lowest possible ranking is 10BASE-T which assumes half-duplex or shared Ethernet operation. The very next ranking is 10BASE-T full-duplex indicating that full-duplex has higher performance than half-duplex. Notice that there is no ranking for 100BASE-T4 full-duplex. This is because this technology is not capable of full-duplex operation. Finally, the highest ranking is 1000BASE-T full-duplex. This ranking scheme has been provided for completeness. It is not assumed that a particular adapter can handle all technologies. In fact, some of these technologies may not have been commercialized. However, they are all listed consistent with the IEEE 802.3 standard.

| 1000BASE-T full-duplex |
| 1000BASE-T |
| 100BASE-T2 full-duplex |
| 100BASE-TX full-duplex |
| 100BASE-T2 |
| 100BASE-T4 |
| 100BASE-TX |
| 10BASE-T full-duplex |
| 10BASE-T |

Table 4—Auto-Negotiation assumes a ranking of priorities. 10BASE-T is at the bottom.

Each device examines each other's technical abilities and determines the lowest common denominator. For example, if an Ethernet adapter can only handle 10BASE-T while a switch port can handle either 10BASE-T or 100BASE-TX, 10BASE-T will be chosen by both. If two Ethernet adapters, one only advertising 10BASE-T and the other only advertising 100BASE-TX connect, there will be no subsequent communication since no compatibility exists.

FLP and NLP

With legacy 10BASE-T systems, a Normal Link Pulse (NLP) is sent between two devices in order to demonstrate that the link is functional in both directions. With Fast Ethernet, the NLP is replaced with a Fast Link Pulse (FLP) that carries the Auto-Negotiation information. The FLP is interpreted by legacy 10BASE-T devices as a NLP so as to provide backward compatibility. If a connection is broken, as evidenced by a lack of link pulses, a series of FLPs will be observed once the link is reestablished. After auto-negotiation is completed, link pulses or link activity are monitored as with legacy 10BASE-T networks.

CONCLUSION

Fast Ethernet provides the path to higher performance Ethernet systems with the ability to interoperate with legacy 10 Mbps systems. Added features such as full-duplex and auto-negotiation also make this migration path attractive.

REFERENCES

Ethernet—The Definitive Guide, Charles E. Spurgeon, 2000, O'Reilly & Associates, Inc.

Switched and Fast Ethernet, Second Edition, Robert Breyer and Sean Riley, 1996, Macmillan Computer Publishing USA

International Standard ISO/IEC 8802-3 ANSI/IEEE STD. 802.3, 2000 Editiion, The Institute of Electrical and Electronic Engineers, Inc.