

This article appears in **The Industrial Ethernet Book**, issue 1, published by Fieldbus.pub. For more information on **The Industrial Ethernet Book**, please contact Adrian Chesney, adrian@fieldbus.com or +44 (0) 1329 846166.

the EXTENSION

A Technical Supplement to control NETWORK

INTRODUCTION TO ETHERNET

Ethernet for Control—Understanding the Basics

INTRODUCTION

There has been much discussion recently regarding the applicability of using Ethernet at various levels of the control hierarchy. Since Ethernet is so prevalent in the office and frequently used as the enterprise network for high-end controllers, it would seem to be a natural to use Ethernet at the control level or even at the device level as proposed by some in our industry. The arguments for its use include low cost, good connectivity and simple migration to higher speed networks. The cry to use “standard” Ethernet for control applications requires an understanding of the basics of Ethernet.

What is standard Ethernet?

We are not sure what standard Ethernet is but it certainly is not the 2.94 Mbps version that came out of Xerox’s Palo Alto Research Center (PARC) in the early 70s. In 1980, Digital Equipment Corporation (DEC), Intel and Xerox published the DIX V1.0 standard which boosted the speed of Ethernet to 10 Mbps while maintaining Ethernet’s thick trunk cabling scheme. In 1982 the DIX V2.0 standard was released which is now commonly referred to as Ethernet II. Xerox then relinquished its trademark.

At the time of the first DIX standard, the Institute of Electrical

and Electronic Engineers (IEEE) was attempting to develop open network standards through the 802 committee. In 1985 the IEEE 802.3 committee published “IEEE 802.3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications.” This technology is called 802.3 CSMA/CD and not Ethernet; however, it is frequently referred to as Ethernet even though the frame definition differs from DIX V2.0. Although 802.3 and DIX frames can coexist on the same cable, interoperability is not assured. Therefore, when discussing “Ethernet,” it is necessary to clarify 802.3 frames or DIX V2.0 frames.

To further confuse issues, standard Ethernet sometimes means an attached protocol—mainly TCP/IP. Ethernet only defines the data link and physical layers of the Open Systems Interconnect (OSI) Reference Model whereas TCP/IP defines the transport and network layers respectively of the same model. Therefore, when the suggestion is made to use standard Ethernet for control does this mean TCP/IP connectivity as well?

APPLICATION
PRESENTATION
SESSION
TRANSPORT
NETWORK
DATA LINK
PHYSICAL

Figure 1—Ethernet defines the lower two layers of the OSI Reference Model.

ETHERNET FRAMES

The two types of Ethernet frames used in industry are similar. The DIX V2.0 frame, frequently referred to as the Ethernet II frame, consists of an eight-byte preamble, six-byte source and destination addresses, a two-byte type field used to identify higher layer protocols, a variable data byte field followed by a four-byte frame check sequence (FCS) field. The IEEE 802.3 frame divides the preamble into a seven-byte preamble followed by a single byte start of frame delimiter (SFD). The two-byte type field now becomes a two-byte length field. The data field now includes an 802.2 logical link control (LLC) field that precedes the actual data. The FCS remains the same.

Preamble

The DIX preamble consists of 64 bits of alternating “1s” and “0s” but ending with two “1s” to indicate that a valid frame is to begin. This creates a 10 Mhz signal that synchronizes the receivers on the network before actual data arrives. Ethernet uses Manchester encoding.

The IEEE redefined the preamble to be seven bytes of preamble, the same as the DIX preamble, followed by a one-byte start of frame delimiter (SFD) which looks like the last byte of the DIX preamble. There is no change in operation between the DIX preamble and the IEEE preamble and SFD byte. Both preambles are not considered part of the frame when calculating the size of the overall frame.

Ethernet II DIX Frame						
64 bits	48 bits			48 bits	16 bits	32 bits
Preamble	Individual/ Group Address Bit	Destination Address		Source Address	Type	Frame Check Sequence

IEEE 802.3 Frame						
56 bits	8 bits	48 bits			48 bits	32 bits
Preamble	SFD	Individual/ Group Address Bit	Globally/ Locally Administered Address Bit	Destination Address	Source Address	Frame Check Sequence

Figure 2—Two types of Ethernet frames are used in industry.

Destination Address

In the DIX standard the first bit of the 48-bit destination address indicates if the address is a multicast address or a physical address. A “0” indicates a unicast transmission to the indicated destination while a “1” indicates a multicast or group address.

The IEEE standard further defines the second bit of the 48-bit destination to indicate if the address is locally administered or globally administered. This bit is a “0” when the address is globally administered; that is, assigned by the Ethernet interface manufacturer.

A 48-bit address of all “1s” is a broadcast address in both DIX and IEEE formats indicating that the transmission is directed to all devices on the network.

Source Address

The 48-bit source address is appended to the transmission as an aid to the higher layer protocols. It is not used for medium access control. To avoid duplicate node IDs for global addresses, the Ethernet adapter manufacturer obtains an Organizationally Unique Identifier (OUI) from the IEEE (for an administration fee). The OUI is 24-bits long and is used as the most significant portion of the 48-bit address. The manufacturer, using

good recordkeeping, will assign sequential numbers to each adapter card he makes thereby creating a worldwide unique address. With 24-bits to work with, a lot of adapters can be produced from a single manufacturer. A list of OUI assignments can be found on the Internet.

Type and Length Field

The original intention of Ethernet was never to use its data link layer as the means for providing guaranteed delivery of data. It was always the intent that a higher layer protocol would do that service. Therefore it was only necessary to identify by number which higher layer protocol was being used through the two-byte field in the DIX frame. Originally, Xerox maintained the assignments and now IEEE provides the administration.

The 802.3 standard does not include the type field but instead defines it as a length field. Per the 802.3 standard, a value in this field of 1518 or less indicates the length of the data field, while values above this may be ignored, discarded or used in a private manner. These out of bound values could then be used to identify higher layer protocols just like DIX frames.

What is important here is that since DIX and IEEE frames are identical in terms of the number of bits and length of fields, both frames can coexist on the same network but may not be able to communicate to one another. Much of the existing TCP/IP software that binds to Ethernet uses DIX frames and not 802.3 frames, so care must be exercised when selecting or developing software or claiming interoperability.

Data Field

A raw Ethernet frame (no encapsulated protocol or LLC) can be up to 1500 bytes long but no less than 46 bytes. This is the DIX frame.

Although the total available length of the IEEE data field is the same as the DIX frame, the LLC header reduces the amount of field available for actual data or payload as it is sometimes referred to. If the LLC header and actual payload are less than 46 bytes, the data field must be padded to 46 bytes to ensure that the transmission is not interpreted as a runt packet or packet fragment.

Frame Check Sequence

Both the DIX and IEEE standard use four bytes to hold the CRC-32 check on the complete frame from destination address all the way to the end of the data field. The receiving station calculates its own CRC-32, checks on the received data and compares the results with the transmitted CRC-32 value for a match indicating a successful reception. Note that there is no inherent mechanism in the Ethernet data link layer protocol to inform the source node that a reception was accepted or rejected due to a failed CRC-32 check. That task is left to the higher layer protocol.

ETHERNET PHYSICAL LAYERS

Although Ethernet was originally designed as a coaxial bus system, alternate physical layers have evolved since the early 80s. The IEEE 802 committee has defined several physical layers and that is why it is important to specify the correct option when selecting Ethernet.

10BASE5

The original Ethernet was configured as a bus system with a thick coaxial cable as the medium. That is what was specified in the 1980 DIX standard. An external transceiver called a medium attachment unit (MAU) clamps at particular points on the cable marked by stripes every 2.5 meters. From the transceiver, an attachment unit interface (AUI) cable connects to an AUI port on the actual Ethernet adapter that fits into the computer. The AUI port is a DB-15 connector. A coaxial segment can be up to 500 meters long and AUI cables are each restricted to 50 meters in length. A total of 100 transceivers can occupy one trunk segment. Individual trunk segments can be cascaded using repeaters up to 2000 meters. In 1985 the IEEE standardized this configuration as 10BASE5 to signify 10 Mbps baseband signaling up to 500 meters in length.

Thick coaxial cable is indeed bulky and its topology is not always convenient to wire in a plant. Troubleshooting a 100-station segment could be a nightmare, so you do not see new 10BASE5 installations. There is no support for this cable with Fast Ethernet technology.

10BASE2

The answer to the bulkiness of 10BASE5 along with its expense was Thinnet or Cheapernet

standardized in 1985 as 10BASE2. Thinnet again was a bus topology but this time with internal transceivers. A thin RG-58/u coaxial cable interconnects up to 30 stations to a maximum length of 185 meters. Segments can be repeated up to 740 meters. BNC style connectors, terminators and taps are used to cable the system. Although easier to install than 10BASE5, the focus on new installations is towards twisted-pair cabling. This cable is likewise not supported by Fast Ethernet.

10BASE-T

In 1990 the IEEE published 10BASE-T after pioneering work was done to introduce twisted-pair cabling and star topology to Ethernet installations. The 10BASE-T Ethernet adapters have internal transceivers and RJ-45 connectors. Usually two-pair unshielded cabling is attached to a hub in a point-to-point fashion. Bus connections are not allowed. The connection between an adapter and hub cannot exceed 100 meters in length. Hub-to-hub connection length can vary depending upon the medium used. If another twisted-pair connection is used, the maximum length is again 100 meters. With Thinnet it is 185 meters and with thick coaxial cable 500 meters.

The star topology is much easier to troubleshoot than a bus system; however, the reliability of the hub now must be considered in the overall reliability of the system. Another reason for the focus on twisted-pair is that development of Fast Ethernet is based on twisted-pair and not coaxial cable providing no migration path for installed coaxial cable.

10BASE-F

The 10BASE-F standard is actually a series of fiber optic standards. Fiber optics provides long distance,

higher-speed migration, noise immunity and electrical isolation. There are three media standards:

10BASE-FL—*This fiber link standard replaces older FOIRL standard.*

10BASE-FB—*This backbone standard is not very popular.*

10BASE-FP—*This passive hub technology is also not popular.*

The 10BASE-FL standard requires a duplex 62.5/125µm fiber optic cable for each link. Transmission distances of up to 2km are possible as well as full-duplex operation.

MEDIUM ACCESS CONTROL

What follows is a discussion of the medium access control protocol for a 10 Mbps half-duplex Ethernet network operating with several nodes.

When a station wants to transmit, it first waits for an absence of a carrier, which would indicate that some other station is transmitting. As soon as silence is detected, the station waiting to transmit continues to defer until the Interframe Gap (IFG) time has expired which is a minimum of 96-bit times (9.6µs). If a carrier still appears to be absent, the station begins to transmit while observing its collision sense circuitry. If no collision is detected, the transmitting station assumes the transmission was sent successfully. If the transmitter detects an early collision, one which occurred during the preamble, the station continues to send the preamble plus 32 bits of data called a jam signal. This ensures that other stations will note the collision as well. After the collision, the transmitting station will backoff from retransmitting based upon a backoff algorithm. If no collisions are detected after 512-bit times (not counting the preamble), the station

is assumed to have acquired the channel and no late collisions should occur on a properly working network. The collision counter is cleared. This 512-bit time (51.2µs) is called the slot time and is critical in the way Ethernet arbitrates access to the cable.

Collision Domain

This slot time defines the upper bound limit of the total propagation delay of a transmitted symbol from one end of the network to the farthest end and back. This includes the time it takes the symbol to travel through cables, repeaters and MAUs and varies with devices used. However, regardless of the path, the resulting propagation delay must be less than the slot time. Therefore the slot time defines Ethernet's maximum network diameter which limits its collision domain. A collision domain that exceeds the maximum network diameter violates Ethernet's medium access control mechanism resulting in unreliable operation.

Collisions can generate runt packets that are less than 512 bits in length. These can be detected by the receiving nodes and discarded accordingly. That is why it is important that a minimum valid Ethernet frame always be sent to distinguish valid packets from packet fragments. A minimum of 46 bytes in the data field ensures that a valid Ethernet frame is 512-bits long. Control messages are typically short so it should be remembered that the shortest Ethernet frame is 64 bytes in length.

If the network diameter is small, collision detection is faster and the resulting collision fragments are smaller. As the network diameter increases more time is lost detecting collisions and the collision fragments get larger.

Increased network diameter aggravates the collision problem. Silence on the line does not necessarily mean a distant transmitter has not already sent a packet down the cable, which will eventually result in a collision.




Figure 3—For proper operation, a collision domain must be within the maximum network diameter.

Collision Detection

A collision is defined as two stations attempting to transmit at the same time. On coaxial cable transceivers, there is circuitry to detect the DC level of the signal on the cable. This is the indicator of a collision. On fiber optic and twisted-pair interfaces with separate receive and transmit circuitry, a collision is detected by the simultaneous receiving and transmitting of data. Remember that we are discussing half-duplex Ethernet that allows either transmitting or receiving but not at the same time. Only transmitters look for collisions and it is their responsibility to reinforce a collision with a jam signal. Receivers only look for valid packets and automatically discard runt packets that are caused by collisions. Once a collision is detected by simultaneous transmitters, these transmitters will follow a backoff algorithm

Backoff Algorithm

When a collision occurs on the network, the colliding transmitters will backoff from retransmitting for a time determined by a backoff algorithm. This algorithm requires each transmitter to wait an integral number of slot times (51.2µs)

before attempting a new transmission sequence. The integer is determined by the equation:

$$0 < r < 2^k \text{ where } k = \min(n, 10)$$

The variable k is actually the number of collisions capped at a maximum of 10. Therefore, r can range from 0 to 1023 when k = 10. The actual value for r is determined by a random process within each Ethernet node. As the number of consecutive collisions increases, the range of possible backoff times increases exponentially. The number of possible retries is also capped but at 16.

For example, assume two stations A and B on the network wanting to transmit. They both wait for an absence of carrier and then wait for the IFG time to expire before initiating a transmission. It does not matter if they are 10 meters or 2500 meters apart. They could both be sensing silence and simultaneously begin to transmit causing a collision at some point. They each sense the collision and back off for either 0 or 1 slot time. The odds are 50-50 they will pick the same value and collide again. If they do, they will now back off for either 0, 1, 2 or 3 slot times. The probability of collision is now 25%. Eventually, one will win in which case its collision timer is cleared to zero while the other collision timer continues to increment until a successful transmission.

A high number of retries indicates a busy network with more stations wanting to transmit than originally assumed. That is why the backoff time range is increased exponentially to provide more possible slot times for the additional stations. At ten retries, it is assumed that 1024 simultaneous transmitters exist. This becomes the

upper bound limit of stations that can coexist on one Ethernet network. Actually this is the logical limit. Physically it may be impossible to have that many stations on one collision domain without violating cabling rules.

BACKOFF RANGE AS A FUNCTION OF COLLISIONS			
Collision on Attempt Number	Estimate of Number of Other Stations	Range of Random Numbers	Range of Backoff Times (µs)
1	1	0.....1	0.....51.2
2	3	0.....3	0.....153.6
3	7	0.....7	0.....358.4
4	15	0.....15	0.....768.0
5	31	0.....31	0.....1587.2
6	63	0.....63	0.....3225.6
7	127	0.....127	0.....6502.4
8	255	0.....255	0.....13056.0
9	511	0.....511	0.....26163.2
10	1023	0...1023	0...52377.6
11	1023	0...1023	0...52377.6
12	1023	0...1023	0...52377.6
13	1032	0...1023	0...52377.6
14	1023	0...1023	0...52377.6
15	1023	0...1023	0...52377.6
16	Too High	N/A	Discard Frame

Table 1—Backoff range increases exponentially with the number of collisions.

Channel Capture

As shown above, the Ethernet backoff algorithm provides a means for peer stations to each gain access to the network. Access is provided to all but in an unpredictable fashion. The question is if access is fair.

Assume the same two stations A and B as before. This time, however, they both have high amounts of data to send and they attempt to send at the same time and collide on the first attempt. They both back off but this time A was successful. A's collision counter is cleared but B's does not clear. If station A has more data to send and it is quick to assemble another packet to send, it might collide with B again. This time B could be selecting higher and higher backoff times as its collision counter continues to increment. However, station A feels it has only

experienced the first collision and will probably select a much lower timeout allowing it to transmit and assemble another packet and could beat station B again in the backoff contest. This phenomenon of channel capture is real and demonstrates that access to the network is neither fair nor predictable. The next time around station B could get the upper hand and limit A's access. If another station C decides to transmit as well, it could beat out station A due to the state of A's collision counter. In actuality a station that was last to arrive could transmit first.

Improving Ethernet's Determinism

There has been much discussion in the literature about implementing methods to improve the determinism of Ethernet. One approach is to incorporate a master/slave protocol such as MODBUS or OPTOMUX on top of Ethernet. In this situation, the slaves only respond to the master's commands thereby controlling the traffic on the cable and thus avoiding collisions. The downside of this approach is that you forfeit the inherent multimaster capability of Ethernet.

Another suggestion is to develop a token-passing protocol that would be implemented in Ethernet's data field. This would have to be developed and its acceptance would have to be sought. The software burden would increase and technologies such as ARCNET already can do this with built-in firmware

transparent to the application program requiring no development.

Others suggest simply increasing the data rate to 100 Mbps by using Fast Ethernet technology. By simply using raw horsepower messages will get through with or without collisions. The collision domain decreases by a factor of 10 when migrating to 100 Mbps Ethernet resulting in a maximum network diameter of only 205 meters, which is a small size network. Of course all nodes would need to be capable of communicating at 100 Mbps which could be a burden for under-powered microcontrollers.

One approach is to avoid collisions altogether by using full-duplex technology and switched hubs. In this scheme each node is paired with a port on the hub. Each node/port arrangement creates its own collision domain separate from all others. There are no collisions with a full-duplex link. The switching hub directs messages to other links by observing the destination address within the frame. Switching hubs are more expensive than non-switched hubs and they introduce more latency by their "store and forward" nature. The switching hub now becomes an integral component of the control strategy.

There is an IEEE 802.1p task group studying schemes that would provide higher priorities to the transmission of time-critical data. This activity is mainly addressing the way multicast frames are sent.




Figure 4—A switching hub, bridge or router is required to interconnect two or more collision domains.

Expanding an Ethernet Network

Expanding an Ethernet network is possible by the use of repeaters while maintaining one collision domain. If expansion is required beyond a collision domain, this can only be accomplished by the use of bridges, switches or routers. To maintain one collision domain, a symbol sent from the extreme end of the network must be able to make a complete round trip within the slot time of 512-bits (51.2µs at 10 Mbps). Calculating the complete propagation delay through adapters, AUI cables, transceivers, trunk cables and repeaters is possible but is also a challenge. Table 2 provides information on the maximum number of MAUs per segment and the maximum segment length. The maximum allowable segment length, as well as the repeaters themselves, has been assigned delay values by the 802.3 specification.

The 802.3 specification discusses ways to interconnect cable segments with repeater sets without exceeding the collision domain. A repeater set is defined as repeater electronics and two or more attached MAUs—one for each segment to be connected. The system designer can use either transmission system model 1 or transmission system model 2. Approach 2 is the detailed approach where exact delay calculations and Interframe Gap shrinkage calculations are made. Approach 1 is the simplified approach, which is not as exacting as approach 2. Approach 1 has been further simplified by creating the 5-4-3 rule.

5-4-3 Rule

The 5-4-3 rule states that a system can have up to five segments in series, with up to four repeaters and no more than three mixing segments. The remaining two segments must be link segments. A mixing segment is defined as a segment that may be connected to more than two transceivers. In other words, a bus segment. Only coaxial cable can be used for a bus segment (we are ignoring 10BASE-FP) while fiber optic and twisted-pair cable can be used as link segments. A link segment can only have two transceivers and it must support full-duplex operation (separate transmit and receive channels) to speed up collision detection. This simplified rule does not address all the possible combinations but it does yield some gross network diameters. For example, all five segments cannot be 10BASE5 or 10BASE2. If all five were 10BASE-T then the diameter would be 500 meters. With fiber optics it is different. You cannot use the maximum segment length for all five segments. In the case of 10BASE-F the maximum diameter is 2500 meters. You need to read the standard to understand this restriction.

The 5-4-3 rule does not address the three repeater configuration which yields four segments. In this

ETHERNET MAXIMUM MEDIA SEGMENT LENGTH		
Media type	Maximum number of MAUs per segment	Maximum segment length (m)
Mixing segment		
10BASE5	100	500 (trunk) 50 (AUI)
10BASE2	30	185
Link segment		
FOIRL	2	1000
10BASE-T	2	100
10BASE-FL	2	2000

Table 2—Expansion rules require that segments be identified as being either mixing or link.

case, all segments can be mixing providing a network diameter of 2000 meters for 10BASE5 and 740 meters for 10BASE2. For other configurations you need to refer to approach 2.

SUMMARY

What has been discussed is the operation of Ethernet's physical and data link layers. This alone does not implement an industrial communication network. What is needed is transport layer for reliable transfers of messages and an application layer which provides the actual control commands and responses. We will discuss these topics in the next issue of the EXTENSION.

REFERENCES

Practical Networking With Ethernet, Charles E. Spurgeon, 1997, International Thomson Computer Press

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, 1996, The Institute of Electrical and Electronic Engineers, Inc.