



Foundation Fieldbus End User Council Australia Inc.
9 Corcoran St Duncraig, WA 6023
P.O.Box Z5546 Perth, WA 6831
AUSTRALIA
ABN 60 120 236 370

FOUNDATION FIELDBUS HIGH SPEED ETHERNET (HSE) AND TCP/IP

Steve Mackay
IDC Technologies
Perth, Australia

Keywords

Fieldbus, High Speed Ethernet, HSE, Industrial Ethernet, TCP/IP

Abstract

The use of the low speed Foundation Fieldbus H1 (31.25kbits/s) has been fairly extensively documented. But an area of growing interest is that of Industrial Ethernet and now the Foundation Fieldbus HSE standard which builds on commercial off the shelf (COTS) Ethernet components. This paper examines the operation and structure of HSE with an emphasis on the lower layers (Ethernet and TCP/IP) and shows how Foundation Fieldbus takes the higher ground in the Fieldbus debate by default due to its use of Ethernet and TCP/IP which have become extremely popular. Whilst there may be a feeling that a considerable amount of this material is irrelevant to someone working with Foundation Fieldbus systems, it is the author's belief that a good knowledge of Ethernet and TCP/IP is essential in designing, troubleshooting and fixing problems with HSE based networks.

1. Introduction

The Fieldbus Foundation's HSE standard is detailed in the IEEE 802.3u specification and uses the standard Fast Ethernet to carry the Foundation Fieldbus H1 services as well as messages specifically created for HSE. The HSE standard is especially beneficial for the transfer of large files and high speed transfer of data such as between PLC's and RTU's. The emphasis on the HSE standard is to use standard off the shelf components and that all H1 functions must be preserved.

This discussion is broken down into:

- Review of the Fundamentals
- Review of Foundation Fieldbus H1 and HSE
- Description of Industrial Ethernet
- Description of TCP/IP
- Typical Challenges and advantages with Ethernet and TCP/IP

- Tying Ethernet & TCP/IP to the Foundation Fieldbus HSE Standard

2. Review of the Fundamentals

The OSI Model is briefly revisited here to ensure that Ethernet, TCP/IP and Foundation Fieldbus HSE are placed in the correct context. It should be realized at the outset that the OSI Reference Model is not a protocol or set of rules for how a protocol should be written but rather an overall framework in which to define protocols. The OSI Model framework specifically and clearly defines the functions or services that have to be provided at each of the seven layers (or levels).

The diagram below shows the seven layers of the OSI Model.

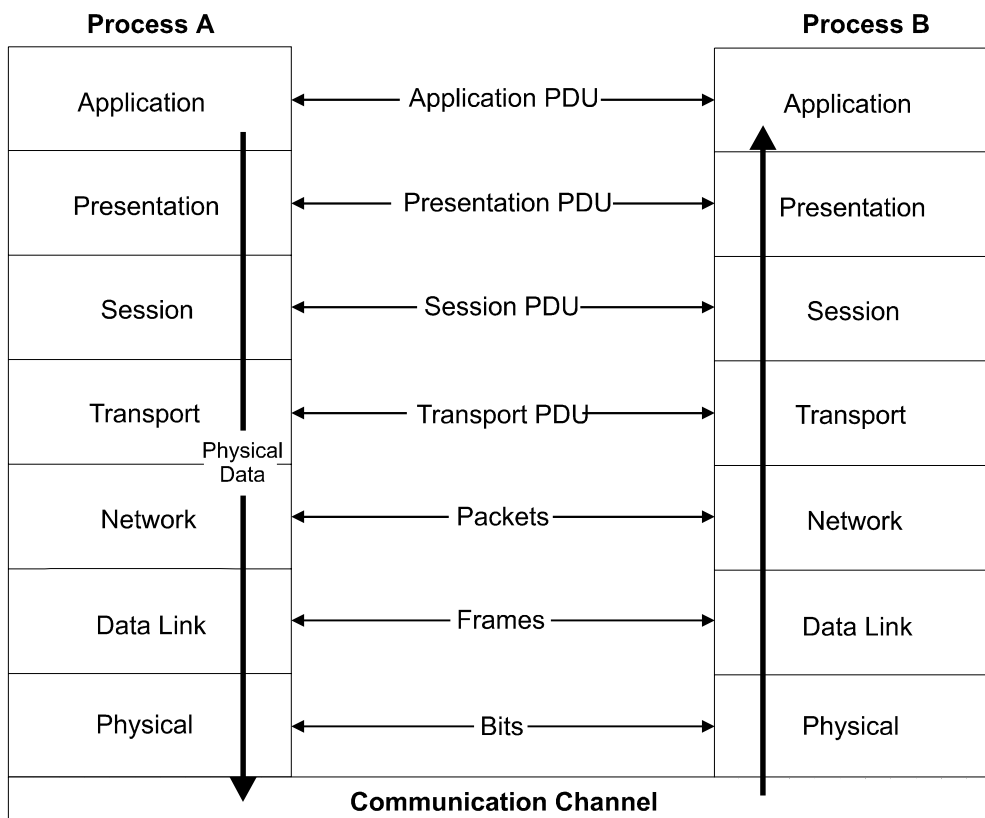


Fig 1. Full Architecture of OSI Model

A brief summary of the seven layers is as follows:

- Application - the provision of network services to the user's application programs. Note: the actual application programs do NOT reside here.
- Presentation – primarily takes care of data representation (including encryption).
- Session - control of the communications (sessions) between the users.
- Transport - the management of the communications between the two end systems.
- Network – primarily responsible for the routing of messages.
- Data Link - responsible for assembling and sending sending a frame of data from one system to another.

- Physical - Defines the electrical signals and mechanical connections at the physical level.

A layer that is “missing” from the OSI Model described above is the User Layer. This is defined specifically for the Foundation Fieldbus implementation and includes such important items such as Function Blocks.

The figure below gives an idea on how transmission of a message is effected by each layer being encapsulated within the layer below it, before it is sent out on the physical data highway. Similarly once the packet (or more strictly speaking – the frame) is received each layer is then stripped off as the packet is pushed to the top where the message is then extracted.

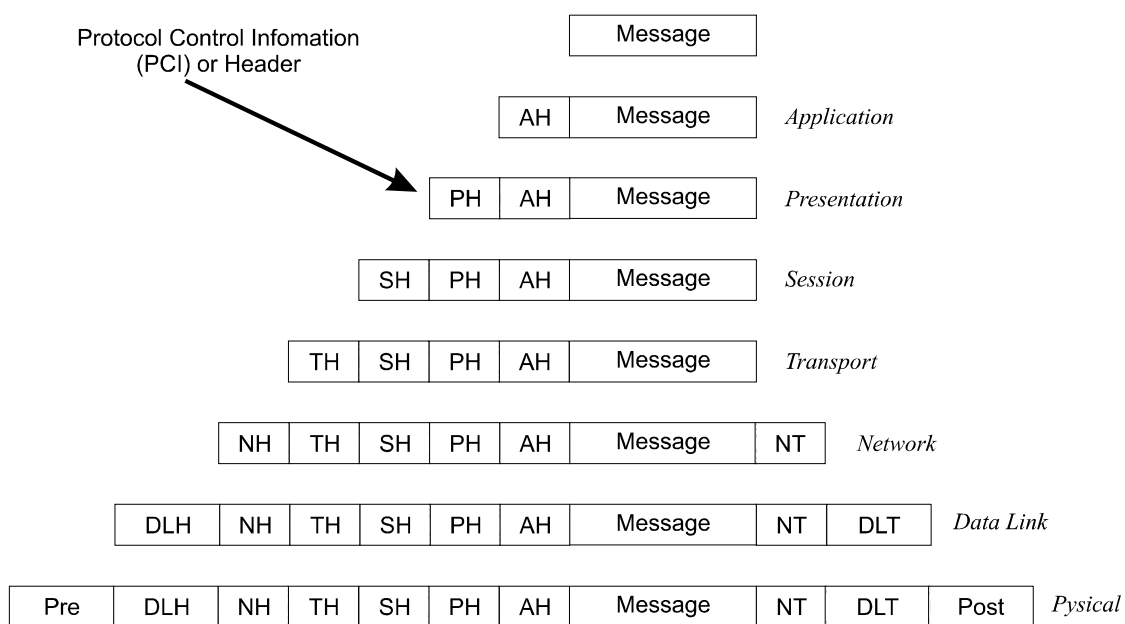


Fig 2. OSI Message Passing

A modified version of the OSI Model is used with TCP/IP and this is discussed later.

3. Review of Foundation Fieldbus H1 and HSE

3.1 H1 Technology Overview

Foundation Fieldbus H1 is a 31.25kbit/s digital multidrop communications standard connecting instruments together. It uses the layers 1,2 and 7 of the OSI Model. There is an additional User Layer (which contains the Function Blocks) but this is not defined by the OSI Model.

The H1 Physical Layer can be either bus powered or non-powered and can be used in an intrinsically safe environment. The type A cable provides distances of up to a 1900 m segment without a repeater. It is possible to extend the distance using up to 4 repeaters per segment allowing $(4+1) * 32$ devices = 160 devices maximum per logical segment.

The H1 Data Link Layer supports three types of communication. These are:

The Publisher/Subscriber method which is a scheduled approach allowing linking of inputs and outputs to allow cyclic transfer of data between the different instruments on the bus.

The Client/Server approach with unscheduled request/response is used for communications between the hosts and instruments.

Event notification or unscheduled multicast is used for alarming and trending.

The Link Active Scheduler (LAS) keeps the time synchronisation and handles the live list of communicating devices. There is also a back up so that if the LAS fails, another link master will become the LAS.

The User Layer comprises the Function Blocks, Device Descriptions and System Management. The most well known feature here is the function blocks which provide a consistent definition of inputs and outputs to allow interoperability of devices between different vendors' products.

3.2 HSE Technology Overview

This provides a high speed bus of 100 Mbit/s (or more) providing a huge increase in bandwidth and in addition providing redundancy. The HSE architecture is effectively an enhanced standard Ethernet model (IEEE 802.3). The HSE Application Layer contains the Dynamic Host Configuration Protocol (DHCP), Simple Network Time Protocol (SNTP) and Simple Network Management Protocol (SNMP). At the User Layer is the HSE Management agent and Function Blocks.

The use of redundancy is especially attractive as it uses off the shelf Ethernet hardware. Use of Ethernet Switches enables multiple paths to be set up so that failure of the primary path allows seamless transfer to the secondary path.

A more detailed description of the Ethernet and TCP/IP protocols is given in the following sections.

4. Description of Industrial Ethernet

4.1 Introduction

Ethernet uses the CSMA/CD access method described later. This gives a system that can operate with little delay, if lightly loaded but the access mechanism can fail completely if too heavily loaded. Ethernet is widely used commercially and the equipment such as Network Interface cards and switches are cheap and widely available. There were initial concerns about Ethernet being unacceptable for industrial use due to its lack of determinism (guaranteed transfer of information) and lack of robustness compared to the more reliable method of token passing. However the economics of Ethernet are so attractive compared to these other technologies that it has become the bus of choice for industrial systems.

4.2 Operation

The 802.3 standard defines a range of cable types that can be used for a network based on this standard. They include coaxial cable, twisted pair cable and fibre optic cable. The IEEE 802.3 standard defines the following:

- 10Base2 10Mbps with 185 m maximum length coaxial cable bus segment
 - 10Base5 10Mbps with 500 m maximum length coaxial cable bus segment
 - 10BaseT 10Mbps with 100 m twin cable to a central hub
 - 10BaseF 10Mbps with twin fibre bus up to 2000 m
- and of course, the Fast Ethernet standards discussed later.

4.3 Medium Access Control (MAC)

Essentially the method used for accessing the cable (or medium) is one of contention (or CSMA/CD which stands for Carrier Sense Multiple Access /Collision Detection) which is why most industrial users were initially cautious about the use of Ethernet as all transfer of data was probabilistic and not guaranteed.

In the idle state, the node merely listens to the bus monitoring all traffic that passes. If a node wishes to transmit information, it will defer while there is any activity on the bus, since this is the carrier sense component of the architecture. At some stage, the bus will become silent, and the node sensing this, will then commence its transmission. It is now in transmit mode, and will both transmit and listen at the same time. This is because there is no guarantee that another node at some other point on the bus has not also started transmitting having recognised the absence of traffic. If two nodes happen to transmit at the same time, there will be a collision of signals. The nodes' transceivers will both quickly detect this collision and will stop transmitting (after sending a brief jam signal to ensure that all nodes stop transmitting). The nodes will each wait a random time before recommencing transmission.

MAC Frame Format

The basic frame for an 802.3 network is shown below. There are eight fields in each frame, as discussed below.

Preamble	Start Delimiter	Destination Address	Source Address	Length	Data	CRC
7 Bytes	1 Byte	2 or 6 Bytes	2 or 6 Bytes	2 Bytes	46 - 1500 Bytes	4 Bytes

Fig 3. MAC Frame Format.

Preamble

This field consists of 7 octets of the data pattern 1010101010, The preamble is used by the receiver to synchronise its clock to the transmitter.

Start Frame Delimiter

This single octet consists of the field 10101011. It enables the receiver to recognise the commencement of the address fields.

Source and Destination Address

These are the physical addresses of both the source and destination nodes. The fields are generally 6 octets long. This means that there are 2^{48} possible physical addresses.

Length

A two octet field that contains the length of the data field. Blue Book Ethernet defines this field to be the Protocol Type; whilst the 802.3 standard uses this field to indicate the length.

Information

The information that has been handed down from the layer above. It is generally a protocol packet.

Pad

There is a minimum length of 64 octets that must be transmitted to ensure that the collision mechanism works. If there is less data than this, this field will be padded out with random octets to ensure that there is a minimum of 64 octets transmitted.

Frame Check Sequence (FCS)

A 32-bit CRC value that is computed in hardware at the transmitter and appended to the frame.

4.4 Fast and Gigabit Ethernet Systems

Although Ethernet with an estimated 200 million installed nodes world-wide is the most popular method of linking computers on a network, its 10MBps speed is too slow for very data intensive or real time applications. Hence Foundation Fieldbus choice of the higher speed Ethernet Systems of Fast and Gigabit Ethernet systems which are described below. The reasons for the success of Fast and Ethernet systems against competing "Ethernet" offerings have been their backwards compatibility with the earlier 10Mbit/s standards.

Fast Ethernet 100 Base-T (100BaseTX and 100BaseFX)

Although there are other 100 Mbit/s approaches possible, these are the preferred method. This uses the existing Ethernet MAC Layer with various enhanced Physical Media Dependent (PMD) layers to improve the speed. There are described in the IEEE 802.3u as follows:

IEEE 802.3u defines three different versions based on the physical media:

- 100Base-TX which uses two pairs of Category 5 UTP or STP
- 100Base-T4 which uses four pairs of Category 3,4 and 5 UTP (not popular)
- 100Base-FX which uses multimode or single mode fibre optic cable.

Fast Ethernet provides a transmission speed of 100 Mbps, ten times faster than that of “ordinary” Ethernet. It does, however, retain the same frame format. It is described by two standards, namely IEEE 802.3u and IEEE 802.3y.

IEEE 802.3u defines three different versions based on the physical media namely 100Base-TX (which uses two pairs of Category 5 UTP or STP), 100Base-T4 (which uses four pairs of wires of Category 3,4 or 5 UTP) and 100Base-FX (which uses multimode or single-mode fiber optic cable).

IEEE 802.3y, on the other hand, defines 100Base-T2 which uses two pairs of wires of Category 3,4 or 5 UTP.

One of the limitations of the 100Base-T systems is the size of the collision domain, which is 250m. This is the maximum sized network in which collisions can be detected, being one tenth of the size of the maximum 10 Mbps network. This limits the distance between our workstation and hub to 100m, the same as for 10 Base-T, but usually only one hub is allowed in a collision domain. This means that networks larger than 200m must be logically connected together by store and forward type devices such as bridges, routers or switches. However, this is not a bad thing, since it segregates the traffic within each collision domain, reducing the number of collisions on the network. The use of bridges and routers for traffic segregation, in this manner, is often done on industrial CSMA/CD networks.

The dominant 100Base-T system is 100Base-TX which accounts for about 95% of all Fast Ethernet shipments.

Gigabit Ethernet

Gigabit Ethernet uses the same 802.3 frame format as 10Mbps and 100Mbps Ethernet systems. It operates at ten times the clock speed of Fast Ethernet at 1Gbps. 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 three different physical layers: 1000Base-LX and 1000Base-SX using fiber and 1000Base-CX using copper.

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 (the time needed to transmit a minimum-sized frame) of 64 bytes used with both 10Mbps and 100Mbps 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 20metres!

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 $4.096\mu\text{s}$ to compensate for the tenfold speed increase. This then gives a collision domain of about 200m.

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. This is shown in Figure 4.

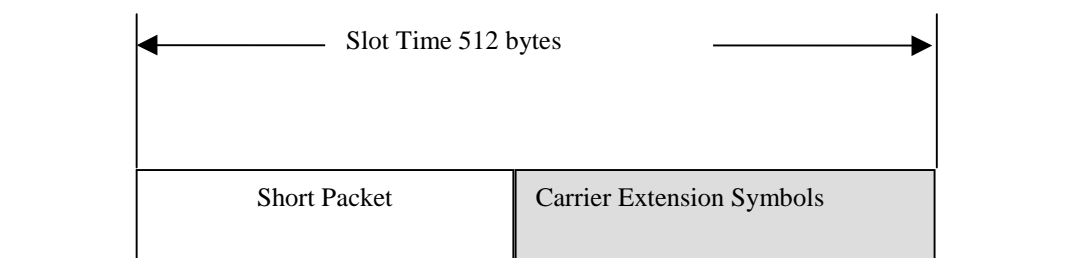


Fig 4. Carrier Extension

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 5.

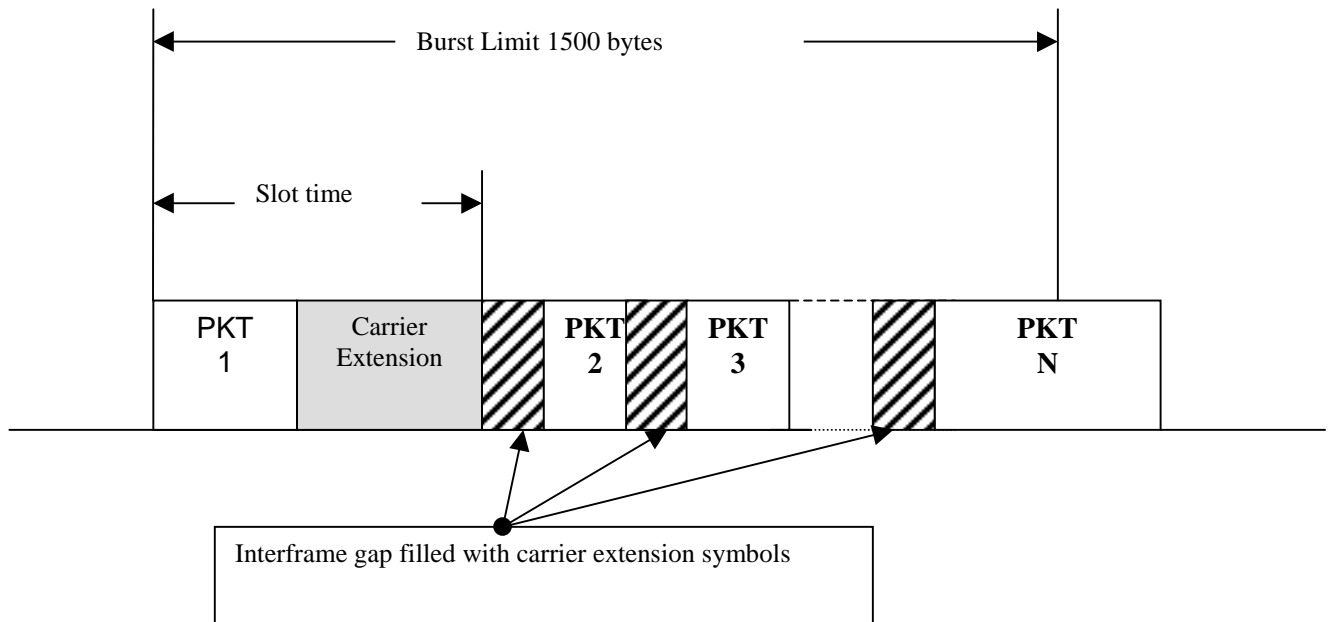


Fig 5. Packet Bursting

5. Description of TCP/IP

The next layers above the Ethernet Level are:

IP Internet Protocol
 TCP Transmission Control Protocol

5.1 Introduction

TCP/IP is the *de facto* global standard for the Network and Transport layer implementation of internetwork applications because of the popularity of the Internet. The Internet (or in its early years known as ARPANet), was part of a military project commissioned by the Advanced Research Projects Agency (ARPA), later known as the Defence Advanced Research Agency or DARPA. The communications model used to construct the system is known as the ARPA model.

Whereas the OSI model was developed in Europe by the International Standards Organisation (ISO), the ARPA model (also known as the DoD model) was developed in the USA by ARPA. Although they were developed by different bodies and at different points in time, both serve as models for a communications infrastructure and hence provide “abstractions” of the same reality. The remarkable degree of similarity is therefore not surprising.

Whereas the OSI model has 7 layers, the ARPA model has 4 layers. The OSI layers map onto the ARPA model as follows.

- The OSI Session, Presentation and Applications layers are contained in the ARPA Process and Application Layer.
- The OSI Transport Layer maps onto the ARPA Host-to-Host Layer (sometimes referred to as the Service Layer).

- The OSI Network Layer maps onto the ARPA Internet Layer.
- The OSI Physical and Data Link Layers map onto the ARPA Network Interface Layer.

OSI LAYER	PROTOCOL IMPLEMENTATION						ARPA LAYER
APPLICATION	File Transfere	Electronic Mail	Terminal Emulation	File Transfer	Client/Server	Network Management	PROCESS AND APPLICATION
PRESENTATION	File Transfer Protocol (FTP)	Simple Mail Transfer Protocol (SMTP)	TELNET Protocol	Trivial File Transfere Protocol (TFTP)	Sun Microsystems. Network file Systems Protocol (NFS)	Simple Network Management Protocol (SNMP)	
SESSION	MIL-STD 1780 RFC 959	MIL-STD 1781 RFC 821	MIL-STD 1782 RFC854	RFC 783	RFC's 1014, 1057 & 1094	RFC 1157	
TRANSPORT	Transmission Control Protocol (TCP) MIL-STD 1778 RFC 793			User Datagram Protocol (UDP) 768		RFC	HOST TO HOST
NETWORK	Address Resolution ARP RFC 826 & RARP RFC 903		Internet Protocol (IP) MIL STD 1777 & RFC 791		Internet Control Message Protocol (ICMP) RFC 792		INTERNET
DATA LINK	Network Interface Cards: Ethernet, Token-Ring, ARCNET, MAN and WAN. RFC 894, 1042, 1201 and others						NETWORK
PHYSICAL	Transmission Media: Twisted pair cable, Coaxial Cable, Fiber Optics, Wirless Media etc. etc.						INTERFACE

The relationship between the two models is depicted in the following figure.

Fig 6. OSI vs. ARPA Models

TCP/IP, or rather- the TCP/IP Protocol Suite- is not limited to the TCP and IP protocols, but consist of a multitude of interrelated protocols that occupy the upper three layers of the ARPA model. TCP/IP does NOT include the bottom Network Interface Layer (typically Ethernet defines this), but depends on it for access to the medium.

5.2 The Internet Layer

This layer is primarily responsible for the routing of packets from one host to another. Each packet contains the address information needed for its routing through the internetwork to the destination host. The dominant protocol at this level is the Internet Protocol (IP).

IP is responsible for the delivery of packets (“datagrams”) between hosts. It is analogous to the postal system, in that it forwards (routes) and delivers datagrams on the basis of IP Addresses attached to the datagrams, in the same way the postal service would process a letter based on the postal address. The IP Address is a 32-bit entity containing both the network address (the “zip code”) and the host address (the “street address”).

The IPv4 address consists of 32 bits, e.g. 11000000011001000110010000000001. Since this number is fine for computers but a little difficult for human beings, it is divided into four octets w, x, y and z. Each octet is converted to its decimal

equivalent. The result of the conversion is written in the format 192.100.100.1. This is known as the “Dotted Decimal” or “Dotted Quad” Notation. As mentioned earlier, one part of the IP address is known as the Network ID or “NetID” while the rest is known as the “HostID”. Originally, IP addresses were allocated in so-called Address Classes. Although the system proved to be problematic, and IP addresses are currently issued “classless”, the legacy of IP address classes remains and has to be understood.

There are, however, several other additional protocols required at this level such as:

- Address Resolution Protocol (ARP), RFC 826. This is used for the translation of an IP address to a hardware (MAC) address, such as required by Ethernet.
- Reverse Address Resolution Protocol (RARP), RFC 903. This is the complement of ARP and translates a hardware address to an IP address.
- Internet Control Message Protocol (ICMP), RFC 792. This is a protocol used for exchanging control or error messages between routers or hosts (eg the Ping command).

5.3 The Host-to-Host Layer

This layer is primarily responsible for data integrity between the sender host and receiver host regardless of the path or distance used to convey the message. It has two protocols associated with it, namely:

User Datagram Protocol (UDP)

The User Datagram Protocol (UDP) is a “connectionless” protocol and does not require a connection to be established between two machines prior to data transmission. It is therefore said to be “unreliable” - the word “unreliable” used here as opposed to “reliable” in the case of TCP and should not be interpreted against its everyday context.

Sending a UDP Datagram involves very little overhead in that there are no synchronization parameters, no priority options, no sequence numbers, no timers, and no retransmission of packets. The header is small, the protocol is streamlined functionally. The only major drawback is that delivery is not guaranteed. UDP is therefore used for communications that involve broadcasts, for general network announcements, or for real-time data.

Transmission Control Protocol (TCP)

Transmission Control Protocol (TCP) is a connection-oriented protocol and is said to be “reliable”, although this word is used in a data communications context. TCP establishes a session between two machines before data is transmitted. Because a connection is set up beforehand, it is possible to verify that all packets are received on the other end and to arrange re-transmission in case of lost packets. Because of all these built-in functions, TCP involves significant additional overhead in terms of processing time and header size.

TCP fragments large chunks of data into smaller segments if necessary, reconstructs the data stream from packets received, issues acknowledgements of data received, provides socket services for multiple connections to ports on remote hosts, performs packet verification and error control, and performs flow control.

The Source and Destination Ports (16 bits each) identify the host processes at each side of the connection. Examples are Post Office Protocol (POP3) at port 110 and Simple Mail Transfer Protocol (SMTP) at port 25. Whereas a destination host is identified by its IP address, the process on that host is identified by its port number. A combination of port number and IP address is called a Socket.

5.4 The Process and Application Layer

This layer provides the user or application programs with interfaces to the TCP/IP stack. Protocols at this level include (but are not limited to) File Transfer Protocol (FTP), Trivial File Transfer Protocol (TFTP), Simple Mail Transfer Protocol (SMTP), Telecommunications Network (TELNET), Post Office Protocol (POP3), Remote Procedure Calls (RPC), Remote Login (RLOGIN), Hypertext Transfer Protocol (HTTP) and Network Time Protocol (NTP) and other Foundation Fieldbus specific programs.

6. Typical Challenges and advantages with Ethernet and TCP/IP

As a result of the above discussions the typical challenges with Ethernet from an industrial perspective are outlined below:

- Lack of determinism
- No Loop Powering of instruments
- High power consumption of interface cards
- No intrinsic Safety Ability
- Large Overhead in the frame and packet sizes
- A large number of the protocol fields are not necessarily relevant to Foundation Fieldbus requirements

But the advantages are easily seen to be:

- High Speed Communications Link of 100 Mbit/s
- Redundancy of the system
- Off the shelf Components
- Easy interface into the remainder of the company IT infrastructure due to use of Ethernet and TCP/IP
- Low cost system due to popularity of Ethernet and TCP/IP.

7. Tying Ethernet & TCP/IP to the Foundation Fieldbus HSE Standard

Although HSE is an enhanced standard Ethernet (802.3) model, it can still use the typical Ethernet & TCP/IP hardware and software components available on the market. HSE is a fault tolerant network allowing redundancy and multiple H1 Fieldbus segments to be interconnected with speeds up to 1 Gbit/s.

An excellent understanding of Ethernet and TCP/IP helps one to install and troubleshoot HSE Networks effectively and efficiently. It is hoped that the previous discussion has gone some way in describing the underlying layer Ethernet and TCP/IP layers and thus in helping one to achieve this.

References

Mackay, S., D. Reynders and E. Wright, "*TCP/IP and Ethernet Networking*" (2000).

Vincent, S.J., "*Foundation Fieldbus High Speed Ethernet Control System*" (2001).