

---

## ***Chapter 8***

### ***Local Area Networks - Applications And Standards***

#### ***A Summary...***

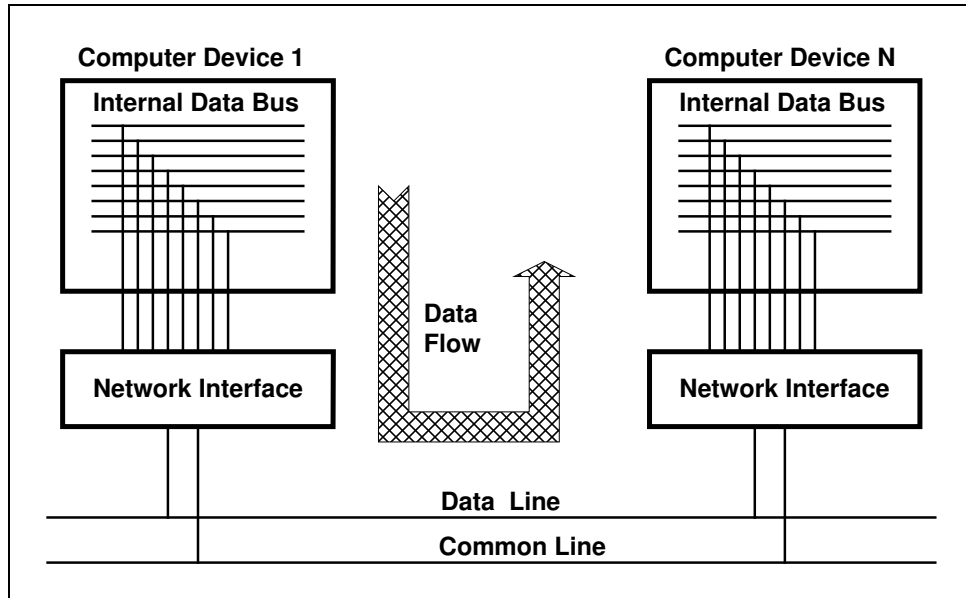
Interfacing to networks. Standards activities in Local Area Networks - the IEEE "802" Committee. The Ethernet network. MAP and TOP networks. File Server systems.

#### ***Read This Chapter If...***

- ◆ You would like to learn about the most common LAN systems
- ◆ You would like to learn about file-servers
- ◆ You need to know how emerging networking standards will affect industry plans for integration

## 8.1 Interfacing Computers to Networks

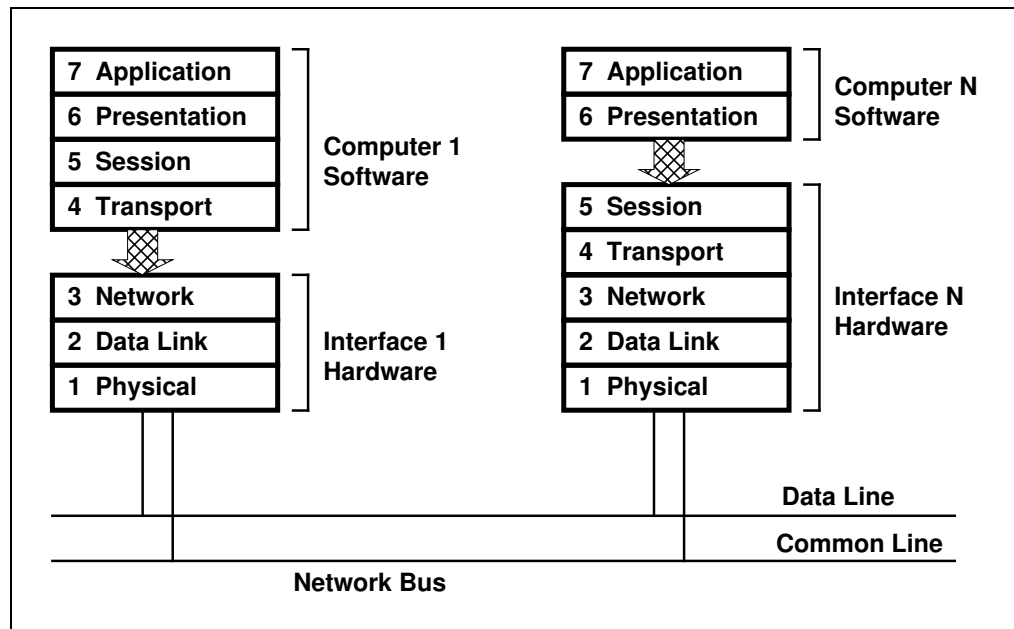
Interfacing the internal computer data bus structure to an external network, that is based upon the OSI 7 layer model, is a complex task. At the most basic level, the interface to the network generally needs to perform parallel/serial and serial/parallel conversion. However an interface may additionally need to perform signal modulation and demodulation, contention resolution, data framing, error checking and correction and other sundry functions. The role of the interface is shown schematically in Figure 8.1.



*Figure 8.1 - Interfacing a Data Bus to a Network Bus*

The exact role of the network interface depends as much on the interaction with its local computer as on the standards chosen for the OSI model. In other words, the network interface may fulfil some or all of the 7 functional layers which are needed to connect a computer to an OSI network. All the functional layers that are not provided by the network interface must be performed by the computer system in software.

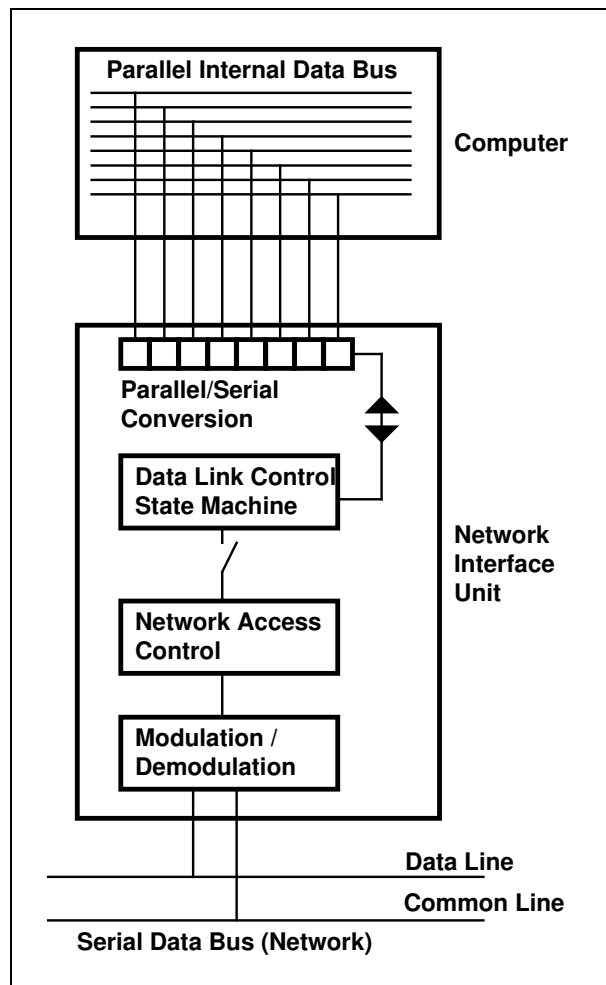
Provided that each computer, its software and interface card are all matched, then the distribution of OSI layer functions (between hardware and software) is transparent to the network. This is shown schematically in Figure 8.2.



*Figure 8.2 - OSI Networking of Devices using Different Interfacing Structures*

After examining Figure 8.2, one may be tempted to ask why all 7 layers of the OSI model are not always implemented "in hardware" on a network interface card. In other words, why should the computer system be burdened with any of the networking layer functions? In theory this would appear to be the sensible solution and in some cases it is actually achieved. However, the complexity of 7 layer protocols may be such that "complete" interfaces can only be implemented using specialised VLSI circuitry. The costs involved in circuit development, coupled with the enormous range of available networks and protocols, have tended to discourage interface manufacturers from providing 7-layer hardware solutions. It is therefore common to find interfaces which provide only the first two layers in hardware. The upper layers are often supplied to users in the form of pre-compiled software libraries which can be linked into end-user application programs.

Despite the wide range of network protocols, there is generally a high degree of commonality between the lower layers (1 and 2) of the OSI model. For example, HDLC is a very common specification for the data link layer of a communications protocol - it is used within a number of different networks. The CSMA/CD contention scheme is another specification that commonly forms part of the physical layer of a number of different networks. There are also a number of modulation and demodulation techniques that are widely used within the physical layer. It is therefore commercially prudent for Original Equipment Manufacturers (OEMs) to implement only the lower layers of a network interface in hardware. These lower levels of the interfacing unit are shown schematically in Figure 8.3.



*Figure 8.3 - Implementing the Lower Levels of a Network Interface in Hardware*

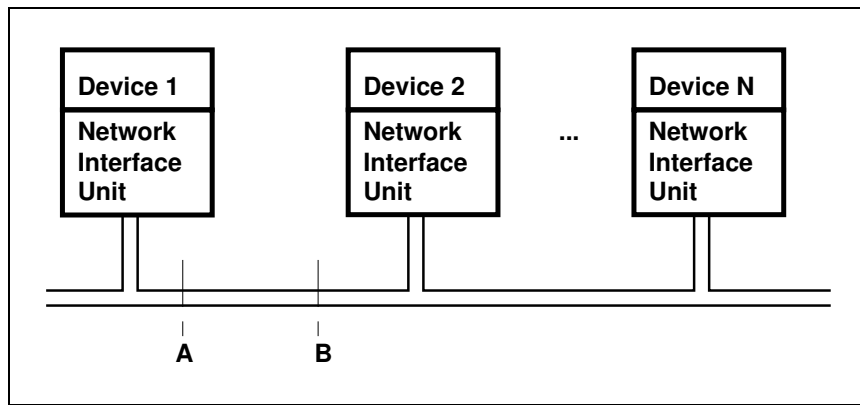
It should be clear from Figure 8.3 that a fundamental requirement of the interfacing hardware is that it must be tailored to each, different computer bus architecture. In the case of general-purpose computers, internal bus architectures have evolved into a manageable number of common systems, so that the problem of diversity is not as great as one might imagine. A limited range of common bus structures are also becoming well accepted in the industrial environment.

However, once we go above the network layer of the OSI model, we tend to find that the specifications for each layer are as diverse as the networks themselves. Additionally, we find that the role of the higher layers is inherently more complex than that of the lower layers. These factors have, in the past, provided good reasons for implementing the upper layers entirely in software on the computers themselves.

The distribution of functions between a network interface and a computer sounds much simpler in theory than it turns out to be in practice - particularly in the industrial environment. The key reason for this is "programmability". Many industrial control systems (most notably CNCs) are marketed in a "closed architecture" form. Such systems are normally only programmable in terms of machine hardware and not for normal data processing and Input/Output functions. Therefore, many industrial systems cannot be "software-adapted" to provide the upper layer functionality needed for OSI communications. Since it is often equally impractical for interface manufacturers to provide the equivalent network functionality in hardware, the net result has been that a substantial proportion of industrial controllers have been unable to connect to a factory bus network. As a result, star networks composed of point to point RS-232 links (each with their own protocol) have remained in place for a good deal longer than has been desirable.

## 8.2 Network Performance - Transfer Rates

Now that we have examined some of the basics of connecting to networks, it is appropriate to examine some of the key issues related to the performance of the LANs themselves. We will use the typical bus network of Figure 8.4 as the basis of our discussion.



*Figure 8.4 - A Bus Type Communications Network*

The most commonly quoted figure for data communications performance is the "link speed" in terms of the bit-rate. This figure can be somewhat misleading if one does not fully appreciate its limited scope. In terms of Figure 8.4, the link speed defines the number of bits per second which flow along the network bus (between A and B say) when data is actually flowing.

In itself, the link speed does not define the average rate at which messages are actually transferred between two network nodes. In other words, the time taken to transfer a file (say) from one node to another is NOT equal to:

$$\frac{\text{file size (in bits)}}{\text{link speed}}$$

There are clearly a number of other factors which will determine the rate at which a message can actually be transferred from one device to another.

The appropriate quantity to use in defining the performance of a network, is called the "Link Response Time" (LRT). This quantifies the amount of time taken to transfer a meaningful message from one network node to another and to receive an acknowledgment message. The factors which influence the LRT include:

- Link speed
- Time taken for a node to access the network medium (contention)
- Delays involved in each device preparing the data into suitable packets or frames,
- Size of network data frames (including overheads such as error checking bits, addressing bits, etc.)
- Number of times a frame has to be re-transmitted in the event that a transmission error has been detected
- Time taken for a receiving node to send an acknowledgment message
- Signal transfer time through the modem in a network interface unit
- the number of frames in each message.

Many of these factors are governed by the protocols that are in use on the network.

In the final analysis, the delays caused by contention for use of the network medium could have the greatest effect on the Link Response Time. This is especially true in the case of the non-deterministic CSMA/CD scheme, where repeated collisions during a busy period could seriously affect the time taken for the transfer of a single message. Moreover, the fact that the CSMA/CD system is non-deterministic makes it difficult to make meaningful calculations on the LRT. On the other hand, in token-passing contention schemes, the maximum delay in accessing the network medium is known and therefore a meaningful value for the LRT can be determined.

The quality of the transmission medium is another factor which will affect the performance of the network. A link which is affected by electro-magnetic interference will require frequent re-transmission of erroneous data frames, thus degrading its response time. The overheads involved in "framing" or "packetising" data for transmission over a network can also affect the Link Response Time. If the actual data field of a transmitted frame is small, then the framing bits (including error checking) can constitute the major proportion of the total number of bits transmitted in a frame. This factor is dependent upon the data link level protocol in use and has an influence on efficiency.

Over and above all of these factors, we need to consider the "device latency" of the network nodes themselves. In other words, the amount of time that it takes each node to gather data (say from disk storage) and assemble it into frames for transmission. Similarly, the amount of time needed to disassemble frames and process incoming data may also be of significance.

The diagram in Figure 8.5 shows some of the time delays involved in transmitting a frame from one node to another. We can derive simple formulae from diagrams such as these, which will allow us to sensibly determine the Link Response Time (and performance) of a computer network.

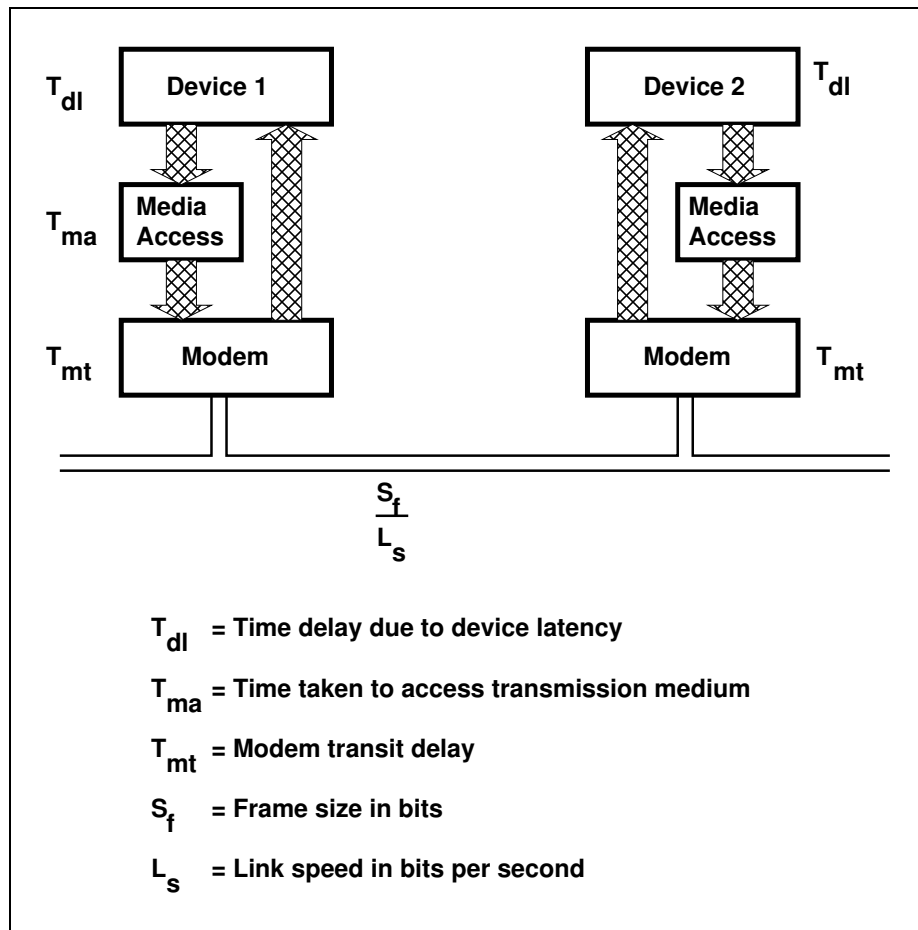


Figure 8.5 - Delays Involved in Transmitting Node to Node

Using information, such as that shown in Figure 8.5, we can attempt to assess the Link Response Time (LRT) of a particular network. In the assessment that follows, we make the assumption that the two communicating nodes are identical and generate the same time delays. In general, this will not be the case and we would need to substitute different parameters for each device.

As a starting point, we can see that the time taken to transmit a frame, of size  $S_t$  from one device to another is given by:

$$\begin{aligned} T_{trans} &= T_{dl} + T_{ma} + T_{mt} + \frac{S_t}{L_s} + T_{mt} + T_{dl} \\ &= T_{ma} + 2 \cdot T_{dl} + 2 \cdot T_{mt} + \frac{S_t}{L_s} \end{aligned}$$

Similarly, the time taken for the receiving node to transmit an acknowledgment (or negative acknowledgment) frame of size  $S_a$  is:

$$T_{ack} = T_{ma} + 2 \cdot T_{dl} + 2 \cdot T_{mt} + \frac{S_a}{L_s}$$

The time taken for the complete transmission and acknowledgment sequence is then:

$$T_{seq} = T_{trans} + T_{ack}$$

If erroneous frames are detected, then  $N_t$  sequences may be required to successfully send one frame:

$$T_{frame} = N_t \cdot T_{seq}$$

A complete message may be comprised of  $N_f$  frames and therefore the Link Response Time is:

$$LRT = N_f \cdot N_t \cdot T_{seq}$$

where

$$T_{seq} = \frac{S_t + S_a}{L_s} + 4 \cdot T_{dl} + 4 \cdot T_{mt} + 2 \cdot T_{ma}$$

Once formulae, such as those above, are established for a specific network, it is possible to make assessments on the realistic (best and worst) times required for transferring files and other information through that network.

### 8.3 Networks Standards Activities - IEEE 802 Committee

We have already seen that there are many networking standards available to fulfil each layer of the Open Systems Interconnection model. Unfortunately, as fate would have it, there are also many, different organisations which are responsible for establishing these standards. Some of the standards for data communications are established by special interest groups and professional engineering bodies, whilst others are established by federal governments in different countries.

Many standards organisations choose to adopt standards defined by other standards organisations, particularly where the cost of establishing complex standards is very high. This is true in many aspects of data communications and so we have a number of standards organisations specifying identical standards. However, each standards body likes to classify standards according to its own coding system. The net result is that it is often difficult to determine which standards are identical in structure and different in name only. A typical example of this is the RS-232C standard, which is also known by its "V.24" nomenclature.

We will now examine a few of the major standards organisations which develop and issue standards related to data communications.

On a global basis, the major standards body is the International Standards Organisation or ISO. This is made up of representatives from the national standards bodies of participating countries. Two of the most notable members of the ISO include the American National Standards Institute (ANSI) and the British Standards Institute (BSI). Like many standards bodies, the ISO is divided into a number of technical committees, which produce standards for a diverse range of entities. The ISO committee which is responsible for information processing systems is known as Technical Committee 97 or TC97. As previously noted, the ISO was responsible for the development of the 7 Layer, Open Systems Interconnection model (which is referred to as ISO82). The ISO is also responsible for the establishment of protocol standards for layers within the OSI model.

The Comité Consultatif Internationale de Telegraphie et Telephonie (CCITT) is another international standards body concerned with data communications. The major role of the CCITT is in the establishment of standards related to public communications networks (rather than Local Area Networks). In particular, CCITT is involved in standards for Public Switched Telephone Networks (PSTN) and Integrated Services Digital Networks (ISDN).

At a federal level, the United States has a number of major standards bodies in addition to the American National Standards Institute. These include the National Bureau of Standards (NBS), which provides standards for information processing equipment purchased by the US government, the Electronics Industries Association (EIA) and the Institute of Electrical and Electronic Engineers (IEEE). The IEEE is a professional engineering body and the EIA is a special interest group established by American electronics organisations. Both these national groups have made significant contributions to the development of standards for Local Area Networks. These standards are often adopted by ANSI and the NBS.

Finally, one European version of the EIA is called the European Computer Manufacturers Association (ECMA). Another special interest group, the ECMA is primarily comprised of members who manufacture computers within Europe.

The data communications standards which are adopted in the majority of countries around the world are based upon a combination of those laid down by:

- *ANSI*
- *BSI*
- *CCITT*
- *ECMA*
- *EIA*
- *IEEE*
- *ISO*
- *NBS.*

A large number of local standards are little more than re-titled versions of standards generated by the above-named bodies.

As we have already noted, Local Area Networks and interfacing are vital to manufacturing data communications. Since a great deal of standardisation work has been carried out for the two, lower layers of the OSI model, it is worthwhile to examine these standards a little more closely. The IEEE committee, referred to as the "802 Committee" was established to work towards the development of LAN standards. Of particular interest are those standards which describe the first (physical) and second (data link) layers of the OSI model.

To begin with, the IEEE has defined the 802.1 standard, which is referred to as the Higher Layer Interface Standard or "HILI". This standard is used to embrace those specifications which are chosen for the physical and data link layers of the OSI model and ensure that their development leads to compatibility with the upper 5 layers.

As far as the physical layer of the OSI model is concerned, the IEEE has defined network topologies and contention schemes with which we are already familiar. That is:

- CSMA/CD Bus (IEEE 802.3)
- Token Passing Bus (IEEE 802.4)
- Token Ring (IEEE 802.5)

These standards define the network topology, contention schemes, communications media, allowable modulation types, physical connectors and plugs, etc.

Up until now we have only looked at the data link layer of the OSI model as a single unit. As it turns out, the data link layer can actually be broken down into two sub-layers. These sub-layers are for:

- Logical Link Control
- Media Access Control.

The Logical Link Control sub-layer (abbreviated LLC) is the one which is actually responsible for framing data bits and maintaining data integrity through the use of Cyclic Redundancy Checks, etc. It is this sub-layer which is described by protocols such as HDLC, BiSync, etc.

The IEEE has defined another standard known as IEEE 802.2 to fulfil the requirements for the Logical Link Control sub-layer. This standard provides for two different optional modes of operation for:

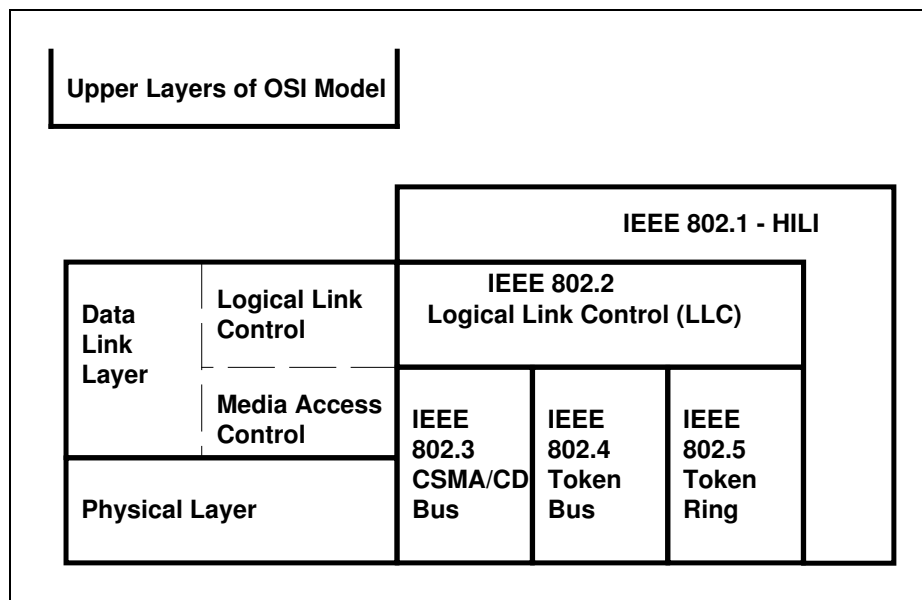
- Connectionless Networks
- Connection-Oriented Networks.

In the connectionless (packet-switched) network mode, the data link layer of the OSI model assumes that the higher layers will provide error control, flow control and sequencing (refer to section 7.7). This connectionless mode can also be used in situations where data integrity is not vital. The connectionless mode of operation is essentially a "lazy" mode of framing data. It minimises demands on hardware in situations where error checking is better performed by other layers.

The connection-oriented mode of the IEEE 802.2 protocol is based upon the HDLC protocol, and generates very similar data frames that contain Cyclic Redundancy Checks (Frame Check Sequences), sequencing information, etc. As its name implies, the system is intended for connection-oriented networks (circuit-switched systems) where the upper layers of the OSI model do not provide error checking of their own. It is also used when a high degree of error checking is required. This mode of operation clearly requires a more sophisticated piece of interfacing hardware in order to function.

As its name implies, the Media Access Control sub-layer (abbreviated MAC) is the one which is responsible for resolving contentions for use of the network media. The specification for the OSI physical layer will determine the type of contention scheme in use and therefore the MAC sub-layer must be matched to that specification. We now find (seemingly contrary to the original aims of the OSI model) that the IEEE 802.3, 802.4 and 802.5 standards effectively cover the entire physical layer plus the MAC sub-layer of the data link layer. As it turns out, despite the modularised intentions of the "layer" system, the lower three layers of the OSI model are interdependent.

The structure of the major IEEE standards and their relationship to the OSI model is shown in Figure 8.6. Another commonly cited IEEE communications standard is IEEE 802.6. However, this is not shown in the diagram because the standard governs Metropolitan Area Networks and is outside the scope of our discussions.



*Figure 8.6 - IEEE 802 Networking Standards*

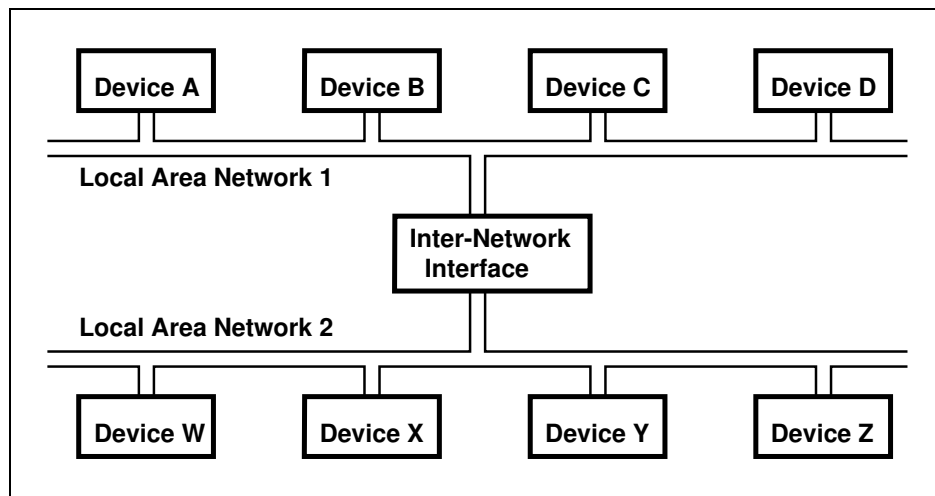
It can be seen from our discussion of IEEE 802.2, that even when we have a single standard, such as this, there is still scope for a number of different options within the specification. Another example is the IEEE 802.4 standard, which specifies a number of different, optional cable types and modulation schemes. We must now note, that in order to make devices communicate with one another on a network, it is not sufficient for them to both simply adhere to the same layer standard - they must also be configured to the same options within each standard specification.

We will look at IEEE specifications a little more closely as we progress through this chapter and see how they are applied within emerging networking standards.

## 8.4 Bridges, Routers and Gateways

We have thus far noted well that standardisation of communications networks has been difficult to realise because of an early lack of support from major computer and control systems manufacturers. We have also noted that differing environmental requirements make it difficult to have one network standard because no single, configuration is optimal for all environments. We therefore need to accept that, in the foreseeable future, we will have the prospect of trying to make different networks communicate with one another.

It is a difficult enough task to make a number of devices meaningfully communicate with one another on a single network. It is even more difficult to make devices which are attached to totally different networks communicate with one another. Schematically, the dilemma is shown in Figure 8.7.




*Figure 8.7 - Interconnecting Different Networks*

If we assume that networks 1 and 2 of Figure 8.7 are totally different in terms of framing, contention, modulation, etc., then we can see the enormity of the problems for the inter-network interface device. A typical example might be where Device B (network 1) wishes to communicate with Device Y (network 2). The inter-network device must simultaneously satisfy the physical requirements and access requirements of both networks, whilst performing a translation of data frames from one form to another. If the two networks are OSI based, then at least there is some scope for breaking this problem down into layers of compatibility and functionality.

There are three, commonly used devices which perform the role of the inter-network interface between OSI systems. These are referred to as:

- Gateways
- Routers
- Bridges.


In physical terms all these devices should be considered as computers with varying degrees of sophistication. The most complex of these systems is the Gateway, which is designed to translate all seven layers from the protocol of one network to the seven layers required by another network. This is shown schematically in Figure 8.8.



*Figure 8.8 - A Gateway Between Two Different OSI Networks*

The upper layers of an OSI network are relatively complex (particularly in the applications layer) and therefore the amount of time required for a Gateway computer to perform a protocol translation may be significant. In addition to the inter-network time delays involved with Gateways, their cost is substantial and clearly they are a "last alternative" solution for communication between totally dissimilar networks.

If the upper layers (4, 5, 6 and 7) of two, OSI networks are the same, then it is possible to use a "Router" to perform protocol translation for the lower 3 layers. Routers can be used to connect a number of such networks together at a common point as shown in Figure 8.9. Packets are "routed" from one network to another based upon the destination address specified within the network layer (3) of the packet.

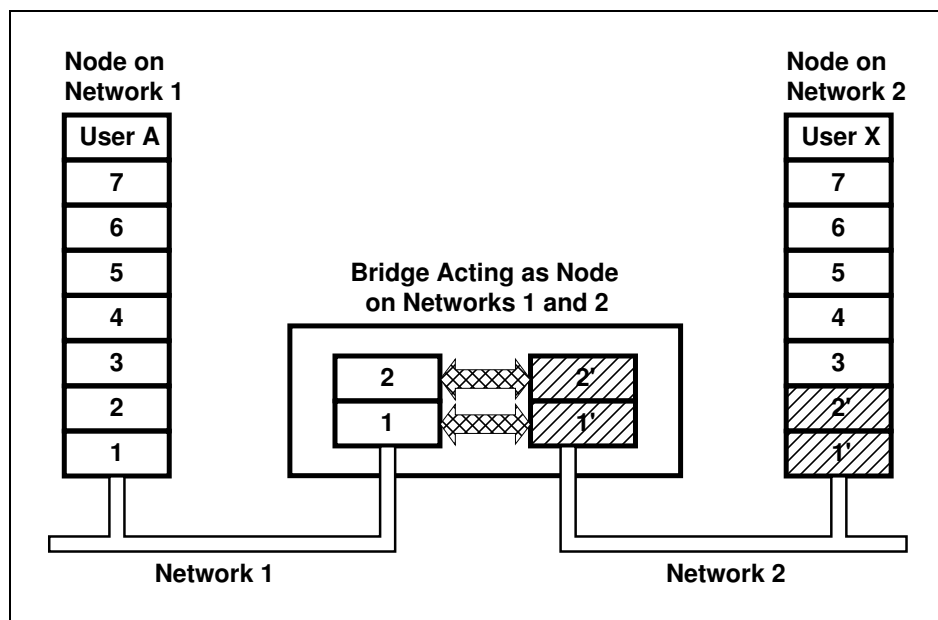


*Figure 8.9 - Routing Networks at a Common Point*

The OSI networks which are easiest to interconnect are those which are completely identical or those which differ only in the lower one or two layers. In these situations a Bridge can be used to interconnect the systems as shown schematically in Figure 8.10.

We now appreciate that whilst in theory it may be desirable to have all devices connected to a common network, in reality this is unlikely to occur. We realise that environmental and performance features are the major reasons for using different network philosophies (leaving aside vendor marketing strategies). We also appreciate that the easiest way to interconnect networks is to ensure that the networks are as similar as possible.

The physical and data link layers of the OSI model are the ones which are most closely related to environmental factors. As a corollary, it is therefore logical to attempt to achieve common, upper layer standards for the OSI model, whilst varying the lower 2 layers to suit differing user requirements. This is in fact the way in which some organisations have approached the problem of industrial network standardisation.



*Figure 8.10 - Bridging Similar OSI Networks*

## 8.5 The Ethernet System

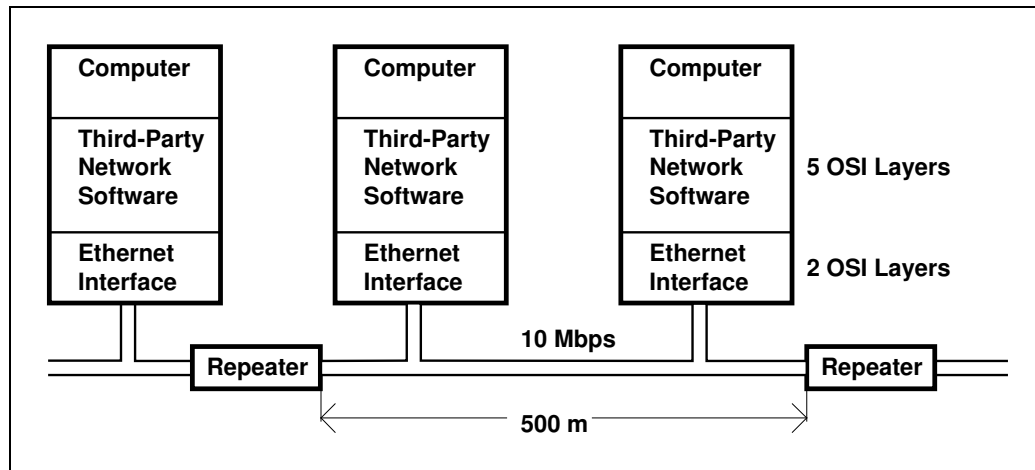
We now move on to examine one of the earliest, and most commonly used, specifications applied in Local Area Networking. The Ethernet system, which was developed by Xerox, in conjunction with Intel and the Digital Equipment Corporation, came into being as an experimental system in the early 1970s.

Ethernet is very frequently and very incorrectly referred to as a Local Area Network. The Ethernet system defines only the lowest, two layers of the OSI model and hence it does not represent a network in itself. In other words, one can't simply buy an "Ethernet Network" and expect to have applications support routines available. The Ethernet specification is however used as the foundation (backbone) for a range of commercial networks that provide the additional, upper five layers of OSI model functionality needed to support communications.

The Ethernet specification for the physical and data link layers of the OSI model is based upon a baseband, CSMA/CD bus. Referring to the IEEE terminology introduced in section 8.3, we can say that Ethernet effectively defines the physical layer and the data link layer of the OSI model, although emphasis is usually given to its Media Access Control. It should also be noted that the development of the IEEE 802.3 specification, for baseband communications on a bus network, was based upon the Ethernet specification and not vice-versa (as one may be tempted to assume). This was essentially done in the early 1980s as an IEEE response to the then growing acceptance of Ethernet as a defacto standard.


Ethernet was designed to provide a bus network of 2500 metres maximum length, established from cable segments of 500 metres maximum length. The cable segments themselves are joined together with "repeaters" which do not interfere with the CSMA/CD contention on the network. Ethernet allows for data transmission rates of up to 10 Mbits per second, with as many as 1024 network nodes. The system is shown schematically in Figure 8.11.

Although the Ethernet system was designed for baseband transmission over coaxial cable, the CSMA/CD contention scheme will function over any multi-access broadcasting medium. Radio, twisted-pair and optic fibre systems have all been successfully used in conjunction with the Ethernet CSMA/CD system.



*Figure 8.11 - Ethernet as the Basis for Networking*

The actual Ethernet data frame is shown in Figure 8.12. It is not unlike the HDLC frame in its form. It consists of a 7 byte preamble, followed by a single byte Starting Frame Delimiter (SFD), two or six byte Destination and Source Addresses (DA and SA), data length specification field, data and padding bits and finally a Frame Check Sequence (FCS). The length of the preamble is designed to allow for receiver synchronisation and consists of alternating 1s and 0s. The padding bits are added if there are insufficient bytes in the data (provided by the LLC) for the protocol to operate.



*Figure 8.12 - The Ethernet Data Frame*

The Ethernet system is relatively simple and it functions extremely well in the office environment. This is why it is in such widespread use. However, it has already been noted that the CSMA/CD system is "non-deterministic" in nature and as such may be undesirable within an industrial environment because message delays cannot be predicted with any certainty.

Another serious shortcoming of the 802.3 standard is the length limitation, brought about by the maximum round-trip propagation delay permissible for collision detection to operate. This restricts the length of the CSMA/CD system. Whilst this is generally an acceptable restriction in a typical office environment, it can present problems in large industrial sites.

In addition to the above factors, we now also appreciate that industrial environments will ultimately need the multi-channel, multi-function facilities of broadband communication systems in the coming years. These factors have led many people to believe that single-channel, baseband systems, such as Ethernet, are not suitable for long term industrial networking strategies. However, as we have seen with RS-232, suitability has never been a major criterion for the proliferation of data communications standards. It may well be that the existing user acceptance of Ethernet based networks may ultimately lead to its adoption and modification for industrial applications.

## 8.6 The General Motors / Boeing MAP & TOP Networks

The General Motors (GM) Corporation was one of the first end-users to launch a major international drive for standardisation of industrial data communications, based upon the OSI network model. The GM initiative is a prime example of how difficult and complex a task it is to achieve standardisation of communication in the factory environment.

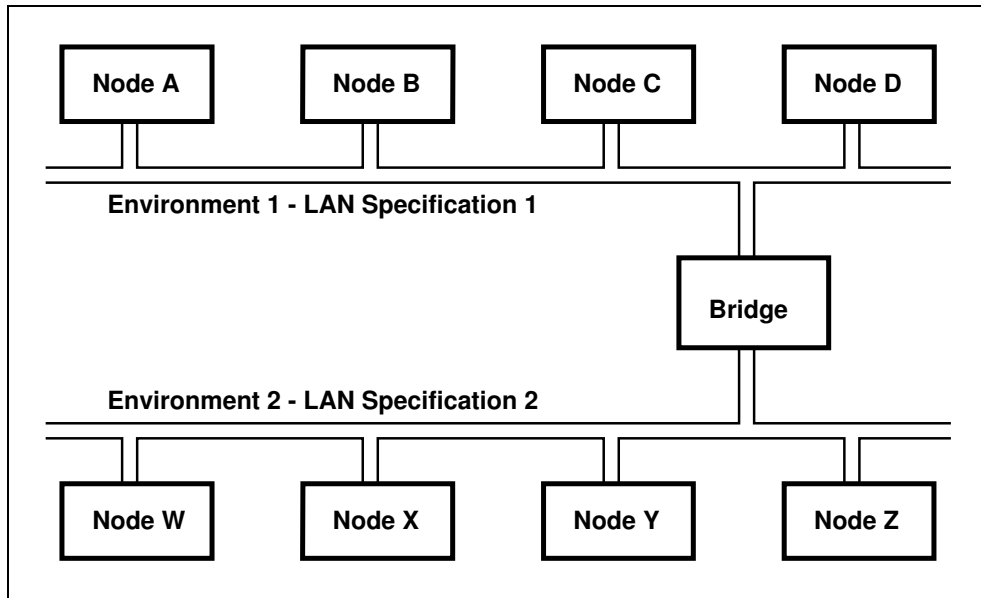
The GM MAP Task Force was established in 1980 with a charter to identify data communications standards that would provide for multi-vendor communications in manufacturing environments. The Task Force was comprised of representatives from 15, different GM divisions to ensure that all plant needs were discussed and that all appropriate control system and computer vendors were involved. It was decided by the Task Force, from the earliest stages of investigation, that all new networking strategies would be based upon the ISO/OSI 7 layer model framework.

It may seem rather ironic that an automobile manufacturer should be the first organisation to launch a major assault on networking standardisation, but it must be remembered that GM was one of the largest users of computer-based equipment in the world when it initiated the Task Force. It should also be noted that the problem of communications standardisation in manufacturing was so large and so chronic that it could not have been tackled without the driving force of such a large organisation (and end-user). Few other organisations in the world had the same level of industrial authority to influence trends in factory network development.

The term "MAP", in the networking context, is an acronym for "Manufacturing Automation Protocol". In short, it is a layer by layer specification for every layer in the OSI communications model. The specification for each layer has been chosen with due consideration for its suitability to the manufacturing environment.

We have already noted that no single networking strategy is optimal for all environments. So how did GM manage to find a universal, "plug-in compatible" networking strategy where others had failed? The simple answer is that they didn't and nor did they attempt to. The objective of MAP was to rationalise standards to the extent where a limited number of identifiable and compatible, standard networks could be established.

From a strategic point of view, MAP development relied upon defining different data communications environments - office, factory, national and international. Networking specifications for each environment could then be systematically identified so that it would ultimately be possible to bridge individual networks (and environments) with low level interfaces. This is shown in Figure 8.13.



*Figure 8.13 - Bridging Complementary Networks*

Although this form of network rationalisation sounds perfectly feasible, the objectives are quite massive in their scope and they could not be achieved through the operation of General Motors on its own. General Motors therefore chose to dedicate its efforts towards the definition of standards for the factory environment.

The work of the GM MAP Task Force was then greatly augmented through the cooperation of the Boeing Corporation, who appreciated the benefits that network standardisation would bring to their manufacturing. Boeing chose to assist GM through the development of a complementary networking specification for the office environment. The acronym chosen for the Boeing specification was TOP, which is an abbreviation for Technical Office Protocol. In terms of the simplistic diagram of Figure 8.13, we could say that LAN specification 1 represents the TOP protocol and LAN specification 2 represents the MAP protocol.

It must be clearly understood that the primary objective of the GM/Boeing MAP/TOP combined Task Force was to rationalise networking standards and not to introduce unnecessary new standards. Wherever possible, the MAP and TOP specifications were both based upon existing standards, established by ISO, IEEE and CCITT for Local and Wide Area Networking. Figure 8.14 illustrates the various standards which are defined in MAP V3.0 and TOP V3.0 for each layer of the OSI model. Note that wherever applicable, these specifications use the ISO nomenclature for standards, rather than their IEEE or CCITT forms.

OSI Layer	MAP V3.0 Specification	TOP V3.0 Specification	OSI Layer
7	ISO 8571 FTAM *ISO 9506 MMS ISO 8649/50 ACSE ISO DP9595/96 Management ISO DP9594 Directory Services	ISO 8571 FTAM *CCITT X.400 MHS ISO 8649/50 ACSE ISO DP9595/96 Management ISO DP9594 Directory Services	7
6	ISO 8822/23 Presentation ISO 8824/25 ASN.1	ISO 8822/23 Presentation ISO 8824/25 ASN.1	6
5	ISO 8326/27 Connection Oriented	ISO 8326/27 Connection Oriented	5
4	ISO 8072/73 Transport Class 4	ISO 8072/73 Transport Class 4	4
3	ISO 8348 CLNS ISO 8473 Network Layer Protocol ISO DIS9542 ES-IS Exchange	ISO 8348 CLNS ISO 8473 Network Layer Protocol ISO DIS9542 ES-IS Exchange	3
2	ISO 8802.2 LLC 1,3	ISO 8802.2 LLC 1,3	2
1	ISO 8802.4 Token Bus	ISO 8802.4 Token Bus *ISO 8802.3 CSMA/CD Bus *ISO 8802.5 Token Ring	1

\* Denotes differences between MAP and TOP specifications

*Figure 8.14 - MAP V3.0 and TOP V3.0 Specifications*

The similarity between the MAP and TOP specifications is self evident, with the only major differences occurring in the first and seventh layers. The MAP network, having its primary role on the factory floor is based upon the deterministic (and more complex) token passing bus arrangement (IEEE 802.4). The TOP network is generally based on the low cost ISO 8802.3 (IEEE 802.3) CSMA/CD bus, but can optionally be run on a token bus or token ring network.

The other major difference between MAP and TOP is in the services which are provided to end-user programs by the application layer. The major objective of modern factory communication is to enable computers, CNC machines, process controllers, PLCs and robots to communicate with one another. It is therefore necessary for the MAP system to provide services which will allow a diverse range of such devices to talk to one another in real time. This is a major task in itself and is over and above the normal file transfer and management services provided by a typical application layer for the office environment. The generic standard which describes this facility is referred to as the Manufacturing Message Specification or MMS.

MMS is a complex standard that defines entities known as Virtual Manufacturing Devices or VMDs. VMDs are used to describe devices such as PLCs, CNCs, etc. The MMS is then tailored, through a set of MMS companion standards, for interaction with VMDs. One companion standard could define (say) robotic devices, another CNC devices and so on. In other words, MMS is the kernel specification around which companion specifications are developed as needs arise. The MMS services to the end-user include the following:

- Virtual Manufacturing Device (VMD) support
- File access
- File management
- Operator communication
- Variable access
- Program invocation management
- Semaphore management
- MMS operating environment management
- Domain management
- Event management
- Journal management.

As with other standards we have already examined, MMS defines a number of different classes of service. This is done so that implementations which do not need all the above services are not burdened with unnecessary overheads. The actual number of services provided depends upon which conformance "class" is selected for network operation. The concept of the VMD is directed specifically towards factory-floor systems. MMS is a very comprehensive specification which is too complex and cumbersome for technical office applications.

The networking requirements of the technical office are not unlike those of any other office and include:

- File transfer and access
- Distributed database interfacing
- Electronic mail systems
- File, directory, print and plot servers
- Document, data and graphics interchange.

These services can be provided by the CCITT X.400 Message Handling Specification (MHS), which is an electronic mail facility that is used in place of MMS in the technical office.

The other important elements, which are common to the applications layers of both MAP and TOP, include:

- FTAM
- ACSE
- Directory Services
- Management.

*FTAM* is an acronym for File Transfer Access and Management. This service, as its name implies, is designed to facilitate a simple common access scheme for network files.

*ACSE* is an acronym for Association Control Service Elements. ACSE allows an application on one network node to set up an association with an application on another network node. This is achieved through the session layer on each node. In order for an association to be set up, it is necessary that every application has a mechanism for discovering the location of every other application on the network. This is done through the Directory Services to the application program.

*Management services* are provided to applications programs so that they can access or configure the current status of any node in the network. Since these services are not only dependent upon hardware, but also upon the required network operation and security, they must be provided by the network vendor.

As one can clearly see, even from a brief discussion of the specifications, MAP and TOP are very sophisticated networks. Having seen these standards, the obvious question to ask is, does manufacturing need this level of sophistication? The answer is probably not in the short-term. However, network standardisation is such a complex task that designers need to think of long-term requirements. As major specifications proliferate, they gain such a momentum that it is very difficult to change direction even after ten or twenty years (eg: RS-232). It is therefore essential that initial specifications attempt to address future issues as far as practical.

The IEEE 802.4 physical layer specification for MAP is another standard which addresses both present and future requirements. This token bus specification defines three, different modulation options comprising:

- Single-channel Phase-Continuous, Frequency Shift Keying (FSK) carrierband at 1 Mbps
- Single-channel Phase-Coherent, Frequency Shift Keying (FSK) carrierband at 5 or 10 Mbps
- Multi-channel Duo-binary, Amplitude Modulated (AM) Phase Shift Keying (PSK) broadband at 10 Mbps.

In terms of future expansion, the logical modulation option is broadband, which allows for multiple voice, video and data channels to exist on the same, semi-rigid coaxial backbone cable. However, the cost of providing broadband communications interfaces and cabling for low cost or unintelligent devices is unacceptably high. MAP therefore provides both the broadband and phase-coherent carrierband options.

The total, MAP/TOP OSI architecture is shown schematically in Figure 8.15. This network begins to resemble our ideal, plug-in compatible network. Bridges are used to interconnect the TOP network and the broadband and carrierband MAP networks. Note also the presence of the "Headend Remodulator" on the broadband MAP network. In order to achieve full-duplex communications within a single cable, broadband system, transmit and receive paths are assigned different frequency bands. Transmitted signals are sent toward the headend device which converts these frequencies into the receive band.

The OSI based MAP/TOP scheme shown in Figure 8.15 does not address all manufacturing communications needs. Whilst a 7 layer OSI network is suitable for CNCs, robots, computers, guided vehicles, etc., this form of interfacing is unacceptable for devices such as sensors. Not only is a 7 layer protocol far too expensive, but it may also be far too slow. For this reason, the MAP protocol also incorporates an Enhanced Performance Architecture (EPA) format. The EPA is a reduced platform (2½ layer scheme) which uses only the physical layer, the data link layer and MMS portion of the applications layer. It is not OSI compatible. The EPA scheme is shown in Figure 8.16, where some devices are equipped with only 2½ layer networking, whilst others provide both 7 layer and 2½ layer networking. Although the 2½ layer networking scheme is faster than a full 7 layer OSI arrangement, it is still unable to satisfy the high-speed demands of real-time continuous process control.

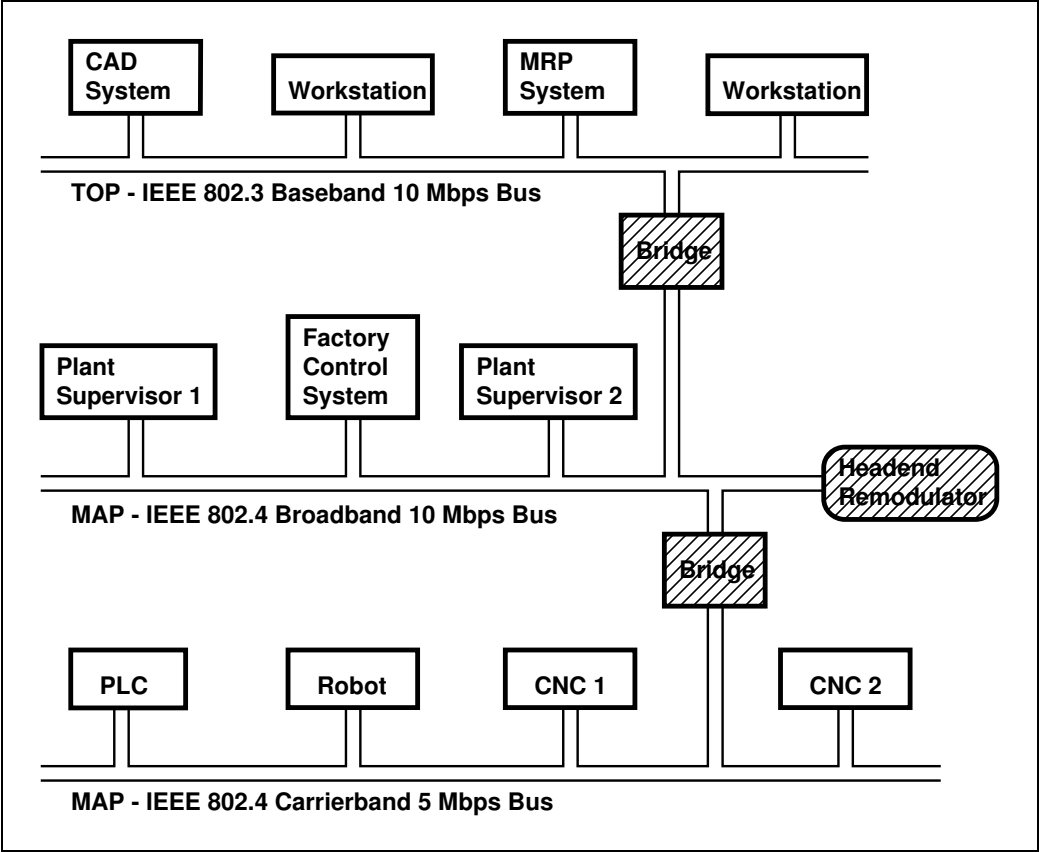


Figure 8.15 - MAP/TOP Architecture for OSI Networking

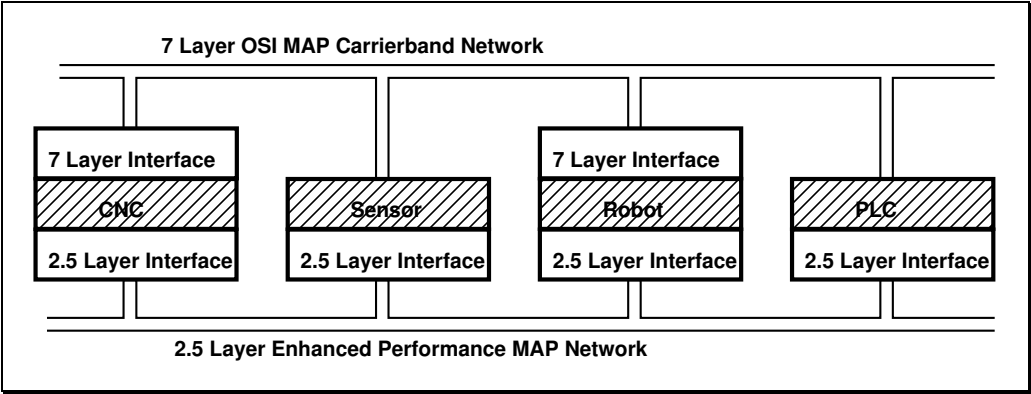


Figure 8.16 - Enhanced Performance Architecture for MAP

An examination of Figures 8.15 and 8.16 leads us to examine the issue of MAP/TOP and OSI network conformance. As one may have gathered from the above discussions, each specification for every layer may offer a number of different options (classes / modes) for services. Two network nodes, which each conform to the same specification for each layer of the OSI model, may not necessarily be able to communicate with one another directly.

A good example of conformance problems is illustrated by the IEEE 802.4 specification, which allows for different modulation techniques, different communications media and so on. Clearly then, the criteria for conformance to an OSI protocol is NOT just the use of equivalent specifications for each layer of the OSI model. In order for two devices to effectively communicate, they must be set up to use matching (or complementary) classes of service for each specification of each layer of the OSI model. A specification for a 7-layer OSI based network that explicitly defines all the relevant options for each standard in every layer of the model is sometimes referred to as a "Profile".

It is evident from the complexity of each specification chosen for each layer of the MAP/TOP system, that the development of interfacing hardware and software needs to be carried out in consultation with specialised MAP/TOP conformance testing centres. The traditional standards bodies alone are not structured to provide such support and testing. An important preliminary role for the MAP/TOP Task Force, Steering Committee and User Groups was therefore to establish conformance testing centres for these complex protocols.

The path to plug-in network compatibility in the manufacturing environment has been a difficult one. Even with the resources of General Motors and Boeing and worldwide support from user groups, standards bodies, computer companies and equipment suppliers there have been many problems. The complexity of an OSI-based industrial networking standard is such that it discourages interface manufacturers from investing in specialised VLSI circuits until specifications become stable. As a result of major specification changes that occurred between MAP/TOP version 2.1 and version 3.0, a great deal of uncertainty was created amongst Original Equipment Manufacturers. This caused the implementation of the MAP/TOP scheme to falter until it was decided by the MAP/TOP committees to effectively freeze the specification at version 3.0 for a number of years. The freeze was seen as a sensible action that would encourage third-party, interface manufacturers re-enter the MAP/TOP product market and thereby assist in the proliferation of the specification.

It is also important to note that the MAP/TOP specification was not immediately embraced as a standard in any country. The General Motors and Boeing MAP/TOP Task Force and their associated user groups expended a great deal of effort in order to have the total specification recognised as a networking standard. However, despite the fact that all the constituents of MAP and TOP were already considered as standards, the total specification was the subject of lengthy reviews by standards related bodies in North America, Europe and Japan. Ten years after the original inception of the MAP concept, there was still a lack of consensus in regard to its adoption as a universal networking standard for manufacturing.

The irony of computer based developments is that sometimes the standards most carefully designed to fulfil an existing need are usurped by less-suitable, defacto standards that proliferate through an explosion of low cost devices in another area. RS-232 is a good example of this phenomenon. It therefore remains to be seen whether the time delay in bringing a MAP/TOP type strategy to fruition is long enough to allow other standards, such as those emerging in office networks to the fore. In either event, the introduction of 7 layer OSI networks into the factory environment is ultimately necessary to minimise the ever increasing cost of equipment integration.

Whilst it may be that in the long term, a MAP/TOP strategy will become the panacea for industrial networking, it must be realised that short term implementation of such an OSI networking scheme is difficult for many manufacturing organisations. The reasons for this are that:

- It is not feasible to provide OSI network interfaces for the majority of programmable controllers built prior to the mid-1980s. Many of these controllers have "closed-architectures", do not have standard internal bus structures and cannot be readily be adapted for communications
- The cost of developing one-off OSI network interfaces for uncommon and exotic industrial controllers is prohibitive
- It is often not practical to retrofit "old technology" systems with new controllers because this requires an unacceptable down-time for production equipment and complete re-commissioning
- The replacement cost of manufacturing equipment is very high and so too is the replacement cost of the industrial control equipment. It is therefore highly unlikely that manufacturers will discard existing equipment simply to achieve network integration.

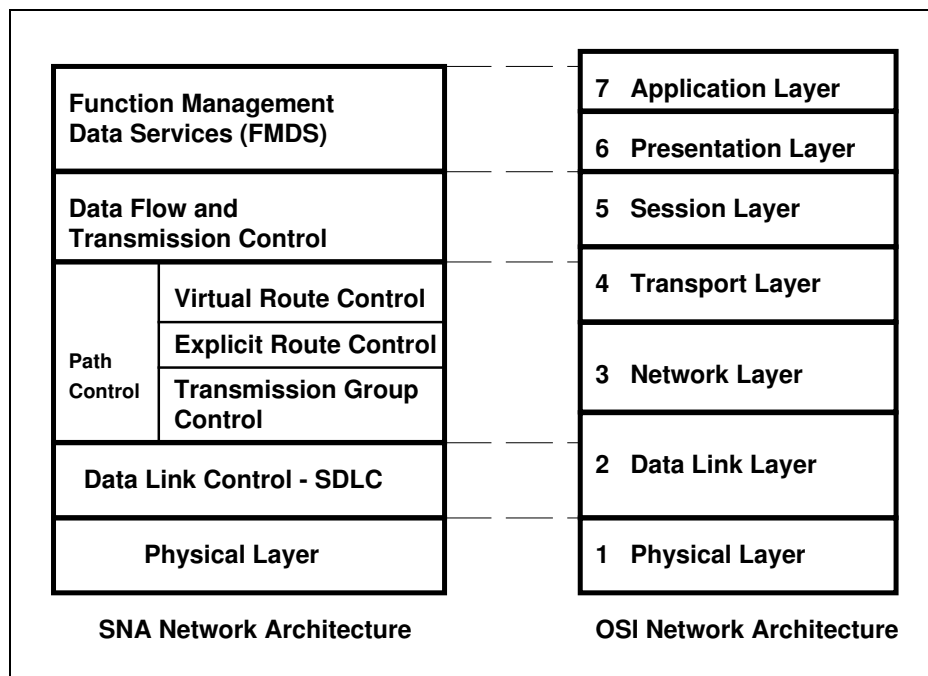
In simple terms, these factors mean that OSI based networks can only be phased in as existing industrial equipment and controllers are phased out and replaced with more modern devices. In the interim period, manufacturers are still faced with the prospect of developing onerous communications protocols, based upon RS-232 type hardware links.

On the other hand, the very nature of technical office computer systems makes OSI networks such as TOP feasible even in the short term. Unlike the shop-floor environment, the technical office only uses a relatively small range of computer architectures of similar vintage and with a limited number of internal bus structures. These devices are generally open-architecture and can readily be networked within an OSI environment. The major impediment to TOP implementation is that there are already many competing office networking strategies that have already been applied.

In the final analysis, regardless of which specifications are chosen for networking in the manufacturing environment, it is to be hoped that the concept of "plug-in compatibility" will be realised.

## 8.7 SNA

The IBM Systems Network Architecture, commonly known as SNA, is very briefly introduced herein because of its widespread use in conjunction with the company's own equipment. SNA is IBM's proprietary protocol architecture. Its structure precedes the OSI 7-layer framework, but it does contain a number of similarities. The OSI model and SNA model are shown side by side in Figure 8.17.



*Figure 8.17 - IBM Systems Network Architecture vs OSI*

The SNA protocol is based upon synchronous serial transmission, with framing and error control carried out by the IBM implementation of HDLC, known as SDLC (Synchronous Data Link Control).

The basic entities within SNA are called Logical Units (LUs) and these can represent either end users or applications programs. Communications within the SNA system is between these Logical Units. It is the Logical Units which are assigned addresses and not the end users as such. Logical Units are therefore referred to as Network Addressable Units or NAUs.

Although the levels of functionality differ between OSI and SNA, the end results are not dissimilar and consist of a large range of services that are available to application programs and end users. SNA not only preceded the OSI framework, but also the push for manufacturing networking, capable of handling industrial control systems. The FMDS layer therefore does not provide services for communications with the majority of manufacturing devices.

Although SNA is IBM's proprietary networking system, there are a range of networking gateways that will enable other computer manufacturers to interface their systems to SNA.

The SNA strategy was not unlike the GM MAP strategy except that it was vendor-driven rather than end-user-driven. The fact that SNA was not universally adopted is another illustration of how difficult it is, even for very large organisations, to develop networking standards that gain widespread acceptance.

## 8.8 File Server and Office Networks

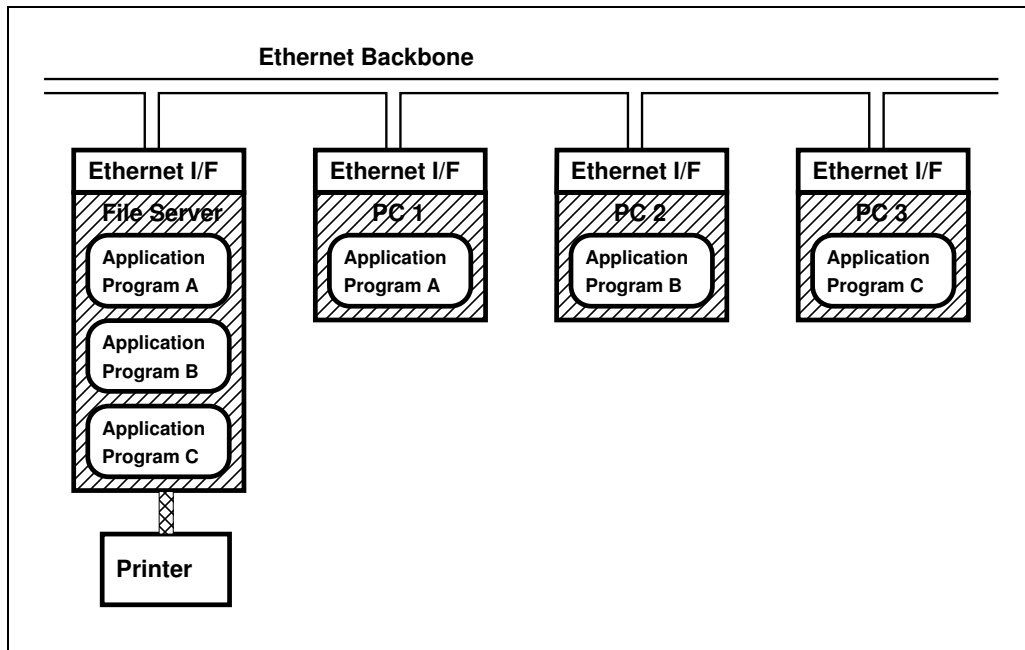
The proliferation of Personal Computers (PCs) and Workstations throughout the office environment has changed the way in which data is processed and stored. Where mainframe computers once performed both processing and storage, PCs and workstations now carry out a significant proportion of both activities.

In many instances, mainframe computers have disappeared altogether to make way for PCs, which more adeptly perform word-processing, desktop-publishing, accounting, spreadsheets, etc. A major problem with PCs and workstations however is that, unless they are networked, they tend to discourage centralised data storage and "resource" sharing. The resources that tend to be most inefficiently utilised with individual PCs are hard-disk capacity and printers. For example, ten stand-alone PCs running the same word-processing package use up ten times the disk space that would be used if the application program could be kept in a centralised store. With stand-alone PCs, individual users tend to have their own printers, and these are idle for a large proportion of the total time spent at a computer terminal.

A number of sophisticated Local Area Networks have been introduced into the office environment to facilitate resource sharing, centralised file storage and distribution between multiple PC workstations. These turn-key systems are called file server networks and are based on the principle that one workstation is used to store and distribute files for other workstations, thus minimising hard-disk and printer requirements. File server systems are commonly based on an Ethernet or Token Ring Network backbone and provide all the software that is necessary to share application programs and data files between nodes with security. A typical PC file server arrangement is shown schematically in Figure 8.18.

In the network of Figure 8.18, the file server PC holds all the applications programs on its own bulk-storage unit (hard disk drive/s). If PC 1 wishes to run a particular application program B (say a word-processor), then it makes a request to the file server to "capture" that program file. The file server down-loads the program, through the network, into the memory of PC 1. The application program executes exactly as if it had been retrieved from a local hard disk unit.

Let us say that the application program on PC 1 is a word-processor. Once the end-user has completed the development of a document, it is transferred back to the file server. If the user-defined security arrangements permit, then other nodes can then call up the document to examine or modify it. The file can also be sent through the network to the file server for printing on the file server's printer.



*Figure 8.18 - A File Server Network*

The file server node is not unlike any of the other workstations on the system, except that it is generally equipped with high capacity hard disk drives so that it can store a wide range of application programs and user data. However, in order for the file server system to operate, each workstation needs to be loaded with a special piece of software that will capture and send files through the backbone network. The software on the file server itself needs to be more comprehensive and take into account file security and access, etc. All the complexities of the network are generally transparent to the end user who only needs to issue simple commands (not unlike those found on any PC Operating System) in order to capture and store files through the network.

The file server itself can either be "dedicated" or "non-dedicated". A dedicated file server is used only for handling network requests for files. A non-dedicated server performs file serving in background mode whilst a user can run an application program in the foreground. However, the trade-off is that the non-dedicated system sacrifices the performance of the entire network for extra functionality in the server. This is clearly undesirable and unnecessary given the low cost of personal computers. It is also possible to use more than one node as a file server for the network. This is referred to as a distributed file server network.

A simpler version of the file server is called a "disk server", which simply provides a centralised disk storage system for all the network nodes. In such a system, each node treats the disk server as though it were another local hard disk drive. In the DOS environment for example, each PC may know the disk server as drive "F:". In order for a node to run a program called "FRED.EXE", which exists on the disk server, the node user would simply key in:

**F:FRED**

The use of the drive prefix "F:" would be a command to the network software that indicated that a file transfer from the disk server had to take place.

Since one of the objectives of the file server system is to minimise the amount of hard-disk space required to run a range of office applications, it is now possible to purchase network PCs which have no internal hard-disk. These PCs can be equipped with special Read Only Memory (ROM) chips, that contain a program which loads the Operating System from the file-server (via the network). This not only saves hard-disk space but (in conjunction with file server security systems) also eliminates the spread of computer viruses in public computer areas such as classrooms. The problem with this option is that it severely restricts the flexibility of any individual PC, since it can no longer run as a stand-alone device.

There is much written about the tremendous savings that file server systems and office LANs can provide through the use of shared printers. While in theory it may sound attractive for any user to be able to direct files through a network to any printer, in practice the scheme has a number of problems. The idealistic arrangements take little cognisance of the fact that in practice:

- printers do not have an unlimited supply of paper
- offices normally use a range of different papers (letterheads, memo sheets, envelopes, etc.)
- printers do jam paper
- people invariably realise part way through the printing of a document that their pagination is incorrect or that they have selected inappropriate fonts and then wish to abandon print-outs.

Although commercially available networks can handle such occurrences, they do illustrate the fact that ultimately there needs to be some human intervention in systems that endeavour to network intelligent mechanical devices (printers) to remote host computers.

The printer problems in the office are completely analogous to those faced by engineers on the factory floor, when networking intelligent production equipment (such as CNCs or robots) to host-design-computers. Even if the network is error-free and the source files are error free, there is no guarantee that the target machines have been fitted with the correct tooling. With all the sophistication that networks can offer, sometimes there is simply no substitute for a person standing directly next to a mechanical device and switching off the power when things go wrong.

File server networks have traditionally suffered from the same problem as all other networks - they required experienced personnel for installation and maintenance. Without such personnel, file server networks could generate far more problems than they resolved. However, if correctly implemented and coordinated, a file-server system and Local Area Network can provide a sound mechanism for information flow and resource sharing in the office environment.

The emergence of new, Windows-based operating systems for personal computers and workstations has also changed the way in which we look upon file server networks. Newer versions of operating systems (such as Microsoft Windows for Work-groups and Windows NT) are being designed with built-in networking functionality and the ability to not only share files but processing power amongst personal computers and workstations. The maintenance problems associated with networks will gradually decline as the network becomes more and more integrated into the computer operating system environment.

What we are now witnessing therefore, is the fact that the network is becoming an integral part of the most prolific operating systems. This has enormous ramifications that will almost certainly spill over into the manufacturing environment. As with most computer based developments, standards are inevitably driven by volume far more effectively than they are by Government standards bodies and committees. So, while endless committees have been arguing over the future of manufacturing standards, the high volume personal computer market may have already made the decision for the manufacturing world.

## **8.9 What Will Our Computers Say Once they are Networked?**

All through our discussions on networking within manufacturing we have discussed the kinds of data that we should be transferring, the kinds of data that we would like to transfer and the kinds of data that we do transfer through communications networks.

In a larger sense however, it is not the communications network that will initiate the flow of information around the manufacturing environment. This kind of systems integration can only occur when applications programs begin to make use of the services which are provided by data communications networks. In order for this to happen, much work will need to be done in standardising the formats in which data is stored and the way in which applications programs access that data.

In the final analysis, we should not look upon the emergence of industrial networking standards as the end to all our integration problems, but rather as the beginning of a completely new set of problems. These much larger problems of software, database and data format standardisation will ultimately need to be overcome in order to integrate management, control and monitoring systems in the factory environment.

