

5

NETWORK ARCHITECTURE AND STANDARDS

INTRODUCTION

Part of the power of local area networks is their ability to support a wide variety of devices. Supporting a wide variety of devices, however, can present substantial compatibility problems. For varying devices to be linked together, the hardware and software of these devices need to be compatible, or else complex interfaces have to be built for meaningful communication to take place. To facilitate this complexity, network architectures are being developed that allow complex networks to be built using a variety of equipment.

WHAT IS NETWORK ARCHITECTURE?

The term most frequently used for the overall design of a network is architecture. The goal of architecture is to achieve a high level of performance at a minimal cost. Although local area networks cannot be classified solely by architecture, the specific combination of elements determine the characteristics of a particular LAN. Table 5.1 depicts the high level objectives of network architecture.

TABLE 5.1 : High-level objectives of network architectures

- * Connectivity. Permit diverse hardware and software products to be connected to form a unified networking system.
- * Modularity. Permit the use of a relatively small set of massproduced general-purpose building blocks in a wide diversity of network devices.
- * Ease of Implementation. Provide a general solution to network communication that can be easily installed in a variety of configurations to meet the needs of all types of users.

- * Ease of Use. Provide communication facilities to network users in a way that frees them from concerns about or knowledge of network structure or implementation.
- * Reliability. Provide appropriate error detection and correction facilities.
- * Ease of Modification. Permit the network to evolve and be easily modified as user needs change or new technologies become available.

Specific protocols tend to be associated with specific topologies to the point that topology is often considered the means of implementing the communication protocol. The control strategy used by a network determines the number and type of connected workstations, the nature of interactions between workstations and the particular transmission link used.

No one architecture is innately superior to another: too many variables affect each individual situation. The most that can be said is that given a specific set of circumstances one network design seems to accommodate more of the user's requirements than another. Ultimately, the architecture of a network is determined by the location of the workstations to be connected, the demands of the information being transmitted, the available resources and the nature of the organisation creating the network.

Network architecture is a fancy term for the way that networking products are constructed. Networking hardware and software is implemented on systems via a mechanism called network or communication architecture. Communication architecture is the "layering" of software based upon the functionality of each layer. It is very similar to an organisation chart at a corporation. At the lowest level of the protocol layers lies the data link access, the software used to talk directly to the hardware. This is usually a cryptic interface and is difficult to implement and maintain. At the next level would be some sort of communication line handler whose job would be to keep messages sorted out and manage connection creation and destruction between machines. The next layer up would be a session control mechanism responsible for the overall message flow control and ensuring that the communication "session" between systems goes smoothly. The remainder of the upward layers is dedicated to direct user programme interaction for specific functions. For example, one layer would be used for communication with programmes desiring remote file access and manipulation, another with programme-to-programme communication, etc. Very few communication architectures do not use layered architectures, and these are somewhat antiquated. The benefits of the layered approach are many, but the most significant one is the ability to change a layer's capabilities without significantly modifying the entire architecture. This feature alone makes a layered network architecture very attractive for companies desiring inter-and intra-systems communication capability.

STANDARDS

To form a viable network, all the diverse elements - equipment, topology, communication links, protocols - must be assembled into a unified system. While the actual design of each component of the communication system can and does vary

depending on the specific system involved, no single component can be selected or designed in isolation. Parts of the system must be balanced for effective communication to occur. If any one component does not interact successfully with the others or is absent, communication cannot occur.

Thousands of possible viable LAN combinations exist. Given that creative engineers love to experiment (inherent in all new technologies is an almost overwhelming desire to play, to see how many different things can be done), how do you coordinate all the pieces?

More importantly, how do you insure that a local area network is able to interconnect diverse types of equipment from multiple sources when there is no common language that computers can use to communicate with each other?

Standards that bring order to the design process are required. Currently, the work of two organisations is relevant to local area networks. They are the International

MHS CCITT X.400	MOTS ISO 10021	DIRECTORY CCITT X.500 ISO 9594	FTAM ISO 8571	CMISE ISO 9595 (SERVICE) ISO 9596 (PROTOCOL)	DTP ISO 10026	JTM ISO 8831 ISO 8832 (PROTOCOL)	VT ISO 9040 (SERVICE) (PROTOCOL)	APPLICATION LAYER
RTSE CCITT X.226 ISO 9066		ROSE CCITT X.229 ISO 9072		CCR ISO 9804 (SERVICE) ISO 9805 (PROTOCOL)				
ACSE CCITT X.227 ISO 8649 (SERVICE) ISO 8850 (PROTOCOL)								
CCITT X.226 ISO 8822 (SERVICE) ISO 8823 (PROTOCOL-CONNECTION MODE) ISO 9576 (PROTOCOL-CONNECTIONLESS MODE)				ASN.1 CCITT X.208 (LANGUAGE) CCITT X.209 (BASIC ENCODING RULES) ISO 8824 (LANGUAGE) ISO 8825 (BASIC ENCODING RULES)				PRESENTATION LAYER
CCITT X.225 ISO 8327 (SERVICE) ISO 8327 (PROTOCOL-CONNECTION MODE) ISO 9548 (PROTOCOL-CONNECTIONLESS MODE)								SESSION LAYER
CCITT X.224 ISO 8072 (SERVICE) ISO 8073 (PROTOCOL-CONNECTION MODE) ISO 8802 (PROTOCOL-CONNECTIONLESS MODE)								TRANSPORT LAYER
INTERNETWORK PROTOCOL ISO 8473 (CONNECTIONLESS MODE)			PACKET LEVEL PROTOCOL CCITT X.25 (CONNECTION MODE) ISO 8208			CCITT Q.931		NETWORK LAYER
LOGICAL LINK CONTROL ISO 8802.2					CCITT X.25 LAPB ISO 7776		CCITT LAPD Q.921	DATA LINK LAYER
CSMA/CD ISO 8802.3	TOKEN BUS ISO 8802.4	TOKEN RING ISO 8802.5	SLOTTED ISO 8802.7	FDDI 9314	CCITT X.21/V.24		CCITT 1.430/1.431	PHYSICAL LAYER
LOCAL AREA NETWORKING					WIDE AREA NETWORKING			

Fig. 5.1 : Some of the standards contained within the OSI seven layers

Standards Organisation (ISO) and the Institute of Electrical and Electronic Engineers' (IEEE) Project 802 committee. Figure 5.1 gives some of the standards contained within the OSI seven layers. Figure 5.2 gives IEEE LAN standards.

802.1 Internetworking			
802.2 Logical link control			
802.3 Medium access	802.3 Medium access	802.5 Medium access	802.6 Medium access
802.3 Physical	802.4 Physical	802.5 Physical	802.6 Physical

Fig. 5.2 : IEEE LAN standards

The Importance of Standards

Local area network standards are following the same path as their wide area network counterparts. Increased activity in the area of standards has resulted from a recognition that computer communication is a field so complex that without rationalisation progress will be severely impeded. There are many computer equipment and system suppliers offering a plethora of communication protocols and conventions, and agreement among them on standardisation would clearly be impossible without the existence of central standards bodies, both national and international.

Users typically have a variety of computing equipment including personal computers, word processors, minicomputers and mainframe data processing systems. Only in rare cases is it possible to interconnect these pieces of equipment without using protocol conversion devices, and even then compatibility is not fully guaranteed. Consumer networks inevitably face difficulties unless there is prior agreement on limiting the choice of equipment to be used. The corporate network can often rely on a strategy to reduce the number of options which may occur and hence the cost of circumventing inconsistencies.

In this context local-area networks are particularly at risk, given their intended role in interconnecting a variety of devices in many cases. Justification for purchasing a LAN or developing a new application using a LAN may depend on smooth integration with existing systems and on the ease with which new systems can be added. Particularly within a small geographical area, information can be transported physically and unless a cheap and efficient replacement for such services can be demonstrated, organisations are unlikely to change their practices. A particular example of this difficulty would be the case of two computers located in the same premises where information is exchanged using tape transfer by messenger or operator. In this case only a link offering high speed and integrity would be contemplated as an alternative. A second example concerns the use of stand-alone word processors. Typically, documents would be transferred between these machines by means of diskette exchange, so here again an electronic link would have to offer the advantages of improved speed and reliability before a capital outlay would be authorised.

At a time when it has been accepted that voice communication is a commodity and, in fact, an essential element of business life, on which large capital sums are frequently expended, it is unfortunate that the same cannot be said of electronic data and text communication. Even the physical connection of a simple microcomputer to the public telephone network can be a task requiring a great deal of effort. Even once the correct interface details have been ascertained, it is not unusual to find that the software used is incompatible with that used in the remote system.

In view of such difficulties there has been an enormous increase in the interest shown by both users and manufacturers in the importance of standards in the last few years.

The Standards Bodies

The organisations which contribute to the standards-making process are usually referred to as the standards bodies. This is an area where until recently even industry observers were confused by both the numbers of 'standards' organisations, and the mechanisms used to formulate and accept standards. Table 5.2 lists a number of acronyms, of the most important of these organisations, which may be familiar to readers of computer and trade journals.

TABLE 5.2 : Main standards bodies

ISO	International Standards Organisation
<i>National standards bodies (member of ISO)</i>	
ANSI	American National Standards Institute
BSI	British Standards Institution
DIN	Deutsches Institut für Normung eV (West German Standards Institute)
AFNOR	Association Française de Normalisation (French Standards Institute)
UNI	Ente Nazionale Italiano di Unificazione (Italian Standards Institute)
JISC	Japanese Industrial Standards Committee
SCC	Standards Council of Canada
<i>Telecommunications standards</i>	
ITU	International Telecommunication Union
CCITT	International Telegraph and Telephone Consultative Committee (Comité Consultatif International Télégraphique et Téléphonique)
<i>Other standards-making bodies</i>	
ECMA	European Computer Manufacturers Association
IEEE	Institute of Electrical and Electronics Engineers (USA)
NBS	National Bureau of Standards (USA)
IEC	International Electrotechnical Commission
IFIP	International Federation for Information Processing

THE OPEN SYSTEMS INTERCONNECTION MODEL

There was a time when there was no such thing as a network architecture. Companies implemented rather rude, crude and socially unacceptable software and hardware communication solutions without any thought as to layering or to the implementation of an architecture. The idea of layering really took off with the introduction of an international standard called the Open Systems Interconnection (OSI) model by the International Standards Organisation (ISO) in 1982 (International Standard 7498).

In 1978, the ISO Technical Committee 97 (this committee handles standardisation of information technology) started subcommittee number 16 (TC97/SC16) to develop an architecture and reference model that would serve as the foundation for future standards activities. From 1978 onwards, they have worked very hard at providing a flexible, reasonable communication architecture that could be implemented on a variety of systems and provide inter- and intra- systems communication capabilities in a variety of environments. Oddly enough, TC97/SC16 has not done most of the work on defining the protocols for each layer of the architecture; other ISO committees have done this, using the model specified by TC97/SC16. All protocols for all layers have not been defined yet, but the model still is highly useful in the definition of how a communication architecture is defined.

In 1978 the International Standards Organisation (ISO) proposed a broad model for mainframe wide area network communication, which they titled "The Reference Model of Open Systems Interconnection".

Open Systems Interconnection refers to the exchange of information among terminal devices, computers, people, networks and processes. The systems are open to one another by virtue of their mutual use of the standards developed from the original reference model. Openness does not imply any particular implementation, technology or interconnection means.

The Open System Interconnection Reference Model is not, by itself, a standard, nor is it a literal description of computer communication. While it defines where to perform tasks, it does not detail how to perform them. Individual services and protocols are not specified. Within the model, communication functions are addressed from the perspective of computer-to-communication network interconnection.

The Reference Model is intended to provide a common basis for coordinating the development of standards aimed at systems interconnection, while allowing existing standards to be placed in perspective within a common framework.

OSI LAYERS

As shown in the Table 5.3, each system is viewed as being composed of an ordered set of subsystems or layers. The layers of the Reference Model are separated by interfaces. Adjacent layers communicate through their common interface. Table 5.4 lists some of the well-known layers.

TABLE 5.3 : The OSI Layers

Layer	Definition
1. Physical	Concerned with transmission of unstructured bit stream over physical link : involves such parameters as signal voltage swing and bit duration; deals with the mechanical, electrical, and procedural characteristics to establish, maintain, and deactivate the physical link (RS-232-C, RS-449, X.21)
2. Data link	Provides for the reliable transfer of data across the physical link; sends blocks of data (frames) with the necessary synchronization, error control, and flow control (HDLC, SDLC, BiSync)
3. Network	Provides upper layers with independence from the data transmission and switching technologies used to connect systems; responsible for establishing, maintaining, and terminating connections (X.25, layer 3)
4. Transport	Provides reliable, transparent transfer of data between and points; provides end-to-end error recovery and flow control
5. Session	Provides the control structure for communication between applications; establishes, manages, and terminates connections (sessions) between cooperating applications
6. Presentation	Performs generally useful transformations on data to provide a standardized application interface and to provide common communications services; examples: encryption, text compression, reformatting
7. Application	Provides services to the users of the OSI environment; examples: transaction server, file transfer protocol, network management

TABLE 5.4 : Some Well-known Layers

OSI	CCITT	ISO	DOD	IEEE 802	ANS X3T9.5
7. Application		Various	Various		
6. Presentation					
5. Session		Session			
4. Transport		Transport (TP)	TCP		
3. Network	X.25	Internet	IP		

Sublayer		Logical link control	Data link
2. Link	LAP-B	Medium access control	Physical
1. Physical	X.21	Physical	

Each layer in the structure provides a defined set of services for the layer above and requests specific services from the layer below. Layers are defined by function: protocols are defined to control the processes managed by each layer. Relationships between layers and the information that must be passed between layers are identified.

Interfaces are located where one layer interacts with another and serve to isolate one layer from the next. Because the mechanisms and functions of the layers can be expected to change as technology develops, the functions of the interfaces are precisely defined, but the format of the data interchanges between layers is not. This permits the characteristics of a layer to change without affecting the rest of the model.

The arrangement of the layers begins with the most concrete - the physical layer, which defines cables - to the abstract - the application layer which defines user interfaces.

Layers one through four are considered system-independent. Protocols associated with these layers are relevant to all systems. Layers five, six and seven are considered system-dependent. Protocols must be defined individually for each different device. Some experts consider the first four layers to specify the communication functions: the remaining, higher, layers specify data processing functions.

To understand the OSI a little better, let us examine what each layer does. Figure 5.3 gives a diagram of the OSI model. You will notice that the model has basically

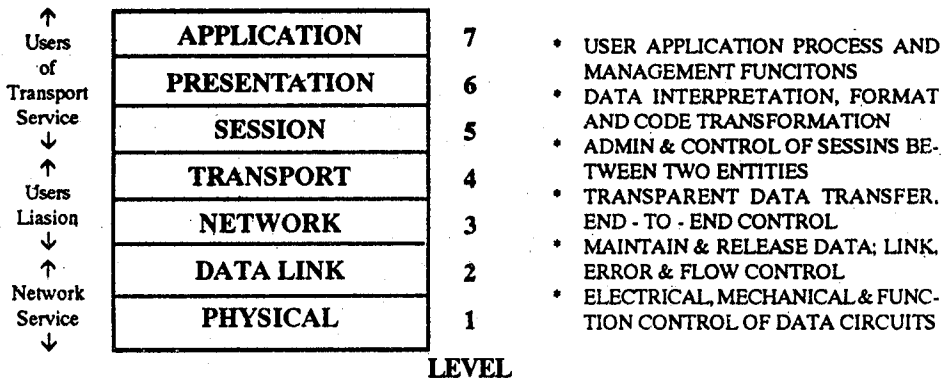


Fig. 5.3 : THE INTERNATIONAL STANDARDS ORGANISATION (ISO) OPEN SYSTEM INTERCONNECT (OSI) MODEL A LAYERED APPROACH TO NETWORK ARCHITECTURE

seven tiers, stacked one upon the other, that reflects a certain function at each layer. User data comes in to the top layer (layer seven) and travels through the various layers of protocols until it finally goes out over the transmission medium (hardware). It then travels to the destination node and begins its travel up the layers of protocols on the remote system until it reaches the destination programme on the remote system. This same ordeal happens on all communicating systems for the duration of communication between nodes.

The following paragraphs define the functionality of each layer of the OSI model:

Layer 1: Physical

This is the touch-and-feel layer. The physical layer provides for the transparent transmission of bit streams from one physical entity to another (or many, as in the case of datagram oriented services such as Ethernet).

The Physical Layer is concerned with transmitting raw data over a channel. It defines the mechanical, electrical, functional and procedural characteristics necessary to establish, maintain and disconnect the physical connection. The layer is designed to accommodate a variety of physical media and differing control procedures.

Aspects of the network affected by physical layer protocols include cables and connectors; electrical signalling methods; and computer and data communication equipment interfaces. Design problems focus on making sure that when one device sends a one, it is received at the destination as a one, not a zero. Typical issues are how many microseconds a bit occupies; how many volts are used to represent a one and how many for a zero; how many pins the network connector has and what each pin is used for; and whether or not transmission may proceed simultaneously in both directions.

Layer 2: Data Link

The Data Link layer handles the transfer of data between the ends of a physical link.

The data link layer aims at taking a raw transmission facility and transforming it into a line that appears free of transmission errors to the network layer. The layer acts to maintain an error-free channel by detecting and possibly correcting errors on the transmission medium.

The data link layer is responsible for providing the functional and procedural means used by higher layers to establish, maintain and terminate communication between users. Protocols associated with the data link layer affect the format of data blocks and address code; detection and recovery from errors; and sequence of transmitted data.

Layer 3: Network

The Network layer handles the routing and switching of information to establish a connection for the transparent delivery of data.

The Network Layer is concerned with the transmission of data through the network. Within the network layer, data from the transport layer is accepted, converted to packet, and routed towards the destination. In some sources, this is called the communication subnet layer because it is sub-system/segment of the network as a whole.

The network layer provides the means to establish, maintain and terminate network connections between systems, and it controls routing of packets. Associated protocols cover administration and management of data; delivery of status messages; regulation of traffic flow; and the division of labour between the Network Interface Units and the host, in particular who should ensure that all packets are correctly received at their destinations.

Layer 4: Transport

The Transport layer provides for error-free delivery of data and also acts as the control area for quality of service requirements for the selected data service.

The Transport Layer provides for the transparent transfer of data between users from source system to destination systems; that is, it manages the end-to-end data flow. Its basic function is to accept data from the session layer, split it up into messages if necessary, and pass these to the network layer. It must ensure that all the pieces arrive correctly at the other end.

Layer 4 represents the vaguely defined boundary between data communication and data processing. Transport layer protocols control distribution of messages, data integrity, loss or duplicate message prevention and proper addressing of user equipment.

Layer 5: Session

Session Layer provides the coordination of communicating processes, between nodes ("virtual" connectivity).

The Session Layer establishes a connection between users: in effect, it provides the user's interface into the network. The user must negotiate with this layer to establish a connection with a process on another machine. Once that connection has been made, the session layer synchronises the dialogue and manages the data exchange; that is, it manages the session, as the dialogue between users generally is called. Technically, a session is a dialogue between two presentation layer processes.

The session layer takes the communication service of the transport layer and adds application-oriented functions. Layer 5 protocols include rules for establishing and ending connections; verifying that proper communication is occurring; and interfacing the network to the operating system.

Layer 6: Presentation

The Presentation layer provides for any format, translation or code conversion necessary to put the data into an intelligible format.

The Presentation Layer transforms information from machine format into that

understandable by user. It performs functions that are requested sufficiently often to warrant finding a general solution for them, rather than letting each user solve the problem.

One major function is translation - between differing file formats, between different terminal formats and between different codes (ASCII to EBCDIC, for instance). The presentation layer is designed to represent information to communicating application-processes in a way that preserves the meaning while resolving syntax differences. Typical presentation layer services include data translation, conversion of file formats and encryption.

Layer 7: Application

The Application layer allows the end application to communicate with the communication architecture by providing the appropriate communication service(s) to the application.

The Application layer provides the user with a window into the system. All exchange of information between the user and the network occurs through this layer. Application services are the only Reference Model services directly comprehensible to users.

Layer 7 protocols are responsible for user and application programme support such as passwords, resource sharing, file transfer and network management.

At each layer, there may be one or more protocols (in the case of layer 2 and above) or communication media (in the case of layer 1) that communicate with a peer protocol or media on the complementary node(s). What this means is that, at any level, there can be more than one way to get data to and from the node; the only requirement is that there be the same peer at the destination node that understands what is sent.

At first this may all seem a bit chaotic and it is - to an extent. It is the job of the communication architect and software/hardware engineers to put the right functionality in the right spots to keep throughput of the network high and the overhead of sending data back and forth low. If you consider each layer to have its own protocol or filter, it is somewhat easier to understand. We can look at communication architectures like a glorified air purifier system. One programme takes a packet of pure air and starts sending it to another programme on a remote node. To insure its purity, each layer of the communication architecture puts a special container around the pure packet of air so that it will not get contaminated on the way. This means that by the time the packet of pure air is out the node, it basically has 6 containers around it (when it is travelling, it is in container 1, the physical layer). When the packet reaches its destination, it travels up the layers, with each protocol removing the container that it knows how to remove, inspecting for damage, and if no damage is present, sending it up the next layer. If a layer finds damage, the packet is thrown away (it is contaminated) and the source node is requested to send another pure packet of air (we cannot have bad air getting to our pure air environment). This is the typical path that most data take when using communication architectures and the layered approach.

To compound misery, communication architectures do much more than send data. Nodes (systems) need to know which nodes are available for access, which services (layers) are active, and in some nodes, if routing is necessary. To do this, many communication architectures keep a database of active nodes, known nodes (nodes the system knows about but that may not necessarily be up and running or available) and nodes that are down.

OPERATION OF THE OSI MODEL IN A NETWORK

To a user, communication appears to be taking place directly between systems, with messages going directly to the destination. Communication is perceived as peer-to-peer, that is, as between corresponding layers in a connected system. Figure 5.4 gives the user's view of communication through the ISO layers. Figure 5.5 gives the system's view of communication through the ISO layers.


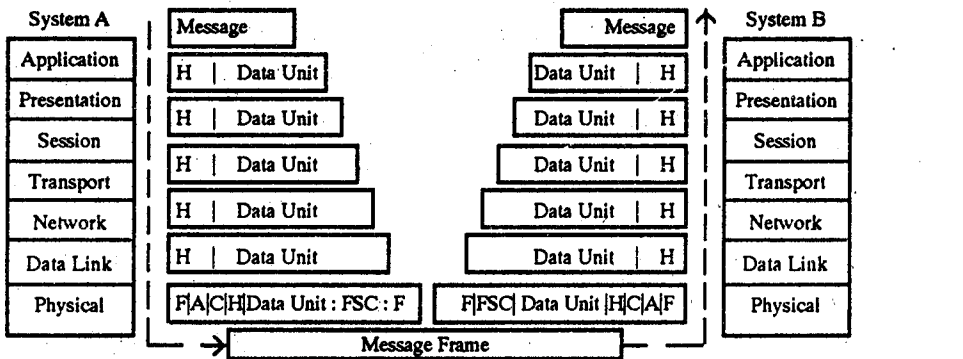


Fig. 5.4 : The User's View of Communications Through the ISO Layers



Each layer manipulates or changes the message in some way : In System A, information is added. In System B, the added information is stripped away until the original message is revealed.

H Message header, added by a layer for its counterpart F Flag C Control
 Data Unit Original message, plus headers A Address FSC

Fig. 5.5 a : The System's View of Communications Through the ISO Layers

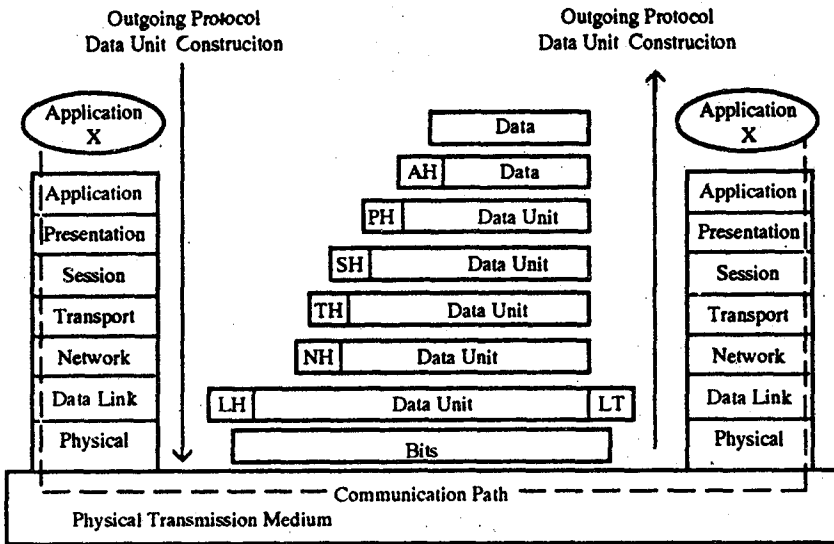


Fig. 5.5 b : The OSI Environment

In actuality, the process is quite complex. When a user sends a message, the original message is embedded in a protocol and passed downward through the interface to the next lower layer. This layer in turn envelops the expanded message within its protocol before forwarding the message. As the message passes downward, each layer adds its protocol and performs its translation/transformation on the message received. A message must pass through all layers between its origination point and the data channel.

Once on the data channel, the message is passed upward through the layers until the requisite layer is reached. As the message moves upward, the surrounding protocols are stripped away and transformations are reversed, layer by layer, until the embedded, original message is revealed.

IEEE 802 PROJECT

A local area network, as defined by the Institute of Electrical and Electronic Engineers' (IEEE) Project 802, is a data communication system allowing a number of independent devices to communicate directly with each other, within a moderately sized geographic area, over a physical communication channel of moderate data rates.

The number and kinds of devices a data local area network should connect, the types of services supported, reliability, and so forth, have been precisely defined in the various 802 committees. Two levels of requirements exist: requirements applicable to networks in general and requirements for specific types of networks. A few of the more general (and hence more relevant for our purposes) requirements have been summarised below. This list is not of all the general requirements, however, nor even of all the relevant requirements.

- **Size:** One LAN should support at least 200 devices and should be able to span at least 2 kilometers. LANs must be capable of being linked to provide service over a greater area.
- **Transmission rate:** Data shall be transmitted through the network at a rate between 1 Megabit per second and 20 Megabits per second.
- **Data communication functions:** The data communication supported should include, but are not restricted to: file transfer and transaction processing; file and database access; terminal support (“dumb” terminals, “smart” terminals, high-speed graphics terminals, etc.); electronic mail; and voicegrams.
- **Attached devices:** Devices interconnected by the LAN should include computers and terminals; mass storage devices; printers, plotters, network and site monitoring and control equipment; bridges and gateways to other networks; telephones; video cameras and monitors; photocopiers; facsimile transceivers.
- **Services:** The LAN should allow a variety of network processes to coexist.
- **Expandability:** Adding or removing devices must be easy. Changes should cause minimal disruption, defined as a transient fault lasting no more than one second.
- **Resource sharing:** When devices need to share LAN facilities, especially the bandwidth of the bus, this sharing must be fair to all devices, even in overload conditions.
- **Reliability:** The LAN should be highly reliable. No more than one packet per year may contain an undetected error.

Very early in the process of defining a LAN and developing standards, the members of the 802 project recognised that no single technology would satisfy all requirements. Specific applications, with their differing priorities, demand different technologies. Therefore, the 802 project divided into several different committees, listed in Figure 5.6, each focusing on creating separate standards. Table 5.5 gives physical layer specifications for IEEE 802 LAN standards.

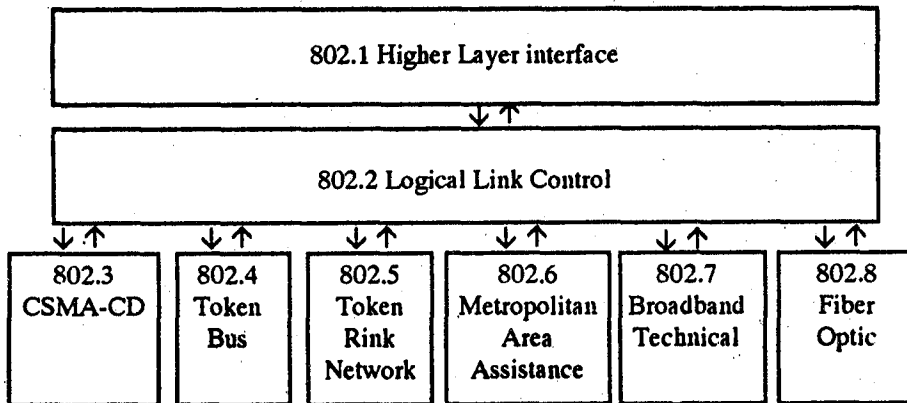


Fig. 5.6 : IEEE 802 Committees

TABLE : 5.5 Physical Layer Specifications for IEEE 802 LAN Standards

	Transmission Medium	Data Signalling Technique	Rate (Mbps)	Maximum Length (m)
IEEE 802.3 (CSMA/CD)				
Original (10 BASE5)	Coaxial Cable (50 ω) (Manchester)	Baseband	10	500
Chaperner (10 BASE2)	Coaxial Cable (50 ω) (Manchester)	Baseband	10	185
10 BASE-T	Unshielded Twisted Pair	Baseband (Manchester)	10	100
Broadband (10 BROAD36)	Coaxial Cable (75 ω)	DPSK	10	3600
IEEE 802.4 (Token Bus)				
Broadband	Coaxial Cable (75 ω)	duobinary AM/PSK	1,5,10	a
Carrierband	Coaxial Cable (75 $\omega\omega$)	FSK	1,5,10	7600
IEEE 802.5 (Token Ring)				
Twisted Pair	Shielded Twisted Pair	Differential Manchester	1,4	b

a = not specified
b = not specified; a maximum of 250 repeaters allowed

802.1 - Higher Layer Interface Standard

The 802.1 committee is not developing standards, but has focused on issues relevant to all other committees such as addressing of messages, internetworking, network management and higher layer interfaces.

802.2 - Logical Link Control Standard

The 802.2 committee has concentrated on functions necessary to provide a reliable communication path between two devices. Levels of service and standard frame format have been defined and accepted by the IEEE as a standard. Currently, the committee is working on network management.

802.3 - CSMA/CD Bus

The 802.3 is aimed at developing a contention bus network. The standards proposed by the committee were virtually identical to Ethernet specifications published by the DEC-Intel-Xerox collaboration: a 10 megabit per second network, which would allow up to 1000 devices to share a baseband coaxial cable. Although the initial proposal was accepted by the IEEE and the ISO as a standard, the committee has been considering changing the type of cable specified to permit use of a thinner wire.

Currently, two other efforts are underway by 802.3. Part of the committee is studying low cost CSMA/CD broadband using a star topology and is in the process of defining how it would work. A second subcommittee has defined a broadband modem which will allow Ethernet to plug into a broadband cable.

802.4 - Token Passing Bus

The 802.4 committee has focused on defining a logical ring on a physical bus so that token passing protocol can be used. Broadband operation with a variety of data rates has been defined.

The specifications have been accepted and are being published as standards. One major problem affecting the usability of a token bus is that the committee is standardising protocols for a complex network that has not been yet widely implemented.

802.5 - Token Passing Ring

The 802.5 has defined a token ring using the star topology to access workstations sequentially. Baseband and broadband versions have been developed. IBM has contributed extensively to this subcommittee. The proposal of the token ring committee was accepted in the fall, 1984.

New Committees

Three new 802 committees were formed during 1984:

- **802.6 - Metropolitan Area Network:** 802.6 is exploring the use of a variety of techniques to send data in a city-wide area. Both CATV and cellular radio are being considered.
- **802.7 - Broadband Technical Assistance:** 802.7 was formed to define broadband LAN standards.
- **802.8 - Fibre Optic:** The 802.8 committee is the most recently formed group. Members have begun exploring the use of fibre optics for a very high speed local area network.

802 Reference Model

Although the 802 Project committee's overall approach is based upon the structure of the International Standards Organisation (ISO) Reference Model of Open System Interconnection, it does differ from the layers defined by the ISO. Moreover, 802 is focused on the lowest ISO layers only, that is, sharing the media. No higher layer protocols are under 802 study.

The 802 Reference Model has three layers, as shown in Figure 5.7.

- **Physical:** concerned with the nature of the transmission medium and the details of device attachment and electrical signalling.

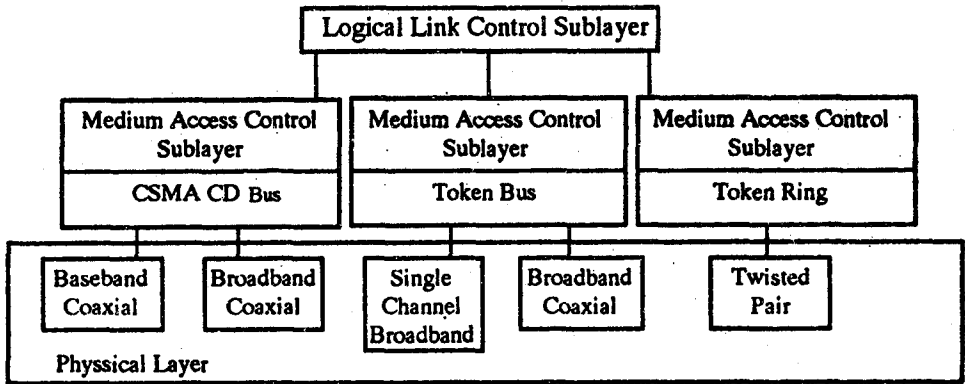


Fig. 5.7 : IEEE 802 Layers

- **Medium Access Control:** focused on methods of sharing a single transmission medium. Typical issues centreon controlling access to the medium, capacity sharing algorithms and station addressing.
- **Logical Link Control:** concerned with providing a reliable communication path between two devices. The relevant protocols cover the flow of frames between stations; establishing, maintaining and terminating communication between devices; and error control.

Figure 5.8 contrasts the 802 layers to the ISO layers. Basically, the 802 Physical Layer is designed to correspond to the ISO Physical Layer. However, rather than being a single layer as in the ISO model, the 802 Physical Layer is itself divided into three parts:

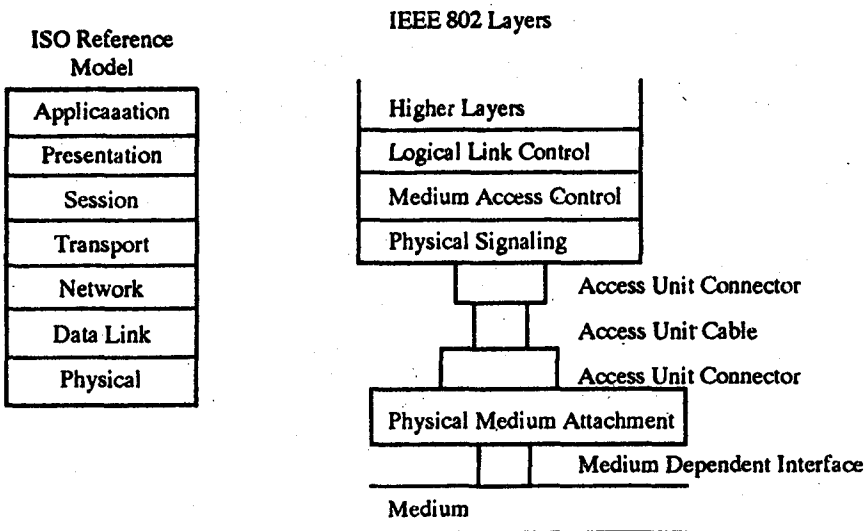


Fig. 5.8 a : IEEE 802 Layers Compared to ISO Layers

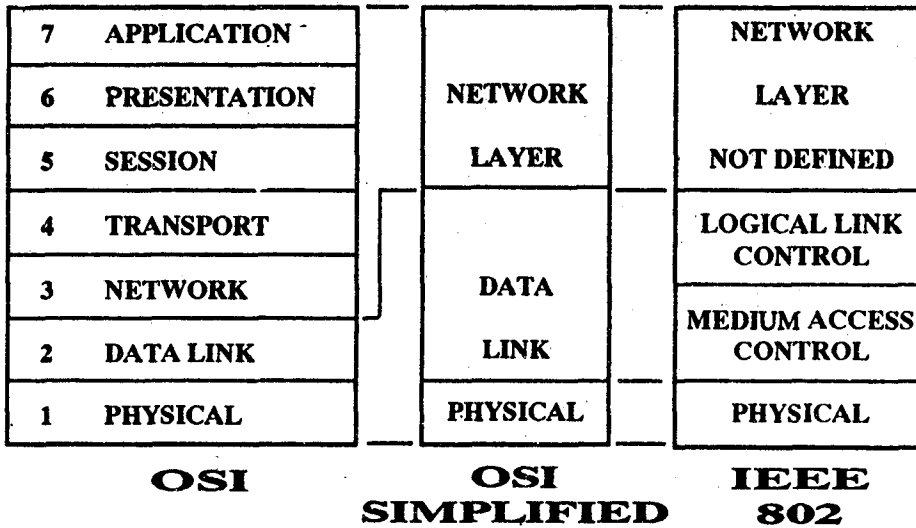


Fig. 5.8 b : OSI and IEEE 802

- The physical signalling sublayer.
- Access-unit interface.
- The physical medium attachment.

This division is intended to permit the most complex part of the layer, the physical signalling sublayer, to be physically combined with the station logic in the network controller.

Jointly, the 802 Medium Access Control Layer and the Logical Link Control Layer correspond to the ISO Data Link Layer. Again, the division has been made to allow combining with other network functions.

ETHERNET

One of the best known and more successful local area networks is Ethernet (see Figure 5.9) developed at Xerox's Palo Alto Research Laboratories. The original Ethernet was designed to link a set of single-user minicomputers that were scattered throughout the research centre. Xerox's immediate goals were to enable the exchange of programmes and data and to provide access to various specialised peripherals.

Ethernet workstations are connected by a single, multidrop, baseband coaxial cable bus using CSMA/CD. As in all contention schemes, the shared channel is a passive broadcast medium with no central control. Access to the channel by stations wishing to transmit is coordinated by the stations themselves.

Workstations are attached to the main bus through a network interface module. The interface buffers and formats messages and subsequently broadcasts the data onto the cable in bursts. The data is in a fixed-length packet containing address information

in the header. Current packet size has been defined as 256 bytes, although the technical specification allows packets ranging from 72 bytes to 1526 bytes.

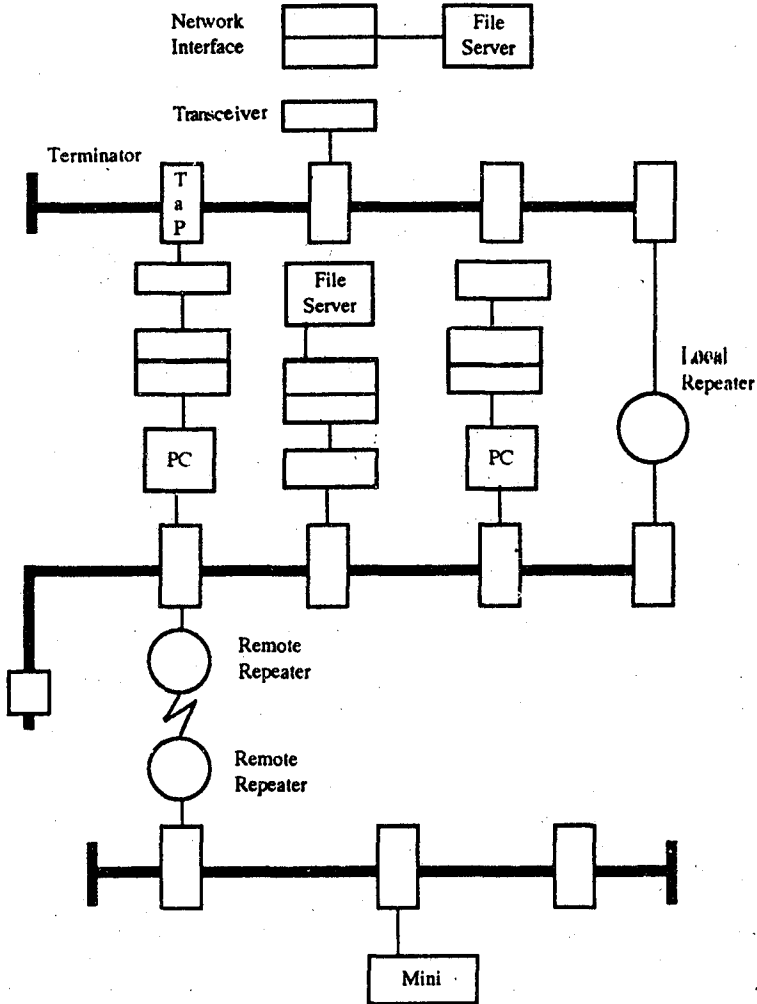


Fig. 5.9 : Ethernet Network

Each station contains address recognition mechanisms, used to identify and accept packets. Every Ethernet workstation, no matter what network it is on, has a unique 48-bit address that is assigned to it and to no other workstation. Hence, when a workstation is moved from one network to another, there is no chance of conflict. This assigning of unique identities has the advantage of flexibility: networks within a company can be physically reconfigured with minimal operating system reconfiguration.

Data on the network moves at a speed of 10 Megabytes per second, over a maximum distance of 2.5 kilometers. No more than 100 workstations can be connected in a 500 meter segment.

Ethernet's strength is that it provides efficient, high-speed resource sharing services within a limited geographic area, at a relatively low cost.

Interest in Ethernet and in local area networks in general, was focused by the 1980 announcement from Digital, Intel and Xerox of a joint project to develop specifications for a local communication network. The project's aim was compatibility, providing sufficient information for various manufacturers such that their widely differing machines could communicate with one another. In effect, the group was establishing a de facto standard.

The attempt was largely successful. The IEEE 802.3 contention bus specification is similar to Ethernet in most details. The two designs are not identical, but are extremely close.

Spurred by these factors, quite a number of vendors have announced hardware and software intended to connect microcomputers into an Ethernet compatible network. At the moment, in fact, the majority of microcomputer LANs use a variation on Ethernet.

- "Ethernet on a chip", that is, an implementation of Ethernet protocols on a single silicon chip, is available from Intel and other chip vendors.
- A fibre optic implementation of Ethernet.
- "Cheapernet", a low cost implementation of Ethernet, is close to being accepted as an IEEE 802.3 standard.

Ethernet Specification

The information in this section is based on the Ethernet Data Link and Physical Layer Specification, Version 1.0 published in 1980 by the Xerox Corporation.

Specification of Ethernet are:

Topology	-	bus
Medium	-	coaxial cable
Access Method	-	CSMA/CD
Speed	-	10 MB
Range	-	2.5 Km
Number of nodes	-	1024
Band	-	baseband

Ethernet is designed to do the following things:

1. To be simple - features which would complicate the design without improving the performance are omitted.
2. To be low cost - in order to be a suitable medium for interconnection of equipment whose cost continues to fall, Ethernet itself should be cheap.

3. To allow compatibility of all Ethernet installations - the specification avoids optional features, thus allowing any Ethernet station to communicate directly with any other, at physical link and data link levels.
4. To allow single nodes, groups or the whole network to be addressed by a transmission.
5. To allow all nodes equal access to the network, on average.
6. To prevent any node interfering with the proper functioning of any other node.
7. To be high-speed - the network should operate at a data rate of 10 Mbs.
8. To be stable - the network performance in terms of data successfully transmitted should not degrade as the amount of data for transmission increases. In other words the system should not clog up as the load increases.
9. To keep delays to a minimum - no data should be kept waiting longer than necessary for transmission.
10. To have a layered architecture - the physical and data link layers specified are completely independent and correspond to the two lowest layers of the OSI model.

Ethernet does not do the following things:

1. Provide full-duplex communication. Only one device can talk at once. The appearance of two-way communication can only be provided by two devices talking alternatively in rapid succession.
2. Provide Error Control. The layers specified only detect bit errors and collisions. Recovery from these and other errors must be handled by the higher layers of the network.
3. Provide Security. There is no encryption or restricted access implied in this specification.
4. Provide variable speeds. The network operates at 10 Mbits/second.
5. Provide a priority control. All nodes have equal access rights to the network.

ARCNET

ARCNET began as a minicomputer distributed processing system offered by Datapoint. It was the first minicomputer local area network to gain a substantial user base. Since 1982, ARCNET has offered interfaces to several microcomputers.

ARC stands for Attached Research Processor Interface units; called a resource interface module (RIM), these are microprocessor-based. The RIMs monitor and control the operation of the network, particularly data transmission, buffer management, error detection, system reconfiguration and related tasks, leaving the attached computers free for user applications.

As Figure 5.10 shows, the RIM consists of four components:

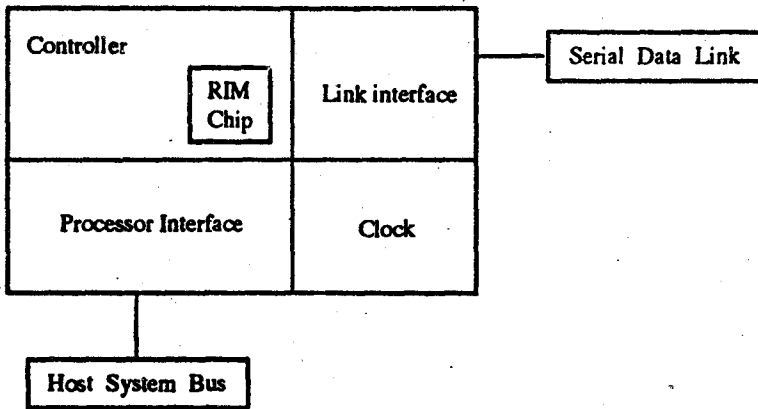


Fig. 5.10 : Resource interface Module

- A controller, consisting of the RIM chip, node ID switch, address and data path information and RAM buffer.
- A clock which provides synchronisation and timing for the network.
- A link interface which connects the controller to the communication media via a serial data link and provides the transmitter and receiver capabilities
- A processor interface which provides the address decoders, bus drivers, buffer access and other functions required to interface the controller to the system bus of the host processor.

Each workstation is connected by baseband coaxial cable to a RIM which is, in turn, connected to a port on a "hub". The hubs serve as amplifiers and connectors for the RIMs within the network. While any one hub can contain a maximum of 16 RIMs, two or more hubs can be joined for larger systems. Although physically the ARCNET resembles a compound star, logically it works as a token passing ring (see Figure 5.11).

To send a message, the workstation writes the message into a RIM buffer and issues a transmit command. The RIM sets a status flag when the token is received and the message is sent. To receive a message, the workstation assigns a RIM buffer to the RIM receiver. Acknowledgement is provided of successful or unsuccessful message receipt.

Control tokens are dispersed among the workstations of the network. All nodes share responsibility for the detection and re-creation of lost tokens, and for recreation of tokens after system reconfiguration. The user is unaware of, and has no control over, details of the token passing.

Each RIM has two addresses. The first is a fixed unique identifier. The second address is a relative one and identifies where on the ring the unit logically falls. When a new RIM becomes active, it sends a reconfiguration burst which terminates all activity

on the network. The RIM that had control of the line releases it; no other RIM claims control. When the RIMs see an idle line, they know the system is being reconfigured. After a quiet period, all workstations reset their id and begin looking for network activity.

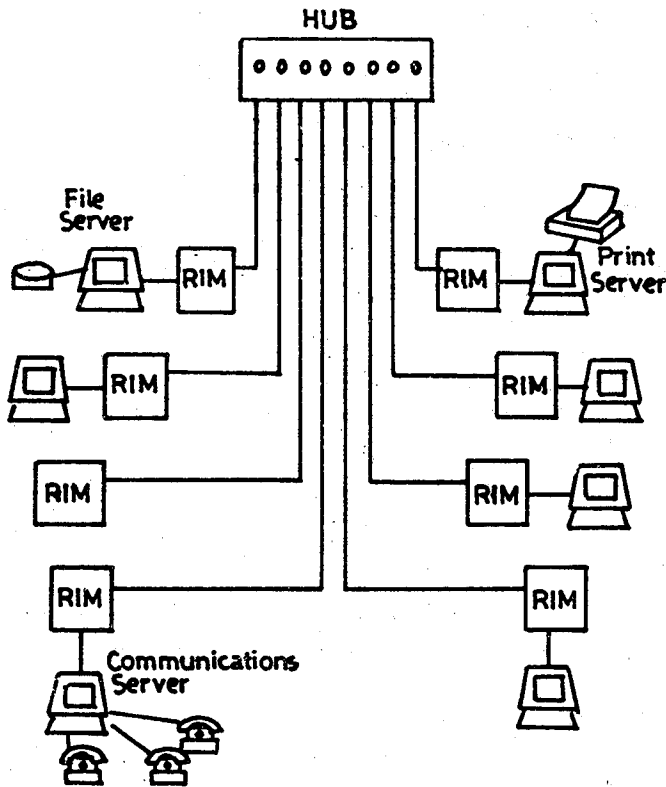


Fig. 5.11 : Arcnet Network

Communication services are restricted to a transfer between a file server and an attached computer. Actual computing is performed separately from where the data are stored, permitting more data to be stored in one location and freeing the computers to process a greater amount of data.

ARCNET imposes few restrictions on the transmission medium. Basically, the receiver at any station must be able to hear the transmitter at any other station. Any two stations must be connected by a single path. Data is transmitted on the network at 2.5 megabits per second. At that rate, the maximum propagation delay between any two stations must be 31 microseconds or less.

Datapoint has made the technical details of its protocol and communication available to the public domain as a way of encouraging other vendors to develop ARC compatible systems. Because only a few relatively inexpensive components are necessary to implement ARC, the appeal is growing.

ARCnet has the following specifications:

Speed	-	2.5 MBPS
Topology	-	Bus or Star
Cable	-	Coaxial
Number of users	-	Depends upon configuration
Maximum segment size	-	100 feet between passive hubs, 2000 feet between active hubs

IBM PC NETWORK

In the fall of 1984, IBM announced a single-wire broadband coaxial CSMA/CD network designed to link IBM PCs. The system is based on an existing network (LocalNet) developed by Sytek, Inc. Although targeted at IBM's three microcomputers - the PC, XT and AT - the network can support other microcomputers that supply a strictly IBM-compatible internal operating system and that are physically bus-compatible with the IBM PC.

The network, shown in Figure 5.12, has two major components:

- A network translator unit, containing a radio frequency modem and one eight-way cable splitter. The translator is the network's head end.
- PC Network Adapter, to be installed in each PC. The adapter card contains most of the network intelligence; it provides encoding and decoding services for the signals and control.

To the basic components must be added the microcomputers (you supply), and broadband coaxial cable, which IBM has thoughtfully made available in the form of cable kits containing all relevant hardware plus pre-cut cable. Short (1 foot), medium (400 foot maximum) and long (800 foot maximum) distance kits with splitters are available.

The minimal network supports up to eight PCs, each of which may be up to 200 feet from the translator's 8-way splitter. Expander kits, containing additional cable taps, are available for users who want to connect more than nine PCs. The expanded network will interconnect up to 72 micros, at a maximum distance of 2000 feet.

The IBM PC Network is, at the moment, one of the few which provides some level of network security. Although logon passwords are not required, files may be protected by a password.

The real lack is the continuing shortage of network software. Despite being based on an existing network, the network operating system is quite limited. Disk caching and print spooling are possible, system management is not.

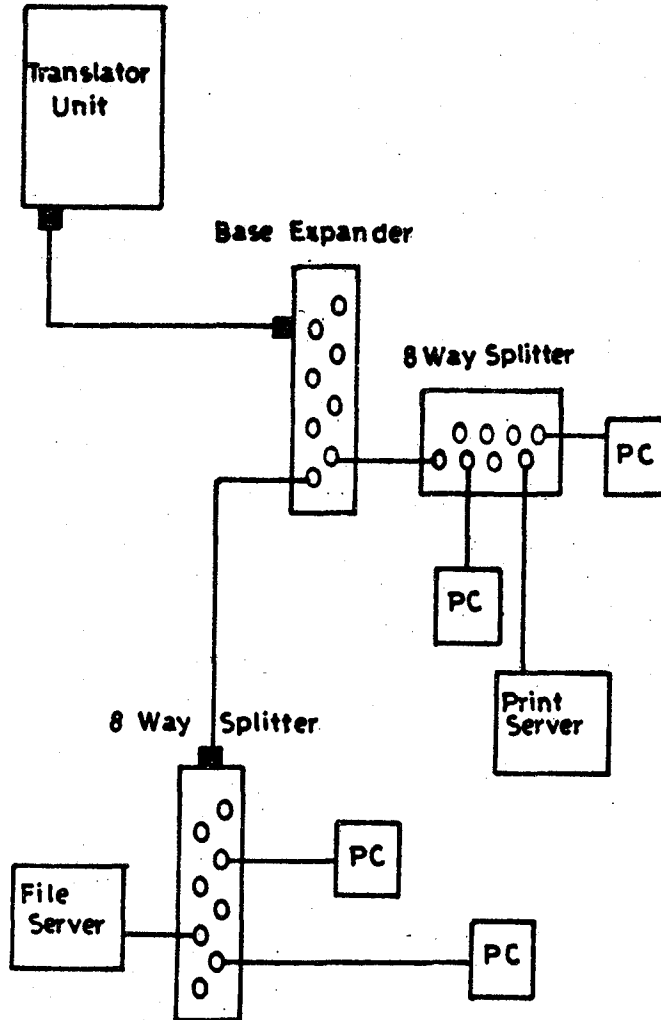


Fig. 5.12 : IBM's PC Network

The PC Network has the following specifications:

Topology	- tree
Medium	- 75 Ohm coaxial cable
Access Method	- CSMA/CD
Speed	- 2 Mbits/second
Range	- 304.8m to translator unit 609.6m between nodes
Number of nodes	- 72
Band	- broadband

IBM TOKEN RING

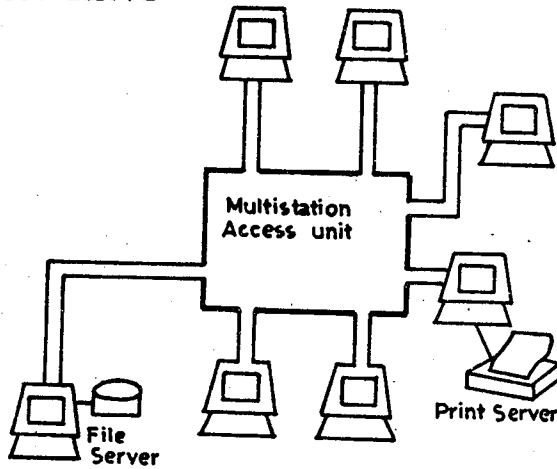


Fig. 5.13 : IBM Token Ring Network

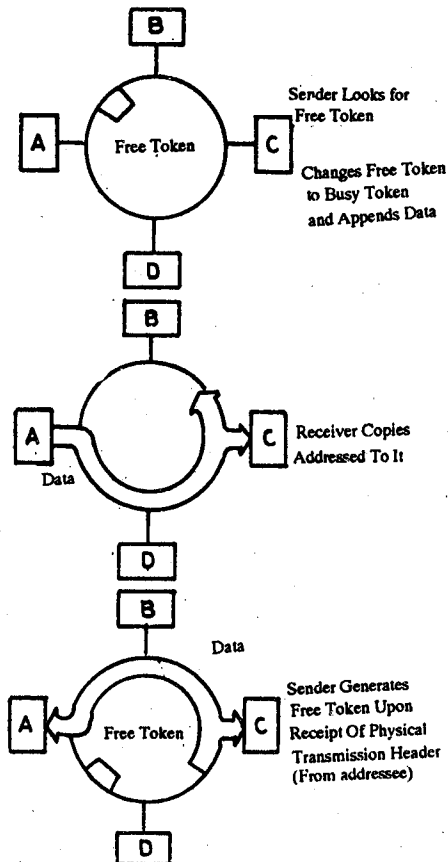


Fig. 5.14 : Token Ring

In October 1985, IBM announced its long expected token ring network. The token ring (see Figures 5.13 and 5.14) is intended as a general purpose local network, aimed at customers who require high-speed communication between intelligent workstations. The initial release supported only microcomputers and gateways to host systems.

Conforming to the IEEE 802.5 standard, the token ring is a 4M bit per second star-wired baseband ring using a token passing protocol. It can support up to 72 microcomputers using twisted pair wiring, or up to 260 micros using shielded data-grade (coaxial) cable. Distance between networked devices is limited to 100 meters on twisted pair, 300 meters on data-grade cable.

Hardware components include the following:

- Network PC Adapter card, which must be installed in each microcomputer connected to the network.
- Network Multistation Access Unit, an eight port wire concentrator which serves as the centre of the star ring.

Required software components includes the most recent version of DOS, the IBM PC operating system. The token ring supports IBM's NetBIOS (Network Basic Input/Output System), programmes written for the IBM PC Network which will run on the token ring.

Other available software includes the following:

- IBM token ring Network/IBM PC Network interconnect programme, which enables the two IBM networks to be coupled. The interconnect programme runs on a micro dedicated as a gateway.
- Asynchronous Communication Server Programme, which provides a gateway, between the token ring, the PC Network and ASCII applications. The programme supports two simultaneous dial-up lines per server.
- Advanced Programme-to-Programme Communication for the Personal Computer, providing a Logical Unit 6.2 Systems Network Architecture (SNA) application programme interface. This interface eventually will enable PCs to communicate as peers with larger IBM systems, a task they cannot perform now.
- Series 1/PC Connect Programme.
- IBM PC Network SNA 3270 Emulation Programme, which has been tested on the token ring.

The IBM Token Ring Network has the following specifications:

Topology	-	ring
Medium	-	shielded or unshielded twisted pair
Access Method	-	token passing
Speed	-	4 Mbs

Range	-	N/A
Number of nodes	-	72 (unshielded), 260 (shielded)
Transmission	-	baseband

SUMMARY

A network architecture defines protocols, message formats and standards to which products must conform in order to connect properly with the network. Architectures are developed by standards organisations, common carriers and computer and network vendors. Network architectures use a layered approach, whereby functions are organised into groups and assigned to specific functional layers in the architecture. Network architectures define the interfaces between layers in a given network node and within the same layer in two different nodes.

OSI provides a generalised model of system interconnection. It encompasses seven layers: application, presentation, session, transport, network, data link and physical. IEEE Project 802 has developed a set of standards for local area networks. These standards specify in detail protocols and data formats for the physical and data link layers. In the IEEE approach, the data link layer is divided into two sublayers: logical link control and media access control. IEEE Project 802 documents families of standards that define the functions performed by the three layers and sublayers described by the architecture.