1

Internetworking and Layered Models

The Internet today is a widespread information infrastructure, but it is inherently an insecure channel for sending messages. When a message (or packet) is sent from one Website to another, the data contained in the message are routed through a number of intermediate sites before reaching its destination. The Internet was designed to accommodate heterogeneous platforms so that people who are using different computers and operating systems can communicate. The history of the Internet is complex and involves many aspects – technological, organisational and community. The Internet concept has been a big step along the path towards electronic commerce, information acquisition and community operations.

Early ARPANET researchers accomplished the initial demonstrations of packet-switching technology. In the late 1970s, the growth of the Internet was recognised and subsequently a growth in the size of the interested research community was accompanied by an increased need for a coordination mechanism. The Defense Advanced Research Projects Agency (DARPA) then formed an International Cooperation Board (ICB) to coordinate activities with some European countries centered on packet satellite research, while the Internet Configuration Control Board (ICCB) assisted DARPA in managing Internet activity. In 1983, DARPA recognised that the continuing growth of the Internet community demanded a restructuring of coordination mechanisms. The ICCB was disbanded and in its place the Internet Activities Board (IAB) was formed from the chairs of the Task Forces. The IAB revitalised the Internet Engineering Task Force (IETF) as a member of the IAB. By 1985, there was a tremendous growth in the more practical engineering side of the Internet. This growth resulted in the creation of a substructure to the IETF in the form of working groups. DARPA was no longer the major player in the funding of the Internet. Since then, there has been a significant decrease in Internet activity at DARPA. The IAB recognised the increasing importance of IETF, and restructured to recognise the Internet Engineering Steering Group (IESG) as the major standards review body. The IAB also restructured to create the Internet Research Task Force (IRTF) along with the IETF.

Internet Security. Edited by M.Y. Rhee
 ISBN 0-470-85285-2

Since the early 1980s, the Internet has grown beyond its primarily research roots, to include both a broad user community and increased commercial activity. This growth in the commercial sector brought increasing concern regarding the standards process. Increased attention was paid to making progress, eventually leading to the formation of the Internet Society in 1991. In 1992, the Internet Activities Board was reorganised and renamed the Internet Architecture board (IAB) operating under the auspices of the Internet Society. The mutually supportive relationship between the new IAB, IESG and IETF led to them taking more responsibility for the approval of standards, along with the provision of services and other measures which would facilitate the work of the IETF.

1.1 Networking Technology

Data signals are transmitted from one device to another using one or more types of transmission media, including twisted-pair cable, coaxial cable and fibre-optic cable. A message to be transmitted is the basic unit of network communications. A message may consist of one or more cells, frames or packets which are the elemental units for network communications. Networking technology includes everything from local area networks (LANs) in a limited geographic area such as a single building, department or campus to wide area networks (WANs) over large geographical areas that may comprise a country, a continent or even the whole world.

1.1.1 Local Area Networks (LANs)

A local area network (LAN) is a communication system that allows a number of independent devices to communicate directly with each other in a limited geographic area such as a single office building, a warehouse or a campus. LANs are standardised by three architectural structures: Ethernet, token ring and fibre distributed data interface (FDDI).

1.1.1.1 Ethernet

Ethernet is a LAN standard originally developed by Xerox and later extended by a joint venture between Digital Equipment Corporation (DEC), Intel Corporation and Xerox. The access mechanism used in an Ethernet is called Carrier Sense Multiple Access with Collision Detection (CSMA/CD). In CSMA/CD, before a station transmits data, it must check the medium where any other station is currently using the medium. If no other station is transmitting, the station can send its data. If two or more stations send data at the same time, it may result in a collision. Therefore, all stations should continuously check the medium to detect any collision. If a collision occurs, all stations ignore the data received. The sending stations wait for a period of time before resending the data. To reduce the possibility of a second collision, the sending stations individually generate a random number that determinates how long the station should wait before resending data.

1.1.1.2 Token Ring

Token ring, a LAN standard originally developed by IBM, uses a logical ring topology. The access method used by CSMA/CD may result in collisions. Therefore, stations may

attempt to send data many times before a transmission captures a perfect link. This redundancy can create delays of indeterminable length if traffic is heavy. There is no way to predict either the occurrence of collisions or the delays produced by multiple stations attempting to capture the link at the same time. Token ring resolves this uncertainty by making stations take turns in sending data.

As an access method, the token is passed from station to station in sequence until it encounters a station with data to send. The station to be sent data waits for the token. The station then captures the token and sends its data frame. This data frame proceeds around the ring and each station regenerates the frame. Each intermediate station examines the destination address, finds that the frame is addressed to another station, and relays it to its neighbouring station. The intended recipient recognises its own address, copies the message, checks for errors and changes four bits in the last byte of the frame to indicate that the address has been recognised and the frame copied. The full packet then continues around the ring until it returns to the station that sent it.

1.1.1.3 Fiber Distributed Data Interface (FDDI)

FDDI is a LAN protocol standardised by ANSI and ITU-T. It supports data rates of 100 Mbps and provides a high-speed alternative to Ethernet and token ring. When FDDI was designed, the data rate of 100 Mbps required fibre-optic cable.

The access method in FDDI is also called token passing. In a token ring network, a station can send only one frame each time it captures the token. In FDDI, the token passing mechanism is slightly different in that access is limited by time. Each station keeps a timer which shows when the token should leave the station. If a station receives the token earlier than the designated time, it can keep the token and send data until the scheduled leaving time. On the other hand, if a station receives the token at the designated time or later than this time, it should let the token pass to the next station and wait for its next turn.

FDDI is implemented as a dual ring. In most cases, data transmission is confined to the primary ring. The secondary ring is provided in case of the primary ring's failure. When a problem occurs on the primary ring, the secondary ring can be activated to complete data circuits and maintain service.

1.1.2 Wide Area Networks (WANs)

A WAN provides long-distance transmission of data, voice, image and video information over large geographical areas that may comprise a country, a continent or even the world. In contrast to LANs (which depend on their own hardware for transmission), WANs can utilise public, leased or private communication devices, usually in combination.

1.1.2.1 PPP

The Point-to-Point Protocol (PPP) is designed to handle the transfer of data using either asynchronous modem links or high-speed synchronous leased lines. The PPP frame uses the following format:

- **Flag field:** Each frame starts with a one-byte flag whose value is 7E(0111 1110). The flag is used for synchronisation at the bit level between the sender and receiver.
- **Address field:** This field has the value of FF(1111 1111).
- **Control field:** This field has the value of 03(0000 0011).
- **Protocol field:** This is a two-byte field whose value is 0021(0000 0000 0010 0001) for TCP/IP.
- **Data field:** The data field ranges up to 1500 bytes.
- **CRC:** This is a two-byte cyclic redundancy check. Cyclic redundancy check (CRC) is implemented in the physical layer for use in the data link layer. A sequence of redundant bits (CRC) is appended to the end of a data unit so that the resulting data unit becomes exactly divisible by a predetermined binary number. At its destination, the incoming data unit is divided by the same number. If there is no remainder, the data unit is accepted. If a remainder exists, the data unit has been damaged in transit and therefore must be rejected.

1.1.2.2 X.25

X.25 is widely used, as the packet switching protocol provided for use in a WAN. It was developed by the ITU-T in 1976. X.25 is an interface between data terminal equipment and data circuit terminating equipment for terminal operations at the packet mode on a public data network.

X.25 defines how a packet mode terminal can be connected to a packet network for the exchange of data. It describes the procedures necessary for establishing connection, data exchange, acknowledgement, flow control and data control.

1.1.2.3 Frame Relay

Frame relay is a WAN protocol designed in response to X.25 deficiencies. X.25 provides extensive error-checking and flow control. Packets are checked for accuracy at each station to which they are routed. Each station keeps a copy of the original frame until it receives confirmation from the next station that the frame has arrived intact. Such station-to-station checking is implemented at the data link layer of the OSI model, but X.25 only checks for errors from source to receiver at the network layer. The source keeps a copy of the original packet until it receives confirmation from the final destination. Much of the traffic on an X.25 network is devoted to error-checking to ensure reliability of service. Frame relay does not provide error-checking or require acknowledgement in the data link layer. Instead, all error-checking is left to the protocols at the network and transport layers, which use the frame relay service. Frame relay only operates at the physical and data link layer.

1.1.2.4 Asynchronous Transfer Mode (ATM)

ATM is a revolutionary idea for restructuring the infrastructure of data communication. It is designed to support the transmission of data, voice and video through a high data-rate transmission medium such as fibre-optic cable. ATM is a protocol for transferring cells. A cell is a small data unit of 53 bytes long, made of a 5-byte header and a 48-byte payload.

The header contains a virtual path identifier (VPI) and a virtual channel identifier (VCI). These two identifiers are used to route the cell through the network to the final destination.

An ATM network is a connection-oriented cell switching network. This means that the unit of data is not a packet as in a packet switching network, or a frame as in a frame relay, but a cell. However, ATM, like X.25 and frame relay, is a connection-oriented network, which means that before two systems can communicate, they must make a connection. To start up a connection, a system uses a 20-byte address. After the connection is established, the combination of VPI/VCI leads a cell from its source to its final destination.

1.2 Connecting Devices

Connecting devices are used to connect the segments of a network together or to connect networks to create an internetwork. These devices are classified into five categories: switches, repeaters, bridges, routers and gateways. Each of these devices except the first one (switches) interacts with protocols at different layers of the OSI model.

Repeaters forward all electrical signals and are active only at the physical layer. Bridges store and forward complete packets and affect the flow control of a single LAN. Bridges are active at the physical and data link layers. Routers provide links between two separate LANs and are active in the physical, data link and network layers. Finally, gateways provide translation services between incompatible LANs or applications, and are active in all layers.

Connection devices that interact with protocols at different layers of the OSI model are shown in Figure 1.1.

1.2.1 Switches

A switched network consists of a series of interlinked switches. Switches are hardware/software devices capable of creating temporary connections between two or more devices to the switch but not to each other. Switching mechanisms are generally classified into three methods: circuit switching, packet switching and message switching.

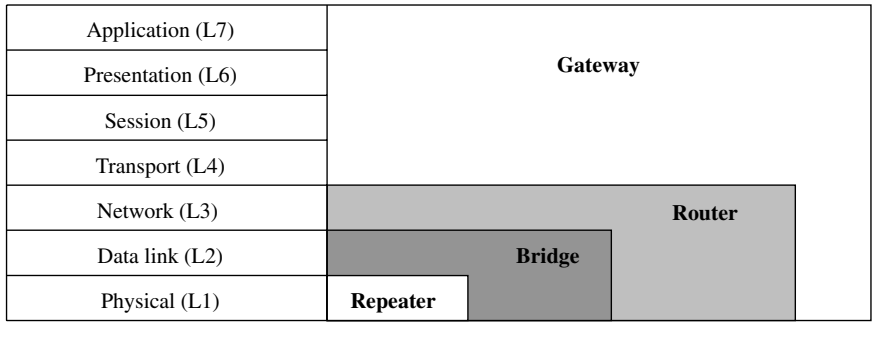

Figure 1.1 Connecting devices.

- Circuit switching creates a direct physical connection between two devices such as telephones or computers. Once a connection is made between two systems, circuit switching creates a dedicated path between two end users. The end users can use the path for as long as they want.
- Packet switching is one way to provide a reasonable solution for data transmission. In a packet-switched network, data are transmitted in discrete units of variable-length blocks called packets. Each packet contains not only data, but also a header with control information. The packets are sent over the network node to node. At each node, the packet is stored briefly before being routed according to the information in its header.

  In the datagram approach to packet switching, each packet is treated independently of all others as though it exists alone. In the virtual circuit approach to packet switching, if a single route is chosen between sender and receiver at the beginning of the session, all packets travel one after another along that route. Although these two approaches seem the same, there exists a fundamental difference between them. In circuit switching, the path between the two end users consists of only one channel. In the virtual circuit, the line is not dedicated to two users. The line is divided into channels and each channel can use one of the channels in a link.
- Message switching is known as the store and forwarding method. In this approach, a computer (or a node) receives a message, stores it until the appropriate route is free, then sends it out. This method has now been phased out.

1.2.2 Repeaters

A repeater is an electronic device that operates on the physical layer only of the OSI model. A repeater boosts the transmission signal from one segment and continues the signal to another segment. Thus, a repeater allows us to extend the physical length of a network. Signals that carry information can travel a limited distance within a network before degradation of the data integrity due to noise. A repeater receives the signal before attenuation, regenerates the original bit pattern and puts the restored copy back on to the link.

1.2.3 Bridges

Bridges operate in both the physical and the data link layers of the OSI model. A single bridge connects different types of networks together and promotes interconnectivity between networks. Bridges divide a large network into smaller segments. Unlike repeaters, bridges contain logic that allows them to keep separate the traffic for each segment. Bridges are smart enough to relay a frame towards the intended recipient so that traffic can be filtered. In fact, this filtering operation makes bridges useful for controlling congestion, isolating problem links and promoting security through this partitioning of traffic.

A bridge can access the physical addresses of all stations connected to it. When a frame enters a bridge, the bridge not only regenerates the signal but also checks the address of the destination and forwards the new copy to the segment to which the address belongs. When a bridge encounters a packet, it reads the address contained in the frame and compares that address with a table of all the stations on both segments. When it finds

a match, it discovers to which segment the station belongs and relays the packet to that segment only.

1.2.4 Routers

Routers operate in the physical, data link and network layers of the OSI model. The Internet is a combination of networks connected by routers. When a datagram goes from a source to a destination, it will probably pass through many routers until it reaches the router attached to the destination network. Routers determine the path a packet should take. Routers relay packets among multiple interconnected networks. In particular, an IP router forwards IP datagrams among the networks to which it connects. A router uses the destination address on a datagram to choose a next-hop to which it forwards the datagram. A packet sent from a station on one network to a station on a neighbouring network goes first to a jointly held router, which switches it over the destination network. In fact, the easiest way to build the Internet is to connect two or more networks with a router. Routers provide connections to many different types of physical networks: Ethernet, token ring, point-to-point links, FDDI and so on.

- The routing module receives an IP packet from the processing module. If the packet is to be forwarded, it should be passed to the routing module. It finds the IP address of the next station along with the interface number from which the packet should be sent. It then sends the packet with information to the fragmentation module. The fragmentation module consults the MTU table to find the maximum transfer unit (MTU) for the specific interface number.
- The routing table is used by the routing module to determine the next-hop address of the packet. Every router keeps a routing table that has one entry for each destination network. The entry consists of the destination network IP address, the shortest distance to reach the destination in hop count, and the next router (next hop) to which the packet should be delivered to reach its final destination. The hop count is the number of networks a packet enters to reach its final destination. A router should have a routing table to consult when a packet is ready to be forwarded. The routing table should specify the optimum path for the packet. The table can be either static or dynamic. A static table is one that is not changed frequently, but a dynamic table is one that is updated automatically when there is a change somewhere in the Internet. Today, the Internet needs dynamic routing tables.
- A metric is a cost assigned for passing through a network. The total metric of a particular router is equal to the sum of the metrics of networks that comprise the route. A router chooses the route with the shortest (smallest value) metric. The metric assigned to each network depends on the type of protocol. The Routing Information Protocol (RIP) treats each network as one hop count. So if a packet passes through 10 networks to reach the destination, the total cost is 10 hop counts. The Open Shortest Path First protocol (OSPF) allows the administrator to assign a cost for passing through a network based on the type of service required. A route through a network can have different metrics (costs). OSPF allows each router to have several routing tables based on the required type of service. The Border Gateway Protocol (BGP) defines the metric

totally differently. The policy criterion in BGP is set by the administrator. The policy defines the paths that should be chosen.

1.2.5 Gateways

Gateways operate over the entire range in all seven layers of the OSI model. Internet routing devices have traditionally been called gateways. A gateway is a protocol converter which connects two or more heterogeneous systems and translates among them. The gateway thus refers to a device that performs protocol translation between devices. A gateway can accept a packet formatted for one protocol and convert it to a packet formatted for another protocol before forwarding it. The gateway understands the protocol used by each network linked into the router and is therefore able to translate from one to another.

1.3 The OSI Model

The Ethernet, originally called the Alto Aloha network, was designed by the Xerox Palo Alto Research Center in 1973 to provide communication for research and development CP/M computers. When in 1976 Xerox started to develop the Ethernet as a 20 Mbps product, the network prototype was called the Xerox Wire. In 1980, when the Digital, Intel and Xerox standard was published to make it a LAN standard at 10 Mbps, Xerox Wire changed its name back to Ethernet. Ethernet became a commercial product in 1980 at 10 Mbps. The IEEE called its Ethernet 802.3 standard CSMA/CD (or carrier sense multiple access with collision detection). As the 802.3 standard evolved, it has acquired such names as Thicknet (IEEE 10Base-5), Thinnet or Cheapernet (10Base-2), Twisted Ethernet (10Base-T) and Fast Ethernet (100Base-T).

The design of Ethernet preceded the development of the seven-layer OSI model. The Open System Interconnect (OSI) model was developed and published in 1982 by the International Organisation for Standardisation (ISO) as a generic model for data communication. The OSI model is useful because it is a broadly based document, widely available and often referenced. Since modularity of communication functions is a key design criterion in the OSI model, vendors who adhere to the standards and guidelines of this model can supply Ethernet-compatible devices, alternative Ethernet channels, higher-performance Ethernet networks and bridging protocols that easily and reliably connect other types of data network to Ethernet.

Since the OSI model was developed after Ethernet and Signaling System #7 (SS7), there are obviously some discrepancies between these three protocols. Yet the functions and processes outlined in the OSI model were already in practice when Ethernet or SS7 was developed. In fact, SS7 networks use point-to-point configurations between signalling points. Due to the point-to-point configurations and the nature of the transmissions, the simple data link layer does not require much complexity.

The OSI reference model specifies the seven layers of functionality, as shown in Figure 1.2. It defines the seven layers from the physical layer (which includes the network adapters), up to the application layer, where application programs can access network services. However, the OSI model does not define the protocols that implement the functions at each layer. The OSI model is still important for compatibility, protocol independence

Layer No.	OSI Layer	Functionality
7	Application	• Provides user interface • System computing and user application process • Of the many application services, this layer provides support for services such as e-mail, remote file access and transfer, message handling services (X.400) to send an e-mail message, directory services (X.500) for distributed database sources and access for global information about various objects and services
6	Presentation	• Data interpretation (compression, encryption, formatting and syntax selection) and code transformations
5	Session	• Administrative control of transmissions and transfers between nodes • Dialogue control between two systems • Synchronisation process by inserting checkpoints into data stream
4	Transport	• Source-to-destination delivery of entire message • Message segmentation at the sending layer and reassembling at the receiving layer • Transfer control by either connectionless or connection-oriented mechanism for delivering packets • Flow control for end-to-end services • Error control based on performing end-to-end rather than a single link
3	Network	• Source-to-destination delivery of individual packets • Routing or switching packets to final destination • Logical addressing to help distinguish the source/destination systems
2	Data Link	• Framing, physical addressing, data flow control, access control and error control
1	Physical	• Physical control of the actual data circuit (electrical, mechanical and optical)

Figure 1.2 ISO/OSI model.

and the future growth of network technology. Implementations of the OSI model stipulate communication between layers on two processors and an interface for interlayer communication on one processor. Physical communication occurs only at layer 1. All other layers communicate downward (or upward) to lower (or higher) levels in steps through protocol stacks.

The following briefly describes the seven layers of the OSI model:

1. *Physical layer.* The physical layer provides the interface with physical media. The interface itself is a mechanical connection from the device to the physical medium used to transmit the digital bit stream. The mechanical specifications do not specify the electrical characteristics of the interface, which will depend on the medium being used and the type of interface. This layer is responsible for converting the digital

data into a bit stream for transmission over the network. The physical layer includes the method of connection used between the network cable and the network adapter, as well as the basic communication stream of data bits over the network cable. The physical layer is responsible for the conversion of the digital data into a bit stream for transmission when using a device such as a modem, and even light, as in fibre optics. For example, when using a modem, digital signals are converted into analogue audible tones which are then transmitted at varying frequencies over the telephone line. The OSI model does not specify the medium, only the operative functionality for a standardised communication protocol. The transmission media layer specifies the physical medium used in constructing the network, including size, thickness and other characteristics.

2. *Data link layer.* The data link layer represents the basic communication link that exists between computers and is responsible for sending frames or packets of data without errors. The software in this layer manages transmissions, error acknowledgement and recovery. The transceivers are mapped data units to data units to provide physical error detection and notification and link activation/deactivation of a logical communication connection. Error control refers to mechanisms to detect and correct errors that occur in the transmission of data frames. Therefore, this layer includes error correction, so when a packet of data is received incorrectly, the data link layer makes system send the data again. The data link layer is also defined in the IEEE 802.2 logical link control specifications.

   Data link control protocols are designed to satisfy a wide variety of data link requirements:

   – High-level Data Link Control (HDLC) developed by the International Organisation for Standardisation (ISO 3309, ISO 4335);
   – Advanced Data Communication Control Procedures (ADCCP) developed by the American National Standards Institute (ANSI X3.66);
   – Link Access Procedure, Balanced (LAP-B) adopted by the CCITT as part of its X.25 packet-switched network standard;
   – Synchronous Data Link Control (SDLC) is not a standard, but is in widespread use. There is practically no difference between HDLC and ADCCP. Both LAP-B and SDLC are subsets of HDLC, but they include several additional features.

3. *Network layer.* The network layer is responsible for data transmission across networks. This layer handles the routing of data between computers. Routing requires some complex and crucial techniques for a packet-switched network design. To accomplish the routing of packets sending from a source and delivering to a destination, a path or route through the network must be selected. This layer translates logical network addressing into physical addresses and manages issues such as frame fragmentation and traffic control. The network layer examines the destination address and determines the link to be used to reach that destination. It is the borderline between hardware and software. At this layer, protocol mechanisms activate data routing by providing network address resolution, flow control in terms of segmentation and blocking and collision control (Ethernet). The network layer also provides service selection,

connection resets and expedited data transfers. The Internet Protocol (IP) runs at this layer.

The IP was originally designed simply to interconnect as many sites as possible without undue burdens on the type of hardware and software at different sites. To address the shortcomings of the IP and to provide more a reliable service, the Transmission Control Protocol (TCP) is stacked on top of the IP to provide end-to-end service. This combination is known as TCP/IP and is used by most Internet sites today to provide a reliable service.

4. *Transport layer.* The transport layer is responsible for ensuring that messages are delivered error-free and in the correct sequence. This layer splits messages into smaller segments if necessary and provides network traffic control of messages. Traffic control is a technique for ensuring that a source does not overwhelm a destination with data. When data is received, a certain amount of processing must take place before the buffer is clear and ready to receive more data. In the absence of flow control, the receiver's buffer may overflow while it is processing old data. The transport layer, therefore, controls data transfer and transmission. This software is called Transmission Control Protocol (TCP), common on most Ethernet networks, or System Packet Exchange (SPE), a corresponding Novell specification for data exchange. Today most Internet sites use the TCP/IP protocol along with ICMP to provide a reliable service.

5. *Session layer.* The session layer controls the network connections between the computers in the network. The session layer recognises nodes on the LAN and sets up tables of source and destination addresses. It establishes a handshake for each session between different nodes. Technically, this layer is responsible for session connection (i.e. for creating, terminating and maintaining network sessions), exception reporting, coordination of send/receive modes and data exchange.

6. *Presentation layer.* The presentation layer is responsible for the data format, which includes the task of hashing the data to reduce the number of bits (hash code) that will be transferred. This layer transfers information from the application software to the network session layer to the operating system. The interface at this layer performs data transformations, data compression, data encryption, data formatting, syntax selection (i.e. ASCII, EBCDIC or other numeric or graphic formats), and device selection and control. It actually translates data from the application layer into the format used when transmitting across the network. On the receiving end, this layer translates the data back into a format that the application layer can understand.

7. *Application layer.* The application layer is the highest layer defined in the OSI model and is responsible for providing user-layer applications and network management functions. This layer supports identification of communicating partners, establishes authority to communicate, transfers information and applies privacy mechanisms and cost allocations. It is usually a complex layer with a client/server, a distributed database, data replication and synchronisation. The application layer supports file services, print services, remote login and e-mail. The application layer is the network system software that supports user-layer applications, such as word or data processing, CAD/CAM, document storage and retrieval and image scanning.

1.4 TCP/IP Model

A protocol is a set of rules governing the way data will be transmitted and received over data communication networks. Protocols are then the rules that determine everything about the way a network operates. Protocols must provide reliable, error-free communication of user data as well as a network management function. Therefore, protocols govern how applications access the network, the way that data from an application is divided into packets for transmission through cable, and which electrical signals represent data on a network cable.

The OSI model, defined by a seven-layer architecture, is partitioned into a vertical set of layers, as illustrated in Figure 1.2. The OSI model is based on open systems and peer-to-peer communications. Each layer performs a related subset of the functions required to communicate with another system. Each system contains seven layers. If a user or application entity A wishes to send a message to another user or application entity B, it invokes the application layer (layer 7). Layer 7 (corresponding to application A) establishes a peer relationship with layer 7 of the target machine (application B), using a layer 7 protocol.

In an effort to standardise a way of looking at network protocols, the TCP/IP four-layer model is created with reference to the seven-layer OSI model, as shown in Figure 1.3. The protocol suite is designed in distinct layers to make it easier to substitute one protocol for another. The protocol suite governs how data is exchanged above and below each protocol

Electronic payment system	**Internet security**
E-cash, Mondex, Proton, Visa Cash, SET, CyberCash, CyberCoin, E-check, First Virtual	SSL, TLS, S/HTTP, IPsec, SOCKS V5, PEM, PGP, S/MIME

OSI model (7 layers)	**TCP/IP model (4 layers)**	**Internet protocol suite**
Application	Application	HTTP, FTP, TFTP, NFS, RPC, XDR, SMTP, POP, IMAP, MIME, SNMP, DNS, RIP, OSPF, BGP, TELNET, Rlogin
Presentation		
Session	Transport	TCP, UDP
Transport		
Network	Internet	IP, ICMP, IGMP, ARP, RARP
Data link	Network access	Ethernet, token ring, FDDI, PPP, X.25, frame replay, ATM
Physical		

Figure 1.3 The TCP/IP model and Internet protocol suite.

layer. When protocols are designed, specifications set out how a protocol exchanges data with a protocol layered above or below it.

Both the OSI model and the TCP/IP layered model are based on many similarities, but there are philosophical and practical differences between the two models. However, they both deal with communications among heterogeneous computers.

Since TCP was developed before the OSI model, the layers in the TCP/IP protocol model do not exactly match those in the OSI model. The important fact is the hierarchical ordering of protocols. The TCP/IP model is made up of four layers: application layer, transport layer, Internet layer and network access layer. These will be discussed below.

1.4.1 Network Access Layer

The network access layer contains protocols that provide access to a communication network. At this layer, systems are interfaced to a variety of networks. One function of this layer is to route data between hosts attached to the same network. The services to be provided are flow control and error control between hosts. The network access layer is invoked either by the Internet layer or the application layer. This layer provides the device drivers that support interactions with communications hardware such as the token ring or Ethernet. The IEEE token ring, referred to as the Newhall ring, is probably the oldest ring control technique and has become the most popular ring access technique in the USA. The Fiber Distributed Data Interface (FDDI) is a standard for a high-speed ring LAN. Like the IEEE 802 standard, FDDI employs the token ring algorithm.

1.4.2 Internet Layer

The Internet layer provides a routing function. Therefore, this layer consists of the procedures required within hosts and gateways to allow data to traverse multiple networks. A gateway connecting two networks relays data between networks using an internetwork protocol. This layer consists of the Internet Protocol (IP) and the Internet Control Message Protocol (ICMP).

1.4.3 Transport Layer

The transport layer delivers data between two processes on different host computers. A protocol entity at this level provides a logical connection between higher-level entities. Possible services include error and flow controls and the ability to deal with control signals not associated with a logical data connection. This layer contains the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP).

1.4.4 Application Layer

This layer contains protocols for resource sharing and remote access. The application layer actually represents the higher-level protocols that are used to provide a direct interface with users or applications. Some of the important application protocols are File Transfer Protocol (FTP) for file transfers, HyperText Transfer Protocol (HTTP) for the World Wide Web, and Simple Network Management Protocol (SNMP) for controlling network devices.

The Domain Naming Service (DNS) is also useful because it is responsible for converting numeric IP addresses into names that can be more easily remembered by users. Many other protocols dealing with the finer details of applications are included in this application layer. These include Simple Mail Transport Protocol (SMTP), Post Office Protocol (POP), Internet Mail Access Protocol (IMAP), Internet Control Message Protocol (ICMP) for e-mail, Privacy Enhanced Mail (PEM), Pretty Good Privacy (PGP) and Secure Multimedia Internet Mail Extensions (S/MIME) for e-mail security. All protocols contained in the TCP/IP suite are fully described in Chapter 2.