how these layers combine to achieve a task such as sending and receiving an e-mail between two computers on separate local area networks (LANs) that are connected via the Internet.

The process starts with the sender typing a message into a PC e-mail application (Figure 2-1). When the user selects "Send", the operating system combines the message with a set of Application layer (Layer 7) instructions that will eventually be read and actioned by the corresponding operating system and application on the receiving computer.

The message plus Layer 7 instructions is then passed to the part of sender's operating system that deals with Layer 6 presentation tasks. These include the translation of data between application layer formats as well as some types of security such as Secure Socket Layer (SSL) encryption. This process continues down through the successive software layers, with the message gathering additional instructions or control elements at each level.

By Layer 3 — the Network layer — the message will be broken down into a sequence of data packets, each carrying a source and destination

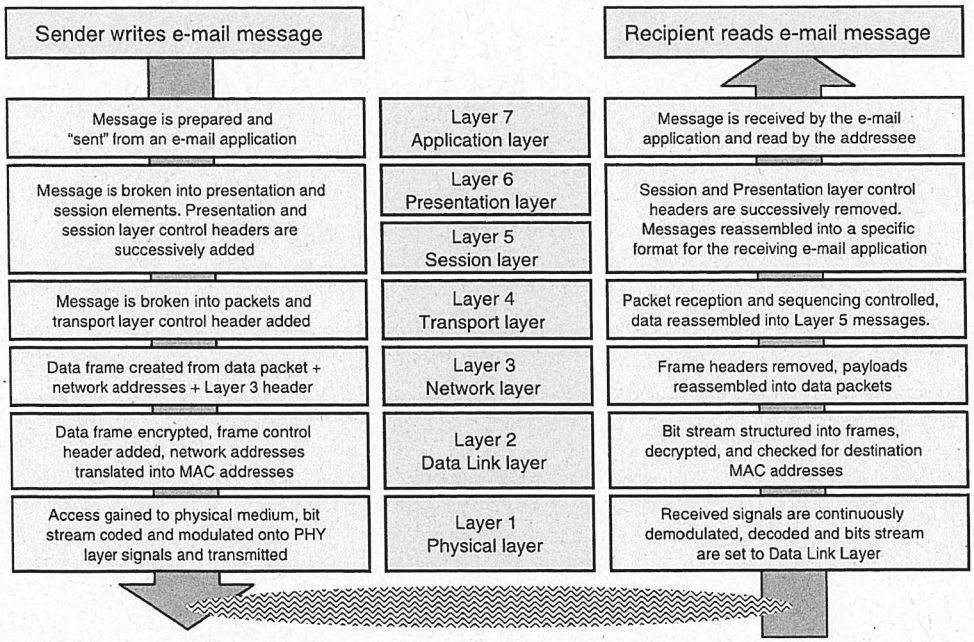

Figure 2-1: The OSI Model in Practice — an E-mail Example

the OSI model. The following sections will give a preliminary introduction to these standards and protocols, while more detailed descriptions will be found in Parts III to V where Local Area (LAN), Personal Area (PAN) and Metropolitan Area (MAN) wireless networking technologies are described respectively.

The next section starts this introductory sketch one layer higher — at the Network layer — not because this layer is specific to wireless networking, but because of the fundamental importance of its addressing and routing functions and of the underlying Internet Protocol (IP).

Network Layer Technologies

The Internet Protocol (IP) is responsible for addressing and routing each data packet within a session or connection set up under the control of transport layer protocols such as TCP or UDP (see Glossary). The heart of the Internet Protocol is the IP address, a 32-bit number that is attached to each data packet and is used by routing software in the network or Internet to establish the source and destination of each packet.

While IP addresses, which are defined at the Network layer, link the billions of devices connected to the Internet into a single virtual network, the actual transmission of data frames between devices relies on the MAC addresses of the network interface cards (NICs), rather than the logical IP addresses of each NIC's host device. Translation between the Layer 3 IP address and the Layer 2 MAC address is achieved using Address Resolution Protocol (ARP), which is described in the Section "Address Resolution Protocol, p. 16".

IP Addressing

The 32-bit IP address is usually presented in "dot decimal" format as a series of four decimal numbers between 0 and 255, for example; 200.100.50.10. This could be expanded in full binary format as 11001000.01100100.00110010.00001010.

As well as identifying a computer or other networked device, the IP address also uniquely identifies the network that the device is connected to. These two parts of the IP address are known as the host ID and the network ID. The network ID is important because it allows a device

used. 10 Base-T is the standard Ethernet, running at 10 Mbps and using an unshielded twisted-pair copper wire (UTP), with a maximum distance of 500 metres between a device and the nearest hub or repeater.

The constant demand for increasing network speed has meant that faster varieties of Ethernet have been progressively developed. 100 Base-T, or Fast Ethernet operates at 100 Mbps and is compatible with 10 Base-T standard Ethernet as it uses the same twisted-pair cabling and CSMA/CD method. The trade-off is with distance between repeaters, a maximum of 205 metres being achievable for 100 Base-T. Fast Ethernet can also use other types of wiring — 100 Base-TX, which is a higher-grade twisted-pair, or 100 Base-FX, which is a two strand fibre-optic cable. Faster speeds to 1 Gbps or 10 Gbps are also available.

The PMD sub-layer is specified separately from the Ethernet standard, and for UTP cabling this is based on the Twisted Pair-Physical Medium Dependent (TP-PMD) specification developed by the ANSI X3T9.5 committee.

The same frame formats and CSMA/CD technology are used in 100 Base-T as in standard 10 Base-T Ethernet, and the 10-fold increase in speed is achieved by increasing the clock speed from 10 MHz to 125 MHz, and reducing the interval between transmitted frames, known as the Inter-Packet Gap (IPG), from 9.6 µs to 0.96 µs. A 125 MHz clock speed is required to deliver a 100 Mbps effective data rate because of the 4B/5B encoding described below.

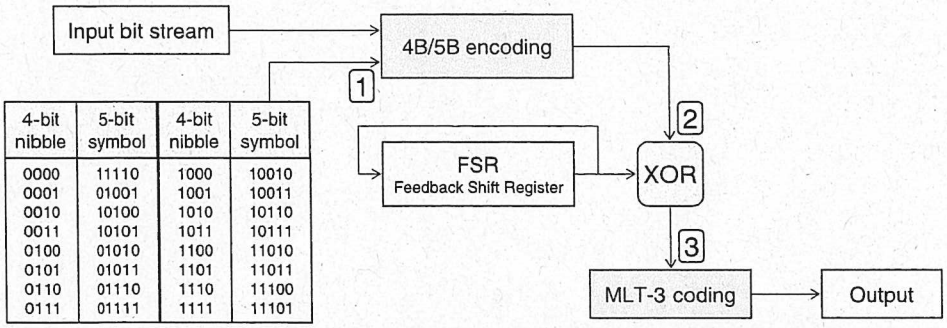

4-bit nibble	5-bit symbol	4-bit nibble	5-bit symbol
0000	11110	1000	10010
0001	01001	1001	10011
0010	10100	1010	10110
0011	10101	1011	10111
0100	01010	1100	11010
0101	01011	1101	11011
0110	01110	1110	11100
0111	01111	1111	11101

Figure 2-7: 100 Base-T Ethernet Data Encoding Scheme

To overcome the inherent low-pass nature of the UTP physical medium, and to ensure that the level of RF emissions above 30 MHz comply with FCC regulations, the 100 Base-T data encoding scheme was designed to bring the peak power in the transmitted data signal down to 31.25 MHz (close to the FCC limit) and to reduce the power in high frequency harmonics at 62.5 MHz, 125 MHz and above.

4B/5B encoding is the first step in the encoding scheme (Figure 2-7). Each 4-bit nibble of input data has a 5th bit added to ensure there are sufficient transitions in the transmitted bit stream to allow the receiver to synchronise for reliable decoding. In the second step an 11-bit Feedback Shift Register (FSR) produces a repeating pseudo-random bit pattern which is XOR'd with the 4B/5B output data stream. The effect of this pseudo-randomisation is to minimise high frequency harmonics in the final transmitted data signal. The same pseudo-random bit stream is used to recover the input data in a second XOR operation at the receiver.

The final step uses an encoding method called Multi-Level Transition 3 (MLT-3) to shape the transmitted waveform in such a way that the centre frequency of the signal is reduced from 125 MHz to 31.25 MHz.

MLT-3 is based on the repeating pattern 1, 0, –1, 0. As shown in Figure 2-8, an input 1-bit causes the output to transition to the next bit in the pattern while an input 0-bit causes no transition, i.e. the output level remaining unchanged. Compared to the Manchester Phase Encoding (MPE) scheme used in 10 Base-T Ethernet, the cycle length of the output signal is reduced by a factor of 4, giving a signal peak at 31.25 MHz instead of 125 MHz. On the physical UTP medium, the 1, 0 and –1 signal levels are represented by line voltages of +0.85, 0.0 and –0.85 Volts.

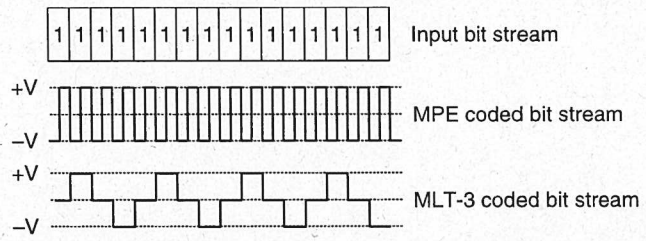

Figure 2-8: Ethernet MPE and Fast Ethernet MLT-3 Encoding

Private IP Addresses

In February 1996, the Network Working Group requested industry comments on RFC 1918, which proposed three sets of so-called private IP addresses (Table 2-3) for use within networks that did not require Internet connectivity. These private addresses were intended to conserve IP address space by enabling many organisations to reuse the same sets of addresses within their private networks. In this situation it did not matter that a computer had an IP address that was not globally unique, provided that that computer did not need to communicate via the Internet.

Table 2-3: Private IP Address Ranges

Class	*Private address range start*	*Private address range end*
A	10.0.0.0	10.255.255.255
B	172.16.0.0	172.31.255.255
C	192.168.0.0	192.168.255.255

Subsequently, the Internet Assigned Numbers Authority (IANA) reserved addresses 169.254.0.0 to 169.254.255.255 for use in Automatic Private IP Addressing (APIPA). If a computer has its TCP/IP configured to obtain an IP address automatically from a DHCP server, but is unable to locate such a server, then the operating system will automatically assign a private IP address from within this range, enabling the computer to communicate within the private network.

Internet Protocol Version 6 (IPv6)

With 32 bits, a total of 2^{32} or 4.29 billion IP addresses are possible — more than enough one would think for all the computers that the human population could possibly want to interconnect.

However, the famous statements that the world demand for computers would not exceed five machines, probably incorrectly attributed to Tom Watson Sr., chairman of IBM in 1943, or the statement of Ken Olsen, founder of Digital Equipment Corporation (DEC), to the 1977 World Future Society convention that "there is no reason for any individual to have a computer in his home", remind us how difficult it is to predict the growth and diversity of computer applications and usage.

match in the table, it forwards the packet to the address associated with that table entry, which may be the address of another network or of a "next-hop" router that will pass the packet along towards its final destination.

If the router can't find a match, it goes through the table again looking at just the network ID part of the address (extracted using the subnet mask as described above). If a match is found, the packet is sent to the associated address or, if not, the router looks for a default next-hop address and sends the packet there. As a final resort, if no default address is set, the router returns a "Host Unreachable" or "Network Unreachable" message to the sending IP address. When this message is received it usually means that somewhere along the line a router has failed.

What happens if, or when, this elegantly simple structure breaks down? Are there packets out there hopping forever around the Internet, passing from router to router and never finding their destination? The IP header includes a control field that prevents this from happening. The time-to-live (TTL) field is initialised by the sender to a certain value, usually 64, and reduced by one each time the packet passes through a router. When TTL get down to zero, the packet is discarded and the sender is notified using an Internet Control Message Protocol (ICMP) "time-out" message.

Building Router Tables

The clever part of a router's job is building its routing table. For simple networks a static table loaded from a start-up file is adequate but, more generally, Dynamic Routing enables tables to be built up by routers sending and receiving broadcast messages.

These can be either ICMP Router Solicitation and Router Advertisement messages which allow neighbouring routers to ask "Who's there?" and respond "I'm here", or more useful RIP (Router Information Protocol) messages, in which a router periodically broadcasts its complete router table onto the network.

Other RIP and ICMP messages allow routers to discover the shortest path to an address, to update their tables if another router spots an inefficient routing and to periodically update routes in response to network availability and traffic conditions.