years. However, the creation of faster workstations, combined with advances in compiler technology, eventually made LISP machines obsolete. Interestingly, the system code for those later machines was also written in LISP. For example, the design philosophy of the Japanese Fifth Generation Project (1981–91) owed a great deal to LISP machines.

The very success of LISP led to several incompatible versions for different computers. Therefore, at a 1981 meeting sponsored by the Advanced Research Projects Administration (**ARPA**), several of the main LISP developers met and formed a committee to explore the possibility of agreeing on a common language. The definition of Common LISP was published in 1984 in a book titled *Common Lisp: The Language*. Several software companies implemented this new version immediately. To this day, LISP continues to be the favorite programming language of AI researchers.

FURTHER READING

McCarthy, John. *LISP 1.5 Programmer's Manual.* Cambridge, Mass.: Massachusetts Institute of Technology, Computation Center and Research Laboratory of Electronics, 1962; 2nd ed., 1973.

Steele, Guy L. *Common LISP: The Language.* Burlington, Mass.: Digital Press, 1984; 2nd ed., Bedford, Mass., 1990.

—*Raúl Rojas*

Local Area Network

A local area network (LAN) interconnects computers in a geographic area of limited size, usually a building or one floor in a building. LANs represent the lowest level in the global connection hierarchy. Their "width" is measured in hundreds of meters, whereas the width (reach) of **metropolitan area networks** (MANs) is measured in kilometers, the width of **wide area networks** (WANs) in hundreds of kilometers, and of the **Internet** in thousands of kilometers. LANs were conceived in the 1970s and became ubiquitous with the success of the **Ethernet** networking cards. LANs are used to connect personal computers and workstations, and to share printers and **file servers**.

The main aspects distinguishing LANs from other types of networks are their topology and the protocols and connection media they use. A popular LAN topology consists of using a single **bus** to connect all computers in the network. The bus is used to broadcast messages to all machines, marking each message with the number of the computer that is the intended recipient. A bus topology is simple and inexpensive, since only a cable has to be laid out in the building. In some cases, the preexisting **twisted-pair** cables for telephones can be used to transmit data. Sending a message can lead to collisions if two computers try to send packets at the same time. Therefore, before sending a packet every computer monitors the line until it is free. If there is a collision even after having done this, both computers back off and wait some time to restart the transmission. This method of collision detection is the one implemented in the **Ethernet** standard.

The first Ethernet networks were installed as part of a research project started by **Robert Metcalfe** (1946–) and David Boggs (1950–) at **Xerox Palo Alto Research Center** (PARC) in the early 1970s. On 22 May 1973, Metcalfe used the term *Ethernet* for the first time, superseding the name *Aloha Alto* that had been used until then. Although that particular date has been called the birthday of Ethernet, Metcalfe has observed that the entire design actually took several years to converge to its final form. The motivation for the development of the Ethernet was to allow entire work groups to share the laser printers that Xerox was developing.

Ethernet is the paradigmatic LAN: simple to install, inexpensive, and as fast as a single user may wish. Ethernet cards cost several thousand dollars when they were first announced, but they are now commodity products packaged in single chips for personal computers.

LANs in organizations are now more sophisticated than the single Ethernet cable mentioned above. Most LANs now have a star topology. Individual machines in a room, for example, are connected to a single cable. This in turn is connected to a switch, which receives all packets from such bus subnetworks (called *segments*). The switch passes packets from one segment to another at full speed, and all packets going out of this area go through a faster line (e.g., a gigabit link) connected to another switch or router. This type of topology is called *switched Ethernet*.

Another alternative LAN topology is the **token ring** network, in which all computers are connected in a ring. Information is passed through the ring, from one computer to the other, but only one computer can send at any time. A special packet, called the *token*, is circulating in the ring until one computer grabs it and starts transmitting. When the machine has finished, it releases the token, allowing other machines to send information. Since in Ethernet and token ring networks there is no bus master allocating use of the transmission medium, they are called *decentralized broadcast networks.*

In both Ethernet and token ring networks, the transmission protocol makes special assumptions about some characteristics of the network (e.g., the maximum packet size or the maximum time that a machine can transmit). In Ethernet networks the maximum delay for a packet to go from one machine to another is bounded because the total length of the Ethernet bus is also limited. These bounds allow LANs to implement protocols that could not be used in heterogeneous networks (such as the Internet), in which a message can take an arbitrary time to arrive at its destination.

Another relevant aspect of LANs is that of the protocols used for transmission. Network protocols are built in layers: one layer of software gives orders only to the layer below, until the physical medium is reached (i.e., the physical layer of the network). The Internet, for example, lets one computer communicate with another by sending **TCP/IP** packets. In the Internet each computer has an identifier that is used to route messages from one computer to the next until the packet reaches its destination. But the TCP/IP packet is itself encapsulated in another protocol for the underlying kind of network. If we have an Ethernet network, the IP packets will be broken into Ethernet packets and they will be sent to the intended destination, where the Ethernet packets will be opened and the original TCP/IP packet will be reconstructed. One may think of network protocols as envelopes: Each time a packet is sent, the layer below puts a new envelope around the information, writes how it should be forwarded, and delivers it to the next layer. This layer in turn puts a new envelope and sends it to the next layer, and so on. Some times the packet does not fit into the envelope of the lower layer and is then cut into pieces that fit into several envelopes. The envelopes are numbered and sent. The recipient collects the envelopes, glues the several pieces together, and reconstructs the packet.

The type of medium used in LANs is also different from that in other types of networks. WANs and MANs are based on **optical fiber**, microwaves, or satellites. LANs are based primarily on twisted-pair and coaxial cabling, although optical cabling is also available at a premium price. The first Ethernet networks used coaxial cable. The cable was connected to one input port of the Ethernet card, and the cable continued through the output port. This made it easy for anybody to disrupt the network intentionally or accidentally just by pulling one of the cables out from the computer. The bus was then severed at one point and the network broke down. Now, computers are connected to the network through transceivers on the wall, which close the bus unless a connector is plugged in. Whereas before, there were two cables from the computer to the transceiver, there is now only a single cable.

Another important LAN technology is **AppleTalk**, introduced by **Apple** in 1984 and originally called the Apple Bus. Apple's Macintosh Office, marketed in that year, consisted of **Macintosh** computers connected to a laser printer through a serial cable. The raw speed of the network was rather limited (100 times lower than Ethernet), but it was very simple to set up. Each Macintosh could connect to the network through the printer port, and a chain would be formed with all machines and the printer. The network could support 16 devices over a maximum total length of 30 meters. The AppleTalk protocol was later ported to Ethernet networks, but the original Macintosh Office was probably the cheapest network technology available in the 1980s: For the end user, the most expensive part of the network was the cable.

Other technologies available for networking in the 1970s were proprietary systems from **Digital Equipment** (DecNet) and from **IBM**. These networks were designed to work only with equipment from Digital Equipment or IBM, respectively. Ethernet was in this sense also revolutionary: It was an open standard that could be adopted by any manufacturer of computers because **Xerox**, of course, was interested in selling laser printers, whereas DEC and IBM wanted to sell computers.

In the future an even lower level in the networking hierarchy is anticipated. *Home LANs* will network all devices together in a household, and *body networks* will network all devices carried by an individual. LANs of the future will therefore be independent of special cabling and will have to rely on wireless transmission. This would also allow the elimination of the many cables between devices such as the **mouse**, **keyboard**, **printer**, and the main computer.

FURTHER READING
Geier, James T. *Wireless LANs: Implementing Interoperable Networks.* Indianapolis, Ind.: Macmillan Technical, 1999.
Stallings, William, ed. *Advances in Local and Metropolitan Area Networks.* Los Alamitos, Calif.: IEEE Computer Society Press, 1994.
Tanenbaum, Andrew S. *Computer Networks.* Englewood Cliffs, N.J.: Prentice Hall, 1981; 3rd ed., Upper Saddle River, N.J., 1996.

—*Raúl Rojas*

Logic Gates

Logic gates are elementary components used to evaluate logical formulas in the **hardware** of computers. Modern computer chips are built from **integrated circuits** containing thousands of gates. **Jack Kilby** (1923–) of **Texas Instruments** developed the **integrated circuit** in 1958 and won the Nobel Prize for physics for this work in 2000.

Binary logic gates work, as the name suggests, with binary logic, which means that there are only two types of signals along the wires, "0" and "1." The voltage on the wire represents the truth value of the formula, with 5 volts (V), for example, as the number 1, and 0V for the number 0.

Many scientists contributed to the idea that hardware can be used to implement logic. **George Boole** (1815–64) used mathematics as a model for logic in 1854. He was interested in finding the smallest possible algebra, so he used only two numbers, 0 for false, and 1 for true. The algebra presented by Boole in ***An Investigation of the Laws of Thought*** substituted the logical AND for multiplication and the logical exclusive OR for addition. This was a pivotal insight because it allowed logic, which until then had been based on rhetoric arguments, to be analyzed using a formal mathematical system.

A necessary step was developing a standard way for philosophers and logicians to write logical formulas. In 1879, Gottlob Frege (1848–1925) developed a method for writing predicate logic formulas in his work *Begriffsschrift.* His notation was dependent on drawing lines connecting parts of the logic formula together. Preferring a text-based approach, Ernst Schroeder (1841–1902) popularized the 1883 algebraic notation of Charles Sanders Pierce (1839–1914) by using it in *Vorlesungen uber die Algebra der Logik* (Lectures on the Algebra of Logic). Giuseppe Peano (1858–1932), wanting to use logic and mathematical symbols in the same formulas, created the modern logic notation by turning letters backward and upside down to form logical symbols. The letter E was used for "there exists" and became ∃, the letter C, for consequence, was used for implication and became ⊃. The German word *alle* contributed the letter A and the symbol ∀. Bertrand Russell (1872–1970) finished this process by using the complete notation in *Principia Mathematica*, of 1910, which formed a basis for most logicians' work of the twentieth century.

A binary logic gate operates on two bits (each with value 1 or 0) and produces a single output bit. There are a total of 16 possible logical operations of two arguments. John Sowa (1940–) names these as shown in the table. The four combinations of two 1-bit arguments are shown in four columns. The output of the AND operation, for example, is only 1 when both inputs are 1. C. S. Pierce analyzed these 16 operations in 1880, organizing logical operators into truth tables and proving that all of them could be calculated using combinations of the single NAND operation. Just as the development of interchangeable parts revolutionized manufacturing, this discovery is a fundamental principle of modern logical circuit design. Every binary circuit that has two inputs and one output is an implementation of one of these operations, and can be constructed entirely of NAND gates and connections. More complex operations, with more binary inputs or outputs, can also be reduced to NAND gates.

Digital logic designers generally use the AND, OR, and NOT operations when designing circuits. An AND