and diesel engines. Gasoline engines require a volatile fuel, so that a flammable homogeneous mixture can be formed prior to combustion. The flame is initiated at the spark plug and propagates steadily into the compressed unburned mixture. With a diesel engine the fuel is injected into the cylinder just before the end of the compression stroke. If the compression ratio is high enough for the fuel to self-ignite, then it will also be hot enough to evaporate the fuel. Diesel fuel can thus be much less volatile, but it does need to self-ignite readily, and in general, large molecules will self-ignite more readily and of course be less volatile. The diesel fuel droplets will evaporate and mix with air, and some of this mixture will be flammable, and within about 1 millisecond of injection the fuel will ignite. The remaining fuel will continue to evaporate and burn, and this provides controlled combustion. In the gasoline engine, self-ignition of the fuel leads to very rapid combustion and a phenomenon known as "knock" or "detonation." As the flame propagates in a gasoline engine, the pressure rises, so the unburned mixture is compressed to a higher temperature, and if the mixture self-ignites, the combustion is very rapid since all of the unburned mixture is flammable and at the point of self-igniting. The rapid pressure rise causes overheating of the engine and failure of the pistons. The earliest study of this was probably by Bertram Hopkinson (1874–1918) at Cambridge University in 1906, who had developed an optical system for recording pressure. He was assisted by Harry Ricardo (later Sir Harry Ricardo, 1885–1974) who was then a student at Cambridge.

Initially, the fuel made by distillation of crude oil or shale was intended for use in lamps, and legislation limited its volatility, so that the more volatile hydrocarbons (later to be an important component of gasoline) were of little use initially. Distillation was in simple pot stills, and the lightest fraction was known as naphtha, with lamp oil or kerosene being less volatile. The remaining fractions were the basis of fuel oil (for diesel engines and heating) and lubricants. Early gasoline engines used wick carburetors, and these needed a volatile fuel that would evaporate without leaving a residue. In the United Kingdom, this led to a specification for "Standard Petrol" that was defined in terms of its specific gravity (0.680 to 0.685). This gasoline would have had a boiling point range of about 50–120°Celsius.

As gasoline quality improved, a higher compression ratio could be used, and this increased efficiency and engine output. As lubricants and materials improved, engine speeds could be increased, and this too led to an increase in power. With less effective air utilization and lower speed range, diesel engines had an output about half that of a gasoline engine. The output of diesel engines was increased by turbocharging, a process by which an exhaust gas turbine drives a compressor that is used to increase the density of the air entering the engine.

The largest changes since the 1970s have been a consequence of emissions legislation to combat air pollution and photochemical smog. In the case of gasoline engines the most significant development has been the three-way catalyst. By operating the engine with a stoichiometric air fuel ratio (the ratio for theoretically complete combustion), the catalyst could simultaneously reduce the nitrogen oxides (NO and NO_2 known collectively as NOx) and oxidize the carbon monoxide (CO) and unburned hydrocarbons (HC). Such control of the air fuel ratio required electronically controlled fuel injection and feedback from an exhaust gas oxygen sensor (lambda sensor). For diesel engines oxidation catalysts are used to reduce HC and CO emissions, but the most visible change has been the elimination of exhaust smoke by the use of diesel particulate filters (traps). Fuel injection and many other functions are now controlled electronically so as to ensure compliance with emissions legislation and then minimize fuel consumption. A four-cylinder, two-liter engine can have a maximum efficiency of about 35 percent, while a similar sized diesel engine will achieve 40 percent. As engine cylinder size increases, heat transfer and other losses reduce, and marine diesel engines, with a cylinder diameter of about one meter, can have efficiencies of over 50 percent.

SEE ALSO *Airplane; Automobile; Electric Motor; Jet Engine; Steam Engine.*

BIBLIOGRAPHY

Clerk, Sir Dugald. *The Gas Engine*. London: Longmans, Green, 1886.

Daniels, Jeff. *Driving Force—The Evolution of the Car Engine*. Yeovil, UK: Haynes, 2002.

Hancock, E. G., ed. *The Technology of Gasoline*. Oxford: Blackwell Scientific, 1985.

Heywood, John B. *Internal Combustion Engine Fundamentals*. New York: McGraw-Hill, 1988.

Reynolds, John. *Engines & Enterprise—The Life and Work of Sir Harry Ricardo*. Thrup, Stroud, UK: Sutton, 1999.

Robert Bosch GmbH. *Bosch Automotive Handbook*. 7th ed. Hoboken, NJ: Wiley, 2007.

Stone, Richard. *Introduction to Internal Combustion Engines*. 4th ed. London: Palgrave Macmillan, 2012.

Wilkinson, P. H. *Aircraft Diesels*. New York: Pitman, 1940.

Richard Stone
Professor of Engineering Science
University of Oxford

INTERNET

In mid-2012, the Internet had over 2.4 billion users, and in 2013 consumers purchased over $1 trillion worth of goods via the Internet. Yet only fifty years ago nothing remotely like the Internet existed, nor was such a high speed and pervasive system of computers and communications facilities even projected as science fiction. How did it come about?

On October 4, 1957, the Soviet Union launched *Sputnik 1*, the first artificial Earth satellite. In the United States, there was near panic that the Soviet Union was so far ahead in technology that the country might lose a war if the Soviets started one. In reaction, President Dwight D. Eisenhower (1890–1969) created a new office in the Department of Defense (DOD) called the Advanced Research Projects Agency (ARPA) to stimulate US technology research.

The American psychologist and computer scientist Joseph Carl Robnett "Lick" Licklider (1915–1990), a veteran of both Harvard University and the Massachusetts Institute of Technology (MIT), became increasingly interested in the interactive use of computers. In 1960 Licklider crystallized his ideas about interactive computing in a paper titled "Man-Computer Symbiosis." In this relationship, people would set goals, formulate hypotheses, and evaluate results. Computers would retrieve information, execute simulations, and display results in forms that people could easily understand. It has been said that the paper "essentially laid out the vision and the agenda that would animate US computer research for most of the next quarter-century" (Waldrop 2001, p. 176).

In 1962 ARPA recruited Licklider to head a new group, the Information Processing Techniques Office (IPTO). Licklider immediately directed funding to the research that was needed to make his vision of human-computer symbiosis a reality. Within a few years, IPTO was funding much of the cutting-edge computer research being done in the United States. IPTO accomplished this mostly by supporting bright university researchers with essentially open-ended contracts. The goal was to develop centers of expertise in advanced computer research at many major US universities. It is reasonable to say that this IPTO funding played a major role in creating academic computer science departments across the country and establishing the United States as the world leader in computer research.

The American psychologist, mathematician, and engineer Robert W. "Bob" Taylor (1932–) became the third IPTO director in 1966. By that time, IPTO was supporting many large and expensive computer centers around the country, each with different computers and unique software systems, focused on distinct aspects of research. Taylor soon saw that every time a researcher proposed a new, unique project, the proposal started with a request for a new computer. He recognized that most of the projects could be carried out on one of the existing IPTO-funded systems if only the new researcher had adequate access to it. However, then-current methods of remote access were inadequate and expensive. Taylor needed to find a way to connect existing computers with a high-performance network that would allow a new remote user the same level of service a local user would have. Late in 1966 Taylor recruited Lawrence G. "Larry" Roberts (1937–), an American electrical engineer, to join IPTO specifically to solve this problem. Roberts had been working at MIT's Lincoln Laboratory on computer graphics and data communication problems, where he acquired "a reputation for being something of a genius" (Hafner and Lyon 1996, p. 45).

NETWORK DESIGN CONCEPTS

Roberts quickly immersed himself in the project. At a meeting of the leaders of the groups IPTO was funding (who were known as principal investigators, or PIs) in early 1967, Roberts discussed his tentative plans; they were mostly greeted with hostility because the plans looked like a large implementation cost would be required for each PI. However, the American computer scientist Wesley A. Clark (1927–) came up with the idea that the communication functions should be performed by small computers developed and maintained by someone else. The research computer would only have to be programmed to interface with its local communication processor, a seemingly much smaller (and less costly) job. Roberts adopted this idea, calling the communication computers Interface Message Processors (IMPs). This decision dispelled some of the antipathy among the PIs for the network idea.

Another effect of the *Sputnik* launch in 1957 was a new vision of warfare of the future. The US military imagined a war in which the Soviets and the Americans launched waves of missiles with nuclear warheads at each other. How would the military decision makers be able to

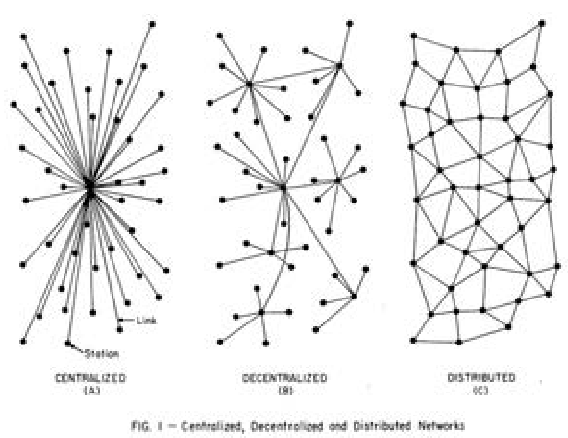

Figure 1. *Network types.* **BARAN, PAUL. *ON DISTRIBUTED COMMUNICATIONS.* RAND CORPORATION, 1964, P.2. REPRINTED WITH PERMISSION.**

command and control countermeasures without a communication system that could survive a first strike? In the early 1960s the US Air Force contracted with the RAND Corporation (RAND is an acronym for *R*esearch *AN*d *Development),* headquartered in Santa Monica, California, to come up with an answer. In 1964 RAND published a report, "On Distributed Communications," in eleven unclassified volumes; the primary author was the American engineer Paul Baran (1926–2011). Figure 1 shows a set of systems connected by communication links in three different ways. Existing computer networks were centralized; if the central point failed, all communication would be disrupted. The telephone network was decentralized; communication in this structure might survive one failure, but it was vulnerable to a small number of failures. Baran proposed that a distributed network was required for communication to survive many failures. Data would move through the network in message blocks, each taking a potentially different path. The US Air Force saw great merit in Baran's ideas but was unable to convince either the American Telephone and Telegraph Company (AT&T) or the DOD's Defense Communications Agency (DCA; now known as the Defense Information Systems Agency) to build it.

Meanwhile, Donald Watts Davies (1924–2000), a Welsh physicist and computer system researcher at Britain's National Physical Laboratory (NPL), was devising a way to make the connection of remote terminals to time-sharing computers less costly. It is typical of such use that data are transmitted in short bursts interspersed with long periods of inactivity. Davies proposed that a high-capacity circuit be shared among many users. A user's burst of information would be bundled into a packet along with the network address to which it should be delivered. When Davies described this idea in an NPL lecture titled "The Future Digital Communication Network" in 1966, one of his listeners told him about Baran's work. Davies looked into Baran's ideas for handling message blocks and adopted some of them. He then convinced the NPL that developing a packet-switching system would be a good investment in Britain's economic growth, but he obtained funding for only a single switch. Work on the project started in 1967.

Roberts presented his preliminary design for the functions an IMP would perform in a paper titled "Multiple-Computer Networks and Intercomputer Communication" in October 1967. Roberts's paper was clear on what the IMPs would do but sketchy about what the network would look like. He suggested it would use leased telephone circuits running at 9,600 bits per second (bps). At the same conference was Roger A. Scantlebury, the technical leader of the team at NPL, who described their packet-switching project. After the presentations, Roberts and Scantlebury got together to discuss these very similar concepts. Roberts learned from Scantlebury about Baran's ideas for making a network reliable by making it distributed. Scantlebury urged Roberts to use a smaller number of faster circuits in his network design, thereby decreasing delay and increasing capacity—the NPL team was thinking about data rates over 1 megabit per second (Mbps).

Returning to Washington, DC, Roberts read Baran's reports and thought about Davies's ideas. Roberts decided the network he would build would use 50,000-bps circuits to connect the IMPs, and that for reliability there would be multiple paths between every pair of IMPs. However, his reason was not to survive a nuclear attack; in the 1960s computer and circuit technology were such that on average, every IMP and every circuit would be out of service a few times each week. Furthermore, the network specification would require round-trip time between every pair of IMPs to be less than a half second, and sustained throughput between any pair of IMPs to be at least 20,000 bps. Roberts turned to the American engineer and computer scientist Leonard "Len" Kleinrock (1934–), an expert in queuing theory applied to communications, for assurance that his system could meet these objectives. In June 1968 ARPA gave Roberts authorization to build the network and, with Kleinrock's assurance, IPTO issued a request for quotation (RFQ) at the end of July.

THE ARPANET

In December 1968 ARPA awarded the contract to Bolt, Beranek and Newman, Inc. (BBN) of Cambridge, Massachusetts. The contract called for delivery of four IMPs, the first to the University of California at Los Angeles (UCLA) on September 1, 1969, and one each month thereafter. If the initial system worked as specified, ARPA could exercise options for up to fifteen more IMPs during 1970 and 1971. Frank Heart led the BBN proposal team, and Robert "Bob" Kahn supplied the team's expertise in communication theory. Two key concepts in the BBN proposal were (1) the use of dynamic adaptive routing in the network and (2) the incorporation of extensive reporting and diagnostic software in each IMP so it could be monitored and modified via the network. Roberts also gave Kleinrock (at UCLA) a contract to analyze the network's performance.

At about the time IPTO issued its RFQ, a group of researchers from the locations where IMPs were to be installed began meeting to discuss how their computers, called *Hosts*, would use the network. These researchers realized that they would need a set of rules to interpret the data they sent and received, and they called these rules *protocols.* Following the general structure used in computer software, they agreed to design the protocols in layers—each layer providing a set of services to the layer above. The protocol design group called itself the Network Working Group (NWG). Its chair was Steven D. ''Steve'' Crocker (1944–), a graduate student at UCLA. BBN assured the NWG that the IMPs would do all the work necessary to ensure that messages were delivered in the same order they were sent and with no errors, so the Host protocols did not need to be concerned with those functions. The first protocol defined by the NWG dealt with establishing a connection between processes in two computers and regulating the flow of data over the connection. A character-protocol for exchanging text (for example, between a terminal and a time-shared operating system) soon followed. Electronic mail (email) was introduced in early 1972 by the American computer programmer Raymond Samuel ''Ray'' Tomlinson, and it was an immediate success.

The contract with BBN was extended to provide for expansion of the network. By August 1972 there were 29 IMPs, as shown in Figure 2. (Note: Terminal interface processors [TIPs] were IMPs that also allowed the direct attachment of terminals.) Roberts decided it was time for a public demonstration of the network, by then called ARPANET, or Advanced Research Projects Agency Network. A TIP was installed at the Washington (DC) Hilton Hotel during the first International Conference on Computer Communication (ICCC), October 24–26, 1972. Conference attendees were invited to sit at terminals connected to the TIP and access programs running on ARPANET Hosts. It was a great success and convinced many skeptics that packet switching was a viable new technology. A movie, *Computer Networks—Heralds of Resource Sharing*, was commissioned to be shown in an adjoining room.

CONNECTING INDEPENDENT NETWORKS

By the time of the ICCC, a number of experimental packet-switching networks were planned, and representatives of most of them were at the conference. Roberts recognized that a logical next step was to connect these independent networks together, and that would require some standardization of Host protocols. He invited the network representatives to meet in the evening and begin discussing how to do this. On hand were people from ARPANET, NPL, a European common market network, the British telephone company network, the French telephone company network, a French research network (Cyclades, which had started in 1972), and the Norwegian telephone company (which had built an experimental network in 1971). The group agreed to design a common Host protocol. The group called itself the International Network Working Group (INWG), and the American computer scientist Vinton Gray ''Vint'' Cerf (1943–) volunteered to be temporary chair/secretary. The French computer scientist Louis Pouzin (1931–) from the Cyclades project argued forcefully that the Host protocol should assume no services from the underlying network other than delivery of packets. It should be up to the users to detect and correct errors, since the network could never be completely successful in doing so. With this approach it should be simple to connect disparate networks with simple gateways interfaced to the separate networks and forwarding packets from one to the other.

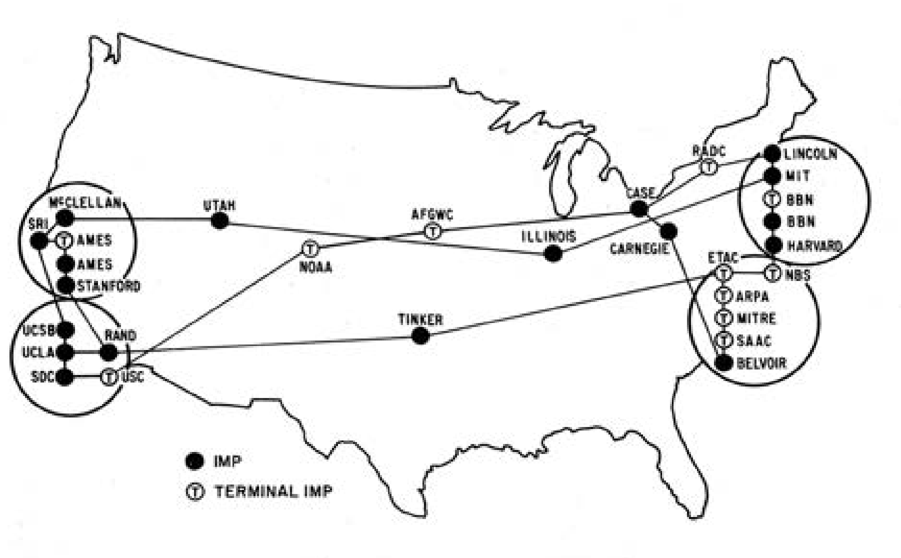

Figure 2. *The ARPANET in August 1972.* COURTESY RAYTHEON BBN TECHNOLOGIES

INWG met again in June 1973 and produced a draft specification of a Host protocol based on the ideas from the previous meeting. The document described a Transmission Control Program (TCP) in each Host and defined interactions between TCPs. Cerf served as the editor of the document, which was circulated for comment. About a month later, Cerf and Bob Kahn produced a refinement of that document, which incorporated ideas from some people (including Kahn) who were not at the INWG meeting. A further edited version of this paper was published as "A Protocol for Packet Network Interconnection" in May 1974. At that time, Cerf was on the Stanford University (California) faculty, and Kahn was in charge of networking at IPTO.

By this time, IPTO was funding research in packet radio networks and packet satellite networks and wanted to connect them to ARPANET. Kahn began funding the development of gateways based on TCP, along with funding the development of TCPs for some of the ARPANET Hosts. On November 22, 1977, there was a demonstration of the combined operation of the three types of networks, interconnected by gateways and using TCP. (Along the way, the P in TCP changed from *Program* to *Protocol,* and the protocol had been split into two layers called TCP and IP [Internet protocol].) In historical terms this demonstration marked the birth of the Internet.

It was clear that since packet-switching networks could be interconnected, the ARPANET need not be the only game in town. In 1975 IPTO had declared ARPANET to be an operational system, and its management had been transferred from ARPA to DCA. After the demo, DCA declared that all ARPANET Hosts must convert to TCP no later than January 1, 1983. There were some problems, but by June 1983 the conversion was complete. A few years later, what had been the ARPANET was split into two separate networks—computer research sites were assigned to one network whose name remained ARPANET, and military users were assigned to a network named MILNET (military network). Each IMP was assigned to one of the two networks, and the two networks were then connected by gateways. At this time the two networks contained over one hundred IMPs.

In June 1984, 141 separate networks had been assigned addresses. Each of these was (or was scheduled to be) connected to the ARPANET/MILNET. With so many networks, a principal function of the gateways was to calculate routes to each network, originating the term *router.* Most of these networks were local area networks (LANs) and most LANs were Ethernets (one of several networking technologies used for LANs). Ethernet was invented by the American electrical engineer Robert M. "Bob" Metcalfe (1946–), who implemented his ideas at Xerox PARC (for Palo Alto Research Center) in 1974. Metcalfe left Xerox to form 3Com in 1979. In 1981 3Com shipped its first Ethernet transceiver and that year began selling adapters for Digital Equipment Corporation's (DEC's) Programmable Data Processor 11 (PDP-11) and VAX (originally the acronym for virtual address extension), and for Sun computers. An Ethernet adapter for the International Business Machines Corporation (IBM) personal computer (PC) followed the next year, and by 1985 over 100,000 had been sold.

NSF AND COMMERCIAL USE

In 1984 the US National Science Foundation (NSF) established the Office of Advanced Scientific Computing to create six supercomputer centers and to make these supercomputers easily available to researchers throughout the United States. The NSF planned to do this by encouraging the establishment of regional networks and to build a high-speed backbone to link these with each other and with the supercomputer centers. The NSF did not intend to build the regional networks itself, but to offer grants of seed money to consortia of universities, which would build the regionals and operate them on a self-sustaining basis. The backbone, called NSFNET (short for NSF Network), would use circuits operating at 1.5 Mbps, thirty times as fast as the ARPANET backbone. In 1987 NSF awarded a five-year contract to a consortium of a Michigan nonprofit (MERIT) for management, IBM for packet switches, and MCI Communications (MCI) for circuits.

In 1986 the American electrical engineer and technology researcher Stephen Wolff accepted the job of division director for Networking and Communications Research and Infrastructure at NSF. Wolff realized that as long as the government owned the networks that formed the backbone of the US Internet, federal policy would prohibit the backbones from carrying commercial traffic. He embarked on a program encouraging the regional networks to allow commercial traffic to begin building commercial demand. By 1989 there were eight regionals in operation, connected to each other via the NSF backbone (see Figure 3). NSF and ARPA agreed to interconnect their networks with TCP gateways, and the NSF backbone, with its much higher capacity, superseded ARPANET as the Internet's backbone.

In 1990 MERIT, IBM, and MCI created a private corporation called Advanced Network Services (ANS) and, with Wolff's blessing, turned over NSFNET operation to ANS. ANS then created a for-profit branch. The regionals could use the ANS for-profit backbone to exchange commercial traffic. However, the regionals were concerned that the monopoly held by ANS would lead to high prices.

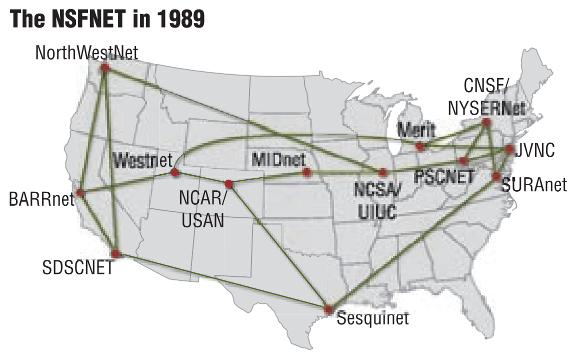

SOURCE: Frazer, Karen. "NSFnet: A Partnership for High-Speed Networking, Final Report, 1987-1995." Merit Network, 1995, p.20.

Figure 3. *Supercomputer networks are SDSCNET, NCAR/USAN, NCSA/UIUC, PCSNET, and JVNC; the other locations are regional networks.*

Therefore, in July 1991 three of the regionals united to form a Commercial Internet Exchange (CIX) with a TCP gateway, which they would jointly operate and where traffic would be exchanged free of charge. This preserved the financial structure users were accustomed to—a charge to connect a computer to the network, but no charges for traffic volume or distance. Soon communication companies such as MCI, AT&T, and Sprint joined the CIX and began building their own packet networks. On April 30, 1995, the NSFNET backbone was terminated; ARPANET had already been shut down in 1990.

Starting at the same time that TCP was defined, two international standards organizations, one controlled by telephone companies, International Telegraph and Telephone Consultative Committee (whose acronym in French is CCITT), and another dominated by computer manufacturers, International Organization for Standardization (ISO), began working on network standardization. To the CCITT, every user device was a terminal and all intelligence would be in the network. They rejected Pouzin's idea of letting reliability be the users' responsibility. They produced the network interface standard called X.25, and during the 1980s, many X.25 networks were built. However, because of the rigid structure of these networks it was very difficult for them to deal with user design innovations. Meanwhile, the ISO developed a layered architectural model and a set of protocols collectively called Standards for Open Systems Interconnection (OSI). By the late 1980s many countries had adopted some portion of OSI as national standards. However, because OSI offered many options, it was possible for two computers to implement OSI protocols and yet not be able to communicate. Because of the problems with X.25 and OSI, many organizations around the world temporarily adopted TCP and the rest of the Internet protocols, with plans to discard them when OSI was more polished.

NETWORKING FOR EVERYONE

One advancement that made it easy for organizations to adopt the Internet was the investment by ARPA in implementing the protocols for several popular operating systems. By the time ARPANET and MILNET split, ARPA had helped write TCP software for most of the major computers used by the DOD. In addition, by 1990 the vast majority of newer systems were running either Unix (officially trademarked by AT&T's Bell Labs as UNIX) or Microsoft Corporation's Windows, and under ARPA's lead there was good Internet software freely available for both. Computer manufacturers had barely begun to provide OSI software with their newest offerings.

In December 1990 Tim Berners-Lee (1955–), an American computer scientist working at the European Organization for Nuclear Research (CERN), demonstrated a system designed to help scientists collaborate by sharing text and multimedia documents. CERN used TCP to link its computers. Above that, Berners-Lee layered a HyperText Transfer Protocol (HTTP). Hypertext is the name given to documents that contain links to other documents. Berners-Lee called the system of linked hypertext documents on multiple computers the WorldWideWeb (Web). HTTP software spread around the Internet, but user interfaces were text-based and hard to use.

Marc Andreessen (1971–), an American programmer and software engineer at the US National Center for Supercomputer Applications began work on Mosaic in 1992. Mosaic was a Web user interface (what is today called a Web browser) based on a graphical display and a pointing device (e.g., mouse). Mosaic established the style for Web browsers still used today. It was made available as a free download at the end of 1990. Forty thousand copies were downloaded in the first month, and over 1 million by the spring of 1994. With Mosaic, one no longer needed computer expertise to use the Web, and the large number of potential viewers inspired a rapid growth in the number of sites offering hypertext documents (websites). In November 1992 there were only twenty-six websites worldwide; there were over ten thousand by August 1995 and millions by 1998.

With the explosive growth in websites after Mosaic was released, there was plenty of information available on the Web, but it was difficult to find any specific item. In 1995 Louis Monier (1956–) and Michael Burrows (1963–) of

DEC created a system called AltaVista to continually search the Web and index almost every word that was found. The index was stored on AltaVista servers, and a high-performance retrieval system could provide a list of every Web document containing any user-specified list of words. AltaVista was made available to the public on December 15, 1995, and recorded 300,000 inquiries its first day. In 1998 daily inquiries numbered 80 million. AltaVista inspired the development of later search engines like Yahoo! and Google. High-performance search engines were the last piece needed for explosive Internet growth. With that growth, X.25 and OSI were forgotten. TCP, HTTP, and the other Internet protocols became the de facto international standards.

Of course, things have not stood still since 1995. Internet service providers have expanded their networks, both in terms of geographic coverage and capacity. The Internet now has Hosts on every continent, even Antarctica. New services have been built on top of TCP, including voice, video, social networking, and commerce. Wi-Fi (the trademarked name for any wireless local area network [WLAN] product) and cellular data services allow mobile access. All this expansion is made possible by the layered design of the protocols, the adaptive routing concepts carried forward from the IMPs to the TCP routers, and the critical decision to put the processing power at the edges of, rather than within, the network.

SEE ALSO *Computers, History of; Email; Mobile Communications and Computing; Personal Computers.*

BIBLIOGRAPHY

Abbate, Janet. *Inventing the Internet.* Cambridge, MA: MIT Press, 1999.

Baran, Paul, et al. *On Distributed Communications.* Santa Monica, CA: RAND Corporation, 1964. Available from http://www.rand.org/about/history/baran-list.html

Blum, Andrew. *Tubes: A Journey to the Center of the Internet.* New York: Ecco, 2012.

Carr, C. Stephen, Stephen D. Crocker, and Vinton G. Cerf. "HOST-HOST Communication Protocol in the ARPA Network." *AFIPS Proceedings, 1970 Spring Joint Computer Conference* 36 (1970): 589–597.

Cerf, Vinton G., and Robert E. Kahn. "A Protocol for Packet Network Interconnection." *IEEE Transactions on Communications* COM-22 5 (May 1974): 637–648.

"Computer Networks—Heralds of Resource Sharing." Internet Archive. Available from https://archive.org/details/ComputerNetworks_TheHeraldsOfResourceSharing

Day, John. *Patterns in Network Architecture: A Return to Fundamentals.* Upper Saddle River, NJ: Prentice Education, 2008.

Hafner, Katie, and Matthew Lyon. *Where Wizards Stay Up Late: The Origins of the Internet.* New York: Simon and Schuster, 1996.

Heart, Frank E., Robert E. Kahn, Severo M. Ornstein, et al. "The Interface Message Processor for the ARPA Computer Network." *AFIPS Proceedings, 1970 Spring Joint Computer Conference* 36 (1970): 551–567.

Kleinrock, Leonard. *Communication Nets: Stochastic Message Flow and Delay.* New York: McGraw-Hill, 1964.

McKenzie, Alexander A. "INWG and the Conception of the Internet: An Eyewitness Account." *IEEE Annals of the History of Computing* 33, no. 1 (January–March 2011): 66–71.

Metcalfe, Robert M., and David R. Boggs. "Ethernet: Distributed Packet Switching for Local Computer Networks." *Communications of the ACM* 19, no. 7 (July 1976): 395–404.

Norberg, Arthur L., Judy E. O'Neill, and Kerry J. Freedman. *Transforming Computer Technology.* Baltimore: Johns Hopkins University Press, 1996.

"Oral Histories: Interviews with Many of the Key Players in the Development of the Internet." Charles Babbage Institute, University of Minnesota. Available from http://www.cbi.umn.edu/oh/

Pouzin, Louis. *INWG Note 42 "Interconnection of Packet Networks."* 1973. Available from http://www.bärwolff.de/public/pouzin-1973-catenet.html

Roberts, Lawrence G. "Multiple Computer Networks and Intercomputer Communication." *Proceedings of the ACM Symposium on Operating System Principles* (Gatlinburg, TN) (1967): 3.1–3.6.

Roberts, Lawrence G., and Barry D. Wessler. "Computer Network Development to Achieve Resource Sharing." *AFIPS Proceedings, 1970 Spring Joint Computer Conference* 36 (1970): 543–549.

Waldrop, M. Mitchell. *The Dream Machine: J. C. R. Licklider and the Revolution That Made Computing Personal.* New York: Viking, 2001.

Alexander A. McKenzie
Bolt, Beranek and Newman, Inc. (BBN) Emeritus