

# the internet: global evolution and challenges

JANET ABBATE

The Internet, a global network of networks, is a remarkably complex technical system built on the creative contributions of scientists around the world from the 1950s to the present. Throughout its evolution, the Internet and other networks have been promoted by governments, researchers, educators, and individuals as tools for meeting a range of human needs. A combination of high-level policy and grassroots improvisation has produced social benefits including easier and more widespread access to computers and information; increased scientific collaboration; economic growth; the formation of *virtual communities* and an increased ability to maintain social ties over long distances; the democratization of content creation; and online political and social activism. The Internet's rapid growth has also spawned technical crises, such as congestion and a scarcity of network addresses, and social dilemmas, including malicious and illegal activities and persistent *digital divides* based on income, location, age, gender, and education. Such problems continue to demand creative solutions from scientists, policy makers, and citizens.

Several general themes characterize the technical development of the Internet. First, from the 1950s to the present there has been a steady increase in the size of data networks and the variety of services they

offer. Rapid growth and diversity have forced network designers to overcome incompatibilities between computer systems and components, manage data traffic to avoid congestion and chaos, and reach international agreement on technical standards. These challenges have led to fundamental advances in research areas such as operating systems and queuing theory. A second trend has been the modeling of network functions as a series of *layers*, each of which behaves according to a standard *protocol*, a set of rules for interaction that is implemented in software or hardware. Layering reduces the complexity of the network system and minimizes the amount of standardization necessary, which makes it easier for networks to join the Internet. A third important feature of the Internet's technical development has been an unusually decentralized and participatory design process. This has opened the system to innovation from a variety of directions and has encouraged informal worldwide collaboration. The following sections describe some of the major milestones in the evolution of the Internet and its predecessors.

## **Beginnings: early terminal networks**

The first electronic digital computers, invented during the World War II and commercialized immediately afterward, were solitary machines: they were not

designed to interact with their human users or to communicate with other computers. Within a few years, however, computer scientists began experimenting with ways to access computers from a distance or transmit data from one machine to another. The data networks of the 1950s and early 1960s were systems to connect terminals to computers, rather than connecting computers to each other. Experiments with terminal networks provided an intriguing research area for computer scientists, but they were also a response to contemporary political and economic realities, including the Cold War and the growth of global economic, transportation, and communication networks.

Computer science research in the United States was largely funded by the military and reflected that country's rivalry with the USSR. For example, an important US development of the 1950s was Project SAGE, a computerized early-warning defense system designed to detect missile attacks. Each SAGE center had an IBM computer that received data through telephone lines from dozens of radar installations and military bases. A key technology developed for SAGE by AT&T Bell Laboratories was the *modem*, which converts digital computer data into analog signals that can be sent over the telephone network. AT&T began to offer modems for general use in 1958, and for several decades modems would provide the chief means of network access for home users.

Demand for terminal networks was driven by another technical milestone of the early 1960s: *time sharing* operating systems. Invented independently in 1959 by Christopher Strachey of the UK and John McCarthy of the US, time sharing allowed multiple users to run programs on a single computer simultaneously. Because the cost of the computer could be shared among a much larger number of users, time sharing made it practical to allow individuals to use a computer interactively for long stretches of time, rather than being restricted to running a single program and receiving the results offline. Commercial time sharing services took advantage of these economies of scale to provide affordable computing to many academic and business customers. By the mid-1960s, commercial time sharing services were developing their own data networks to give their customers low-cost access to their computers.

Global capitalism and the growth of transportation and communication systems provided the impetus for large-scale commercial terminal networks. In the early 1960s, data-intensive industries, such as aviation and stock trading, built cooperative networks to enable firms to share a common pool of information.

For example, in the early 1960s American Airlines and IBM created the SABRE on-line reservation system (based on IBM's work on SAGE), which connected 2,000 terminals across the United States to a central computer. Similarly, the US National Association of Securities Dealers Automated Quotation System (NASDAQ) created a network for stock quotations in 1970. In an early example of international collaboration in networking, a cooperative of airlines called SITA (Société Internationale de Télécommunications Aéronautiques) built a network in 1969 using the *packet switching* technique (see below). The SITA network handled traffic for 175 airlines through computer centers in Amsterdam, Brussels, Frankfurt, Hong Kong, London, Madrid, New York, Paris, and Rome (SITA, 2006). Such financial and commercial networks helped accelerate the integration of the global economy.

### Research networks

Terminal networks were based on a relatively simple hub-and-spoke model that connected numerous users to a single central computer resource. More complex networks involving multiple computers were built by computer scientists from the late 1960s to the late 1970s. Experimenting with new technologies, researchers aimed to break the barriers to sharing data between dissimilar computer systems. Scientists and their government sponsors saw a threefold promise in networking: the ability to share scarce and expensive computers, which would increase access while decreasing costs; the ability to share data and work collaboratively with colleagues in other locations; and the opportunity to advance the theory and practice of computer science.

Three of the most important early research networks were the ARPANET (US, 1969), the NPL Mark I (UK, 1969), and CYCLADES (France, 1972). A key innovation of these experimental networks was a communications technique called *packet switching*. Previous communication systems, such as the telephone and the terminal networks, provided dedicated circuits between the two ends of a connection. In contrast, a packet switching network divides the data to be transmitted into small units called *packets* that are sent out individually, sharing the network circuits with packets from other connections. Packet switching allows communications links to be used more efficiently, thus conserving an expensive resource. In addition, packets from the same connection can be sent to their destination by different routes, making it possible to distribute traffic among multiple links or respond to a breakdown in

one part of the network by routing traffic elsewhere. This flexibility helps prevent congestion and increases the reliability of the network.

The concept of packet switching was invented independently in early 1960s by Paul Baran of the US and Donald Davies of the UK; Davies put the technique into practice in the one-node Mark I network at the National Physical Laboratory. In the US, the Defense Advanced Research Projects Agency (DARPA) sponsored the first large-scale packet switching network, ARPANET. One of the theorists contributing to this project was Leonard Kleinrock, who developed some of the first methods for analyzing packet network behavior. In France, Louis Pouzin pioneered connectionless or *datagram* networking techniques in the packet-switched CYCLADES network. Datagram networks were simpler than connection-oriented networks such as ARPANET, and this simplicity made it more feasible to interconnect different networks—an important step toward developing a worldwide Internet. As Pouzin noted: "The more sophisticated a network, the less likely it is going to interface properly with another." (Pouzin 1975, 429.) Experiments in *internetworking* (connecting multiple networks) were already taking place by the early 1970s. For example, the NPL network was connected to CYCLADES in 1974, and in 1976 both CYCLADES and NPL were connected with the new European Informatics Network. EIN had grown out of a 1971 science and technology study group of the European Economic Community (now the European Union), which recommended the building of a multinational network to help member countries share computer resources and promote computer science research. By 1976 the EIN was providing network service to ten countries, with hubs in Italy, France, Switzerland, and the United Kingdom (Laws and Hathway 1978). The convergence of networking systems thus mirrored the political convergence of the cooperating states.

A number of experimental techniques besides packet switching were featured in the ARPANET. This network connected researchers across the United States working in areas such as time sharing, artificial intelligence, and graphics; because of generous government funding and the large pool of computer science talent involved, the ARPANET builders were able to experiment with promising but extremely challenging techniques. For example, rather than limiting the network to a single type of computer, as had most other experiments in computer-to-computer communication, the ARPANET included a variety of extremely diverse computers. This drove the team of computer scientists, graduate students, and

industry engineers to find ways of bridging the incompatibilities between computers, and their hard work made it much easier to build the next generation of networks. The ARPANET also had a *distributed* topology featuring many switching nodes with multiple interconnections, rather than a single central node. Distributed communications, first described by Baran (1964), could spread out the traffic load and potentially increase reliability by creating multiple paths between any two computers. However, adopting this untried technique greatly increased the complexity of the routing system, forcing the ARPANET designers to analyze and manage unexpected network behavior. In another risky move, the network design called for the routing operations to be decentralized and *adaptive*: each node would make its routing decisions independently and would change its behavior in response to changes in traffic conditions or network configuration (for example, if an adjacent node became disabled). The ARPANET's decentralized design and autonomous routing behavior increased the difficulty of analyzing network behavior; at the same time, these techniques would contribute to the future success of the Internet, because they would allow the network to grow without being limited by a central bottleneck.

One of the most novel features of the ARPANET project was not technical but organizational: an informal, decentralized decision-making process. The network software was developed by a loose confederation of researchers and students called the Network Working Group. Any member of the group could suggest a new feature by circulating a *Request For Comments*; after a period of discussion and trial implementations, the suggestion would be modified, abandoned, or adopted by consensus as a network standard. This collaborative process continues to be used for Internet standards (Bradner 1996) and has helped the system grow and adapt by encouraging free debate and wide participation in its technical development.

By far the most successful application of the early research networks was electronic mail, which became a standard service in the early 1970s. The popularity of email came as a surprise to the ARPANET builders, who had expected that research-oriented networks would focus on sophisticated, computationally-intensive applications such as mathematics or graphics. While email was adopted in part because it was simple to use, its popularity also reflected the realization that scientific research depended as much on human collaboration as on access to machines. Email provided an unprecedented opportunity for ongoing interaction with remote colleagues.

Though they were not open to the general public, the early research networks went beyond providing computer access for a small group of scientists. They produced solutions to formidable technical obstacles and established vital resources for future innovation, including standard techniques and a community of researchers and engineers experienced in networking (Quarterman 1990). Early efforts to build multi-national networks and internets also sowed the seeds of global cooperation, without which today's Internet could not exist.

#### **Expanding access: proprietary, public, and grassroots networks**

In the mid-1970s, the emergence of research networks was paralleled by three other trends: proprietary networking systems offered by computer manufacturers; public data networks built by national telecommunications carriers (PTTs); and grassroots networks that were improvised by individuals with little funding. Companies such as IBM had provided limited networking capabilities since the 1960s, but after the research networks had demonstrated the viability of packet switching, computer firms began offering their own packet-switching technologies. Widely used systems included IBM's Systems Network Architecture (1974), Xerox Network Services (1975), and Digital Equipment Corporation's DECNET (1975). Unlike research networks, these proprietary systems had many corporate users. Corporate networks enabled businesses to be both more distributed—because branch operations could access the data they needed to operate independently—and more centralized, because data from far-flung operations could be instantly monitored by the head office. Thus computer networking reflected and augmented the trend toward economic globalization that accelerated in the 1980s and beyond.

While proprietary systems provided a vital service to organizations with many computers from the same manufacturer, these networks were generally not compatible with computers from rival manufacturers. This could be a problem within a single organization and certainly raised an obstacle to building a national or international network. In addition, these commercial systems were under the control of private corporations and did not adhere to publicly established technical standards. This was of particular concern outside the United States, where most of the large computer manufacturers were located. To provide the public with an alternative, in 1974–75 the national telecommunications carriers in Europe, Canada, and Japan announced plans to build data networks that

would be available to any user, regardless of the brand of computer they used.

The PTTs' vision of data networking, modeled on the phone system, included not only universal access but also international connections. Realizing that this would require agreement on a shared network protocol, in 1975–76 the Consultative Committee on International Telegraphy and Telephony of the International Telecommunications Union developed a packet-switching network standard called X.25. X.25 provided a reliable connection called a *virtual circuit* between two points on a network, allowing terminal users to access online resources without having to install complex networking software. Early adopters of the new standard included Canada's Datapac network (1977), France's Transpac (1978), Japan's DDX (1979), the British Post Office's PSS (1980), and the multinational Euronet (1979). While X.25 was later superseded by other technologies such as *frame relay*, it provided a base for the rapid development of public networks around the world and avoided the chaos of competing incompatible standards. Another influential standards effort in the late 1970s was the Open Systems Interconnection model created by the International Standards Organization. This defined the functions for seven layers of network services, ranging from low-level hardware connections to high-level applications and user interfaces. Although there was much debate over these standards (Abbate 1999), adopting a common model helped computer scientists and manufacturers move closer to creating fully interoperable network systems.

Public data networks provided the first online access for much of the world's population. They also sponsored new types of content and services that made data networks relevant to non-technical users. For example, in the early 1980s France Telecom achieved widespread public use of its Transpac network by offering the innovative Minitel system: a free terminal, given to customers in place of a telephone directory, with access to a free online directory and a variety of paid services. Minitel was in use for almost three decades and served nearly half the French population. With payments securely handled by the phone company, Minitel provided some of the world's first e-commerce, including airline and train ticketing, mail-order retail, banking and stock trading, information services, and message boards (McGrath 2004).

The development of public data networks reflected an emerging view—by both individual users and the highest levels of government—that access to computer communications was a public good, a resource that would be necessary for full citizenship in the twenty-

first century. In serving this mission, public data networks were complemented by a third trend of this period: improvised grassroots networks. These low-cost networks used existing software and simple dial-up connections to exchange mail and discussion lists among an informal community of users. The most well-known were USENET, which was established in 1979 using UNIX protocols, and BITNET, created in 1981 using IBM protocols. These networks played an important role in providing communication to people who had no access to formal networking infrastructure.

### Designing the Internet

How did these disparate data communications systems become united into the global network that we know as the Internet? While some connections between networks were established in the 1970s, design incompatibilities generally limited their services to the exchange of mail and news. The technologies that allow the full range of network services to be shared seamlessly across systems were initially created for the ARPANET. DARPA's explorations in internetworking stemmed from its desire to connect the ARPANET with two new networks it had built, which extended packet switching techniques to radio and satellite communications. Since these media did not have the same technical characteristics as telephone lines—radio links were unreliable; satellites introduced delays—existing techniques such as X.25 or the original ARPANET protocols were not suitable for such a diverse interconnected system. In the early 1970s, therefore, DARPA started an Internet Program to develop a more comprehensive solution.

Another technical development that helped drive the demand for internetworking was *local area networks*. Ethernet, the most influential of these, was invented in 1973 by Robert Metcalfe, drawing on an earlier network called Alohanet that was created by Norman Abramson, Frank Kuo, and Richard Binder (Metcalfe 1996; Abramson 1970). Ethernet and Alohanet pioneered a technique called *random access* that allowed many users to share a communication channel without the need for complex routing procedures.<sup>1</sup> The simplicity of the random access design helped make LANs affordable for a broad range of users. Ethernet became formally standardized and commercially available in the early 1980s and was widely adopted by universities, businesses, and other organizations. Another popular LAN system, token ring, was invented by IBM researchers in Zurich and commercialized in 1985. The popularity of LANs would create many new networks that potentially could be interconnected; but, like the packet radio network,

these random access systems could not guarantee a reliable connection, and therefore would not work well with existing wide-area network protocols. A new system was needed.

The Internet Program was led by Vinton Cerf and Robert Kahn, with the collaboration of computer scientists from around the world. In addition to US researchers at DARPA, Stanford, the University of Southern California, the University of Hawaii, BBN, and Xerox PARC, Cerf and Kahn consulted networking experts from University College London, the NPL and CYCLADES groups, and the International Network Working Group (Cerf 1990). The INWG had been founded in 1972 and included representatives from many national PTTs that were planning to build packet-switching networks. By sharing concerns and pooling ideas, this inclusive team was able to design a system that could serve users with diverse infrastructural resources and networking needs.

The Internet architecture had two main elements. The first was a set of protocols called TCP/IP, or Transmission Control Protocol and Internet Protocol (Cerf and Kahn 1974).<sup>2</sup> TCP was an example of a *host protocol*, whose function is to set up and manage a connection between two computers (*hosts*) across a network. The insight behind TCP was that the host protocol could guarantee a reliable connection between hosts even if they were connected by an unreliable network, such as a packet radio or Ethernet system. By lowering the requirement for reliability in the network, the use of TCP opened the Internet to many more networks than it might otherwise have accommodated. To ensure dependable connections, TCP was designed to verify the safe arrival of packets, using confirmation messages called *acknowledgments*; compensate for errors by retransmitting lost or damaged packets; and control the rate of data flow between the hosts by limiting the number of packets in transit. In contrast, the Internet Protocol performed a much simpler set of tasks that allowed packets to be passed from machine to machine as they made their way through the network. IP became the common language of the Internet, the only required protocol for a network wishing to join: member networks had the freedom to choose among multiple protocols for other layers of the system (though in practice most eventually adopted TCP for their host protocol). Reflecting the diverse needs and preferences of the experts who participated in its design, the Internet architecture accommodated variation and local autonomy among its member networks.

The second creative element was the use of special computers called *gateways* as the interface between

<sup>1</sup> Metcalfe's improved version of the random access system was called Carrier Sense Multiple Access with Collision Detection (CSMA/CD).

<sup>2</sup> Originally there was a single protocol, TCP; it was split into two protocols, TCP and IP, in 1980.



different networks (Cerf 1979). Gateways are now commonly known as *routers*; as the name implies, they determine the route that packets should take to get from one network to another. A network would direct non-local packets to a nearby gateway, which would forward the packets to their destination network. By dividing routing responsibility between networks and gateways, this architecture made the Internet easier to scale up: individual networks did not have to know the topology of the whole Internet, only how to reach the nearest gateway; gateways needed to know how to reach all the networks in the Internet, but not how to reach individual hosts within a network.

Another notable invention that would make the worldwide growth of the Internet manageable was the Domain Name System, created in 1984 by Paul Mockapetris (Cerf 1993; Leiner et al 1997). One challenge of communicating across a large network is the need to know the address of the computer at the far end. While human beings usually refer to computers by names (such as "darpa"), the computers in the network identify each other by numerical addresses. In the original ARPANET, the names and addresses of all the host computers had been kept in a large file, which had to be frequently updated and distributed to all the hosts. Clearly, this mechanism would not scale up well for a network of thousands or millions of computers. The Domain Name System decentralized the task of finding addresses by creating groups of names called *domains* (such as .com or .org) and special computers called *name servers* that would maintain databases of the addresses that corresponded to each domain name. To find an address, the host would simply query the appropriate name server. The new system also made it possible to decentralize the authority to assign names, so that, for example, each country could control its own domain.

#### The World Wide Web and other applications

The Internet architecture made it possible to build a worldwide data communications infrastructure, but it did not directly address the question of content. In the 1980s, almost all content on the Internet was plain text. It was relatively difficult for users to locate information they wanted; the user had to know in advance the address of the site hosting the data, since there were no search engines or links between sites. The breakthrough that transformed how Internet content was created, displayed, and found was the World Wide Web.

The World Wide Web was the brainchild of Tim Berners-Lee, a British researcher at CERN, the

international physics laboratory in Geneva. He envisioned the Internet as a collaborative space where people could share information of all kinds. In his proposed system, users could create pages of content on computers called *web servers*, and the web pages could be viewed with a program called a *browser*. The Web would be able to handle multimedia as well as text, and Web pages could be connected by *hyperlinks*, so that people could navigate between sites based on meaningful relationships between the ideas on different pages. This would create a web of connections based on content, rather than infrastructure. Berners-Lee formulated his ideas in 1989, and he and collaborator Robert Cailliau created the first operational version of the Web in 1990. The technical underpinnings of the new system included *html* (hypertext markup language, used to create web pages), *http* (hypertext transfer protocol, used to transmit web page data), and the *url* (uniform resource locator, a way of addressing web pages).

The Web was popular with the physicists who used it at CERN, and they spread it to other research sites. At one such site, the US National Center for Supercomputer Applications, Marc Andreessen led the development of an improved browser called Mosaic in 1993. Mosaic could run on personal computers as well as on larger machines, and NCSA made the browser freely available over the Internet, which led to a flood of interest in the Web. By 1994 there were estimated to be a million or more copies of Mosaic in use (Schatz and Hardin 1994).

The Web's hyperlinks were designed to solve a long-standing problem for Internet users: how to find information within such a large system? To address this need, various finding aids were developed in the 1990s. One of the earliest tools for searching the Internet was Archie (1990), which sent queries to computers on the Internet and gathered listings of publicly available files. Gopher (1991) was a listing system specifically for the Web, while Yahoo (1994) was a directory of Web pages organized by themes. Yahoo's staff categorized Web pages by hand, rather than automatically; given the vast amount of data accumulating on the Web, however, a variety of new services tried to automate searching. The most successful of these *search engines* was Google (1998). Search engines transformed the way users find information on the Web, allowing them to search a vast number of sources for a particular topic rather than having to know in advance which sources might have relevant information.

Like the Internet itself, the Web was designed to be flexible, expandable, and decentralized, inviting people

to invent new ways of using it. The spread of the World Wide Web coincided with the transition in 1995 of the US Internet backbone from government to private-sector control. This removed many barriers to commercial use of the Internet and ushered in the "dot-com" boom of the 1990s, in which huge amounts of capital were invested in e-commerce schemes. While the dot-com bubble burst in 2000, it was significant in creating a popular understanding of the Internet as an economic engine and not merely a technical novelty. The beginning of the twenty-first century also saw the proliferation of *social media* that provided new ways for people to interact and share information and entertainment online. These included weblogs (1997), wikis (1995), file sharing (1999), podcasting (2004), social networking sites, and a variety of multi-player games.

#### **The Internet and society: successes and challenges**

After half a century of research and innovation, the Internet was firmly established as a widely available resource offering an array of potential benefits. Users had greater access to information of all kinds, and governments and businesses had a new platform for providing information and services. E-commerce brought economic growth, greater choices for consumers, and opportunities for producers in disadvantaged areas to reach new markets. A variety of communications options, from email to elaborate social networking sites, made it easier for friends and family to stay in touch over long distances and for strangers to form "virtual communities" around common interests. Grassroots organizers adopted the Internet for political and social activism and used it to mobilize worldwide responses to natural disasters and human rights abuses. Users of all ages embraced the Internet as a medium for personal expression, and new applications helped democratize the technology by making it easier for ordinary people to independently produce and disseminate news, information, opinion, and entertainment.

However, many challenges remained as the Internet entered the twenty-first century. Users faced abusive practices such as spam (unwanted commercial email), viruses, identity theft, and break-ins. Technical experts responded with solutions that attempted to minimize these ongoing dangers, providing anti-virus systems, filters, secure web transactions, and improved security systems. But other issues were too divisive for a technical solution to satisfy conflicting public opinion, especially when activities crossed national boundaries. Some governments severely limited and closely

monitored the online activities of their citizens; while human rights groups protested this as censorship and intimidating surveillance, the governments in question asserted their right to protect public safety and morality. Other groups complained that the Internet was *too* open to objectionable or illegal content such as child pornography or pirated songs, movies, and software. Filters and copyright protection devices provided means to restrict the flow of such information, but these devices were themselves controversial. Internet governance was another thorny issue, with many of the world's nations calling for a more international, less US-dominated mechanism for managing the Internet's name and address system. Another technical issue with political ramifications was the proposed transition from the old Internet Protocol, called IPv4, to a new protocol called IPv6 that would provide a much larger number of addresses (Bradner and Mankin 1995); this was in part a response to the fact that the United States held a disproportionate share of the IPv4 addresses. Ipv6 was proposed as an Internet standard in 1994, but due to technical and political disagreements the protocol was still only used for a tiny percentage of Internet traffic 15 years later (DeNardis 2009). Given these many obstacles, the Internet's decentralized, consensus-based development process continued to work remarkably well to keep the system thriving amid rapid growth and change.

Perhaps most troubling was the persistent inequality of access to the Internet and its opportunities for economic development, political participation, government transparency, and the growth of local science and technology. Significant gaps remained between rich and poor regions, urban and rural citizens, young and old. The United Nations reported in 2007 that the global digital divide was still enormous: "Over half the population in developed regions were using the Internet in 2005, compared to 9 per cent in developing regions and 1 per cent in the 50 least developed countries." (UN, 2007, 32.) To help address this issue, the UN and International Telecommunications Union sponsored a two-part World Summit on the Information Society in Geneva (2003) and Tunis (2005) to devise a plan of action to bring access to information and communication technologies to all of the world's people (WSIS 2008). Computer scientists also devoted their ingenuity to making the Internet more accessible to the world's poor. For example, in 2001 a group of Indian computer scientists reversed the paradigm of expensive, energy-consuming personal computers by creating the Simputer: a simple, low-cost, low-energy computer

3 "Internationalizing" the governance of the Internet was a central issue at the UN-sponsored World Summit on the Information Society in 2005.

4 The creators and trustees of the Simputer project were Vijay Chandru, Swami Manohar, Ramesh Hariharan, V. Vinay, Vinay Deshpande, Shashank Garg, and Mark Mathias (<http://www.simputer.org/simputer/people/trustees.php>).

that would provide a multilingual interface and could be shared among the residents of a village (Sterling 2001). Similarly, Nicholas Negroponte initiated the One Laptop Per Child project in 2005 to serve educational needs in developing countries. To help fit the technology to local needs, lead designer Mary Lou Jepsen invented an inexpensive, power-efficient screen readable in outdoor light, and software designer Walter Bender created an intuitive graphical user interface (One Laptop Per Child 2008; Roush 2008). The Stockholm Challenge, an annual event since 1995, showcases hundreds of innovative projects from around the world that use ICTs to promote development (Stockholm Challenge 2008).

No longer simply the domain of scientists, pushing the frontiers of the Internet increasingly involves

social as well as technical innovation and the collaboration of researchers, businesses, civil society organizations, governments, and ordinary people. The values guiding the Internet's social and technical development have been complementary: increasing access, accommodating diversity, decentralizing authority, making decisions by consensus with a wide range of participants, and allowing users to take an active role in adding features to the network. On the technical side, these goals have been achieved through layered architecture, open protocols, and a collaborative process for approving design changes, while social goals have been advanced through government leadership and the inspiration of individuals who saw the Internet's potential for communication, cooperation, and self-expression.



## Bibliography

- Abbate, Janet. *Inventing the Internet*. Cambridge: MIT Press, 1999.
- Abramson, Norman. "The ALOHA System —, Another Alternative for Computer Communications." *Proceedings, AFIPS Fall Joint Computer Conference*. Montvale, NJ: AFIPS Press, 1970, 281–285.
- Baran, Paul. *On Distributed Communications*. Santa Monica, CA: RAND Corporation, 1964.
- Berners-Lee, Tim. *Weaving the Web*. New York: HarperCollins, 1999.
- Bradner, Scott, and A. Mankin. "The Recommendation for the IP Next Generation Protocol." *Network Working Group Request for Comments 1752*, January 1995. Available on the Internet at <http://www.rfc-editor.org/>
- Bradner, Scott. "The Internet Standards Process—Revision 3." *Network Working Group Request for Comments 2026*, October 1996. Available on the Internet at <http://www.rfc-editor.org/>
- Campbell-Kelly, Martin, and William Aspray. *Computer: A History of the Information Machine*. New York: Basic Books, 1996.
- Cerf, Vinton G. "DARPA Activities in Packet Network Interconnection." In K. G. Beauchamp, ed. *Interlinking of Computer Networks*. Dordrecht, Holland: D. Reidel, 1979.
- Cerf, Vinton G. Oral history interview by Judy O'Neill (Reston, VA, April 24, 1990), OH 191. Minneapolis, MN: The Charles Babbage Institute, University of Minnesota, 1990. Available on the Internet at <http://www.cbi.umn.edu/oh/display.phtml?id=118>
- , "How the Internet Came to Be." In Bernard Aboba, ed. *The Online User's Encyclopedia*. Addison-Wesley, 1993.
- Cerf, Vinton G., and Robert E. KAHN. "A Protocol for Packet Network Intercommunication." *IEEE Transactions on Communications COM-22*, May 1974, 637–648.
- Denardis, Laura. *Protocol Politics: The Globalization of Internet Governance*. Cambridge: MIT Press, 2009.
- Laws, J., and V. Hathway. "Experience From Two Forms of Inter-Network Connection." In K. G. Beauchamp, ed. *Interlinking of Computer Networks*. NATO, 1978, 273–284.
- Leiner, Barry M., Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff. 1997. "A Brief History of the Internet." Available on the Internet at <http://www.isoc.org/internet-history>. Revised February 1997.
- McGrath, Dermot. "Minitel: The Old New Thing." *Wired*, January 18, 2004. Available on the Internet at <http://www.wired.com/science/discoveries/news/2001/04/42943>.
- Metcalfe, Robert M. *Packet Communication*. San Jose: Peer-to-Peer Communications, 1996.
- One Laptop Per Child. <http://www.laptop.org/> (accessed September 5, 2008).
- Pouzin, Louis. "Presentation and Major Design Aspects of the CYCLADES Computer Network." In R. L. Grimsdale and F. F. Kuo, eds. *Computer Communication Networks*. Leyden: Noordhoff, 1975.
- Quarterman, John S. *The Matrix: Computer Networks and Conferencing Systems Worldwide*. Burlington, MA: Digital Press, 1990.
- Roush, Wade. "One Laptop Per Child Foundation No Longer a Disruptive Force, Bender Fears." *Xconomy*, April 24, 2008. Available on the Internet at <http://www.xconomy.com/boston/2008/04/24/>
- Schatz, Bruce R., and Joseph B. Hardin. "NCSA Mosaic and the World Wide Web: Global Hypermedia Protocols for the Internet." *Science* 265 (1994), 895–901.
- SITA. "SITA's history and milestones." [http://www.sita.aero/News\\_Centre/Corporate\\_profile/History/](http://www.sita.aero/News_Centre/Corporate_profile/History/) (updated 21 July 2006, accessed September 5, 2008).
- Sterling, Bruce. "The Year In Ideas: A to Z; Simputer." *The New York Times*, December 9, 2001.
- Stockholm Challenge. [www.challenge.stockholm.se](http://www.challenge.stockholm.se) <<http://www.challenge.stockholm.se>> / (accessed September 5, 2008).
- United Nations. *The Millennium Development Goals Report 2007*. New York: United Nations, 2007.
- World Summit on the Information Society. <http://www.itu.int/wsis/> (accessed September 5, 2008).