The Internet After Thirty Years

Peter J. Denning

5/21/97
7/4/97 revised

This essay is based on the article "The ARPANET after twenty years", *American Scientist 77*, 6 (November-December 1989), as an installment in the series of columns, *The Science of Computing*.

The ARPANET began operation in 1969 with four nodes as an experiment in resource sharing among computers. By 1971 fifteen nodes were operating and by 1973 thirty-seven nodes. In 1977 it started using the Internet Protocol (IP), a universal connector of networks. Twenty years later its prolific offspring, the Internet, had burgeoned into a world-wide research network of over 60,000 nodes. ARPA officially disbanded the ARPANET in 1989. In 1997, nearly thirty years after the ARPANET was founded, the Internet counted over 20 million computers and 50 million users. The ARPANET influenced the design of many other networks in business, education, and government. It demonstrated the speed and reliability of packet-switching networks. Its protocols were the models for international standards.

Yet the significance of the Internet lies not in its technology, but in the profound alterations it has produced in human practices. Network design is no longer the sole province of engineers. The Internet has become an information superhighway, an enormous digital library, a nervous system connecting scientists and their instruments world-wide, an electronic marketplace for commerce, a world-wide corporate and individual identity projector, a nurturing ground for relationships, a multimedia entertainment center, a virtual university, a breeding ground for conspiracies, and a challenge to the power of governments. It is evolving at the intersection of science, engineering, business, libraries, art, entertainment, education, and politics. It is a space in which the best and worst of humanity appear.

The changes in our use of computers which began thirty years ago are, in retrospect, nothing short of revolutionary. In what follows, I will discuss the origins of the ARPANET and its evolution into the Internet; reflect on its influence our practices; and speculate about the issues that network designers will face next.

Beginnings

The ARPANET story begins in the late 1950s during the era of intercontinental ballistic missiles. The Department of Defense, concerned about the ability of US forces to survive a nuclear first-strike, gave high priority to the durability of the communication network. They commissioned a series of investigations by Paul Baran of the Rand Corporation. Baran concluded that the strongest communication system would be a distributed network of computers with: (1) redundant links; (2) no central control; (3) all messages broken into equal-size packets; (4) variable routing of packets depending on the availability of links and nodes; and (5) automatic reconfiguration of routing tables immediately after the loss of a link or node. Baran's reports became public in 1964 and are available from Rand Corporation's Web site [1].

Meanwhile, Larry Roberts of MIT's Lincoln Laboratory, inspired by conversations with J. C. R. Licklider of the Defense Department's Advanced Research Projects Agency (ARPA), dedicated himself to building networks that made sharing of computers and data simple and economical. In 1965, inspired by Licklider and Roberts, Donald Davies of the National Physical Laboratory in England proposed a packet-switched computer network; it would use telephone trunk lines ranging in speed from 100 kilobits per second to 1.5 megabits per second, 128-byte "packets", switching computers that could process 10,000 packets per second, and special interface computers connecting mainframe "hosts" to the packet network without alterations to their operating systems. From his own experiments in 1966 with direct-dialed telephone links between computers, Roberts concluded that the packet-switching proposals of Baran and Davies would overcome slow and unreliable telephone circuits and would be cheaper. Leonard Kleinrock of UCLA soon produced analytic models of packet-switched networks that could be used to guide a design.

Meanwhile, Robert Taylor had succeeded Licklider at ARPA. Taylor, a psychologist by training, brought an interest in human relationships as a new dimension of computer networking. Previous ARPA projects had created a variety of powerful computational centers at different institutions, each with its own user community and potential to be a national resource. Taylor was interested in the benefits that might arise if these user communities would interact and collaborate as well as share their resources. He envisioned a fast, robust network connecting the centers. In 1967, he persuaded Roberts to come to ARPA and head up the network project. Roberts presented a detailed proposal for the network at the first symposium on operating systems principles in late 1967. The next year, ARPA awarded a contract to a group headed by Frank Heart at Bolt Beranek and Newman (BBN) to build the first interface message processors (IMPs), computers as proposed by Davies to connect mainframes and their operating systems with the network. The first four IMPs were delivered by the end of 1969, and the first packet-switched network was operating by the beginning of 1970. The first public demonstration of this network was organized by Robert Kahn of BBN at the International Conference on Computer Communications in 1972. Kahn succeeded Taylor at ARPA a couple of years later.

Although electronic mail was not among the early goals of the ARPANET, by 1971 mail accounted for most of the traffic. Most users thought of the network as a way of communicating with colleagues and a tool of collaboration.

By the mid 1970s, it was clear that the ARPANET would not be alone. Networks were being developed by other governments and major companies. IBM marketed a network technology for connecting the IBM computers owned by their customers. Kahn, now at ARPA, concluded that there would soon be a need to connect networks together. He asked Vinton Cerf to design a new protocol that would permit users to interconnect programs on machines attached to different networks. By 1977, Cerf completed the design of a matched pair of protocols called TCP (transport control protocol) and IP (Internet protocol). IP routed packets across multiple networks. TCP converted messages into streams of packets, reassembled them into messages at the receiver, and recovered them in case of packet-loss in the underlying network. These two protocols provided highly reliable end-to-end communication in a network of networks. They evolved into protocols approved for world-wide use by the International Standards Organization.

During the 1970s, a variety of European networking projects imitated and improved on the ARPANET technology. The Consultative Committee for International Telegraphy and Telephony (CCITT) devised a protocol that simulated the traditional end-to-end voice circuit on an underlying packet-switched network; designated X.25, this protocol was approved as a standard in 1975 and is widely used in Europe. Some X.25 service has been available in the US since the early 1980s.

You can read more about these developments in a special issue of IEEE *Proceedings*, edited by Kahn in 1978 [15]. This is a collection of sixteen papers on all aspects of packet networks including the original ARPANET, packet radio (precursor of today's cellular telephones), local networks such as Ethernet, and social implications. Some historical notes were given by Denning [7] and a detailed history by Katie Hafner and Matthew Lyon [14].

Networked Communities [6]

The many networks, public and private, that grew up in the 1970s and 1980s have been chronicled and cataloged by John Quarterman and Josiah Hoskins in an article and a book [18,19]. The ARPANET and company networks had been designed by central agencies seeking to connect members of existing communities. Another network grew up spontaneously during this time, without guidance from a central authority. It connected Unix computers that had proliferated on many campuses and research labs in the 1970s. A protocol called UUCP (Unix to Unix copy) was designed to allow one Unix computer to dial up another via a modem and then transfer mail. Mail transport programs routed mail to local mailboxes or relayed it to other computers specified in the UUCP address. People in the USENET, as the UUCP network came to be called, invented newsgroups, a practice that has become a major part of the Internet today. Without a central coordinating body, individual system administrators selected the shortest paths to distribute mail and news. This placed major demands on some sites, consuming many CPU cycles and amassing large telephone bills, forcing many of them to drop out as relays.

In a parallel development, users of IBM computers banded together and created BITNET. Each new site had to obtain a leased telephone line to the nearest existing BITNET site, and had to agree to relay BITNET mail to other computers that eventually connected to it. Served by a coordinating organization, BITNET did not have the congestion and load problems of USENET.

The USENET, BITNET, and X.25 networks could not be connected to the ARPANET because of a government policy limiting ARPANET to government agencies and their contractors. The turning point that eventually brought them all together was the CSNET project, which was created in 1981 by Larry Landweber (University of Wisconsin), David Farber (University of Delaware), Anthony Hearn (Rand Corporation), and Peter Denning (Purdue University) under a five-year, $5M grant from the National Science Foundation [14]. The purpose of CSNET was to link all the computer science departments and industry labs engaged in computing research. It provided TCP/IP interfaces with USENET, BITNET, and X.25, and established nameserver databases to enable any computing researcher to locate any other. This interconnectivity was enabled by a landmark agreement between NSF and ARPA: ARPA allowed NSF grantees and affiliated industry research labs access to ARPANET, as long as no commercial traffic flowed through ARPANET. Armed with this agreement and its own experience in managing networking through the CSNET project, NSF went on to sponsor NSFNET, a high-speed backbone connecting its supercomputing research centers. NSF also stimulated the growth of regional networks that could attach to the NSFNET backbone. By 1990, the NSFNET had become the backbone of the modern Internet. In 1996, NSF handed over its management to commercial Internet Service Providers (ISPs).

The rapid expansion of the Internet during the early 1980s prompted ARPA to devise a hierarchical system of translating domain names into IP addresses. The original method -- storing all these associations in a single file that was copied regularly to all the routing servers -- had become unwieldy. Introduced in 1984, the new Domain Name Service (DNS) distributed the job of associating domain names with IP addresses among a large number of nameservers, one associated with each domain and subdomain.

By 1989, the process of unifying the many community networks within the Internet was complete and the ARPANET was officially disbanded. Few users noticed.

In that same year, Tim Berners-Lee of CERN, the Zurich research center for high-energy physics, completed the design of a network-based system for traversing hypertext links. He wanted to make it easy for his colleagues in high-energy physics to obtain documents cited in other documents. He devised the universal resource locator (URL) to name documents, the hypertext transport protocol (HTTP) to transfer them, and the hypertext markup language (HTML) to identify text-strings that were active hyperlinks within a document. He called this system of Internet-wide linked documents the World Wide Web [3,4]. Two years later, at the National Center for Supercomputing Applications (NCSA), Marc Andreeson designed the Mosaic browser, a simple multi-media interface for HTML documents and the HTTP. When this was made public in 1992, it catapulted the WWW into world-wide prominence and greatly accelerated the interest in and growth of the Internet. Andreeson founded Netscape in 1994, a company that is now at the center of the WWW.

Expanding the Space of Human Action

Imagine that we brought Henry Ford back to show him today's automobiles. He would not be so surprised by changes in design: cars still have four wheels, steering, front-mounted internal-combustion engines, transmissions, and the like. But he would be greatly surprised by the changes in human practices that have grown up around the automobile -- for example, hot rods, strip malls, drive-in fast food, cars as status symbols, licensing of drivers, rush hours, the interstate highway system, traffic reports on the radio, and much more.

Alexander Graham Bell would be little surprised by the design of telephone instruments and switching systems -- handsets, carbon microphones, dialing mechanisms, crossbar switches, telephone exchanges, and operator services. But he would be amazed by the changes in human customs that have grown up around the phone -- telephone credit cards, the Home Shopping Network, universal phone numbers, cell phones, "prestige" exchanges like New York's Butterfield 8 or Beverly Hills's 271, public outcries over changes in area codes, electronic funds transfers, telemarketing, faxes, telephone pornography, and much more.

Edison would doubtless be little surprised by modern light bulbs and generators, but he would be astonished by international power grids, night baseball, radio and television, lava lamps, electronics, computers, and much more. For that matter, can you imagine trying to explain frequent flier miles to Orville Wright?

These examples illustrate the beginnings and ends of evolutionary processes connecting an inventor's initial declarations to widespread practices many years later. Although a technology does not drive human beings to adopt new practices, it shapes the space of possibilities in which they can act: people are drawn to technologies that expand the space of their actions and relationships. When that draw is present, the process passes through six stages:

Declarations
Prototypes
Tools
Industries
Widespread practices

Infrastructure

The movement through these stages is not smooth and well-delineated. It is best described as a drift buffeted by many events that make it impossible to predict where it will end up. The time scale for the drift from the first to the last stage is long -- one or two generations, or 20 to 50 years [5].

The Internet illustrates this process [6]. In and around 1965 the first declarations were made; they took the form of design proposals and the commitments of funds within ARPA. By 1970 the first prototypes were operating in the early ARPANET. By 1977, the first Internet tools were in place: TCP/IP protocols, electronic mail, file transport (FTP), remote login (telnet), telephone login, and newsgroups. By 1985, we could discern the early stages of industries, in the form of organizations that provided connection and coordination services, notably CSNET, BITNET, and the loose federation USENET. There was also one commercial Internet Service Provider (ISP), GTE Telenet, which offered X.25 service. By 1990 the Internet was a widespread practice; many people listed e-mail addresses on their calling cards and stationery and many commercial service providers had joined the industry. Moreover, the Clinton administration embraced the "information superhighway" as a national infrastructure for every home and school. By 1997, the Internet, and its multi-media subset the World Wide Web, was an important (but still fragile) part of the infrastructure. Most commercial advertisements included "http" addresses, Internet tools were bundled with personal computers, e-mail was a standard part of doing business, schools everywhere had computers in the classroom, virus detection and eradication was a major industry, many education offerings were available on web pages, scientific journals were being distributed electronically, and telephone and video were broadcast on the Internet. A host of cultural and political controversies had arisen.

It is easy, in retrospect, to describe the different stages; but it is not so easy to predict when they will happen and what practices will constitute them. Electronic mail was not mentioned among the original goals of the ARPANET, and yet within two years, as we have seen, it was the major source of traffic. Even at the founding of CSNET in 1981, after a decade of electronic mail experience with the ARPANET, the NSF did not want to base its argument for the new network on the demand for electronic mail. Today electronic mail is accepted as a sufficient reason for networks. Connectivity also emerged unexpectedly as a driving concern. Interruptions in the flow of electronic mail are now considered major disasters, as we witnessed in the Internet Worm incident of November 1988 [8]. High speed personal workstations became increasingly inexpensive and powerful and are now individual nodes in the networks. Electronic publication and digital libraries have emerged as industries in their own right, threatening traditional print publishing and placing heavy demands on networks to move manuscripts. The now ubiquitous fax combines the widespread practice of sharing paper documents with the wide reach of the telephone network, facilitating transactions across time zones, borders, and languages. The WWW, sometimes called the "killer application of the 1990s" [2], was a complete surprise.

These events, surprising to most, were foreseen by some visionaries. Licklider and Taylor foresaw the importance of electronic mail; but they knew they could not convince their colleagues that spending millions of dollars to facilitate e-mail was the wave of the future. Computer security experts in the 1960s anticipated most of the attacks and intrusions that have since occurred; they just could not convince their counterparts to spend millions of dollars to add these functions to already-expensive operating systems. Ted Nelson and Doug Englebart saw world-wide hyperlinked structures as a basis for sharing intellectual property in the 1970s [11,17]; they just could not convince their colleagues that their schemes would become practical. Fernando Flores and Terry Winograd described software to support transactions between customers and performers in the 1980s [21,22]; their colleagues were skeptical of its universality as a communication pattern and the worth of retrofitting it into networks and operating systems. All these visionaries have been vindicated: e-mail is a common practice, information warfare and computer security are buzz words, the hyperlinked world wide web is familiar to every schoolchild, and workflow management has become a major industry. What methods did these visionaries use for observing the movements of the world? Can the rest of us learn from them?

In looking ahead to see where the drift may be taking us, there are two things we can do. The first is to look for major breakdowns -- problems and interruptions that beset many people who demand solutions. Then look for practices that some people or groups use that would resolve the breakdowns [12]. These practices will be marginal or anomalous: they will not appear to be mainstream, they will not fit in with prevailing wisdom or common sense, they will be deprecated or dismissed. But they cannot be ignored, for they contain the seeds of the resolutions of the breakdowns. The visionaries see this clearly and offer captivating stories about how the anomalous practices may evolve into solutions. The visionary leaders found enterprises that eventually turn the anomalies into mainstream practices. Here are some of the major controversies today that have arisen against the backdrop of the Internet; they are the breeding grounds for the innovations of tomorrow:

- **Domain Names and Trademarks:** Legal disputes now exist between established companies who want their Internet domain name to match their registered company name and others who took the name a few years before. Some con artists use names similar to a company's name in order to trick people into accessing the surrogate Web site instead of the real one. These disputes show that the protection and projection of company identities is an increasing concern [12].
- **Intellectual Property:** Copying is easy and cheap in the Internet; anyone can be a publisher or distributor. Entertainment companies seek strict international agreements regarding the copying and transmission of copyrighted images and sounds. Educational institutions and scientific researchers seek international accords that permit them to freely copy and disseminate their digital works. The two groups argue about whether viewing Web pages is an infringement

of copyright. These debates signal a big change in our definitions of intellectual property.

- **Telepresence:** Many people place great stock in technologies that permit people to simulate their presence elsewhere [16]; Gordon Bell and Jim Gray think that these technologies will be the next "killer applications" after Web browsers [2]. Urban planners argue that telepresence technologies will promote telecommuting and save enormous sums now lost to traffic congestion. Environmentalists worry that telepresence will allow people to flee the cities and eventually ravage the remote areas now visited by nature lovers. Travel agents say that telepresence may improve business because people want to meet their new friends and colleagues from the Internet. These disputes foreshadow significant changes in the way people form and maintain relationships.
- **Software Engineering of Safety-Critical Systems:** In 1968, a group of researchers meeting at a NATO conference declared a software crisis and founded the discipline of software engineering to address it. Thirty years later, software engineers were no closer to being able to systematically design and deliver reliable, dependable, and safe systems. This is not the fault of the tools and methods, but of the large increase in the size and distributed complexity of applications -- such as international funds-transfer for banking, fly-by-wire aircraft, air traffic control, medical diagnostic and conferencing systems, railway signaling systems, or electric power grid control systems. These concerns signal a new profession for licensed software engineers and new, rigorous education processes for certification.
- **Protection, Security, Authentication, and Law Enforcement:** Companies now depend almost completely on information technology to store business data and records, and to distribute reports around the company. They do not want intruders or saboteurs to steal, disclose, destroy, or damage their proprietary data. They want mechanisms that can distinguish authentic customers from thieves and impostors and trace Internet connections. Meanwhile, police increasingly encounter crimes where criminals use Internet facilities to perform their acts remotely, hide their locations, and conceal their identities and plans. In the US, Presidential proposals for key recovery systems that allow law enforcers access to encryption keys meet with stiff resistance from civil libertarians. This debate exposes a deep conflict concerning public safety and police powers.
- **National Borders:** Governments can no longer control the flow of information because too many individuals have been given the power by the Internet to decide for themselves what information to move and when. Most export controls do not work for software or data. It is often difficult to ascertain where a business transaction has occurred, which means that taxes cannot be collected. It is also hard to tell whose laws apply to a given transaction since the parties may be in different

jurisdictions. This foretells a decline in the power of public institutions and possibly an era of social unrest.

- **Trust:** To control costs and lower customer complaints, many managers have turned to sophisticated systems that monitor detailed actions of employees, measure productivity, control access, and create audit records. Employees, however, interpret these mechanisms as a means of surveillance, an institutionalization of the distrust the mechanisms are supposed to render unnecessary, and even a deprivation of dignity. These disputes signal a renewal of trustworthiness as a virtue.
- **Distance Education:** Hypesters have declared that the Internet is the greatest educational tool yet devised by humankind. They use such buzz words as virtual universities and distance delivery. Some have even declared that education will eventually be so automated that the consumer-student can have access to the world's great professors, recorded and imitated by intelligent agents, at a fraction of today's college tuition. An increasing number of critics say this isn't so: most learning occurs in real situations under the guidance of an experienced expert, something the Internet cannot do with its simulators [10,20]. In the end, the issue is who controls the curriculum: the consumer or the faculty? This debate foreshadows a decline in the social power of universities and a rise in private, for-profit educational services.
- **Virtual Reality:** Many people are enamored of computer simulations which are so realistic that a person cannot distinguish from the real thing. Its supporters call it the greatest source of education innovations ever known. Its detractors say it decouples people from reality and teaches them to be unsociable or, even worse, antisocial. Both supporters and detractors tacitly agree that virtual reality is a powerful technology, but this controversy signals the birth of a major new industry that will force the current Internet into the background.
- **Computer Modeling:** Computer models forecast future environmental, economic, and political conditions; more and more, policy planners turn to them to set public laws and regulations. Many critics object, saying that the chaos, capriciousness, and uncertainty of the real world cannot be accurately captured by any centralized model. Thus far, the critics are on the losing side. These disputes signal the development of new industries marketing computer models as planning and forecasting aids for an increasingly uncertain world.
- **Artificial Intelligence:** For the most part, the ambitious goals of the 1950s have not been met [22]. Now most AI people say their job is to augment human intelligence, not imitate it. Whether or not a machine can think is no longer on most people's minds. Still there are numerous conflicts over whether AI-inspired approaches will yield better designs than careful engineering approaches. These arguments foretell a new attention to the design and utility of computing in support of everyday practice.

- **Locus of Computing Research:** Some observers believe that computer technologists have lost their edge, that they are refining old inventions more than creating new ones. Others admire the surge of innovation occurring as people in other disciplines, especially the humanities, become programmers enough to enable them to incorporate computers into their own practices. Is the best computing research being done in other disciplines? Will computer science pull inward and focus solely on technology development, leaving design in the hands of outsiders? Will federal forbearance continue, or will there be fewer and fewer research dollars for unfettered investigations? These disputes signal new social contracts between universities, government, and industry for research supported by public monies [9].

The second thing we can do about the future is to look for deep trends, ones that are o big and powerful but have so much inertia that they cannot be deflected very much. Then extrapolate their consequences. Some of our visionary leaders have told us what principles and methods they use. Here are three:

- **Moore's and Gilder's Laws.** Intel founder Gordon Moore once noted empirically that the speed and capacity of microprocessors and memory chips doubles every 18 months [2]. The experts give Moore's law another twenty years until the size of transistors and wires will be too small to function. What then? Researchers are already foraging for major new paradigms of computation. Neural computing seeks silicon devices whose structure imitates nature, enabling direct simulation of biological entities -- for example, artificial eyes and ears and maybe an artificial brain. Biological computing aims to grow organic memories, neurons, and, one day, computers. DNA computing aims to encode large combinatorial problems as strands of DNA for which the solution will emerge by chemical reaction. Quantum computing aims to encode data in quantum waveforms that can coexist in the same "space-time" and compute through their interactions and interferences. No one knows how these explorations will turn out. But they all continue the trend toward moving massive computing power into ever-smaller packages at the site of the application, making it portable and taking it away from the center of the network -- what George Gilder calls the "law of the telecosm."
- **Grove's Strategic Inflection Points.** Intel Chairman Andy Grove has observed that businesses and industries encounter inflection points -- points at which the business can either expand or decline [13]. They are often induced by technologies that shift some aspect of production, competition, or supply by a factor of 10 or more. The successful companies anticipate these inflections and prepare to change the business to meet them. It takes patience and practice to learn the necessary skills of observation.
- **Flores's Three Ages.** Business Design Associates Chairman Fernando Flores has divided the next fifty years into three eras [12]. During the information era, which he says is now coming to an end, business

communication, transactions, and marketing were perceived as forms of information transfer; the Internet, client-server architecture, and database system are typical of the kinds of technologies invented in this era. The second era, an age of convenience, is now under way and will last about twenty years. It focuses on maximizing customer satisfaction by catering to individual preferences, desires, and tastes. It is pervaded by a concern for making business convenient for consumers. Collaboration, workflow, and just-in-time manufacturing technologies will dominate in this era. The third era will focus on the reputation and identity of the business in a world of many imponderables. These uncertainties will be the result of cheap, fast, global communication and the increased confusion people will experience as they pursue their own conveniences. The technologies of this era will focus on measuring aspects of identities, projecting them, and building stable social realities.

Conclusion

The ARPANET began operation in 1969 with four nodes as an experiment in resource sharing among computers. It evolved into a world-wide research network of over 60,000 nodes, influencing the design of other networks in business, education, and government. It demonstrated the speed and reliability of packet-switching networks. ARPANET protocols have served as the models for international standards; they also spawned the Internet, which numbers at least 20 million nodes. And yet the significance of the Internet lies not in its technology, but in the profound alterations networking has produced in human practices. Network designers and engineers must learn to ply their trade in a world of many clashing discourses and conflicting interests, where technologies will be judged more for their ability to enable shifts of social power and less by the traditional criteria of utility and efficiency.

References

1. P. Baran. 1964. Collected technical reports are available from http://www.rand.org/publications/RM/baran.list.html.

2. G. Bell and J. Gray. 1997. “The revolution yet to happen.” In *Beyond Calculation: The Next 50 Years of Computing* (P. Denning & R. Metcalfe, eds.). Copernicus, and imprint of Springer-Verlag.

3. T. Berners-Lee. 1996. "The Web Maestro: An Interview with Tim Berners-Lee." *Technology Review*. July.

4. T. Berners-Lee. 1996. "WWW: Past, Present, and Future." *IEEE Computer 29*, 10 (October). 69-77.

5. J. Birnbaum. 1985. Toward the domestication of microelectronics. *Communications of ACM 28*, 11 (November). 1225-1235.

6. P. Denning. 1989. Worldnet. *American Scientist 77*, 5 (September-October). 432-434.

7. P. Denning. 1989. "The ARPANET after twenty years." *American Scientist 77*, 6 (November-December). 530-534.

8. P. Denning. 1989. The Internet worm. *American Scientist 77*, 2 (March-April). 126-128.

9. P. Denning. 1997. "A new social contract for research." *Communications of ACM 40*, 2 (February).

10. P. Denning. 1997. "How we will learn." In *Beyond Calculation: The Next 50 Years of Computing* (P. Denning & R. Metcalfe, eds.). Copernicus, and imprint of Springer-Verlag.

11. D. Englebart. 1962. "Augmenting human intellect: A conceptual framework." Stanford Research Institute (October).

12. F. Flores. 1997. "The leaders of the future." In *Beyond Calculation: The Next 50 Years of Computing* (P. Denning & R. Metcalfe, eds.). Copernicus, and imprint of Springer-Verlag.

13. A. Groves. 1996. *Only the Paranoid Survive*. Bantam Doubleday Dell.

14. K. Hafner and M. Lyon. 1996. *Where Wizards Stay Up Late: The Origins of the Internet.* Simon & Schuster.

15. R. Kahn, ed. 1978. Special Issue on Packet Networks. IEEE *Proceedings 66*, 11 (November).

16. W. Mitchell and O. Strimpel. 1997. "There and not there." In *Beyond Calculation: The Next 50 Years of Computing* (P. Denning & R. Metcalfe, eds.). Copernicus, and imprint of Springer-Verlag.

17. T. Nelson. 1987. *Literary Machines*. Published by the author.

18. J. Quarterman and J. Hoskins. 1986. Notable computer networks. *Communications of ACM 29*, 10 (October). 932-971.

19. J. Quarterman. 1990. *The Matrix: Networks and Conferencing Systems World-wide*. Digital Press.

20. C. Stoll. 1995. *Silicon Snake Oil: Second Thoughts on the Information Highway*. Doubleday.

21. T. Winograd. 1988. Introduction to Language/Action Perspective. *ACM Trans. Office Information Systems 6*, 2 (April). 83-86.

22. T. Winograd and F. Flores. 1987. *Understanding Computers and Cognition*. Addison-Wesley.