

THE SCIENCE OF COMPUTING

THE ARPANET AFTER TWENTY YEARS

Peter J. Denning

ABSTRACT: The ARPANET began operation in 1969 with four nodes as an experiment in resource sharing among computers. It evolved into a worldwide 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. Its protocols have served as the models for international standards. And yet the significance of the ARPANET lies not in its technology, but in the profound alterations networking has produced in human practices. Network designers must now turn their attention to the discourses of scientific technology, business, education, and government that are being mixed together in the milieu of networking, and in particular the conflicts and misunderstandings that arise from the different world views of these discourses.

In the fall of 1969, the first node of the computer network known as the ARPANET was installed at UCLA. By December of that year, four nodes were operating, by 1971 fifteen nodes, and by 1973 thirty-seven nodes. Today, this network has evolved into a collection of networks called the Research Internet spanning over 60,000 nodes. Worldwide networking, including fax over telephone lines, now embraces millions of nodes. Although we may be inclined to interpret these developments by saying that Worldnet is emerging, it is more accurate to say that Worldnet is here and our awareness is emerging (1).

The changes in our use of computers begun 20 years ago are, in retrospect, nothing short of revolutionary. I would like to discuss the origins of the ARPANET, reflect on its influence our practices, and speculate about the issues that will be faced by designers of networks.

The ARPANET story begins in the late 1950s, during the early development of intercontinental ballistic missiles. The Department of Defense was concerned about the ability of US forces to survive a nuclear first strike, and it was obvious that this depended on the durability of our communication network. Paul Baran of the Rand Corporation undertook a series of investigations of this question, concluding that the strongest communication system would be a distributed network of computers having several properties: it would have sufficient redundancy so that the loss of subsets of links and nodes would not isolate any of the still-functioning nodes; there would be no central control; signals would traverse a series of nodes from source to destination, the exact route being determined by the set of working nodes and links at a particular time; and each node would contain routing information and could automatically reconfigure that information within a short time after the loss of a link or

node. Further, Baran proposed that messages be broken into units of equal size and that the network route these message units along a functioning path to their destination, where they would be reassembled into coherent wholes. Baran's reports became public in 1964.

Meanwhile, Larry Roberts of MIT's Lincoln Laboratory, enticed by visions articulated by J. C. R. Licklider of the Defense Department's Advanced Projects Research Agency (ARPA), decided to devote himself to realizing the potential of networking: sharing the resources of one computer easily and economically with another. Inspired by Licklider and Roberts, Donald Davies of the National Physical Laboratory in England proposed in 1965 a computer network using telephone trunk lines ranging in speed from 100 kilobits per second to 1.5 megabits per second, messages broken into "packets" of 128 bytes each, switching computers that could process 10,000 packets per second, and special interface computers that would connect mainframe "hosts" to the packet network without requiring alterations in the hosts' operating systems. From his own experiments in 1966 with direct-dialed telephone links between computers, Roberts concluded that the packet-switching architecture of the proposals of Baran and Davies would be required to overcome slow and unreliable telephone circuits and would, moreover, be cheaper. Leonard Kleinrock of UCLA had produced analytic models of packet-switched networks that could be used to guide a design.

At the same time these developments were taking place, Robert Taylor, who had succeeded Licklider at ARPA, had become interested in computer networking from a different perspective. Previous ARPA projects had created a variety of powerful computational centers at different institutions. Each had established its own user community and had become a potential 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 network to connect the centers that would be fast and robust under failures and that would work with the operating systems of the many vendors whose computers were in use at the various 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 Neumann (BBN) to build the first interface message processors (IMPs), computers as proposed by Davies to translate between messages and packets. 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.

The ARPANET began operation in 1969 with four nodes; it has evolved into a worldwide research network of over 60,000 nodes

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

By the mid 1970s, it was clear that research networking was growing rapidly and that ARPANET would need to connect to other networks. This realization inspired a reworking of the original end-to-end protocol, which was called NCP (network control protocol), producing in its place a matched pair of protocols called TCP (transport control protocol) and IP (internet protocol). IP would be responsible for routing packets across multiple networks and TCP for converting messages into streams of packets and reassembling them into messages with few errors despite loss of packets the underlying network. These two protocols provided highly reliable end-to-end communication in a network of networks, eventually exercising a significant influence on the protocols now approved for worldwide use by the International Standards Organization.

Various "community networks" began to appear around 1980; notable examples are BITNET connecting IBM machines, CSNET connecting computers in the computer science research community, USENET connecting UNIX sites by telephone, and internal networks within companies such as IBM and DEC. In 1984 the National Science Foundation started connecting its supercomputing centers with a high-bandwidth network called NSFNET, which now serves as a backbone for the community networks and the Research Internet. As the Internet grew, the original method of naming nodes became unwieldy; a hierarchical naming system that allowed each "domain" to select its own internal addresses was introduced in 1984.

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 today. Some X.25 service has been available in the US since the early 1980s.

Networking has had an effect on the world more profound than that of the more spectacular and expensive Apollo moon missions

If you are interested in reading more about these developments, I recommend a special issue of IEEE *Proceedings*, which contains 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 (2). I also recommend an article by John Quarterman and Josiah Hoskins (3).

The ARPANET was officially disbanded earlier this year, but because the Research Internet already taken over its functions, few users noticed. The current administration and Congress are planning a further expansion of networking through an organization to be called

the National Research and Education Network (1).

In the remainder of this essay I would like to consider these events in a way that reveals why this twenty-year-old invention, networking, should have had such an effect on the world, an effect more profound than that of the more spectacular and expensive Apollo moon missions. My analysis is intended to give some guidance as we consider how to design networks in the future, observing the progress of a discourse from its birth in the declarations of a few people through major shifts in practices that they could not have anticipated. (A discourse here means a nearly transparent mode of thinking, speaking, and acting that transcends individuals and extends over a long period of time.) I will illustrate such a progress with three examples.

First, suppose we brought back Henry Ford for a look at today's automobiles. He would be little surprised by changes in design: cars still have four wheels, steering, front-mounted 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, annual sales of millions of cars, the interstate highway system and the intracity systems in places like Los Angeles, nationwide trucking, cars as status symbols, multicar families, state licensing of drivers, rush hours, traffic congestion reports on the radio, and much more (4).

Second, suppose we brought back Alexander Graham Bell to see our telephones. He would be little surprised by the design of instruments and switching systems -- handsets, carbon microphones, dialing mechanisms, crossbar switchers, operator services, and the like. But he would be greatly surprised by the changes in human practices that have grown up around the telephone -- telephones in every home, office, and hotel room, car phones, phone booths, international direct dialing, news services, multinational corporations, electronic fund transfers, telemarketing, ordering by phone, fax, telephone pornography, and much more.

Third, suppose we brought back Thomas Edison. He would be little surprised by what we would present to him in the design of light

bulbs and electric generators. But he would be greatly surprised by the changes in human practices that have grown up around electricity - international distribution of power, total dependence on electric power in the developed countries, radio and television industries, electronics, computers, and much more.

A careful examination discerns five major stages in the progress of a technological discourse (my analysis is guided by conversations with Fernando Flores and a paper by Joel Birnbaum (5):

- declarations
- prototypes
- tools
- industries
- widespread practices

The passage through these stages is not smooth and regular, but rather is best characterized as a *drift* affected by many events that make it impossible to predict what the practices will ultimately be. The time scale for the drift from the first to the last stage is long -- one or two generations, or 20 to 50 years.

The ARPANET illustrates the drift within the computer science research community, with new stages at roughly five-year intervals. Around 1965 the first design proposals were put forth. By 1970 the first prototypes were operating in the early ARPANET. The first tools were in place in 1975; these included electronic mail, file transport, remote login, and telephone login. Industries were emerging by 1980: community networks such as CSNET, BITNET, and USENET and also the commercial networks such as GTE Telenet. By 1985 widespread practices had evolved around the network, such as linking of workstations on every desk through local networks into the Research Internet, alterations of office practices around workstations and in word processing, shifts in the responsibilities of secretaries, collaborations over networks, setting up of electronic bulletin boards, and attacks by intruders, worms, and viruses.

There have been major surprises as well that altered the drift's direction. 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. Nonetheless, at the founding of CSNET in 1980,

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 facilities. 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 as major disasters, as we witnessed in the Internet Worm incident of November 1988 (6). High speed personal workstations became increasingly cheap and powerful and are now individual nodes in the networks. Electronic publication has emerged as an industry in its own right, placing heavy demands on networks to move manuscripts from authors to editors to printers. Facsimile transmission -- the now ubiquitous fax -- also emerged as an independent industry. By combining the widespread practice of sharing paper documents with the wide reach of the telephone network, it facilitates international coordination of actions despite wide differences in time zones and allows exchanges between clients who can read but not speak each other's languages. Few workstation or telephone designers dare exclude fax. Thus the significance of the ARPANET and its derivatives lies not in the networking technology but in the fundamental shifts in human practices that have resulted -- the new discourse invented by a few individuals in the mid-1960s.

The significance of the ARPANET and its derivatives lies not in the networking technology but in the fundamental shifts in human practices that have resulted

A central question that arises from the interpretation of the ARPANET as part of a new technological discourse is, what other discourses will inevitably come together in the arena created by networking, and how can the design of networking in the future accommodate them? I see four major discourses: scientific technology, business, higher education, and government.

The discourse of scientific technology looks ahead to a high-tech world of scientific research, featuring by the year 2000 supercomputers with 1 to 10 teraflops performance, networks with 1 to 3 gigabits bandwidth, portable computers and smart cards linked by radio into the world network, and in every workstation 3D animated graphics, high definition TV screens, audio, video, fax, voice input, and speech output. This discourse has a darker side. It views the world, including people, as a collection of resources to be acquired, used, optimized, and discarded when no longer needed. It views situations, including those that involve the human condition, as "problems" for which technological and procedural "solutions" are to be found; unable to admit that some problems may be insoluble, this discourse labels such problems as "intractable" but ultimately solvable given sufficient knowledge and resources.

Future design questions will not be purely technological, but thoroughly intertwined with the human practices that arise around networks of computers

The discourse of business is concerned with attitudes and practices for working together, the acquisition of power in the marketplace, the completion of transactions over distances large and small. It talks about global markets, personalized products and services, a worldwide information infrastructure consisting of networks and workstations, a conviction that business success implies mastery of networking, a concern for the effects of rapid communication on business practices (e.g., chaotic change), and a concern for how networking and computing will affect ordinary business practices. It focuses on financial performance, market share, quality of product and service, and productivity.

The discourse of higher education holds that knowledge encompasses a structured set of information, teaching is the transmission of a subset of this information into the minds of students, and research is the discovery of new information already existing in the world. It has

institutionalized a system of rewards that reflects the high value it places on individual (academic) freedom and accomplishment: emphasis on research over teaching, a concern to identify the unique personal contribution of each participant in joint work, a focus of research within rather than across disciplines, a distrust of students collaborating on homework, and disregard for skills needed to work effectively in organizations. This discourse is baffled by complaints of students who say that they graduate without practical competence in their disciplines, without the ability to learn new subjects, and without a sense that their research is relevant to the world.

The discourse of government includes a concern for competing in international markets, maintaining the national research lead and developing a faster manufacturing capability, a desire to be world leader in all areas, and a suspicion of multinational cooperative ventures.

By recognizing that these discourses will mix together in the world of networking, we can see that opportunities for better design will arise from our learning the concerns and blind spots of each discourse. We can also anticipate the conflicts and misunderstandings that may arise.

I would like to close with three examples of such conflicts. The first concerns the role of electronic mail. Business users see electronic mail as a generalization of fax: they talk about machines hooked to their telephone lines that allow the exchange of messages with an addressing protocol like "mail to name@phone-number." The business view is rooted in two widespread practices: the use of telephones, which are everywhere and understood by everyone, and the sharing of paper documents, as witnessed by the phenomenal success of Federal Express. On the other hand, network engineers see fax as the next technology to integrate into electronic mail. Electronic mail like that in the ARPANET -- that is, text files that can be exchanged and edited -- is not as deeply ingrained in business practices. In my opinion, network designers who fail to take into account the power of existing business practices will be surprised to find that fax-inspired technology will win out over ARPANET-inspired technology. Those who do make allowances for business practices will devise means of

combining the best features of fax with those of electronic mail.

A second example concerns network vulnerabilities. Business, science, and government users have a deep concern for the integrity and privacy of information entrusted to computers and databases. They worry about intruders, worms, and viruses. In contrast, concern for academic freedom at the universities has produced muted public statements that seem to indicate a lack of willingness to take measures to foster respect for network security in students. Moreover, the scientific technology discourse inclines those who participate in it to argue that protocols for authentication, secrecy, and error recovery provide a complete basis for a "network immune system"; this discourse is blind to the need for introducing new practices in a world where widespread cooperation is essential. Network designers will have to reconcile these divergent concerns.

My third example of possible conflict involves trust. Many managers in government and business are concerned that employees not

abuse their privileges in computer systems by releasing organizational information assets to outsiders; to allay this feeling of distrust and their fears of external attacks, managers propose increasingly complicated access controls and auditing mechanisms and call for "trusted computer systems." These same mechanisms appear to employees as means of surveillance and monitoring, an institutionalization of the distrust the mechanisms are supposed to render unnecessary, and even a deprivation of dignity. If human practices external to a system of computers and networks generate distrust among those who must coordinate action, how can monitoring, auditing, and access control mechanisms restore trust?

These three examples illustrate the types of questions network designers must face in the years ahead, questions that are not purely technological but are thoroughly intertwined with the human practices that arise around networks of computers.

References

1. P. Denning. 1989. Worldnet. *American Scientist* 77, 4 (September-October). 432-434.
2. R. Kahn, ed. 1978. Special Issue on Packet Networks. *IEEE Proceedings* 66, 11 (November). .Re J. Quarterman and J. Hoskins. 1986. Notable computer networks. *Communications of ACM* 29, 10 (October). 932-971.
3. T. Winograd. 1988. Introduction to Language / Action Perspective. *ACM Trans. Office Information Systems* 6, 2 (April). 83-86.
4. J. Birnbaum. 1985. Toward the domestication of microelectronics. *Communications of ACM* 28, 11 (November). 1225-1235.
5. P. Denning. 1989. The internet worm. *American Scientist* 77, 2 (March-April). 126-128.