

Why the Arpanet Was Built

Stephen J. Lukasik

Georgia Institute of Technology

The who, what, when, and how of the Arpanet is usually told in heroic terms—Licklider's vision, the fervor of his disciples, the dedication of computer scientists and engineers, the work of graduate students, and so forth. Told by one of the key actors in this salient part of US and Internet history, this article addresses why the Arpanet was built.

The who, what, when, and how of the Arpanet is usually told in heroic terms. J.C.R. Licklider's vision, the fervor of his disciples, the dedication of computer scientists and engineers, the work of graduate students, and the attraction of the Arpanet to early participants carries with it a sense of inevitability. But why the Arpanet was built is less frequently addressed. Writing from the viewpoint of the person who signed most of the checks for Arpanet's development, in this article, I detail the rationale for investing US Department of Defense resources for research and development of the first operational packet-switched network. The goal was to exploit new computer technologies to meet the needs of military command and control against nuclear threats, achieve survivable control of US nuclear forces, and improve military tactical and management decision making. Although not central to the decision to pursue networking, it was recognized that these capabilities were common to non-defense needs.

The relationships among the various events recounted here can be better understood in terms of their organization into six sequences of events. ARPA's existence and sole purpose was to respond to new national security concerns requiring high-level visibility. In this case, it was the command and control of military forces, especially those deriving from the existence of nuclear weapons and deterring their use. The translation of these needs to programs and capabilities were the joint responsibility of the technical community and the agency's managers. This included both network theory and engineering the critical experiments central to making practical progress. The next steps were dependent on the transfer of the R&D results to users, both in the DoD and related US government

agencies, as well as to international research and operations communities.

Numerous people made substantial contributions to this endeavor. At the very beginning, networks of people (US friends and foreign adversaries) provided assistance and pressure to succeed. In the development mid-phase, the number of people contributing their expertise, covering a number of technical fields, grew rapidly. And as the bounds of both the research and user networks expanded beyond the US, the people participating in networking continued to grow exponentially.

The individuals identified in Figure 1 and in the references provide chronological tie-points in a network of events and people, some tightly coupled, some quite loosely coupled, and some completely accidental. Events are tied to the textual discussion with letters corresponding to those bracketed in the text.

The Cold War

The first step on the journey to networking computers was taken on 29 August 1949 when the Soviet Union detonated its first nuclear weapon, a major event in the postwar confrontation between the US and the Soviet Union.^{1,[A]} Because aircraft were the only way to deliver a nuclear weapon to a distant target at that point, air defense became an immediate US concern.

The US took two actions. In 1948, the Air Force had already studied the need for an air defense early-warning network of radar stations and had an 85-station network of limited capability underway. In December 1949, the Air Force Science Advisory Board convened an Air Defense System Engineering Committee (ADSEC), chaired by Massachusetts Institute of Technology professor




Figure 1. Chronological development of the Arpanet. The points highlight when individuals, organizations, and random chance intervened. Events are tied to the textual discussion with letters corresponding to those bracketed in the text.

George Valley, with the intention of recommending a system that would overcome the many deficiencies of the ad hoc radar network hurriedly being assembled (tellingly called *Lashup*).² The second was on 31 January 1950 when President Truman decided to accelerate the development of fusion weapons to deter Soviet use of their less powerful fission weapons.³

MIT's Radiation Laboratory had been a major center of radar development in WWII and some of the participants in that work, in the now-renamed Research Laboratory for Electronics, naturally turned to the new problem. Current air defense technology was the same as in WWII. Data from radars scanning air sectors was posted by people with grease pencils writing in reverse on transparent boards. From this display, tactical air defense commanders determined bomber routes, projected positions, available response time, and the targets to be attacked. Decisions included alerting anti-aircraft units, dispatching fighter aircraft and guiding them to intercept the incoming bomber streams, providing warning to civilians to take shelter, and alerting civil defense workers.

In early 1950, ADSEC recommended automating this process. Radar data would be digitized, transmitted over telephone lines, and analyzed by the new digital computers being developed. The MIT Whirlwind computer, funded by the Office of Naval Research, was seen as pointing to the technology needed for such a system. Although the calculations involved in correlating and displaying radar data were simple, slide-rule level, the mass of data and the need for real-time display for the command and control of response forces dictated something more powerful.

The urgency of US air defense increased when North Korea invaded South Korea on 25 June 1950 and when China intervened on 26 November of that year, thus bringing the specter of nuclear war with the communist world even closer.

Discussions between the Air Force and MIT to establish an Air Defense Development Center in the mode of the Radiation Laboratory proceeded. MIT convened a study group, Project Charles, consisting of its faculty and a number of outside experts to vet the proposal. The result was to approve the idea for an off-campus classified Project Lincoln,

named for its proposed location near Lincoln, Massachusetts. A significant demonstration occurred on 20 April 1951 when a “bomber” was intercepted in real time by a “fighter” using actual radar data transmitted over telephone lines and analyzed by Whirlwind. The result of this effort was the Semi-Automatic Ground Environment (SAGE) air defense system.

The launch of an earth satellite by the Soviet Union on 4 October 1957 was another major development in the Cold War confrontation.^[B] If the USSR could launch satellites of that size, they could also deliver missiles armed with nuclear warheads to the US. Besides initiating a space race to accompany the advancing arms race, the event resulted in rethinking the US policy for science and technology and in substantial increases in research support for universities and the education of students in science and engineering. Business-as-usual was put aside. The best minds in the country were enlisted in the equivalent of the WWII Manhattan Project, and the national security establishment went into crisis mode.

One US response, DoD Directive 5105.15 dated 7 February 1958, established the Advanced Research Projects Agency (ARPA, now DARPA.) The directive is marvelously general: “to be responsible for the direction or performance of such advanced projects in the field of research and development as the Secretary of Defense shall, from time to time designate by individual projects or by that category.” Over the years, this has come to be interpreted as seeing that technological surprise would never be repeated, a rationale and guiding principle the agency has unwaveringly followed.^[I] What preventing technological surprise means, in practice, is that the agency does the surprising, not the other way around.

Over the next 10 years, ARPA developed a modus operandi that still distinguishes it from other parts of the federal government. It is quite small: typically approximately 150 employees, with only modest internal structure. Program managers, who typically are recruited for only a few years, are masters of their subjects, equal to the most expert specialists with whom they work. The agency does not have a large budget. When adjusted for inflation, the current one is only a factor of two larger than its first budgets. What distinguished the agency from other parts of government was its freewheeling style,

flexibility in management and contractual arrangements, enforcement of the highest level of excellence from its contractors, and openness to technical ideas from all directions.⁴ It sought results that redefined how one looked at the problems it tackled—game-changing solutions to what today are called “DARPA-hard” problems. Modest evolutionary improvements were left to others.

ARPA’s early years were characterized by “Presidential Issues,” assignments coming to it directly as the result of White House concerns about its most critical technical concerns. ARPA initiated programs in large boosters to close the missile gap highlighted in the 1960 Presidential election; reconnaissance, communication, navigation, meteorological, and missile launch warning satellites; ballistic missile defense; a worldwide satellite tracking system; establishing a scientific basis in seismology for negotiating and monitoring nuclear test ban treaties; counterinsurgency and unconventional warfare; and creating an interdisciplinary science of materials from the separate disciplines of physics, chemistry, metallurgy, ceramics, and engineering. Shorter-term efforts were in solid propellant chemistry, energy conversion, Presidential protection following the Kennedy assassination, and other matters that drew on its contracting flexibility, speed, expertise, and innovative approaches.

The Command and Control of Military Forces

ARPA’s history and character lies at the heart of networking. To appreciate why things are as they are, a deeper look at their origin is needed. The history of ARPA for the period 1958–1974 covers its development and has been used extensively to aid in reconstructing the events I recount here.⁵

Following the agency’s establishment in 1958 as the presumed savior of dysfunctional US national space efforts in the DoD, it was decided that the US national space effort should not be lodged in the DoD.^[K] Instead, in 1959 it was moved to a new nondefense agency, NASA, created around a greatly expanded National Advisory Committee on Aeronautics (NACA). By this time, the White House was having second thoughts after ARPA’s first director, Roy Johnson, publicly disagreed with it over the civil space decision. Similarly, the scientific establishment that had initially promoted the agency’s formation now decided it would rather have its funding from other sources

than the DoD.^[H,J] ARPA space programs were transferred either to NASA or returned to the military services.

The services were relieved. They had been concerned from ARPA's beginning that a strong R&D arm under the direct control of the Secretary of Defense through his immediate staff in the Office of the Secretary (OSD) would threaten their weapon development prerogative. The joint chiefs of staff (JCS), reflecting this position, opposed the agency's formation.^[H] Now the services returned to the attack with renewed vigor, arguing that ARPA was never intended to be a permanent agency, but simply a now-completed transitional phase in the reorganization of US space efforts in response to the Soviet satellite launch. The appointment of the second ARPA director, Army Brigadier General Austin Betts, was seen as a holding action until the agency's fate could be decided.^[J] In that uncertain period, ARPA took on "filler" tasks that relied as much on its statutory funding flexibility as on its technical capabilities.

The new Secretary of Defense, Robert McNamara, and his new Director of Defense Research and Engineering (DDR&E) Harold Brown decided to retain ARPA as part of OSD. Betts was replaced by Jack Ruina, earlier an MIT professor, as the agency's third director in 1961.^[J] Its project in Command and Control Research, the foundation for all its subsequent information technology research, was one of the assignments of convenience during the Betts period. An expensive backup computer at the System Development Corporation (SDC) for the SAGE air defense system, the AN/FSQ 32XD1A, was declared surplus in 1963 as concern shifted to missile defense.^[M] The DDR&E was interested in war gaming as the DoD sought to understand the dynamics of nuclear war to establish requirements for weapon systems development and procurement. SDC found itself not only with a spare computer, but also with a decreasing role in air defense software development for the Air Force. So the problem of the large valuable surplus computer, its software, and the people to run it was turned over to ARPA to address OSD's system requirements needs.

With the merging of the Departments of War (the Army) and Navy, the formation of the Air Force as a separate service, and the management of all three under a civilian Secretary of Defense by the National Security Act of 1947, tensions that persist to the present

emerged. The armed services' command and control functions were central, for they provided the basis for what military forces should and should not do. OSD welcomed the opportunity to gain deeper insight into the issues with the intention of making military capabilities more useful to the President and the Secretary of Defense. Unlike the other marriages of convenience in the early days of ARPA, the command and control assignment took on a life of its own. The original ARPA program description called for the agency to "support research on the conceptual aspect of command and control and to provide a better understanding of organizational, informational, and man-machine relationships."⁶

The event that turned this small assignment into the major national capability that evolved from it was when Jack Ruina hired J.C.R. Licklider (Lick) in October 1962 to run the program.^[N] A psychologist working on human factors in WWII, a participant in the Project Charles study that laid out the conceptual foundations for the SAGE networked air defense system, a professor of psychoacoustics at MIT, and a computer researcher at Bolt Beranek and Newman (BBN) before going to ARPA, Lick saw the command and control problem as having two parts: the machine processing of information and the presenting of that information to humans in a form suitable for use in making decisions.

Another incident played a role in formulating the ARPA program. The Defense Science Board, noting the success with which the physical sciences were being applied to national security problems, pushed for some DoD attention to the behavioral sciences beyond the narrow focus in the department on human factors. An appropriation for work in the behavioral sciences, small in size but large in establishing legitimacy for such work, was received by ARPA near the end of the Ruina directorship. Licklider's office provided a convenient place from which to administer it. Thus, Lick had both a technical focus on information technology and a charter that could cover the decision side of his integrated vision.

The ARPA environment provided ideal ground for making fundamental advances in command and control. Although previously the agency's work was carefully watched by the Secretary of Defense, the White House, and the President's Science

Advisory Committee, in this endeavor Lick was left alone. In his words,

[The Director] seemed too busy, he was just relieved to get somebody to run the office. . . . I talked with him periodically [and] he would make suggestions about directions of things, but pretty much let me do what I wanted to do.

There was a belief in the heads of a number of people—a small number—that people could become very much more effective in their thinking and decision making if they had the support of a computer system, good displays, and so forth, good data bases, computation at your command. It was the kind of image that we were working toward the realization of . . . It really wasn't a command and control research program [anymore]. It was an interactive computing program. And my belief was, and still is, you can't really do command and control outside the framework of such a thing . . . of course, that wasn't believed by people in the command and control field.

What Lick and his successor, Ivan Sutherland, did was initiate a number of activities that ultimately served both purposes well. Lick expanded the command and control research program at SDC into a high-quality computer science program based around Centers of Excellence. MIT was one such center, and one step in the direction of computer support for user interaction with data was time-sharing. A centerpiece was Project MAC—variously understood to be an acronym for Machine-Aided Cognition, Man and Computer, and Multiple-Access Computer.

A contemporary perspective on time-sharing was that the speed of computers was so great compared to the speed with which people think that they could be shared, apparently simultaneously, among several users. But users could only interact with data in their host computer and only with other people at their same institution who were connected to that computer. One could have several terminals connected to different computers, but the transfer of data from one computer to another was awkward. The need to aid cognition by facilitating the interaction of people and data as broadly as possible—what Lick called the *Great Intergalactic Network*—was being explored by a number of people in different places but was not yet fully articulated.^{7,8} This combination of motives, interactive cognition, and management efficiency persisted through the Information Processing Techniques (IPT) program in different areas.

Lick made several direct efforts at networking, although no one quite knew how to turn networking into an academic research area, and certainly the military phrase “command and control” was not part of the explanation. One exploratory effort at UCLA looked into networking separate computers to function as a virtual parallel processor. A much larger effort was to capitalize on the presence of three large IBM mainframes on the UCLA campus. The Western Data Processing Center (WDPC) in the Business School served a number of business schools in the western part of the US to help introduce computation into their curricula. A second was in the Health Sciences School and was supported by the National Institutes of Health. The third served the rest of the UCLA community, including the ARPA-supported work. The plan was for a large project, with IBM 360/40s to serve as the network front-ends for the host mainframes. The project, which was to include substantial IBM contributions in personnel and hardware as well as ARPA support, was submitted for funding in February 1965. It collapsed when UCLA decided not to renew the basic WDPC contract with IBM. By this time, Sutherland had replaced Lick as director of IPT. A noteworthy feature of Sutherland's thinking was that the network was to be robust, on the ground that a DoD-like network should be able to “survive a bombing raid.”⁹

Starting from this state of development, and fully implemented by his successors, Lick initiated a major change in thinking about networked computers, a change much broader than simply connecting computers, a capability implicit in some defense systems and applications at the time. Existing computer networks were functionally specific networks, special-purpose networks designed for a particular application. They were in that sense much like the special-purpose computers of that period, or application-specific chips today, optimized to do certain things. What Lick conceived was a general-purpose network that could be used for as many decision-support purposes as its users had the imagination to conceive. To Lick, “network” had three meanings: the network of hardware, wires, and software; the man-machine network, where each performed functions for which each was best equipped to reach a larger goal; and the network of people, with their skills, experiences, and information resources that could be focused on larger cognitive goals requiring interaction to

solve problems too large or complex for an individual.

Not only did Lick see information technology and behavioral and cognitive science issues as connected, but in his view research and military needs also had much in common. He notes in his April 1963 memo,

I am hoping that there will be, in our individual [research] efforts, enough evident advantage in cooperative programming and operation to lead us to solve the problems and, thus, to bring into being the technology that the military needs. . . . many of the problems will be essentially as important, in the research context as in the military context.

It was more a matter of luck than management that the JCS and the armed services did not realize the potential implications of ARPA's work in command and control. If they had seen the ARPA effort as eventually changing the way military forces would operate and communicate based on, to use the current jargon, *net-centric warfare*, they would likely have objected to so fundamental an intrusion into the essence of their missions.

Even before the first Soviet nuclear weapon test, the newly established Air Force saw the need to address its strategic role in the projection of US power. In 1948, at the beginning of the Cold War, Soviet aims in Europe were becoming clear. The Air Force recognized the need to enlist the US scientific and technical community. One part of this process was the establishment of the RAND Corporation.¹⁰ This new "think tank" looked into issues such as bombing accuracy, air-to-air combat, and the operation of air defense centers. This work, involving both mathematical analysis and analog simulation facilities, was done in the System Research Laboratory.¹¹ SRL eventually undertook software production for the SAGE system but had to spin it off to a new organization, the SDC, when it became too large an effort for RAND to manage. RAND did, however, continue technology efforts such as a pen tablet and handwriting recognition to facilitate man-machine interaction.

When Lick left ARPA at the end of 1964, the integrated character of the original command and control program was lost and the Behavioral Science Command and Control Research office was divided in two. The information technology side of command and control stayed with the now-renamed IPT office, while the behavioral charter shifted to

The technical solution to avoiding decapitation involved new ways of routing and switching in decentralized and distributed communication systems that could survive damage from an attack.

an independent Behavioral Science Office. The latter office did not initially focus on decision making. The head of the Behavioral Science Office had been in ARPA's counterinsurgency field office in Thailand and moved the program into anthropology and political science. Later, when the office did focus on decision making, it was in decision theory. An early application supported political negotiations for transfer of the Panama Canal from US to Panamanian sovereignty. The new behavioral science effort was intended to extend the ARPA research approach to other parts of the social sciences, a theme that continued in ARPA for a number of years, despite the controversy the DoD's presence engendered.

The Arpanet from Concept to Realization

By 1959, U2 flights over the USSR established that the Soviets were deploying intercontinental ballistic missiles (ICBMs). An alleged missile gap was a prominent issue in the 1960 Presidential campaign. The concern was that US air defenses were powerless against the high speeds and consequently short decision times implied by ICBMs. ARPA was directed to examine technical options for a ballistic missile defense system to avoid a situation prone to instability from operational errors and misreading intent by either side. The 1962 Cuban missile crisis was a case in point.

In addition, the DoD was concerned with how the national command authority could survive a first strike and still have sufficient connectivity to order a retaliatory second strike.

Were the Soviets to believe a decapitating attack was impossible, the logic went that they would be deterred from initiating it.¹² The technical solution to avoiding decapitation involved new ways of routing and switching in decentralized and distributed communication systems that could survive damage from an attack.

RAND was at the forefront of these policy and technology studies. Paul Baran joined the organization in 1959 from Hughes Aircraft and synthesized many ideas for reliable priority communication into a proposal for a distributed system based on what he called “message blocks.”^[C] The theoretical work was largely completed by 1962 and published in 1964.^{13,[D,E]} The fifth report of this series summarizes 11 routing concepts examined between 1952 and 1961. The ninth report on security, secrecy, and tamper-free considerations addresses architectural issues of network security—deficiencies in the Internet we are only now coming to regret. As a high-water mark in thwarting malicious users, Baran notes that his system “assumes a potential infiltration by enemy agents having access to portions of the system and the cryptographic keybases.”

RAND made a formal recommendation to the Air Force in 1965 to deploy Baran’s system. Based on an outside review chaired by Mitre, the Air Force concurred.^[F] At that point, the organizational structure of DoD intruded. Under the Defense Reorganization Act of 1949, the Defense Communication Agency (DCA), not the Air Force, had the responsibility of developing, deploying, and managing DoD communication systems. DCA at that time was unfamiliar with digital technology. It was assisted by AT&T, a major provider of communication services and deeply committed to analog voice circuits and their switching. The Air Force decided giving the Baran/RAND proposal to DCA for implementation was a plan for disaster and the proposal was withdrawn.^{14,[G]}

The system Baran and his team proposed looked remarkably like the Arpanet eventually did, but it had many features still sorely lacking in the public Internet 40 years later, such as the redundancy to withstand heavy attacks yet fail gracefully as links were severed; high reliability; security of the information within it; and the ability to defeat malicious users inside and outside the system. Baran’s network was intended to connect the national command authority to US nuclear forces. How well the system would

have scaled, how easily it could have been adapted to new technology, and how well it would have performed in a nuclear attack will never be known. Baran’s work is impressive for the scope of the design specifications. It was the first packet-switched network “invented.” ARPA’s later effort was a second pass by DoD at packet switching. Its network efforts met with greater acceptance, largely because they were not tightly coupled to the defense command and control community. This decoupling was reciprocal. The bulk of the ARPA contractors involved were unaware of, nor did they particularly care, why the DoD was supporting their research.¹⁵

One of the networking efforts independent of the earlier DoD work was that of Len Kleinrock, a graduate student in the MIT Electrical Engineering and Computer Science Department. In his May 1961 PhD thesis proposal, “Information Flow in Large Communication Nets,”^{16,[P]} he examined questions of scalability, performance evaluation, control, routing, and related issues in distributed networks. His thesis was completed in 1962 and was published in 1964,^[Q] coincident with the publication of Baran’s 1960–1962 work at RAND.¹⁷ Like Baran, Kleinrock recognized that information would have to be transmitted as blocks of bits. Baran and Kleinrock later assisted ARPA in synthesizing the current technical thinking into the Arpanet design.¹⁸

The next step in realizing Lick’s vision occurred when Bob Taylor joined the agency from NASA in 1965.^[O] Lick’s university Centers of Excellence were central to the development of his approach to command and control: time-sharing; flexible and convenient interaction with computers through text, graphics, pointing, and handwriting; and artificial intelligence. Taylor added an additional push into networking, and later IPT directors added language, speech, and image understanding to the mix that enabled the power of computers to expand into tasks where humans could be augmented or replaced.

Another illuminating story of the ARPA management environment is Taylor’s account of securing “seed money” to study networking—some \$1 million in 1965—on the basis of a 20-minute discussion with the then director of ARPA, Charles Herzfeld.^[R] I recall a later conversation with Taylor, in December 1966, following another seminal event, securing the cooperation of his

principal investigators to pursue Lick's networking ideas.

The IPT contractors, for the most part academics, but including not-for-profit organizations such as RAND, SRI, and SDC, would meet periodically to hear the research undertaken by the graduate students who were the office's intellectual "shock troops" and to encourage the sharing of ideas among them. Networking was new, few were working in the area, and the matter was seen by the IPT principal investigators either as engineering, not computer science, or a diversion of funds that might be better spent to expand their own, self-evidently more important, areas of research. Taylor's relief in getting them to embrace networking and adapt their computer science research to a networked environment might be difficult to understand when the normal buyer-seller relationship can be characterized as "He who pays the piper calls the tune." But ARPA operated on the basis of partnership with, and consensus among, its world-class contractors. It was not about to pull rank and tell national experts they were wrong, a position that would be inconsistent with why they were part of the ARPA program.

Taylor had the ball rolling and at least a grudging team, but he needed someone to translate the dream into reality. For this, he recruited Larry Roberts. Roberts was at Lincoln Laboratory where, in June 1966, he was working on an implementation of networking, a packet connection between the TX-2 computer at Lincoln Laboratory, and a Q32/PDP-1 computer at SDC in Santa Monica. An account of the work (a message protocol for the link) was presented at the Fall American Federation of Information Processing Societies (AFIPS) Conference in October 1966.^{19,[5]} Roberts was reluctant to leave to run the ARPA program, but did so when Taylor enlisted Herzfeld to do some arm-twisting at Lincoln Lab. The story attests to the extraordinary effort ARPA spent on recruiting the best people for its programs.

The period 1965-1968 was seminal for the synthesis in Roberts' mind of then-current ideas about distributed networks. His contribution, amounting to an independent "invention,"²⁰ was in combining his own work with that of others and converting it into specific decisions to turn the ideas into a working prototype network. He consulted with others working on the subject. Baran was one contributor, bringing not only his team's technical contributions from his

earlier work on digital networks but highlighting the importance of the survivability of defense communication networks under severe stress.

In a 1999 presentation, Roberts recalled that

Packet switching was new and radical in the 1960s. In order to plan to spend millions of dollars and stake my reputation, I needed to understand that it would work. Without Kleinrock's work on Networks and Queuing Theory, I could never have taken such a radical step. All the communications community argued that it couldn't work. This [Kleinrock's] book was critical to my standing up to them and betting that it would work.²¹

In his review of the history of networking, Roberts notes, "the discussions and interest Licklider spawned had an important motivating impact on the initiators of the first two actual network projects, Donald Davies and myself."²² Donald Davies headed the UK Data Communication Network program at the National Physical Laboratory (NPL).²³ Commenting on the relationship between the US and the UK work, Roberts noted,

It was in good part due to this [Licklider] influence that I decided, in November 1964, that computer networks were an important problem for which a new communications system was required. Evidently Donald Davies of the National Physical Laboratory (NPL) in the United Kingdom had been seized by the same conviction, partially as a result of a seminar he sponsored in autumn 1965 which I attended with many M.I.T. Project MAC people.

The next need for Arpanet funding came in May 1968 for a contract to design and build 20 interface message processors (IMPs). The IMPs sat between time-shared host computers to which a user was connected and the public telecommunications network. In fairness to the voice-circuit-oriented communication industry at the time, packet switching did seem to require substantial overhead in software, computing time, and the use of a number of different computers (managed by separate organizations), as well as a dismaying number of interfaces to navigate to transmit a single ASCII character. Roberts clearly understood how rapidly microelectronics technology was moving, his earlier network rationales notwithstanding, and that eventually the

cost of computation, subject to economies of scale, would become trivial compared to communication and its dependence on capital investments.

It was a high-risk gamble that fit ARPA's approach to difficult but important areas of national security. It involved having faith in creative people allowed to operate with a minimum of bureaucratic micromanagement and in the dedicated work of many others in and out of government. Even this computer-enthusiast physicist must own up to raising an eyebrow when hearing the communication protocols involved in packet switching explained.

By this time, Eberhardt Rechtin replaced Herzfeld in November 1967 as the agency's director, and I became ARPA's deputy director. In a presentation for the ARPA FY 69 budget, Roberts brought Herzfeld's foundling back to the front office. Despite the continued commitment of the agency to high-risk programs to prevent technological surprise, by 1968 the environment for ARPA was quite different. Although the need for survivable networks had not changed, political forces were tightening the screws on ARPA. The agency's budget was shrinking, partly to help pay for the Vietnam War and partly because opponents to the war in the Congress were using every opportunity to cut DoD programs. "Fraud, waste, and abuse" was a frequently invoked litany of sins. Lack of relevance to national security was another. Senator William Proxmire periodically identified projects for a Golden Fleece Award, projects with scientifically accurate but easy-to-ridicule titles. Senator Mike Mansfield was soon to demand "relevance statements" for every R&D line item in the DoD.

As ARPA's deputy director, the mind-numbing chore of overseeing the preparation of these statements fell to me. The enormous amount of time with word-by-word review enhanced my normal interest in applying ideas to help fix problems and the experience exerted a strong influence on my subsequent actions as director.

So the question of the degree of risk in the proposed Arpanet arose, both technical and political. Too little risk and the task could be left to others, but too much and perhaps it was a poor investment of limited funds. Roberts was asked if the concept could be tested less expensively at a single site with simulated IMPs and synthetic network traffic. We agreed that such an approach, lacking the real-world accidents and user errors that

a network would have to contend with, would slow down the program with little likelihood that what would be learned would be relevant to understanding real traffic in real networks. The program was approved but structured with an initial four-node test network to examine routing algorithms and other potential show-stoppers before committing to the larger roll-out requested.^[1] This was what led to the first Arpanet configuration, the UCLA-SRI-UCSB-Utah net. The four sites were selected not only to economize on leased lines but because each had an area of expertise to contribute to this first phase of the implementation. The first node was at UCLA, where Kleinrock was now on the faculty, to provide a capability to measure network performance and to test his ideas on packet delay and network congestion. The other three nodes similarly made important contributions to the four-node test network.

BBN was selected as the contractor to modify Honeywell 516 minicomputers to serve as IMPs.²⁴ The contract was let in December 1968 with the first delivery 1 September 1969. The first IMP was delivered to UCLA three days early, ruining yet another weekend for Kleinrock's graduate students.^[1] By December 1969, all four of the test sites were operational.^[1] Finding no reason to change our plans, the Arpanet implementation of the remaining 16 nodes proceeded seamlessly. Seven additional nodes became operational between April and July of 1970. I do not recall Roberts asking for approval to proceed, but he did not have to. The way ARPA operated, he would have come back for guidance only had he felt it was needed.

Maps of the early development of the Arpanet tell a story of the demand side of networking. Beyond the growth in number of nodes, they reveal something of the nature of the demand for networking by user communities (see Table 1).

Eb Rechtin left ARPA in December 1969 to become principal deputy to the DDR&E, and I became ARPA's acting director and then director, eventually leaving at the end of 1974. In addition to the several histories cited, Arpanet's history as part of the broader IPT program is recounted by Arthur Nordberg and Judy O'Neill.²⁵ The history of time-sharing, graphics, and artificial intelligence are also used to illustrate the way in which the innovative atmosphere of ARPA put in place many other parts of information technology as understood today.

Three trends in the growth of networking can be seen.²⁶ The first is a build-out at a linear rate to meet the needs of ARPA/IPT networking R&D, supported for the first two years by the initial IMP procurement and the provision of network service by ARPA as a free good. This was largely completed by September 1973. The number of ARPA R&D nodes and contractors then began to saturate. The next trend in the Arpanet's growth was driven by military R&D sites and a few closely related non-DoD sites; the NASA/Ames node, for example, followed from the decision to locate ARPA's Illiac IV parallel processor there. As early as March 1972, the number of new nodes added for internal ARPA purposes was exceeded by the number of military sites exploring the utility of networking technology for their constituencies and missions. The third trend is the saturation in overall growth of nodes, although users and traffic continued to increase.

A 1974 Mitre study of Arpanet usage showed that about three-quarters of the traffic was email.^{16j} This was somewhat different from Roberts' expectations on how the net would increase the efficiency with which computer scientists would work. At a time when every calculation had to be programmed for a specific computer, the early belief was that users would send their input data to the site best able to handle it. In the 1966 Thomas Marill and Lawrence Roberts paper,¹⁹ the authors argued that users with large matrices to invert would send the problem to sites that already had matrix inverters. Those having a problem requiring natural language understanding would send their calculation to a machine already programmed to do that. But in any event, the ensemble of researchers utilizing the Arpanet, supported by the growing utility of the network, became a self-organizing collection of users that evolved in different directions from what was anticipated. These other uses, for example, included the packet voice work that eventually became what is now voice over IP (VoIP).

One expects R&D programs to have unexpected results; that is the nature of any exploratory enterprise, especially programs in an agency that explained itself in terms of high-risk and high-payoff goals. The result was the Licklider vision of cooperative problem solving, not management efficiency in the utilization of hardware and software, became what the Arpanet was about.

Having developed and demonstrated the utility of networking for facilitating group

Table 1. Growth of early Arpanet nodes by type of sponsoring organization.

Date	IPT R&D	ARPA R&D (Non-IPT)	Military R&D	Non-DoD R&D	Total
January 1970	4	0	0	0	4
June 1970	9	0	0	0	9
December 1970	13	0	0	0	13
September 1971	15	0	1	1	17
March 1972	17	0	6	2	25
August 1972	17	2	7	3	28
September 1973	20	3	11	9	41
June 1974	21	3	14	9	45
July 1975	24	3	18	10	55
July 1976	24	3	21	10	58
July 1977	24	3	24	10	61

efforts, IPT's next step was to break the tether to fixed locations defined by the telephone system. Of obvious importance to military forces is mobility. This reasoning led to the ARPA packet radio net (PRNET). Like so much in the IPT networking program, there were the same set of mixed motivations. Anyone spending DoD money would see the need for extending network connectivity to mobile users. The researchers' goals—deriving from the Licklider vision of networking as encompassing everyone, at all places and all times, with their information, and with the widest possible set of input and output options—were completely congruent with military needs.

Not only did the PRNET have to function independently of the Arpanet, but to meet the intent of networking, it had to connect to the Arpanet so users in both domains could communicate and so the contents of one network were available to users of the other. Hence, the PRNET introduced the need for internetwork protocols.

The technical aspects of the PRNET were detailed by Robert Kahn and his colleagues in 1978.²⁷ The experimental PRNET employed spread spectrum transmissions, common for military applications requiring covertness, but not authorized for civilian use at the time. The spectrum allocation for the experiment was approved by the Interdepartmental Radio Advisory Council (IRAC), the nation's frequency coordination body. But when it came to doing a packet-radio field experiment in the San Francisco area in 1974, the Federal Aviation Administration raised objections for reasons of flight safety at area airports. The minor flap, remedied by a new

spectrum allocation, was one of misunderstanding, or perhaps inattention, but is illustrative of the problems encountered when innovation meets the real world of hardware, regulations, and organizational missions.

Breakout from the ARPA “Beachhead”

In the late 1960s, John Foster, the DDR&E to whom the agency reported, felt ARPA had slipped into a role of being “DoD’s NSF,” and after almost 10 years in existence, there seemed to be little coming out. Upon Eb Rechtin becoming director in 1967, addressing this concern became an ARPA management priority. At that time, the US was planning the deployment of the Sentinel missile defense system. The DDR&E decided to augment the Army’s technical capability by transferring ARPA’s Defender program, including current funding and staff, to the Army Ballistic Missile Defense Agency. In his first year in ARPA, Rechtin negotiated this transfer, amounting to approximately 25 percent of the total agency’s budget and more than half its strategic technology programs and personnel.²⁸

Beyond addressing ballistic missile defense and OSD’s perception of the agency’s value, the Defender program transfer underlined a need for its staff to shift some of its attention from the input stage, formulating new programs to address new problems and threats and managing high-quality projects, to the output stage, moving their results to those who could use them. Good research was not enough. Successful research had to have more than potential utility. Its actual utility for national security had to be demonstrated to Executive Branch budget managers and Congressional appropriators through “rubber on the road.” It was the beginning of a coupling of the 6.1-funded “Research” programs in ARPA to the 6.2 “Exploratory Development” category.²⁹ It also led to closer and more positive relations with the military services that had previously seen ARPA as a thorn and a waste rather than an effective partner in national security. This was more natural for the 6.2-funded offices in the agency, but was a less familiar perspective for the research offices.

Although networking was shifting from a managed and fully subsidized part of IPT, and was being used by other parts of government, those organizations were still engaged in networking R&D but focused on their own mission areas. There were no purely operational users who cared only about the quality and reliability of the service.

The Arpanet was providing email service to its contractors and supporting a networking research community. The Mitre study documenting that did not sit well with OSD auditors. It meant ARPA had created an operational communication system under their very noses, without following the rules dictating the DCA as having that mandate. Had ARPA said initially the net would be a mail system utilizing packet switching, it should, to conform to DoD policy, have gone to DCA to get it built—notwithstanding the fact that no one at the time knew how to do that. The request would have engendered much study and debate culminating, if approved at all, in a system that would likely have been late and over-budget, the same concerns that dissuaded Baran and the Air Force in 1965 from utilizing DCA to construct their proposed message-block system.

In the long run, the auditor’s recognition of ARPA’s violation of policy had a positive effect. It took ARPA, an R&D agency, out of the business of managing an operational network far from its research objectives and required creation of a management infrastructure dealing with details that were distractions to the agency. It also forced DCA to pay attention to packet switching. DCA handled the Arpanet management task well and was soon converting to the new packet-switched environment for its own programs.^[Z]

In 1973, I was prodding the non-IPT offices in ARPA to incorporate the network idea into their thinking. One approach was simply to use the net for internal ARPA operations. The flexibility of email for management traffic and creating records of transactions was appealing, not simply to boost networking but because it worked better than what we had: seven-ply carbon paper, copiers, phones, and snail mail.³⁰ The normal competitive instincts for the director’s attention, and hence funding for their projects, made easy access to the front office through the net attractive.

Another opportunity to use the Arpanet for operational purposes arose from the Nuclear Test Detection office, which I directed when I joined ARPA in 1966. In 1963 the US entered into a Limited Test Ban Treaty banning nuclear tests in the atmosphere, in space, and underwater.^[W] But the inability to distinguish between underground explosions and earthquakes prevented an agreement banning those tests pending further research on how to verify compliance,

which is central to arms-control agreements. Because the nature of earth movements for explosions and earthquakes are inherently different, this research sought to substantiate their differences as recorded by instruments at teleseismic distances and to determine under what circumstances the two could be reliably separated.³¹ This ARPA program, Vela, was to provide data that US policy makers could use in negotiations for a treaty banning all nuclear weapon tests.

Central to the Vela program were arrays of seismometers. Through beam forming, one can improve the signal-to-noise ratio to significantly improve detection in the presence of seismic noise arising from natural and man-made sources. The first of these arrays built by ARPA was the Large Aperture Seismic Array (LASA) in Montana in 1965.^[X] It comprised 345 short period and 16 three-component long period instruments with an aperture of 200 km.³²

The convenience of developing seismic array technology in the US had a downside for studying teleseismic discrimination. LASA was too close to the only well-defined source of nuclear explosions available, those at the Nevada Test Site. Experience gained in LASA's design and operation was used in the installation of a second large array, the Norwegian Seismic Array (NORSAR), placed at a teleseismic distance from both US weapon tests and those at the Soviet test site in Kazakhstan.³³ It was developed in conjunction with the Royal Norwegian Council for Scientific and Industrial Research and installed by the Norwegian Defense Research Establishment. NORSAR consisted of 132 short-period instruments and 22 long period instruments. The array was approved in January 1967 and became operational in February 1971.^[a,b] LASA and NORSAR recorded signals from nuclear events in the USSR, which allowed corroboration of data and the introduction of multiarray analysis to reduce the detection threshold below a seismic magnitude 4.2.

Data from the arrays were analyzed at the Seismic Array Analysis Center (SAAC) in Alexandria, Virginia, where a Terminal Interface Processor (TIP), essentially a standalone IMP that allowed terminals direct connection to the Arpanet, was installed by August 1972.^{34,[Y]} The installation of a TIP at NORSAR occurred at the same time.^[c] This center could now perform multiarray analysis to compare signals from different azimuths. Getting near-real-time NORSAR data to the

SAAC turned out to be more of a bureaucratic problem than anticipated, not because of technical issues but because it required interacting with a number of new entities that neither knew nor cared about packet switching or arms control.

The connection of Norway to the SAAC provided an opportunity to add UK network researchers to the Arpanet. Davies' independent network research at the NPL net, consisting of about 10 sites in the London area, was approved in 1968, and went into operation in 1970.^[f,g,h] It was first reported upon in 1972.³⁵

When international borders are crossed, national common carrier regulations are subject to international agreements that are far more complex than those involved in dealing with a government-friendly monopoly carrier. New bureaucratic problems surfaced, which Peter Kirstein reports on in detail.³⁶ Although the technical issues proved to be tractable, the regulatory and political roadblocks were more persistent, overcome only after considerable delays. Because the UK in the early 1970s was moving toward entry into the European Common Market, their political emphasis was on links to the European Informatics Network in preference to those with the US. Further impediments were bureaucratic. To connect the UK network to the Arpanet required the use of a TIP furnished by ARPA. On arrival at London's Heathrow Airport, however, it was impounded by Customs and Excise for duty to be paid and an additional £5,000 was levied for Value Added Tax, neither of which had been anticipated. The UK Science Research Council (SRC) held that this connection did not present a particularly fruitful opportunity and declined to provide the funding needed.

Another complication arose when the Scandinavian Tanum satellite ground station came online, which allowed the retrieval of NORSAR data at 50 kilobits per second (Kbps) for the first time. This upgrade of bandwidth obviated further need for the 9.6-Kbps cable line that passed through London en route to Norway. A new London-Norway link was needed to connect to the now expanded Norway-Washington link, but the tariff proposed by the carriers acting under the umbrella International Telegraph and Telephone Consultative Commission (CCITT) was beyond any possible budget provision. The French wanted the tariff to be priced at the number of multiples of 2.4 Kbps

voice circuits to achieve the Arpanet speed of 50 Kbps. The British Post Office offered to provide the link free for a year. Later, an appeal to the British Post Office to provide a 50 Kbps link from their satellite station at Goonhilly Downs was approved. They provided the use of the satellite link on a research basis and thus circumvented the French roadblock.

Kirstein explains that because of the policy to favor European over US links, it was necessary to shift the UK–Arpanet connection from NPL to University College London (UCL). This idea was contained in a proposal Kirstein submitted to the SRC at the end of 1971. It was turned down then as being too speculative and uncertain, despite successful operation in the US for several years. The Department of Trade and Industry insisted on statements of interest from industry. After nine months of study, their conclusion was that “one would gain more from a two-week visit to the US than from a physical link.”

Nevertheless, the lack of official UK enthusiasm had several positive results. The UK link had to be for research purposes (an ARPA condition as well) and the net had to be developed based on real traffic, thus opening up participation by a broad set of UK academic researchers. Originally, UK access was not completely open, requiring a password, but this requirement was eventually dropped. The first packet was transmitted from UCL to the University of Southern California Information Sciences Institute (USC/ISI) in Marina del Rey on 25 July 1973.^[i]

The duty and VAT arguments were not settled until 1976. The UCL appeal was initially rejected, but more senior levels of the treasury ruled—after being threatened with export of all equipment back to DoD and reimport under the US–UK Exchange of Forces Agreement—that “the equipment that you have imported, and any future equipment brought under the same agreement would be free of duty and VAT.” This had the beneficial result of preventing other UK government departments from taking over the project because of the duty and VAT they would otherwise have to pay. The UK Ministry of Defence, through this period, was most helpful and supported the UCL research on network protocols with the condition that MoD unclassified research networks could be connected. The international logjam was eventually broken and the British Post Office, British Library, Department of Trade and Industry, the MoD Laboratories

in Aldermaston and Malvern, and the SRC joined the enterprise. The diplomatic highlight of the Arpanet was when Queen Elizabeth opened the Malvern link to the US in 1976, thus marking, as Kirstein notes, the first involvement of a head of state in a computer network.

Aside from the bureaucratic gymnastics arising from adding overseas sites to the Arpanet, the experience stimulated two additional dimensions to ARPA’s networking R&D. The first was the addition of satellite links that enabled the creation of a global network. (Fiber optic cables were not in the picture at this time.) The second was the need, already seen in the concurrent packet-radio network development noted earlier, to interconnect networks. The need to connect NORSAR provided the opportunity to generalize the network from a US to a global entity and was central to the Arpanet’s eventual commercial exploitation. Irwin Mark Jacobs explained:

The Atlantic Packet Satellite Experiment serves as the development vehicle and experimental test-bed for the overall Packet Satellite Program. The experiment includes both the experimental satellite network, called SATNET, and the supporting development and measurement activities and facilities. . . .

The basic philosophy of the experimental program and the SATNET development efforts has been to achieve a flexible experimental environment which could serve as a measurement vehicle for a spectrum of specific system realizations and applications. Further, development and measurement emphasis has been placed on those system aspects that are least understood and most in need of feasibility demonstration.³⁷

These developments marked the transition of the academic computer scientists’ Arpanet to a global facility encompassing operational as well as R&D users, embraced by the communications authorities of three countries, and with packet technology demonstrated as technically viable in a variety of transmission modes. The role of ARPA in networking technology had entered its next phase.

Related to the scope and timeframe of this memoir is the technology of wireless local area networks (LANs). In 1979, I became the chief scientist of the Federal Communications Commission. FCC Chairman Charlie Ferris brought me in, based on my ARPA experience, to speed up the FCC’s approval of

new communication technology. The organizational challenge was similar to that facing DoD following the launch of Sputnik in 1957. Some of what I accomplished at the FCC in the next few years included commission approval of the first spectrum allocation for cellular telephony, authorization of direct broadcast satellites, and technical rules for the use of unlicensed spread spectrum technology. The story of unlicensed wireless policy is a similar mix of the technology, bureaucracy, and politics seen in networking.³⁸

Lessons from the Arpanet Development

In 1983, the Arpanet was divided into two parts, a research network that continued to be called Arpanet, and Milnet serving military users. In 1989, the Arpanet was closed down. ARPA's intent had been to demonstrate the utility of packet switching for military command and control, and in that, it succeeded brilliantly. As a consequence unanticipated by its sponsor, the Arpanet and its successor, the Internet, further demonstrated the general utility of networking for "command and control" far beyond the needs of the DoD, just as Lick anticipated.

Absent the subsequent metamorphosis of the Arpanet into today's Internet, the story of the ARPA experiment would be a footnote in a technical history of military command and control. But because of its evolution into a public-access service based on the same technology and protocols, the Arpanet development process deserves comment.

The Arpanet, as is typical of the output of any process, bears the imprint of its origin. Central was its open nature based on "rough consensus and running code," a process recently described by Steve Crocker.³⁹ Open source code, open platforms, wikis, social networks, text messaging, and the like are typical of the academic culture that produced the Arpanet. But that openness has not scaled well to the global community.

A frequent mantra of those concerned with correcting the problems introduced into the Internet from its Arpanet origin is the notion of public-private partnerships. This professed need follows from the private ownership of most Internet facilities and services. It reflects the joint responsibilities of both private owners and governments for correcting its security flaws. Public-private partnerships work in many cases, the Arpanet development being a successful case in point. The partnership between the DoD and the

ARPA contractor community worked because the partners were moderately homogeneous in their level of knowledge and their willingness to cooperate with one another. Although the networking community had its differences (largely technical disagreements), the broader DoD community that might have opposed ARPA's work in command and control was buffered from the performers. The mediator, ARPA, managed the interface, and was trusted by both sides.

There is nothing comparable today for addressing Internet security. There is no trusted mediator with technical expertise and the resources needed to fix global problems. The Internet has grown so large that the number of interested parties is probably too numerous to find enough common ground. Even if by some miracle governments and industries worldwide were to agree on a course of action to remove threats to the Internet and the services dependent on it, substate groups—such as criminals benefiting from the current insecurity—are unlikely to cooperate. Neither will heads of state forgo the intelligence-collection opportunities it provides, nor will military services abandon the precision of cyberweapons that can be directed against national infrastructures and other military targets.

The adoption of packet switching by common carriers and the infusion of digital technology into so many aspects of life have been commercially successful, but the current quality of network service reflects the process through which this occurred, just as the original net reflected ARPA processes. In the US, the FCC, which might have provided such guidance to protect users, has been largely absent after authorizing the resale of transmission capacity by nonregulated carriers. There were two factors influencing the FCC's lack of Internet regulation. Primarily, the commission made a considered decision to not regulate what it called "valued-added services," preferring the deregulatory view that regulation is more apt to constrain innovation than assist it. Although the FCC recognized the impending information revolution and did not want to impede it, like most people, it did not have an adequate understanding of the technology and its policy implications to enable it to act responsibly in terms of protecting users.⁴⁰ Thus, it is a question whether we could have avoided the insecure network we now have.⁴¹

This is not unlike the current problem with the proliferation of nuclear weapon states.

In 1939, nuclear weapons seemed like a good idea. By 1949, they had become a nightmare, but the nightmare had the Internet as a consequence. In the case of networking, the technical community and the DoD focused on what could be done at the time. We did not attempt to follow all possible moves ahead and plan how we would prevent unexpected and undesirable consequences. To attempt to do so before proceeding would have paralyzed us to inaction, assuring that nothing positive would be accomplished. On balance, ARPA and its developers protected the nation from one kind of technological surprise, only to create another. Similar problems of technology confront us today, in dealing with energy and climate change, economic development and poverty, and genetics and its potential for misuse.

One can speculate about what would have happened without an ARPA or the accidents that caused it to improve capabilities for the command and control of military forces. Quite possibly, we'd have something similar to what we have today. Computers would still need connecting, carriers would still want to sell transmission capacity, and industry would still want to sell the services thereby enabled and the hardware and software to provide them. The resultant network technology might be structured as proprietary islands of computers and users, perhaps like today's separate cellular networks. Today we might be looking to public-private partnerships to bring about a more interconnected state of affairs. The early ARPA decision to keep protocol developments open and in the hands of developers acting for the user community as a whole would have been more problematic in the face of private-sector competition.

Acknowledgments

The following colleagues have reviewed this manuscript at various stages and provided encouragement and further insights: Paul Baran, Larry Roberts, Peter Kirstein, Charlie Kline, Len Kleinrock, Barry Wessler, Bob Balzer, Steve Crocker, Charlie Herzfeld, Eric Willis, Mike Marcus, Ed Feigenbaum, and Bob Kahn. Their help is gratefully acknowledged.

References and Notes

1. Bracketed letters refer to events indicated on the summary timeline in Figure 1.

2. R. Buder, *The Invention that Changed the World*, Touchstone/Simon & Schuster, 1997.
3. R.E. Powaski, *March to Armageddon*, Oxford Univ. Press, 1989.
4. A program manager recalls, "When I first met the two-star Air Force general who was to become the co-sponsor of my National Software Works project, his first words were, 'I've always admired ARPA,' and I expected him to say something about our technical prowess. He finished the sentence with '... you are able to spend money so effectively.'"
5. The history is based on official documents and interviews with ARPA technical and management personnel. See the Richard J. Barber Associates report, *The Advanced Research Projects Agency, 1958-1974*, supported by the Advanced Research Projects Agency of the Department of Defense under contract number MDA 903-74-0-0096. It is available through the National Technical Information Service (www.ntis.gov).
6. These and other quotes and details are taken from the Barber and Associates history. When serving as the agency's deputy director, I found traces of this early history in the agency's files and was surprised to find we were originally directed to work on command and control, not nurturing computer science in universities to prevent technological surprise, as I had been led to believe. Military command and control was not in the forefront of the minds of the ARPA researchers by this time nor was command and control research a recognized driver in ARPA's computer and behavioral science programs. It was seen simply as one area of potential application.
7. J.C.R. Licklider, "Memorandum for Members of the Intergalactic Computer Network," ARPA Behavioral Science Command & Control Research, 23 Apr. 1963; www.kurzweilai.com/articles/art0366.html?printable=1 posted December 11, 2001.
8. M.M. Waldrop, *The Dream Machine: J.C.R. Licklider and the Revolution That Made Computing Personal*, Viking, 2001. In his bibliography, Waldrop notes his use of the oral histories of the University of Minnesota Charles Babbage Institute.
9. This account was provided by C. Al Irvine, the leader of the UCLA effort and the author of the three-center network proposal, to Steve Crocker.
10. The first staff members were a distinguished group: Henry Kissinger, John Forbes Nash, Stephen Dole, Herman Kahn, Harold Brown, and Donald Rumsfeld.
11. W.H. Ware, "RAND and the Information Revolution," RAND, 2008.

12. This argument depends on a presumption of rationality of the attacker. Concerns over irrational attackers introduce further complication into the decision processes.
13. The RAND work was published in a series of reports as the design grew in detail. In 1964, the team published them as a single series titled *On Distributed Communications*. These can be found at www.rand.org/about/history/baran.list.html. The documents in the series are P. Baran, *I. Introduction to Distributed Communication Networks*, RM-3420-PR; S.P. Boehm and P. Baran, *II. Digital Simulation of Hot-Potato Routing in a Broadband Distributed Communication Network*, RM-3103-PR; J.W. Smith, *III. Determination of Path-Lengths in a Distributed Network*, RM-3578-PR; P. Baran, *IV. Priority, Precedence, and Overload*, RM-3638-PR; P. Baran, *V. History, Alternative Approaches, and Comparisons*, RM-3097-PR; P. Baran, *VI. Mini-Cost Microwave*, RM-3762-PR; Paul Baran, *VII. Tentative Engineering Specification and Preliminary Design for a High-Data-Rate Distributed Network Switching Node*, RM-3763-PR; P. Baran, *VIII. The Multiplexing Station*, RM-3764-PR; P. Baran, *IX. Security, Secrecy, and Tamper-Free Considerations*, RM-3765-PR; P. Baran, *X. Cost Estimate*, RM-3766-PR; and P. Baran, *XI. Summary Overview*, RM-3767-PR. See also Baran's account in the Willis Ware RAND publication.
14. P. Baran, private communication with S. Lukasik, 2008. See also J. Abbate, *Inventing the Internet*, MIT Press, 2000, chap. 1.
15. Early workers in the Arpanet project differ in their view of what they were engaged in, with few recognizing the fundamental military purposes of the program. All have various recollections of the scope of their expectations. An IPT Principal Investigator, working outside the networking area, noted ruefully, "We were all non-visionary."
16. "Proposal for a Ph.D. Thesis, Information Flow in Large Communication Nets," 31 May 1961; www.cs.ucla.edu/~lk/LK/Bib/REPORT/PhD.
17. L. Kleinrock, *Communication Nets: Stochastic Message Flow and Delay*, McGraw-Hill, 1964; reprinted by Dover Publications, 1972.
18. P. Baran, "The Beginnings of Packet Switching: Some Underlying Concepts," *IEEE Comm. Magazine*, July 2002, pp. 42–48.
19. T. Marill and L.G. Roberts, "Toward a Cooperative Network of Time-Shared Computers," Nov. 1966; www.packet.cc/fies/toward-coop-net.html.
20. "Invention" is used here with some trepidation, given the plethora of those who would be pleased to accept such a mantle. In the case of any large and diverse enterprise, the essence of the creation is not the individual but the group. In this regard, the precedent of the 2007 Nobel Peace Prize Committee to recognize as awardees more than 450 lead authors, 800 contributing authors, and 2,500 expert reviewers of the Intergovernmental Panel on Climate Change is relevant. Nevertheless, the penchant for seeking simple causes for complex effects is strong.
21. L.G. Roberts, "The First Theory of Packet Networks," presentation at SIGCOMM 99; www.lk.cs.ucla.edu/first_packet_network_theory.html.
22. L.G. Roberts, "The Evolution of Packet Switching," *Proc IEEE*, vol. 66, no. 11, Nov. 1978, pp. 1307–1313. The independence of the UK conception of packet switching from that in the US is unclear from this, but that is immaterial in view of the continuing technical exchanges between the UK and the ARPA group that provided important stimulation to both efforts. The drive to connect computers together between 1962 and 1964 is unsurprising, with Baran, Kleinrock, Licklider, Davies, and Roberts converging on the idea almost simultaneously. This was the period when digital computers shifted from being one-of-a-kind devices to a commodity product of increasing size and power. Connecting them was an obvious next step. The devil, as always, was in the details.
23. In view of the natural interest in priority, Len Kleinrock, in his historical review, "Creating a Mathematical Theory of Computing Networks," *Operations Research*, vol. 50, no. 1, 2002, recounts the following chronology regarding the NPL work. In D.W. Davies et al., "Proposal for a Digital Communication Network," unpublished memo, June 1966, www.archive.org/details/NationalPhysicalLaboratoruProposal-ForADigitalCommunicationNetwork, he described his proposed design for what he termed a "packet switched network." The first open publication was in D.W. Davies et al., "A Digital Communication Network for Computers Giving Rapid Response at Remote Terminals," *Proc. ACM Symp. Operating System Principles*, ACM Press, 1967. This network, which became operational in 1970, was described by R.A. Scantlebury, "A Model for the Local Area of a Data Communication Network Objectives and Hardware Organization," *Proc. 1st ACM Sym. Problems in Organization of Data Comm. Systems*, ACM Press, 1969.
24. Industry giants such as IBM and AT&T declined to bid, testimony to the high-risk nature of the ARPA network technology.
25. A.L. Nordberg, J.E. O'Neill, and K.J. Freedman, *Transforming Computer Technology: Information Processing in the Pentagon, 1962–1986*, Johns Hopkins Univ. Press, 2000.

26. The maps used are not referenced since they can be found on the Internet in profusion.
27. R.E. Kahn et al., "Advances in Packet Radio Technology." *Proc. IEEE*, vol. 66, no. 11, 1978, pp. 1468–1496.
28. To head off what was seen as a large new weapons program and a destabilizing technology in the offensive strategic nuclear balance, the US and the Soviet Union signed an Anti-Ballistic Missile Defense Treaty in 1972 providing for a limitation of defensive systems to a missile field protecting the national command authority and one protecting a single missile field to serve as a survivable deterrent force. This formalized the strategic doctrine of mutually assured destruction—defense had become offense, and offense was defense. This formulation was coined by J. Newhouse and quoted by A. Wohlstetter in "Bishops, Statesmen, and Other Strategists on the Bombing of Innocents," *Commentary*, vol. 75, no. 6, 1983, pp. 15–35.
29. This move to transfer research is even more strongly reflected in the current DARPA program. The agency now receives 6.3 Advanced Development funds, thus helping push its results even further downstream.
30. My experience in using email to manage the agency is detailed in the Waldrop's *The Dream Machine* and in K. Hafner and M. Lyon, *Where Wizards Stay Up Late: The Origins of the Internet*, Simon & Schuster, 1996. It was not smooth. Quirks computer scientists worked past or fixed themselves were a much bigger problem to a physicist for whom writing Fortran programs was a high-water mark in computing expertise.
31. The key word here is *teleseismic*, meaning at a large distance from the nuclear event. Although nuclear tests in a small country can be monitored from immediately outside its borders, in large countries such the Soviet Union and the US, where suspicion of the integrity of data from internal sites for monitoring adherence to a test ban could not be satisfied, treaty verification had to be based on data collected outside the borders of both parties.
32. The LASA design, and much of the signal processing research, was the responsibility of the MIT Lincoln Laboratory under Paul E. Green, Jr. It is described in the Nuclear Test Detection issue of *Proc. IEEE*, vol. 65, no. 12, 1977.
33. At the time, the arms control community was urging the installation of a network of 12 LASA-type arrays around the world. The ARPA position was that the detection threshold and discrimination capability of seismic arrays had not yet been established to the point where the US could safely enter into a treaty banning underground nuclear tests. The Norwegian array was thus a continuation of the ARPA Vela research program. It thus undercut the deployment of what would have been an unproven treaty verification system.
34. Later renamed the Seismic Data Analysis Center (SDAC); both acronyms are used on Arpanet maps as appropriate.
35. D.L.A. Barber et al., "Operating Experience with the NPL Network," *Proc. ACM Symp. Computer Networks*, IRIA, 1972.
36. P.T. Kirstein, "Early Experiences with the Arpanet and Internet in the United Kingdom," *IEEE Annals of the History of Computing*, vol. 21, no. 1, 1999, www.cs.ucl.ac.uk/jon/arpa/internet-history.html. See also V.G. Cerf and P.T. Kirstein, *Proc. IEEE*, vol. 66, no. 11, 1978, pp. 1386–1404, for technical details.
37. I.M. Jacobs, R. Binder, and E.V. Hoverstein, "General Purpose Packet Satellite Networks," *Proc. IEEE*, vol. 66, no. 11, 1978, pp. 1448–1467.
38. See *Info*, vol. 12, no. 5, 2009, edited by C.L. Jackson at www.emeraldinsight.com/1463-6697.htm, particularly the papers by M.J. Marcus "Wi-Fi and Bluetooth—The Path from Carter and Reagan-Era Faith in Deregulation to Widespread Products Impacting our World," and S.J. Lukasik, "Unleashing Innovation: Making the FCC User-Friendly."
39. S.D. Crocker, "How the Internet Got its Rules," *New York Times*, 7 Apr. 2009.
40. S.J. Lukasik, "Unleashing Innovation: Making the FCC User-Friendly," *Info*, vol. 12, no. 5, 2009, pp. 76–85.
41. The FCC has recently inserted itself into Internet regulation but not to address its security flaws; instead, it has taken up the political issue of net neutrality.



Stephen J. Lukasik is a distinguished senior research fellow at the Center for International Strategy, Technology, and Policy at the Georgia Institute of Technology. His recent work deals with vulnerabilities of industrialized societies to terrorist attack. Lukasik received a PhD in physics from the Massachusetts Institute of Technology. He is a member of the International Institute for Strategic Studies, American Physical Society, and American Association for the Advancement of Science. Contact him at steve@gnsi.org.

cn Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.