


DARPA AND THE INTERNET REVOLUTION

By Mitch Waldrop

In the best bureaucratic tradition, DARPA – or ARPA, as it was called at the time – involvement in the creation of the Internet began with a memo.



29 Oct 69	2100	LOADED OP. PROGRAM CSK	
		FOR BEN BARKER	
		BBV	
	22:30	Talked to SRI	CSK
		Host to Host	
		Left op. program	CSK
		running after sending	
		a host dead message	
		to imp.	

Dated April 23, 1963, the memo was dictated as its author, Joseph Carl Robnett Licklider, was rushing to catch an airplane. No surprise there: Licklider was spending a lot of his time on airplanes in those days. The previous fall, he had come to the Pentagon to organize the Information Processing Techniques Office (IPTO), ARPA's first effort to fund research into "command and control" – that is, computing. And he had been crisscrossing the country ever since, energetically assembling a network of principal investigators scattered from the Rand Corporation in Santa Monica, Calif., to MIT in Cambridge, Mass.

Licklider's task might have been easier if he had been pursuing a more conventional line of computing research – improvements in database management, say, or fast-turnaround batch-processing systems. He could have just commissioned work from mainstream companies like IBM, who would have been more than happy to participate. But in fact, with his bosses' approval, Licklider was pushing a radically different vision of computing.


His inspiration had come from Project Lincoln, which had begun back in 1951 when the Air Force commissioned MIT to design a state-of-the-

Above, left: A sketch of the ARPANET in December 1969. The nodes at UCLA and the Stanford Research Institute (SRI) are among those depicted. Above, right: An excerpt from the ARPANET logbook. The information pertaining to the first ARPANET transmission is highlighted.

art, early-warning network to guard against a Soviet nuclear bomber attack. The idea – radical at the time – was to create a system in which all the radar surveillance, target tracking, and other operations would be coordinated by computers, which in turn would be based on a highly experimental MIT machine known as Whirlwind: the first “real-time” computer capable of responding to events as fast as they occurred. Project Lincoln would eventually result in a continent-spanning system of 23 centers that each housed up to 50 human radar operators, plus two redundant real-time computers capable of tracking up to 400 airplanes at once. This Semi-Automatic Ground Environment (SAGE) system would also include the world’s first long-distance network, which allowed the computers to transfer data among the 23 centers over telephone lines.

Licklider, who was then a professor of experimental psychology at MIT, had led a team of young psychologists working on the human factors aspects of the SAGE radar operator’s console. And something about it had obviously stirred his imagination. By 1957, he was giving talks about a “Truly SAGE System” that would be focused not on national security, but enhancing the power of the mind. In place of the 23 air-defense centers, he imagined a nationwide network of “thinking centers,” with responsive, real-time computers that contained vast libraries covering every subject imaginable. And in place of the radar consoles, he imagined a multitude of interactive terminals, each capable of displaying text, equations, pictures, diagrams, or any other form of information. By 1958, Licklider had begun to talk about this vision as a “symbiosis” of men and machines, each preeminent in its own sphere – rote algorithms for computers, creative heuristics for humans – but together far more powerful than either could be separately. By 1960, in his classic article “Man-Computer Symbiosis,” he had written down these ideas in detail – in effect, laying out a research agenda for how to make his vision a reality. And now, at ARPA, he was using the Pentagon’s money to implement that agenda.

In retrospect, the program that Licklider put together in those early months would evolve into perhaps the most successful federal research program in history. The list of researchers he funded now reads like a Who’s

Interface Message Processor	
1965	Bolt, Beranek and Newman, United States
 <p>Credit: Bolt, Beranek & Newman</p> <p>IMP development team, c. 1965</p>	<p>The Interface Message Processor (IMP) was the first packet router for the ARPANET, the predecessor of today’s Internet. Inside was a Honeywell 516 minicomputer with only 6,000 words of software to monitor network status and gather statistics. The first ARPANET transmission occurred between the University of California in Los Angeles and Stanford Research Institute in Menlo Park, California, at 22:30 PST on October 29, 1969.</p> <p><i>Gift of Bolt, Beranek and Newman, Inc., X105.82</i></p>


The Interface Message Processor was a key refinement to the decentralized nature of ARPANET, routing information so that data could avoid being sent through every network node’s main computer.

Who of the greatest U.S. computer pioneers, from Allen Newell, Herbert Simon, and Alan Perlis at Carnegie Tech (now Carnegie Mellon University) to John McCarthy at Stanford University and Douglas Engelbart at the Stanford Research Institute (now SRI International). And the technologies they developed, from computer graphics and artificial intelligence to the mouse-controlled user interface, laid the foundations for computing as we know it today. Particularly important was Licklider’s flagship effort, MIT’s Project MAC, the world’s first large-scale experiment in personal computing. The project managers couldn’t hope to give anyone a stand-alone personal computer, of course, not with the cheapest machines still costing hundreds of thousands of dollars. But they could scatter dozens of remote terminals around the campus, and in people’s homes. And then through the technology of time-sharing, they could tell their big, central machine to dole out little slices of processing time very, very rapidly, so that each user would feel as if it were responding to him or her as an individual, in real time. By the mid-1960s, Project MAC would evolve into the world’s first online community, complete with online bulletin boards, e-mail, virtual friendships, an open-source software exchange – and hackers.

Licklider’s research program was so successful, in fact, that it’s now hard for us to remember just how visionary it was. IBM and the other major computer manufacturers were going in a completely different direction at the time, emphasizing punch cards and batch-processing machines suited to the needs of the business world. Mainstream computer engineers tended to see the ARPA approach as totally wrong-headed. Use precious computer cycles just to help people think? What a waste of resources!

Likewise, the handful of researchers who were going in Licklider’s direction were isolated voices in the wilderness, scattered at universities and think tanks all over the country. His job now was to seek them out, nurture their work with ARPA cash, and forge them into a self-sustaining community that could carry on after he was gone.

Thus his memo on April 23, 1963, which he addressed to “the members and affiliates of the Intergalactic Computer Network” – that is, his principal investigators. Although the IPTO program had gotten off to a good start, he told the investigators, he saw a real danger that it might never realize its full potential – that the nascent ARPA community might never become anything more than a high-tech Tower of Babel, in which widely



A map of the ARPANET in December 1970.

scattered enclaves produced incompatible machines, incompatible languages, and incompatible software. Geography almost guaranteed it, he wrote.

So they would have to transcend geography. They would have to take all their time-sharing computers, once the machines became operational, and link them into a national system. “If such a network as I envisage nebulously could be brought into operation,” Licklider wrote, “we would have at least four large computers, perhaps six or eight small computers, and a great assortment of disc files and magnetic tape units – not to mention the remote consoles and Teletype stations – all churning away.”

Leave aside the primitive technology and the laughably small number of machines: The vision that lay behind that sentence is still a pretty good description of the Internet we have today. Indeed, Licklider’s Intergalactic Network memo would soon become the inspiration for the Internet’s direct precursor, the ARPANET.

Because the necessary technology wasn’t even close to being ready, Licklider himself could not get an ARPA network started during his time there. (He left in 1964, first going to IBM and then back to MIT.) Nor could his successor, Ivan Sutherland. It wasn’t until 1966 that the third IPTO director, Robert Taylor, decided that the time was right. Not only was the technology much further along by that point, but the proliferation of incompatible systems was getting just as bad as Licklider had feared. Taylor had been forced to install three different Teletype terminals in his office at the Pentagon just to connect to the three major ARPA-funded, time-sharing systems. As he said in a later interview: “Anyone in that context would have quickly thought, ‘Hey, wait a minute – why can’t I get to any of these places from one terminal?’”

Getting the money to launch a networking project proved to be no problem at all. ARPA Director Charles Herzfeld had been good friends with Licklider during the latter’s tenure in IPTO, and understood the significance of computer networking quite well; he listened to Taylor’s pitch for about 20 minutes and approved \$1 million on the spot.

Next, with money in hand, Taylor brought in Larry Roberts from Lincoln Laboratory (the successor to Project Lincoln) to manage the networking project. And by early 1967, Roberts had made three key technical decisions that would define the architecture of the new network – and that still define the Internet today.

First, since nobody was going to give the agency a few billion dollars to string its own wires across the country, ARPA would have to move the data through AT&T’s telephone system. Unfortunately, that system’s basic dial-up process was far too cumbersome and slow for computer-speed communications. So instead, Roberts decided ARPA would make a series of massive long-distance calls, and just never hang up. More precisely, the agency would go to AT&T and lease a series of high-capacity phone lines linking one ARPA site to the next, so that the computers would always be connected.

Second, Roberts decided that digital messages would not be sent through the network as a continuous stream of bits. Instead, like a long letter written on a series of postcards, each message would be broken into segments of fixed length – “packets,” in modern parlance. The idea was to safeguard against static and distortion in the line, which could easily garble the bits and reduce the message to gibberish. The packets wouldn’t eliminate the noise. But they would isolate the errors and give the system a chance to fix them, either through sophisticated “error-correcting” codes, or by asking the original sender for a new copy. (The



A 1969 Interface Message Processor.

idea of using packets for networking had originated several years earlier with Paul Baran at the Rand Corporation, and had been mathematically analyzed in the 1962 Ph.D. thesis of Roberts' MIT classmate, Leonard Kleinrock, who was now at UCLA.)

Finally, Roberts decided to make the network completely decentralized, with no one master computer responsible for sorting the packets and routing them to their destination. Such a Grand Central Station approach would have been much simpler to implement, Roberts knew. But one blown transistor could have taken the whole network down. So instead, he decreed that the ARPA sites would be linked in a complex pattern rather like the Interstate highway map, and that each site would share the routing responsibilities equally. That is, the computer there would read the digital address on each packet as it came in, accept that packet if the address was local, or else send it off again on the next stage of its journey. This “packet-switching” approach would mean more complexity and more programming during the set-up phase. But the final system would be far more robust: No one failure could bring it down.

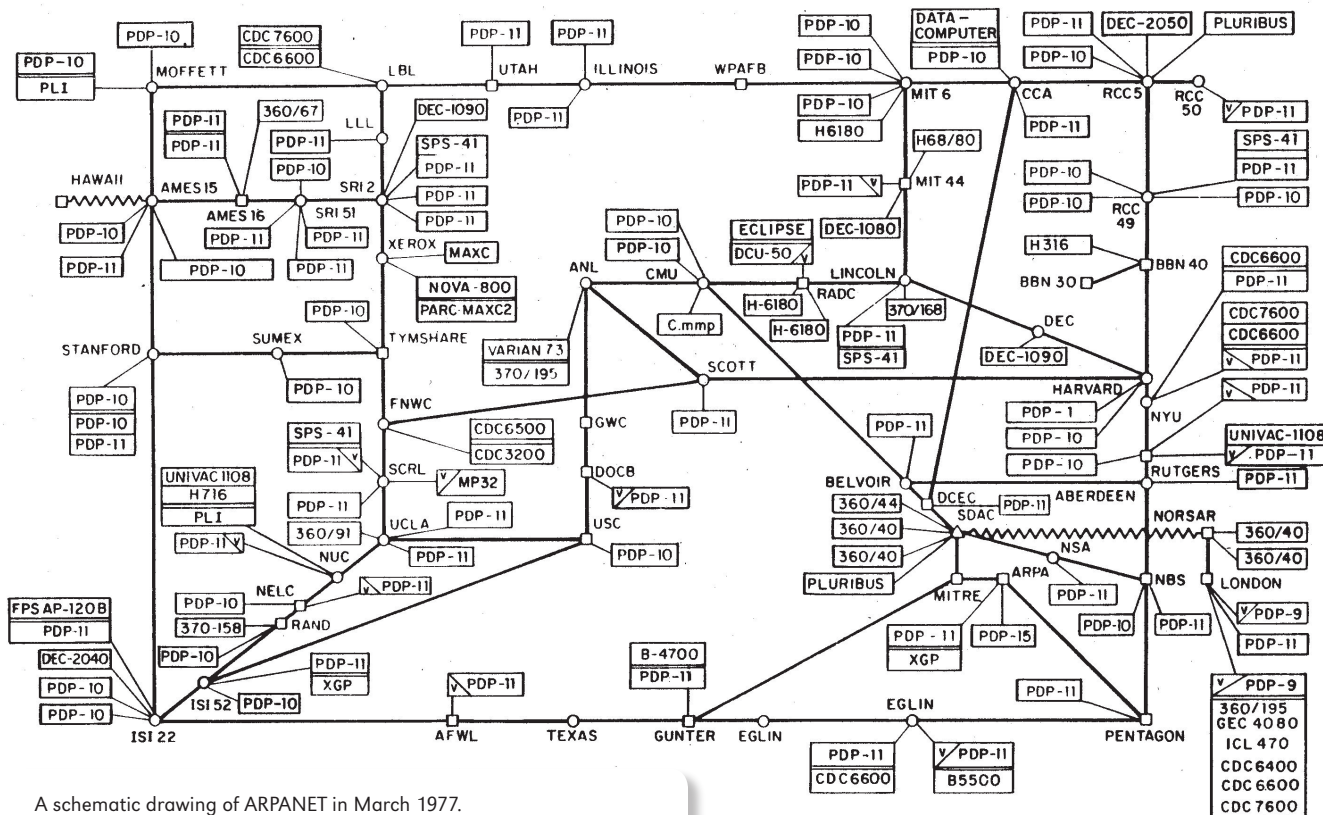
Then in the spring of 1967, at the suggestion of Wesley Clark of Washington University, Roberts added a key refinement to this scheme. Instead of asking each site to run the packets right through their main computers, which was like trying to run an Interstate highway right through the main street of every little town in its path, the network

would be the digital equivalent of a limited-access highway. Just outside each town would be an “interchange” in the form of a small computer – dubbed the Interface Message Processor (IMP) – that would handle all the routing chores. The result would be a clean interface: ARPA would take responsibility for designing and implementing the network proper – meaning the information highways and the digital interchanges – while the researchers at each site would focus on the comparatively simple task of linking their central computer to IMP.

Roberts later discovered that a very similar networking scheme had recently been developed by Donald Davies and his group at the National Physical Laboratory in the United Kingdom, although they never got the funding to implement it. Fortunately, however, he had no problem with funding at ARPA; the leaders of the agency continued to be very supportive. In late December 1968, Roberts awarded a contract to build the IMPs to Bolt Beranek and Newman (BBN) of Cambridge, Mass. And, in a heroic effort that has since been chronicled in books such as *Where Wizards Stay Up Late*, and *The Dream Machine*, BBN delivered the first of the machines in just nine months. Kleinrock and his research group at UCLA saw that first IMP installed in their laboratory on Labor Day, 1969, making them the first node of the new ARPANET.

Installations at the other ARPA sites soon followed, with Roberts, who took over from Taylor as IPTO director in the fall of 1969, constantly pushing the researchers to make full use of the network. That job became much easier after 1971, when BBN engineer Ray Tomlinson devised an e-mail program for the ARPANET. (Tomlinson, who had gotten the idea from local e-mail utilities that had long since been implemented at Project MAC and the other ARPA time-sharing sites, also came up with an elegant new way to define e-mail addresses: take the “username” that the person typed when logging in to his or her host computer, and link it to that computer’s “host name” on the network with an “at” symbol: (username@hostname). E-mail quickly became by far the most popular application on the network – even at ARPA headquarters, where director Steven Lukasik was a particularly avid user.

In 1972, Roberts hired Robert Kahn to oversee IPTO’s new networking initiatives, which by then included experiments in packet-switched communications via both radio and satellite. Kahn, who had been a member of the original IMP team at BBN, soon found himself tripping over one of those trivial-seeming issues that ultimately turn out to be the most profound of all. How was he going to get all these satellite networks and radio networks talking to one another, and to the ARPANET? The various networks were optimized for very different environments. The ARPANET, for example, lived in a world of comparative stability, with packets flowing over fixed, reliable, land-based telephone lines. But packet radio lived in a world of chaos, with mobile transceivers forever moving in and out of range, getting cut off by hills and tunnels, and generally losing packets at every turn. The two systems had major incompatibilities in transmission speed, packet length – almost any parameter you could name.



A schematic drawing of ARPANET in March 1977.


Kahn realized that the best way to integrate these networks was to start by disintegrating them. Take the satellite portion, say, and make it into a completely separate network, in much the same way that the airlines are completely separate from the highway system. Give it its own IMPs, its own software, its own transmission protocols, everything. Then connect it to the ARPANET through a kind of gateway – or in modern parlance, a “router” – a specialized computer whose sole job would be to translate ARPANET packets into satellite packets and vice versa. As long as both sides met the interface standards – which were still to be defined – neither side would have to know anything at all about the internal details of the other. You could simply plug them together through a gateway, like plugging an appliance into a standard electric socket, and the packets would flow as needed. Moreover, you could just as easily plug in a packet radio network, or any other type of network that might come along. (A prime example of the latter was Ethernet, a high-speed, packet-switching, local-area network that was developed at the Xerox Palo Alto Research Center in 1974, inspired in large part by the ARPANET.) The result, Kahn realized, would be a system that was completely open: a network of networks that could, in principle, accommodate anyone.

Kahn worked out the basic architecture of this “inter-networking” scheme in collaboration with Vinton Cerf of Stanford University, who had been a graduate student in Kleinrock’s UCLA group when the first IMP was installed there. And he gave Cerf the job of defining those basic interface standards.

Cerf’s solution, the simplest he could think of, was inspired by an analogy to the postal system. As things stood, trying to send a packet from one network into another was like trying to send a postcard written in Japanese Kanji characters through a post office in the United States: Nobody there would know how to read it, and it would go straight to the dead-letter office. But this happened only because postcards traveled naked, as it were, never changing their outward appearance to fit with changing conditions. So instead, Cerf wondered, why not arrange for each postcard to be mailed inside an envelope addressed in the local language – and then, when it crossed the border, moved into a new envelope addressed in the new local language? End of problem: The local mail sorters would read Kanji in Japan, English in the United States, Arabic in the Middle East, and so on.

Now, to translate this idea into networking terms, Cerf reasoned, imagine that every computer started with a specific, universal protocol for addressing its packets. (This would be the software equivalent of having everybody agree to address their postcards on the right-hand side, using the Roman alphabet, with the name, street, city, state, country, and so on in a standard order.) But before launching the packets straight out into the local network, which might not understand the universal protocol, the computer would wrap each one inside an “envelope” of extra bits that the local network would understand. In that form, each packet could sail through its home network until it reached the gateway to the next network in line – whereupon the gateway computer would take each packet out of its envelope, read the address written in the universal

This July 1977 map represents a much-expanded ARPANET.



protocol, wrap the packet in a new envelope appropriate to its new surroundings, and send it on its way.

The universal protocol itself would have to deal with a number of practicalities, including an “inter-networking protocol” that would encode such things as a packet’s ultimate address, and a transmission control scheme that would allow the destination computer to request replacements for packets that had been lost in transit. Thus the modern name: Transmission Control Protocol/Internet Protocol, or TCP/IP. So Cerf’s networking research group at Stanford actually spent a year working out the details of the idea. Then in 1974, he and Kahn published a paper entitled, “A Protocol for Packet Network Interconnection,” the first architectural description of how the Internet would function as a network of networks, with TCP/IP holding it all together. (That paper is why Kahn and Cerf are so often hailed today as the inventors of the Internet, to the extent that any two people can be singled out for that honor – and why, in 2005, they were presented with the Presidential Medal of Freedom.)

The challenge now was to get this “ARPA Internet” working for real. In 1976, Kahn persuaded Cerf to join DARPA – as it had been renamed in 1972 – to take on that challenge. (Kahn himself would serve as IPTO director from 1979 to 1986.) During the next several years, Cerf oversaw the development of operational versions of the protocol running on multiple hardware platforms, as well as a series of increasingly ambitious field demonstrations. By March 1982, TCP/IP version 4 was deemed reliable enough for the Department of Defense to make it

the standard for all military computer networking. And on Jan. 1, 1983 – a date that arguably marks the real-world birth of the Internet – the ARPANET itself switched over to TCP/IP.

By that point, moreover, the ARPANET was no longer alone. Since TCP/IP was in the public domain – having been developed at taxpayer expense – other networks were beginning to use it as well, and to link with one another in an expanding Internet. Easily the most significant of these TCP/IP-based networks was the National Science Foundation’s NSFNet, launched in 1986. NSF’s original intent was to link academic researchers to a new system of supercomputer centers, which it had announced the previous year. But because NSFNet was the first network available to every researcher on every U.S. campus, its usage expanded exponentially. By decade’s end, NSFNet had become the de facto U.S. backbone of the fast-emerging Internet, and had laid the foundation for the Internet’s explosive worldwide growth in the 1990s.

Indeed, NSFNet effectively made the ARPANET itself obsolete. DARPA had been pulling back in any case. The ARPANET had long since ceased to be a research project, and the agency had no wish to fund long-term operations; that was not what DARPA was for. And with NSFNet – along with the just-emerging commercial networks – available to DARPA-funded researchers along with everyone else, there was no longer any real need for it.

In 1989, the ARPANET was formally ended. But the larger Internet it had given birth to was just beginning.