

Two Early Interactive Computer Network Experiments

David Hemmendinger
Department of Computer Science
Union College
Schenectady, NY 12308
hemmendd@union.edu

Abstract

Two early networking experiments joined a time-sharing computer at the System Development Corporation with systems at the Stanford Research Institute briefly in 1963 and at MIT Lincoln Laboratory in 1966-1967. Both were influenced by J. C. R. Licklider's interest in resource-sharing and included experiments with the interactive use of remote programs.

Keywords C2: Networking, Remote Applications

Introduction

The second half of the 1960s saw considerable research in computer networking in the USA and in England. The best-known network is the Arpanet, which began operating in 1969, while in England, Donald Davies planned a national network, the first small version of which was roughly contemporaneous with the initial Arpanet tests.¹ Both of these networks used packet-switching, proposed by Davies and by Paul Baran of the RAND Corporation, which became the norm for long-distance computer networks. There were earlier network experiments, however, most of which have received little documentation. This article describes two of them, both using a time-sharing computer at the System Development Corporation (SDC) in Santa Monica, California. The first was a short-lived connection in 1963 between SDC and the Stanford Research Institute (SRI), while the second was a better-known experiment that connected SDC with the TX-2 computer at MIT Lincoln Laboratory (LL) in Cambridge, Massachusetts in 1966-67. It is often cited as the first long-distance network connection, and an online book calls it "the seminal experiment," though it misdates it as 1965.² One reason for its being better-known may be that it was conducted by Larry Roberts, who then became the head of the Arpanet project. It is interesting, however, that in each case, the results of the experiment were never published except in internal technical reports and memoranda.

Background: resource sharing

In the late 1950s, several researchers proposed versions of time sharing to allow multiple teletype or other terminal connections to a computer to be active concurrently.^{3,4} This would allow several users to test and debug programs on line, and more generally, to allow interactive sessions while still using computer time efficiently. Time-sharing systems appeared in the early 1960s, including one at SDC on its AN/FSQ-32 (Q-32) system.⁵ J. C. R. Licklider, director of the ARPA Information Processing Technology Office (IPTO), 1962-64, was a major proponent of time sharing and used his position to fund time-sharing projects. His interest in it was part of a larger interest in exploring human-computer interaction, and it came to include an interest in networking computers together for the sake of resource-sharing and collaborative work.

Computers were networked to exchange information at least by the late 1950s, when the SAGE (Semi-Automatic Ground Environment) system used computer communication to coordinate radar information.⁶ These networked computers did not give operators the ability to invoke user-level programs remotely. Rather, they transmitted digital information by telephone and radio, using a fixed set of programs. The network experiments that this article describes, by contrast, allowed a system user to invoke any suitable program on the remote computer.

In an April, 1963 memo to members of the "Intergalactic Computer Network," Licklider imagined working at one computer site and bringing in programs from other computers. While his notion of importing binary programs from remote computers to his was problematic, since they would be architecture-specific, he also wrote

When the computer operated the programs for me, I suppose that the activity took place in the computer at SDC, which is where we have been assuming I was. However, I would just as soon leave that on the level of inference. With a sophisticated network-control system, I would not decide whether to send the data and have them worked on by programs somewhere else, or bring in programs and have them work on my data. I have no great objection to making that decision, for a while at any rate, but, in principle, it seems better for the computer, or the network, somehow, to do that.⁷

Histories of the ARPA IPTO office such as *Transforming Computer Technology* by Norberg and O'Neill⁸ describe its work as being guided by Licklider's vision of resource sharing and collaborative work, a vision that his successors, Ivan Sutherland and then Robert Taylor shared. Time sharing, one part of the realization of that vision, had its origins in the interests of academic computer centers in promoting effective computer use in the write-test-debug cycle of program development. Initially, according to Norberg and O'Neill,⁹ a computer network was a time-shared computer with remote terminals, such as the SDC Q-32 and the MIT Compatible Time Sharing System had. Unlike the development of time sharing, computer networking as communication among multiple computers appears to have been initiated largely through IPTO funding and encouragement to connect computers together. It is perhaps not surprising that networking was promoted in a top-down fashion. While one institution could become more efficient by time sharing its computer, networking computers at several institutions might be seen as letting others use one's scarce resource. In fact, reports on the early reception of the idea of an Arpanet indicate just that. Robert Taylor reported about a 1967 ARPA network meeting saying, "Most of the people I talked to were not initially enamored with the idea. I think some of the people saw it initially as an opportunity for someone else to come in and use their cycles."¹⁰ Later, as we know, people became enthusiastic about the opportunity for collaborative work, confirming Licklider's expectations.

Other histories of the development of the Arpanet and the Internet such as Abbate's¹¹ and Hafner and Lyon's popular account¹² support the notion that computer networking was spurred by the ready availability of IPTO funding. The remainder of this study will look at how that worked in two instances, and will look briefly at some other possible early network projects for which evidence is less clear.

Sources

Several articles on the SDC Q-32 time-sharing system mention that it had a network connection, but there are only a few published mentions of its use, none detailed. These include a book and a lecture by Charles Bourne, along with an SRI technical report. The 1966 SDC-LL connection is described in the FJCC article by Marill and Roberts. Both experiments are also mentioned in *A History of Personal Workstations*¹³ and in a UCLA Ph.D. dissertation.¹⁴

Unfortunately there is no comprehensive archive of SDC technical documents. The Charles Babbage Institute has some in its archives of the Burroughs Corporation, which acquired SDC, but none describe network experiments, though they are mentioned in an annual research report for 1966. More is available from Lincoln Laboratory, both in online quarterly research reports at the Defense Technical Information Center, which are quite sketchy, and in some documents in the LL archives. I have also consulted several oral histories by Robert Taylor and Larry Roberts, and I have benefited from emailed correspondence with people who worked on the two

projects.

The SRI-SDC experiment

The SDC Q-32 ran an early time-sharing system (TSS) that went into operation in mid-1963. The 1964 paper that presented the TSS¹⁵ very briefly mentions a network connection with a Control Data CDC 160A at SRI, over three hundred miles away. The actual connection was between the 160A and a PDP-1 that was the communications front-end for the Q-32. The network was used by an SRI program to do full-text searches of a bibliographic database stored on the Q-32. Its designer, Charles Bourne, worked on information retrieval with Engelbart and was later director of the Univ. of California Institute of Library Research. He described the project in an SRI technical report,¹⁶ the 1999 Conrad Memorial Lecture,¹⁷ and a 2003 book.¹⁸ It was a proof-of-concept system intended to demonstrate the feasibility both of full-text search for strings of words or bibliographic terms, and was part of Douglas Engelbart's Augmenting Human Intellect project that was supported in part by IPTO. According to a later account by Engelbart,¹⁹ they had proposed to use an interactive workstation connected to a local computer, but Licklider asked them instead to put the database on the SDC Q-32 system, and build a networked connection to it (IPTO had also funded the Q-32 time-sharing system). This interactive computer-computer network ran the search software on the remote Q-32 with the SRI CDC machine providing the user interface. The search program with its remote connection was demonstrated successfully on a small set of documents in late 1963, though some of its planned capabilities, such as Boolean operators in search strings, were not implemented. Its demonstration was supposed to have used 50 records at SDC, but defects in many of them meant that only seven records were actually used. Work on the system was apparently abandoned shortly after the demonstration, as the project lost funding. In his lecture, Bourne said that SRI projects "went where the money was" and that they didn't find funding agencies for the bibliographic project after SRI funds went to projects related to the Cold War. According his book, although he was the project manager, he had forgotten about the never-named project until he was conducting research for the book over three decades later.

Engelbart's reference to Bourne's short-lived project is the only other published mention of it that I have found. He added to his account, "For various reasons, not uncommon in pioneering ventures, that first year was unproductive relative to the purposes and plan of our project." He also said that, although Licklider was willing to continue support, he couldn't offer enough funding. Figure 1 is a photograph of that interactive workstation, from Bourne's report, which also mentions using a five-key "chordal" keyset to allow one hand to type while the other used a light pen.²⁰ It appears to be similar to what Engelbart used in his 1968 computer mouse demonstration.²¹

According to Bourne, the network ran at 2000 bits per second, using an AT&T 201A modem (on a 2 kHz line, according to Schwartz)²². If it used the then-new ASCII encoding, it would have run at about 200 characters/second. Although there are no further details of the network implementation in published accounts, interviews of programmers at both SDC and SRI provide further information.

Clark Weissman, one of the Q-32 system programmers, said that the network between SDC and SRI used two full-duplex telephone lines, one for commands to the Q-32 (e.g. 'login', 'load', 'go') and the other for data to the program that the commands started.²³ (He noted that this arrangement was similar to CPU design, which separates data and control.) Both lines were connected to the PDP-1 that handled all terminal I/O to the system. As he remarked, the flaw of this design was its requiring two lines, but the virtue of the design was its also requiring minimal modification of the operating system. Since command and data streams were separate, the operating-system executive (command interpreter) could operate without change, except that it would have to be able to treat the two phone lines in effect as separate terminals, so that it would first

accept input from the command line, and when commands started running a program on the Q-32, that program would be assigned the data-line "terminal" for its input and output. Since the PDP-1 managed the terminals, any changes would have been made to its operating system rather than to the Q-32 system.



Figure 1: Control Data CDC-160A workstation at SRI (Courtesy of SRI International).

Len Chaitin, the SRI programmer who worked with Bourne to implement the full-text search program, also recalls the two-line design.²⁴ He and Weissman both said that the network connection worked well; Chaitin said that a small amount of bit-checking was done in hardware. He said that the network design was intended to be as simple as possible so that it might be used by others later, though that didn't happen. It is striking, however, in view of the difficulties that others found in using transcontinental telephone lines, that this simple arrangement was effective over a distance of several hundred miles.

I have found no SDC documents that describe either the SRI-SDC experiment or any further uses of the 2000 bps link. It is still listed in the 1965 SDC annual research report,²⁵ which also cites a final report on its design,²⁶ but that document is not in the Babbage Institute SDC collection nor have I found it elsewhere.

Statements of "first" need to be carefully made, and they are generally not the most important characteristic of the object of the claim. Nonetheless, the 1963 SDC-SRI network appears to be the first one that was designed and employed for interactive use, albeit briefly. There were earlier computer-computer networks, e.g. in the SAGE system, that were used for data transfer, and it is possible that people found clever ways to use them interactively, but they were not designed for that purpose. More important, the SRI-SDC network was an early instance of a network set up to demonstrate the possibility of harnessing dissimilar computers in a single project that might take

advantage of the strengths of each — data storage on the Q-32 and interactive terminal use on the CDC system. It is also an early example of how Licklider's interest in networking directly led to an experiment.

The LL – SDC experiment

A 1967 review of the TSS performance²⁷ calls the SDC-SRI network obsolete, and refers to a connection between the Q-32 and the TX-2 at Lincoln Laboratory. This is the one described by Thomas Marill and Larry Roberts,²⁸ often called the first interactive computer network connection. It is thus surprising that although their paper gives details of a planned network experiment, its results weren't published. There have been just a few mentions of them, such as by Roberts in a retrospective on the Arpanet,²⁹ which said that communication was possible, but that the telephone lines were too slow and unreliable. That remark is unlike what participants recalled about the SDC-SRI link, which was that it was slow but reasonably reliable; perhaps it reflects higher expectations of networked computing three years later.

The origin of the Marill-Roberts network experiment is unclear. According to Norberg and O'Neill's study of U.S. military support for information processing,³⁰ Marill, who had just started the Computer Corporation of America (CCA), was looking for a contract in 1965 and gave IPTO a proposal for a network study, which IPTO set up as a subcontract under its Lincoln Lab contract.³¹ Other accounts, such as the DARPA history of the Arpanet,³² just refer to the project as being a subcontract of the ARPA contract with LL.

Marill's role in computer networking appears limited to this one experiment. He was a psychologist who had worked with Licklider for his MIT Ph.D. He subsequently joined Licklider at Bolt Beranek and Newman (BBN), where he worked on the psychology of perception, information retrieval, time-sharing, and artificial intelligence.³³ At CCA, Marill later worked on a variety of projects, including databases, information retrieval, and image recognition. In the late 1980s and 1990s, he was a researcher at the MIT Artificial Intelligence Laboratory, working on robotics and perception.

Roberts is well-known as the head of the Arpanet project, which he took over at the end of 1966, at the request of Robert Taylor, the third IPTO director.³⁴ He had attended a conference on the future of computing at Hot Springs, VA in late 1964, where conversations with Licklider persuaded him of the importance of computer networking.³⁵ He remarked about that meeting in an interview, "Well, I was interested in communications, but I didn't have any strong direction at that point, or before that point,"³⁶ and at Lincoln Lab, he continued to work on computer graphics. He knew Marill, and in 1965-66, was contract manager for Marill's project.

The Arpanet Collection³⁷ of the online Internet Archive,³⁸ has material collected by Katie Hafner for her Internet book. It includes a set of letters that Marill gave her about his contract with Lincoln Lab.³⁹ His May, 1965 proposal to LL for a networking project elicits a non-committal response, leading him to write to Roberts in more detail in June about his reasons for wanting to do the study, in which he proposes to connect three computers. That is followed by a Laboratory request in July that he submit a formal proposal. It is not clear from the letters who had initiated the May proposal. It is unlikely that Marill would have proposed to use the SDC Q-32 without having been given reason to think that it would be available, either as a response to a proposal that he gave the IPTO or in a request that it had initiated. In the last letter in the collection (November 1965), Marill asks Roberts for a six-month contract extension, referring to a discussion that they had had a week earlier.

Marill's June, 1966 report³¹ was written during that six-month period; a third draft is dated 28 February,⁴⁰ and is nearly identical with the final version. Its contents reflect the early state of networking then. Like the plan for the later Arpanet, it presents resource-sharing as a principal reason for having a network: to make specialized systems or programs accessible from any time-

sharing system. It is presented largely as a feasibility study, and devotes considerable space to a discussion of character coding and transmission and of available modems, with their costs. It explores telephone line options and recommends several, with the choice depending on usage. The one that the later Marill-Roberts paper proposed to use, a 4kHz Western Union voice-grade line with a 1200 bits/second modem, was its recommendation for up to 50 hours/month of use.

The most interesting part of the report is its software section. Like the later Marill-Roberts paper, it outlined three strategies. Two required little change to the operating system. The first "quick-and-dirty" approach needed only to let a program communicate with two terminals, the user's and the remote computer. A user would start a program on the local system, and it would communicate with the remote computer first to log in and then to run a program on it, exchanging data with it as a user might do directly. As Marill remarked, the only other system change would be to provide a command to let the user program dial a telephone number.

The report added that without further change, such a connection would operate at teletype rates (11 characters/second). A higher-speed modem and line could increase the rate; Marill suggests that 100 characters/second would not strain the current hardware, though their processing time might degrade service on other lines. In fact, even steady transmission at that rate would have required no more than 1-2% of the SDC PDP-1 computational capacity to process the input, according to the description of interrupt-handling in the PDP-1 manual.⁴¹

A second method would use separate lines for commands and data. With a single channel, the program or the operating system must parse all input to distinguish commands and data. With two channels, the system need not examine the data line at all, since its content would go directly to a program. While using two lines might seem costly, Marill suggested that it need not be since only one of them would need to be high-speed, and the command connection would generally be made only briefly. We should recall that just such a two-line design was used for the SDC-SRI network. In his final section of recommendations, Marill mentions the SDC-SRI network as a data-only link that was supplemented by a second line for commands, so he was evidently aware that the two-line design had been implemented.

The third method, mentioned as speculation, is to use a single high-speed line for both commands and data. It would need a protocol to distinguish commands from data, and Marill suggests only that possibly every n th bit or character might belong to a command stream, so that the operating system would only need to count bits or characters to be able to recognize commands — a rather awkward design.

The three computers that Marill proposed to use were the TX-2, the Q-32, and an IBM 7094 run by the MIT Project MAC for research in time sharing. He discusses the advantages of resource-sharing between LL and Project MAC, but he doesn't say what led him to expect to use the 7094, and it never became part of his network experiment. According to Norberg and O'Neill, "the MAC people were not interested in participating;"⁴² another account says that IBM didn't want to get involved.⁴³ Marill's report says that the MAC 7094 had an IBM 7750 computer as a communications front end (like the PDP-1 at SDC), which had 1200 bps lines, and his plan required no changes to MAC hardware or software. Once the experiment added a communication protocol, however, system software would need some change to be able to interpret it, and it is possible that that was a basis for non-participation, but we have no definite information. A 1967 memo by Marill just says, "for various reasons, this never came to pass."⁴⁴

Lincoln Laboratory quarterly reports for 1966 and 1967 provide brief documentation of the rather slow progress in setting up the network during the year after the APEX operating system for the TX-2 was in use.

The hardware, software, and operating requirements of a low-speed, multiplexed data-terminal sequence for TX-2 are being studied. Initial requirements are for a system which can communicate with other computers via "dialed-up" or leased phone

lines at either teletype or 2000-bit/sec rates.⁴⁵ (Nov. 1965 - Jan. 1966)

The logic design for the in-out unit required to interface telephone-line data terminals with TX-2 has been completed. ... Western Union Broadband Switching Service, a data set for 1200-bit-per-second asynchronous operation, and an automatic answering unit are being obtained for the data connection to System Development Corporation in Santa Monica, California.⁴⁶ (Feb.-April 1966)

The [TX-2] computer now has a low-speed data channel which will be used for communication with remote consoles and with other computers. The first phase of the software required to use this hardware with the APEX time-sharing system is now being checked out. APEX has been changed to allow two new consoles, and the basic routines for sending and receiving streams of characters are ready. The first remote connection will be with the FSQ-32 computer at the System Development Corporation in Santa Monica, California.⁴⁷ (Aug.-Oct. 1966)

The network link between the TX-2 and the AN/FSQ 32 at SDC has been exercised with mixed success. Significant demonstrations are planned for the next quarter.⁴⁸ (Nov. 1966 - Jan. 1967)

The network link between the TX-2 and the AN/FSQ-32 at SDC has been used for several demonstrations. A distributed program which uses the Lincoln Reckoner at TX-2 and LISP at SDC has been created and run. Statements typed in at TX-2 are parsed at SDC and then given to the TX-2 Reckoner for calculation. Some statistics on network operation are being gathered.⁴⁹ (Feb.-April 1967)

The principal addition to Marill's CCA technical report that the Marill-Roberts paper provides is the description of the communication protocol to be used. It broke messages into blocks of up to 119 characters, each with a start-character denoting a message for system or user, a stop-character, and a checksum. There were also 'ACK' and 'NACK' characters for acknowledgments and resend-requests, and several other control characters. Roberts later referred to these message blocks as packets, though of course there was no packet-switching.⁵⁰ This protocol, with relatively small message blocks, would have permitted error-recovery with the retransmission of a block or of a few blocks, without the expense of having to retransmit an entire long message, though at the cost of an addition of a small number of control characters, an overhead of a few percent at the most. In a 1988 interview, Roberts remarked that it was natural to think in such terms: "We were all thinking of blocks That's the way computers worked."⁵¹

A CCA memo in the Lincoln Lab archive proposed a more efficient protocol, which would let a second message block be transmitted without waiting for the previous one to be acknowledged, although the second block would not be terminated before that acknowledgment arrived. If the receiver reported an error in the first block, the second would be terminated immediately and the first resent.⁵² This protocol would permit successive message blocks to be transmitted without delay as long as each 'ACK' for a block was received while the next block was being sent. As the memo author, Bill Mann, now recalls,⁵³ the SDC people in charge of their end of the network preferred not to implement this more complex protocol. Instead, they used the simpler one that is in the Marill-Roberts paper and is also documented in a CCA memo.⁵⁴ Mann said that he objected to its inefficiency, and Roberts' reply was that the project was just a proof-of-concept test and that performance didn't matter, an interesting contrast to Roberts' later comments about the network's being too slow to be practical, though of course both responses are appropriate in their contexts.

A June, 1966 memo to Roberts⁵⁵ outlines a proposed schedule for the experiment: in June and July, set up the TX-2 phone hardware, modify the APEX operating system to link to the Q-32, run a test program that has the modem at SDC echo characters back to the TX-2, and write a program to make a TX-2 terminal look like a Q-32 console, and in August run a demonstration

TX-2 program that uses an SDC program "in a meaningful way." It mentions further phases of the project, and remarks that the timetable may need revision. There is no indication whether that was necessary.

Two progress reports describe further plans and what was realized. The first is about a January 1967 meeting in which Marill reported on network plans and lists 16 experiments.⁴⁴ Five are to be TX-2 programs that use the Q-32, four are Q-32 programs that will use the TX-2, and seven are to use a DEC 338 graphics terminal (a modified PDP-8) at Harvard, to connect to one of the other two computers. The most ambitious ones are to use graphics displays, and the memo notes that these tests will require interim measures that violate networking principles by requiring changes in the remote programs and by not using a standard display language. Both the Q-32 and the TX-2 had graphics projects (including some by Roberts), making them candidates for remote experimentation.

The report commented that three TX-2 programs that used the Q-32 had been run, but that flaws in the modems made them unreliable. Work on three counterpart programs on the Q-32 had started, but none of the remaining ones, including the graphics programs, had been started. A DEC 338 didn't become part of the network experiment, though one was set up later in Roberts' ARPA office.⁵⁶

One of the three TX-2 programs that used the Q-32 was a console simulator that let one run remote programs. A second let a TX-2 user run TINT, an interactive JOVIAL language interpreter, as if it were local to the TX-2. The third, AT (Algebraic Translator), was the most interesting: it took mathematical expressions in standard infix form from the Lincoln-Writer, a special typing terminal, sent them to a LISP program on the Q-32, which translated them into prefix notation and returned them to the TX-2 for evaluation by the Lincoln Reckoner program (called Basic Translator here), a calculator for scalar, vector, and matrix mathematics.⁵⁷ Marill credits Weissman with the idea for this program.⁵⁸ According to an April 1967 progress report,⁵⁹ these three programs ran reliably. Six of the remaining ones were canceled and another six, were being reconsidered, and only a Q-32 console simulator for the TX-2 was still being planned, though deemphasized. There is no indication that any more programs were completed.

There is less information about the SDC side of the network. The Charles Babbage Institute collection of Burroughs Corporation papers includes some technical memos from SDC, which Burroughs had acquired, but there does not appear to be anything among them about the network. The SDC 1966 annual research report⁶⁰ has a section by Weissman on programming systems, which mentions two network projects. One, for a network of SDS 940 computers at the University of California at Berkeley, BBN, and SRI as well as at SDC, was under design, but apparently never implemented, while the other was the TX-2 connection. The report emphasized its allowing a system to invoke programs on another, dissimilar one. It added that this method "could be the only way for the data processing community to truly share the programs and work of others." A section on time-sharing networks lists the Western Union 1200 bits/second line that Marill described, and referred to collaboration with CCA in modifying both the SDC PDP-1 and the TX-2 operating systems, which is consistent with the CCA documents. Its explanation of the lack of Q-32 programs to use the TX-2 was that the resources of the APEX operating system on the latter were largely devoted to display programs, and that protocols needed to be developed for such programs. The Q-32 was due to be replaced by an IBM 360 in 1967, and though the report suggested that the network project would be transferred to the new system, there is no indication that that happened.

Clark Weissman has said that the network did use separate lines for commands and data, a slow one for the former and a fast one for the latter.²³ This would have been a natural thing for SDC to have done, having already used that simple design for its earlier network. There is no mention of a two-line network in extant SDC documents, however, and among the Lincoln Lab

material, the only mention of what could be a two-line network is the reference to a low-speed channel for connections to other computers in the November 1966 LL progress report, above, which also refers to there being two new consoles on the TX-2. This is scant support, though if it means that the two consoles were for the low-speed and the 1200 bits/sec telephone lines, it would be consistent with what Weissman recalls.

Weissman also recalls that Roberts did not like the cost of two cross-country lines, and didn't think that the two-line model was viable for networking. It is thus quite possible that a two-line network was initially tried, but wasn't developed to the point of being documented in progress reports since it was due to be replaced by a single line with a protocol to separate commands and data.

It is not clear how long the LL-SDC network continued to be used. Roberts had left for ARPA at the end of 1966, though in a paper presented at the first Symposium on Operating System Principles in Gatlinburg, TN in October 1967,⁶¹ he mentioned that the network "is now utilized by users to increase their capability." There are no documents about the network in the Lincoln Laboratory archive for the rest of 1967. A February 1968 CCA memo again described the three programs that had been implemented and gave examples of the algebraic translator program that used the Q-32 Lisp, though by then the network was not in use since SDC had retired the Q-32.

Some of the earlier documents refer to plans to set up automatic telephone dialing from the TX-2. Marill's memo on the January meeting⁴⁴ says that the dialer wasn't yet ready, but a March 1967 CCA memo⁶² reports that it has been done, and that the TX-2 can dial out when a network call is made and can also receive calls. Figure 2 is a block diagram of the TX-2 system, though it doesn't indicate if it has automatic dialing capability.

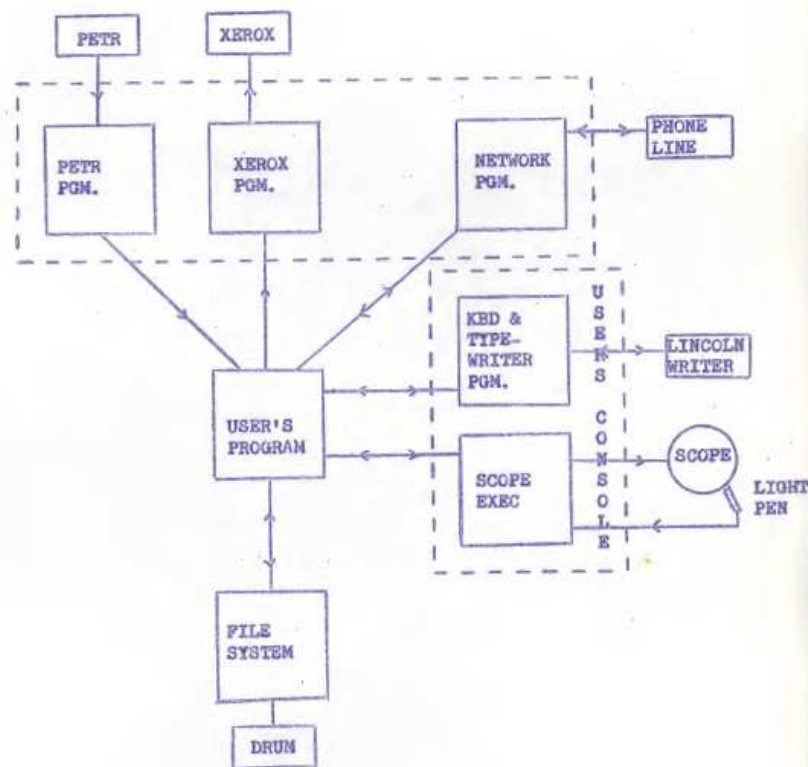


Figure 2. The network and TX-2 system (Courtesy of the MIT Lincoln Laboratory archives).⁶²

Marill's last progress report,⁵⁹ has an example of a session with automatic dialing. It shows the AT program being run on the TX-2, followed by his explanation of it (boldface indicates user input; the rest is output either from the APEX system (e.g. "DIALING SDC"), or from the Q-32 (e.g., the "Lisp 1.5" line):

```
LOG IN CCA
      CCA IN  425 PGS , THU 16 FEB 67 1320.13
      BTR  2

AT
DIALING SDC
CONNECTED
1967 FEB 16 1021.9 LISP 1.5 M2.6 7JA13 RL1447
      (EX CD)
      TLST
      BOPS
      UOPS
      (COMP EVBT TEST)

X=7.012345
(SIN X)*(SIN X) +(COS X)*(COS X)
      1.0000000
```

Typing in "CCA" logs us in at 1320.13 Eastern time.

Typing "AT" causes this program to start. It dials up SDC, logs in, loads LISP, transmits a LISP program for compilation by the LISP compiler, waits for the compilation to take place, and then waits for input from the Lincoln Writer.

LISP was loaded on the Q-32 at 1021.9 (Pacific time), or 1.8 minutes after log in on the TX-2. This 1.8 minutes includes the time to select the desired program, "AT," and to type this selection in (I don't recall that any one was rushing at this point). By 1325.1 Eastern time, or 3.2 minutes later, the LISP program to be compiled had been transmitted, and the compilation had taken place.

We now typed in a value for X, and a function to be evaluated. We do not have timing information here but estimate that the answer was typed out in less than one second. This includes the time to ship the algebraic string to the Q-32, have it compiled into Basic Translator [Reckoner] code, have the code shipped back, have the function evaluated, and the answer typed.

The CCA memos list documentation of the network and gathering statistics on its performance as tasks to be done, but there is no indication that they were carried out. The project did have two outcomes. It gave an example of a simple network connection and message protocol that could work, though its use didn't continue. Like the SRI-SDC network, it also demonstrated the use of remote computational resources. In Marill's example, once the program had been compiled, it gave results in a second or so, not a bad performance for transcontinental communication. It would be interesting to know how often there were network errors, as it would be also for the earlier SRI-SDC connection, but I have not found any record of that information.

Both SDC and LL documents refer to using graphics displays that they each had. The participants did not develop the necessary protocols, which ideally would have required a device-independent format for messages, to be translated into device-specific code at each end. Although this ambitious goal was not reached, what was achieved is noteworthy as an early example of the resource sharing that the Arpanet was also set up to do several years later.

Other proposed experiments

In their chapter on early IPTO networking contracts, Norberg and O'Neill mention several other proposed experiments, and this section adds a little more information. The first contract, in 1964, was with the University of California, Los Angeles, and was to join its main computing center, one at the health sciences center, and its Western Data Processing Center, a local consortium. It was also to involve the SDC Q-32 system, but the three UCLA centers did not want to collaborate with one another,⁹ and despite initial ARPA optimism, nothing ensued.

The 1966 proposal to connect SDS 940 computers at UC Berkeley, BBN, and SRI, mentioned earlier, was apparently intended to benefit from the convenience of networking one computer type. Although it did not involve an SDC system, a group met at SDC in June, 1966 to produce a design document, edited by Wayne Lichtenberger.⁶³ I have not been able to find a copy of the document, and its editor does not recall it, though he said that its topic was consistent with what he was working on then.⁶⁴ There is no indication that this proposed project was implemented, and I have not found any mention of Scientific Data Systems (SDS) participation.

IPTO also supported a project at Carnegie-Mellon (CMU) and Princeton Universities and IBM to network IBM 360/67s. A report on it⁶⁵ cites the Lichtenberger report as describing a similar project, presumably because both were intended to network a single computer model. The CMU project was implemented in 1968, at least as a prototype. It used point-to-point connections, unlike the Arpanet. The Network Working Group's Request for Comments (RFC 33) on the Arpanet host-to-host protocol mentions it as not having influenced the Arpanet, adding that

early time-sharing studies at the University of California at Berkeley, MIT, Lincoln Laboratory, and System Development Corporation (all ARPA sponsored) have had considerable influence on the design of the network. In some sense, the ARPA network of time-shared computers is a natural extension of earlier time-sharing concepts.⁶⁶

Possible 1965 experiment

Several online and published accounts, including the Computer History Museum's Internet timeline, give 1965 as the date of the SDC-LL experiment.⁶⁷ Although this is a mistake, there are also two oral-history interviews with Robert Taylor that mention a 1965 experiment. In a 2008 interview in the Computer History Museum's collection,⁶⁸ Taylor said that in 1965 he asked Larry Roberts to conduct a "bit reliability experiment" to help assess the feasibility of networking time-sharing systems, and that Roberts subcontracted the work to Thomas Marill. The experiment was to use SDC and Lincoln Lab time-sharing systems since, Taylor said, he was funding both. A 1989 oral history interview with him at the Babbage Institute¹⁰ has a similar account: that in 1965 he had asked Roberts about some network reliability issues, and that Roberts had had the CCA "send some bits back and forth over some phone lines" between SDC and Lincoln Laboratory. He emphasized in a recent email⁶⁹ that this test was not to connect the Q-32 and the TX-2 for networking applications, but only to use the two systems "as senders and receivers of bits" briefly, just long enough to give an idea of the reliability of cross-country transmissions. Taylor's interviews don't give a more precise date for this work. As of late 1965, he believed that cross-country transmissions were feasible, he said in his email. He thought that this belief was based on what came from SDC and Lincoln Laboratory, a view that is consistent with his role in initiating the Arpanet project in February 1966.¹⁰

Larry Roberts described this network test in email,⁷⁰ saying that in 1965, he had needed to determine the burst error rate on telephone lines to decide on an appropriate size for network message blocks. He said that he thought it most likely that he hadn't used a modem at all (the TX-2 didn't get one until mid-1966), but instead used an analog-to-digital converter to connect a phone line to the TX-2, a technique he had previously used to help Amar Bose measure the frequency

response of his living room in connection with loudspeaker design. He thought that he reported his results to Taylor.

None of the 1965 or 1966 Lincoln Laboratory reports mention this bit-transmission experiment. These recollections thus aren't enough to provide a definitive account of what may have been done in 1965. By both Taylor's and Roberts' accounts, the test would have been quite minor and brief, likely to have been too short-lived to have been documented. (Retrospectively, we are fortunate both that the outcome of any bit-shipping test suggested that a transcontinental connection was feasible, and also that its use led to a desire for something faster.)

Conclusion

I have described three network experiments that were motivated by the ARPA project of the first IPTO directors, J. C. R. Licklider, Ivan Sutherland, and Robert Taylor, to develop resource-sharing and collaboration among computer centers. They were short-lived and lightly documented. The 1963 SDC-SRI link is interesting as a first experiment, but largely as a "what might have been": a simple design that worked and that could have provided experience with linking dissimilar computers and using remote resources had it continued. It also gave SDC staff experience with networking, though there is little information about how that experience contributed to the later SDC-LL experiments.

Why is this experiment not better known, since it used one of the first time-sharing systems? One answer is that it ended after its initial demonstration. Another is, as Clark Weissman wrote in his email, that the SDC system was intended as support for the R&D done at sites like SRI and Lincoln Lab, although it did also do research on information-retrieval. According to him, the Q-32 system was a national resource that was used by a number of university researchers, generally through remote dial-up connections. While SDC people did publish several papers on time-sharing issues such as scheduling, they did not publish much on details such as the two-phone-line network setup.

The other two experiments used the 1966-1967 SDC-Lincoln Lab link. One was the trial of the simple communication protocol with the transcontinental telephone line. It led Roberts to become dissatisfied with its speed, cost, and reliability, and thus contributed to finding a better approach. The third experiment was the actual use of the Q-32 Lisp software by the Lincoln Reckoner, as a demonstration of the utility of such networking. Although nothing was published about this use of remote software, it is arguably the most important of the three, as a example of automatically invoking software on a remote computer and using its results locally. As far as I am aware, there were no further demonstrations of such use until after the Arpanet was set up in 1969.

These IPTO-funded experiments were largely the result of IPTO directors' promotion of computer networking. Some participants, such as Roberts, became interested through contact with Licklider, while others were pursuing their own projects: Bourne's work on information retrieval and Marill's search for a contract. To a large degree, the initial motivation for conducting network experiments was top-down; several accounts say that researchers with large computers were not immediately eager to share them. Histories that trace computer networking to the visions of a few people are thus more appropriate here than in many areas. Contingency certainly played a role as well — the confluence of events and needs that led Baran and Davies to develop packet-switching, Wesley Clark's idea of the IMP (Interface Network Processor) to simplify the Arpanet implementation, and the influence of Davies' group on Roberts at the 1967 Gatlinburg meeting when he presented the idea of the Arpanet, for example,

Perhaps the most important reason that the Marill-Roberts paper on their planned experiment is frequently cited as an account of the first long-distance computer network is that Roberts headed the Arpanet project shortly after its publication. As with so many questions of what was

the first X, the story is complicated. The early and briefly successful SDC-SRI network may have influenced a possible initial two-line design of the later network with Lincoln Laboratory, which could then have been used to investigate how to use a single telephone line with a communication protocol — an example of how a project that did not itself continue might yet contribute to a subsequent one.

Finally, a report of the graphics section of Lincoln Laboratory, where Roberts worked, describes the use of the SDC Lisp by the Lincoln Reckoner and says

The network contractor [CCA] has been asked to investigate the service facilities needed to make this link useful to general users. Until the problems of documentation, service facilities and instructional methods are solved, it is unlikely that a remote networked computer will be of much practical value to the general user.⁷¹

We have not yet completely solved the problems of documentation and of instructional methods, but the operation of the Arpanet and of the Internet that followed it have shown that their facilities have indeed been of considerable practical value to a great many users.

Acknowledgments

I thank Ed Coffman, Bill Mann, Larry Roberts, Robert Taylor, and Clark Weissman for discussions of networking, and Charles Bourne and Len Chaitin for information about the SRI project. The staff of the Charles Babbage Institute provided a copy of the fifty-year old CCA technical report, and the MIT Lincoln Lab archivist, Nora Zaldivar, provided the CCA memos. Comments by anonymous reviewers helped to improve the article.

References

1. Martin Campbell-Kelly, "Data communications at the National Physical Laboratory (1965-1975)," *Annals of the History of Computing* **9**(3/4) (1988).
2. James Pelkey, *A History of Computer Communications, 1968-1988*, <http://www.historyofcomputercommunications.info> (2007).
3. J. McCarthy, *A Time Sharing Operator Program for our Projected IBM 709*, Massachusetts Institute of Technology memo (January 1959). to Philip M. Morse.
4. C. Strachey, "Time sharing in large, fast computers," *Proceedings of the IFIP Congress*, pp. 336-341, Paris, France (June 1959).
5. D. Hemmendinger, "Messaging in the early SDC time-sharing system," *Annals of the History of Computing* **36**(1), pp. 52-57 (Jan-Mar 2014). Discusses some interactive uses of the SDC system.
6. Robert R. Everett, Charles A. Zraket, and Herbert D. Benington, "SAGE—A data-processing system for air defense," *Annals of the History of Computing* **5**(4), pp. 330-339 (October, 1983).
7. J. C. R. Licklider, Memorandum for members and affiliates of the Intergalactic Computer Network: Topics for discussion at the forthcoming meeting,, Advanced Research Projects Agency memo, Washington, D. C. (23 April 1963).
8. Arthur L. Norberg and Judy E. O'Neill, *Transforming Computer Technology: Information Processing for the Pentagon, 1962-1986*, The Johns Hopkins University Press, Baltimore, MD (1996).
9. Norberg and O'Neill, p. 156.
10. Robert W. Taylor, "Oral history interview with Robert W. Taylor," CBI OH 154, Charles Babbage Institute (February, 1989). <http://purl.umn.edu/107666>. Conducted by William Aspray.

11. Janet Abbate, *Inventing the Internet*, The MIT Press, Cambridge, MA (1999). pp. 48-49.
12. Katie Hafner and Matthew Lyon, *Where Wizards Stay Up Late*, Simon & Schuster, New York (1996). p. 68.
13. in *A History of Personal Workstations*, ed. Adele Goldberg, Addison-Wesley, Reading, MA (1988).
14. Gerald D. Cole, "Computer Network Measurements: Techniques and Experiments," UCLA-ENG 7165, University of California, Los Angeles (1971).
15. Jules I. Schwartz, Edward G. Coffman, and Clark Weissman, "A general-purpose time-sharing system" in *Proceedings of the April 21-23, 1964, Spring Joint Computer Conference*, pp. 397-411, Association of Computing Machinery (1964). DOI 10.1145/1464122.1464163.
16. Charles P. Bourne, "Research on computer augmented information management," Report No. ESD-TDR-64-177, Electronic Systems Division, US Air Force Systems Command, Bedford, MA (November, 1963). NTIS Report No. AD-432 098.
17. Charles P. Bourne, "40 years of database distribution and use: An overview and observation," *Miles Conrad Memorial Lectures*, National Federation of Advanced Information Services (February, 1999). <http://www.nfais.org/page/47-charles-p-bourne-1999>.
18. Charles P. Bourne and Trudi Bellardo Hahn, *A History of Online Information Services, 1963-1976*, MIT Press, Cambridge, MA (2003).
19. Douglas Engelbart, "The Augmented Knowledge Workshop" in *A History of Personal Workstations*, ed. Adele Goldberg, pp. 191-192, Addison-Wesley, Reading, MA (1988).
20. Charles P. Bourne, *Research on computer augmented information management*. p. 43.
21. D. Engelbart, *The Demo* (1968). <http://web.stanford.edu/dept/SUL/library/extra4/sloan/mousesite/1968Demo>.
22. Jules I. Schwartz, "The SDC time-sharing system, part 1," *Datamation* **10**(11), pp. 28-31 (November, 1964).
23. Email, 10-14 February 2013.
24. L. Chaitin, telephone interview by D. Hemmendinger, 18 March 2013.
25. Claude Baum, ed., *Research and Technology Division Report for 1965*, System Development Corporation, Santa Monica, CA (1966). Technical Memorandum TM-530/009/00.
26. Y. S. Loy, *System and logical design of a 2000 BPS data terminal unit for Q-32/PDP-1 operation, final report*. SDC document TM-2552, Santa Monica, CA (July 19, 1965).
27. Jules I. Schwartz and Clark Weissman, "The SDC time-sharing system revisited" in *Proceedings of the 1967 22nd National Conference*, pp. 263-271, Association of Computing Machinery (1967). DOI 10.1145/800196.805996.
28. Thomas Marill and Lawrence G. Roberts, "Toward a cooperative network of time-shared computers" in *Proceedings of the AFIPS Fall Joint Computer Conference*, pp. 425-431, ACM, New York, NY (November, 1966).
29. Larry G. Roberts, "The ARPANET and computer networks" in *A History of Personal Workstations*, ed. Adele Goldberg, Addison-Wesley, Reading, MA (1988). 145.
30. Norberg and O'Neill, p. 158. Hafner and Lyon's book has a similar account, though neither book cites a source. The BBN report, the ARPANET, *A History of the ARPANET: The First Decade*, DARPA, 1981, by F. Heart et al., says only that in 1965 Marill and CCA were

given the networking-study subcontract and that the CCA study reported in 1966 was done in late 1965. It adds, "Later in 1966, CCA received another contract to carry out the linking of the Q-32 and the TX-2,"

31. Thomas Marill, "A cooperative network of time-sharing computers," Technical Report No. 11, Computer Corporation of America, Cambridge, MA (June, 1966).
32. F. Heart, A. McKenzie, J. McQuillan, and D. Walden, *A History of the ARPANET: The first decade*, Defense Advanced Research Projects Agency, Arlington, VA (April, 1981). p. III-10.
33. David Walden and Raymond Nickerson, eds., *A Culture of Innovation*, Waterside Publishing, East Sandwich, MA (2011).
34. Taylor had tried to get Roberts to run the Arpanet project in early 1966, but Roberts didn't want to leave Lincoln Lab, where he was working on computer graphics. In his oral histories at both the Babbage Institute and the Computer History Museum, he describes how, in the fall of 1966, he put pressure on the head of the Lincoln Laboratory, which depended on ARPA funding, to persuade Roberts to take the job. Taylor said that he "blackmailed him into fame", a remark to which Roberts refers in "The Arpanet and computer networks", (Goldberg, *op. cit.*, p. 145).
35. Norberg and O'Neill, p. 33.
36. Lawrence G. Roberts, "Interview, 1988-06," James L. Pelkey Collection: History of Computer Communications, Lot X5671.2010, Accession 102746626, Computer History Museum. p. 3.
37. <http://archive.org/details/arpanet>.
38. <http://archive.org>.
39. <http://archive.org/details/MailToKatieHafner>.
40. Thomas Marill, "A nationwide cooperative computer network," Technical Report No. 11, draft 3, Computer Corporation of America, Cambridge, MA (28 February, 1966). MIT Lincoln Laboratory Archive.
41. Digital Equipment Corporation, *Programmed Data Processor-1*, Maynard, MA (1960).
42. *op. cit.*, p. 158.
43. Bernard Williams, "Coercing the network," *Dr. Dobbs Journal* **26**(10) (October 2001). <http://www.drdoobbs.com/coercing-the-network/199200749>.
44. Thomas Marill, Progress Report on Networking, Computer Corporation of America memo (14 February 1967). To J. L. Mitchell.
45. Lincoln Laboratory, *Quarterly Technical Summary, General Research*, Lexington, MA (February, 1966). . p. 6.
46. Lincoln Laboratory, *Quarterly Technical Summary, General Research*, Lexington, MA (May, 1966). pp. 5-6.
47. Lincoln Laboratory, *Quarterly Technical Summary, General Research*, Lexington, MA (November, 1966). p. 15.
48. Lincoln Laboratory, *Quarterly Technical Summary, General Research*, Lexington, MA (February, 1967). p. 8.
49. Lincoln Laboratory, *Quarterly Technical Summary, General Research*, Lexington, MA (May, 1967). p. 6.
50. D. W. Davies, K. A. Bartlett, R. A. Scantlebury, and P. T. Wilkinson, "A digital communication network for computers giving rapid response at remote terminals," *Proceedings of*

the first ACM symposium on Operating System Principles, ACM, New York (November, 1967). Roberts adopted the term "packet" from the presentation of this paper. Discussions with Scantlebury also persuaded Roberts to use 50 Kbps phone lines for the Arpanet instead of the much slower ones that he'd planned to use.

51. Lawrence G. Roberts, "Interview, 1988-06," James L. Pelkey Collection: History of Computer Communications, Lot X5671.2010, Accession 102746626, Computer History Museum. p. 7.
52. William Mann, Message format and protocol, Computer Corporation of America memo (23 May 1966). MIT Lincoln Laboratory Archive.
53. William Mann, email, 26 June 2014.
54. Hallum Murray, Q-32/TX-2 message format - revision II, Computer Corporation of America memo (27 July 1966). To L. G. Roberts.
55. Thomas Marill, Schedule for networking demonstrations - Phase 1A, Computer Corporation of America memo (13 June 1966). To Larry Roberts.
56. Robert M. Gray, "'A Survey of Linear Predictive Coding: Part I of Linear Predictive Coding and the Internet Protocol,'" *Foundations and Trends in Signal Processing* **3**(3), pp. 1-147, Now Publishers (2010). Reports that in 1967, Danny Cohen had connected a DEC 338 in Roberts' Pentagon office to the TX-2. It adds that the point-to-point telephone line was "not yet networking". Another view is in a message by Cohen (11 April 1999) in an Internet mailing list that mentions working on this connection and says, "It sure was the first ARPAnetworking" (<http://seclists.org/interesting-people/1999/Apr/52>). There is no report, however, on uses of the link.
57. A. N. Stowe, R. A. Wiesen, D. B. Yntema, and J. W. Forgie, "The Lincoln Reckoner: An operation-oriented on-line facility with distributed control," *Proc. of the Fall Joint Computer Conference* **29**, pp. 433-444 (1966).
58. Weissman said that the program was basically one that he wrote for his book, *Lisp Primer: A Self-tutor for Q-32 Lisp 1.5*, SDC, 1965, and that this arrangement worked well, providing a demonstration of the value of networking for program-sharing.
59. Thomas Marill, Progress Report on Networking, Computer Corporation of America memo (3 April 1967). To J. Mitchell.
60. Claude Baum, ed., *Research and Technology Division Report for 1966*, System Development Corporation, Santa Monica, CA (1967). Technical Memorandum TM-530/010/00.
61. Lawrence G. Roberts, "Multiple computer networks and intercomputer communication," *Proceedings of the first ACM symposium on Operating System Principles*, pp. 3.1-3.6, ACM, New York (November, 1967). DOI 10.1145/800001.811680. Written in June, 1967, according to Roberts' online version.
62. Hallam Murray, Functional description of the network system, Computer Corporation of America memo (30 March 1967). To J. Mitchell.
63. "Tentative specifications for a network of time-shared computers," Document No. 40.10.130, Advanced Research Projects Agency, Washington, D. C. (September 9 1966.).
64. W. Lichtenberger to D. Hemmendinger, email, 23 August 2013.
65. Ronald M. Rutledge, Albin L. Vareha, Lee C. Varian, Allan H. Weis, Salomon F. Seroussi, James W. Mayer, Joan F. Jaffe, and Mary Anne K. Angell, "An interactive network of time-sharing computers," *ACM '69: Proceedings of the 1969 24th National Conference*, Association of Computing Machinery, New York (1969).

66. S. Crocker, S. Carr, and V. Cerf, "New host-host protocol," Network Working Group Request for Comments:33 (12 February 1970).
<https://tools.ietf.org/html/rfc33>.
67. Computer History Museum, *Internet History, 1962-1992* (accessed January 2015.). Other sources include Pelkey's online book (*op. cit.*), Lambert, et al. *The Internet: A Historical Encyclopedia*, ABC-CLIO, 2005, and Larry Roberts' Internet chronology (<http://packet.cc/internet.html>), which says "*This experiment was the first time two computer talked to each other.*".
68. Robert W. Taylor, "Oral history of Robert W. Taylor," Oral Histories Online, Lot X5059.2009, accession 102702015, Computer History Museum (October, 2008). Conducted by Paul McJones.
69. Email, 26 February 2013.
70. Email, 30 October 2014.
71. Lincoln Laboratory, *Semiannual Technical Summary, Graphics*, Lexington, MA (May, 1967). p. 2.