Chapter 3
Advances in Communications

AT SRI:

1969
First Computer Network Connection
–With UCLA–

1976
First Internetwork Connection
–Two Dissimilar Networks–

1977
First Internetwork Connection
–Three Dissimilar Networks–

Communications is a broad and immensely critical part of our existence. In this chapter we will examine only the technical side of this huge preoccupation we have to talk to one another and to learn the state of our world. SRI has been engaged in some form of communications research almost from its beginning but here we will cover just three important aspects that have had significant impact on our lives and certainly upon mine, for this is my chosen area of work.

Computer Networking

Rossotti's had been a well-known Mid-Peninsula "watering hole" since the days when it was a stagecoach stop between San Francisco and Monterey. It was in the second bank of foothills west of San Francisco Bay, had a casual atmosphere and some outdoor seating—a good place for a special event. It would be all right to park SRI's "bread truck" van alongside and run a few wires to one of the tables in the courtyard. It was far enough from SRI to look remote but close enough to have good radio contact to SRI through the Stanford University field site repeater. That's why this venue was chosen to mark the occasion of the first internet transmission on August 27, 1976.[1]

The van was an SRI-outfitted mobile radio lab that contained the equipment to make it a portable node on the emerging Packet Radio Network (PRNET), sponsored at SRI by the DoD's Advanced Research Projects Agency (ARPA) starting about 1973. Placing a terminal on one of the wooden courtyard tables and connecting it to the van, a number of SRI people who had gathered for the celebration filed a weekly Packet Radio Program report, representing all the Program's contractors, to ARPA. Although the testing of such a connection had been going on for several months, this long email report was, in a ceremonial sense, the first internet transmission; that is, the first formal use of the internet protocol called TCP.[2] That protocol convention was designed to carry information over dissimilar networks, in this case the PRNET, through a gateway at SRI, across the wire-based ARPANET, to a set of hosts distributed around the United States. This small, understated but deliberate episode is indicative of but one of SRI's early contributions to the field of digital communications. But let's return even earlier to the beginning of computer networking itself.

[1] In the formative days of computer networking, terms like internet and intranet were in common use with their respective meanings. The lower case usage here purposely reflects that language, with the Internet arriving some undefined time later.

[2] TCP is the acronym for Transmission Control Protocol, network software that establishes, operates, and closes a reliable virtual circuit across dissimilar networks. Although it is still in use today, the overhead for this type of connection was deemed excessive for some types of traffic. This recognition soon led to development of a companion transaction protocol called the Internet Protocol (IP). Together they comprise the transport system of today's Internet.

The First Computer Network Connection

SRI's first computer network involvement arose when the notions of such networking were first being created. Accounts of that creation and the people who worked to develop computer networking at organizations like ARPA, the University of California, Los Angeles (UCLA) and Santa Barbara (UCSB), the University of Utah, Bolt, Beranek, and Newman (BBN), and SRI can be found in *Casting the Net*,[A] *Inventing the Internet*,[B] and other books on the history of the Internet. Although the early formulations of networking were carried out as part of other ARPA programs, in the spring of 1968 SRI was awarded the first contract under a new program area at ARPA specifically called ARPANET.[C] It was given to SRI engineer Elmer Shapiro to study the "design and specification of a computer network." The report of the 4-month study mentioned rudiments of concepts that today are part of the lore of packet networks.[D] Figure 3-1 from that report defines a few network nodal terms and shows the three types of links for which protocols had to be written: IMP-IMP, IMP-host, and host-host.[3] From the flexibility of using any workable set of switches in host-to-host communications, and its transparency to the end user (Shapiro called it a "transparent pipe"), came the term virtual circuits. Messages passing through the network had various compartments, some relevant to switches and some to hosts. This convention evolved into protocol layering. Protocols were discussed, but packets per se and the concept of a network virtual terminal, a specification allowing the many different types of terminals to meet just one interface standard, were adopted later. Hosts were assumed to be of the timeshare nature, permitting many simultaneous users not all of which were local. Because Shapiro was under contract to ARPA, ARPA program manager Larry Roberts asked him to get the relevant interface and packet transport software going. Shapiro had also suggested to Roberts that a more formal group be established to govern the design of the new network.[E]

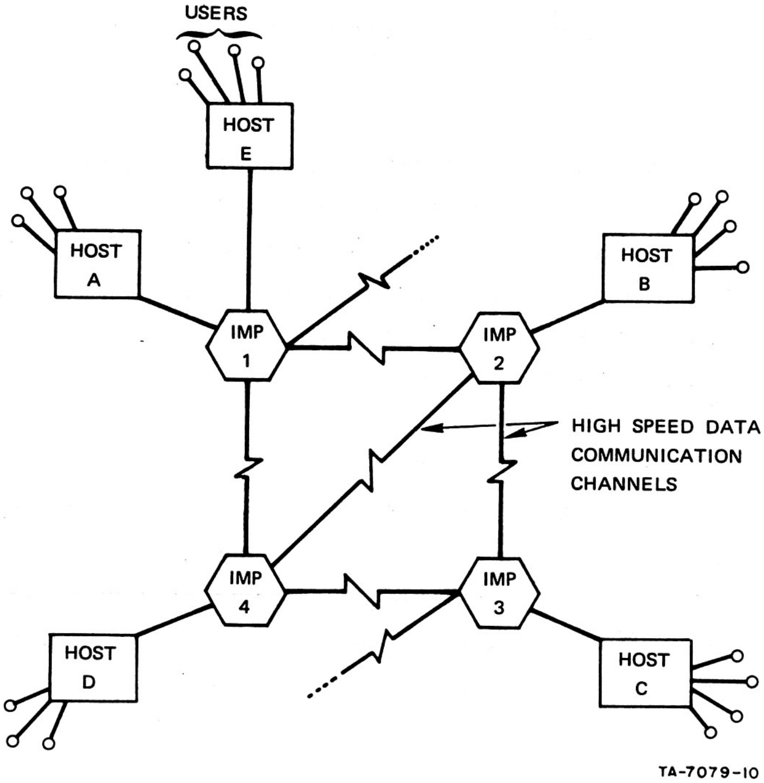

Figure 3-1. Network nodal terms as shown in Shapiro's early network study. (E.B. Shapiro, op. cit.)

Accordingly, in mid or late 1968 a meeting was called of people from what were to become the first four ARPANET sites and posed the problems to be addressed. Shapiro chaired that first meeting of what was to become the dominant voice in the technical design of the network, the Network Working Group (NWG). That first gathering also included Steve Crocker (UCLA), who would soon become the chair and formative leader of the NWG; Steve Carr (University of Utah); Ron Stoughton (University of California, Santa Barbara); and Jeff Rulifson (SRI). Crocker described the occasion as a "seminal" meeting.[F] But it was undoubtedly the fortuitous absence of established conventional

[3] The term IMP, used in Figure 3-1, stands for Interface Message Processor. It was the host computer's access point and the ARPANET's basic packet switch. Although initially each site had to interface its own host to the IMP, the first standard interface software was soon developed by Bolt, Beranek, and Newman (BBN) as defined in their so-called 1822 Interface. The building of the IMPs was awarded to BBN in December 1968.

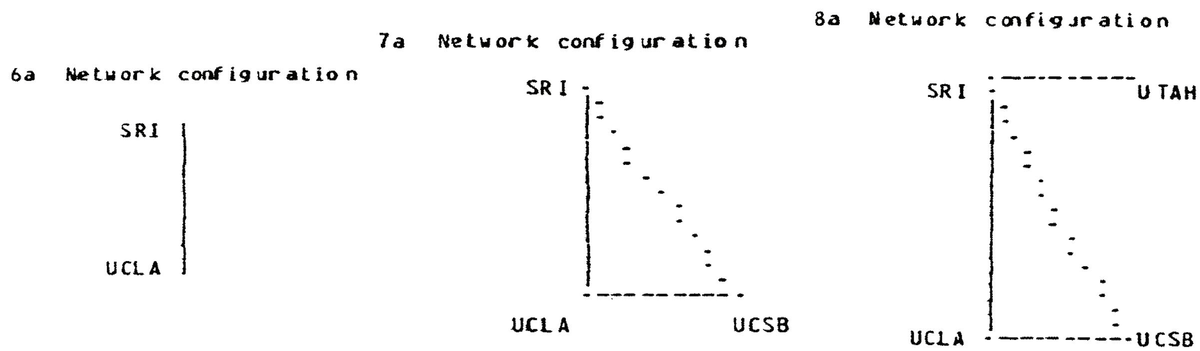

Figure 3-2. The first three ARPANET configurations—actual text images from RFC 4.

network professionals that allowed the tabling of fresh new ideas. This now respected meeting set the stage for the high level at which this totally new networking world would be pursued.

The medium of interchange between these pioneers and the place where concepts were hammered out came to be called Requests for Comments (RFCs), a name and process still in use in the Internet with nearly 4,000 RFCs having been issued. That unassuming name, suggested by Shapiro to Crocker,[G] seemed to look for and await more professional guidance. As it turned out, the professionals these young pioneers awaited didn't exist, and RFCs became the process for setting network standards and remain so to this day.

It is worth noting that RFCs 2, 4, and 5 were written by SRI people: two from Douglas Engelbart's lab and one, by Shapiro, from the Information Systems Lab. Bill Duvall outlined, in RFC 2, the initial procedure for host-host interaction. These first awkward network interactions were patterned after a regular terminal-to-host dialogue and became the basis for the first connection between two networked computers. Shapiro wrote RFC 4 in March 1969, an implementation schedule appropriate to his ARPA contract and his leadership role in the NWG. This RFC showed the first three configurations of this evolving network, which are reproduced in Figure 3-2 just as they appeared there. Rulifson of SRI wrote RFC 5, the definition of a language called DEL that was to be the first, albeit rudimentary, realization of a host-independent computer language.

On October 1, 1969, SRI received the second ARPANET packet switch or IMP. UCLA, in its capacity as Network Measurement Center, had received the first IMP a month earlier. Its arrival offered the important opportunity to create, with UCLA, the first linkage on a general-purpose digital network, which occurred in October 1969.[4] These first sites had to build their own host-to-IMP interface, a major effort that took several weeks. The next task was to develop the host-host connection software. For background there was Duvall's link-level terminal-host protocol outlined the previous April in RFC 2. The first such program was a simple procedure, limited to one character at a time, and as a result there was a brief "hiccup" in that very first interaction. It was late at night when the SDS Sigma 7 at UCLA and the SDS 940 at SRI could be unloaded. The natural first step was to have a terminal on one host log on to the other, connected but distant host. To complicate things a bit, the machines used two different operating systems, SEX at UCLA and GENIE at SRI. Therefore each machine had to convert each character received to its local equivalent, meaning the programmers watching had to convert the hex to an equivalent alpha letter. After correctly receiving an *L*, an *O*, and a *G*, the SRI machine, as was its custom for its attached terminals, automatically returned a *GIN*, thus violating the one-character-at-a-time rule of that first interplay. This caused the UCLA machine to respond with a *?*, to which the SRI host automatically also returned a *?*. This loop continued until one machine crashed or was shut down. Duvall couldn't remember which.[H] The multi-character correction was made and the network was off and running.

An SRI Quarterly Progress Report to ARPA for the period indicates the matter-of-fact meaning of the new connectivity:

[4] The exact date of the first transmission is unknown, but it was certainly in October 1969. Charles Kline, then at UCLA, finds in his notes that it occurred before 29 October (see *San Jose Mercury News*, front page article by David Plotnikoff, September 9, 1999).

"The Interface Message Processor (IMP) and communication terminal for the ARPA Computer Network have been installed and the hardware interface between the IMP and the 940 is complete and in operation.

Under control of EOM instruction from the 940, the interface will transmit and receive Network messages directly to and from buffer in 940 core. Interrupts signal the 940 when transmission is complete. The interface was designed to operate with either 940 core or external core, and.....

Test programs have been developed to check out communication with the IMP and with other Network host computers by "looping" messages. A preliminary operating system that allows one remote user to log in to our system over the Network and simultaneously allows us to log in to any other available system on the Network has been written and is in operation."[1]

In truth this interim report reads like thousands of others at SRI, where doing things for the first time is simply taken for granted. But to their credit, these people were seeking the fundamentals of networking, concepts and practices that would provide not just an ad hoc solution to two computers talking, but a foundation for the best evolution of this new technology. That the work of these pioneers, represented in general by the NWG, should seem compelling is mostly the consequence of some talented young researchers being given some new resources, some workable deadlines, and, most important, a clean slate.

The SRI people became involved, I believe, because the program managers at ARPA saw that the visionary computer-based, distributed collaboration work being developed in SRI's Augmentation Research Center (ARC) anticipated rich computer networking (see the following box and Chapter 2). Perhaps the most important reason that the second IMP came to SRI was Engelbart's willingness to host the Network Information Center (NIC). Participation in this effort was certainly consistent with his notion that the ARC was an information-creating and -wielding resource and… there were apparently no other volunteers. Thus the NIC became part of the ARC and, through the initiative of the first NIC director, Dick Watson, helped motivate and, in some cases, contribute technically to the evolving network protocols and higher level services. The NIC also became responsible for distributing, both online and physically, the technical designs and procedures for joining this first-ever packet network. It also began to assign and distribute network host names and addresses. These service roles came together in the NIC's practice of writing and distributing three formal documents:

- The ARPANET Directory, which listed the network address and telephone number of all registered users
- The ARPANET Resource Handbook, which was a compendium of the hardware and software of each host and the research interests and procedures and policies of its owner
- The DoD Protocol Handbook, which contained all the standard network protocols as well as relevant reference material.

The NIC performed these roles for about 22 years, from 1970 until October 1992, long after the creation of the Internet. Elizabeth "Jake" Feinler led the NIC over most of that time. The NIC also borrowed an idea from the Massachusetts Institute of Technology (MIT) for an informal geographical locator program and built the first online directory service, WHOIS. This innovation, developed by Ken Harrenstien and others, enabled visitors to obtain from the NIC server the network-related addresses and particulars of any registered host user. In its first year, in the mid 1970s, WHOIS received over

In the early 1970s, an SRI editor, Shirley Hentzell, worked online with Douglas Engelbart on a proposal that was also being composed jointly and simultaneously at two separate locations; that is, the same text was viewed at the same time by those at SRI and by a Purdue University professor in Indiana. The terminal showed the text in one window, with suggested insertions being made as we watched, and the other windows showed comments from Engelbart at SRI and the professor at Purdue. It was not only magic, it was a type of collaborative interaction only now coming into commercial development thirty years later.

28,000 visits from a still embryonic network.

But these anecdotes are getting ahead of the story of the network's genesis. Shapiro recalled a very interesting aspect of its early use.[J] The first few sites had been connected together. Utah had been chosen so as to develop, with the California sites, an interstate configuration that would ensure AT&T's participation in the network. But after some initial experimentation, the network wasn't being used very much, to the disappointment of Larry Roberts, the ARPA manager leading the ARPANET implementation. The network functionality was confined mostly to file transfer (FTPs) and some terminal access (TELNET) traffic, but even their collective affiliations through ARPA weren't enough to cause much traffic between the contractor sites. Certainly, computer resource sharing, the initial prime motivation for the network, was not materializing as an important need. Roberts, in a very pragmatic ploy that looks brilliant in retrospect, suggested that traffic between the contractor sites might very well be a condition for continued contract funding. Moreover, the amount of traffic might influence the amount of future work. Clearly, Roberts's pressing people to explore the utility of the new network not only now seems justified, but is also evidence that can only be called visionary.

By most appraisals, the first serious and widespread demand for networking came with the almost cavalier introduction of cross-network electronic mail by Ray Tomlinson of BBN in 1971.[5] It should be mentioned, however, that electronic messaging and file exchange, with symbolic addressing, was a fact of life within the SRI ARC as early as the fall of 1970. It was a feature of the NLS Journal system and when network addressing came along, the expansion to other machines was as natural as specifying a person and the machine on which their mailbox program resided.[6]

As for the continuing design of the ARPANET, SRI's role receded as the 1970s wound down. Shapiro, who had been recommended to ARPA by Engelbart, participated only temporarily. He had helped ARPA (Roberts) write its first bid invitation for network hardware (the IMP), helped evaluate the proposals, helped set up the NWG, and then turned to other pursuits. Members of the ARC such as Watson continued to help demonstrate the efficacy of packet switching and the new network, often through the use of applications.[K] But the ARC would elect to play a relative minor role in developing network technology except for that centered in the operation of the NIC. The NIC had an important role in the issuance of network host names and addresses and as the repository of the RFCs and other standards for network connection, access, and interchange. With this important exception, SRI largely withdrew from the later technical evolution of the ARPANET itself.

The First Internet Transmissions

The second notable area of computer networking development at SRI arose within a couple of other Engineering Research Group laboratories during the 1970s. Several visionaries at DARPA,[7] most notably Roberts and Bob Kahn, had seen the military need for a mobile, wireless version of the embryonic ARPANET. SRI and DARPA had discussed the possibility of a transportable, possibly handheld terminal or switching node for such a network rather than the massive, steel-encased IMPs of the early fixed network. Some researchers in SRI's Information Sciences Division attempted to develop a conceptual design for such a device, but they weren't able to convince DARPA that they were up to the task, in part because they had little radio experience. In response, SRI moved the project to the Telecommunications Sciences Center, which had splendid radio capability. By that time DARPA had formed a team of contractors in what came to be called the Packet Radio Program, which was created to develop a wireless adjunct to the evolving ARPANET. Its resulting network would be called the Packet Radio Net (PRNET). Members of that DARPA

[5] Like so many internet innovations, electronic mail has many fathers. It was present on some time-sharing systems by at least 1970, but Tomlinson's program, including its "@"convention, set the stage for network-wide rather than host-wide email.

[6] As a component of the NLS Journal, a generalized email was a natural tool of the community of users in SRI's ARC as early as August 1970. When the ARPANET attachment came, extending it outside the ARC was a simple modification carried out in 1970 (Doug Engelbart, personal communication to the Stanford University History Department, available at www-sul.stanford.edu/depts/hasrg/histsci/ssvoral/engelbart/main4-ntb.html). A rudimentary form of mailboxes was implemented on Multics by Tom Van Vleck in 1969 (see www.multicians.org).

[7] With the addition of "Defense" to its name ARPA became DARPA in 1972.

program were BBN in Boston, Collins Radio in Dallas, Network Analysis on Long Island, UCLA, and SRI. Because of a good understanding of radio and some background in high-frequency (HF) networking, SRI was chosen as system engineer and technical director (SETD) and integrator for DARPA's packet radio effort, a position it maintained for over a decade. (A comprehensive description of packet radio technology can be found elsewhere.[L, M])

The introduction of a radio segment to supplement the ARPANET was a natural outgrowth of the military context in which a great deal of U.S. research is done. Ultimately, military use of this new, interactive digital technology would depend on adaptation to the reality that the military is inherently mobile and may be deployed to any point on earth. Thus a radio network, particularly one that served a mobile population, was needed, and it turned out to be intrinsically different from the existing fixed, wired one. This clear difference, along with the need for the two networks to work well in tandem, led to the notion of a communication software structure that would effectively bind these disparate networks together as though they were one.

One technical insight needs to be inserted here to help understand how disparate packet networks can still function together. In most communications networks only the source and destination terminals are visible to network users, and all resources that lie in between are normally of little interest to users as long as they fulfill their role. In circuit switching, common to the telephone networks of the day, once a physical pathway was chosen, the same route is maintained for the whole session. When such circuits are leased, the connections may even be "hardwired;" that is, never changed. In packet switching, where subunits of a message may travel entirely different routes from source to destination, the exact role of intervening resources would not even normally be known. Thus, there arose the concept of a virtual circuit, where the only defining network nodes lie at the ends and the intervening ones are neither specified nor known by either the users or the network providers. This switching concept had been part of the basic ARPANET design, and it was to be extended to this amalgam of wire and radio networks, and thus to the world of internets.

The clear differences between the wire-based ARPANET, the mobile radio-based PRNET, and eventually satellite networks led Kahn, then leading the networking efforts at DARPA, and Vint Cerf of Stanford to design the first end-to-end protocol that could span dissimilar packet networks. The essence of such a construct began to emerge as Kahn attended a seminar held by Cerf at Stanford in the summer of 1973 and mentioned the existence of these dissimilar networks.[N, O] After some airing in the internet working group and elsewhere, the first encapsulation of such a protocol came together for them over a weekend in October 1973 at Palo Alto's Rickey's Hotel. They continued to rework it through 1973 and published the protocol in May 1974.[P] They named it the Transmission Control Program, or TCP, even though today the "P" is always referred to as protocol. With modifications, some of them very important, TCP is still in use today as the basis for packet transport in the worldwide Internet.

Following the introduction of TCP, DARPA let contracts for three separate implementations: one to Stanford University, one to BBN, and one to University College, London. The first, clearly "buggy" specification appeared in December 1974, when Stanford produced RFC 675. BBN had an in-house version working reliably about a year later and began exchanging TCP traffic with Stanford on an intranet basis. Jim Mathis, a student of Cerf's at Stanford, started to implement the Stanford group's protocol in 1975. In summer 1976, Mathis joined SRI, where he completed a version that would run specifically on the much more modest hosts of the PRNET (Digital's LSI-11 microcomputer). In the meantime Cerf, now a program manager at DARPA, was trying his best to inculcate the DoD with the virtues of packet switching and TCP for its future data networks.

Some of Cerf's effervescence led to an independent effort to implement TCP in a new DoD network called AUTODIN II, which was intended to be an upgraded version of an ancient paper tape system that had none of the desirable attributes of the ARPANET. The Defense Communication Agency (DCA) contracted with SRI in 1975 for yet another specification of TCP for AUTODIN II. Jon Postel,[8] Larry Garlick, and Raphael Rom

[8] Jon Postel was a member of SRI's ARC until spring 1977. Later, at USC's Information Sciences Institute, he became one of the most important and influential designers of the Internet.

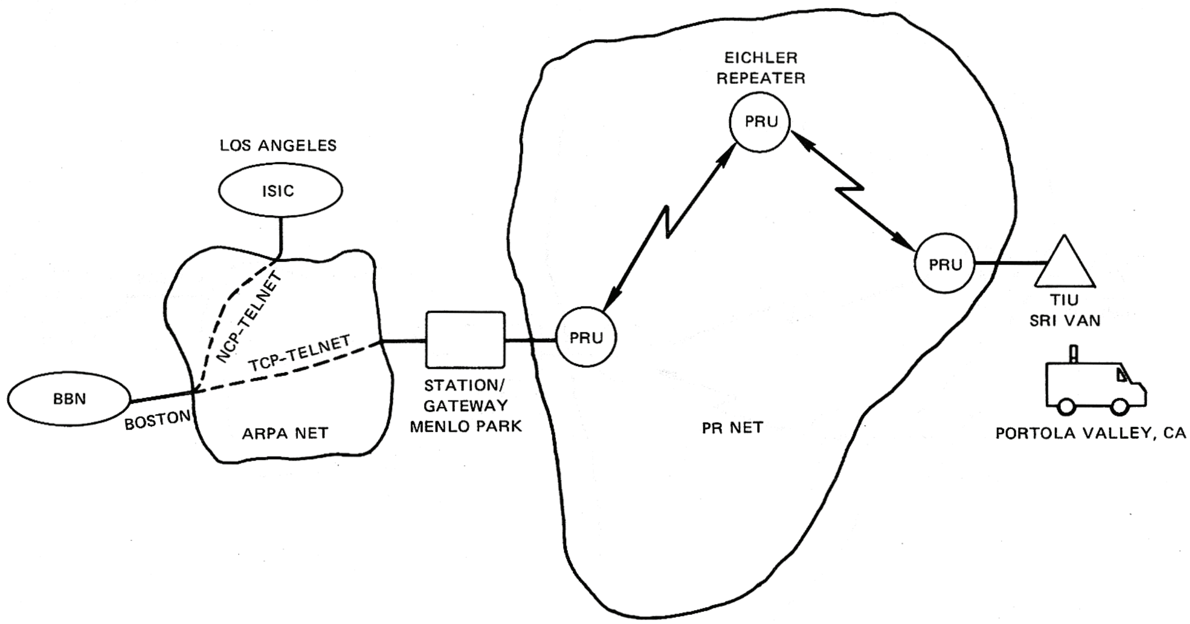

Figure 3-3. The network map and site of the first two-network internet transmission on August 27, 1976.

submitted that design in July 1976.[Q] During this work, some of the limitations of TCP became apparent, and the critical need for a less reliable or datagram protocol was conveyed to DCA.[9] For a variety of reasons, including the accommodation of packet speech, this concern was later addressed by the internet community, and the datagram, or transactional Internet protocol (IP), was separated out from TCP in 1978.[10] The standard to which TCP and IP were to subscribe was issued by Postel in RFCs 791–793 in September 1981.

With this background, we can revisit SRI's role in the first truly *internet* transmission. As mentioned at the outset, the first testing of TCP across dissimilar networks started in late summer 1976. The first trials stayed one radio hop from the Packet Radio station (the PRNET's controlling node) where the bidirectional ARPANET gateway software, built by Ginny Strazisar at BBN, was located. During July and August the SRI team tested and tuned Mathis's version of TCP for better accuracy and speed. In August 1976 a terminal attached to an LSI-11 "host" running TCP, which was in turn attached to the PRNET, which was in turn connected to an ARPANET gateway, first accessed an ARPANET host. This trial was the first time, at least in any formal sense, that dissimilar networks were bridged by TCP and was thus clearly a two-network internet connection. That specific network configuration is shown in Figure 3-3, taken (as Figure 1) from a packet radio progress report written at that time.[R] Figure 3-4 shows the aforementioned demonstration using TCP to convey the lengthy weekly Packet Radio Program report. Other TCP connections would soon follow.[11]

To fulfill the assumed need for a network of global reach, DARPA next moved to include a

[9] Raphael Rom, personal communication, February 10, 2000. Also, a pressing need for a less than infallible protocol was to convey speech over such packet systems. In other ways, too, the need for a low-delay, best-efforts protocol became clear.

[10] According to Abbate (Janet Abbate, *Inventing the Internet*, MIT Press, 1999, p. 130), Cerf, Cohen, and Postel decided to create IP during a meeting in January 1978. Cerf (personal communication, January 15, 2002) attributes most of the IP initiative to Cohen and says that packet speech was an important justification. IP was to be a best-efforts, no end-to-end retransmission protocol, and therefore to have smaller delay variance.

[11] An expected part of the DARPA work was to demonstrate progress and give evidence of this new networking capability. Therefore, TCP, spanning the PRNET and the ARPANET, would be demonstrated in May 1977 between the SRI van and hosts at ISIC and the SRI-KL host. Also, on August 11, 1977, a TELNET connection would be established between the van and the Naval Ocean Systems Center in San Diego for Admiral Stansfield Turner (Director, Central Intelligence Agency) and William Perry (DDR&E). On September 19, 1977, a single LSI-11 microcomputer, containing TCP, connected four independent terminals through a packet radio to four different ARPANET hosts, essentially all of the ones that were running TCP servers at the time.

Figure 3-4. (left) Participants in the First Two-Network Internet Transmission (From the left: Don Cone, Army observer Mike Berishinski, Nicki Geannacopulos, Dave Retz, Ron Kunzelman, Jim McClurg, and Jim Mathis.) (right) Nicki Geannacopulos entering the packet radio weekly report. (Don Nielson took the pictures.)

third packet network, one that was satellite-based. Within a year, DARPA was ready to test all three networks together. On November 22, 1977, what has also come to be regarded by many as the first internet transmission was sent from the SRI mobile packet radio van to a host computer at the University of Southern California (USC) by way of London! The route is shown in Figure 3-5.[S] The SRI van, used as a mobile test node throughout the Packet Radio Program and now a historic vehicle, is shown in Figure 3-6.[T]

Two interesting demonstrations were used repeatedly at the time to illustrate the robustness of this new concept of networking. To illustrate the flow of traffic between the mobile van and some distant network host, a character generator would grind out continuous alphanumeric sequences that formed patterns on a CRT in which errors would be obvious. While this source was moving at high speed in the SRI van, the signal would sometimes be interrupted owing to shielding of the radio signal (for example, when the van passed beneath an underpass). The flow would stop momentarily and then resume with no errors being observed. Error-detection using cyclic redundancy checks, applied at the end of each transmitted packet, was used to verify reception accuracy. These checks, along with the end-to-end ordering and retransmission properties of TCP, would not permit delivery of altered packets even though packets were frequently lost. A similar procedure was to withdraw from the packet radio its critical synthesizer card. This would terminate the character flow, but reinserting the card would restart the flow. Thus, for a variety of reasons, the traffic would stop or could be interrupted, but no errors were ever observed. Those demonstrations were splendid evidence that each packet could have sanctity, even in a tough environment of intermittent propagation and noise. This exciting capability was certainly foreign to those circuit-oriented engineers who saw reliable mobile radio data systems as some sort of oxymoron.

SRI's final contribution to the development of TCP came as the protocol was being incorporated into Berkeley UNIX, the host operating system that was dominating the research-oriented Internet in the early 1980s. Bill Croft, a member of the Telecommunication Sciences Center at SRI, faced the problem of installing networking software on a number of PDP-11 minicomputers at SRI and elsewhere. In porting the Berkeley UNIX TCP/IP software from the VAX to the PDP-11, he found and repaired "dozens of bugs" in the original UNIX code. Croft's corrected version of TCP/IP, now an integral part of UNIX, became an important vehicle in the rapid spread of networking.[U]

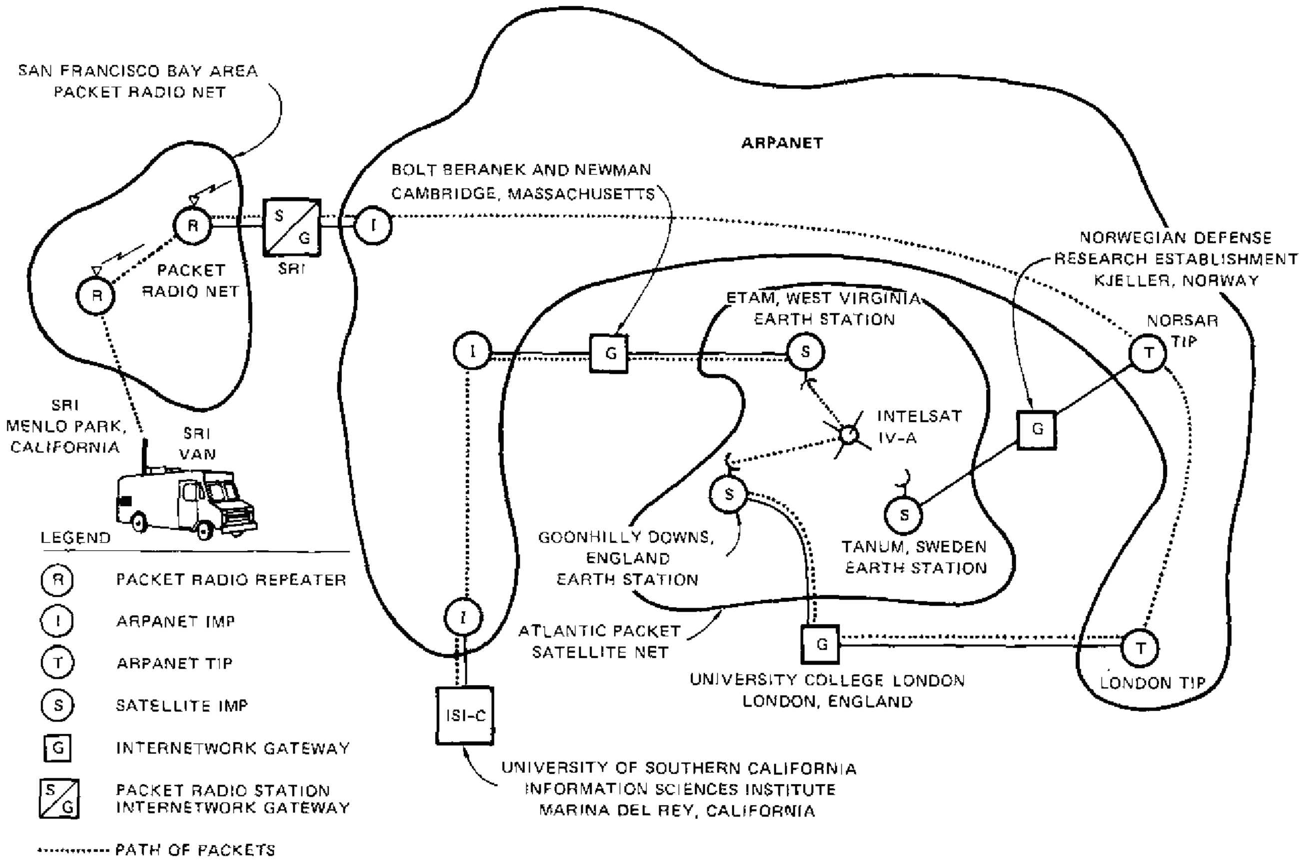

Figure 3-5. Route of the first three-network internet transmission on 22 November 1977.

Thus, the beginnings of internetworking, driven by the needs of the U.S. military for both mobile and longer range digital radio systems, were the result of broad collaboration both within and outside SRI.[12]

The Introduction of Packet Speech

When the ARPANET was perhaps 5 years old and before the development of internet protocols, Kahn at DARPA requested that a group of contractors explore how the new network could handle normal telephone traffic. Given the initial focus on reliable data transmission, it was not clear whether the variability in interpacket delay would permit

[12] Other than project leaders Donald Nielson and Ronald Kunzelman, important SRI contributors to these early internet efforts were Jim Mathis, Stan Fralick, John Leung, Don Cone, Dave Retz, Jim McClurg, Russ Wolfram, Mike Placko, Keith Klemba, Janet Tornow, Don Alves, Bud Sargent, and Vince Sherville. With the advent of other PRNET and internet design issues and a range of internet services such as speech, these SRI staff members were joined by Earl Craighill, Nachum Shacham, Jan Edl, Jose Garcia-Luna, David Beyer, Barbara Denny, Diane Lee, Mark Lewis, Raphael Rom, Andy Poggio, Bill Zaumen, Andy Poggio, John Hight, Richard Ogier, Ed Kozel, and Zaw-Sing Su.

Figure 3-6. The SRI packet radio van. Inset shows two packet radios, an LSI-11 microcomputer containing TCP plus terminal interfaces, and a DataMedia terminal.

the smooth flow required by a voice call. In 1974 because of the narrow bandwidth of the initial circuits comprising the net, Kahn initiated the Network Speech Compression Program. This program resulted in the choice of some speech compression algorithms, and these were first tried over the ARPANET. In 1976 SRI's Earl Craighill and Tom Magill, both of whom had been working on the speech program, convinced DARPA to let them try speech on the Bay Area PRNET. By this time the internet protocol, TCP, was also being tested, and thus speech experiments also began on an internet basis.

Because the SRI van was an easily outfitted facility in which packet radio and internet equipment had already been installed, it became the first mobile node for packet speech experiments (see Figures 3-7 and 3-8). In addition to the challenges of mobile data transport, transporting natural-sounding speech focused on the importance of delay variance. Innovations were needed in variable rate encoding, new buffering strategies, and rapid rerouting of packets whenever the route-in-use failed. All these techniques were needed to help smooth the flow of speech. Thus, internet speech connections were being made as early as 1977–1978, about the same time as the Internet itself was becoming a reality. It was mentioned earlier that the handling of packet speech was one reason for wanting a less-than-reliable or transaction protocol. In packet-based speech a wide variance in the delay of individual packets, stemming from repeated end-to-end packet retransmissions, is worse than the occasional loss of one. Thus, packet speech, anticipated to be a common future internet service, became a motivating factor in the creation of IP.

Taking Packet Radio to the Field

In the 1980s SRI continued its leadership in the development of packet radio systems. Through a variety of new network terminal devices and a larger supply of packet radios, SRI placed packet radios aboard ships, airplanes, and helicopters to demonstrate that this technology could fit into the operational environment of all of the military services. Several concurrent and concatenated programs at DARPA continued such development as much larger, more dynamic, and more survivable military networks were envisaged. In 1983 DARPA began a mobile digital network effort called the Survivable Radio Network Program or SURAN. SRI continued as the integrator and technical director in this new work, which again involved many contractors. Over a period of 7–8 years development continued on robust algorithms and protocols for routing amid node failures, more secure network access and management procedures, improved adaptation in larger and larger networks, and finally methods of implanting those changes on available hardware. The packet radio equipment became smaller and more portable; such devices included lower cost packet radios from Hazeltine Corp. and new, smaller terminal-network interface units designed and built at SRI. In addition to conducting large and expensive field tests, SRI also operated a large laboratory testbed to accomplish the same kind

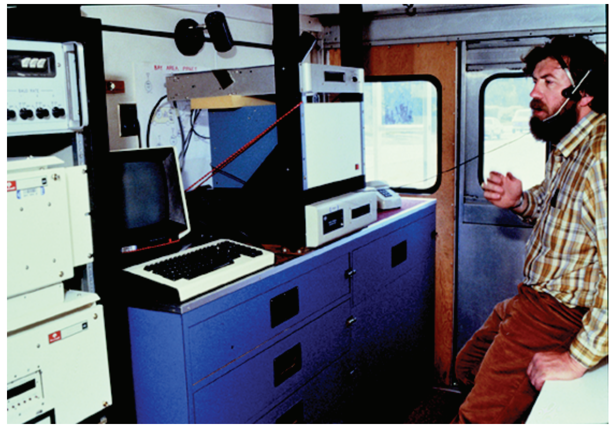

Figure 3-7. SRI's speech packet project leader Earl Craighill in the SRI Van (which housed the speech encoding and packet-forming equipment).

Figure 3-8. SRI's Jan Edl demonstrating speech transmission over the internet. (The Mickey Mouse phone was deliberately used to illustrate that the speech equipment hardware and software were designed to accommodate a "standard," off-the-shelf telephone.)

of evaluation.

Another related area of work unfolding in the 1980s consisted of applying these new digital network-based systems to the information needs of military operations. SRI's role in networking also came to include developing and demonstrating how to use such a network. Demonstrations occurred as early as 1980 at Ft. Sill, Oklahoma, where a packet radio network was used in Army fire control exercises. The topological clarity of the packet radio connectivity made the PRNET control station the focal point for observing much of the entire operational exercise.

But the largest and most influential use of both packet radio and the ARPANET in military operations began at Ft. Bragg, North Carolina, in 1979. With mainly DARPA funding and under the leadership of SRI's Mike Frankel, SRI undertook activities there that showed, for the first time, the power of interactive computing and digital messaging systems in military exercises. In a team effort, SRI and military personnel explored, on a limited scale, the utility of these computer-based networks in relation to existing military communications and information processing systems. Because of a lack of encryption, packet radio, like its DARPA hosts, was limited to unclassified roles. A collaborative effort between the Army and DARPA to test interactive computing and digital messaging went on for over 7 years and clearly affected how the Army approached its information needs. The first phase of the program was called Army Data Distribution System (ADDS).[13] A continuing exploration of how the online digital network could offer new and better capability was called ADDCOMPE. Its specific objective was to explore new concepts in command and control. Could emerging digital systems actually increase the availability of critical information even as a battlefield became more decentralized? More flexible and robust packet-switched communications, enabling new distributed information processing technologies, in fact, did just that and SRI was leading the military through that change.

In a concurrent program, appropriately called TACTICS, SRI also introduced a set of military-required application programs that expressed these new ways to view and move information. Application programs such as SITMAP, an online, dynamic depiction of the military situation in map form using standard symbols, showed the Army how to keep much more apprised of a situation. New messaging and a new Tactical Reporting System were also developed and used in normal operational exercises. All these new capabilities made the power of distributed information technology evident and started the Army down a revolutionary path in its information systems that continues to this day. Prompted by power and miniaturization of this technology, now so well illustrated by the Internet, all military services are rapidly moving toward modern communications and information systems that can support the increased mobility and effectiveness demanded of a worldwide deployable force (see box on Accommodation.)

Another good example of how the new technology changed operational capability came from SRI's presence among the airborne units at Ft. Bragg. Since their mission was to be ready to fly anywhere on short notice, the ability to rapidly identify the correct materiel to take with them was crucial. SRI staff members working on site at Ft. Bragg noticed the time-consuming and precise practice of loading military transport aircraft. Realizing that new software methods might speed up this task, SRI suggested a method that led to a totally new airload planning system. Using SRI expert-system technology and working closely with highly experienced Army and Air Force loadmasters, the SRI team wrote a new program. SRI designed and built a new software expert system that was orders of magnitude faster than existing methods and accurate enough for widespread military use, for example, in the 1991 Gulf War.[14]

The last of the packet radio initiatives at DARPA was a 5-year program that sought to increase the available bandwidth of the network to accommodate the digital distribution of essential military information such as imagery. Begun in 1995, the Global Mobile Information

[13] ADDS had unclassified roles in many Army and joint field exercises in the United States and overseas. Domestic examples were Solid Shield '81, Gallant Knight '82-'83, Gallant Eagle '82, and Brim Frost '83. Overseas examples were demonstrations with the British at the SHAPE Technical Center in the Netherlands linked to England and at U.S. headquarters in Heidelberg. Both of these used packet radio networks that linked to existing European military networks in ways that showed the true flexibility of interneting (Douglas Hagan, SRI International, personal email communication, May 26, 2002).

[14] The program was called the Automated Air-Load Planning System (AALPS), also mentioned in Chapter 4 on artificial intelligence. Its project leader was Debra Anderson.

TECHNICAL ACCOMMODATION BY THE MILITARY

As one might expect, the introduction of new information technology to the military is no easy task when time-tested methods have become standard operating procedures. However, with skepticism came some willingness to try as long as…. The two images below depict first reluctance and then accommodation. The picture on the left is a veteran Army colonel gingerly exploring the utility of some of the earliest programs at Ft. Bragg in 1981. On the right is a marine corporal with a completely mobile, handheld computer that enables a continuous account of his situation and those around him. SRI also designed and built the latter in its ongoing exploration of the equipment and software that increases the awareness and capability of our soldiers or any collaborating group. That current system is called INCON, a wireless automatic information distribution system now undergoing field-testing. The accommodation period between these two images is about 20 years.

Systems Program, or GLOMO, delved broadly into the functionality needed in wideband distributed systems. On the radio side this effort included new antennas and CMOS digital radios. On the software side, it required new distributed file systems. The most visible change was the replacement of special-purpose, quasi-milspec hardware by equipment from the rapidly evolving commercial sector. In the area of terminal equipment, commercial products included conventional laptop computers, then miniature laptops, and now wireless personal digital assistants (PDAs). SRI was DARPA's lead technical director and integrator of the demonstrated technology for GLOMO, playing the same role it had played in similar communications programs for over 30 years.

Attempts at Commercialization

In the 1980s, when packet radio was a working, demonstrable system, SRI called on a few commercial organizations to test their interest in its technology. One application, easy to envision, was the use of packet radio technology as a telephone and data communications overlay on the deteriorated or non-existent telephone systems in underdeveloped countries. Another was its use in oilfields or other places where distributed monitoring and control were critical. Control of a fleet of mobile trucks was another. Although SRI suggested all these applications to companies such as Motorola, none aroused interest. Today, of course, wireless is an exploding technology that enables, for example, a customer to track a FedEx package from pickup to delivery.

The first company to venture into the world of commercial, area-coverage, wireless networking was Metricom. With money from Paul Allen of Microsoft and leadership from Paul Baran, one of the originators of packet transport, the company was formed in 1985 and grew during the early period of the dot-com surge. But its offering began with a paltry 28 kilobits per second (kbps) bandwidth (not upgraded until 1999) and a high $75 per month fee. Also, because Metricom targeted consumers over vertically integrated businesses, customer growth was slow. More important, by ultimately extending its scope of operations to

17 cities nationwide and incurring the associated high capitalization costs, Metricom eventually outran its inventory of paying customers, slightly more that 40,000 at the beginning of 2001. With a billion dollars in debt, it filed for bankruptcy in July 2001 and offered its Ricochet Network assets for sale a month or so later. Other than having created similar technology, SRI had no role in Metricom.

In 1995 a number of former SRI packet radio engineers, led by David Beyer, banded together to form another packet radio technology company, Rooftop Communications.[15] Its aim was to provide residences and businesses with wireless, broadband access to the Internet. Like packet radio, each node was both an access point and a relay station that helped form a more accessible network; that is, every node did not have to have line-of-sight connectivity to a high, prominent, centralizing node. Unlike Metricom, which had a half-dozen or so company-owned nodes per square mile, Rooftop nodes exist only at each subscriber's site and act as both network routers and access points. By 1998 Rooftop had developed the first commercial self-configuring, multi-hop wireless IP routers.[16] Self-configuring in this case means that each router (switch) determines which other routers to use for a given packet address. Rooftop's routers are also able to change their routing neighbors dynamically as environmental conditions change. According to Dave Beyer, the Rooftop protocols were designed from scratch but arose out of the "deep experience" he and other company staff members had gathered at SRI. Protocol issues like the hidden terminals, the density of nodes required for full connectivity, the ability to scale the network to a much larger size, and supporting real-time traffic all drew on SRI experience. Software development methods built and used at Rooftop also drew on SRI experience, as did the staff's ability to gain an overall system architecture view similar to that required for the role as system engineer for the DARPA packet radio project.

Recognizing the value of this informed approach to wireless system design, Nokia Telecommunications purchased Rooftop in September 1999 for $57 million in cash and stock. SRI as an organization had no stake in Rooftop.

Recently, however, SRI has built a software system called "PacketHop" that embodies much of what packet radio offered 20 years earlier. PacketHop:

- Builds a radio-based infrastructure that is self-forming, self-organizing, and self-healing
- Offers true peer-to-peer communications
- Has multi-hop, multiple-path algorithms that reroute around problem areas
- Applies to a wide range of transmission media
- Is scalable to thousands of nodes
- Offers guaranteed packet delivery through packet management functions.

As of this writing the technology has been licensed to three concerns. Speedcom Wireless of Sarasota, Florida has a worldwide non-exclusive agreement that extends for 6 years and gives SRI some equity in the enterprise and a seat on an advisory panel. Associated with this licensing is a reminder of the lag time of new technology. In about 1974, before SRI received its first Collins-built packet radios in about mid-1975, SRI handcrafted a set of radios itself and took them to Hawaii to run over the same links as the radio-based University-based Aloha Net. Recently, an article in the Hawaii Business Magazine[V] indicated that a Hawaiian company called Landmark Networks (now called FireTide) is the second licensee of SRI's "PacketHop mesh technology." The company will work with the University of Hawaii to create in the Honolulu area what is perhaps the most flexible wireless network yet. Almost 30 years has elapsed between initial use and commercial application. The third licensee is SRI's own spin-off company Packet-Hop.

Yet another company emerged in 2000 using packet radio technology. It is called SkyPilot and two of its earliest members were a former SRI researcher and manager, Mark Rich, and researcher Bernadine Yetso. It also uses the mesh network technology inherent in packet

[15] Besides CEO David Beyer, other members of the 20-person firm who once worked at SRI were Thane Frivold, Darren Lancaster, and John Hight. JJ Garcia-Luna-Aceves and Bich Nguyen worked part time and Ed Kozel served on its board of directors.

[16] Others exploring this multi-hop routing were the radio amateurs, Tetherless Access Ltd., whose products failed to reach the market because of technical and management failures, and Metricom, whose routers were not commercially available.

radio. As of this writing, SkyPilot is in a beta-testing phase and completed a $24 million second round of financing.

A final startup using the Packet radio mesh technology was an SRI spin-off called PacketHop. Ambatipudi Sastri and Michael Brown transferred to the new company in 2003. It has a developing role in public safety networks.

A Concluding Perspective

SRI was clearly in the forefront of the new networking age. As with most technical "revolutions," the deployment of this new digital technology took time. Important factors in that wait were the decreasing price-to-performance ratio of electronic equipment and acceptance by the military and the public of real-time online operations. When the technology was developed 25 years ago, almost no one, at least among companies serving the consumer community, was interested in digital radio systems even when they offered the characteristics of the internets and intranets of today. But the revolution was coming. In 1985 SRI researchers helped define the methods for multimedia electronic mail.[W] In the late 1980s and early 1990s the SRI NIC's host-counting measurements were the first to verify the exponential growth of the Internet. Those measurements continued until well beyond the millionth connected host. In spite of attempts by standards bodies to create a replacement for TCP/IP, it remains, after almost 25 years, the protocol for all Internet traffic.

Today, one can easily see and feel the growing communications revolution. With functionality that was already available in the packet radio networks at the genesis of the Internet, you can now, in examples of pervasive alacrity, establish instant messaging, send email, teleconference, co-edit drawings, examine medical records, or obtain comparative prices or order goods on wide variety of wired or wireless Internet terminals, including cellular phones. Even the recent acceleration in the exporting U.S. jobs overseas is abetted by these high-speed digital communications networks, which from the beginning included purposeful design attributes such as "distance- and location independence." The savings in labor far outweigh the costs of very complete and capable computer-based communications systems.

With the recent explosion of wide-bandwidth (now called broadband) wireless services, enabled by a profusion of widely and frequently distributed network entry points and a bevy of wireless-equipped portable computers, there seems to be another, perhaps over-exuberant "land rush" for ubiquitous connectivity encompassing both the Internet and new kinds of interpersonal communications. "WiFi" is marching in as the medium of choice for the frequent traveler; it is also esteemed for its ability to rid the home and the workplace of unsightly and tethering wires.

Though the change around us is abundantly evident, this new way of communicating is so fundamental that its collective capabilities and uses are impossible to predict. After almost 3 decades, it is still an expansive and wild frontier. What is clear is that with its burgeoning capacity, broad connectivity, and innovative services, the boom in both Internet and personal communications will forever change the way we conduct our business, civil, and personal lives.

Communications Aids for the Deaf and Blind

Helping the Deaf to Communicate

In the late 1950s the director of SRI's Communications Laboratory, Ray Vincent, received a telephone call from the director of the University of California's Lick Observatory, located atop Mount Hamilton east of San Jose. The observatory needed a device to work with its cameras that would help take the atmosphere-induced scintillation out of photographs of heavenly objects. The observatory director also suggested that a physicist by the name of Bob Weitbrecht would be a good person to work on it. Weitbrecht was a tall, lanky man with a shuffling gate. He had some practical technical skills, and just happened to be deaf. He joined the SRI Communications Laboratory in 1958 and, with others in the lab, successfully built and demonstrated a control system for an imaging plate that moved quickly and accurately enough to compensate for the apparent movement of the position and size of a star (see

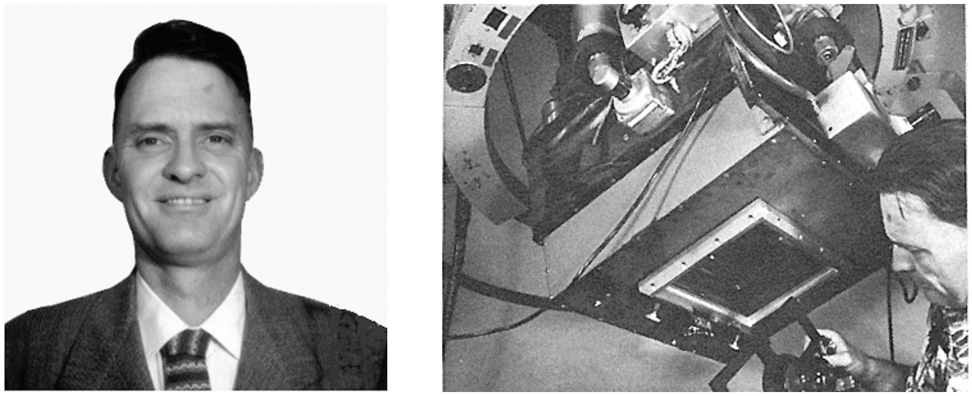

Figure 3-9. Bob Weitbrecht and his Lick Observatory camera.

Figure 3-9). The concept helped extend the useful life of telescopes in many of the world's observatories.[17] But rather than astronomy, it is Weitbrecht's and SRI's contribution to communications for the deaf that we wish to recount here.

Until the 1960s members of the deaf community had no means of interacting with each other or others except through face-to-face signing or lip reading or the use of third parties as interpreters. The ubiquitous telephone system, so useful to others, became a taunting symbol to the deaf. In some instances and at considerable cost, teletypewriters, often surplus equipment from the telephone companies, were connected to leased-direct (i.e., nonswitchable) lines to enable deaf people to link up for a "conversation." In 1964 a deaf orthodontist in Los Angeles, Dr. James C. Marsters, shipped a teletype to Weitbrecht and asked him to create a way to use such a machine over the normal switched telephone system. Resourceful person that he was, Weitbrecht saw that two such machines could link through a telephone if they could simply "whistle" their signals rather than using the hard-wired, direct current pulsing they were designed for. Devices that do such signaling are called modems, standing for the modulation-demodulation task they perform.

These teletypewriters used a five-level character code called Baudot, and so it was natural that Weitbrecht, using such machines, would design his audio signaling system to use those same five levels. SRI applied for a patent for the new device in August 1966, and the patent was issued in two versions in 1970 and 1973. These modems were then made and sold by a small company Weitbrecht started, Applied Communications Corp., in Belmont, California. Using Weitbrecht's modem along with the widely available free teletypes, a deaf person could be in communication for about $300. These connections were then as easily made as any other telephone call. Eventually, Weitbrecht left SRI to devote himself full time to providing this important service through his company. Also, to expand the availability of teletypes, a national organization, Teletypewriters for the Deaf, Inc., now Telecommunications for the Deaf, was formed in the Washington D.C. area. As many as 40,000–50,000 of these machines were in use in the deaf community by the 1980s, and the teletype museum indicates that up to 1 million may be in use today.[18]

But other technologies were being developed in the computer access world that would one day become inimical to the Weitbrecht modem. Ironically, some of the strongest evidence of that conflict would become apparent within SRI as staff members tried to adapt computer-related devices to deaf communications. In the 1960s Bell Laboratories was developing its own modem to help meet the increasing need for access to new time-shared computers using what were essentially modern electronic versions of teletypewriters. The new signaling standard for this modem, however, adopted the eight-level ASCII character set, rather than the five-level code, and used full-duplex rather than half-duplex

[17] The camera was installed on Lick Observatory's 75-year-old, 36-inch refractory telescope in 1961. There it would be used for very long exposures on very faint stars as dim as 19th magnitude that would not have been possible before. The SRI camera was also being designed into a new 60-inch Navy telescope in Arizona.

[18] This information was first accessed on www.deafexpo.org/tty_museum.htm on July 29, 1998. The TTY Museum is now owned and operated by CSD, a nonprofit deaf organization (see www.c-s-d.org).

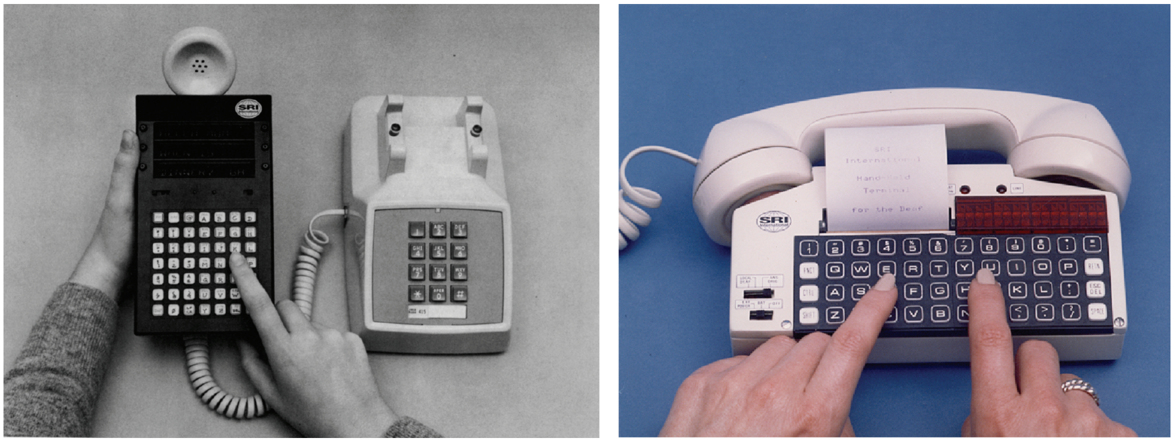

Figure 3-10. Functional demonstration models of small, portable terminals for the deaf (1980).

signaling; that is, both ends could be sending and receiving at the same time. This major new standard, eventually called the Bell 103 standard, was issued in 1967, and it was, of course, incompatible with the Baudot convention.

Another staff member, John Crandall, who joined the SRI Communications Laboratory after Weitbrecht's departure, had a deaf child and, with his wife, cared for other deaf children. Crandall's views on the need for communications among the deaf were welcomed by some other Laboratory staff members who were trying to design and build a portable, handheld digital radio-terminal combination for remote and wireless access to computers (hosts) on the ARPANET. A number of Lab people[19] began to consider the notion of a very small, equivalent teletype, but one that would "speak" both the new Bell 103 standard and the Weitbrecht convention. Their strategy was clear: make a highly portable device that would double as a computer access terminal *and* as a teletype-compatible communications device. They believed the new device would lead the members of the deaf community into the developing, commodity world of inexpensive computer terminals, familiarize them with the emerging uses of computers, including email, and, it was hoped, introduce them as well to the emerging job markets in the computing industry.

The Communications Laboratory staff members knew they could build a small, portable terminal for the deaf, and to help demonstrate that opportunity, the Lab Director, Donald Nielson, had a mockup built. After about 2 years of using the mockup and putting the idea forward, SRI won a contract from the Department of Health, Education, and Welfare's (HEW) Office of Education in 1977 to build several demonstration models.[X] The results are shown in Figure 3-10.[20]

These devices helped illustrate to both the deaf community and the government agencies that served them the need for portable and compatible communications terminals. Ultimately, two SRI staff members, Nielson and Ken Harrenstien, testified before the California State Public Utilities Commission (PUC) about the need for such terminals to use both the new Bell standard *and* the Baudot convention. This exposure helped lead to the subsidization of terminals for the deaf through a small state tax on telephone usage. Unfortunately, the affirmation of the SRI recommendation by the PUC meant that Weitbrecht's company, which restricted its products to the five-level standard, would ultimately face a weaker market.[21]

[19] These included, in the beginning, Stan Fralick and Donald Nielson. The later, more fully developed designs and implementations shown in Figure 2, were created by Dave Fylstra, Bill Ross, Roy Stehle, Russ Wolfram, and Jim Gaddie.

[20] In light of present-day communications technology, perhaps the most striking aspect of these innovative terminals was that they were acoustically or inductively coupled to the telephone. Today they would have modular jacks and would plug into any telephone outlet. It was not until 1968 that the Federal Communications Commission (FCC) issued what was called the "Carterphone decision," which for the first time sanctioned the direct electrical (rather than acoustic) connection of communications device to the phone line. Because such direct connection by consumers had been prohibited by the telephone companies, when these terminals were developed there was no electrical or signaling standard.

[21] To the credit of Weitbrecht and those backing his company, now Weitbrecht Communications, Inc., it is still a successful company over 20 years later. After Weitbrecht's death, the company moved from Belmont to Santa Monica, California (www.weitbrechtcom.com). But the new digital

Once a highly portable terminal had been built, the next step was to introduce the deaf to the benefits of email that had been evident on the ARPANET since the early 1970s. Therefore, in 1977 SRI began developing the concept for an email and conferencing demonstration system called Deafnet. Deafnet started as an attempt by HEW to make the deaf population aware of the advantages of email. Ultimately, its goal would expand to introducing pilot email systems to be operated by deaf organizations and individuals themselves.[Y] Funded in 1978, the SRI project first built a demonstration system that was installed at Gallaudet University in Washington D.C., the nation's university for the deaf. The email system offered service to both ASCII and Baudot terminals, as well as direct conferencing for up to four clients or end terminals. Handling both types of modems was made possible by the clever engineering of SRI's Russ Wolfram, who designed an "intelligent" modem that automatically determined the type of signal being received at the server. The email software was a modified version of that in use on the DoD's ARPANET.[Z]

The number of users at Gallaudet and in surrounding Washington D.C. amounted to about 250 enthusiastic clients. Given that small user base and the operating cost of computers of that day, the service, without subsidy, began to look too expensive. Again, taking advantage of the emerging ARPANET, the network and email coverage was expanded to include mail hosts at Boston and SRI as early as September 1979. Later in the project, smaller, more affordable computer systems were designed that would better fit the budgets of the deaf community. To control costs and to obtain a broad support network for the email servers, a RadioShack TRS-80 computer, shown in Figure 3-11, was selected as the server. A lot of work went into porting the Deafnet program to this smaller machine. Though it worked, a nationwide email and conferencing system was still predicted to cost each user between $50 and $75 per month for moderate usage, and amount considered still too expensive. A larger population was needed to share the capital and operating costs, and there was simply no way to make that available at the time.

To ensure success of the email service in the deaf community, those involved in the project considered it essential for the businesses to be owned and operated by deaf people. The government primed each service by providing the equipment and paying SRI and other organizations to train staff members and help keep the service running for a while. Those who volunteered to start and operate the service were for the most part enthusiastic and were truly pioneers in this fledgling email technology.

SRI, under HEW sponsorship, assembled systems for 9 of an initial 20 metropolitan areas around the country in a so-called dissemination phase of Deafnet. As it turned out, the demonstration project was successful, but the dissemination project was not—only Minneapolis-St. Paul was still in operation a year later. Several factors contributed to the lack of great success in the dissemination phase: the requirement by the federal government that the

communications and computer technologies are also

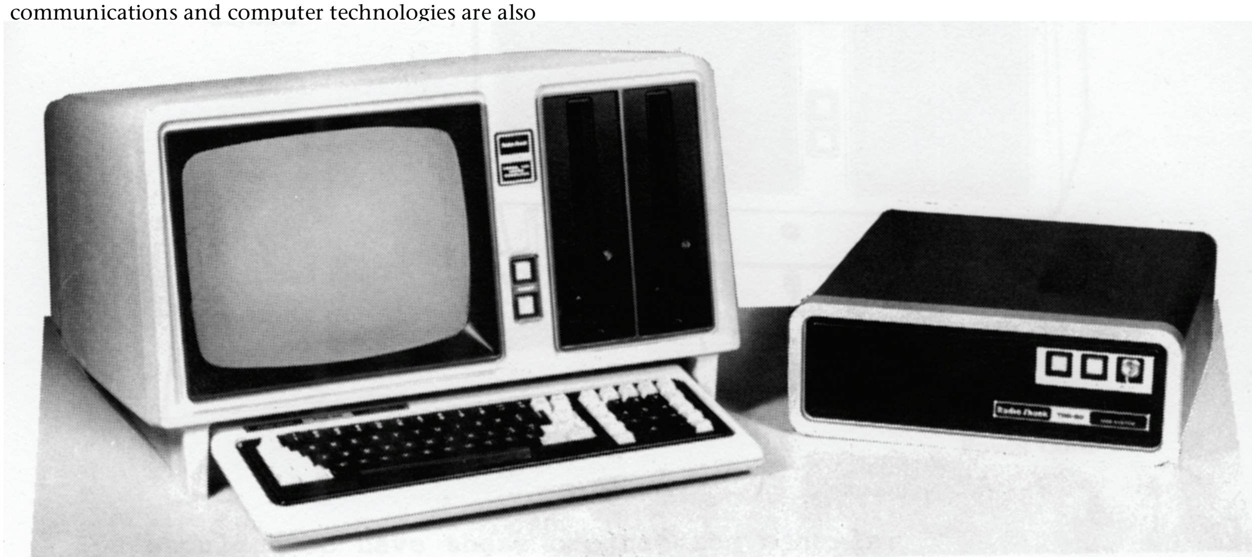

Figure 3-11. TRS-80, the email computer server chosen for dissemination to metropolitan areas.

system become self-supporting before there was a sufficient user base; the high variability and poor reliability of the old teletypewriters as computer peripherals; the requirement for continuous 24-hour service; some high marginal operating costs; and the surprisingly unenthusiastic adoption of the service by the deaf community at many sites. Many deaf people were just not ready to accept the new technology. Beyond technical aversion, many individuals in the deaf community at the time did not see the need for a large email community when the number of people they interacted with was, intrinsically, quite limited. The intelligent modem that could communicate in either Baudot or ASCII was technically successful but less commercially successful than it might have been, largely because of the displacement of the need for it by the rapid acceleration of personal computing and the ASCII code.

To this point, around 1980, all the initiatives in SRI's work for the deaf community had come from the technical or engineering side of SRI. During the dissemination phase of Deafnet, an SRI staff member from the social

SRI'S ROLE IN THE DEVELOPMENT OF ACOUSTIC MODEMS

SRI and its staff made at least three contributions to the development of the telephone-based modulation/demodulation devices called modems. Out of a desire to serve the needs of the deaf community or to support the impending revolution in computer access, three SRI staff members designed, built, and demonstrated modems for which SRI received patents or were involved in new ventures to produce them. Bob Weitbrecht, the deaf engineer mentioned in this section, was probably the first person to conceive of an acoustic modem, and SRI filed his patent in August 1966. Only 6 months later, with almost no interaction with Weitbrecht, SRI's John Van Geen also filed a patent for a Teletype Communication System. Both patents were first issued in 1970.

Van Geen's modem design bore little resemblance to Weitbrecht's. Born of the frustration of costly nation-spanning TWX lines, Van Geen's patent centered on dialup connections and the emerging Bell 103, 300-baud standard for computer access. Modems had existed for a few years, but all were housed within a terminal, which was permanently wired to a set of phone lines. To link to a remote computer required the purchase of a "permanent" telephone connection. More important, at that time it was illegal for anyone but the telephone companies to directly connect anything to their lines. These constraints were severe. Long lead times and high costs set difficult thresholds of use. Van Geen saw the advantages of being able to use any phone and take advantage of the switching system to connect from any phone to, in this case, one of the first timesharing hosts at Dartmouth University. His patented design enabled error-free connections over very long distances (as much as 6,000 miles), well beyond the capabilities of existing modems. To explain how this design and its contribution found fulfillment, a third person must be introduced.

In June 1958 J. Reid Anderson came to SRI from Bell Laboratories to work mainly in the area of magnetic materials. After about 5 years at SRI, seeking relief for his entrepreneurial itch, Anderson started a company called Scientific Products, Inc. One of his first products was an electronic metronome. After a time, and quite incidentally, he met John Van Geen and learned about his modem. Seeing the opportunity, Anderson, with an accountant named Ray Jacobson, formed Anderson-Jacobson in January 1967. The new company's mission, approved by SRI, was to build acoustically coupled modems following Van Geen's design. Van Geen received shares of Anderson-Jacobson stock, and he and SRI received 1% and 3% royalties, respectively, on all their modem sales. Note that this happened long before SRI formally encouraged such arrangements.

During the late 1960s, pressure was mounting to remove the restrictions on the direct connection of "foreign" equipment to the phone lines. In 1968 the FCC permitted the first such connections, although only through a telephone company interface. Since Anderson-Jacobson modems could be acoustically coupled, they were not subject to the restrictions, and they would come to lead the modem industry. Modems, directly connectable after 1977, would become the universal umbilical cord that tied computers to their networks. For some years Van Geen consulted for Anderson-Jacobson, and he was designing a 1200-baud FSK modem and a 2400-baud PSK modem when Anderson and Jacobson had a falling out.

Anderson then left the company along with most of its technical talent. Van Geen departed as well. Anderson later founded Verbatim to manufacture computer diskettes. Anderson-Jacobson was sold in 1988 to CXR Telecom of Fremont, California, and its modem products were incorporated into the offering of the acquiring company.

sciences, Teresa Middleton, whose daughter was deaf, became its project leader. The later stages of the handheld terminal development effort and the majority of the Deafnet projects benefited enormously from the talents of yet another deaf SRI staff member, accomplished computer programmer, Ken Harrenstien, who had been hired by Elizabeth Feinler as a member of the ARPANET's NIC. His technical talents gave him great insight into the terminal project, and because of his particular expertise in the major timesharing software systems of the day[22], Harrenstien became the lead technical person on the design of the Deafnet email system.[23]

It was clear that the deaf community couldn't receive the benefits of a handheld terminal unless some company would accept it as a product. Toward that end SRI ultimately gave demonstrations to about 10 potential commercial manufacturers, hoping to interest them in manufacturing a small, portable terminal resembling the one SRI had built. Although a couple of companies proceeded with their own devices, including Plantronics, none chose to license SRI's technology. The market size was not impressive until the California PUC mandated a price subsidy, which led a few manufacturers to venture into the field. But clearly the general market for small terminals would eventually prosper, particularly with the advent of small microprocessors. Indicative of the era of the SRI models, every LCD character on one of our models cost $40, so 16 characters seemed quite enough. A number of portable terminals for the deaf were developed after the SRI device had been publicized, one of which is shown in Figure 3-12.

Clearly, the largest impact of these projects was raising awareness within the deaf community of the emergence of new technologies from which they would benefit. Even today email is fulfilling only a fraction of its potential, so imagine the difficulty in accepting it in the 1970s. SRI and HEW clearly had an impact. In the earliest days, when the SRI Communications Laboratory director was

Figure 3-12. An early portable terminal built commercially for the deaf community.

just beginning to seek funds for the development of the portable terminal and carrying a mockup to help convey the concepts, the President of the National Association of the Deaf literally embraced Nielson when he saw the communications and mobility that such a device offered deaf people.

Finally, apart from the deaf community, there is the matter of modems and their impact on the world. Because the world's telephone systems are ubiquitous, they became the natural points of departure for the digital highways of the future. By enabling the network connection of computer terminals and their hosts to existing networks, modems accelerated the explosion of computer networking. Today, modems are a necessary part of all computers with network access. The box on page 3-20 and Figure 3-13 recount the important modem developments at SRI.

Helping the Blind to Read—The OPTACON

The Optacon is a device that permits the blind to read arbitrary text. It does so by translating the image of textual characters into magnified tactile images conveyed through a set of vibrating reeds felt with one's fingers. As depicted in Figure 3-14, the user moves a wand, containing a light source and a set of optical sensors, across a printed page. This movement by one hand generates a flow of enlarged characters across the index finger of the other hand. Blind people can be trained to "read" such a character stream at speeds as great as 60 words per minute. "Optacon" is an acronym for optical-to-tactile converter, and the first Optacon was designed and built at SRI and Stanford University in the late 1960s.

[22] A measure of Harrenstein's talent came when, in the 1981 conversion of the ARPANET's communications protocols (from NCP to TCP), MIT couldn't get their large timeshare hosts working. Harrenstien was called back to Cambridge and had them up and running in a few days.

[23] Other major contributors to Deafnet were Earl Craighill, Hal Huntley, Dan Allan, Raphael Rom, and, of course, the dissemination project leader, Teresa Middleton.

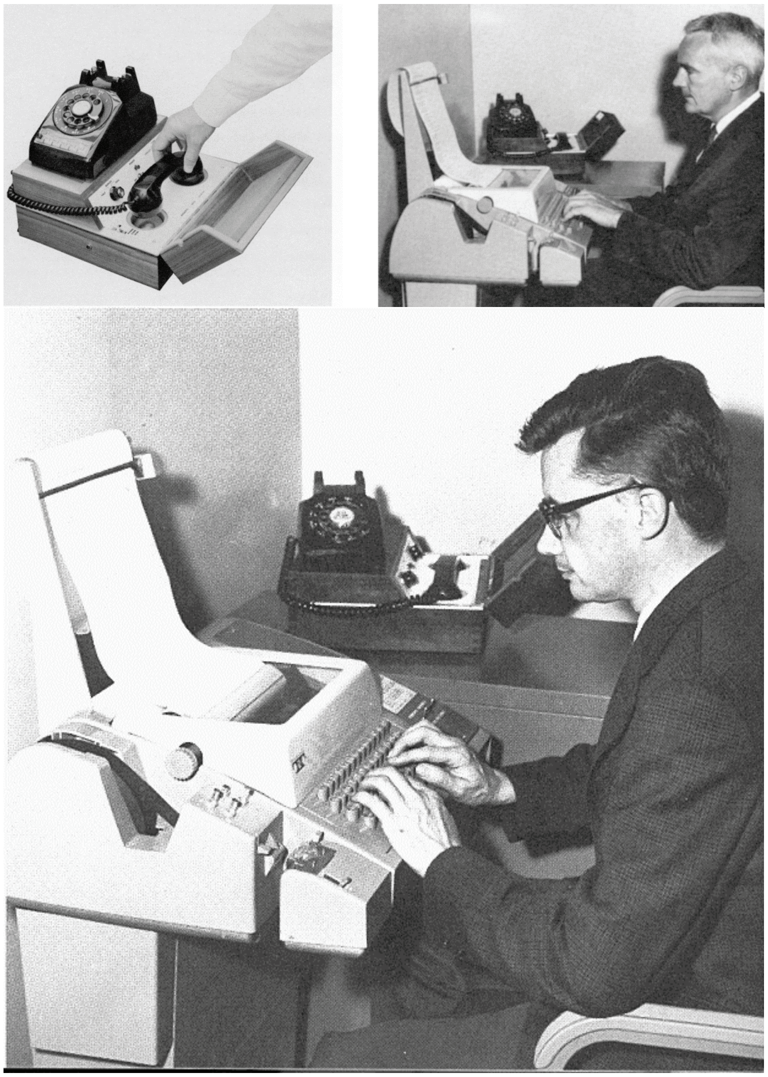

Figure 3-13. The Anderson-Jacobson acoustically coupled modem (upper left) and Reid Anderson (right) and John Van Geen (below) using it (about 1967).

The idea came to Dr. John Linvill, a professor at Stanford, in 1962 when he was traveling with his family in Germany. He had visited an IBM research center near Stuttgart and had seen a high-speed printer in which carbon traces were accelerated against the rapidly moving paper by a line of vibrating pins. His perspective on this new printing method was influenced substantially by its implications for his daughter, Candy, who had been blind since she was 2 years old. He wondered whether Candy could actually feel the impulses of the carbon ink against the paper. Could something be made that converted text into impulses, and could a blind person be trained to "read" them? If so, such a device would vastly increase the printed material accessible to the blind and improve their employability. More immediately, it would clearly ease his wife Marjorie's burden of spending 4 hours each day preparing material in Braille for Candy, who attended a regular school. Linvill began mulling over the possibilities and filling up notebooks of ideas about vibrating pins.

Jim Bliss was enrolled in a Master's program at Stanford in 1956 and working at SRI. He graduated in 1958 and went to MIT to study circuit theory. Because of his advisor there, he soon became interested in devices that would aid the blind. As a thesis he explored whether our kinesthetic senses would support some reasonable level of communications. He had built a sort of reverse eight-key typewriter in which the keys vibrated in three dimensions in response to a symbol input. Bliss also met a blind person, John Dupress, who would help him obtain funding when he graduated from MIT and returned to SRI. At SRI Bliss formed a small group doing research in tactile communication and vision. In 1961 he met John Linvill by chance at Stanford. Linvill started the conversation by asking, "Well, what are you up to?" Bliss answered, "Well, I'm making a gadget to teach reading to blind people. It's a Times Square display."[AA] Over the next 5 years, with SRI funding from the Department of Education (DoEd) and the National Aeronautics and Space Administration (NASA) and the Office of Naval Research (ONR) at Stanford, Bliss and Linvill developed the Optacon.

The Optacon had two custom integrated circuits, one for the 144-phototransistor light sensor and one for the vibrating reeds. The former was the first integrated circuit to be built in Stanford's new Integrated Circuits Lab. From Stanford also came the piezoelectric method of causing an array of tiny metal rods to vibrate in patterns that could carry information. At SRI Bliss and Hew Crane were developing a general-purpose, computer-controlled tactile matrix capable of transmitting static or dynamic tactile images to potentially any part of the body.[BB] The goal was to explore the best methods of

Figure 3-14. Candace Linvill using the Optacon to read a text (about 1968).

tactile or electrical expression to determine the "channel capacity" of various dermal surfaces. This work was sponsored in part by NASA because of the difficulty in voice communications under the high noise of liftoff.

SRI was using computer-generated character streams to explore the question of how well a blind person could learn to read the flowing characters as they moved across the array of reeds. The 12-year-old girl in Figure 3-14 was able to read 5 words per minute after about 90 minutes and 20 words per minute after about 30 hours, where the rate leveled off. Those rates are slow compared to rates of reading Braille with its compressed characters, but the approach had the advantage of working over a wide range of symbols.

The first Optacon emerged from these efforts in 1966, and about 10 units in all were made in the SRI and Stanford laboratories. But then the DoEd requested 50 Optacons for a field trial, and neither SRI nor Stanford was prepared to produce them on that scale. They tried to interest local manufacturing companies but found no takers. Therefore, in 1970, Bliss joined Linvill and two others at Stanford to form Telesensory Systems, which built the 50 units at $5,000 each (see Figure 3-15).

Bliss became the first president of Telesensory Systems and led its expansion into a variety of products for the sensory disabled. By 1991, Telesensory Systems had sales of $30 million, 200 employees, and a worldwide market. The original Optacon sold well for about 15 years, and then it was redesigned. By the end of 1996, when Telesensory announced its discontinuance of the Optacon, nearly 20,000 Optacons had been sold in over 30 countries. It helped people on all continents find better access to the workplace, the ultimate rehabilitation. Like in all important advances, the uses of the Optacon couldn't all be anticipated. Take the case of a blind author by the name of Deborah Kent Stein. Here is her account of getting an Optacon in summer 1977:

> Without a doubt reading with the Optacon was slow. Through steady practice I built my speed to about 100 words per minute, compared with my Braille-reading speed of 250 words per minute or more. But reading speed was not the issue. What mattered was access, and the Optacon provided that. Books, newspapers, magazines, catalogues, bills, record jackets, and the recipes on boxes of cake mix—the barriers were down, and suddenly everything was within reach. For the first time friends lent me their favorite books, sent me clippings, and dared to share their private thoughts in typewritten letters.
>
> 'So what's the first thing that machine helped you do?' my aunt asked when I brought the Optacon home from Philadelphia. 'I cleaned out my purse,' I told her. It was true. I didn't plunge straight into the latest bestseller. I emptied my purse onto the couch and sorted through several weeks' accumulation of receipts, theater programs, ticket stubs, and random scraps. In the past I would have had to wait for the opportune moment with some patient friend or paid reader who could help me weed out the debris. Perhaps I might simply have taken the

Figure 3-15. A blind SRI programmer, Bob Stearns, using the Optacon (about 1974).

matter into my own hands, dumping everything into the wastebasket and hoping I wasn't losing some crucial phone number or appointment slip. Now, with the Optacon, I could check each questionable paper and dispose of it as I saw fit, on my own time, without having to let anyone else glimpse the rat's nest my purse had become.

I have had a Kurzweil scanner since 1990. I no longer use the Optacon for reading full-length books as I often did in the past. But the scanner has never replaced the Optacon in any other regard. They are both tools for accessing print, but each has its own unique strengths and limitations. The scanner can read quickly through large blocks of standard print. It enables me to store material on diskette for future reference, thus building up a small library of books and articles. But the scanner has strong views on what standard print really is. Poor to moderately well-xeroxed copies, most newsprint, all faxes, print that is unusually small or exceptionally large—all call forth the maddening message: 'Page too difficult, may be upside down!' Pages with more than one column may be read accurately, as long as the space between the columns isn't too narrow. Italicized words often turn into strings of 'unrecognized characters.' And anything handwritten, no matter how clearly, is totally out of bounds.

With the Optacon, on the other hand, the only limits are my time and patience. With a bit of both I can read virtually anything. Cursive handwriting is the only holdout; I can usually read handwriting if people print. I can also examine charts and tables and can puzzle out simple line drawings and maps. The underlying fact is that the scanner interprets what it perceives, often in its own idiosyncratic fashion. The Optacon shows me what is on the page and allows me to interpret for myself.

When I got the Optacon twenty years ago, I believed it would be available to blind people for as long as civilization endured. I never imagined that the company that created and marketed this extraordinary instrument would one day renounce it as obsolete. But by the mid 1980's TSI (the descendant of Telesensory) had moved on to other, more lucrative products. It promoted the Optacon, even the newest model, with waning enthusiasm. In 1996 came the dreaded proclamation. The Optacon would no longer be manufactured. Old machines will be serviced 'until the turn of the century,' unless the parts run out sooner. The Optacon is an essential part of my life. In my work as a freelance writer I turn to it a hundred times in the course of the day—to check a page number for a footnote, to make sure the margins are correct on a printed page, to check whether my printer needs a fresh ribbon.

Beyond my working life the Optacon is just as important. I can browse through gift catalogues before Christmas and birthdays. I can sort the mail and read the pieces that are addressed to me. I can use the dictionary, the encyclopedia, and even the Yellow Pages. Without the Optacon I could not do any of these things independently. Each of these small but necessary tasks, plus dozens and dozens more, could be done only with another person's assistance.

The Optacon has given blind people a level of autonomy and flexibility unparalleled in history. Yet that gift is being withdrawn. That sense of freedom, that knowledge that print poses no barriers, may be lost to future generations. As a devoted Optacon user I belong to a minority within the blind community. We spend a lot of time worrying, raging, strategizing, and mourning. We stockpile used machines, buying them up at every opportunity. With renewed hope we pursue each rumor that another company will buy up parts, will service old machines, will build new ones. We tell each other that something has to be done. We try to carry that message to the world.[CC]

As time progressed and the general population aged, low or poor vision became a much more common problem than total blindness. Telesensory introduced products for people with poor vision, too. The Optacon itself also faced the relentless advance of technology in two ways: the introduction of automatic optical character recognition (OCR) hardware and the increased availability of online (as opposed to printed) information affected the market for the Optacon. Telesensory sold the rights to the Optacon in 1998 to another company with products for the blind, Blazie Engineering of Maryland. Bliss left Telesensory in 1994 and started another company, Sensory Technologies, which later became JBliss Imaging Systems. This company concentrates on bringing low-vision software to the personal computer.

Two other derivatives of this SRI effort are worth mentioning. In 1970 the SRI lab used the camera portion of the Optacon to capture and transmit a digitized image of printed characters over the phone to a remote computer that would convert the letters to words and speak them back over the phone to the blind reader. SRI also developed a two-finger version of the Optacon for HEW. Sometime later, in 1975, SRI's Jim Baer led an effort for NASA to design a wrist-worn pager or alarm device that would alert a deaf-blind person. Because of an internal government conflict, it was never released for commercialization.

And what of Candace Linvill (Figure 3-14), the person for whom the Optacon was invented? Today, she has a PhD in psychology from Stanford University and works in a hospital in Redwood City, California. She is proof of the potential of blind people, of the fulfillment that can come to those who work in their behalf, and of the promise of technology to help.

Postal Automation

Americans' appetite for mail has come a long way since 1639, when Richard Fairbanks' tavern in Boston was named the repository for overseas mail. For hundreds of years, up to the mid-1960s, mail was handled essentially the same way as in colonial days. But beginning in the 1950s and for at least the past 40 years, the U.S. postal system has been under continual siege by the voluminous and overwhelming demand for mail services. When the Post Office Department underwent reorganization in 1971 to emerge as the more independent U.S. Postal Service (USPS), it had reached a point of precarious inability to handle the nation's mail. Since then, and even with the family of new, private mail carriers to handle urgent mail and packages, pressure on the USPS has continued. This independent government agency now delivers about 200 billion pieces of mail a year, over 40% of all the world's mail.[DD] As one indicator of its size, in its effort to maintain steady delivery times under the onslaught, the USPS has become the dominant U.S. air cargo shipper, accounting for nearly 45% of all domestic shipments in 1998.[EE]

Meeting this challenge and trying to change a history of labor-intensive pricing clearly required automation. If the USPS has been able to meet its obligation, and it marginally has, its success has been due to an almost continuous commitment to automation, at least in the interior of its mail handling practices. The ends, the pickup and delivery parts, are still costly and historical vestiges that the USPS has seemed unwilling to touch.[24]

SRI has been an active technology developer and supplier to the USPS since the early 1980s, performing basic research, studies, and system design and prototype development. These efforts have drawn on and integrated a broad range of SRI technologies, including imaging, optical character recognition (OCR), artificial intelligence, computing and network architectures, electromechanical control, and design for ease of manufacture. This section describes a few SRI technologies that have given the Postal Service more efficient, cost-effective, and novel postal and intelligent materials-handling services.

The Automatic Machine Reading of Zip Codes

The zip code was clearly an attempt on the part of the USPS to speed the sorting and handling of all types of mail. But unless that code was amenable to repeated automatic machine reading along its journey, its impact would be modest. In the early 1970s the Postal Service (the Post Office until 1971) issued a request for proposals for a high-speed, non-contact method to imprint on letters some form of machine-readable address. Since the zip code was normally written on the envelope by hand or machine, that first machine reading of such varied text was anything but reliable, particularly after it suffered the vagaries of mail handling. The USPS sought a way to read a zip code once, then encode the destination address in binary form and print it onto each letter. At SRI Fred Kamphoefner's Engineering Sciences Laboratory had been working on nonimpact

[24] It was about 1975–76 and the Post Office was drowning in red ink. Postmaster Benjamin Bailar was asking for another postage increase. As a researcher in information technology, I saw a way for the Post Office to become more capital intensive rather than labor intensive; to attack the most expensive part of mail handling. The notion was to place printing terminals in homes and offices that could act as sources and receptors of mail. The reality was that transaction mail (bills, etc.) was the only profitable segment of mail and the most displaceable; that is, the most eligible for a complete bypass of the postal system. I pointed this out to the director of Strategic Planning for the Post Office. A combination of an inexpensive ($50) printing phone and the cooperation of the phone companies to deliver (with verification) transaction mail to these terminals at night was either a threat or perhaps an opportunity if the Post Office wanted to go into that business. He dismissed me with the conjecture that I didn't understand letter carrier unions. Their then current notion was to open envelopes, scan their contents, electronically transmit them to a destination post office, print and re-envelope them, and in doing so ignore 84% of the cost of handling mail, which, of course, lies at the ends of the process.

printing methods since the mid-1960s,[25] so it submitted a bid for this work and won the contract.

SRI's ink jet technology had been used to imprint machine-readable bar codes on a variety of objects, including the backs of checks and credit card sales slips. Once the readable code was imprinted, it enabled the rapid automatic sorting of such transactions and, for both checks and credit slips, became widely used. In each case it was necessary at the first handling to lift the relevant information from the check or sales slip, either by hand or by OCR. For the USPS, this meant reading the zip code. Within a short time and using its ink jet printing technology, SRI delivered an experimental system that was successfully demonstrated at the USPS engineering facility in Washington D.C. The Postal Service, convinced of the utility of this approach to automatic handling, awarded a contract to A.B. Dick to build and install printers across the postal system. As it turned out, A.B. Dick had been developing a production version of the same SRI ink jet technique, which it had asked SRI to review in mid-1972. Although it would then not be involved in USPS hardware for some time, SRI continued to develop inks for such printers for both the Postal Service and A.B. Dick. The ink jet-barcode process is still in use on mail today.

Figure 3-16. The RCP and its designers: Greg Myers, Talia Shaham, and Wayne Cruz.

As an aside, after the USPS contract ended, SRI tried to convince several major mailers of magazines to print the address directly on the cover itself, not on a label that had to be attached. This approach would have displaced another SRI printing technology, the high-speed, label-printing Videograph, (see Chapter 7). As it turned out, no one was interested at the time, although the practice is common today.

Address Processing and Recognition Technology

The total automation of mail handling begins with the absolutely essential task of finding and reading the address of the intended recipient. Over the years SRI has developed methods for the location and interpretation of addresses and bar codes on letters, flats (flat mail pieces larger than letters), and parcels. SRI postal address reading technologies include the novel use of gray-scale image processing and contextual constraints in the OCR process to successfully locate and read mail-piece addresses in spite of poor-quality printing, extraneous advertising, and background interference. SRI's research efforts have resulted in real-time systems that locate the address block and recognize its words. These recognition modules are part of the USPS Recognition Coprocessor (RCP), a system that was deployed nationwide in 1997 (see Figure 3-16). The system, which is implemented as a networked array of Pentium-class computers and processors, processes letter mail at up to 15 mail pieces per second and when first deployed demonstrated a 7% to 12% improvement in the correct interpretation of addresses. Each percentage point improvement in the processing rate leads to annual savings of $9.5 million per year by the USPS. SRI's contribution directly accounts for about one-

[25] See discussion of SRI's relevant contributions to high-speed printing in Chapter 7.

half of those savings or approximately $50 million per year.

The RCP system augments the processing of existing mail sorting equipment by first identifying the location of the destination address and separating the address into lines of text, words, and characters. These characters are then “read” by OCR software and compared against directories of valid postal addresses. The result of the comparison is a delivery point code (11 digits consisting of the ZIP+4 code and 2 digits identifying the specific delivery point), which, just as in the case of zip codes described above, is sprayed onto the mail piece as a bar code. A closer look at the components of the RCP will clarify how it works.

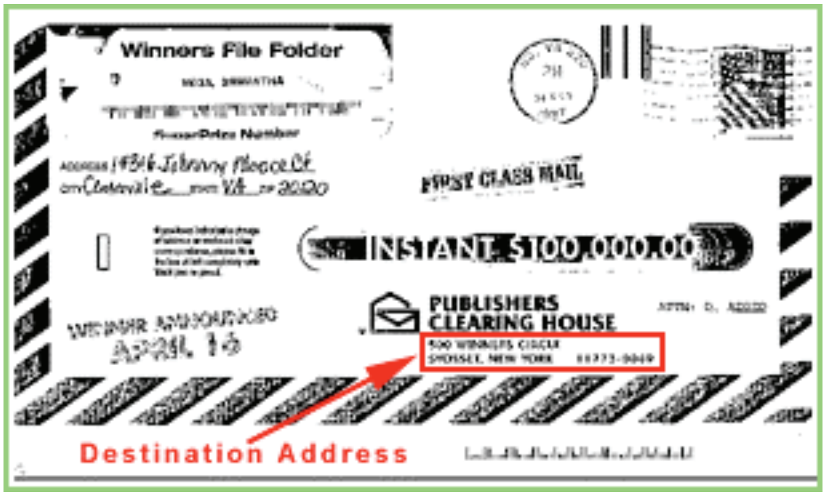

Figure 3-17. A challenging address to identify.

To process a scanned image of a mail piece, an automated system must first locate the destination address and then perform OCR, using existing address directories to interpret the recognized characters as an address. SRI’s methods for address block location (ABL) automatically find the destination address on the mail piece and segment it into individual lines of text. Algorithms can distinguish the address from background patterns and lines preprinted on forms and labels, extraneous advertising, and speckle noise. The ABL technology was implemented as an all-software module in the USPS RCP for letter mail; it has also been implemented in a combination of digital hardware and software for processing images of flat mail. The module has been tested and tuned at the USPS facilities for the last 7 years. SRI has also added the ability to distinguish machine-printed from hand-printed and script addresses, information that can be used to automatically select the appropriate OCR module for each individual mail piece. Figure 3-17 provides an example of the challenges faced in address identification.

To help read the characters in the address, SRI has developed a novel binarization technique for extracting text from gray-scale images of scanned mail pieces or other documents. Conventional binarization methods may perform poorly if the background is non-uniform, or if the contrast between the foreground and background is small or variable, or if noise is present in the image. SRI’s technique, specifically designed to work with images of text, can accommodate local variations in intensity without picking up extraneous interference. In addition, it can detect printed characters with poor contrast. This technique has been implemented in software as well as in custom digital hardware that operates at 40 million pixels per second. Figure 3-18 shows the use of the SRI technique on address text printed on patterned backgrounds.

To gain high reliability in character and word recognition, the SRI technique uses contextual knowledge about addresses. This context-based OCR technique combines the characters recognized by the OCR engine into words that are part of an intelligible address, thus taking full advantage of the interword contextual constraints inherent in an address. This SRI approach is much more powerful than conventional recognition methods and has been implemented and deployed in the USPS RCP.

As an example of context-based “word” recognition, Figure 3-19 shows a poorly printed address line on top. Just below is the character hypothesis produced by the OCR subsystem. Subsequent arcane numbers reflect processing steps and the correct city-state-Zip combination is found at the bottom. It could not have been found so quickly and reliably by examining combinations of characters alone.

As pointed out earlier, mail is imprinted with zip addresses encoded as barcodes. SRI developed a Wide-Area Bar-Code Reader

Figure 3-18. Results of SRI binarization technique for extracting text from patterned backgrounds.

(WABCR) system for letter mail that reads the so-called POSTNET bar code (which contains Zip code information). Based on a charge-coupled linear-array scanner, the WABCR can locate the bar code anywhere on a letter mail piece. Thus, bulk mailers can print the bar code as part of the address block, which can be variously positioned on the envelope. Since the bar code pattern may also be printed over advertising and patterned backgrounds, extracting it with traditional processing methods, such as with fixed thresholds, is not reliable. To make such reading possible, SRI developed gray-level image processing techniques to locate and interpret the bar code. The system demonstrates an improved immunity to the presence of text and background interference; the gray-level processing also gives the system a high tolerance for poorly printed and skewed bar-code patterns.

The Processor Horserace

Throughout the development of computers, and perhaps even today, a race is under way whose outcome is starting to be clear, at least for a time. The race is between special-purpose hardware and software, designed that way to meet some real-time or complex computational need, and the relentless march in the capabilities of commodity processors such as the Pentium. The difficult and demanding tasks of the USPS were witness to this race in countless ways. During the development of the RCP mentioned above, several design directions were pursued to meet the timing or difficulty demands of the automatic handling of mail. In the early 1980s the only way to automatically find and read an address was to employ a VAX minicomputer from DEC. Even then each piece took an incredible 20 minutes! That totally useless outcome might have been grounds for discontinuing the pursuit of automated mail handling unless you had faith in Moore's Law and thus believed that processing power would someday be adequate. After the VAX, several

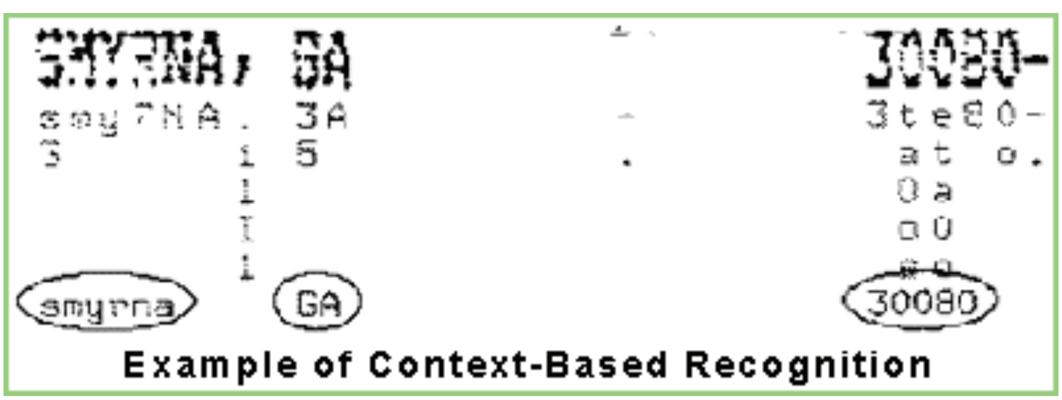

Figure 3-19. An example of context-based address recognition.

A SORTER…SORT OF!

One of the sorting tasks in mail handling is to separate letters from so-called "flats," which are large, oversized, sometimes-stiff envelopes. SRI had been working on an automatic system for this type of sorting. The input was a continuous flow of vertically standing envelopes and flats of all sizes. To estimate the size of the piece first in line, a light source was placed to one side of the stream so that shadows from the individual pieces of mail were created. An image recognition system was built that was able to determine the size of the next piece to be sorted. To physically move the piece for sorting, a vacuum system attached to the end of a robotic arm was created. Its vacuum interface to the piece could be adjusted in area and vacuum power to match the size of each piece of mail. The entire process was controlled by a microprocessor. One Saturday, Greg Myers, one of the engineers on the project, was bringing some friends by SRI to show off the sorter. Without first checking the software state of the machine, he switched it on in what turned out to be a random state left from the last person who was modifying it. To his embarrassment and the glee of his visitors, the sorter proceeded to throw letters and flats all over the room. The invisible hand of software governs both success and failure.

special-purpose machines were built using the 6800 or 68000 microprocessors common at the time. In this case the struggle for adequacy has been won by the commodity processor. Today, in 800 USPS locations across the country, the RCPs are sorting 15 pieces of mail each second. The processor consists of 10 Pentiums running in parallel, all controlled by an eleventh. It is an enormously powerful and reliable system for which SRI designed and built the software and Northrup Corporation built the hardware shown in Figure 3-16.

Given the tremendous number of letters that are handled by the USPS every day, it seems aggressive, if not presumptuous, to say that an SRI address processing software recognizes and parses the address on every one of them. It just happens to be true.

Over at least a couple of decades SRI has undertaken a wide variety of tasks for the U.S. postal system. Almost all of them are difficult and the progress toward a successful outcome is not always linear. The box depicts a small, humorous example along the way.

Endnotes

[A] Peter H. Salus, *CASTING the NET—From ARPANET to INTERNET and Beyond*, Addison-Wesley, 1995.

[B] Janet Abbate, *Inventing the Internet*, MIT Press, 1999.

[C] Peter H. Salas, op. cit., page 25.

[D] E.B. Shapiro, *A Study of Computer Network Design Parameters*, Final Report to ARPA, SRI Project 7016, December 1968.

[E] Peter H. Salas, op. cit., page 28.

[F] Internet RFC 1000 and the ARPANET Completion Report.

[G] Janet Abbate, op. cit., page 74.

[H] William Duvall, personal communication, May 10, 1999.

[I] Douglas Engelbart, SRI Quarterly Progress Report 5 Covering the Period 8 August 1969 through 7 November 1969, SRI report to ARPA and Rome Air Development Center, November 26, 1969.

[J] Elmer Shapiro, personal communication, December 7, 1998.

[K] Richard Watson, "Notes on a Conversation with Bob Kahn on the ICCC," RFC 372, July 12, 1972.

[L] R.E. Kahn, Steven Gronemeyer, Jerry Burchfiel, and Ronald Kunzelman, Advances in Packet Radio Technology, Proc. IEEE, 66(11), 1468-1496, November 1978.

[M] Barry Leiner, Donald Nielson, and Fouad Tobagi, eds., Special Issue on Packet Radio

Networks, Proc. IEEE, 75(1) 1–176, January 1987.

[N] Janet Abbate, op. cit., page 127.

[O] Vinton G. Cerf, personal communication, January 15, 2002; additional comments by Cerf are available at www.wcom.com/about_the_computer/cerfs_up/qa/answers.phtml.

[P] Vinton G Cerf, and Robert E. Kahn, A Protocol for Packet Network Interconnection, IEEE Transactions on Communications, Comm-22 (5), 637-648, May 1974.

[Q] Jon Postel, Larry Garlick, and Raphael Rom, *Transmission Control Protocol Specification*, SRI report to the Defense Communications Agency, July 15, 1976.

[R] Ronald C..Kunzelman, Michael A. Placko, and Russell T. Wolfram, *Progress Report on Packet Radio Experimental Network*, SRI Quarterly Technical Report 5 to DARPA, September 1977. This report also documented multiple PRNET terminals communicating simultaneously with multiple APRANET hosts.

[S] R.C. Kunzelman, V.D. Cone, K.S. Klemba, J.E. Mathis, J.L. McClurg, and D.L. Nielson, *Progress Report on Packet Radio Experimental Network*, Quarterly Technical Report 10 to DARPA covering the period from November 1977 through January 1978, February 1978.

[T] Don Nielson, "The SRI Van and Computer Internetworking," *CORE*, a publication of the Computer History Museum, February 2002.

[U] Bill Croft, personal communication, December 12, 1999. Bill Joy led the installation of TCP/IP in UNIX (v4.2 and v4.3) while still at the University of California, Berkeley. Croft eventually left SRI to work for Joy one of the founders of SUN Microsystems.

[V] Kelli Abe Trifonovitch, Plans for a Wireless Waikiki Could Put Hawaii on the Technology Map, *Hawaii Business Magazine*, December 2, 2002.

[W] A. Poggio, J.J. Garcia-Luna Aceves, E.J. Craighill, D. Moran, L. Aguilar, D. Worthington, and J. Hight, CCWS: A Computer-Based, Multimedia Information System, IEEE Computer Magazine, 18(10), 92-103, October 1985. (The first two authors were guest editors of this special issue on multimedia communications.)

[X] David J. Fylstra and William C. Ross, *A Communication Device for the Deaf or Mute Person*, Final Report on SRI Project 6964 with the Dept. of Education's Office of Special Education and Rehabilitation Services, August 1980.

[Y] Dan Allan, Earl Craighill, Shuel Oreu, Charles Jackson, Susan Russell, Harold Huntley, and Jane Wilson, *A Nationwide Communications System for the Hearing Impaired – Strategies Toward Commercial Implementation*, Final Report to the National Telecommunications and Information Administration, October 1981.

[Z] Kenneth Harrenstein, Earl Craighill, David Fylstra, Harold Huntley, William Ross, and Susan Russell, *Deafnet: A Distributed Communications Service for the Deaf*, Final Report on SRI Project 7883 to the Dept. of Education, May 1982.

[AA] *San Jose Mercury News*, August 14, 1988.

[BB] James C. Bliss and Hewitt D. Crane, Touch As a Means of Communications, *SRI Journal*, Feature Issue No. 5 on Human Senses and Systems, 2-15, January 1969. This article reflects how extensively Bliss and Crane had been looking at the various means for tactile sensing and control.

[CC] The web site of the National Federation of the Blind: http://www.nfb.org/bm/bm98/bm980506.htm.

[DD] According to Hoover's Online business information source (www.hovers.com).

[EE] Colography Group, an Atlanta-based research firm that specializes in the airfreight and air-express business, (www.colography.com).