Chapter 2
Innovations in Computing

Computing Research at SRI

Because of its importance to SRI and the world, research into computers, including their interconnection and some of their underlying technologies, will comprise the next three chapters of this book. SRI has been engaged in various aspects of computing research since 1950 when the design of practical computers was in its very infancy. SRI involvement began at that time with an enormous challenge: to build a reliable, real-time, transaction-oriented, error-free computer that could input checks and update bank account statements. The successful completion of that work over the ensuing five years also led to two other notable SRI computer research activities that will also be described in this chapter: automatic handwriting recognition and computer technology. To appreciate SRI's contributions to computing also requires the depiction of its role in the conversion of computing from a detached, batch-oriented convention to an interactive, highly responsive one that now dominates nearly all of computing. Rounding out this Chapter will be the important story of SRI's work in automatic speech recognition, including its entry into the commercial marketplace. All of this work, including that covered in the following two chapters, has left an important impact on the inevitable influx of computing power into our lives.

ERMA—The First Banking Computer

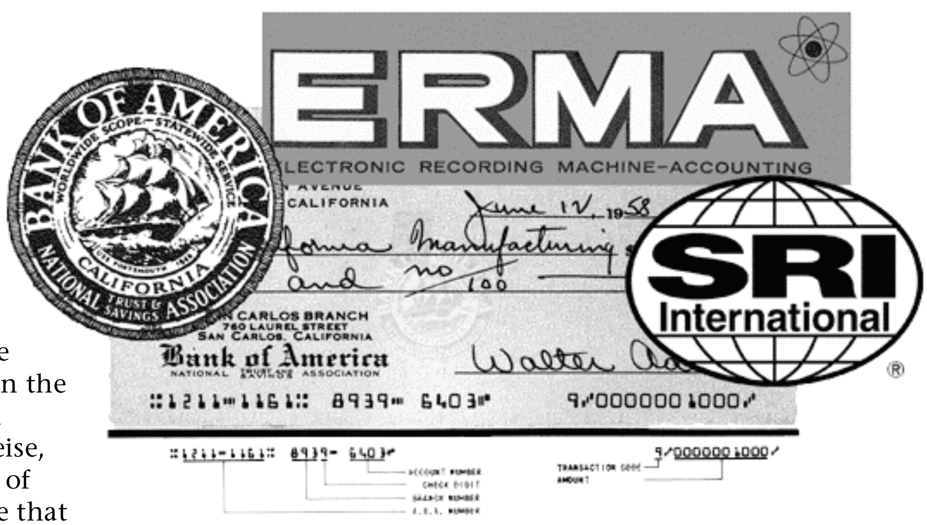

The Genesis

SRI's involvement in electronic banking began almost accidentally.[A] In the spring of 1950, SRI Vice President Weldon "Hoot" Gibson and a couple of colleagues were promoting SRI at the San Francisco headquarters of the Emporium retail stores. When the meeting ended early, Gibson decided to call on S. Clark Beise, an acquaintance at the Bank of America (BofA). He told Beise that although it was a "shot in the dark," he thought the Bank ought to be looking at possible electronic banking applications. He also expressed doubt about whether the BofA would ever get what it needed from IBM or other providers of banking equipment. What Gibson didn't know was that the BofA was already laboring under the rapidly growing volume of checks that, handled manually, was threatening to prevent completion of daily account updates.[1] Gibson gave Beise the telephone number of the head of SRI's

[1] In the late 1940s, then BofA President Mario Giannini, the founder's son, scoured the banking equipment industry to find help managing what in 1950 came to be 4.6 million accounts. He found no one ready or willing to meet his challenge. When Senior Vice President Clark Beise informed him of the SRI alternative, he authorized Beise to secretly contract with SRI. When Giannini died in 1954, Beise became President, ensuring the continuation of ERMA at SRI. ("The ERMA Chronicles," Bank of America Technology Magazine, 156, June 2000.)

Engineering Group, Dr. Tom Morrin, and then left. A few days later BofA contacted Morrin to initiate the innovative partnership with SRI that resulted, over the next 6 years, in development of automated check handling, the first computer for banking applications, and a magnetic ink reading system still in use today.

In 1950 BofA was the largest bank in the world, and Beise was a senior vice president and one of the few members of the Bank's upper management actively leading the search for machine-based innovations. He also was unique among senior management in realizing that check handling, not creating new business accounts, would be the major factor limiting growth. Automation was essential to keeping pace with the Bank's growing business, and the business equipment manufacturers Beise had talked to were not interested in the investment needed to create a whole new electronic banking system. But SRI could act as the Bank's R&D arm to show the manufacturers such a system could be built and perhaps to sell them the prototype.

At that time the BofA's checking accounts were growing at a rate of 23,000 per month, banks were being forced to close their doors by 2:00 PM to do the proofing and processing necessary to finish daily posting, and the situation was only going to get worse. As a measure of the cumbersome process, consider that unless a check was deposited at the bank branch where it was drawn, it had to be sorted by hand and individually entered into an adding machine at least six times during the clearing process. Business machine companies were unwilling to invest in new bookkeeping machines that might hurt sales of existing equipment, and the few computers that existed were viewed as calculators rather than record-keeping and accounting systems.

Assessing the Feasibility

The first meeting of SRI and BofA staff took place in June 1950, when the Bank's Frank Dana arranged through Morrin for Oliver Whitby and Joseph Lovewell to visit the Palo Alto Branch to view banking procedures in action and to assess the magnitude of the check-handling problem. This visit was followed by a number of meetings between SRI and the Bank wherein they agreed that together they would address the following basic functions: credit and debit all accounts, maintain a record of all transactions, retain a constant record of customer current balances to be printed as needed, respond to stop-payment and hold orders on checks, and notify the operator if a check caused the account in question to be overdrawn.

Of course, strict accuracy was mandatory. Errors in arithmetic or in the integrity of an account simply could not go undetected. Also, the bookkeeping system had to be fast enough to make sure that all accounts were posted and reconciled each day. The system was not intended to assume the role of proofing the checks and sorting them by account or for distribution to the bank of origin. Operators would key in the values of the checks, and the system would then compute the new balance.

In July SRI was asked to embark on a feasibility study of an electronic bookkeeping machine to handle the five functions mentioned above, stressing the following performance attributes: speed, as dictated by the banking time schedule; handling of all the information needed for storage, processing, and printing; and the ability to provide up-to-date balance information for customers. The work was designated as "client private" so that only SRI and BofA staff members working on the project would be aware of it.

Thus the "whiz kids," as the BofA called the SRI project team members, went to work weighing the requirements of the new bookkeeping system against the available technology. Recall that only a few computers existed anywhere, and they had nothing to do with accounting. Transistors were just becoming available but were very unreliable, and the only magnetic media that would be large and fast enough were magnetic drums. But new available technology was always a point of discussion. From the BofA's perspective, the nature of the check should be changed as little as possible because checks were the interface between the bank and its customers, about whom the Bank was solicitous. But SRI immediately noticed that the blank checks of the day had no identification on them and were, in fact, passed around for anyone to use. Because accounts were kept alphabetically, opening a new account required an awkward reshuffling of the account list. At SRI's urging, the Bank agreed to require checks with preprinted account numbers. Each new account, then, would simply be added to the end of an account list, leaving all preceding numbers intact. The BofA's acceptance of this approach made it possible to combine the formerly separate processes of proofing and bookkeeping.

In September 1950 SRI completed a written report on the feasibility study. The report stated that an automatic bookkeeping system could be built that would satisfy the BofA's requirements. SRI called the proposed system an electronic recording machine, or ERM, and Morrin suggested a three-phase approach to system development:

- Study the banking procedures external to the machine
- Create the general logical design
- Build the system and test it.

The last step would not be done by SRI, but by an equipment manufacturer. In mid-November 1950, the BofA awarded SRI $15,000 over 6 months to complete the first two phases. Even with a $5,000 payment added in April 1951, the compensation seems minimal for such a large undertaking. SRI's interim report was delivered on 30 April, as promised.

The ERM was outlined as a large computer-like system for bookkeeping only. Checks would be proofed and sorted at the branches and then sent to the ERM for posting to individual accounts identified by numbers that would be optically readable on each check. All transactions would be recorded and printed by account number, and a statement reflecting them would be sent to customers once a month. The following numbers reflect the ambitious goals: Each ERM would handle 32,000 accounts and on the average process about 48,000 items per day. (ERM processing begins after the amount of each check has been added at the bottom in machine-readable form.) SRI's estimated cost for each machine ranged from $530,000 for a minimal configuration to $830,000 for a complete one.

After the numbers were refined a bit, the BofA took the bookkeeping system descriptions to a number of manufacturers, but only Burroughs expressed any interest. Because of existing product lines, Burroughs suggested that it convert its bookkeeping machine to an ERM and also offer a proof machine and a printer. After a few meetings, however, Burroughs withdrew its interest, saying the costs would be double SRI's estimate of about $1 million.

Whether this interplay was just a test of the SRI estimates, as Morrin suggested, will never be known, but Beise soon asked SRI to proceed with an ERM prototype to give the manufacturers evidence that such a system could be built for about the amount that SRI had estimated.[2] Although building prototypes was usually outside the role of a research institute, Morrin reluctantly agreed to proceed through design and construction. In January 1952 SRI signed a contract for $875,000 to develop the prototype.

SRI's Reluctant Continuation

The working relationship between SRI and the BofA was very close. SRI was responsible for the technology and worked with the Bank to examine the banking process and discover how it might be modified commensurate with what the equipment could or could not do. One of the most important contributors from the Bank was Charles Conroy, Assistant Vice President. He worked daily with the SRI designers formulating the essential compromises between what the banking system needed and what technology could deliver at a workable price. The BofA had also brought other people on board, not directly part of the ERM team, to help it prepare for the advent of computers, be they ERM or not. The need was becoming evident to everyone.

The next few years brought a whirlwind of innovation and yet a continuing dilemma. The SRI team wanted the performance potential of new technology but had to use components that were proven reliable and affordable. For example, transistor-implemented logic, though fast, small, low power, and potentially reliable, could not be used because the early manufacturing methods left them quite unreliable. Even the keyboard circuits that encoded each key had to be built. High-speed printers, needed for statements, were both experimental and preponderantly mechanical. In addition, magnetic tape drives for large-scale

[2] In a letter to the authors of the Harvard case-study (Op. cit.,[A]) leader, SRI's Tom Morrin wrote the following: "Mr. Beise told me…after his retirement…[that] in making the decision about Burroughs, et al., he called David Sarnoff at RCA, for whom we had done lot of work in color television, transistor circuitry, special tubes, etc., and was told that SRI was the only one to do the job, but to double our cost estimates. He knew, of course, that in advanced research there are no really good ways to estimate costs. One has to decide, as Beise did, as you go along whether it will pay out."

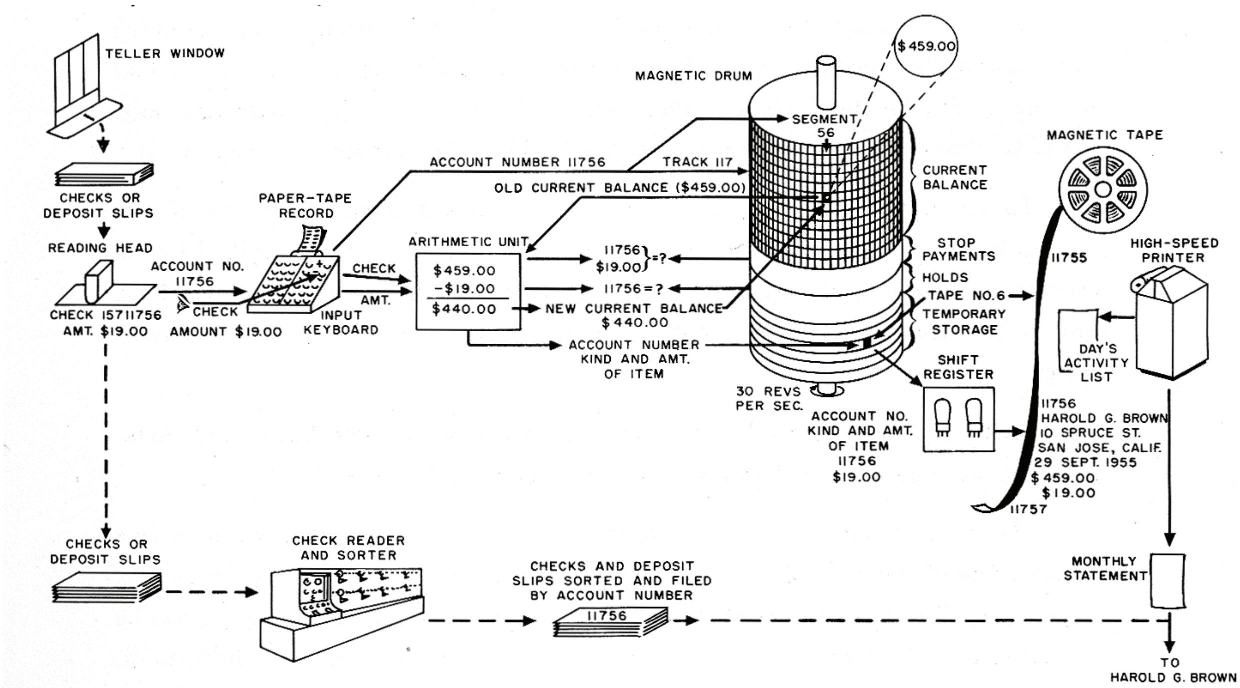

Figure 2-1. The ERMA check handling and accounting process.

storage had to be specially developed (in this case, by ElectroData, a subsidiary of Burroughs).

As it developed from 1950 through 1955, the ERM used magnetic drums and tapes for storage and vacuum tubes, silicon diodes, and relays for logic. Unlike other (batch-mode) machines of the period, the ERM made account data available on line to validate inputs and to respond to inquiries about account status. To understand the chronology of ERM, we must appreciate that stored-program computers were only a few years old and still in a laboratory state when SRI began its work. The first real-time computer, the MIT "Whirlwind," was also being assembled at MIT. It used magnetic core memory, but that technology was not available to SRI. Clearly, no existing system was close to the functionality needed by the ERM…it was simply a whole new machine! As the machine evolved, so did its name. ERM was difficult to say and the Bank wanted a name more easily communicated. As a result, "accounting" was appended to ERM and it acquired its more familiar and approachable title, ERMA.

To understand the rudiments of ERMA, it helps to view it in terms of the trajectory a check takes through the system. In Figure 2-1 the lower dashed line shows the path of the physical check and the accounting relevant to the check is depicted along the upper route. Here a number of tests are made against the account, including its present balance. A report of transactions and the canceled checks are returned once a month to the holder of the account.

One of the early technical obstacles was the check reader/sorter. According to Bill McGuigan, then an assistant director of Engineering at SRI, Beise made it clear to Morrin that the BofA did not want to have to deal with punched cards.[B] Transferring all the check's information onto another medium would just encumber the process. Thus the check or deposit slip itself had to be read, at least for the account number and perhaps the check number. Because they were handwritten, the amounts stated on the checks required human entry. SRI was then faced with building a check reading and sorting system that would be infallible. The check-reading task was addressed by Ken Eldredge, Fred Kamphoefner, Phil Merritt, and others, while the mechanical design of the sorter was done by Alonzo W. "Bill" Noon. Noon's team produced a sorter able to handle about 10 checks per second with errors of less than 1 per 100,000. (see Figure 2-2) The early success of this sorter showed that Beise's request for a cardless system could, in fact, be met.[3,4]

[3] According to project supervisor Jerre Noe, when Thomas J. Watson of IBM visited the project and saw the high-speed check sorter working so well, he became noticeably concerned about this alternative to his punched-card realm.

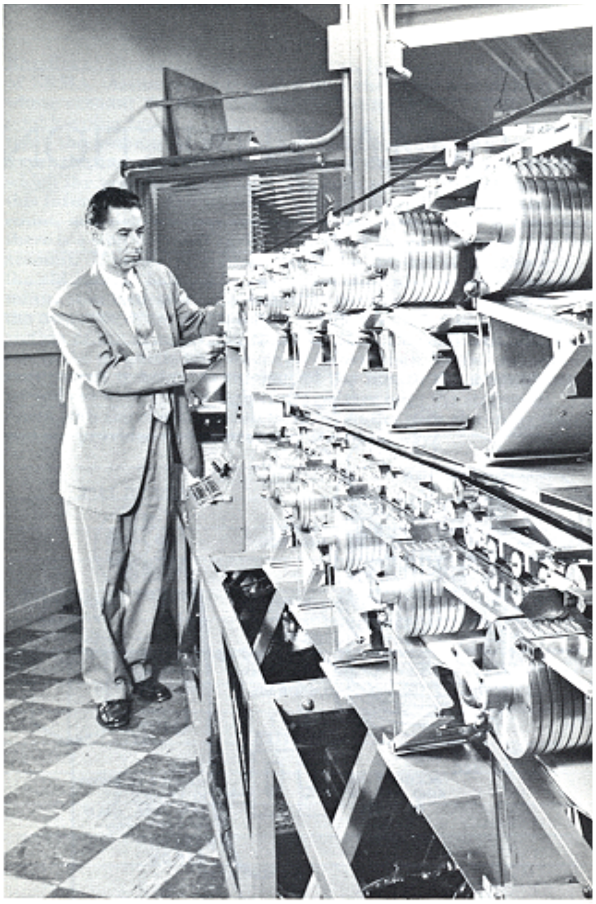

Figure 2-2. The check sorter for ERMA with Ken Eldredge, the creator of MICR.

Figure 2-3. The ERMA that SRI demonstrated to the BofA in September 1955. (The control console is at upper left and the input machine is at mid-left. The SRI staff are (from the left) Dr. John Blickensderfer, Joyce Trainer, Mari Watson, and Maury Mills.)

The final ERMA computer, part of which is shown in Fig. 2-3, contained more than a million feet of wiring, 8,000 vacuum tubes, 34,000 diodes, 5 input consoles with electronic reading devices, 2 magnetic memory drums, a check sorter, a high-speed printer, a power control panel, a maintenance board, 24 racks holding 1,500 electrical packages and 500 relay packages, 12 magnetic tape drives for 2,400-foot tape reels, and a refrigeration system. ERMA weighed about 25 tons, used more than 80 kW of power, and had to be cooled by an air conditioning system.[c]

Following the successful demonstration of ERMA at SRI, the BofA engaged General Electric Corporation (GE) to manufacture 40 machines for installation in California. SRI worked with GE to transfer the basic processing algorithms to the new architecture. The GE version of ERMA used transistor logic and magnetic core memory, both still in the experimental stage when the SRI machine was being designed.

Creating Readable Magnetic Fonts

A second major innovation in the project was the preprinting of checks with individual account numbers that could be read automatically with great fidelity as well as by humans. Solutions such as fluorescent ink that was invisible to the eye were initially considered, but they to be too vulnerable to pen and pencil marks that would produce reading errors and, more importantly, would require that banks change many of their bright check cancellation inks. During the project, the BofA also considered an optical reading system from a company in Arlington, Virginia, that offered one of the first optical character readers. But both fluorescent and non-fluorescent optical methods were too prone to errors from overwriting.

By about mid-1952, SRI had decided the solution was to create a machine-readable font using magnetic ink. By late 1952 the check owner's account number was being imprinted on the back of the check in both magnetic barcode and optical characters so that the number would be readable by both humans and

[4] According to Fred Kamphoefner (personal communication, February 10, 2001), Watson also said his "engineers told [him] that it wasn't possible!" At the time the SRI sorter was sorting stacks of used checks that varied by up to two to one in length, width, and thickness, plus an occasional IBM card and an old dollar bill.

When the BofA chose September 1955 for a final demonstration of ERMA, SRI's project team began working multiple shifts to create a minimal but working version of an automated bookkeeping machine. But the check sorter continued to have problems right up to the last moment. The appointed "ERMA Day," September 22, 1955, was carefully managed by BofA. The demonstration was held at SRI in Menlo Park, and reporters were bused there from San Francisco to prevent leaks. Beise of BofA and Morrin of SRI, without mentioning either organization, described the great contribution this system would make to banking. But before the demonstration, and because of the tenuous reliability of the check sorter, a cue from the SRI engineers to Morrin was needed to show the sorter was indeed working. After a slight delay, the sign was given, and the demonstration went perfectly, and the reporters left to file their stories. That evening, at a private showing for SRI employees, the sorter spewed checks all over the room! (*IEEE Annals of the History of Computing*, 1993, op. cit.)

machines. The account number was also preprinted on the front bottom of the check for convenience. This technique was used throughout the remainder of the SRI project, including in the prototype demonstrated in September 1955. The days of the barcode were numbered, however, partly because of esthetics but mostly because of a realization that magnetic-based Arabic characters could be made machine-readable. This innovation was to arise out of a companion SRI project from the Bank having to do with traveler's checks.

While traveler's checks didn't have an account number, they did have a unique serial number. Traveler's checks were processed in much smaller volumes than conventional checks, and the BofA's processing method involved first creating a punched card to be processed. The ability to automatically read the check serial number was helpful and would, incidentally, solve the identical reading problem faced in ERMA. To avoid the optical problems in reading both checks, Eldredge, by then director of SRI's Control Systems Laboratory, devised several test patterns and magnetic-ink fonts that proved to be reliably read by both humans and machines. (see the traveler's check font in Figure 2-4) Phil Merritt, then a graduate student at Stanford University, built the recognition circuitry. By applying some new matched filtering techniques he had picked up, Merritt designed and tested a second, more normal looking font (the left font in Figure 2-4). The third and final font is the stylized one that now appears on all checks, consisting of the 10 integers plus a few ancillary characters. Given a readable Arabic font, the account number could be moved to the front of the check, where it remains today on all checks. Figure 2-4 shows the three different fonts:

- SRI's traveler's check font is shown in the upper righthand corner of the check (partially obliterated by a purposeful stain) and, in draft form, in the top center of the lower portion. Immediately below that is a profile of the sum of magnetic flux in a vertical slit scanned across the character.
- The first SRI font for ERMA is shown in the left inset, along with traces that represent the derivative of the sum of a vertical slit scanning each magnetic character.
- The final ERMA font that became the

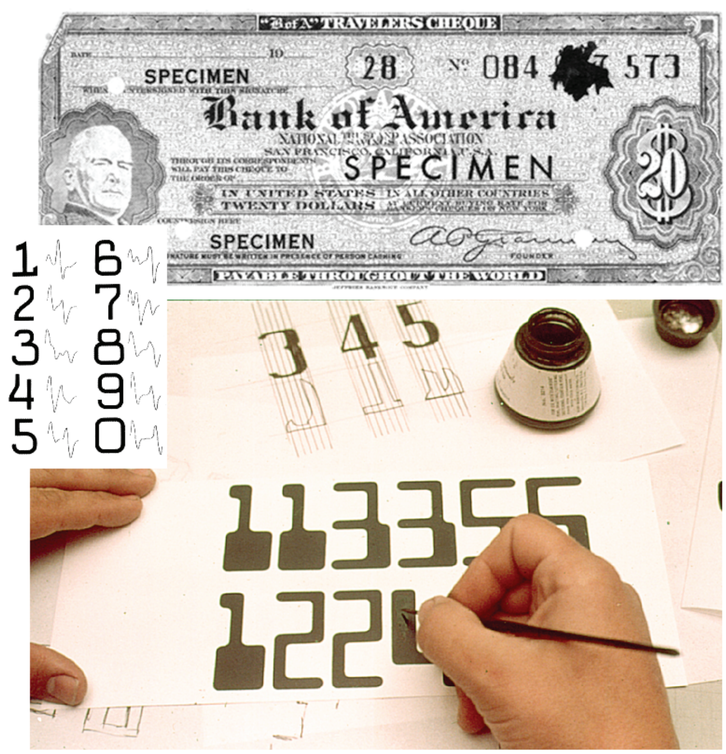

Figure 2-4. Evolution of machine-readable fonts for checks.

international standard is at the bottom.

The third font, shown at the bottom of Figure 2-4, was, as it appears to be, the work of a committee, specifically the Bank Management Committee of the American Bankers' Association (ABA), which began deliberating about the font in July 1956. Several manufacturers offered input into its peculiarities. SRI called the font MICR, for magnetic ink character reading, but after the ABA's acceptance, it also became known as the ABA Common Machine Language. It has since evolved into an international standard. Fortuitously, and perhaps because of its implied importance, the U.S. Patent Office awarded SRI and Eldredge Patent No. 3,000,000 for a magnetic-ink-based character reading process. The magnetic ink character reading worked so well that it was common to demonstrate its resilience by marking and then wadding up a check, unraveling it, and putting it through the reader to show that it still could be read.

Traveler's checks were easier to work with than commercial bank checks because they were all the same size and were issued in only a few denominations. So-called normal checks of that day might vary as much as two to one in length, width, and thickness.[5] The MICR font for checking account numbers, while more difficult than other fonts for humans to read, has been extremely reliably read by machines. The BofA's Al Zipf, who was quickly becoming the Bank's best expert in data processing, helped transform MICR into an ABA standard. He later led the transition of ERMA from prototype to production.

The Production of ERMA

Once SRI had proved the feasibility of processing checks directly without creating any surrogate paper, it was time for ERMA to go into production. According to Jerre Noe, 24 companies had at least some interest in bidding on the job, in contrast with mild interest from Burroughs alone before SRI demonstrated the prototype. In the interim, technology had advanced, SRI had proved the existence of a working system, and probably most important, a clear market for such machines had emerged.[6] GE won the competition to produce the system commercially.[7] When GE's Industrial Computer Section began to consider the potential market for ERMA in mid-1956, it asked SRI's business group to examine the banking and other industries to estimate the total market for a system such as ERMA. This examination was necessary because ERMA was GE's first venture into the computer field. After the award, SRI provided technical support for ERMA production design at GE until 1957 or 1958. The total cost to the BofA for ERMA was about $10 million, including the cost of many large components billed directly to BofA to avoid markup. SRI billings to GE over the transition totaled less than $1 million.[D] That total included a $5,500 contract awarded to SRI in August 1957 to "dispose of the Mark I ERMA equipment."

The BofA installed the first production ERMA system in its Foxworthy-Plummer Branch in San Jose in January 1959. The units included a sorter/reader, a computer, magnetic tape units, and a high-speed printer. Over the next 2 years, 32 systems were installed, and by 1966, 12 regional ERMA centers served all but 21 of the Bank's 900 branches. The centers then handled more than 33,000 accounts each hour and 750 million pieces of paper each year, about the number they were predicted to handle by 1970. Without such a system as ERMA, the BofA would clearly not have been able to meet the demand for account updates, no matter what assumptions were made about manual check processing.

The SRI ERMA Team

In almost every respect the ERMA project was a team effort, in part because so many different

[5] Fred Kamphoefner, personal communication, June 8, 1999. The BofA at first wanted to preserve its clients' freedom to choose the kinds of checks they used, including one in the shape of a fish. Although the Bank soon started to compromise, the SRI sorting machine was made to handle checks that varied two to one in width, length, and thickness.

[6] SRI conducted two studies for the BofA on the commercialization potential of ERMA. The reports on these studies, submitted on 1 April and 18 July 1955, looked at the market nationwide (estimating a market of about 600), suggested several marketing options, and stated criteria for evaluating the bidding manufacturers.

[7] According to Oliver Whitby (personal communication, August 2000), SRI helped to evaluate the four main bidders: IBM, RCA, GE, and Texas Instruments. SRI and BofA technical people liked TI, but GE's size and attractive financial proposal persuaded the BofA's decision makers. Curiously, GE's upper management was not in favor of entering the computer business and didn't learn of the arrangement until it was time to sign the contract. The proposal had come from GE's Palo Alto laboratory.

critical skills were needed. Morrin had overall project responsibility and interfaced with BofA. Noe, then an assistant director of SRI's Engineering Group, assumed technical leadership of ERMA. Eldredge, Whitby, and Dr. Byron Bennett played supervisory roles. Their collective job was to make the many components of the system "play together."[E] Most project participants shared the feeling that the effort was historic, although no one could foresee the exact shape or pace of computer applications and computer research. The technology was raw and in flux. Design decisions had to be taken with great uncertainty about the future of new hardware devices, and there were no tools for logic design or programming. Project team members had to experiment with new devices and develop their own design techniques, and yet preserve commercial security throughout the project. All these factors made ERMA an exciting experience for young engineers and fueled the effort that was needed for its success.[8]

[8] The principal participants in the ERMA project and their roles were:

Overall Direction	Tom Morrin (SRI) and Clark Beise (BofA)
Technical Direction	Dr. Jerre Noe (SRI) and Charles Conroy (BofA)
Major Supervisors	Drs. Byron Bennett,. Oliver Whitby, and Ken Eldredge (all SRI)
For the Computer	Richard Melville (manager of construction), Milt Adams, Dr. Frank Clelland, Howard Zeidler (all SRI)
Logical Design	Bonnar Cox, Jack Goldberg, Dr. William Kautz (all SRI)
Physical Wiring Design	Roy Amara, George Barnard, Dr. John Blickensderfer (all SRI)
Quality Control	Bruce Clark (SRI)
Other Engineers	John Boysen, Rolfe Folsom, Alfred Fuller, Keith Henderson, Robert Leo, Maurice Mills, Robert Rowe, Dennis Finnigan (all SRI)
For Paper Handling/ Reading	Fred Kamphoefner, Paul Wendt (all SRI)
Mechanical Design	B.J. O'Connor, A.W. Noon (all SRI)
Electronic Design	Mendole Marsh, Phil Merritt, C.M. Steele (all SRI)
Magnetic-Ink Development	Sam Graf (SRI)
Other Engineers/ Technicians	F.C. Bequaret, John Boyson, Tom Drewek, Bernard Elspas, Rolfe Folsom, Al Fuller, Ken Gardiner, Willard Guthoerl, Keith Henderson, Tatsu Hori, A.E. Kaehler, Mitchell Matovich, Ron Presnell, and Benjamin Wolfe (all SRI)

The ERMA Epilogue

ERMA was the first computer used for commercial check processing and the first system in which the actual check was the input medium. Because of SRI's system insight and because the BofA was large enough to establish a convention, banks began, after ERMA's introduction, to preprint individualized checks for each customer and to identify customer accounts by accession number rather than alphabetically by customer name. ERMA gave the BofA a leadership position in the banking industry's use of computers and allowed GE to enter the market for computers for data processing.[9]

A few later developments provide footnotes to ERMA. In a 1960 report to BofA's Board of Directors, Zipf noted that in 2 years ERMA had cut BofA's check processing costs almost in half, a savings of about $6 million a year.[10] The last GE ERMA machine was unplugged in June 1970. Many years later, an ERMA that had been relegated to a warehouse in Livermore, California, was refurbished to operating condition by BofA volunteers and placed on permanent display in Concord, California.

Over time the BofA has collected some of the accolades regarding ERMA. Examples include the following:

- Federal Reserve - "Bank of America established the dominant design of modern electronic banking."
- Professor James L. McKenney, Harvard University - "The development of the MICR line, which enable checks to be sorted and processed at high speeds, has been recognized as one of the great breakthroughs in banking."
- *American Banker*, 1984 - "I don't know how we would have gotten on without MICR. It was one of the enormous leaps in technology that preserved banking."

[9] According to Dan Evans, BofA, (series of articles in BofA's *Technology World*, issues 156-159, June-September 2000), two other major banks were looking into automation. Chase Manhattan was working with MIT, and First National City Bank had a partnership with ITT Labs in Belgium. Neither of these materialized. Also, unknown to the BofA until after the fact, the bid from GE was unsanctioned by GE's leadership, who wanted no part of the data processing industry. When GE won, management reluctantly agreed to support the BofA contract. The company did abandon the computer field 14 years later, when the relevant division of GE was sold to Honeywell.
[10] Dan Evans, op. cit.

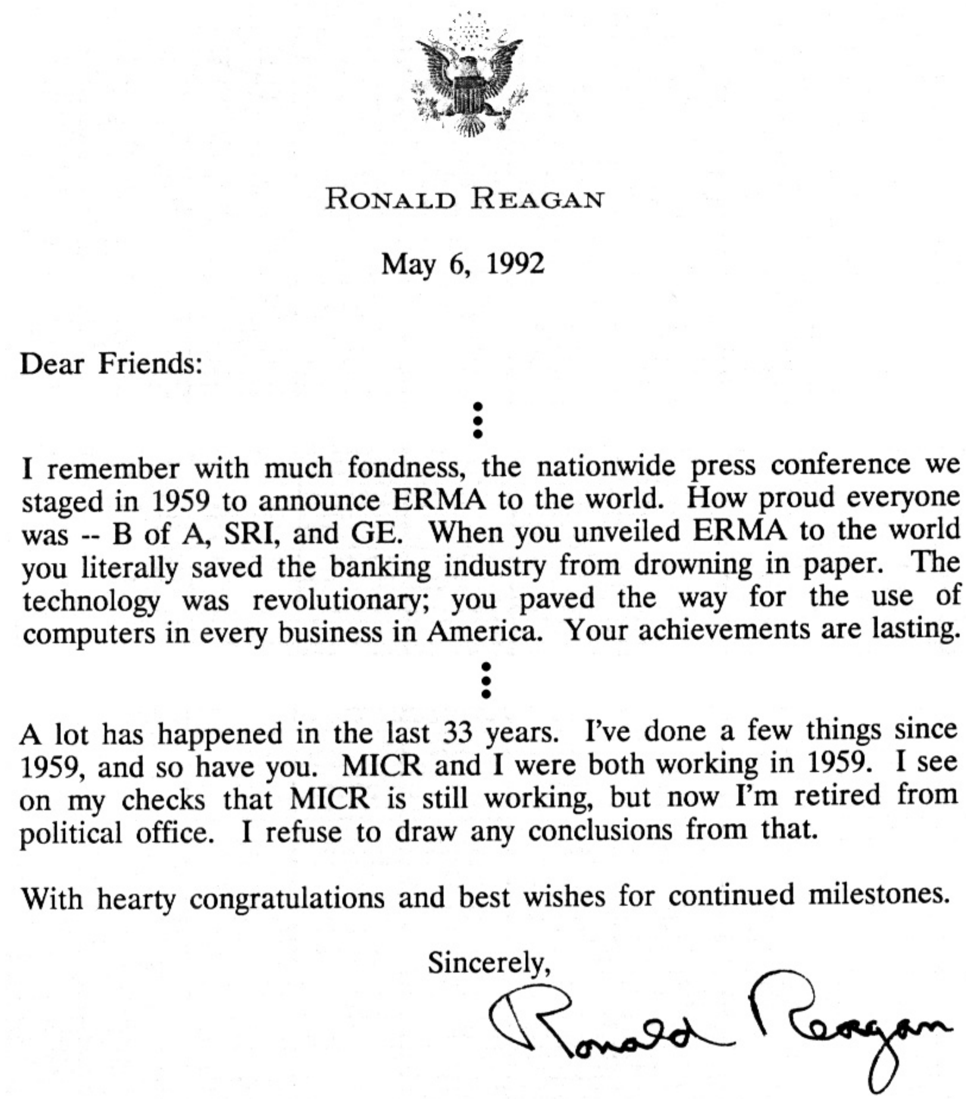

RONALD REAGAN

May 6, 1992

Dear Friends:

⋮

I remember with much fondness, the nationwide press conference we staged in 1959 to announce ERMA to the world. How proud everyone was -- B of A, SRI, and GE. When you unveiled ERMA to the world you literally saved the banking industry from drowning in paper. The technology was revolutionary; you paved the way for the use of computers in every business in America. Your achievements are lasting.

⋮

A lot has happened in the last 33 years. I've done a few things since 1959, and so have you. MICR and I were both working in 1959. I see on my checks that MICR is still working, but now I'm retired from political office. I refuse to draw any conclusions from that.

With hearty congratulations and best wishes for continued milestones.

Sincerely,

Ronald Reagan

Figure 2-5. Excerpt of a letter from President Reagan on the occasion of the thirtieth anniversary of ERMA going online.

- President Ronald Reagan, who knew of ERMA from his days with GE - see the letter in Figure 2-5.

Beise became an SRI Board member in 1959[11], and SRI went on to help the Federal Reserve System specify, purchase, install, and verify ERMA-like machines for its huge check processing needs. Bonnar "Bart" Cox was instrumental in this assistance, and by 1959 he was leading the testing of the first installed Federal Reserve systems for accuracy and capacity. The check reading and processing systems developed under ERMA, by then available in the marketplace, were up to the task. Finally, at a celebration in March 2001 honoring the SRI ERMA team with SRI's Gibson Achievement Award, a representative of the BofA, retired Senior Vice President Duncan Knowles, revealed an important fact. The power of machine processing, specifically as evidenced by ERMA, gave the BofA the encouragement it needed in 1958 to start the first nationally accepted credit card system, BankAmericard. Later renamed Visa, it is used worldwide, with over 1 billion cards issued.

[11] Clark Beise was one of several SRI Board members who deliberated over the terms of the SRI separation from Stanford in 1970.

Automatic Handwriting Recognition—The Signature Pen

The genesis of SRI's so-called signature pen occurred in the fertile mind of Hew Crane in the waning days of the ERMA project in the late 1950s. Knowing that most of the important information on an individual check is handwritten, Crane thought there must be a way to automatically encode that information so that computers could understand it. He also considered the possibility that the idiosyncrasies of handwriting that made it difficult for a machine to read it might give computers a means to identify who was writing a particular check. Crane also reviewed the issue of whether the pen or the writing surface should contain the instrumentation needed to gather the data from which information could be processed. Nearly 50 years later, the automatic recognition of handwriting has traversed both of these approaches but, with the appearance of small personal digital assistants (PDAs), it is the pad and not the pen that has become the critical information-gathering component. We will briefly see how that course transpired at SRI.[F]

It was the late 1950s and though a rudimentary digitizing surface existed (the RAND Tablet), Crane first concentrated on the pen. This choice left the writing surface quite arbitrary and so it could include checks or forms. SRI's first demonstration system was centered on a direction-sensitive shaft or pen, programmed for the recognition of the digits 0 through 9.[G] It was built with the help of George Eilers and it simply sensed four cardinal directions used in the sequence of writing a digit.[12] As long as the pen was held in the prescribed orientation and the same sequence was used, it did a reasonable job. Those stipulations, however, proved to be unacceptable limitations.

In the context of check writing, Crane also saw the value in being able to record the dynamics of the pen's use to recognize not just what was written, but to identify the pen's user as well. While the potential interest centered on checks, barely over the horizon lay an even more important signature verification need, the ubiquitous credit card slip. But it was still early, not just in the growth of credit card use but also in the development of small powerful, but more importantly, affordable processors. Those reasons, plus the lack of a sponsor to help perfect the pen, would relegate its exploration to the shelf for over a decade.

In the early 1970s, a large supplier of paper forms began to worry about the ushering in of the "paperless society." Having seen an IRE paper by Crane, they called to learn the status of the so-called signature pen. The company was interested in online form input, and while that contract didn't result in anything significant, it did rekindle interest at SRI in technology of using handwriting as a means of computer input.

By the mid-1970s SRI had built a new version of the pen that used an array of strain gauges to measure the writing force in three dimensions. Xebec Systems sponsored an exploration of the pen for recognition of all alphanumeric characters. But, again, the design was impractical for recognizing characters because it imposed too many restrictions on the user in terms of writing directions (e.g., left-handers).[13] However, that approach did enable the measurement of how the pen was held and so this second version of the pen shifted toward its potential as a signature verification system.[14] By then, the burgeoning use of credit cards was expected to provide an excellent reason for such verification and it was easy to imagine using such a pen at every point-of-sale (POS) terminal. VISA and others became clients about this time but again the SRI pen effort would be stalled by the economics of the credit card world. One reason was that the generalized cost of the signature verification, including false rejections, was too high in light of the vast majority of credit card drafts being of low value. It also turned out that credit card forgery was but a small increment on the marginal costs.[H]

---

[12] SRI received a patent in 1964 for the character recognition circuit used in this system. A number of other pen patents would follow in 1975, 1977, and 1978 as the detection and recognition systems advanced.

[13] One could imagine a recognition system that learned a given users writing style but processor and memory costs were so high at that time that the per user equipment cost would have been prohibitive.

[14] Besides Crane, others working on this improved version of the signature pen were Dan Wolf and John Ostrem.

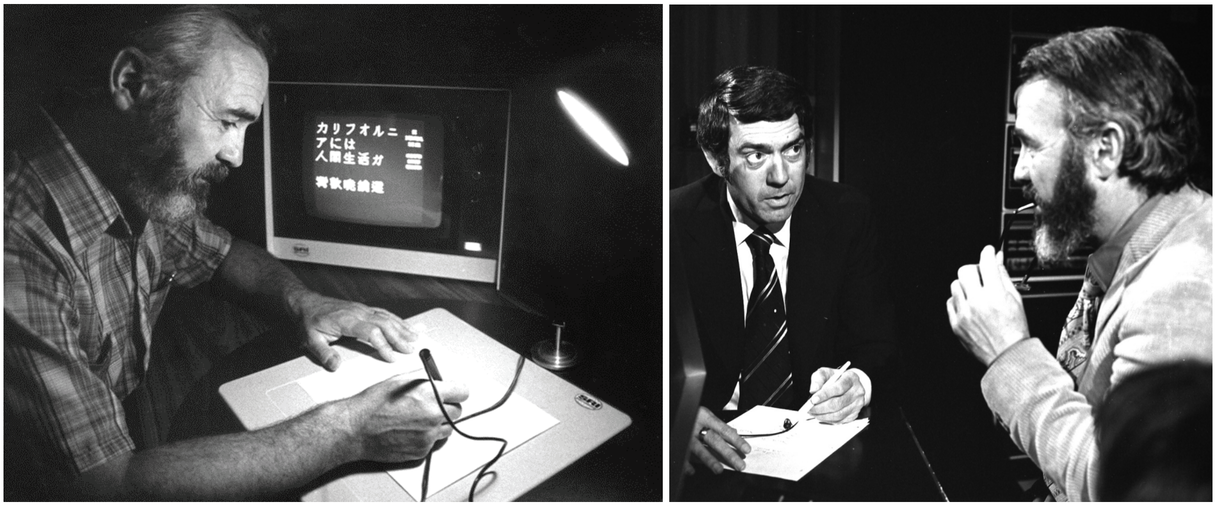

Figure 2-6. Hew Crane using the kanji handwriter recognition pen and (right) demonstrating it for CBS's Dan Rather in April 1976.

But another area of potential use arose, driven by the influx of computers into Asia. Their written language, Kanji, consisted of a huge number of characters that was next to impossible to represent on a normal-sized keyboard. A handwriting recognition system thus seemed a promising opportunity for the emerging East Asian market.

According to accounts by Crane and Earle Jones, an SRI division director, Crane's team obtained a combination of internal SRI support and project funding from the Taiwan Ministry of Telecommunications to build a trial system that could recognize a few hundred Kanji characters. This time the device was primarily a tablet digitizer and they called it Handwriter (see Figure 2-6).[15] Japan was chosen as the target market because of its use of Kanji and its phonetic kana alphabet. Through a chance meeting between Don Scheuch, then Senior Vice President of Engineering, and a business acquaintance, an experienced Chinese-American entrepreneur named James Dao was identified and offered the chance to form and lead a new company. After considering the intellectual property and traveling to Japan to evaluate interest there, he agreed to license the Handwriter technology in exchange for 25% of the initial stock of the company. The company, called Communications Intelligence Corporation (CIC), was formed in 1981 with Lester Hogan (of Motorola and Fairchild Semiconductor) as Chairman of the Board.

About $2 million was raised from individual investors and venture capitalists in the Bay Area, and once the new company had some money, several SRI staff members left to join CIC.[16] Contracts were to be sought from large Japanese corporations, and Jones was selected to join CIC and seek orders from Japanese companies such as Seiko, Mitsubishi, and C. Itoh. Part of this effort also involved CIC's manufacture for Apple of the SRI/CIC Handwriter, calling it, naturally, MacHandwriter. It was sold only in Japan but, surprisingly, the Japanese market for handwritten Kanji input simply did not develop.

At CIC development work on more general handwriting recognitions systems continued throughout the 1980s with ultimately over $10 million in investment capital being raised. A shakeout of handwriting recognition companies in the early 1990s initially weakened CIC and it filed for Chapter 11 in 1994. But its new leadership downsized the company, chose to abandon hardware development, concentrated on software for the PDA market, and eventually created some of the industry's best handwriting recognition software. As increasingly capable PDAs began to appear, CIC was able to secure licenses for its Jot™ recognizer from Microsoft in 1999, for use with its small computer operating system, Windows CE, and from Palm for the Palm Pilot. Jot is still licensed by Palm including its use in PalmOS 5.2 and in a new PDA-phone combination called pdQ.

[15] The wire shown in the figure is the lead to a pressure sensor inside the stylus.

[16] Included were John Ostrem, Peter Edberg, and David Foyt with Hew Crane participating part time.

CIC has become a successful company and over the years has licensed its handwriting recognizer to over a dozen companies including Microsoft, Casio, Hitachi, Handspring, Palm, HP, IBM, Compaq, Nortel, Ericsson, Mitsubishi, Fujitsu, NEC, Symbol Technology, and Telxon. Not to forget signature verification, earlier in May 1998, CIC also licensed its Sign-it™ signature verification software to Microsoft for use in Word 97, to Adobe for Acrobat in 1999, and in the fall of 1999 shipped the first biometric signature verification system for Palm organizers. Since President Clinton signed e-signature legislation in 2000, CIC has become an industry leader in handwritten e-signatures.

SRI's road to success for handwriting input and signature verification was a bit long and troubled but ultimately worth it. SRI sold its shares of CIC in 1992 and 1996 for a combined value of just over $4 million, some of which went toward defraying expenses when CIC was housed on the SRI campus.

The Origins of Personal Computing

Though he left SRI not quite 30 years ago, Douglas C. Engelbart is, almost without question, still SRI's most famous talent. His vision, played out over his nearly 2 decades at SRI, helped revolutionize the way people viewed the purpose and utility of an emerging new class of machines called computers. His stay at SRI was both enabling and at times frustrating to him. In both interviews and various books written about him and his contributions, he has spoken openly of both those conditions. In this account, you will find an SRI perspective of his story, including some detailed insight on those conditions. Doing so will require more detail than other projects described in this book, but I will endeavor to tell it fairly and accurately, including those frustrations he faced and his sometime awkward relations with his managers.[17] By concentrating on his SRI world, some of the wondrous and ubiquitous consequences of his work will be left to others who are relating his impacts on the world.[I] Just to be clear, I confess an enormous respect for how he changed the face of computing.

The above "section" title would almost certainly not be Engelbart's choice, since it reflects only part of and probably even a diversion from what he was seeking during his years at SRI. The tremendous changes that personal computing has brought seem to blind us to a greater good that he sees. He believed that personal computers, to the extent they remained "personal," were contrary to what he was trying, above all else, to show: that the computer could provide the basis for a set of tools that would raise *organizational,* not just personal, effectiveness.[J] Though what he and his laboratory showed the world were the first truly comprehensive and intimate ways that computers could serve their users, they were but a way station on the road to Engelbart's vision. Engelbart saw how computers could not just aid individuals, but in fact alter or, he hoped, profoundly alter the way we work together. That was his quest when he began

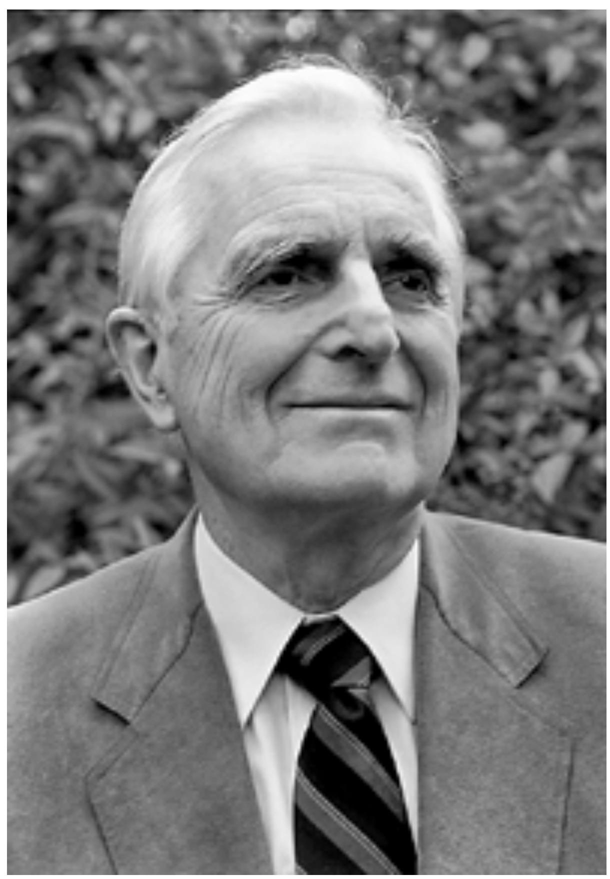

Figure 2-7. Dr. Douglas Engelbart

[17] So as not to detract too much from the evolution of Engelbart's ideas and their outcome, his interactions with SRI management, including the ultimate demise of his SRI group, is detailed in Appendix G.

almost 50 years ago and it remains his quest today. But as it has turned out, the responsiveness he and his people designed and demanded from a computer could not help but foster the power of what personal computing meant, no matter how subordinate or incidental that power was to his vision. Their innovations were so commodious that they have benefited both individuals and in most cases the groups to which they belong. Engelbart and too few of his SRI associates are now belatedly receiving the acclaim they deserve.

The Genesis

Engelbart (see Figure 2-7) started his professional life in a fairly ordinary way. He graduated as a good student from Oregon State University in 1948 with a degree in electrical engineering and moved directly into the application of his acquired skills by delving into a wide range of electrical and electronic problems for the National Aeronautics and Space Administration (NASA) Ames Research Center (then NACA) in Mountain View, California. He was to get married while working there and unfolding before him was the conventional path of job security and family. But even at 25 Engelbart soon came to burn inside for some larger contribution he might make to society. Through deliberate effort, he landed on the notion that people needed some way to deal more effectively with the increasing complexity in their lives. With that general goal and in a stroke that can only be called prescient, he imagined that we could individually and collectively enhance our intellectual pursuits and solve more complex problems through a continuous, almost intimate interaction with computers.

Never mind that computers were hardly in existence in 1950. Even those creators trying to bring them into being, saw them as enlarged calculators or mathematical problem solvers to which were brought periodically very specific and well-formulated programs. Engelbart's vision revealed them as few others' did, as the means to facilitate our thinking in a responsive and interactive way. While he was aware of the very few others in the country that were probing, philosophically, the nature of human-computer interaction, his take on the roles of the user and the machine were his own. In retrospect his vision was truly a defining concept for the information age, which was still nowhere on the horizon. He would come to learn just how unorthodox his concepts were.

If Engelbart had an epiphany on that December day in 1950 when he began to see how computers could give wings to his dream, it was influenced greatly by his earlier absorption of a 1945 article by Dr. Vannevar Bush, the preeminent Massachusetts Institute of Technology (MIT) scientist.[K] Bush had a vision of a means by which information could be organized associatively to make more naturally retrievable the rapidly expanding body of man's knowledge. Engelbart found Bush's concepts appealing and consistent with what he was thinking. By 1951, he could also see a way to get there.[L]

His now focused journey began with Engelbart quitting his job at NASA in 1951 and going to the University of California, Berkeley, largely because a computer was under construction there. At first he was a part of the research staff, but soon realized that he needed to again enroll in classes to better his understanding of what was developing. By 1955 he had received his doctorate, and the Berkeley computer was still not running. Engelbart stayed on as an assistant professor until 1956 and then had a brief fling at starting a new company to produce some of the digital devices he had developed at the University. He became President and Director of Research at Digital Techniques. But the company never got off the ground. Its failure was due in part, ironically, to an investigative report that a potential investor, at the recommendation of Engelbart, had commissioned at SRI. The plasma devices that Engelbart was pursuing, the report said, would one day be eclipsed by semiconductor devices. Though Engelbart wasn't necessarily in disagreement with that outcome, he was trying to use his Berkeley device patents and the company to ultimately finance his augmentation dream. This turn of events just meant he had to find another avenue to get there.[M]

Engelbart and SRI

Engelbart was not unfamiliar with SRI. He had heard about a computer that had been built at SRI for the banking industry. Jerre

Noe, the division head of the group to which Engelbart would ultimately come, recalls that he met Engelbart at a lecture he gave in Berkeley, while Engelbart was still there. But according to Engelbart's own recollection, he saw at SRI the freedom to define a goal and pursue the means to reach it.[N] So, he called on SRI and then waited for about 3 months before receiving an offer. The group he joined in September 1957 was, interestingly, involved in the aftermath of the huge ERMA banking computer project described earlier. That is what had earlier called his attention to SRI, but by the time Engelbart actually joined the group, it was not at all certain what lay ahead for its members.[18]

Characteristically, Engelbart dived into his new job with verve and creativity. He began designing electronic components, and his work later resulted in SRI receiving a series of patents in the areas of magnetic cores and electronic discharge devices. In 1958, just a year after coming to SRI, he was promoted by his boss, Reid Anderson, to a Senior Research status in recognition of his contributions.[O] Perhaps his earliest evidence of information retrieval at SRI appeared in the *SRI Journal*.[P] With colleague Charles Bourne he wrote about the challenges presented by the deluge of technical information being published. The article was not very revealing of his vision, but it asked a lot of questions, in fact, no fewer than 123 questions, many thoughtful, without suggesting a single answer.

In a series of internal SRI memoranda from January to December of 1960 we can see the vocabulary and consequences of Engelbart's vision unfold. He was exploring his emerging vision in a number of petitions to his management for investment money. His first request, on January 21, 1960, had a telling theme: "The dynamic utilization of automatic information-handling equipment for everyday personal use." He spent a bit over $10,000 on that first exploration of issues surrounding personal rather than group use. In May 1960, however, he switched directions a bit and suggested that SRI become involved in teaching machines. In July he authored a request that zeroed in on his pet project, "seeking to improve the intellectual capability of humans." It was a request to develop a Stanford seminar, in collaboration with staff members in the SRI Economics Division, and a request for $1,500 in capital funds.

In September 1960 Engelbart offered a long treatise on the notion of teaching machines, directed at building physical skills through sets of cognitive linkages. On December 5 he wrote an informal summary of his earlier writings on "Augmented Human Intellect." This summary became Exhibit A in an otherwise short proposal to the Air Force Office of Scientific Research (AFOSR) requesting $27,000 to match equal SRI funding for an exploration of his ideas.[Q] The proposal was entitled, "Augmented Human Intellect Study." The seemingly futuristic title was not particularly unusual. For example, J.C.R. Licklider, who was recruited to head an information technology office at DoD's relatively new Advanced Research Projects Agency (ARPA), in 1962 launched the long-running ARPA Project MAC, which stood for machine-aided cognition. Licklider had been writing about man-computer symbiosis in 1960. Another example was the theme of the 1961 Joint Computer Conference in Los Angeles: "Extending Man's Intellect." Yet in the AFOSR proposal, the difficulty even Engelbart had in describing his research aims is evident in this vague excerpt from the proposal's Brief:

> The vocabulary doesn't seem to exist with which to communicate briefly the subject range involved in this multidisciplinary study. The best summary of the proposed research would seem to be a statement that we are going to try to determine what the range of subject matter must be within which to give rational, overall consideration to the problem of making humans more effective as problem solvers, and to try to provide perspective for a coordinated attack upon this problem.

In any case SRI was awarded a grant. As far as I can tell, this gave Engelbart his first external money to refine his vision unbridled; to try to express for others the passion he had inside.[19] A condition of the grant was that SRI

[18] While that uncertainty may sound ominous, it is commonplace at a research institute where securing new projects is part of life. In this case, however, the end of a large, long-term, and ground-breaking project presented an unusual challenge.

[19] Research in Individual Information Handling Problems, SRI Project 3578, Contract No. AF 49(638)-1024, March 1 1961 to March 31 1962, $26,924. It is curious that 1

had to match the amount with its own internal funds. So from 1960 through 1962, SRI spent discretionary investment money of about $60,000 to meet that condition and to pay for the writing of the requisite proposals.

By fall 1962 Engelbart had finished the project and written the now famous AFOSR report that offered his first augmentation concepts and described the role the computer could play in them. While not a blueprint for the future, it was a crown of conceptualization. Here are some of the early concepts:

- For the most part our intellectual effectiveness is not limited by our intelligence.
- Humans, together with personalized artifacts such as computers and their attendant information handling displays and storage, form a powerful system for the solving of complex problems.
- In various ways our intellectual activity can be directly aided by the capabilities and offerings of a digital computer, which can provide:
  - An ability to supplement human memory with that of a computer
  - A work environment with visual access to that memory as well as other libraries and processes, forming whole new kinds of encyclopedias
  - Associative linking possibilities between processes and documents
  - Devices that enable the user to interact effectively with a displaying computer for document generation and modification faster than a typewriter can achieve.
- Problems are more easily solved if they are decomposed into a hierarchy in which each layer has a defined capability.

month before this contract was awarded to Engelbart, SRI was awarded $30,634 for another AFOSR project entitled "Structuring Complex Man-Machine Systems" (SRI Project 3546). The AFOSR contracting agent (not technical monitor) for both projects was Rowena Swanson. In spite of the connotations its title implies today, SRI Project 3546 dealt only with systems analysis, not man-machine interaction.

- A systems engineering approach is appropriate for this human-machine and human-team integration.
- This augmentation of intellect does not have to await an understanding of mental processes; we can begin at once, adapting and building as we go.

Perhaps it was his familiarity with radar screens from his Navy days that helped Doug see a medium for this new human-computer interaction, and he imagined that computers would one day have enough power to aid an intellect. In the report he openly acknowledges the influence of and insights from the earlier-mentioned article by Vannevar Bush, "As We May Think."[R] What Engelbart saw that Bush could not see was the potential of the computer as the means to deliver what Doug came to think of as supplementation of the intellect: ways that humans, quite generally, could rise above their intrinsic intellectual and collaborative limitations.

In June 1961 Engelbart had sought and received an additional $27,000 from AFOSR for a second year of work.[20] As his ideas were maturing, it is instructive to look at how he expressed his continuing need for SRI to fund his concept development. His 1962 petition for Internal Research and Development (IR&D) funds was titled "Program on Human Effectiveness." It "proposed research…aimed at refining and developing in greater detail the initial conceptual framework for augmenting the human intellect." This proposal expressed his vision in the most general terms. For 1965, his last year of SRI IR&D support (at least for a while), he received $17,468 for something much more specific: "To develop and evaluate an experimental computer-aided text-manipulation system suitable for a wide range of applications…." This specificity was no doubt a result of the crystallization of at least one tangible means to achieve the goals he sought.

Engelbart used the AFOSR report to introduce his ideas to other potential funding agencies and was fairly successful at ARPA and NASA. But while he was receiving both internal SRI and client support during the early 1960s, he still didn't have a computer.

[20] This amount, with the SRI matching funds, would cover between one and two people for the year.

The first money from ARPA in 1963 wasn't enough to buy one and still have enough left over to do appreciable work. At ARPA Licklider insisted that Engelbart link up to a timeshare machine at the Systems Development Corporation (SDC) in Los Angeles via what Licklider wanted to call the California Net.[21] Though that option was limiting and frustrating to Doug, he had little choice and therefore needed at least a minicomputer in Menlo Park to make the connection and the associated interaction experiments run as best they could. According to Dave Brown, a friend and acquaintance of Doug's at Berkeley and later SRI, Engelbart had become a friend of Finley Carter, SRI's president at the time. Brown relates that Doug used that access to Carter to obtain a computer funded by SRI capital investment. Unfortunately, Doug doesn't recall this happening.[22] By whatever way it was acquired, the first small computer he used was a $100,000 CDC 160A. Early proposals had requested money to lease that specific machine, and project files show its costs were shared among the active projects using it. Though shared, it at least enabled Engelbart to begin experimenting on the human-machine interaction methods he wanted to create.

Over those formative years of 1960–65 and in addition to whatever external money was forthcoming, SRI invested about $120,000 of IR&D funds in Engelbart's ideas. A portion of this money may have gone toward the $54,000 in required AFOSR matching funds. Noe, Engelbart's SRI division director in Doug's early years at SRI, commented recently that he personally believed in the future Engelbart espoused and therefore approved his requests for internal funds. In any case, Doug received enough investment money from SRI (1960–65) and in contracts from Harold Wooster at AFOSR (1961–62), Licklider at ARPA (early 1963)[S], and Bob Taylor at NASA (in 1964)[T] to eventually win a place at the revered table of ARPA contractors funded by the Information Processing Techniques Office (IPTO). That position meant years of continued support as ARPA moved aggressively into this new field of computing.[23]

This aggregation of support should not be misinterpreted, however, for it took an enormous amount of proselytizing and selling by Engelbart to make all this happen. But it was a propitious time at ARPA. It was beginning its revolutionary venture into computer sharing and networking. For more than a decade the internal and external funding was sufficient for Doug to pursue his vision and to increase his laboratory to a truly talented team of over 40 people. The Human Effectiveness Program was renamed the Augmented Human Intellect Research Center in 1969, which became the Augmentation Research Center (ARC) by 1971.

Though SRI's internal support was an important enabler to this growth, during 1963–64 Engelbart encountered what he saw as resistance from his manager within SRI. From the comments of one of his contemporaries, it seems that his managers, and to some extent his clients, understood what he was trying to do only in some general sense. This lack of precision was likely due to some combination of their lack of vision and Doug's inability to explain what he was trying to accomplish (see Appendix G).[U, 24] Though in Engelbart's

[21] While that moniker of "net" today sounds pedestrian and may imply a switching fabric, it was, in fact, nothing of the sort—just a long distance link to a timeshare host in El Segundo. Also, according to SRI's Dave Brown, Engelbart's trip to Project MAC let Engelbart understand the importance of having one's own machine. From letters and visits between DARPA and SRI in the February–March 1963 time frame, Doug's manager, Roy Amara, and John Wensley who Amara had assigned to work with Doug, both were responding to Licklider's remote connection preference rather than supporting Doug's need for a dedicated machine.

[22] Douglas Engelbart, personal communication, February 22, 2001. On the other hand, a conversation with SRI's Bud Rorden on July 13, 2004 revealed that he had used the 160A for a totally different project dealing with satellite data. If this had been a dedicated project-provided machine, that kind of sharing wouldn't have occurred.

[23] Engelbart had tried to obtain funding from NASA, National Institutes of Health (NIH), ARPA, and other agencies before MIT psychologist Licklider, who appeared at ARPA in October 1962, the same month that the AFOSR concept report was published, would fund Engelbart beyond just paper studies. Licklider formed the new office in information processing called IPTO. But it was Bob Taylor, who had funded Engelbart earlier at NASA and who came to IPTO under its next director, Ivan Sutherland, who would buy Engelbart and his group their first state-of-the-art, dedicated timesharing computer in 1966, an SDS-940. Sutherland became head of IPTO in 1964 and Taylor in 1966.

[24] In fairness, the notion of augmenting knowledge work is intrinsically vague. But as a 1963 ARPA contract was winding down and SRI was seeking an extension, Licklider, in a March 1964 letter to Doug's supervisor,

recollection this was a trying time, it didn't seem to appreciably affect his internal financial support or his eventual ability to grow to laboratory size. Not until the mid 1970s, when the external support began to erode seriously, did SRI management again choose to intervene (more on this later).

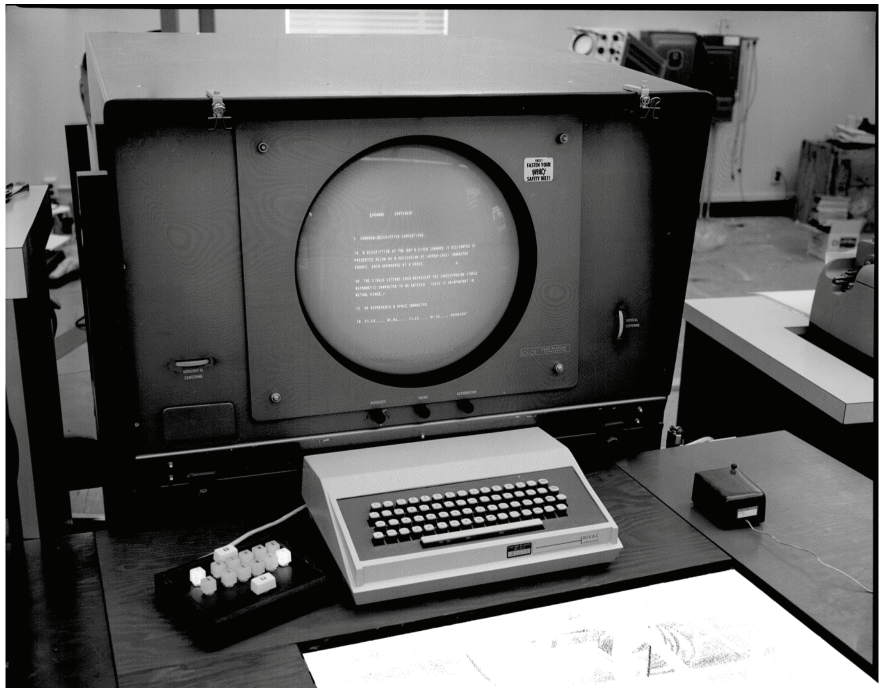

Figure 2-8. A June 1965 version of an SRI workstation.

The Evolution of the NLS Workstation[25]

Though the software underlying a computing system carries most of its critical functionality, perhaps the clearest embodiment of the course being pursued by Engelbart and his laboratory lay in the NLS workstation. How best could the promise of real-time interactive computing be realized? Uncommon in this use at the time, a cathode-ray screen was clearly a good way to present dynamic information but the burning question was how to interact with it. Let's look briefly at the evolution of the workstation.

Although ARPA was the first to provide serious funding for workstation implementation in early 1963, it was NASA money a little over a year later that helped provide a critical mass. For some time those two contracts plus the ongoing AFOSR grant provided the resources to explore these first facets of real-time human-computer interaction.[26] Much of the first ARPA money had gone to modifying the CDC 160A, both to provide interaction tools such as a cathode ray tube (CRT) display and to manage the link to SDC. But the goal of improving user productivity, at least in the beginning, didn't actually require the link and its distant resource. Because Engelbart always wanted his own timeshare machine, the link to SDC was always a costly inconvenience and distraction for him. While not large, the 160A would be the first platform for exploring a more intimate interaction. The final report on the first NASA project reveals the state of development at that time and here are some reflections from it.[V]

The major workstation implementation goal surrounded the question of how best to view the information residing within or accessible by a computer and how such information (in the early stages just text) could be generated and manipulated. While the CRT was the early and probably only

Roy Amara, made it clear that more money was conditioned upon SRI being able to demonstrate what was happening. Engelbart's proposal for that extension contains both vague phrases like "bootstrapping my own work" and "computer-aided work" and specific concepts probably not then adequately appreciated. Where specific, it mentioned aspects of text editing in a 10-page document (the first widely used, interactive text editor, EMACS, was still more than a decade away) and "associative linking" whose enormous significance probably was not yet recognized outside Doug's fertile mind. Also, at this point, the interactive tools were not yet settled on and there was no mention of the mouse as a pointing device. All was in flux and concepts common today, not yet formulated.

25 NLS was the software that gave life to functionality in the ARC. Since the thesis of the work was online interaction, as opposed to the dominant practice of off-line or batch computing, NLS was simply an online System.

26 This claim, of course, applies to that type of interaction that has the general utility we now associate with virtually all computers. Another, detailed and chronologically oriented account can be found in a book reporting on a 1978 conference held in Palo Alto on personal workstations: Adele Goldberg, ed., *A History of Personal Workstations*, ACM Press and Addison-Wesley, 1988.

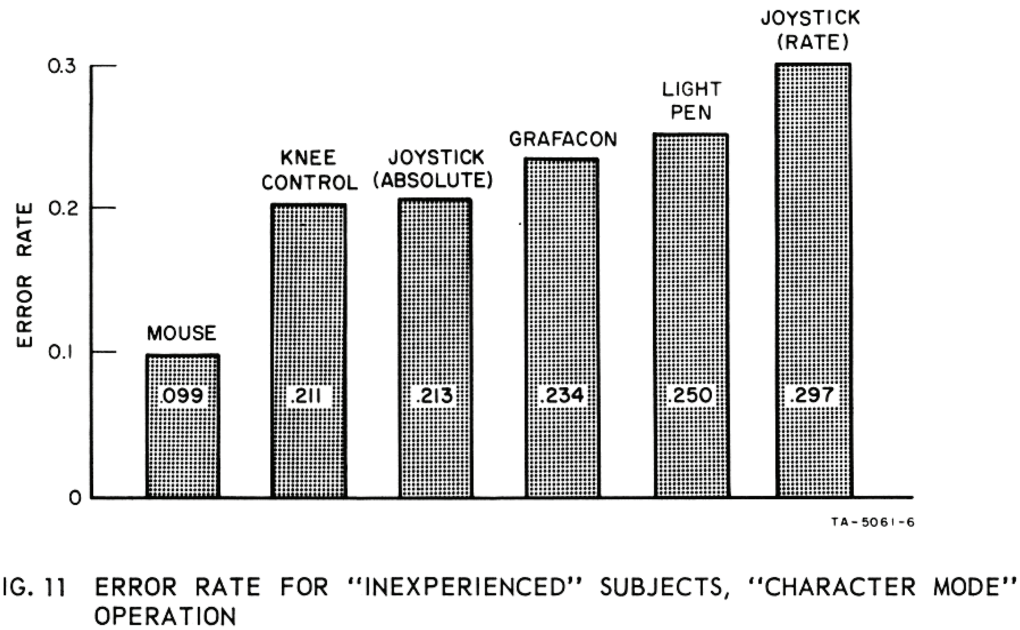

Figure 2-9. Results of one phase of the testing of pointing systems for a computer display.[v]

possible choice for a display, adapting it to textual material with comfortable viewing conditions took some engineering. The tube shown in Figure 2-8 as part of a 1965 workstation gives some idea of the cumbersome state of displays for this kind of environment. The working part of the display consisted of 16 lines of 63 characters when it was under normal use, but shifted to 13 lines of 40 characters each when taking photographs or movies.

The mouse at the right was manifested on the screen as a small, character-sized "+" that is not visible in the display in Figure 2-8. Explaining the rest of the workstation shown is best done through an excerpt from the report:

> 2b4 Within comfortable reach of the user's right hand is a device called the "mouse," which we developed for evaluation (along with others, such as a light pen, Grafacon, joystick, etc.) as a means for selecting those displayed text entities upon which the commands are to operate.
>
> 2b4a As the mouse is moved over the surface of the table, its position is constantly being monitored by the computer, which displays a special tracking cross, which we call the "bug," on the screen in a position corresponding to that of the mouse on the table.

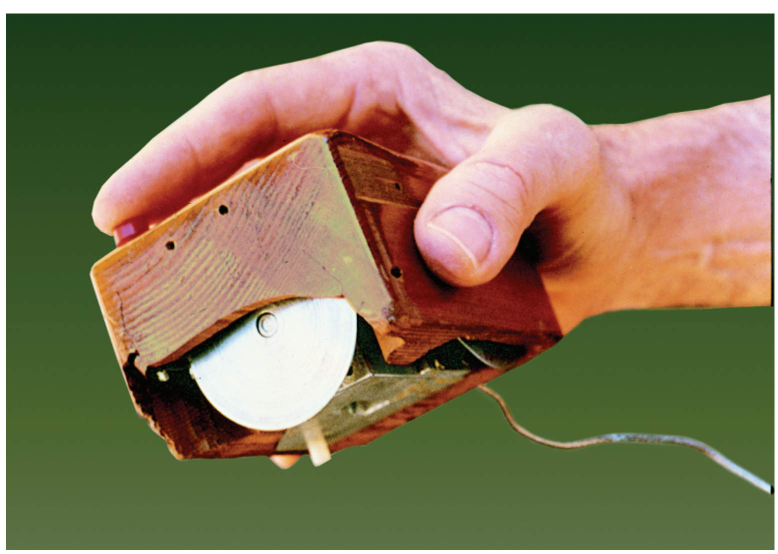

Figure 2-10. Engelbart holding the first working version of the mouse.

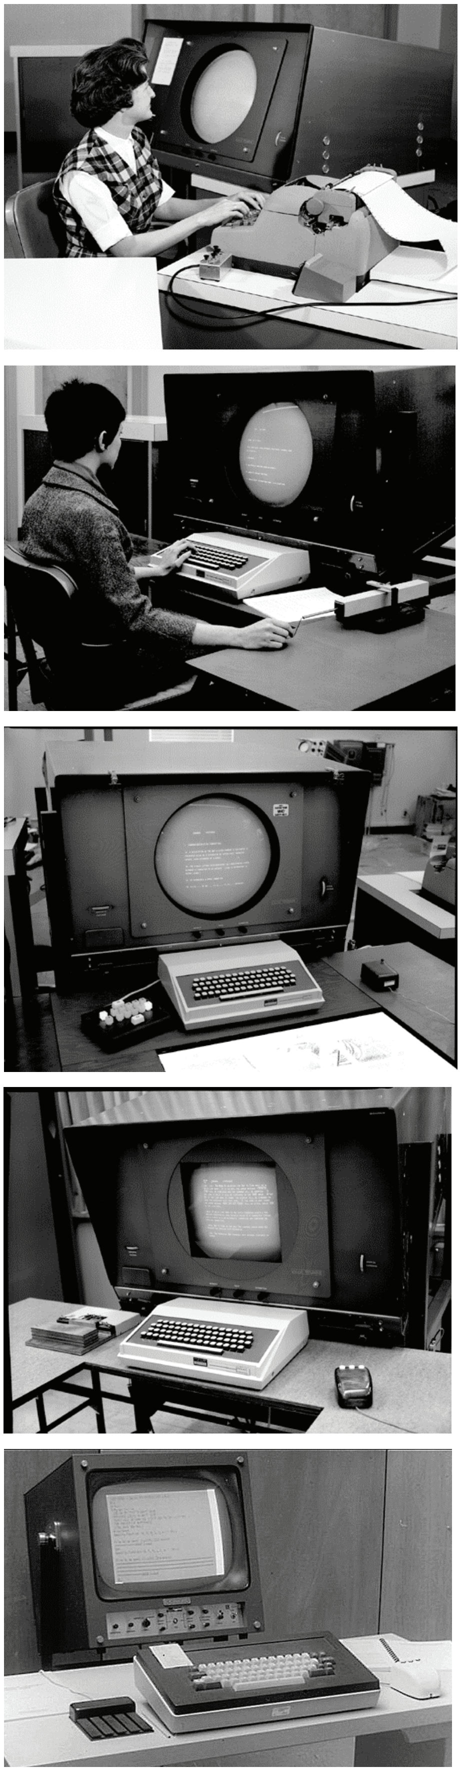

March 1964

March 1965

June 1965

March 1967

1968

Figure 2-11. Evolution of the ARC Workstation during the 1960s

> 2b4b A user soon finds it very easy to keep his eyes on the screen and cause the bug to move about upon it as quickly and naturally as if he were pointing his finger (but with less fatigue).

The paragraph headings used in the above quotation are direct evidence of the group's use of bootstrapping. The entire report was composed, edited, and printed using the workstation described, and the headings were part of the accounting system used in the text editing system built into NLS. That the mouse was selected over the other pointing systems mentioned is evident in just one of a number of results in the NASA report shown in Figure 2-9. The mouse was not only used for selection of text in word processing but also for a variety of other functions on the screen. The first working version of the mouse is shown in Figure 2-10.

Workstation command or "operator" selection could be performed in one of two ways. The keyboard could be used to enter one of a set of commands relevant to the task immediately at hand. The 14-button control box on the left was a shortcut for the most frequently used operators or editing commands. This device was later replaced by the five-finger, chordal keyset that was more flexible but also required a steeper learning curve. Its five keys could, using straight binary representation, input any of the 32 characters in the standard Teletype alphabet. While the keyset never caught on, Engelbart still swears by it and uses it to this day. Many other devices had been tried to help perfect the interaction, but in the ARC it ultimately came down to the standard keyboard, the keyset, and the mouse. The evolution of the workstation at SRI from 1964 to its first public debut in 1968 is shown in Figure 2-11. In Engelbart's mind the concepts behind this development involved a partnership between the tools created for augmentation and the newly acquired skills of the user. The new tools were not intended to ease the user into decreasing states of cognitive involvement. Bootstrapping was consistent with this partnership.

As mentioned at the outset, this second partner, the user, was expected to evolve too. Engelbart used the diagram in Figure 2-12 to explain this partnership during a 1986 talk.[W] Notice that the left-hand side is not just a

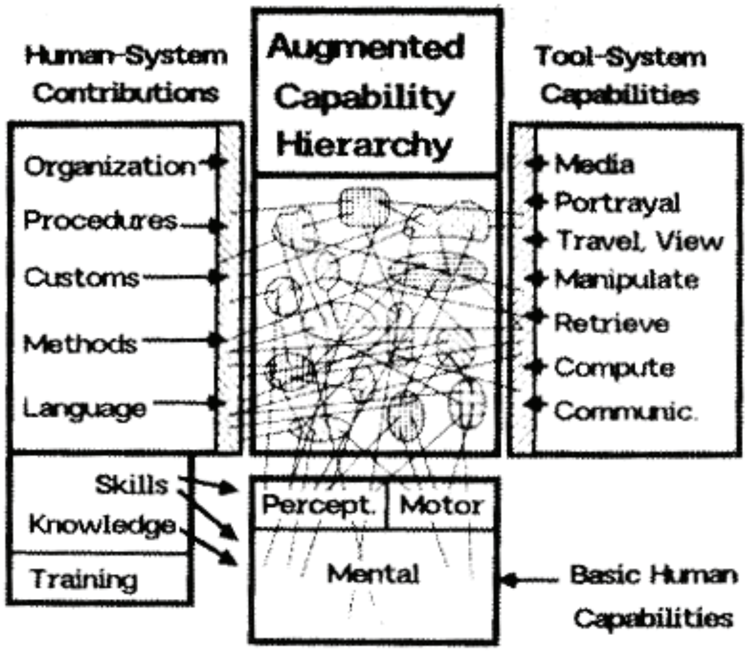

Figure 2-12. Engelbart's conceptual diagram of the roles of the human and his or her machine.

single "user" but a more generalized, human system with its cultural underpinnings. The middle or interaction space is where the consequences of the partnership arise. The failure of the human side to develop as quickly as the tool side gave Engelbart considerable concern. By the time of this talk he had begrudgingly accepted that desirable changes in the human side might take generations, and that timely co-evolution of the two would be unlikely.

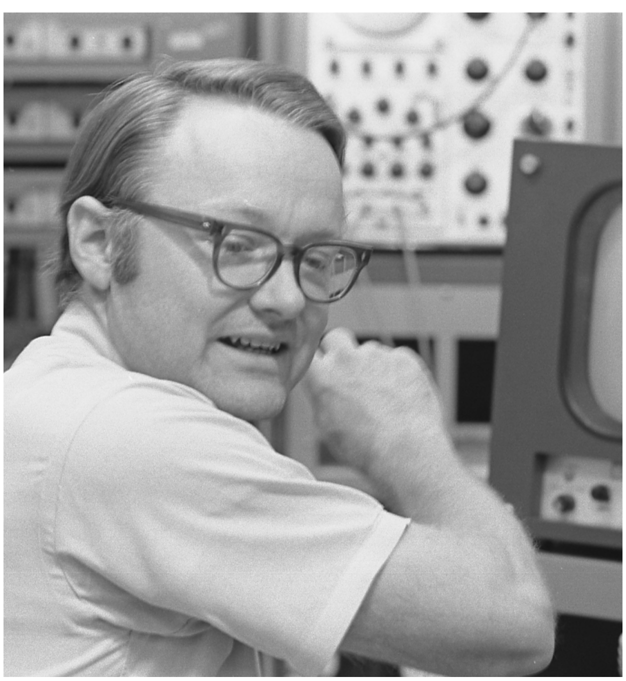

Figure 2-13. Bill English, architect of much of the ARC online hardware and interactive systems.

A Culminating, Watershed Event

In late 1968 a demonstration of huge significance in the development of computing took place. Any lingering doubt in the minds of the few who knew what Engelbart was pursuing would soon evaporate as a part of Doug's vision assumed striking clarity. Computer-aficionado Alan Kay, now at Hewlett Packard, said that it not only changed the way he viewed computing, it changed his life. Charles Irby, a software engineer not at SRI at the time, said it was one of the most impressive things he had ever seen. The place was San Francisco's Civic Auditorium and the occasion was one of the largest computer conferences of the day, the Fall Joint Computer Conference for 1968.[27] On this first day of the conference, December 9, Engelbart had requested a special hour and a half session not just to describe their work but to demonstrate NLS, the ARC's online system. Given that the very large room was filled to overflowing with perhaps 2,000–3,000 people, the courage of his decision was evident. After several months of preparation, his group had set up a live demonstration of what they had created, a system that connected two remote collaborators to an extent never seen before.

Remember this was a time of punched cards, batch processing, and the equivalent of priestly sentinels that kept common folk at arms length from the expensive computing machines that executed their programs. Engelbart was in the Civic Auditorium giving the world its first glimpse of the mouse, a five-fingered keyset, and a CRT as an interactive display. Don Andrews, Jeff Rulifson, and Bill Paxton were taking turns at an identical station at SRI in Menlo Park, where the computing resources were. A large and touchy light-valve projector was used to

[27] Held twice a year the Joint Computer Conferences were sponsored by the Association for Computing Machinery, the IEEE, and the Computer Society.

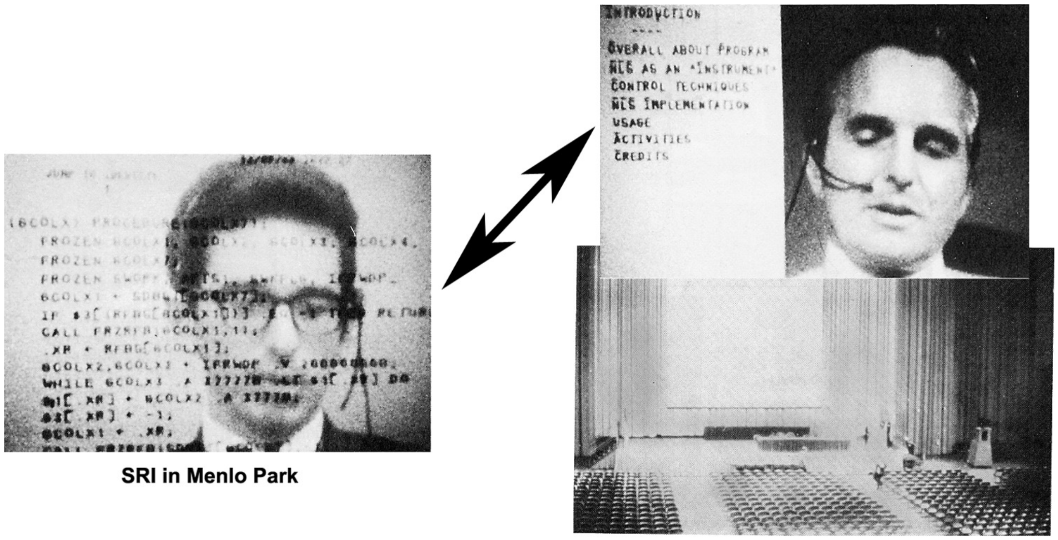

SRI in Menlo Park

Civic Auditorium in San Francisco

Figure 2-14. Images from the 1968 demonstraton: Rulifson (left) at SRI and Engelbart in San Francisco. (*A History of Personal Workstations,* Edited by Adele Goldberg, ACM Press, 1988.)

show everyone in the auditorium Engelbart's NLS screen. Since the normal telephone network couldn't handle the communications capacity needed, a special microwave link was established with a relay on Skyline Drive to cover the 35 or so miles between the sites. Bill English (see Figure 2-13) deserves special credit for his inventive abilities to make all the functions work, both in the evolution of the NLS workstation and particularly here in this prodigious undertaking. He was backstage in San Francisco pulling all the strings.

There, for the first time publicly, Engelbart displayed and explained what NLS and real-time, personally responsive computing could do.[X] He demonstrated online editing and hypertext jumps from a word to its supporting or ancillary information; opened a visual-plus-audio window with colleagues at SRI; and opened a cooperative session in which two separated people could read and edit the same document at the same time, each looking at an identical screen (see Figure 2-14). The crowd was at once mesmerized and sat or stood in stark disbelief.[Y] But at the end, there was a standing ovation for what was truly a watershed moment in the history of computing.[28]

Other Innovations in Computer-Aided Knowledge Work

While the San Francisco demonstration offered a glimpse into a new world of information and communication technology and would become world famous in that development community, that shouldn't imply its recognition or acceptance by a world not yet ready for such capability. Nevertheless, the ARC continued to operate another 9 years. Innovations in NLS continued, and the Center, committed to Engelbart's guiding vision, produced many of the functions we now consider essential to computing. As mentioned, during the first half of the 1960s the group built a startling array of interactive computing features. Moreover, they wisely used these features each day to test and refine them. As mentioned, Engelbart called this practice bootstrapping and it was likely the first instantiation of this approach in computer development: Use the tools you are building to perfect the tools you are building. Engelbart's biographical sketch at the Bootstrap Institute's web page (www.bootstrap.org) provides the following compilation of innovations developed in the ARC:

[28] For 2 days after the demonstration, the ARC operated an open house in specially set up rooms in the Civic Center. Some 2,000 visitors came during this period to learn more about the concepts and actually use one of the two consoles connected to the SDS 940 at SRI (*SRI Intercom* 105, January 29, 1969).

- The mouse
- Dimensional display editing
- In-file object addressing, linking
- Hypermedia
- Outline processing
- Flexible view control
- Multiple windows
- Cross-file editing
- Integrated hypermedia email
- Hypermedia publishing
- Document version control
- Shared-screen teleconferencing
- Computer-aided meetings
- Formatting directives
- Context-sensitive help
- Distributed client-server architecture
- Uniform command syntax
- Universal "user interface" front-end module
- Multi-tool integration
- Grammar-driven command language interpreter
- Protocols for virtual terminals
- Remote procedure call protocols
- Compiled "command meta language"

While this list is imposing, for Engelbart these accomplishments were just the beginning, just the means toward the all-important end of collaborative or group problem solving. Workers at Xerox Palo Alto Research Center (PARC), including a sizable contingent of expatriates from SRI's ARC, added features such as the desktop metaphor, the bit-mapped graphic display, and the graphical user interface (GUI). An important difference was that the thrust of the work at PARC veered somewhat away from Engelbart's conceptions. At PARC the strategy was to make everything easy for the user and to let the machine do whatever it could to enable that. In contrast, as mentioned earlier, Engelbart's philosophy expected the user to change, too. He was always looking for the optimal combinations of human and machine, even if the human had to train substantially to get there.[Z] This difference in philosophy cannot be overemphasized. As these two main branches of the tree of personal computing continued to grow, one could be labeled user-friendly and the other the growth of human-machine proficiency. That schism also existed for a time within PARC between those SRI expatriates who wanted to continue NLS running on a network of minicomputers and those who focused on individual machines and a newer, easier language, called SmallTalk. The latter group won out, but some of them had to move to Apple to see their dream played out.[AA]

The Mouse, Its Acceptance, and Licensing

Between about 1972 and 1987, when the patent expired, SRI licensed the mouse on a nonexclusive basis to perhaps as many as a dozen different mouse purveyors and manufacturers.[29] The term "mouse purveyors" must be used because many of the computer supply houses such as Apple bought their units rather than manufacturing them. Licenses were all nonexclusive and early ones were offered for a $500 initial fee and 2% of the net selling price. By the time that personal computers were hitting the street in serious numbers, however, SRI's mouse patent was nearing the end of its life. For example, the Macintosh was introduced in 1984, 3 years before the SRI patent expired. Thus, Xerox and then Steve Jobs of Apple Computer each bought a paid-up license for $45,000. The known licensees are listed in the following table, roughly in the order in which they were licensed. According to Bonnar Cox, who was handling this licensing at the time, these licenses produced, collectively, about $150,000 for SRI over perhaps 5–6 years when returns were noticeable.

Others who considered licensing but whose participation is uncertain are Logitech (1983) and Televideo (which used Mouse Systems' optical mouse). As a measure of the ubiquity of this ergonomically effective device, Logitech announced on its web site in September 2003 that it had shipped its 500-millionth mouse![BB]

[29] The mouse patent, No. 3,541,541, was issued November 17, 1970 and ran until that date in 1987.

Nonexclusive Licensees of the SRI Mouse and Approximate License Issue Date			
Computer Displays, Inc. (later Adage) Boston, MA	1970	Apple Computer Cupertino, CA	1981
Cybermex Stanford, CA	By 1972	Kinetronics Arlington, MA	1982
Century Tool San Carlos, CA	1977	3G Company, Inc. Gaston, OR	1982
Tymshare, Inc. Cupertino, CA	1978	Mouse Systems Corp. Sunnyvale, CA	1983,4
Alps Electric Co. (for Microsoft, Tandy, Wang, etc.) Yokahama, Japan	~1980	WICO Corp. Niles, IL	1983
Hawley Labs (Mouse House Division) Berkeley, CA	1981	Fujitsu Limited Tokyo, Japan	1985
Xerox Corp. Palo Alto, CA	1981	Product Associates* Redwood City, CA	?

*(Martin Hardy, ARC alumnus)

Networking

One aspect of the ARC's work receives less publicity than its other contributions is computer networking. This topic is discussed in greater length in Chapter 3, but a bit more deserves to be related here. Bob Taylor was hired at ARPA in about 1965 to help create networking among computers. Though a number of small individual networks of computers had arisen, he was frustrated by the number of different terminals he was forced to have in his Pentagon office to use this collection of networks. Taylor clearly saw the benefits of a single-access terminal and persuaded Larry Roberts, who had done similar networking at Lincoln Labs, to join him. Roberts soon became one of those pivotal players at ARPA who saw the advantages of computer networking and worked tirelessly to bring it to life. In the periodic meetings of his ARPA contractor community, Larry would remind those gathered of the natural benefits of connecting computers together. One of the biggest reasons was something called resource sharing. Bob Kahn, a member of that community and later a director of IPTO, was an important advocate, suggesting that good connectivity between time-shared hosts would let major sites each specialize in some important aspect of computing. Some would have elaborate programs in graphics, others in simulation or data processing. The concept seemed economically sound: each time-share host would need only enough expensive memory to serve one chosen specialty, and all users on the network would gain remote operational access to those programs. Users who had a graphics problem would simply log in at the University of Utah, regardless of where they were working.

In this new world of interactive computing and remote, cross-network access, the ARC was a unique site. But at that moment there was another reason to attach it to this new network. Early in the discussions of networks at ARPA, the need became apparent for a place on the network that knew where all the resources resided. There was also a need for a central repository of the common, correct software to be used to obtain access and other functionality on the network. That place was called the Network Information Center or NIC. The concept of such a place surfaced during the 1967 ARPA principal investigator meeting where these networking concepts were first discussed. Probably because the NIC didn't sound "researchy" enough and because they resisted sharing their machines with remote users, other contractors seemed reluctant to volunteer to run the NIC. But Engelbart saw the NIC differently, as a meeting ground for a family of network users, and volunteered SRI for that role, even though no one in the ARC had an appropriate background.

In these early years of networking, Engelbart's initiative had some profound consequences for his group and SRI:

- The first though temporary leadership of the Network Working Group and a design role in many of the earliest networking innovations, such as:
  - The network virtual terminal—This innovation reduced complexity by enabling network access terminals to use a single interface specification rather than a different one for every type of host on the network. The same concept also benefited the network hosts.[CC]
  - The network virtual circuit—The circuits freed packet traffic from the need to follow a consistent or *a priori* route.[DD]
  - Some of the early designs of communications protocols—Telnet, the ARPANET's remote terminal access protocol, was first advanced in RFC 97 in early 1971.[30]
- The receipt from ARPA of the second packet switch, called an IMP, in early October 1969.
- The first computer-computer communications, with the University of California, Los Angeles (UCLA), in October 1969.
- The operation of the NIC, a network protocol repository and the point of issue for network names and addresses for the ARPANET, then Internet, a role it held for 21 years.[31]
- The first network locator directory, called WHOIS.

The diagram in Figure 2-15 indicates the conceptual state of networking in the ARC. This sketch was taken from an internal administrative presentation made by Engelbart in June 1973. The free flow of information, access to central repositories, and network-wide services are evident. The diagram could depict the Internet today, 30 years later.

One of the Internet's earliest and most notable contributors spent a portion of his networking days in the ARC. His name was Jon Postel, and he was in the ARC from September 1974 to March 1977. He was an early critic of the internet protocol, TCP, as originally specified.[32] While at SRI he saw the need for two protocols instead of the single TCP version. Within 6 months after leaving SRI, Postel's advocacy helped spawn the simpler IP protocol, which was faster and eliminated the end-to-end reliability negotiations of TCP. Postel became one of the Internet's true pioneers, the principal architect of many of its functional protocols, and a pivotal contributor to the domain-naming concept that gave all hosts a text-based name along with their numeric network addresses. Most of Postel's best-known contributions were made at USC's Information Sciences Institute after he left SRI, but he got some of his early grounding in the ARC. He died in October 1998 from complications in a heart operation, and his contributions to the Internet were acknowledged worldwide.

The "ARC" Develops a Leak

One problem faced by people who are way ahead of their time is being cast as irrelevant. Since truly dramatic visions are by nature not widely held, Engelbart and his Center began to lose their sponsorship in the early to mid-1970s. The NLS software that embodied the concepts of the ARC was idiosyncratic; that is, it was not easily mapped into operations different than those found in the ARC. In addition, and perhaps most importantly, it was a moving target. Continual change is not uncommon at SRI or any research site. Nevertheless, when government agencies such as the DoD were unable to adapt NLS to their own world or their world to NLS, it became harder and harder for them to continue their sponsorship. Even ARPA must at some point transition its innovations to its DoD "clients" or abandon them for other pursuits.

[30] Several ARC people contributed to Telnet development, including Richard Watson, John Melvin, and Jim White.
[31] Richard Watson provided the early guidance for the NIC but for most of the two decades that it offered services to first the ARPA and then the general Internet community, it was under the leadership of Elizabeth Feinler, known throughout her world as "Jake".

[32] See the Chapter 3 on computer networking for a discussion on SRI's own specification of TCP.

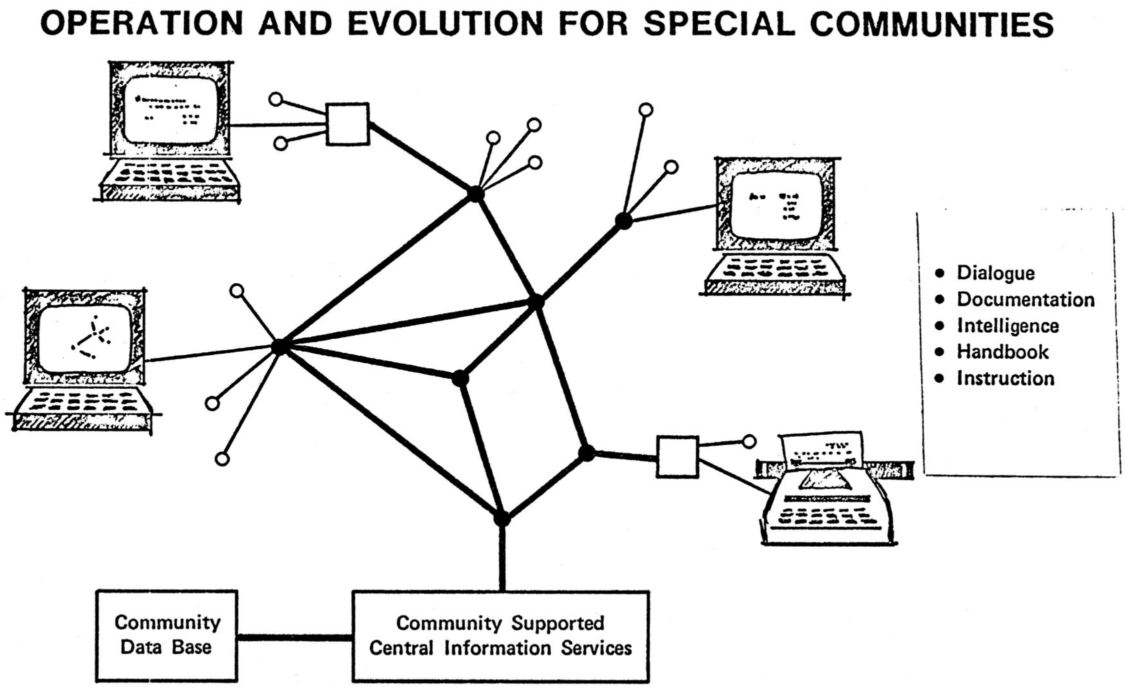

Figure 2-15. Conceptual network functionality from an SRI internal presentation by Engelbart dated June 13, 1973.

Another reason for atrophy may have been a decrease in innovation at the ARC caused by the departure of key staff members. In a real sense the 1968 demonstration proved to be an emotional pinnacle for what the system could do. A revolutionary research system had been built and demonstrated, but neither its continued incremental improvement nor its transition to new hardware was as exciting to some as starting anew with some other opportunity. In addition, the technology was changing, and Engelbart, in the opinion of some ARC staff members, was not. Smaller machines that started to appear affordable for one person seemed on the way. But Engelbart thought it best to continue to rely on large time-share machines, in particular the soon to be released DEC PDP-10.

For all these reasons, around 1971, many of the key ARC staff members started to leave for Xerox PARC. There a number of bright people, some of whose lives had been altered by the 1968 demonstration, had assembled under Taylor, the former head the ARPA information technology office that had funded Engelbart. Bill English, Bill Duvall, Jeff Rulifson, Charles Irby, Bill Paxton, and Don Wallace were among those who carried at least the functionality of the Engelbart vision to PARC where the next significant milestones in the development of personal computing would occur.

In spite of these departures, Engelbart and his ARC would continue for several more years. By 1972 the Center had grown to over 40 people and the administrative load that Doug deplored got only worse. With the encouragement of management and for the first time, he appointed two assistant managers. While that helped administratively, by the mid-1970s there were more troubling signs on the horizon, in particular, the rising inability to get or retain contracts that would fund the research that in his mind he had only begun to explore.

One clear evidence of this decrease in sponsorship occurred in late 1975 as Engelbart submitted a SRI budget that requested some relaxation of the ARC's financial goals so that some of his new, but unsponsored ideas could be developed. This act was essentially a request for other labs in his division to partially subsidize his operation and, as it turned out, they were willing to do. But the coming year, 1976, wasn't a particularly healthy one for the division as a whole and not only weren't the subsidies fully forthcoming, pressure arose on the entire division to reduce costs. Such pressure, plus the ARC's existing turnover

ENGELBART'S INSIGHT INTO THE SCALABILITY OF ELECTRONIC CIRCUITS

There were reports and at least one paper revealing Engelbart's conviction that scaling to microcircuits would yield greater logic power, but a good feeling for that time comes from a speech by colleague Jack Goldberg at SRI's tribute to Engelbart when he received the SRI 2000 Gibson Award: "In 1958–59, while waiting for funds to pursue his primary interest, Engelbart led a project at SRI called the Philosophy of Logic Realization. It was at a time when there were many physical schemes for doing digital logic—optics, magnetics, and solid state. Transistors had questionable reliability, and they came one device per package. Engelbart applied dimensional analysis to see how well various digital logic schemes could scale to micro levels. I remember clearly Engelbart telling me then—before the invention of integrated circuits— that the scalability of charge-controlled devices, such as transistors, would result in an inexorable march to micro-level devices. I think this insight convinced him that the computing power needed for individual human augmentation would be available eventually."

problems, meant that by the end of 1976 the professional count was down to perhaps 15 plus a smaller number of programmers and research analysts providing NIC services. The ARC was trying to retain revenue by selling NLS-centered services on its time-share hosts. The revenue stream from such subscriptions and their associated expenses were enough out of the mainstream of SRI's accounting system that the Center's financial performance was consistently in doubt.

For better or for worse, these are the times when research managers are compelled to act. Engelbart was asked to step down from the leadership of the ARC in late 1976. How all this played out is given in more detail in Appendix G, but in this instance, management, faced with the ARC's falling income stream and potential market value in its software and services, decided to sell it to the growing Tymshare Corporation at the beginning of 1978. Some ARC staff, including the NIC, transferred into the communication activity at SRI but the sale of the ARC and its intellectual property meant that Engelbart, along with many of his remaining staff, opted to move to Tymshare and thus end his two fruitful decades at SRI.

An Epilogue: Legacies of the ARC, Intended or Not, Accepted or Not

Although personal computing originated at SRI under Engelbart's vision, he still believes that personal computing, whatever that term means, was only a piece of his dream, just one necessary milestone along the way. But in the 7–8 years from his early work at SRI to the spectacular demonstration at the 1968 Fall Joint Computer Conference in San Francisco, Engelbart, English, Rulifson, Paxton, Duvall, and many other members of the ARC defined, for the first time anywhere, what we think of today as personal computing. Their collective accomplishments forever changed the way computing was viewed. Though he clung perhaps too long to the behemoth time-share computers, Engelbart also saw back in 1959, before it became evident in Moore's Law, that computing resources would one day become affordable enough that time-sharing would no longer be necessary (see text box). Perhaps unwittingly, Engelbart's cost-benefit notion of micro-miniaturization would be advanced by our predilections simply to own things, including our own computer. In any case the stage is now set to go beyond simple affordability or ownership into a world of rich networking where we can truly be indifferent about where the actual computing takes place or the information resides.

While it may seem a fine point, Engelbart in effect worried that personal computing, as he saw it, for example, in the developments at Xerox PARC, would result more in subrogating or automating intelligence rather than *augmenting* it. Certainly, it can't be detrimental to hide the arcane complexities of our interactions with machines, for they can only detract from some higher purpose to which the man-machine combination can aspire. I have always believed that it is the first role of a personal computer's power to bring about a more natural interaction. While that is not likely to appear until speech understanding is perfected, we have come a long way. But what illuminates Engelbart's vision best is what is still missing. Doug is right in that all this deference to the human does not, in itself, lead to the tackling of more complex and meaningful problems, and it certainly does not necessarily demand more from the user. The direction that personal computing is taking is to help make us more versatile in what we can do easily, but no where is there a staircase of computing sophistication that takes a user to a higher effective plane of thought—which is, after all, what augmenting intelligence really means.

This is a story still unfolding. Though it may take concerted effort to confirm it, Engelbart's concept of how computers can truly augment our own intellect is probably developing around us. One only has to look at the young people of today, who are comfortable in gaining benefit from a personal computer and who have little trouble with a computer-stimulated intuition, to see how important the gravitation of computer power toward the individual has been. To them the means, the machine and its associated networking, are becoming transparent to the services available. And it is often the use of computer models that opens a more complete understanding of many complex processes such as the human genome and its proteome.

But what of Engelbart's Holy Grail, elevating computer-empowerment groups toward unprecedented accomplishments? That too is arriving. While he remains discouraged by society's blindness to his vision, there are plenty of examples, probably countless ones, of computers enabling groups to function in ways unachievable before. Consider just one case. In the field of product development, the time to market or to break-even is today's yardstick of choice.

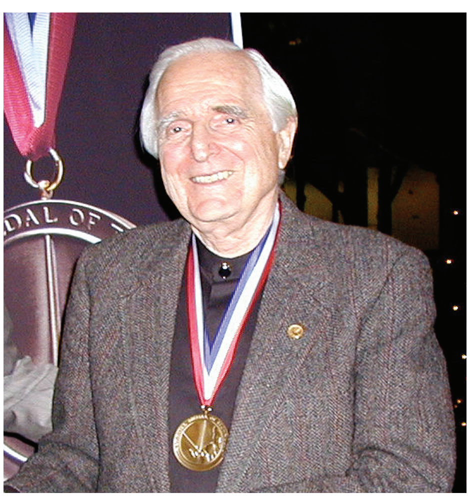

Figure 2-16. Engelbart Wearing the National Medal of Technology for 2000

Competition has demanded this time urgency and, since time is money, more rapid development sets the stage for lower costs. Rapid prototyping is a process whereby teams, representing the talent needed to bring a product to market, are now joined together through computers and their supporting networks to cut the time and cost of product development. Production engineers, market specialists, profitability managers, and life-cycle maintenance people all join the early-stage designers to shorten the product development cycle. And their physical location is virtually immaterial. Development durations that used to be measured in years with sequential handoff are now measured in months or weeks. These are nothing more than Engelbart's envisioned groups, collaborating through networked computers, across space, to accomplish what would have been impossible before. Such acceleration is the *sine qua non* of today's successful companies and good evidence to support his thesis.

Other examples veritably flow from the Internet. When we have the tools to locate what we seek more efficiently, it will become a true extension of our awareness and our individual and collective capabilities. The groundwork for that extension was laid in the online access system started in the ARC. Today, scientists from all over the world can link to the Hubble telescope in space. One of the world's largest and most sophisticated ionospheric research radars, soon to be built by SRI, will be reachable from any Internet site. That kind of collective, real-time access, together with the rapid sharing of information, promote great insights, extend collaboration, and offer a continued acceleration of man's quest for knowledge. Engelbart's desire for the co-evolution of both humans and machines, in their interdependent partnership, is still valid, however, man's preference will continue to place the greater change in his machines than in himself. Nevertheless, in this man-computer partnership, our *collective* abilities and knowledge will continue to accelerate just as his vision portrayed.

Recognition for Engelbart's Achievements

Although neither Engelbart nor SRI ever gleaned much compensation for the innovations mentioned in this story, Engelbart has received some individual recognition. Over the last decade or so Engelbart has been honored with over 30 awards, including:

- The MIT-Lemelson prize with its $500,000 award for innovation
- The Turing Award, the highest honor in computer science worldwide
- The IEEE Von Neumann Award for contributions to information technology
- The National Medal of Technology, the highest U.S. award for innovation (see Figure 2-16)

As evidence of the respect I have for what Doug and the ARC accomplished in the last half of the 1960s, I had a major hand, with Kinney Thiele, in his nomination for the last two awards on this list.

In addition to Engelbart's personal laurels, the major professional computing society, the Association for Computing Machinery or ACM, did acknowledge two others. ACM gave its 1990 Software Systems Award to Engelbart, English, and Rulifson for the development of NLS, a full quarter-century after it was created.

Though appreciative of this worldwide recognition of his achievements, Engelbart continues to believe that the honors are somewhat misplaced. He still fights for the principles of group, organizational, and societal enablement and problem solving that he voiced almost 40 years ago. Perhaps it is enough to say that, without question, there is much to be honored and respected and yet much left to do.

Computing Science at SRI

It was late in 1954 and the incredible excitement and stress of building one of the world's first real-time computing systems, ERMA, was in full swing.[33] But in spite of that preoccupation and the newness of computers like ERMA, there was growing realization that they were emerging as broadly useful, general-purpose tools. As evidence, Jerre Noe, then Assistant Director of Engineering at SRI, decided to start a Computer Laboratory and named Byron Bennett as its first head.

Some members of that new lab had been responsible for designing the logic circuits of ERMA and then, as the project was winding down, could see the need for more rigor and formalism in computer design. Jack Goldberg, Bonnar "Bart" Cox, and Bill Kautz had handled most of the logic design of ERMA and, according to Goldberg, their work had been all technology and no science. It seemed to them that there should be better ways to approach the design of these new and increasingly capable digital tools. Also, since the individual logic components were so unreliable, the design would somehow have to take those failures into account. So, in 1956 Kautz started an SRI program to try to define mathematical methods for computer design.

By 1957 the new program, part of what was by then called the Computer Techniques Laboratory, was concentrating on computer hardware design and reliability. This work continued for at least a decade or so, and in March 1969 the group gained a stronger identity under Goldberg's leadership (see Figure 2-17). Around 1970 the research focus began to

Figure 2-17. Jack Goldberg, first director of the Computer Science Lab.

[33] See the first section of this chapter.

shift to software, in which the lab would eventually make its greatest contributions. Though the group continued to grow, it didn't achieve laboratory status as the Computer Science Lab (CSL) until 1977. Given that the CSL is still a vibrant, technically excellent laboratory, if its history is measured from its start in 1957, it is one of the oldest continuously operating groups at SRI. Of great consequence, however, is the dedication of its staff to becoming a world-class software laboratory. This quest often means trying to arrive at not just a narrow, specific result, but, where possible, defining whole new classes of engineering approaches or solutions. This competency, in a field that is still very vital and expanding, has enabled CSL to also be one of the most financially successful labs at SRI. From its wide range of contributions to computing science and technology, only four areas will be presented here: magnetic logic, fault-tolerant computing, software formal methods, and information security.

All-Magnetic Logic

In the 1950s there was an increasing realization that computers would serve well in control and switching systems such as for telephones. But the unreliability of the logic elements of the day, including vacuum tubes and early semiconductors, led to a perceived need for a highly reliable logic element. Hew Crane believed, from work he had done at RCA Labs, that he could build a computer using only magnetic elements. Except for the wire itself, they were the most reliable components in a computing system at that time. If a logic system that used only magnetic cores and copper wire could be built, it would be essentially trouble-free.

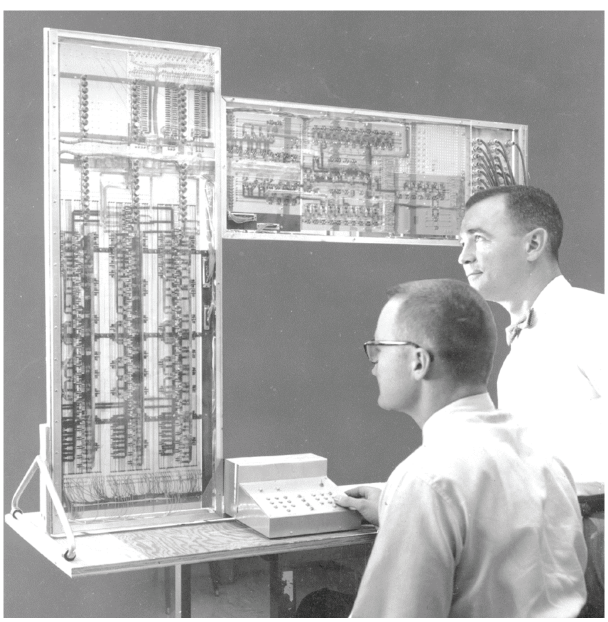

Figure 2-18. Bill English (seated) and Hew Crane with the world's first and perhaps only all-magnetic-logic arithmetic unit (about 1961).

The first major problem facing the designers of an all-magnetic computer was the need for a two-state element that was also capable of inter-element gain; that is, a binary element that could be used to control the state of a similar element without the need for an intermediate amplifier. Some major computer manufacturers, such as Burroughs, had tried magnetic elements but, in the end, all opted for much less reliable vacuum tube diodes. The answer attempted at SRI was something called a multi-aperture core. By the late 1950s Burroughs and then AMP Inc. began supporting SRI's work. With SRI assistance AMP built and delivered many all-magnetic logic systems to handle critical functions. For example, all-magnetic logic systems were used to control one of the largest railroad switching yards in the world for the National Railroad in Toronto. These systems seemed to find a role wherever reliability was paramount, including in aerospace applications. The AMP logic systems were used in, for example, the docking radar for the Gemini capsule, a flight programmer for the Agena rocket, and communications satellite access control. Parts of the New York subway system also used such units (see box).

The SRI work culminated in the early 1960s in a project under U.S. Air Force sponsorship that resulted in the all-magnetic arithmetic unit shown in Figure 2-18. Even though the reliability of semiconductor devices would soon improve, this unit attracted a lot of attention.[EE, 34] According to an estimate retired AMP vice president Sweeney gave to Crane, AMP sold, in today's dollars, in the vicinity of $300 million worth of magnetic logic devices and systems.[FF]

[34] Others participating in the all-magnetic computer were Jim Baer, Doug Engelbart, Bill English, Herb Heckler, Ed Van de Riet, and Joe Hunt. SRI recently donated this system to the Computer History Museum, Mountain View, CA.

Joe Sweeney, the vice president of research and engineering at AMP, Inc., was one of the sponsors of the SRI work on all-magnetic logic. He described going to see a part of the New York subway control system 10 years after its installation. Down several levels under Grand Central Station, he and subway technicians found the control box covered with a decade of dust and debris from passing trains. According to the employees, the control box had never been opened. Sweeney vacuumed away the dust, opened the box, and inspected the innards. Pleased with what they saw, they closed the box and continued to forget about it.

Fault-Tolerance, Distributed Systems, and Formal Software Methods

The role of fault-tolerance in increasing the reliability of a logic system probably had roots at SRI in the late 1950s, when researchers were confronting the unreliable building of monolithic circuits. The device population in such circuits was increasing dramatically and the failure of one logic gate or memory cell could render the whole integrated circuit useless. To compensate for failures, could circuits be "overbuilt" and then self-organized to provide some guaranteed level of performance? That line of reasoning would take SRI innovators along two notable directions, one in the direction of artificial intelligence (see the opening of Chapter 4) and a very different one that we will briefly delve into here. This second approach directly explored the concept of a new field in logic and computer design called fault-tolerance. Though the subject had been under discussion for some time, it was the early 1970s before the first project to deal comprehensively with fault-tolerant systems was undertaken by the Computer Science Group for ARPA and the Office of Naval Research (ONR).[GG]

The purpose of this first study was to assess the state of the art in fault-tolerance for ARPA and estimate its effectiveness in the design of new computers. Evaluations were made of existing systems, concepts, and theory along with any new approaches that might be in the offing at SRI or elsewhere. This kind of project is excellent in helping form the solid base on which new, more promising concepts can be built, and that is what happened at SRI. The study considered both hardware and the emerging software approaches to fault-tolerance. It concluded that fault-tolerance theory and technology could significantly improve system reliability without dramatically increasing system cost. There were, however, important gaps in the art of designing and implementing software that would support fault-tolerant hardware. These included a lack of real time system diagnostics and the inability of the operating systems of the day to provide the hardware-software integration that rapid system restoration demands. The study also examined the role in reliability of a more formal approach to system specification. Finally, the authors advanced a novel approach to hardware-software unification using reconfigurable computer memory, which would be much cheaper than massive redundancy.

That latter insight was simultaneously being explored elsewhere within CSL as its staff began to investigate the question of highly dependable computing for NASA. Part of NASA's mission in the early 1970s was to promote the technologies that would make both aircraft and spacecraft more reliable. It was clear by then that computers would find their way into aircraft; the question was whether they could be designed and built to be continuously available. About 1973 NASA asked SRI to use all it knew about fault-tolerant computing and build an experimental computer that could control the safety-critical functions of airplanes. An ultra-dependable controller was vital. What happened next is a hallmark of a truly competent research organization.

The need for ultra-reliability could have been met in more than one way, of course. One possible approach was to use redundant computers, some of which would operate in "hot" standby mode in case the principal machine failed. (Tandem Computer was founded on that principle in 1974.) But the staff of CSL took another, more lofty approach to the reliability question; namely, to integrate a number of separate processors (in this case five) tightly networked together. It was called a software-implemented fault tolerant computer, or SIFT for short (see Figure 2-19).[HH] Again note that the emphasis was deliberately not on duplicate hardware and software running in the

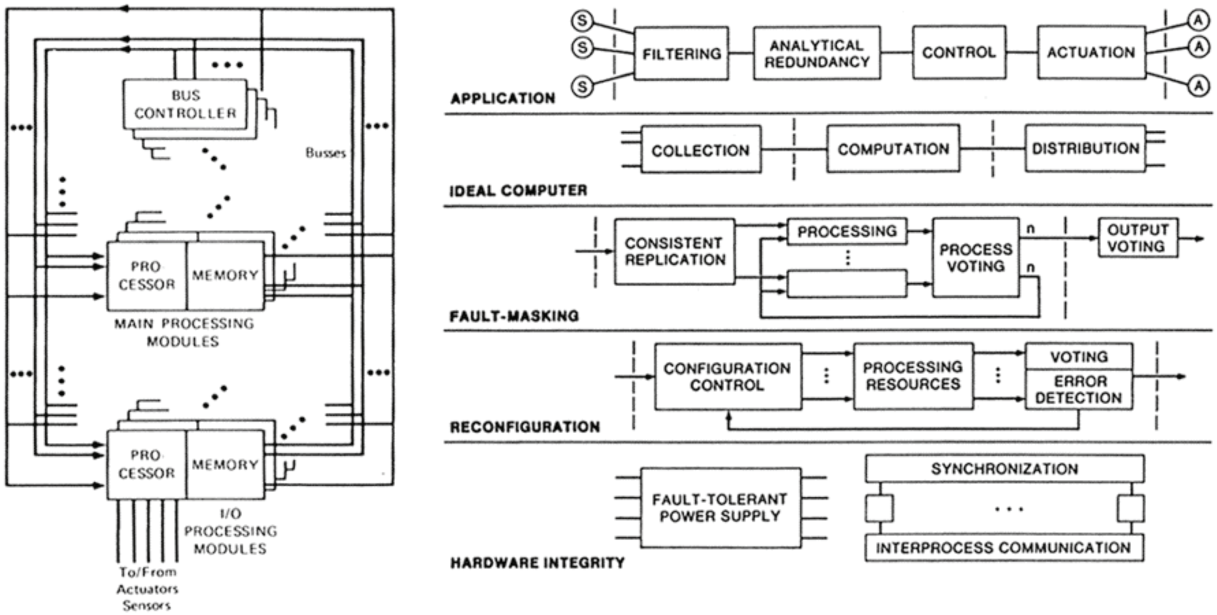

Figure 2-19. The physical (left) and the logical (right) structure of SIFT. (Notice the hierarchical form of the logical structure.)

shadows, but on gaining the requisite reliability from more richly redundant hardware and, especially, highly accurate and resilient software. This integrated multiple processor approach was a beachhead in an emerging field that came to be known as distributed computing.

This focus on tightly coupled processors plus the fact that the Lab was immersed in the notions of provable software correctness ensured that the problem of fault-tolerance would be tackled fundamentally. There were no rules for how multiple processors would even determine their individual role with respect to the instantaneous load, let alone assess their individual or collective health. In exploring this new ground, the CSL decided that each task would be multiply (redundantly) executed, and a vote on correctness would be taken before a task was invoked. As a result, the separate processors needed to be only loosely synchronized. Given that approach and the need for correctness, one of the first decisions to be made was whether any of the processors could or should believe the professed state of another.

So, in pursuing the SIFT design, CSL's Dr. Robert Shostak and Marshall Pease conducted one of the first forays into the question of consistency in "distributed systems." They both formulated the issue and provided a general solution for how to deal with it.[II] The issue was popularized by the Lab's Dr. Leslie Lamport when he called this potential for inconsistency in distributed systems the Byzantine Generals Problem.[JJ] The name refers to a condition in systems with multiply redundant processors where states that should be identical have been computed or communicated inconsistently. The implication is that to be able to detect the presence of a faulty processor, due either to its own error or that of its communications channel, requires a certain minimum number of like participating devices. By first tackling and then solving the more general problem of inconsistency, the CSL researchers laid new and fundamental ground for distributed systems, an expanding field in information science.

The other consequence of taking the SIFT design along a higher developmental path was the use of software formal proofs. Most people are unaware that software can be written to a certain level of uniformity using so-called specification languages that make it possible to prove that the resulting code will execute exactly as specified. This approach to software design is seldom used because the processes are difficult to understand and the proofs required to show that the code complies are computationally intensive. The hope has been that the computers would become powerful enough to absorb this complexity and computational intensity or that the provable languages would become closer to a natural, readable language that mere mortals could use. Neither of these states has been achieved, but CSL has contributed as much as or more than

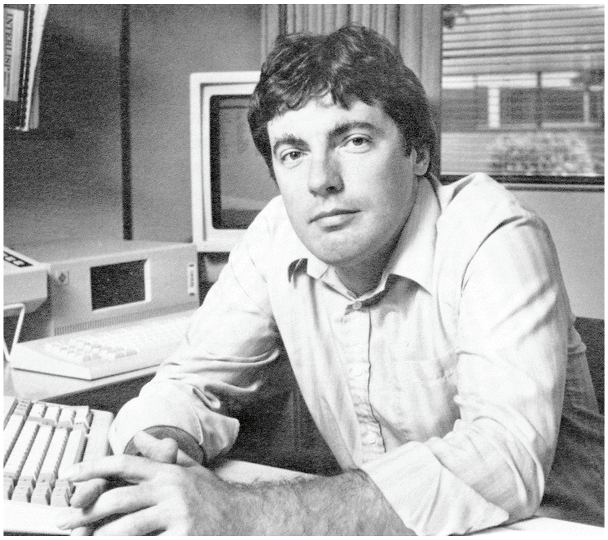

Figure 2-20. John Rushby, one of the world's leading authorities on formal methods.

any other group in the world toward achieving these goals for software. In hardware design, formally based analysis techniques are now commonly used to analyze digital circuits before they are fabricated.

The purpose of formal software methods, of course, is to improve the reliability of software. Reliability is beneficial for all software, but critical for contexts in which software malfunctions endanger peoples' lives. Furthermore, since software now accounts for the bulk of digital system costs, reducing its almost ubiquitous bugs would be a huge cost advantage. John Rushby, one of CSL's most distinguished researchers, loves to point out that because software doesn't wear out, the only software errors are design errors (see Figure 2-20). He reasons that, like other engineering disciplines, software development needs a mathematical basis for its design. Other engineering disciplines have long ago developed beyond a trial-and-error approach but, alas, virtually all software development finds itself still in that stage. Rushby's colleague and the current (2004) director of CSL, Pat Lincoln, expresses the need similarly: "Bridge designers use finite element analysis and computer-aided design tools to test their intended structure before they ever begin construction. Software engineers, more often than not, build systems without knowing if they will, in the end, collapse." Lincoln believes there are underlying mathematical principles of software, just as there is a physics of bridges, and eventually software engineers will be able to know at design time if a software system will work correctly. He asserts that, "Formal methods will one day be applied to all software. It will be embedded into compilers. Standard engineering practice will include mathematical analysis of designs, at every stage."

Even given current technology, applying formal proofs to all but the simplest of programs has been daunting. The SIFT computer was complex enough that, given the power of machines at that time, its software was never completely proven. Yet for more tractable problems, such as the guaranteed synchronization of multiple system clocks, formal methods have worked. They have also been successfully applied to the aforementioned, complexity-driven technology, integrated circuit hardware. In the end the SIFT work established the basis for what is now called the "state machine" approach to fault-tolerant systems, and the SIFT computer, built by Bendix, was installed at NASA Langley's Airlab in 1982, where it ran for about a decade. Not all of the original SIFT design features were realized, however.[KK] Tasks had to be randomly assigned to processors rather than dynamically by need. Voting also had to be scheduled rather than being transparent to the application programs.

One powerful concept in software design that can make a program amenable to verification is the use of functional hierarchies. This concept is actually one of decomposition, of structuring functionality so that by a series of layered abstractions one can go from simple or atomic operators whose functionality can be easily verified, to increasingly comprehensive operators that directly and unambiguously depend on that same lower level functionality. This concept was used in SIFT, for example (see Figure 2-19). But an obvious difficulty in this approach is defining each step in the abstraction with enough rigor to ensure verification as the abstraction increases. In security, for example, the hierarchies must span from the lowest level operator that can be ensured to what usually amounts to a higher-level, but still unambiguous security policy.

Using this approach to software construction, CSL developed two of the most recognized and useful tools in the verification of software programs: Hierarchical Development Methodology (HDM) and PVS. HDM and its successor, Enhanced-HDM, were among the first programs anywhere for the formal specification and verification of software. They used the powerful concept of formal abstraction and relied on hierarchies in the software design process. Also, behind HDM is a specification language called Special (and, subsequently, E-Special). Reflecting a continued need for innovation in this field, the evolution between HDM and its successor was radical. Because of the need for the inviolability of critical national security software, development of both these programs was sponsored by the U.S. National Security Agency.

Development of PVS, on the other hand, was supported by several hundred thousand dollars of SRI internal funds. Although Natarajan Shankar, John Rushby, and Sam Owre began work on it 1990, PVS is still a current expression of the state of the art in program specification and verification. The acronym originally stood for, among other things, the "Peoples" Verification System to indicate the desire to broaden the usability of such programs. But it has come to be known, like SRI itself, only by its acronym. The goal in building PVS was to create and make available a more usable specification and verification system than, say, E-HDM. PVS also includes a necessary software specification language and a theorem prover. Today, with the participation of Leonardo DeMoura and Ashish Tiwari, a suite of utilities has also been built around PVS. These include some automated decision procedures and something called SAL, which stands for Symbolic Analysis Laboratory. SAL performs the abstraction needed in the formal characterization of concurrent systems and has been applied to such different areas as automobile control systems (GM cruise control) and a biological model of diabetes. Finally, PVS is unusual in being licensed in a LINUX-like way, essentially free to anyone under terms that any improvements they make will be fed back to SRI so that configuration control can be maintained and the package will evolve to everyone's benefit. With this kind of accessibility, PVS has very likely become the most widely used formal software creation program in the world.

In a somewhat related vein, CSL researchers such as Joseph Goguen and Jose Meseguer have been the instigators of research into what they call modern ultra-high-level software languages. These are languages that can express functionality and variables in a manner much closer to a specification than can ordinary languages such as FORTRAN or the C family. "High level" in this case means that functionality comes from more expressive statements that avoid most of the detail of more common languages. They reflect a belief that detail is a hobgoblin of accurate software, and the less the author has to deal with it the better. The first ultra-high-level language developed by CSL was called OBJ. It was intended to bring together several advances in programming style: logic programming with its equational form,[35] a declarative approach to execution,[36] and object-orientation for the simplicity of equivalent representation of a wide range of variables, constructs, and functions. OBJ and its descendent, MAUDE, now the responsibility of Carolyn Talcott and Steven Eker, have been licensed for research and development use in hundreds of laboratories worldwide. They are part of the search for the next generation of more accurate and easily programmed software languages.

As evidence of the general utility of the CSL work and as a peak at the future, lab director Pat Lincoln and some of the above CSL staff are now applying formal methods to new disciplines including molecular biology. Here one of the great benefits of a diverse research institute comes into play. With the collaboration of SRI biologists,[37] formal methods are enabling a new approach to understanding biological systems. Using some of the above automated reasoning tools such as PVS, SAL, and MAUDE, CSL has integrated them with a number of other analytical tools to form BioSPICE.[38] This collection is another open-source platform being made available to molecular biology labs worldwide.

[35] This means that statements lend themselves to replacement by equivalences, and functionality can thus be verified.

[36] Software languages can be usefully grouped into two types: imperative languages such as FORTRAN, C, Ada, and Cobol, wherein all actions to achieve an outcome are explicitly programmed; and declarative languages such as Prolog and LISP, wherein a computational goal is specified but the method for achieving that goal is left to the language. For a variety of reasons declarative languages are still not accepted for "industrial" or critical uses.

[37] Including Keith Lauderoute, Merrill Knapp, Larry Toll, and Analisa D'Andrea.

[38] This term is an intentional analog to the integrated circuit characterization program, SPICE, built at the University of California Berkeley, which stands for simulation program with integrated circuit emphasis.

Information Security

Perhaps the second greatest shortcoming of software systems, after poor reliability, has been their vulnerability to outside influence. This flaw applies not just to programs themselves, but to the data that surround them and the accessibility of both to unauthorized people. The Internet and the broad and simplified access it offers to users, regardless of their location, has underscored a problem that has been part of software from its beginning, particularly since the creation of time-shared, remotely accessible machines. The rush toward innovation in and employment of computers has never been able to wait for the fundamental conventions and constraints that are necessary for secure systems. Even worse, once the systems are in place and pervasive, both hardware and software, there are no practical ways security can be retrofitted. Remarkably, this lack of focus or discipline in the manufacture of general computing systems has to this day prevented the introduction of a system capable of holding multiple, independent levels of secure information guaranteed to remain separated. The two major aspects of security are protection against (1) unauthorized access to or modification of information and (2) denial of a computing or communications service.

The first major initiatives into computer security, both at SRI and elsewhere, can be tied to the introduction of a number of time-share computing concepts in the late 1960s. Machines were expensive and sharing them among many users was a way to spread costs. But multiple users on a single processor obviously raised questions of privacy, even if memory was somehow segregated. These concerns were of particular interest to the military, with its multilevel conventions of secrecy and the corresponding need to prove that one level of secure access on a machine couldn't compromise information at a higher level on that same machine.

With this motivation and a repertoire of tools and knowledge about formal software specifications such as HDM, CSL created one of the first provably secure computer system designs, the Provably Secure Operating System (PSOS). PSOS had a hierarchically structured, highly regulated (strongly typed), capability-based hardware-software architecture.[39, LL] Computing systems based on PSOS's notion of strongly typed objects were eventually implemented by Honeywell in the 1970s and by its spinout, Secure Computing Corporation, in about 1985.[40] But because the computing market was growing faster than these specially configured systems could be sold, they fell outside the mainstream, and as a consequence, their overall cost remained high. In addition, people's ability to rationalize away even known system vulnerabilities in favor of more rapid and less costly implementation, has retarded growth in the market for secure systems and stalled the necessary standards work needed to make them pervasive.

Beyond multi-user hosts, one of the next most vulnerable points in the landscape of software systems is a database. Securing a database requires controlling access, and a primary issue is whether it is possible to isolate certain information from some users who might have authorized access to different or limited segments of that information space. In 1984 Dorothy Denning, a staff member of CSL at the time, started a long-term project for the US Air Force Rome Laboratory on the design of a multiaccess database called "Secure Data-Views." The goal was to build a relational database that could meet the criteria for the National Security Agency's highest security level, Class A1. Again, an attempt was made to create a mathematical model that defined the behavior of the database, particularly with respect to its security properties. From this model would come design specifications that could be verified, thus ensuring that only certain parts of a database could be accessed. As the name implies, controls were applied not to individual data elements per se, but to the views of the database that were authorized. A simple example would be a database of employees in which their salaries were viewable by only selected parties. This approach to database security was taken over by CSL's Teresa Lunt and came to be called SeaView.[MM] The SeaView security model is used today in research settings and in the design of multilevel

[39] CSL researchers Goguen and Meseguer, under the NSA-sponsored PSOS project, defined the highly influential "noninterference" standard for multiple levels of security on a single computer system.

[40] Honeywell's offering was the Secure Ada Target (SAT) and Logical Coprocessor Kernel (LOCK). The division of Honeywell that produced the system was spun out as Secure Computing Corporation, whose main secure system offering today is a firewall known as Sidewinder.

secure databases for both the military and commercial sectors. It is capable of providing the highest government security classification, Class A1, when implemented over a Class A1 operating system kernel.

The third and final area of computer security to which CSL has made notable contributions is in the real-time detection of unauthorized computer or network access. In comparison to the upfront designs mentioned earlier, proposed solutions for this problem, intended to fit over a wide range of multiple-access systems, are still very complex but of a different type. This complexity has at least two dimensions: (1) the variation in the hardware and software systems that are already in existence but need an overlay of access monitoring and control, and (2) the variety of procedural models for access authorization within the normal office or company workflow.

CSL's work on real-time access detection began with a project in 1983 for a nameless government agency that wanted to explore the computing wanderlust of internal as well as external people. The solution involved statistical profiling of users to enable a system to sound an alert whenever significant departures from those profiles were detected. This work was also conceived and led by Denning. Building on this effort and the advent of expert systems technology, SRI designed a new scheme called Intrusion Detection Expert System (IDES). The easily expressible rules of that technology made it adaptable to a wide variety of administrative procedures in which security policy was expressed. It was intended to run in the background of host computers, monitoring the access to important open processes that defined machine operation. The intention was to detect unusual behavior in the machine that might tip off the presence of unauthorized access. From this monitoring IDES would also adaptively learn normal user behavior patterns to provide a basis for the statistical detection of abnormal events. Such real-time insights could both record relevant process information and alert human operators. The rule-based nature of IDES could also permit easy modification as experience with the program suggested the need.

Work on IDES continued for almost 9 years under sponsorship of the Navy and Air Force. A second-generation system called NIDES extended the audit-record capability of IDES and improved interfaces for a security officer and the maintainer of the rule-base. In related efforts NIDES was used to extend real-time observations to the profile software applications rather than just user descriptions. This program was installed for service testing in several machines around the United States, including an experimental phase on the main computers of the FBI.

The current and state-of-the-art embodiment of unauthorized access detection programs at SRI is EMERALD (a somewhat stilted acronym for Event Monitoring Enabling Responses to Anomalous Live Disturbances).[NN] The EMERALD development program, created and led by Phil Porras, applies many of the techniques developed in NIDES to the area of communications networks. SRI was issued a patent in 2002 on the underlying event-monitoring technology in EMERALD and is now seeking to package it in a form releasable as a commercial product.[41]

CSL also created one of the first means to establish private links between arbitrary points on the Internet, requiring only that the end terminals have the appropriate encryption and other security software. This system, called ENCLAVES, was demonstrated in 1994–5. This general capability is now quite common under the name of virtual private networks.

One additional CSL effort in the area of security and reliability should be mentioned. The burgeoning of computer use is bound to have consequences that are not all beneficent. Everyone experiences negative impacts that, given human nature, get more attention and linger longer in the memory than the positive ones. One member of the CSL has done more than any other person anywhere to chronicle those unfortunate occurrences: Peter Neumann, who joined SRI in 1971. Neumann has been near the center of the research, advocacy, and controversy surrounding computer security ever since. In about 1985, as a personal endeavor and under the auspices of the ACM,[42] he began what has come to be called the Risks Forum, an online, broadly participatory collection of computer mishaps. Because the foibles of computer use touch so many lives, the Forum became one of the most popular sites on the

[41] U.S. Patent No. 6,484,203 covers computer-automated hierarchical event monitoring and analysis within an enterprise network, including the deployment of monitors. It also covers detection of events through the analysis of a variety of network traffic data.
[42] The Association for Computing Machinery.

new Internet. Thousands of contributors, from disgruntled neophytes to Nobel laureates, have eagerly related their frustrations or their observations on computer misuse or mishaps. The collective account in the Risks Forum comprises perhaps the single most comprehensive anthology of computer-related mishaps in existence. As such it is also the best harbinger of the risks we face in future computer use.[OO] Though these Neumann archives are sometimes overly inclusive, thousands of relevant accounts are given: from the pranksterish commandeering of the Rose Bowl scoreboard by California Institute of Technology students in 1984[PP] to countless, more catastrophic events such as the illegal 1989 intrusion into the California Department of Motor Vehicles computer to retrieve the address of an eventual murder victim. In spite of all this collected awareness, progress toward achieving security in the rapidly emerging information infrastructure is still slow.

One addendum to this review of computer and system security centered in CSL actually lies nearly totally outside it. The business consulting groups in SRI began to conduct security evaluations of companies in about 1970. By spring 1987, largely under the leadership of Donn Parker, some 85 companies had been reviewed and some 1,600 cases of computer abuse had been logged. Parker has also written several books on computer crime and its prevention. That experience base enabled SRI to begin at that time a large multiclient service it called the International Information Integrity Institute (I-4).[QQ] "Integrity" has a broader implication than security alone, including such mishaps as accidental destruction or misplacement. Focus on that larger topic has had broad appeal, and I-4 continues to this day. However, as a result of some organizational turbulence involving the demise of SRI's Business Group,[RR] it is now operated by a separate company housed on the SRI campus.

The Ability of Machines to Listen

Humans have a propensity to talk, and speaking is a more developed need, perhaps, than listening. Speech is overwhelmingly our preferred mode of interacting with other people. But that entire predilection is interrupted when we face a machine. For centuries we have not been able to expect much in this regard from a machine, but that expectation has already changed. Machines in the new class we loosely call computers are creeping more deeply and completely into our lives. Our somewhat misguided tendency to grant them anthropomorphic skills is abetted by their extraordinary processing capacities, and they will become better and better at listening to us. [43]

Whether you love it or hate it, the recognition of speech by machines is now with us to stay. Though computers that recognize speech may often seem impersonal and unresponsive, because they save their owners a lot of money, time, or convenience, they will only become more prevalent. Like most technical achievements, the arrival of speech recognition seems much more rapid than was actually the case. Finally, for reasons that will become evident, this section could be subtitled "The Benefits and Risks of R&D Commercialization."

SRI's Entry into Automatic Speech Recognition

Perhaps the first time serious money was expended for research on automatic speech recognition was in 1971 when DARPA funded a program of Speech Understanding Research (SUR). Cordell Green, a former member of SRI's Artificial Intelligence Center who had been called to active duty, initiated the program at DARPA. About half a dozen organizations, including SRI, participated in the research to learn whether a computer could recognize perhaps a thousand words, spoken by a limited set of speakers, under calm and quiet conditions. At the end of 5 years several of the recognition programs were tested, but under some unanticipated rules; that is, rules not established at the outset of the research. None measured up, and SUR was terminated.

SRI also made an excursion into speaker identification in the early 1970s. For the Law

[43] This section has been modified from the first printing with the help of Patti Price, Mike Cohen, and Hy Murveit.

Enforcement Assistance Agency of the Department of Justice, SRI looked into the use of so-called voiceprints for identifying the person speaking.[SS] SRI built a prototype system and acted as an advisor to Rockwell when it received a contract to produce a commercial version of the SRI prototype. While the system was semiautomatic and worked in some uses, it had to be operated by a trained person, a phonetician.

Speech research resumed at DARPA in 1984 as part of an ambitious artificial intelligence (AI)-centered program known as the Strategic Computing Program. Many of the SUR research centers, such as Carnegie Mellon University (CMU), Bolt, Beranek, and Newman (BBN), SRI, and MIT renewed their efforts under this new program, along with both industry and European researchers. This time evaluations were conducted each year and the better techniques were quickly propagated among the participants. This effort continued into the late 1990s and was pivotal in the development of commercial-quality speech recognition systems that have made their way into the marketplace.

In this second round of funding SRI became one of the top-tier organizations developing what would emerge as unprecedented and successful recognition techniques. Whereas in the earlier SUR program SRI's Artificial Intelligence Center (AIC) worked with SDC to build a recognizer based largely on language properties, the second speech recognition effort was centered in the Sensory Sciences Research Laboratory (led at the time by Hew Crane), which had no AI expertise at all. Before 1984 the lab's research interest in speech had concentrated on the exploration of the physical models of the vocal tract and other elements of speech phenomenology.

As it turned out, the second technical direction that automatic speech recognition took didn't favor an AI approach anyway, at least initially. SRI's approach became much more statistical and followed an approach called hidden-Markov modeling (HMM). Here word recognition is based on the previous one or more words spoken and the likelihood associated with various word sequence transitions. As computer power increased, this approach enabled greater word-transition flexibility and thus ever-greater vocabularies.[44] Prosody, the intonation and duration patterns of speech, was also modeled. The speed and simplicity of this overall approach, plus the advent of more powerful workstations, led to the commercial systems in use today.

The Second Venture into Speech Recognition

The leader of the second SRI speech program was Don Bell. Beginning with the arrival of Dr. Hy Murveit from the University of California, Berkeley, Bell gathered a world-class group of researchers in the Sensory Sciences Lab. Their innovations in the intricate algorithms enabled speaker independence, vocabularies that climbed toward 100,000 words, continuous uninterrupted speech, and tricks for computational acceleration helped them perform repeatedly at or near the top of the annual competitions that DARPA held to promote the success of speech recognition technology. The difficulty of meeting DARPA's goals is hard to overestimate, and the climb was not easy. It would take about a decade before the collective DARPA community would produce sufficient recognition accuracy for a usable product.

The development of a more accurate recognizer required that a number of linguistic and speech modeling research questions be answered:

- How could models of grammar, words, phonemes etc., be efficiently and robustly combined?
- Could natural language grammars help constrain word sequences and still handle the ungrammaticalities of spontaneous speech?
- Would the successes in English generalize to other languages?
- Could other pattern-matching techniques or hardware accelerators improve speed and/or accuracy?

To determine not just the words spoken but also their meaning, SRI's AIC collaborated

[44] The term for this difficulty in speech recognition tasks is perplexity, and its definition is related to entropy from information theory. It is approximately the geometric mean of the word branching factor after the application of a language model.

with the Sensory Sciences Lab in a new effort integrating speech and natural language. The natural language (NL) understanding research had gone on in the AIC for years, but had to be modified considerably because people don't speak the way they write. SRI's resulting system was developed in the air travel planning domain and was called ATIS (Air Travel Information System). The other DARPA-sponsored sites then adopted ATIS as a benchmarking task. One innovation of the new technology was a "progressive search" that allowed the same system to use the most detailed contexts (slow but accurate) or to back off to fewer, less detailed models when speed was necessary. SRI won the last ATIS competition held by DARPA, and Hy Murveit bought trophies for the SRI team.

By late 1992 the SRI HMM recognizing program, DECIPHER, was working well enough that Bell began to explore its commercial utility. In that regard, he consulted several specialists and began to look at the market in a number of speech recognition applications. In 1993 a new lab director, Patti Price, selected one of those consultants, Ron Croen, to help continue the quest for commercialization, including licensing.[45]

But the overly optimistic promises and subsequent failure of earlier recognition systems made for dictation had soured the commercial world on the subject, and SRI was unable to sell a license to the technology. A few small companies were already in the business, but the new DARPA technology was not yet in commercial practice, although it soon would be.

Then, a number of researchers at the Lab, including Murveit and Dr. Mike Cohen, began to assert that the best way to commercialize SRI's speech technology was by starting a company focused on telephone-based applications.

The Commercialization

Also at this point, and totally coincidentally, new upper management at SRI decided that commercialization of technology was to be important, perhaps dominant, in SRI's future (see Appendix B), and the Lab and division-level entrepreneurial efforts got caught up in this new corporate emphasis. Consequently, after reviewing the results of a market study or two, the chairman of the SRI Board, Paul Cook, decided in late 1994 that SRI should start its own company in speech recognition. For the first time in its history, SRI invested serious money, $1 million, to start a commercial enterprise, the Corona Corporation, to which SRI transferred:

- A worldwide exclusive license to its core technology in speech recognition, including demonstration systems based on DECIPHER
- A toolkit for the rapid development of commercial applications
- Four key personnel
- An obligation to transfer future research results exclusively
- A reputation for high-quality benchmark results in the DARPA tests
- The commitment of time and personnel to achieve the transfer of the technology.

Three key technical personnel, Murveit, Cohen, and Peter Monaco, together with Ron Croen, founded the new company with Croen as president. Cohen continued half-time at the SRI lab for the first year to help with technology transfer and to ease the transition. (The turmoil of the spinout resulted in the loss of two or three other people who chose neither to remain at SRI nor to join Corona.)

The founding team moved into a separate building on the SRI campus and carved out a new area for speech recognition: call-center telephony. In 1994, speech recognition was nearly exclusively used for dictation. The call-center strategy allowed the fledgling company to grow a new market without going head to head with IBM and Dragon Systems, who dominated the dictation market. Ten years later, it is call-center applications that now dominate speech recognition.

In its first year the new company gained first-round financing from the Peninsula's ripe venture capital world, increased in size, and changed its name to Nuance Communications, Inc. Subsequently, second- and third-round financing became available, the company moved off the SRI campus, and

[45] Dr. Price also reorganized the lab as the Speech Technology and Research Laboratory (STAR), which is its present name.

as of this writing in 2004 has perhaps 300 staff members.

To gauge the quality of the recognition software that Nuance produced, consider its first major client and deliverable. SRI's Business Group had persuaded Schwab Discount Brokerage to examine the ability of automatic speech recognition to reduce the cost of Schwab's telephone quotation system. SRI, Nuance, and Schwab explored and eventually built the first over-the-phone stock quotation system. As a measure of its performance, it could (and still can) respond automatically to oral requests for any of 11,000 stocks listed in Schwab's catalog, which can be expressed in 3 million different ways, with an accuracy exceeding 95%. By the end of 1996 this system was saving Schwab about $1 million per month. Since then Nuance has provided similar systems to a dozen or more other commercial clients. More recently, the company has gone public and become one of the leading suppliers of speech and speaker recognition systems, speech synthesis, and development tools.

While Nuance represented a commercialization success for SRI, the STAR Laboratory, which had made it all possible, faced a struggle. Spawning Nuance had some severe negative consequences that became an object lesson for subsequent spinouts:

- In the formation of the spinout or its associated turmoil, the lab eventually lost all but one of its senior staff and several other talented engineers to Nuance or other employers.
- The SRI license to Nuance was so broad and exclusive that it prevented the lab from engaging in any commercial marketing with the exception of nonexclusive licensing in language training. With very few exceptions, only projects with government sponsorship were available after the license, although the lab could use Nuance technology in the nonexclusive area.
- The commercial outlet for the STAR Lab as the long-term research arm for the new company never materialized.
- Some new technology, such as speaker identification, noise mitigation, and recognition models for other languages, continued to flow from the lab, with no remuneration for SRI, and most often with the continued loss of people as the technology was transferred to Nuance. In the years after the spinout, the president of SRI at the time and his equity manager were focused on making their first large commercial investment, and dismissed attempts to mitigate internal difficulties the spinout process created.[46]
- The long-term drain of people from the Lab to Nuance became an ongoing challenge, abetted by the government's continuing de-emphasis of speech research. Although new Lab leadership and a new user interface program were formed, retaining personnel became more difficult. Ironically, the license's foreclosure of new commercial opportunities also became a disincentive to some new Lab hires.

To continue to maintain a vital group, Price focused the Lab on the language education carve-out, an area initiated much earlier by Jared Bernstein. She convinced Broderbund, a northern Bay Area software company, to partner with SRI in a government-sponsored dual use (government and commercial) Technology Reinvestment Program (TRP) in language education.

Broderbund had shown early successes in some of its educationally oriented products. But this arrangement faltered badly when, halfway through the several-year product development, Broderbund abandoned its role in the program. Strangely, in a downsizing intended to support the company's stock price, its leaders showed they were oblivious to the existence of Broderbund's obligations under the alliance even though it involved millions of dollars. The Lab, however, was able to revise the project toward an even stronger result: rather than one speech-enabled language education product with Broderbund, the effort would produce tools to make it easier to speech-enable language education software by the government and by anyone licensing the tools. The resilient

[46] That indifference to the technology-generating laboratory in the course of commercialization is now completely gone under the more equitable policies of the present SRI administration.

Laboratory also launched a special recognition engine for computer-based education tools as the term of the Nuance license agreement was expiring. Under the current lab director, Dr. Kristin Precoda, the STAR Lab has also found a new area in which to excel: speech translation. The war in Afghanistan created a desperate need for soldiers to be able to speak to Afghans. The lab has built bilateral translation software that effectively translates one of the leading Afghan dialects, Urdu, into English. Under Precoda's and Dr. Horacio Franco's leadership, about a dozen stalwarts still offer world-class speech technology.

Regarding Nuance, it initially concentrated on telephone-based speech recognition applications. But at one point it broadened its efforts to include voice-activated Web interaction. Because of the ubiquity of the telephone and the burgeoning use of cell phones, Nuance foresaw a speech-enabled, browser-like interface that could become not only a means to retrieve information from the worldwide complement of online sources, but also a communications device that can enhance all telephone voice traffic. The company calls this product Voyager, and the telephone-accessible part of the World Wide Web, the "Voice Web." Nuance has since retreated a bit from this initiative to return to its most successful area, voice-based servers and their accompanying service and support.

All in all, Nuance Communications has done quite well in building a leading position in the telephone-based speech recognition market. It has delivered, both directly and through a growing number of industrial partnerships, a large number of recognition and speaker identification systems worldwide.[47] While SRI's stake in the commercialization of Nuance has yielded a large financial return, it came at the expense, perhaps unnecessarily, of diminishing SRI's research base for at least the duration of the license agreement. The STAR Lab and the Nuance spinout, two different worlds, with very different objectives and reward systems, have intersected with important consequences.

Watchwords for Commercialization at SRI

Given that SRI remains a nonprofit research institute, it is useful to recall that its commercialization ventures are intended first to strengthen the research base and, if possible, to enrich SRI's compensation package so it can compete successfully with other recruitment efforts in Silicon Valley. But some serious questions need to be answered to safeguard SRI's interests in the commercialization process:

- Does the technology to be commercialized represent most or all of that practiced by the inventing lab?
- What are the consequences if key personnel of the lab participate in the startup at the outset? What happens if they stay with the startup?
- Does the license leave the lab with adequate latitude to pursue continued funding in its research specialty?
- Does the license permit, encourage, or require an ongoing relationship between the lab and the startup, one that is beneficial to both?
- Do the managers with authority over both the lab and the startup have the best interests of both organizations in mind?

These questions are obviously posed from the perspective of the research organization. Left on its own, the startup will, by its nature, have little concern for the continuing success of the enabling laboratory. It has investors to satisfy. And if the lab continues to be technically strong and advances the state of its field beyond the technology licensed to the startup, then new licensing ventures with others can become aggravating competition that threatens the startup. How is this thicket to be negotiated?

The answer to these questions depends on why SRI wants to become an incubator of technology that can be commercialized and what fraction of its operation that effort represents. The announced reason for such efforts is, in the long run, to strengthen the research base at the Institute and thereby provide greater benefit to the public interest it is chartered to serve. That cause, derived from its nonprofit status, means that all excess

[47] In 2005 Nuance merged with competitor ScanSoft to form a more powerful market presence.

income is reinvested and in some instances into labs that have the potential for commercializing their output. It is hoped that this cycle of innovation would grow with increasing ability to help all parts of the Institute.

If the work of an inventing lab is narrowly defined and the intellectual property to be commercialized cannot avoid encompassing everything that the lab and its staff members do, care is needed to design a licensing agreement that leaves the lab viable. Possible options are to 1) restrict the field of use of the licensed technology, 2) restrict the time of exclusivity, and 3) require a flow of funds back to the lab for longer term R&D.

If none of these are possible, then the lab personnel should all be given the opportunity to transfer to the licensee, and SRI should abandon work in that area rather than leaving researchers without continuing opportunities.

If the technology to be licensed is only a part of a laboratory's work, then the license conditions can be more flexible, but care must be taken to preserve the lab's ability to remain viable. If researchers do transfer, then the structure of the equity offerings to those that leave and those that stay should be sufficiently equal to make it difficult to decide which path to take. This approach seems fair, and SRI tried to follow it as Nuance was being created. It is now clear, however, that persistently equitable arrangements can be difficult to achieve.

Endnotes

[A] Some of the material presented here is drawn from two sources: Fisher and McKenney, The Development of the ERMA Banking System: Lessons from History, *IEEE Annals of the History of Computing*, 15(1), 1993; James L. McKenny, *Waves of Change*, Harvard Business School Press, 1995.

[B] William D. McGuigan, personal communication, May 27, 2000.

[C] SRI's *Research for Industry*, 7(9), October 1955.

[D] Tom Morrin, personal leter to Prof. James L. McKenney, Harvard Business School, June 26, 1992.

[E] SRI's *Research for Industry*, 7(9), October 1955.

[F] Some of this story was taken from autobiographical material of Hew Crane entitled: "*Reflections on a 50-Year Career*" dated October 12, 1997.

[G] H.D. Crane, "Sequence Detection Using All-Magnetic Circuits," *IRE Trans. on Electronic Computers*, Vol. EC 9, 155-160, June 1960.

[H] Earle Jones, personal communication on August 2, 2004.

[I] For the most recent and complete record, see Thierry Bardini, *Bootstrapping*, Stanford University Press, 2000.

[J] Douglas Engelbart and Harvey Lehtman, Working Together, *BYTE*, page 245, December 1988.

[K] Vannevar Bush, As We May Think, *The Atlantic Monthly*, 176 (1), 101-108, July 1945.

[L] Ken Jordan, "The Click Heard Round the World," *Wired Magazine*, 12(1), January 2004.

[M] Stanford oral history interview found at www-sul.stanford.edu/depts/hasrg/histsci/ssvoral/engelbart/engfmst2-ntb.html

[N] Ibid.

[O] See Chapter 3 on Reid Anderson's own entrepreneurial exploits.

[P] Charles R. Bourne and Douglas C. Engelbart, Facets of the Technical Information Problem, *SRI Journal*, 2(1), 2-8, 1958.

[Q] SRI Proposal No. EU 60-251 to the Air Force Office of Scientific Research, December 13, 1960. The contract technical monitor was Harold Wooster.

[R] Vannevar Bush, op cit.

[S] Research Aimed at Computer Augmentation of a Programmer, SRI Project 4385, Contract No. SD-163, $195,000, February 1963 to February 1964.

[T] Computer-Aided Human Control of Computer Display, SRI Project 5061, Contract No. NAS1-3988, $85,626, June 1964 to July 1965.

[U] Personal email communication from Jack Goldberg, November 15, 2001.

[V] W.K. English, D.C. Engelbart, and Bonnie Huddart, *Computer-Aided Display Control*, SRI Report to NASA, July 1965.

[W] Adele Goldberg, ed., *A History of Personal Workstations*, 187–232, Addison-Wesley for the ACM Press, 1988.

[X] Douglas C. Engelbart and William K. English, A research center for augmenting human intellect, *AFIPS Conference Proceedings - Fall Joint Computer Conference*, San Francisco, December 9–11,1968.

[Y] Ken Jordan, op. cit.

[Z] Thierry Bardini, op. cit. Contains extended coverage of this transition of ideas and the attainments at PARC. The notion is also mentioned in Adele Goldberg, op. cit.

[AA] Ibid.

[BB] Michael Kanellos, CNET News.com, September 18, 2003 or www.logitech.com as of July 9, 2004.

[CC] Richard Watson, personal communication, December 9, 1998.

[DD] Implicit in the "transparent pipe" mentioned in: Elmer Shapiro, *A Study of Computer Network Design Parameters*, SRI Project for ARPA, December 1968.

[EE] Hewitt D. Crane, A High-Speed Logic System Using Magnetic Elements and Connecting Wire Only, *Proceedings of the IRE*, 47(1), 63-73, January 1959 and D.R. Bennion, H.D. Crane, and D. Nitzan, *Digital Magnetic Logic*, published by McGraw-Hill, 1969. Also from 1997 autobiographical material of Hew Crane entitled: "*Reflections on a 50-Year Career*."

[FF] Hew Crane, personal communication, January 24, 2002. The sales estimate was based it on material Crane received from retired AMP Corporation vice president, Joe Sweeney.

[GG] P.G. Neumann, J. Goldberg, K.N. Levitt, and J.H. Wensley, *A Study of Fault-Tolerant Computing*, SRI Final Report to ARPA and the Office of Naval Research, July 31, 1973.

[HH] J.H. Wensley, Leslie Lamport, Jack Goldberg, Milton W. Green, Karl N. Levitt, P.M. Melliar-Smith, Robert E. Shostak, and Charles B. Weinstock, SIFT: Design and Analysis of a Fault-Tolerant Computer for Aircraft Control, *Proceedings of the IEEE*, 66(10), October 1978.

[II] Jack Goldberg, email communication dated July 30, 2004.

[JJ] L. Lamport, R. Shostak, and M. Pease, The Byzantine Generals Problem, *ACM Transactions on Programming Languages and Systems*, 4(3), 382–401, July 1982. According to Goldberg, this paper set off a rash of technical papers and interchanges that lasted for more than a decade.

[KK] Daniel L. Palumbo and Ricky W. Butler, A Performance Evaluation of the Software-Implemented Fault-Tolerant Computer, *J. of Guidance, Control, and Dynamics*, 9(2,) p. 175, March–April 1986.

[LL] Peter G. Neumann, Robert S. Boyer, Richard J. Feirtag, Karl N. Levitt, and L. Robinson, *A Provably Secure Operating System: The System, Its Applications, and Proofs*, SRI report to the National Security Agency, May 1980.

[MM] T. F. Lunt, D. Denning, R. R. Schell, M. Heckman, W. R. Shockley, The SeaView Security Model, *IEEE Trans. on Software Engineering* (TOSE), 16(6), 1990.

[NN] P.A. Porras and P.G. Neumann, "EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances," *Proceedings of the Nineteenth National Computer Security Conference*, Baltimore, MD, 353-365, October 22-25, 1997.

[OO] Peter G. Neumann, *Computer-Related Risks*, Addison-Wesley/ACM Press, 1995.

[PP] According to the Computer Risks Forum, the radio control of a computer, previously spliced into their scoreboard control circuits, caused the UCLA-Illinois game score to be replaced with a CalTech-MIT score of 38 to 9.

[QQ] I-4 Is Underway, *SRI Journal*, 7(1), Spring 1987.

[RR] See Appendix B.

[SS] SRI's *Investments in Tomorrow*, No. 9, Autumn 1973. The voiceprints, called "sonograms," essentially plotted acoustic energy as a function of frequency and time.