Oral History of Lester D. “Les” Earnest

Interviewed by:
Dag Spicer

Recorded: November 26, 2012
Mountain View, California

CHM Reference number: X6709.2013

Dag Spicer: Well, it’s November 26, 2012, and we’re here with Les Earnest from Stanford. Welcome, Les. We’re delighted to have you here.

Lester (Les) D. Earnest: Thank you.

Spicer: I wanted to start out at the beginning and learn a bit about your childhood and your early influences. What did your parents do, for example, and those kind of things?

Earnest: Well, my dad started out as a radio engineer, self-taught. And then he went to San Diego State University, where he met my mother. Well, it wasn’t San Diego State University then, it was some normal school. But they hit it off and I was the result. My dad promptly came up the chain and at the beginning of World War II was made general manager of a defense plant, first in Des Moines, Iowa, then in Louisville, Kentucky. So he moved there during the war and came back to my hometown, which is San Diego. And my dad became manager of the parks and recreation department for San Diego. [He] initiated the development of the Mission Bay Park, which was quiet a large undertaking. My mother was a schoolteacher initially. But while raising two kids, she earned a Ph.D. in her spare time.

Spicer: Wow. Amazing.

Earnest: And then became a college professor. Very popular one. Matter of fact, San Diego State has a website devoted to her. And her former students still talk about her a lot. So that’s once over lightly. I grew up as a bicycling beach boy. Made frequent trips to go body surfing. Got into a fair amount of trouble along the way, but had a good time. I eventually got into Caltech on a scholarship. Then ran wild and flunked out. <*laughs*>

Spicer: Now, were you more technically-oriented, or into the humanities in your teenage years?

Earnest: I was thinking, well, I was influenced by my freshman chemistry teacher at Caltech, who was Linus Pauling.

Spicer: Wow.

Earnest: And I signed up to be a chemist only to discover I wasn’t really cut out for it. I was really a lot more interested in electronics. Course, there were no computers to speak of then. So I changed direction and went that way. Then I was having trouble with my draft board. As a result of a clerical error,

I was no longer classified as a student but as 1-A, and they eventually drafted me in my senior year, halfway through.

Spicer: Now, 1-A is the most eligible category?

Earnest: Exactly.

Spicer: Yeah.

Earnest: So fortunately, my dad had enough political clout to get that reversed, which was pretty challenging. But I knew I had to serve my time. And we were in the Korean War. I didn’t relish the idea of going into the Army. I found a program for which I was ideally suited. You had to have an engineering degree and poor eyesight.

Spicer: <*laughs*>

Earnest: I qualified. And ended up as an aviation electronics officer. And in a research lab in suburban Philadelphia, which sort of set my career from there on.

Spicer: And how did it do that?

Earnest: Well, they had, for one thing, the world’s largest and fastest analog computer called Typhoon.

Spicer: Hum.

Earnest: But it needed check solutions from a digital computer. And I was more interested in digital. The computer we had was an IBM card-programmed calculator, the so-called, CPC. Which was really an electromechanical machine, but it did have an electronic arithmetic unit. And we used that to simulate aircraft flight, aircraft and missile flights, as check solutions on the big machine.

Spicer: Were these trajectories you were doing?

Earnest: Yeah.

Earnest: With various kinds of guidance systems, and seeing which ones worked.

Spicer: Right.

Earnest: That sort of thing. I ended up spending three and a half years in the Navy. And when I got out, I had a lot of choices for jobs because I was one of the very few people who was an experienced programmer. There were only probably a couple hundred of us in the world then. That was 1956. And looking around it seemed to me that the hottest technology at that time was being developed at MIT. They had just invented the magnetic core memory, which revolutionized things for a number of years. And then of course disks had been developed for the first time by IBM a short time earlier. So things were coming together. And MIT was developing the so-called SAGE air defense system.

Spicer: Right.

Earnest: SAGE stood for Semi-Automatic Ground Environment.

Spicer: Mm-hm. Sorry. From the Navy, where did you go? Right after the Navy?

Earnest: To MIT.

Spicer: To MIT. Okay.

Earnest: MIT Lincoln Lab in particular.

Spicer: Okay. Right.

Earnest: Which was in Lexington, Massachusetts. And that was a lot of fun, in a way. It was really high-tech stuff. It involved computer networking, very complex, six levels, of computer networking. And my responsibility was weapons integration, since I had come out of a military air situation. It put me onto developing guidance systems for managing interceptors and Bomarc missiles. And also I designed the console interface for intercept directors, the people who manage these interceptions.

Spicer: Mm-hm.

Earnest: However, early on I learned that the system didn’t really work. My office mate, when I arrived, was working on processing radar data. So I promptly asked him what they did about jammers, bombers that did radar jamming, which produced hundreds of returns from a given direction, at all different

distances. And his answer was, “We don’t do that.” I thought he was kidding. <*laughs*> Because by that time, all bombers were equipped to jam.

Spicer: Right.

Earnest: But I discovered that they really didn’t deal with it. They pretended that jammers didn’t exist. And they set up all of their tests without jamming. So that it would appear to work. And they kept the deficiencies of the system classified so that the public, and they hoped, Congress, would not know. And kept it going for 25 years, to the enormous profit of the contractors involved. Which I found rather disturbing. But there wasn’t much I could do about it.

Spicer: How could something so basic be swept under the rug by so many people?

Earnest: Well, MIT’s original motivation, I believe, was they had a solution looking for a problem, namely a computer project that had originally been funded by the Navy to flight simulations for training. That was the original purpose of the Whirlwind computer that they had developed. But they eventually concluded that it couldn’t do that job very well. And the Navy decided to terminate their funding. So they had this really hot computer project, and it was about to die.

Spicer: This was the Whirlwind?

Earnest: Yes.

Spicer: Okay.

Earnest: So they put together a committee to look for potential applications. And considered a number of them. One of them that they looked into was air traffic control. It could’ve done a good job at that, I believe, but that part of the government wasn’t buying.

Spicer: Right.

Earnest: And the Air Force was interested, but they surely understood that pretending jamming doesn’t exist <*laughs*> is not sensible. But they went along with it. I do not understand the politics of that whole situation. But here we had a—well, but it ended up over 25 years, I estimate, costing hundreds of billions of dollars in modern currency for a system that would have been about as effective as the Maginot Line was in World War II.

Spicer: Right. Now, let me play the devil’s advocate here and say could this system have worked in a political sense, even if it were a technical boondoggle?

Earnest: You mean if we had convinced the Soviets—

Spicer: That it worked and that it was a credible defense.

Earnest: <*laughs*>

Spicer: Did they assume it worked?

Earnest: I don’t think so. They knew about jamming.

Spicer: Right.

Earnest: And I would be surprised. There’s no way to tell exactly what was going on in there. I do recall seeing some intelligence reports in that era claiming that Soviets were building a very similar system. Those were pretty thoroughly lies designed to support the construction of SAGE.

Spicer: Another gap—

Earnest: Yeah.

Spicer: —that we had to be worried about.

Earnest: Right.

Spicer: Supposedly. Or a fictitious gap.

Earnest: Yeah.

Spicer: Yeah.

Earnest: So anyway, I did my part. It worked even though the tracking wouldn’t really. Radar tracking. And I eventually was, well, I might mention that one of the last things I did, I was assigned to obtain permission to use nuclear warheads on Bomarc missiles. Now, that was fundamentally an insane idea. That is shooting off nukes in our own skies in order to take out bombers coming in.

Spicer: Yeah.

Earnest: Nevertheless, we succeeded in getting permission. The fundamental requirements were that we had to prove that the probability of a missile being launched by accident on a given day was less than a very small number. And also had to prove that it would take at least two co-conspirators to do it. And we were able to show that sort of. Under hindsight, we overlooked something. This is with 50 years’ hindsight. We did not consider the possibility that a malevolent programmer would do it, could do it. Now, if somebody had tried and was a good chance it wouldn’t have worked because a number of people scrutinized the code, and the chance it would’ve been caught, but that’s not a certainty. And as the internet has since taught us, there are malevolent programmers. There weren’t back then. We all thought programmers were all honest, upright and altruistic. That was why we never considered that possibility.

Spicer: That feeling lasted until well into the Unix era too, I think, until maybe the PC and the internet certainly—

Earnest: Yeah.

Spicer: —changed that. Yeah. Hmm. Tell us a bit. I read a story about a colleague of yours that did an analysis of the likelihood of the launcher erecting and firing. Do you remember, there was random noise? If both formats were—

Earnest: Yeah. Okay. Well, that was an outgrowth of our study on nuclear warheads. One guy in my group noticed that there were two phone lines that were used to send launch commands from the central computer to the missile site. And there was a black box at the missile site that could detect if the primary line went bad. It would switch to the backup. However, if both lines went bad, it would stay connected to the backup. And we figured out it would then amplify whatever noise was on the line and produce a stream of random bits. So one of the guys in my group then did a Markov analysis to figure out on average how long it would take to see a fire command for one of the missiles at the site. And it was something like two and a half minutes.

Spicer: That’s frightening.

Earnest: Yes.

Spicer: <*laughs*>

Earnest: However, he went on to show that would cause the missile to erect and prepare to fire.

Spicer: Right.

Earnest: But in order for it to fire it had to get a full set of commands, heading, speed, altitude, so on. He was able to show that the probability of that happening for that missile within the five minute or so window was very improbable. So that effect would be the missile would erect and then abort. So I then wrote a report called “Inadvertent Erection of the IM99A.” And as luck would have it, couple weeks later it actually happened at one of the Bomarc sites, and scared hell out of everyone.

Spicer: Wow.

Earnest: However, that, as I recall, was not made public. And I was assigned to share a committee to fix it, which was pretty easy. So that problem sort of went away.

Spicer: Tell us a bit about modems, because you’re right at the era, and you’re working on the systems from which modems developed—

Earnest: Right.

Spicer: —SAGE being one.

Earnest: Well, the first modem was developed by a guy named Jack Harrington, who was then at the Cambridge Research Center of the Air Force, I believe it was. What he was trying to do was not to connect computers together but to make radar data remotely visible. The bandwidth of radar returns was too high to go over phone lines, so what he did was made a little digitizer that would recognize a blip and generate a packet, digital packet, that told where it was, that could then be sent over a phone line to a remote display. But then he subsequently moved to MIT Lincoln Lab to work on SAGE. And this idea was expanded to being into computer networking. And there were a bunch of different networks. It also could pass words that were misspelled. That used a packet radio system, was one way. And then there were network connections between adjacent direction centers, that is different SAGE centers, talked to each other about flights that were near the border, between them. And they also then forwarded information, general information, about what’s going on to the so-called command centers, a higher level

computer made out of the same computer hardware as the direction center. And then ultimately, they communicated to NORAD headquarters, which was buried in Cheyenne Mountain near Colorado Springs, Colorado.

Spicer: Right.

Earnest: And so there’s whole lot of networking, all packetized. And these were all so-called star networks, that is all the connections came together at a central computer, which then coordinated the switching and all that.

Spicer: Right.

Spicer: So then this is I guess around the late ‘50s we’re talking about, SAGE, early ‘60s?

Earnest: SAGE went more or less operational in 1959.

Spicer: Okay.

Earnest: Yeah.

Spicer: Now, at the same time, John McCarthy is beginning to think about timesharing?

Earnest: Yes. Well, he saw SAGE, and he decided he wanted to have something like that for his research in artificial intelligence, in particular for development of the Lisp programming language. And so he first set up an experiment in which a regular batch processing system had one interactive terminal connected to it so that you could interactively program and debug while the batch stream was still going on in the background. That was, I forget, using a 709 or maybe a 704, IBM machines. But he also wrote up his idea for timesharing, general-purpose timesharing. SAGE was a timesharing system of a sort. They built it as a real-time system, which it was. But since the processing of radar data was interleaved with processing of interactive work with the people running it, it functioned in effect as a timesharing system and people really liked that idea that you could get responses in a few seconds to actions you took, instead of hours to days, which is the way the batch processing systems of that era worked. So a number of people talked about the idea of making interactive general-purpose systems. I was one of them. But we didn’t figure out a way to do it. McCarthy did. And that, in turn, inspired a bunch of people in the MIT community to try to build a system like that. And in short order, there were four or five timesharing systems going around there.

Spicer: Now, if I remember correctly, SAGE had no interrupts. It was a single program that ran through and gave time slices out to various people and processes.

Earnest: That’s correct. It was one big loop that took about two and a half seconds, maybe three seconds, to finish one pass. So that was the approximate response time. And, of course, it had to be error-free. If there was anything wrong with it, the system would crash. And we got stuck in some interesting loops. But a lot of work went into making that not happen.

Spicer: So tell us about the first spell checker and connect the dots from when you left MIT, to when you began working on spell check?

Earnest: No. Well, by that time I was working for MITRE Corporation, which was a spin-off from MIT, a nonprofit that had been set up essentially to get MIT out of the SAGE business. I’d like to think that this was a decision on their part to get out of the mess that they had sort of created.

Spicer: Was it just too big a project for MIT? And were not the right—

Earnest: Well, I don’t know. I was not a party to that decision.

Spicer: Hmm.

Earnest: But I think they figured out that it didn’t work, and it was a giant fraud. They wanted to get out of it.

Spicer: Right. Hmm.

Earnest: I’d like to think that was why they made that decision.

Spicer: Right. Because there was a similar tension with John von Neumann when he tried to build the IAS machine, at Princeton. People were almost hostile to him. They would essentially say, “This is a place of theory and thought and here you are getting your hands all greasy building these machines. They don’t really belong in an academic environment.” So maybe that’s unique to the Institute for Advanced Study, which is kind of theoretical.

Earnest: Well, that was a machine that was well ahead of its time, and I had the good fortune to get to see it in the mid-50s when I was still in the Navy. We took a trip over there and it looked a bit kludgy, that is, they were vacuum tubes that were out in space held by stiff wires where they had modified something.

Spicer: Right. <*laughs*>

Earnest: Anyway, yeah. As I say, I was not a party to the major decision-making at that time. But MIT did get out of it. There had been two corporations set up specifically to work on SAGE, namely, MITRE, and then a spin-off from the RAND Corporation called System Development Corporation, SDC. And their responsibility was programming. Major undertaking.

Spicer: Yeah.

Earnest: But let’s see. You mentioned spelling checkers.

Spicer: Mm-hm.

Earnest: And that was a sort of side effect to a project. I had, while I was working at MIT and MITRE, I took some graduate courses at MIT with the intention of getting a master’s. And at one point I had to face up that I had to do a thesis. And I had to work full-time on it. They wouldn’t accept part-time work. So I arranged to take a semester off from work. And undertook a project to develop a cursive writing recognizer. That is, cursive handwriting.

Spicer: Was there a need driving this or was it a pure research idea?

Earnest: Well, the idea we had was if people could do computer input just by writing on a piece of paper, that could be an inexpensive and effective way to get information in.

Spicer: Hmm. Okay.

Earnest: Simpler in many ways than a typewriter or keyboard. So I set out to do that, and [in] the initial project, I set up an emulator of a computer that worked not just with bit vectors but with bitmaps, two-dimensional arrays of bits with 0 representing white and 1 representing black. So I designed this computer that processed these bitmaps and got it to be able to extract features from the drawing or handwriting. In fact, at one point, [I] had it so you could do a drawing and it would then decompose it and turn it into a dot-to-dot puzzle. That is, if you connect the dots in order, it reconstructs the original drawing. That worked. But getting it to abstract from handwriting was a more challenging thing, because

the lines cross and merge and so forth. And I was having trouble with that. At the end of the semester I hadn’t succeeded in solving that problem. They urged me to continue, but at that point I had a wife and three children and didn’t think I could afford to do that much longer. So I convinced them that they should give me a master’s degree based on what I had done, and succeeded. Shortly after that, one of my advisers offered me access to a bigger computer. It was located at Lincoln Lab, called TX-2. And I took the offer. It allowed me to work evenings and weekends on this project on my own. And I had a lot of fun doing that. Eventually [I] got it to work and I won a free trip to Europe to present a paper on that, as a result. Then in order to make that work, I had to create a spelling checker. And so I had created this dictionary, if you want to call it that, that I took with me when I later moved to Stanford to help start what became the Stanford Artificial Intelligence Lab or SAIL. And shortly after I got here, I knew that I needed a spelling checker because of my rotten spelling. So I think I mentioned this, didn’t I? Didn’t I already tell this story? No. Sorry.

Spicer: Okay.

Earnest: That was in our previous conversation. <*laughs*>

Spicer: Okay.

Earnest: Okay. What I did was I got a student to take this list of words and process it against a text file in sort of batch mode and produce a list of words that it did not recognize and the page and line number in which each occurred. But then it offered no help in figuring out the correct spelling. So it was only mildly popular. I should mention that one limitation was that the list of words I had included no plurals, only singulars. So I had to devise a suffix-stripping scheme that would turn plural words into singular. And that worked. It also could pass words that were misspelled. But that, in practice, didn’t seem to happen much. But anyway, that worked after a fashion. Then later—

Spicer: I’m sorry. What machine would this have run on?

Earnest: This was running initially on our PDP-6, DEC PDP-6, which was the first commercial timesharing system. And then subsequently on the DEC-10 computer. Well, then later I got another student, Ralph Gorin, to develop a better version that was interactive, and that offered alternative correct spellings that you could choose from. And that became very popular. It was called SPELL. And a very short time later, the ARPANET began working and all of our software, we put it up for free. Take it. <*laughs*> And people gobbled it up from all over the world and started using it. Then about a decade later, when personal computers began to appear, the text editors nearly all had spelling checkers. People had gotten used to it by then, and it's now all over the world.

Spicer: Did people build on your dictionary, for example?

Earnest: I think they each started off on their own. I will readily admit that my dictionary was a little weird. I discovered that the book I took it from, which was the 30,000 most common English words, had given strong weight to the King James version of the Bible, so there were a bunch of archaic terms there, thee, thy, thou and so forth. On the other hands, there were many missing modern terms, such as television and radar and so on, so I stuck those in. But once there were online collections of documents, it's very easy to create word lists that are much more complete and that's what happened. People just take a document, decompose it into its individual words and you've got a dictionary.

Spicer: What year did you put this out on the ARPANET?

Earnest: 1971. All of the other spelling checkers were, I believe, developed much later. There's been a bunch of histories written about spelling checkers. Nearest I can tell, all of them are wrong. Many people claim to have invented it. That's what happens.

Spicer: Yes, I see a lot of that in computer history. What was your next step? What did you work on next after the dictionary and the cursive writing?

Earnest: Well, let's see. There was all kinds of stuff going on in my lab.

Spicer: You were at Stanford at this point?

Earnest: Yeah. One of the things I worked on was, I started the first successful hand/eye project there. That is, using a television camera to identify objects sitting on a table, using a mechanical arm to manipulate them. We started out playing with blocks, sort of like a two-year-old and stacked them, this sort of thing. That worked. Actually, MIT beat us to that one. They were a couple of months ahead of us in getting theirs to work.

Spicer: This falls under the category of machine vision, I guess?

Earnest: Yes. We called it hand/eye coordination. Then another chance encounter led me to start another robotics project. I was visiting a former Caltech teammate on our illustrious football team. I played four years at that. Then he had, in the corner of his—a mechanical engineer—had a four wheeled cart that he had started off originally to look into the control of a moon rover. There's a two-and-a-half second round trip delay in communicating from the earth to the moon, and the original plan for moon exploration was to send a robot vehicle and steer it from the earth. But he was concerned that it might not be possible to steer a vehicle with a two-and-a-half second delay reliably and he was right. That is, when you're steering now and it won't take effect till a little later, you have a tendency to lose control and go in random directions. However, another colleague of his developed an improved version of remote

control in which your steering commands are integrated and there's a dot shown on the same screen as what the camera sees, showing where you will be when a change in direction now will take effect. With that scheme, it worked. Now that all became moot when President Kennedy decided that we should send men to the moon. I happen to think he was wrong. That was a gigantically expensive thing and we could have done a much better job of exploring the moon at a tiny fraction of the cost had we sent robots, or remotely steered vehicles. But that was a political decision. We wanted to one-up the Russians and we did. That then started the whole astronaut program, which in my view, has been a gigantic waste of money, and always will be. Now they want to send people to Mars. It's an even stupider idea. Robots are doing an excellent job of exploring Mars and other parts of the solar system and they will continue to do so. Meanwhile the politicians will continue doing their thing.

Spicer: Right. So this was the Stanford Cart you were talking about, what became the Stanford Cart.

Earnest: Okay, so my idea was to borrow that and turn it into a robot road vehicle, that is, one that would steer using visual information from a television camera.

Spicer: Would it be self-guided?

Earnest: Yeah, that was my goal. It turned out to be somewhat unreachable at that time. I recruited a electrical engineering student, Rod Schmidt, to first of all build a television transmitter, so that the television image could be sent to the central computer, the PDP-10 by that time, and then [a] radio link in the other direction that would allow our computer to steer the vehicle. In order to make that all work, I had a big ramp built down our front steps so that our vehicle could get from the lab down to the road and wander about. Well, the first thing that happened was, it became enormously popular among our graduate students as a toy, because they could sit at a desk, seeing what the camera sees and steer it around the neighborhood, and had a good time doing that.

Spicer: So this is being piloted by a human.

Earnest: By a human, yes.

Spicer: Remotely.

Earnest: Right. Okay now, my goal was to be able to process the visual information in a timely manner and generate steering commands. At that point, John McCarthy got interested. As soon as we got it working, he got interested and sort of took the project away from me, and we eventually discovered that our computer was not fast enough to steer it. It was fast enough to steer it a few centimeters, a centimeter per minute, something like that, very slow. And then later, another graduate student, Hans

Moravec—well, I should mention that Rod Schmidt did get a PhD out of it, and then Hans Moravec took over and made a more elaborate system that could navigate through a place with obstacles of various sorts and find a path around them.

Spicer: Was there any communication between you and the SRI people with Shakey?

Earnest: Well, yeah, we both used to visit each other periodically and had a lot of friendly discussions and of course, some of them—well, Nils Nilsson, who was part of the Shakey project, eventually came to our group and continued. Now later on, we donated the Cart to the Computer History Museum and you have it on display. We had earlier actually donated it to the original Computer Museum in Boston and they had it on display also. The next generation of SAIL people, led by Sebastian Thrun, took on this goal again and they succeeded. That is, they won, sometime in the 2000s, the race across the desert with a computer-controlled vehicle, in which case the computer was on board the vehicle,. It was not a remotely controlled thing.

Spicer: And these vehicles are autonomous too, I believe. They drive themselves.

Earnest: Yeah, they drive themselves. Well, that was my goal too, but it didn't get there.

Spicer: You didn't have the horsepower at the time, probably.

Earnest: Right. We had a physically gigantic computer doing the job. I mean, it filled a room, but it just wasn't fast enough to do it.

Spicer: So there's a connection between the Stanford Cart and the autonomous vehicles that Sebastian worked on?

Earnest: Yeah. Well, sort of indirect, that is, Hans Moravec I mentioned, having developed a scheme, then when he got his PhD, he went off to Carnegie Mellon University and helped establish their robotics group there. And then Sebastian came out of there to Stanford, and then took up this idea. So there's over the years been a lot of friendly cooperation among all the universities who were interested in AI and robotics and so on.

Spicer: One of your jobs during all this time is that you're also an administrator.

Earnest: Yeah, that was my main job, was being a bureaucrat.

Spicer: But you were able to indulge some research interests as well?

Earnest: Yeah.

Spicer: Is there anything you want to tell us about running SAIL and what it was like to work with John McCarthy for example?

Earnest: Well, John was, as you know, a brilliant innovator and he had no taste for administration at all. So he sort of shucked all that off onto me. It worked out most of the time. We would occasionally clash over minor details, but everyone in the lab seemed to be having a good time. A lot of cooperation, exchange of ideas. Some of our graduate students came up with the idea of building a new kind of computer, modeled after the DEC-10 computer, but much faster. They put together a computer aided design system called SUDS, Stanford University Drawing System, that allowed you to do both logic drawings and printed circuit card layouts and you could cross-check to make sure that they both described the same thing.

Spicer: I think Curt Widdoes used that for the S1, the Livermore supercomputer.

Earnest: True. In fact, it was used by a lot of companies. DEC used it as their primary design vehicle for a decade or more.

Spicer: This is a Stanford-created program?

Earnest: Yeah.

Spicer: Did it ever get commercialized?

Earnest: Well, it certainly was commercialized by DEC and also by Valid Logic, Curt Widdoes' group, and by Sun. The original Sun workstation was designed using SUDS and in fact, even after the company had been going for some time, when I returned to Stanford, I found that Andy Bechtolsheim was still designing the new Sun workstations at SAIL. I finally convinced him that he should start using a Sun workstation instead. He reluctantly switched.

Spicer: Were there any projects, over the 20 plus years you were at SAIL, that stand out as really interesting or innovative, creative?

Earnest: Well, a whole lot. Well, when I arrived, there was some of the pioneering work on speech understanding being done by Raj Reddy and that, as you know, has matured considerably in recent years. Now we have robots talking to us on the phone all the time, sometimes annoyingly, but it reduces the costs of human interaction in some contexts. So speech understanding, robotics, higher mental functions, was one group run by Ken Colby, where they did a number of experiments on—well, they were the first to pass a Turing test, but in an odd sort of way. They built a computer model of a paranoid, a program called Parry, that you could interview and they tested it on a neuro-psychiatrist conducting interviews with real paranoids, over Teletype, and with Parry. In general, they couldn't tell which was the real person and which was the computer.

Spicer: Was it Vint Cerf who did a parody of that—although it doesn't sound like Vint Cerf—where they had Parry talking to ELIZA?

Earnest: Yeah. Well, a lot of us did that. I played that game too and it was sort of fun. The conversation was not that interesting. Well it turns out that Ken Colby, the guy that created Parry, also had a hand in creating ELIZA. Anyway, he later wandered off to UCLA and continued his work there.

Spicer: What's happening with ARPA now, in '67, '68 when you're invited to join an ARPA committee? Is that right?

Earnest: Oh, you're talking about the networking, yeah. Well, I mentioned that back when I was doing the handwriting recognition, I was sharing the [MIT] TX-2 computer with a number of people who were working on PhDs: Larry Roberts, Ivan Sutherland, who was working on Sketchpad, Len Kleinrock, who was doing network simulations, and incidentally, that work was the basis of the ARPANET. He proved that a sparsely connected network could work with packetized communications. Now subsequently, Ivan Sutherland was recruited by [J.C.R.] Licklider to be the second person to run the ARPA IPTO, their Information Processing Techniques Office. At that same time, I was recruited to move, first to the Central Intelligence Agency—I spent a year there, waste of time—then to work for the Joint Chiefs of Staff on an overall military control system. And that was about the time that Ivan moved to Washington and we met and socialized a bit. Then he tried to recruit me to join him at ARPA. I told him that my goal was to get as far from the Pentagon as possible. So he thought about that a minute and said, "Maybe I've got something for you." He then sent me to Stanford and put in a good word for me there and I drove out, interviewed and was offered the position that I've been talking about, which was a whole lot of fun.

Anyway, in the meantime, Licklider had initiated the idea of a general-purpose computer [network] system and Ivan subsequently recruited Bob Taylor to take over his office. Then he in turn coerced Larry Roberts into joining, to run the networking project. Those of us who had been working together on the TX-2 computer all sort of joined in. I represented Stanford and Len Kleinrock, who by that time was at UCLA, was involved. In fact, he volunteered to do the initial testing of the network. Larry put together performance specifications and solicited bids, got about a dozen of them I think. And those of us on the

committee reviewed them. We had our final meeting on that in Monterrey in November, 1968. The curious thing, I advocated hiring BBN to do it. I knew a number of the people there who'd come from Lincoln Lab. The committee as a whole recommended Raytheon. Somehow, two months later, I found that BBN had been hired. I never learned exactly how that happened, but it was a good choice. They built the network and made it work.

Spicer: You said once in one of your writings, the first ARPANET transmission was not in 1969, it was in '71. Is that right?

Earnest: Well, what was put together initially was a four-node test bed. It was not a network. It was to develop the packet switching software and make it work. Now those who are involved like to say that that was the beginning of the network. It was the beginning of the development in my view. It really didn't work until around 1971. In my view, it began working when the first set of host-to-host protocols were implemented. That was file transfer and what was called Telnet, that is, remote login to another computer. That was around 1971. Then shortly after that, in 1972, we got email, which should have been in from the beginning. That was an oversight of our committee. We regarded email as frivolous, but it was added and became the dynamite application. So that all happened in the late '60s, into early '70s. It really didn't take off until around 1972. They began working on it I think in 1971.

Spicer: How did having the ARPANET around change your work?

Earnest: Well, it was much easier to collaborate with people from afar and people started writing joint papers together, often between our group and the MIT AI group and the Carnegie Mellon group. There was a lot of stuff going on interactively. So it made some things a lot easier. And then ARPA was still kind of a closed group, that is, only people who were ARPA contractors were allowed in. So most people were seeking ways to expand the coverage. In our group, we attempted to make, we did make, a system called Dialnet, which offered all the services of ARPANET, but using a dial-up modem, so that you could call another computer and log in on it, do file transfers, do email, all that stuff. We offered to give it away. We actually debugged it with a group at UC Santa Cruz, so it worked between our two systems. John McCarthy and I went to the first West Coast Computer Faire, gave a talk on it, and offered it to anybody who wanted it. There were no takers. I was surprised. Another presentation at that same conference turned out to be rather successful. That was when Apple put forth I think it was the Apple II at that point.

Spicer: I think it was the Apple 1.

Earnest: Apple 1? Could have been, yeah.

Spicer: Tell us about the Finger program.

Earnest: Since computer time in that era was extremely valuable, and there was only a limited amount of it, people ended up working at SAIL 24 hours a day, seven days a week. Since I was the principal bureaucrat, I had to interact with everyone from time to time and it was kind of awkward to figure out when I could catch someone. For a time, I actually tried living a 25-hour day, that is, each day I would go to bed one hour later and get up one hour later, and kept working around the clock that way. I was also bicycle-commuting and I found that riding to work at 3 a.m. on a cold, rainy night was not too pleasant. Also, getting up just in time for a dinner party was a little strange. I didn't do that too long. Then I decided that what I needed was a tool to figure out what phase other people were in. The computer music people, we had sort of pushed off to the late night [shift], midnight to 8 a.m., so I sort of knew when they were going to be around, but I eventually wrote a program called Finger to tell me—well, first of all, traditional programs identified people just by their login ID, not real names. I wanted to use real names, so I wrote a program that you could ask about a list of people and it would tell you, first of all if they were running currently and not only that, but what they were running, what program. If not, if they had stopped interacting, it would tell you how long the terminal had been idle, and if they were logged off, it would tell you when they last logged off, so you could estimate when they might re-appear. Then after a bit, somebody asked for another feature, which is a file where they could post what their plans were, like if they were going on vacation or what time they were expecting to be back or things of that sort, so we added that.

I called this program Finger because it replaced—people had been using a program called Who that would let you run down the list, try to figure out who these people were. It showed only their login IDs, not their full real names. Finger also showed if someone was on a terminal and where it was. You could figure out which room they were in, that sort of thing. That turned out to be very useful to me and also, as usual, we put the software up on our public file area, and other people with DEC-10 systems around the world started using it. But they then started asking for a network version, so that they could see not only the people on their computer, but anywhere else on the network. So we added that feature. You could look all over. Next thing I knew, people were using this in a completely different way than what I intended. They would put together lists of their friends, wherever they were and could then ask about all those, and they could tell which ones were online so they could send email or texting if they were right there. They also started using the plan file in an unintended way, by using it to put up technical ideas or political statements or whatever, in other words, what is now called a blog. It became, in effect, the first social networking system, although nobody used that term then. That didn't come into use until 25 years later. The network version of Finger went up in 1975 and it still exists actually. It's not so widely used today. Some people use it to distribute their private key for—

Spicer: PGP [Pretty Good Privacy]?

Earnest: Yeah, for script to graphic communication. But anyway, that was sort of an interesting transformation from original intent through something else.

Spicer: Tell us about the vending machine at SAIL.

Earnest: We were located far away from food and we were running 24 hours a day. The nearest food was a beer garden with good greasy hamburgers, but they were only open at certain hours. So I put in an honor system food facility in which we purchased stuff and put it out for people to take and they were supposed to put money in a box to pay for it. That worked quite well for about a year, but suddenly, we started losing money big time, couldn't figure out what was going on. So I negotiated a deal. Canteen Corporation had an exclusive contract with Stanford for providing vending machines and we had originally tried to use them to provide some food. However, they only serviced it about once a week and within a day, it would be empty, so that didn't work out too well. Anyway, I negotiated with them to rent us a vending machine and then we built a computer interface to let our main computer control it, then would send people out to buy various kinds of food or get it delivered and that was quite popular. It was controlled by a terminal. You had to have a name and password. And you could then buy whatever you wanted. And it would automatically bill you via email.

Spicer: So, the way you got food out of this machine was to walk up to a terminal and select A5, or whatever.

Earnest: Right. It was originally a Model 33 Teletype. And then we later got a display terminal that it was little less noisy. And we sold all kinds of stuff. Incidentally, I noticed a—I put it in with one feature you could ask for double or nothing. That is, either it was free or it would double the price. It was honest fifty, fifty. And I noticed a cultural difference among the computer scientists; almost none of them ever went for double or nothing. The computer music folks almost all did for reasons I haven’t figured out. But it—if you didn’t gamble, it would still give it to you in once case out of a hundred and twenty-eight. So, there was a slight bias against gambling, as there should have been. We also—it also sold beer. But it knew how old you were. And it would not—if you were under age and attempted to buy beer, it would just say, “Sorry, kid.”

Spicer: You had underage people at your—?

Earnest: Oh, yeah.

Spicer: Really?

Earnest: We had—we invited in a bunch of high school students who were very bright and let them play. And that reminds me of another offshoot. One of our first spin-offs was a group who went to start a timesharing service in Seattle. They were not funding it. It was done by other people. But they took all of the SAIL software with them to provide this service. And it ran for I guess a couple of years, but eventually failed. But they also carried on our tradition of inviting bright high school students in to play.

And among those was a kid named Bill Gates and another, Paul Allen, who learned to program computers there in that facility using the SAIL software.

Spicer: Which facility is this, sorry?

Earnest: The name of the company? Computer Corporation, something Computer Corporation.

Spicer: Okay, and they were a service bureau, essentially?

Earnest: Yeah. So, that sort of predisposed them to pursue that line. And, as you know, a few years later they dropped out of—Gates dropped out of Harvard and started this little company called Microsoft. It’s possible he might have done that without having been introduced to computers earlier, but probably not. And he turned out to be a rather astute negotiator with IBM, took them for a ride.

Spicer: How long—when did you leave SAIL?

Earnest: The first time was in 1980.

Spicer: Okay. And that’s when you went to found Imagen?

Earnest: Yeah. I decided to complete my education by doing a Silicon Valley start up. And I learned more than I wanted to know about venture capitalists. I spent many days trying to round up capital to start. At that time, I had a beard and was wearing Hawaiian shirts. I found that that seemed to be a communication barrier with some of them. So, I eventually, after six or seven of them—talking to them, I shaved off my beard, started wearing a coat and tie. It didn’t seem to make much difference. What I didn’t know was that at that time they were mostly out on Sand Hill Road and had lunch with each other. So, if you’d talked to four or five of them and haven’t gotten a bite, you might as well forget about it. But I went on through about the thirty-five venture capitalists. They would usually send their junior people out to look at us. We showed them a computer printer that printed much more beautiful stuff using laser printers, than anything on the market. So, they would go back to the senior partners enthusiastic, and then would come back with a list of questions. Is IBM in this business? No. Is DEC in this business? No. Is Apple in this business? No. What makes you think there is a business? Well, the reason they weren’t was because we were about two years ahead of the competition. But we were unable to get any venture capital. Meanwhile, these same venture capitalists were giving away many millions of dollars—not giving away, investing in “me-too” disk makers, nearly all of whom went belly up. But we were unable to get a nickel out of them. We eventually got—rounded up a hundred thousand dollars from friends and family and decided to bootstrap, which was extremely painful. We in fairly short or—I did this with Luis Trabb Pardo, who’s one of our graduate students who had put together, using a Canon laser printer and

microcomputer, had done a demonstration system. So, he and I started this company. We realized neither of us knew anything about marketing. So, we recruited a guy who did have some background, Bob Wallace. But then we ran out of money. So, we went on zero salaries for six months, which is kind of painful. We were finally ready to deliver. I was doing the manufacturing, contracting out through various locals. And we had a set of orders. The first one was from SRI, people who knew us. In fact, all of the orders were from people who knew me. This was all done by word of mouth. So, we delivered—we had to deliver one by the end of the fiscal year, that is June 30th to SRI. And did so. But the software wasn’t working, yet.

Spicer: This is 1980?

Earnest: 1981—yeah, 1981. And so, they were getting antsy because they had paid a bunch of money, and it didn’t work. But I managed—we managed to get it all together and showed that it did work a couple of months later. So, then come January, I set out to deliver five of them. Well, actually, approaching January, I ordered five more printers from Canon. And they said very interesting. What is your credit? Well, we had none. I ended up having to put up my house as backing, which was a little scary. So, just after New Year’s of 1982, I assembled five of our minicomputers that went with the system, put them in a big box, flew to New York with them, rented a truck, went to the Canon warehouse and picked up five printers, took two of them to the airport to be shipped off, took the other three, first to Yale, delivered it. It worked. I brought the software along, then to MIT. It worked. Then to Bell Labs the third day, it worked. Then I flew to Vanderbilt University and put one in there the next day, and then to Purdue the day after that and put it in. And then flew home. Sixty days later, we actually had money. So, we began to then make a second batch.

Spicer: What was the actual product, because you got the laser engines from Canon?

Earnest: Yeah, the laser from Canon. We built a computer system. I should mention that the computer we used was the Sun workstation computer. I had negotiated with Andy Bechtolsheim to use that. That was before Sun Microsystems was formed. And I also had negotiated a license with Stanford to take the technology that we had developed in the lab. So, we paid royalties to both of them. So, after sixty days, we had some money. And we started building a second batch. And went through that cycle a few times and actually had enough money to run more or less continually. So, that was a big step forward. And we had a million dollars in sales the first year, three million the second year, ten million the third year, all with a profit. If we ever had a loss, we’d have been out of business. And so, we were on our way, we thought.

At that point, the venture capitalists took notice and expressed an interest in investing but with some provisions. First of all, they wanted a real manager to run the company, none of these flaky academics. So, we agreed to hire a real manager from Hewlett Packard. And he came on board. And one of the first things he did was to fire me. You’ve got to get rid of old kings because they’ll cause trouble. And he then fired our marketing guy who had, incidentally, replaced me as president in a political gambit. But he

then—the new guy, Patrick Welch, signed a contract with Hewlett Packard that had only two things wrong with it. One was it required almost the entire R and D effort of the company to develop a new product that HP could sell. And it had to be ready in one year or there was a horrendous penalty. On the other hand, if HP decided they didn’t want it, there was no penalty. He signed this contract without showing it to the corporate lawyer, kept it a secret. And you can expect what the outcome was at the end of the year. HP said “never mind.” They entered the business.

Spicer: And dominated it.

Earnest: Yeah. Right. It was—the interesting question is was that hiring a plant, or was he just incredibly stupid. And I can’t prove it one way or another. But the net effect was we had to sell the company to a competitor, QMS, in Alabama. And they, in turn, got gobbled up by a Japanese laser printer company, and so on. But anyway, since I came out smelling like a rose, I had enough money as a result of this that I didn’t need to work anymore. But it could have been a much more successful thing if we had not been jerked around. Sad, but educational.

Spicer: How long were you at Imagen?

Earnest: Four years.

Spicer: Okay, and then back to SAIL?

Earnest: Well, back to Stanford. Now, in the meantime, McCarthy had dismantled SAIL. It didn’t exist. He asked me to return, and I agreed. So, I started working with him on a multiprocessor LISP system. We had a multiprocessor system in the basement called Gang of Four. And that never quite panned out. But meanwhile, Nils Nilsson had taken over has chairman of the computer science department. And he asked me to help him out. So, [I] agreed to work half time as associate chair of the department and did so. Although, I should confess that I was not quite up to speed. I—as a result of the stress of putting Imagen together, I’d gained an enormous amount of weight. And a side effect of that was apparently that I had sleep apnea. I couldn’t sleep very well. I read an article about that, went to my doctor and said I think I have sleep apnea. She said you couldn’t possibly. You’re much too lively. She dismissed it. It turns out that she was wrong. I went ahead fifteen years struggling before I finally insisted on a test and had it big time. So, anyway, I was in kind of bad shape then. And eventually, the funding went away for McCarthy’s project. And I went away, too.

Spicer: So, when did you leave Stanford, then?

Earnest: 1988.

Spicer: Okay, and what have you been doing since then?

Earnest: World travels, giving politicians a hard time, writing occasional articles refuting historical nonsense in the computer field. Well, I should mention that I also have been heavily involved since 1973 in bicycle racing administration and trying to reform it.

Spicer: Were you surprised about Lance Armstrong?

Earnest: Was I surprised?

Spicer: Were you surprised?

Earnest: Not at all. As soon as I learned that he was consulting with Dr. Ferrari in Italy, whose specialty was blood doping, I knew what was going on, couldn’t prove it. And we muddled ahead in that mode. It’s kind of sad. He was, of course, not the only one who was doing it. It was widespread in the sport. And—

Spicer: I wonder what it is about cycling that—you don’t hear about doping in a lot of other sports. Certain sports, sure, but cycling seems to have taken the lead for—

Earnest: Well, actually—

Spicer: Weightlifting, maybe. I don’t know.

Earnest: It was initiated in distance running around 1976 in the Olympics. A guy was clearly blood doping. It has been used in a number of sports. And it’s going to be a technological race into the future. Before long, we will have genetic engineering to deal with, in which parents can have their offspring designed to be superb basketball players or football players or whatever. Now, should that be allowed? I suspect we will have to do it. I don’t know a way to prevent that.

Spicer: These things often start with—I’m thinking of plastic surgery as a parallel. They start as a means of correcting a deficit. And then they become—they just seem to propagate into the general population.

Earnest: Well, back on the bicycle racing theme, one of the early things I attempted to do there was to change the rule on helmets. I observed that many people were seriously injured by head injuries. And there were beginning to appear some helmets that could do a good job of protecting you. So, I attempted to get that rule tightened. And it was extremely unpopular. People have a tradition of wearing what we

called leather hairnets. These are padded leather straps that go over your head and provide almost no protection. And I fought that one and lost. They booted me off the board of directors of the national organization because of my advocacy of strong helmet rules. As I left, I wrote a report refuting all known arguments against a strong helmet. And as luck would have it, when the next board of directors meeting was held near a place where another rider died in a velodrome crash wearing an inadequate helmet. And they got the message finally and passed a strong helmet rule. And then two years later I got my seat back when everybody had figured out that was the right thing. Then that rule slowly percolated around the world. So, now it’s generally required in bicycle racing everywhere. That I think of as a real accomplishment.

Spicer: Yeah, that’s—you should be proud of that.

Earnest: Anyway, but then another thing I attempted to do was bring all the different racing or—bike racing organizations in the U.S. into one organization. I initiated a reorganization that was initially blocked because I was proposing democratic representation of all the participants. And there were a number of commercial interests who did not like that idea. Eventually, they had their way. They took my bylaw proposal and distorted it so that a majority of the board of directors represented commercial interests who made up less than one percent of the participants. That’s rather un-democratic. And they got away with it. We are still—they are still operating in that mode today with the result that USA Cycling is run by a San Francisco investment banker. And he and his colleagues run it to suit their interests. My next project will be to attempt to unseat them. But that’s going to be a rough go because there’s big money on the other side. It will require a change in federal law, namely the Ted Stevens Olympic and Amateur Sports Act, under which the U.S. Olympic Committee operates. It needs to be fixed to reduce the monopoly control of the central organization to require democratic representation of participants. But that’s not going to be an easy one.

Spicer: No. Is there anything else you’d like to chat about or tell us about you future plans?

Earnest: About future plans, well for me that’s my next project. We haven’t talked about all the other things that spun off from SAIL.

Spicer: Okay, yeah we can do that. So, why don’t we do that? Let’s look at some SAIL spin offs. You mentioned Chowning, and Kay, and Taylor…

Earnest: Yeah, another spin off, of course, was Sun Microsystems. Andy Bechtolsheim was one of the co-founders. And of course Bill Joy from Berkeley provided the software for that. They were pretty successful for a time. Of course, they’ve now been gobbled up by Oracle. Then another spin off of sorts was Cisco, more of a rip-off than a spin-off in that all of their start up technology they swiped from Stanford, pretending that they had created it all and got away with it, basically. Stanford eventually

negotiated a license with them in which they paid their debts in the form of equipment discounts which in a real sense just—equipment discounts to Stanford, which is a real sales enabler. But I learned a few things from that. Let’s see what other companies were—well, Apple got a big boost from our people. Jef Raskin started the Macintosh project.

Spicer: He was at SAIL?

Earnest: What?

Spicer: He was at SAIL?

Earnest: Yeah, mm-hmm.

Spicer: Oh, I didn’t know that.

Earnest: Yeah.

Spicer: In the music group, or somewhere—?

Earnest: No. I forgot exactly what he was doing. He did, actually, while he was at SAIL, he did one of the earliest start-ups of sorts. That is he built a model airplane kit and sold it using our facilities and design.

Spicer: You mentioned Alan Kay and Larry Tesler going to PARC.

Earnest: Yeah, they went to PARC. And, of course, Alan was part of the Alto development and also developed the Smalltalk language while he was there. Larry Tesler, who originally helped me develop Pub, this document compiler with spreadsheets and so forth, did the first WYSIWYG text editor, what you see is what you get. He did that at PARC.

Spicer: Do you remember what that was called?

Earnest: I should, but—

Spicer: It’s not Bravo?

Earnest: I think it is, yeah. And then he moved on to Apple and developed the software for their Lisa computer, which didn’t—was too expensive to be very popular. And then he was in charge of their first tablet of the Newton, which eventually died. It did include a cursive handwriting recognizer, which—

Spicer: How did it compare to yours?

Earnest: It was about the same, that is not quite good enough.

Spicer: It’s a formidable problem.

Earnest: Yeah, it is.

Spicer: There’s so much variation.

Earnest: And they actually started from an easier point, that is they worked with the pen trajectory, which tells you a lot more about the handwriting than just the resulting image. So, it should have been better than mine, but it wasn’t. Curiously, even though I worked with Larry a lot, I never mentioned to him that I had done a handwriting system. He could have had it for free, but they contracted the recognizer out to a Russian group for a million or two. But it wasn’t good enough. Anyway, that was an oddity. Let’s see. What else was there back there?

Spicer: The news service, we didn’t talk about that.

Earnest: That was initiated by John McCarthy to build a news service. We initially contracted with Associated Press. And they sold us their news service at the same rates they sold to college newspapers, really cheap. They didn’t realize that we were doing something quite different, namely as the stories came in—and then we later added The New York Times. As the stories came in, you could either view them directly on your screen, the two news wires, or—but they were also being indexed according to—well, we had a list of uninteresting words like a, an, the. Things occur in every story. But we indexed everything else so that you could then ask for messages that used some Boolean combination of words. This was the same scheme that I had used on the search engine many years earlier. And it worked pretty well being able to extract things that, from the last couple of weeks, that used some combination of words. And you could also leave a standing request so that if any story appeared with some combination of words, it would notify you immediately by email.

Spicer: That’s pretty advanced stuff.

Earnest: Yeah, and that was quite popular.

Spicer: I know Google has a thing called Alerts, Google Alerts where you put in a keyword or phrase that you want it to look for, and it just sits there in the background and sends you stuff as it comes in.

Earnest: Uh huh. Right, similar idea, right.

Spicer: It’s very useful, actually.

Earnest: Well, we initially made our news service accessible to anybody who logged in on our computer, but I think Associated Press heard about that. And we had to change so that just people who were our real—

Spicer: Stanford people.

Earnest: Stanford people could do it. Well, actually, we had a lot of outside accounts from other research groups, but we didn’t talk about that. Anyway, it did play a minor role in a couple of subsequent incidents. There was the Three Mile Island nuclear problem in which the Lawrence Livermore Lab people had an emergency response unit that was supposed to be involved in dealing with that problem. And they found that they couldn’t keep track of what was going on. So, we set them up with a news service account. And that enabled them to keep track of what was developing. And then, I think it was a little while later. Maybe it was earlier. There was the Tiananmen Square incident in China. At that time, there were a bunch of Chinese students in the U.S. who—and, as you know, there was a sort of news blackout in China. And the students were trying to keep people, friends and family, informed. There were lots of fax machines in China. There was no Internet access at all. So, what they were doing was taking what news they could find and faxing it home. So, what I did was a set up a standing request for anything mentioning China, and then funneled it out to their distribution list so that they could send it home. And that was fairly successful. Although, I later learned that some of those involved—they were—in the students here, there were some that were tied to the wrong side of the political spectrum. And they reported on those who were sending the stuff home. And so, they ended up in jail, some of them, as a result of all that, which is unfortunate. But, anyway, an interesting development.

Spicer: That would have been ’89, I guess.

Earnest: Yeah.

Spicer: Tiananmen Square. Okay, well I think I’m done. I don’t know if you have any more thoughts, or think we’re done.

Earnest: Not off hand.

END OF INTERVIEW

