

An Interview with  
STEPHEN CROCKER

OH 233

Conducted by Judy O'Neill

on

24 October 1991

Glenwood, MD

Charles Babbage Institute  
Center for the History of Information Processing

University of Minnesota, Minneapolis  
Copyright, Charles Babbage Institute

Stephen Crocker Interview  
24 October 1991

Abstract

Crocker discusses computer networks, artificial intelligence research, and his work at the Defense Advanced Research Projects Agency (DARPA). Crocker details his interactions with other DARPA personnel including Cordell Green, Barry Wessler, Larry Roberts, Robert Kahn, and David Russell. He discusses his work as a program manager

in the Information Processing Techniques Office (IPTO). Crocker describes the effect of the Mansfield amendment, which specified research needed to be relevant to the military, on DARPA projects. He also discusses software development.

STEPHEN CROCKER INTERVIEW

DATE: 24 OCTOBER 1991

INTERVIEWER: JUDY O'NEILL

LOCATION: Glenwood, MD

O'NEILL: My primary interest is your time at ARPA, but I'd like to start with your education and how you got involved in computing.

CROCKER: I grew up in Los Angeles, principally, and finished high school there in 1961. I had been exposed to computers while I was in high school, and I got a summer job programming when I finished high school.

Throughout college I spent a lot of time programming, partly as work and partly for fun. In 1965 I joined an ARPA-supported project. It was, in fact, a precursor to the ARPANET venture. ARPA had been engaged in a couple of prior attempts at building networks, some of which are well-documented. This one was not, because it eventually floundered, but it was entirely on-campus at UCLA and involved three existing computer centers of the classic "big iron" type. They were all IBM 7090 or 7094 and had an interesting political history behind them. One was the campus's main computing facility; one was an NIH-supported facility in the medical center and was one of the mainstays of NIH's program of big centers of excellence in biomedical computing; and the third was an adjunct of the business school. In the mid-1950s IBM sponsored two major computer centers, one at MIT and one at UCLA., to stimulate the use of computing in graduate schools of business. The one at UCLA was called Western Data Processing Center. The three computer centers at UCLA grew up with independent staffs and with quite independent cultures.

In 1965 a three-node network was designed that involved connecting these three facilities, which had similar hardware but grossly different cultures and grossly different missions and conceptions of themselves. WDPC, the business school center, for example, discouraged machine language programming and allocated each user two minutes, which was reasonable for a broad class of users doing simple tasks.

In contrast, the main computing facility was set-up by numerical analysts who had done a lot of original research on a home-built machine. Their approach accommodated programming in octal and running for many hours or even all weekend. ARPA decided to support a project intended to link these systems together. In retrospect, it's quite fascinating, because IBM had just introduced 360s, and the project was to use low end machines, 360/40s. These were the second-to-the-bottom of the low-end machines, and these were to be the machines that would be linked to each other, with each then acting as front end to a 7094. Thus there was to be a triangle of 360/40s, and attached to each node would be the 7094s. The function of these 360/40s was to serve as a batch job submittal mechanism, allowing users interactively to prepare programs to be submitted to the 7094s, resembling TSO which came years later. This configuration was also intended to provide a modest amount of interactive computing and direct connection to the 7094s. Similar ideas eventually came to pass in the ARPANET in a different way with TIPS, IMPS, and remote job entry workstation. We set about to build an operating system for the 360/40s from scratch that made use of the hardware in a somewhat different way. The programs were intended to be transparent and relocatable.

I was a student working at the bottom of the totem pole, and I learned quite a bit. The project was staffed by about ten people.

One of the issues was that the project was conducted entirely by WDPC, so from the other two centers' points of view, it was WDPC's project, not a combined effort, though we knew we were serving all three centers. The project leader had been a WDPC employee, and at the outset of the project, he solicited volunteers by drawing a line in the sand and saying, "Those people who are going to join the project, step across now." He intended that the group would think of itself as independent of WDPC and thus not captured by the partisan politics on computer center versus the others. Nonetheless, I think the other groups' view was that WDPC had taken over the effort. Ivan Sutherland and Bob Taylor from ARPA came to review it, and it became clear that we had a very tight schedule and a lot of things to accomplish: we had a deadline of February this was in early December of 1965. February was also the deadline for an increment of funding, and a decision was to be made about whether to continue the project. For

the next couple of days we were energized and fully committed to getting all of the code written and running. The Director of the WDPC, who was the titular head of this project but uninvolved on a day-to-day basis, decided an academic institution should not have to comply with such orders. He fired the top four or five people (out of a ten-person project) and gutted it. He called up ARPA and said, "We aren't going to do that." The project then died fairly quickly. The funds that were left over in the contract were transferred to Professor Gerald Estrin in the Computer Science Department, whom I eventually worked for, and the contract was redirected, as often happens, into a related but different effort focusing on measurement of computer systems.

O'NEILL: And that's the Snuper Computer work?

CROCKER: And that was the Snuper Computer work, that's right.

O'NEILL: One of the things that I keep running across that maybe you can help me clear up is the idea that the network in California was focused on load-sharing -- load-sharing between the 7090s or the 7094s or whatever they were. Is that part of the redirection that you're talking about in December?

CROCKER: No, that's not part of the redirection.

The notion of load sharing reappears in the early justifications for the ARPANET itself, and this project was not the ARPANET; it was distinct. As I recall, there was indeed discussion on moving the load from one 7094 to another in order to balance out loads and so forth. No serious work was done on developing regimes for doing that, for considering the incompatibilities of programs from one machine to the other, or for what was required to write a program that would actually move from one to the other. Had we built the system, it's likely some people would have experimented with that and tried to make some progress. As we found out on the ARPANET, although load sharing was one of the justifications and an announced goal, load sharing is both different from what it might appear and harder than it might seem. The term "load sharing" conveys an image of automatically moving a program from one

node to another when the queue builds up too high on one machine. That doesn't happen in any direct sense. But, in another sense, there's an enormous amount of load sharing. We discovered this much later when we had many DEC System 20s or their earlier versions, PDP-10s, running a TENEX operating system on the ARPANET. We brought those on-line to support the research community and give people accounts on them. As fast as we plugged them in, the loads skyrocketed on them. That was an effective but very simplistic form of load sharing: whenever the load was too high on a machine, not daily or hourly, but when the load was consistently high, one could plug in another machine and move files. We could typically off-load one machine and load-up another in just a few hours.

O'NEILL: Now this is what was going on at ISI?

CROCKER: Yes. ISI was a prime example of that phenomenon, but it could be done anywhere. In a sense, it validated the idea that networking makes possible load sharing, though in a simplistic, low-tech fashion.

O'NEILL: Okay, so you had been working on this project as a programmer, part of this ten-person team. Three or four of the people were removed. What's left of the DARPA funds was reassigned to Estrin, and you went to work for Estrin.

CROCKER: Yes.

O'NEILL: Okay, and now we're at 1966.

CROCKER: As I recall, I worked for Estrin the summer and fall of 1966, and at the end of January 1967, I went to MIT as a graduate student. I completed part of my graduate work at MIT and returned to UCLA a year and a half later.

O'NEILL: Okay. That's one of the we're confused, because your name appears on a publication from MIT lab, and we couldn't figure out what you were doing there.



CROCKER: The chess program.

O'NEILL: Yes, right. Exactly.

CROCKER: Right. I didn't finish a degree at MIT, but the AI lab there was one of the most exciting places in the world, and even though I spent only a year and a half, there I learned quite a bit about computing systems in general, AI, and the MIT culture.

O'NEILL: Why did you decide to go out there [to MIT]?

CROCKER: I visited there in September of 1966. A professor had encouraged me to look at schools outside California, and so I visited 9 schools in a three-week period.

O'NEILL: So, you were interested in computer science and in 1966 visited nine campuses.

CROCKER: Yes: Harvard, MIT, Yale, Princeton, CMU, the University of Wisconsin, the University of Michigan, the University of Illinois, and Cornell.

I was actually interested in artificial intelligence, and that was not a common subject. During my senior year of high school, I had taken courses at UCLA, which had a program allowing seniors to take classes. The first course I took was psychology. I'd had a little exposure to computing by that time, and I had started to look at the world from a computing point of view. I remember asking, "Could you program up some of these psych theories and see if you get that kind of behavior?" because they were just groping toward things they called axioms of stimulus-response connections and so forth. That question caused the professor to send out a letter on this, though he didn't tell me about it until he got responses. One response came from the System Development Corporation. Someone was

teaching third and fourth graders to program, and he invited me to visit, and I did. The other response came from the Rand Corporation. It was an interesting letter; it discussed where to go to study questions like that and what the job prospects were, and the writer invited me to come see him. I visited him at Rand and I was quite impressed because he and a couple of his colleagues had taken questions like that and were pursuing them fairly seriously.

He gave me a couple of papers to read. This was early 1961, around January. It was amazing to me because he really had worked on the very question I had asked. One of the papers was called the General Problem Solver by Newell, Shaw, and him, and the other was by Feigenbaum on Elementary Perceiver and Memorizer. I read the papers and studied them, and about six months later I realized what I had fallen into by having some time with Herb Simon. It was amazing.

Among other things, that episode meant that when I went looking for a graduate school some years later, I went to CMU to follow-up. It was a close call between CMU and MIT. At MIT I had met Minsky and the AI lab, and it seemed to be a pretty exciting place. Anyway, I had AI on my brain.

O'NEILL: Okay. So you were at MIT for about a year and a half, and then you went back to UCLA. Did you go back to the graduate program as a student at UCLA?

CROCKER: I came back originally just to work for the summer in 1968, and at the end of the summer I found myself having some trepidation about returning to MIT. I talked to Estrin, who was a great person—a good friend and a good mentor—and he said, "Well, if you don't want to go back to MIT you can transfer here." He offered an excellent arrangement: skip a master's program, don't bother with classes, write a thesis, and enough work on projects that I could support myself.

O'NEILL: Did UCLA have an AI program of any description?

CROCKER: No, but it offered quite a bit of freedom to do what you wanted, so it was build from scratch. My childhood had been oriented toward mathematics and puzzle solving, and I'd done a fair amount of system-type programming. So of the range of things in AI, I found myself most comfortable with the theorem-proving systems. I had taken quite a bit of linguistics on the side as an undergraduate. I had some understanding of transformational grammar, but I did not see myself as competitive in those areas or in other areas. It was not likely that I was going to make big things happen there. At the end of the summer of 1968, Estrin made his offer, and I thought about it for a while.

O'NEILL: So, when you were back at UCLA then, is that when you started working for Kleinrock or was it still Estrin doing this Snuper computer?

CROCKER: It was still Estrin doing Snuper computer. Things were happening at a fairly interesting pace. In the summer of 1968, I was back at UCLA but only for the summer. ARPA had a tradition of getting the principal investigators of major laboratories together for a PI meeting once a year. There were a small number of projects that were fairly well-funded, and the people running these projects were the stars in the field. So only two people went from MIT, the heads of the project. One group at Carnegie was represented, and Newell and Simon were the heads of that project. With all of the principal investigators together, there were only about 20 to 30 people. These were fun meetings, and Slotnik at Illinois suggested that it would be interesting to clone the PI meetings and have a graduate student conference. That was accepted by the group, and Slotnik arranged to host the first one. That took place in early July of 1968. And my best friend, Vint Cerf, was at UCLA I had introduced him to Estrin.

O'NEILL: You had gone to high school together?

CROCKER: We had gone to high school, and then he went to Stanford and I went to UCLA He had finished at Stanford, moved back to Los Angeles, taken a year off to work, and was looking at graduate schools. He had applied to Irvine, which I had been excited about before Thompkins encouraged me to look elsewhere as well. So I thought

perhaps Vint and I would be at Irvine together. Cerf got turned down at Irvine, which was an amazing mistake. Cerf has an unblemished record in multiple dimensions. He's an organized person and very talented, and it was hard to understand what they were thinking. At that time, he was eager to stay in the Los Angeles area, so I introduced him to Estrin. Soon thereafter he was a graduate student at UCLA. Then a year and a half or two years later I came back to work with him for the summer.

Vint was selected as one of the two graduate students that Estrin sent to this meeting. I remember wishing I could go, but I knew I was just at UCLA for the summer. Then, at the last minute, Minsky chose me as one of his two students. The other student from MIT was Pat Winston, who is Director of the AI lab now. Pat was from the AI lab, and then there were two more from Project Mac. So about 30 of us showed up for a no-agenda meeting, to get to know each other and talk.

It was an amazing group of people. We just cooked for about three days. I went into this with childish arrogance, thinking maybe ten percent of these people would be smart or interesting. At the end of three days, we all knew each other and who was working on what; we had sort of sized each other up. Going down the list of names, it surprised me that about half of those present had been interesting, worthwhile contributors. Also interesting was the match between where I thought good work was being done and the people and the places represented. It was a pretty good match, and I thought, "Somebody's got this figured out. Somebody actually understands." The person running the meeting, Barry Wessler, was the only Program Manager in the office. I went up to Barry and said, "When are you leaving?" basically implying that I wanted his job. He just looked at me as if to say, "Where'd you come from?"

O'NEILL: So, in July of 1968 you had this first graduate student meeting.

CROCKER: Yes. I was spending the summer at UCLA. At the end of the summer, I decided to transfer back to UCLA. ARPA, at that time, was working hard on the ARPANET and had prepared an RFP for the initial sub-net. Now, Barry

Wessler, who had of course been in the middle of it all back in the office, tried to talk to us graduate students about network issues and very few of us, if any, were interested.

O'NEILL: So, at this meeting he was already talking about it?

CROCKER: Yes, because they had I don't remember whether the RFP was on the street or whether it was just about to hit the street

O'NEILL: It hadn't hit the street in the summer. It was in the fall. They had been working on it for awhile.

CROCKER: Shortly thereafter, I was doing some part-time work with old friends, the principals of the prior project back in the on-campus network. They were now working for a small company.

O'NEILL: Jacobi Systems.

CROCKER: Jacobi Systems, yes. Did I send you a lot of stuff?

O'NEILL: No, I talked to Vint Cerf, and he

CROCKER: I see. Well, yes, I dragged Vint into that as well.

O'NEILL: and the fact that you had prepared a proposal in response to the RFP.

CROCKER: That's right. All right, so you're way ahead of me here. Talk to Vint; he knows my whole life.

O'NEILL: He said talk to you.

CROCKER: That's right. Well, if you talked to Vint, then I don't have so many things to follow. We worked on a proposal that did not win that was immediately following, at the end of the summer or in the fall, I guess; we did a little bit of work with them while we were working part-time on the outside.

O'NEILL: Had you gotten interested in working on this at Jacobi Systems because of the stuff Wessler had been talking about?

CROCKER: No. In fact, I was working at Jacobi Systems because I had bumped into Al Irvine, who was the key person, and they were building an interactive interpreter of FORTRAN and BASIC programs, and it also was going to be a front-end for Univac 1108 mainframes. It showed some of the same ideas but in a purely commercial operation unrelated to networking. It was dial up: the model was going to dial up to a small front-end machine that was going to be an Interdata Model 3 and Model 4 combination that would provide easy-to-use interactive computing capability, though of very modest power.

O'NEILL: So, it was more remote computing than networking.

CROCKER: It was more remote computing, yes. It wasn't networking as we know it. Then this network RFP came out and, like many groups, they decided that they ought to have a try at it. I think the proposal was credible. I never reviewed the other proposals, but I have no trouble believing that BBN's was far superior. They were working in the area and understood it, but I thought ours was reasonable. It would have been difficult work, but it could have succeeded. So that was another sort of side venture. The attention factor within the ARPA research community (or the IPTO research community, more properly) to the network was certainly not very high in the summer of 1968, at least as measured at the graduate student level. Then in the fall or late summer, the strategy was to expand the network from west to east. I'm not sure what the reasoning was, but I believe it was thought that the west coast establishments were easier to work with, that it was easier to get them to do things besides what they wanted to do.

So the original plan was that UCLA, UCSB, SRI, and University of Utah would be the first four sites. Elmer Shapiro from SRI called together representatives from these first four sites.

O'NEILL: Now, this is still in the summer of 1968

CROCKER: The latter part of the summer, I think. And I believe both Vint and I went. As I recall, the meeting was at Santa Barbara. I'm not 100 percent sure. We didn't know how all this was going to shape up. We had heard that there was going to be a network and it was going to come to us somehow, and we started asking very basic questions. It wasn't clear what the division of labor was, and it wasn't clear what the network was going to look like. So we asked some very basic questions. We were all system programmers; had access and, in some cases, direct responsibility for the operating systems, and we understood that we were going to have to have protocols, although I don't know that we used the word. And then we were going to have to have some means of addressing things and so forth, and we started to ask questions of that level, and the answers were not forthcoming in any way. In fact, my recollection is that not only weren't there any answers, but there wasn't even a strong acknowledgement that this was the set of questions that we needed. It was rather vague. A few of us were attuned to one another and in a very informal way decided that we should keep talking about this, that we had a common cause in that this thing was going to visit on all of us.

O'NEILL: Were these people you had met at the

CROCKER: Some yes, some no. Let's see if I've got the timing right: I may have the timing backwards, because I remember that Steve Carr at Utah was important, and he hadn't been at the graduate students' conference. I remember being surprised about that. Jeff Rulifson, a student at Stanford, was on the staff at Engelbart's operation at SRI; he was clearly a central participant.

O'NEILL: Was it significant that you were students, or was it just that you were system programmers, so you had an interest in these issues, you had worked on these projects. Does the fact that you were students play a role?

CROCKER: Several conditions are intertwined here: We were students because these were research projects, and the funding was for research. Research means a professor receives the money and pays graduate students. Another aspect is that because we thought of ourselves as researchers, we tended to ask larger questions. Instead of trying to make the work as short and simple as possible, we tended to consider grander aspects, such as how to do it right, how to get the full generality or the maximum flexibility out of something. So, in that sense, there was some sharing of the vision, and that probably influenced the approach to these problems.

We decided that we would continue to talk to each other and made appointments to visit one another a kind of rotating group. We said that the first effect of networks, which were designed to make it possible to collaborate at a distance without all this travel, was to increase travel enormously! We met at SRI, at Utah, and we must have had a meeting at UCLA, I suppose, although it would be less memorable for me. But getting to visit these places was fun.

O'NEILL: Well, how did you get approval you didn't need approval?

CROCKER: Yes, we needed approval, which we got from our professors or bosses. The first meeting was sanctioned and they'd sent us to it. Since the network was coming and they didn't have a plan for it, they saw our initiative as solving the problem: all they had to do was approve.

O'NEILL: So you went off to Utah.

CROCKER: Yes. And that was part of the wonder of the ARPA environment that you took initiative and got to make something happen. ARPA liked that, too; that's how things get done. A few graduate students together was the cheapest possible way to make things happen.



O'NEILL: Yes. When did this group become more formal?

CROCKER: We met off and on through the fall of 1968 and early 1969 and outlined some ideas for ourselves. Among our ideas was the notion that there would be a big difference in computing power and also that 50 kilobits per second Host-IMP might be only modest bandwidth for some applications. We didn't have the low-level details yet because the contract hadn't been let at first, and then we didn't have Host-IMP specs. So we were concerned with what are now called higher-level protocol issues. We fastened on this kind of nifty idea that you might have a computer locally that had, for example, graphics capabilities, and you were interacting with an application program remotely that had some fairly sophisticated way of interacting. If you wanted to transport the interaction mode, download it. So we set about trying to design a language in which front-end programs would be written, with the idea that when you start up this remote program, the first thing that would happen is it would send down the wire an application program that you would run that would be your agent — your user agent, in modern terms.

The first thing that would happen when you started up a program was that we'd send, and that would be a language for which you had an interpreter. So we set about trying to define a standard user agent language. At one point we called it the Network Interchange Language (NIL) and at a different time, Decode-Encode Language (DEL). We wanted the language simple enough so that it could be interpreted efficiently or compiled efficiently, and we toyed with whether there ought to be constraints on the language rather than its being an arbitrary Turing complete language.

O'NEILL: So this was a standardized user interface?

CROCKER: This was a standardized language for writing user interfaces. We already knew that there wasn't such a thing as a standardized user interface, and I don't think it occurred to anybody for longer than a millisecond to try to say, "This is the standard interface." SRI was doing work in Engelbart's project in augmentation of human intellect.

They had their idea of what an interesting user interface was. Lincoln Laboratories had extensive experience with the TX-2 with very sophisticated interactions and a lot of famous things Sketchpad and so forth and the hardware had not standardized anywhere, so you might or might not have had a DEC 340 scope, and if you did then you wanted to take advantage of it. In those days, even byte sizes weren't standardized. That's interesting to note, because in the original protocols we tried to determine whether or not there was a common unit of data. We thought there were 18-bit machines, 36-bit machines, 32-bit machines, and somebody pointed out that there was a 75-bit machine in Israel, which I never knew much about. The original protocols performed transferring in bits. Later, when the host-level protocol was revised, this shifted quite naturally into bytes as a unit of transfer. But this gives an idea of how few standards there were. We were all programmers and so we standardized in a way to provide the flexibility to make it look like anything one wanted. That led us into Decode-Encode Language or a Network Interchange Language as a big thing we were thinking about. We met just prior to that in Utah and agreed that we had been talking long enough, that we'd better write something down. We divided up writing assignments; it was a little initial burst of things that had to come off. RFC-1 is dated April 1969. Except that there's one, RFC-4, that's dated March 1969.

O'NEILL: This was an internal thing that the group itself decided, there was no external directive?

CROCKER: Yes. We had gotten no external directive, and that was an absolutely central concern, that we had no idea whose toes we were stepping on, if anybody's in fact, I envisioned that a group of professionals would see to the things that we were concerned about, and that they would undoubtedly come from the east. We had this sort of magical view of "the east."

O'NEILL: But you were never aware of any other group actually working on this, or being contacted in any way?

CROCKER: No, in fact there wasn't. Our first contact with BBN after they won the contract was on Valentine's Day, 1969. Representatives from the host sites were invited there. Bob Kahn was the chief thinker of topology issues and

network regimes. Will Crowther was the chief network programmer, Severo Ornstein was the chief hardware person, and Dave Walden was the junior programmer assisting Will Crowther. Frank Heart was running this operation.

There were a couple of other important participants Ben Barker, if I remember.

O'NEILL: Yes, he worked with Ornstein.

CROCKER: One thing that wasn't well defined was how much of the total problem they were responsible for. Early on, they discussed tailoring the software inside the IMP with code that would be specific to each host. So there was some fuzziness in what we now call the lower layers of the protocol. Later, the Host-IMP interface was defined, the so-called 1822 interface, which set the precise spec as to what you could expect from the IMPs and then everything else was left to us. But prior to that, there was no discussion at all. Even after it was known what we were doing, there wasn't any direction except the implicit one, as best I recall: keep doing it. During that period, though I don't remember the precise date, I think Kleinrock took over the contract from Estrin. We were just transferred whole sale. They had offices next to each other.

O'NEILL: So for you there was really no difference, it was just a different PI.

CROCKER: Yes. Very little difference. Estrin remained my thesis advisor, and Kleinrock was on my committee, and they had somewhat different styles but it was all very positive "keep doing it" sorts of things.

O'NEILL: Okay, so then through the middle of 1969, as this group is working and deciding to write things down and starting the RFCs, you had not had direct contact with Larry Roberts or Barry Wessler?

CROCKER: Not that I can recall. I had met Barry at the graduate student conference. I certainly knew who Larry was, but I don't actually remember our first meeting.

O'NEILL: Okay, so now the IMP shows up.

CROCKER: A critical point was how we wrote the RFC and the reason they're called Requests For Comments. I stayed up all night trying to get the wording right to initiate this series of notes. I was most concerned that we not tread on someone's toes, someone assigned officially to oversee these things. So I used wording such as "these are unofficial and they are just to stimulate conversation, and anybody is able to write down anything, and they have no status." There was also an issue as to whether these were formal publications, and so we said they were not.

To emphasize the informality, we called them requests for comments. And that's what led to an RFC. If you had told me then that 20 years later, we'd be publishing RFCs on-line through this medium that we were then anticipating, I would have pulled out a copy of Sorcerer's Apprentice and said, "You've been reading something."

O'NEILL: Yes. It's a very successful way to get comment like that.

CROCKER: That was the model. We proceeded, and it became clear that no one else was working on these and in fact, it was an open group; there wasn't anything closed about it. As new sites were identified as likely to join, we added them to mailing lists. In fact, many early RFCs were changes to the mailing list. That was stopped.

O'NEILL: Yes, you can see the mailing list as it starts to grow.

CROCKER: The mailing list started to grow, and every new mailing list was a new RFC, and we changed that after a while.

O'NEILL: Okay, by November of 1969 you characterized I think it was in the introduction to RFC-1000 that you wrote that ARPA interest was intense by that point, and that you had a demonstration with Roberts and Wessler.

CROCKER: November 21. You bet. Yes, it's a big deal to have a site visit.

O'NEILL: Was this the first site visit? I mean, it's only November, right? You got your machine in September.

CROCKER: I think it was the first one since the machine showed up, and I don't remember if there had been one before that, but we were keen to have something running, without question. We had struggled at UCLA with how to connect the Sigma 7 to the network to the IMP and we called up SDS (perhaps it was XDS by that time), and they wanted a large sum - \$19,000, as I recall - and a lot of time to build a custom interface. We didn't have the time; we were under pressure. Mike Wingfield, another graduate student, took it on. Jon Postel and Mike Wingfield, among others, were graduate students in this same group with Vint and me. Wingfield said he could do it in six weeks, and it would cost four or five thousand dollars. Estrin turned him loose. Wingfield is a classic, clean, organized, careful guy, and in five and a half weeks he had built a box with colored buttons and lights on it. He showed me how it worked. You could sequence a bit at a time through this thing, and I plugged it in and it worked - it was fine. Some of the others and I were responsible for the software and we were, of course, running late. The first IMP was due September 1, which was Labor Day. When the IMP showed up at UCLA, it worked instantaneously.

O'NEILL: Yes, surprisingly fast. I'm interested in what that kind of feedback, what kind of interest ARPA - namely Roberts and/or Wessler, whoever it is - starts to show. I mean, they show up in November. Have you . . .

CROCKER: Their support was enormous. The emergence of the Network Working Group, as we had started to call ourselves, represented a complete solution to the problem that they had not attended to of how to get people to pay attention to the higher level protocols. It was exactly the kind of solution that they like in that it involved a community of people, so that there was a mixture of competitive ideas. Individuals could do things, but also there weren't armed camps saying, "You've got to do it my way or not it all." People were talking to each other. And all this was happening without anybody having to summon people to a meeting or let formal contracts.

O'NEILL: Yes, that was one of my other questions. All of this work was going on

CROCKER: They were enthusiastic; the work was being done within the existing budgets and contracts. I started to get heavily involved. As I gravitated toward getting into this, I started to visit a lot more places, and my travel costs started to go up. There was a contract renewal in progress, and Kleinrock asked me what my plans were for the next couple years. I made up, it seems to me, an outlandishly large budget. It was \$30K a year or so, and he wrote it into the budget. It was clear that everybody was supportive. They seemed to be quite comfortable. RFCs were coming out.

Barry Wessler did intercede at a critical point, later in August of 1970. We had worked on a host-host protocol, and there was a feature that I had been pushing for inside the host-host protocol. The issue was how to log on: what port number to connect to at first? The solution chosen was to connect to a well-known port number, and from that connection you got back, briefly, the name of another port. You shut down that connection and connected to the second port.

I had a more general model, again, seeing a special case as an instance of a more general situation. The more general situation is that you've got three computers call them A, B and C and A and B are talking, and B and C are talking. Now, B decides that its going to take everything from A and pass it to C, and everything from C and pass it to A, and that it would be far more efficient to simply introduce A and C and tell them its intention and have them talk to each other. The purpose of telling its intentions is to reduce it to a model of behavior that B has entirely under its control. Thus there are no access control issues, no authentication is sued. A is, presumably, engaged comfortably in an interaction with B and B is engaged in an interaction with C, and so it is entirely within B's purview to act in this way. It could shift everything from A to C and back. That being the case, it should be okay for it to unilaterally decide to introduce A and C together. I wanted to build into the host-level protocol a transaction, an operation, that says these two are introduced. I called it dynamic reconnection. I wanted to do it in mid-flight, basically, saying just from this point on. One of the visions there was that you might have one machine that handled log-in, like if you

have a service of some sort, and so you connect to this server and it engages in initial interaction to decide if it knows you and log in. Then it may pass you off to another place, which could be anywhere. So that was going to be the way to solve the initial connection problem for Telnet and FTP and so forth. Some tricky things happen, though, if you actually try to take that idea and go down the path. Suppose C is engaged in a discussion with D at the same time and it decides to drop out of the connection between B and D. So now B is trying to tell A and C to talk; C is trying to tell B and D to talk. It has a logical solution A and D should wind up talking, and B and C just drop out. Suppose D and A are in the middle of a discussion.

O'NEILL: It gets more complicated very quickly.

CROCKER: It gets more complicated. Instead of backing away from it, I designed a solution in which B sends a token out to A and C, and if it reaches C while C is in the state of having sent its own token out, then the token just passes on and looks for the end. If there's actually a circuit and the token comes back around the wrong way, then everybody knows that nobody wants to say anything in the future and they just drop it. This was vaguely related to the firing squad problem in terms of having images to draw on, you know, connecting machines together and sending tokens around. So, I thought it was kind of neat. Other people thought it was extremely complicated and that they couldn't implement it, and there was an outcry. Barry Wessler called me up and exercised his authority as ARPA Program Manager and said, "Drop this feature." I was surprised, but that was the end of that feature. That left the initial connection protocol undefined, and then it was resolved in a much clumsier, I thought, ad hoc fashion, although effective for the purpose of opening a connection.

O'NEILL: Okay, so that was the singular time when you got that sort of input from ARPA?

CROCKER: Yes, that was quite singular. I took it on faith that it was appropriate for him to say. We didn't have my idea working, so I couldn't claim it was the right approach. The subject came up because you were asking about interactions with ARPA, and so on. This incident was noteworthy because it was out of character but representative

that they had the authority, and it meant that people were talking and had some cognizance of what was happening. Now that I think about it, we had a meeting in spring. Some representatives from MIT came, and I remember that I used what I had learned at MIT about the Multics system as a model of what operating systems ought to be able to do and process communication. I was surprised when the Multics folks complained most vigorously about the protocol and that they couldn't implement it. I wondered why they couldn't implement it, since I'd used Multics as the model. Some years later, I understood that theory and practice in Multics were not the same, and that those present didn't have access to the lower rings where they would have to have had access to do these kinds of things.

O'NEILL: You were talking to people who weren't at the level who would be implementing?

CROCKER: Well, they were implementing, but they were restricted in what they could implement because of the way the system was structured. Their solutions base didn't match the theory that I had drawn from. Cordell Green came to that meeting. He was at ARPA then. He had been at the AI lab at SRI, and he had been a student at Stanford, and he had gotten his Ph.D. He had been at the graduate students' conference. He was involved with AI work, so I was interested in what he was doing. He had ROTC obligations, and when he went into the Army he was assigned to ARPA. He came out to this meeting with MIT and other folks.

O'NEILL: Now, he came out as an SRI

CROCKER: No, he came as an ARPA representative. I'm trying to dredge the ARPA interactions here, and he watched this meeting disintegrate. We called people from all over the country, and the meeting lasted one half a day. At noon we shut it down; we were getting nowhere. It was a fiasco, a nonproductive event. He and I talked later. He was very mild-mannered and was supportive at a distance. I was very embarrassed about having the meeting fall apart, and he didn't exacerbate the embarrassment. He was a bit empathetic. The failure of the meeting didn't bring to an end the network working group.



O'NEILL: It was just a bad meeting.

CROCKER: It was just a bad meeting, and there seemed to be some perspective about that.

O'NEILL: Okay. You mentioned that you talked to Barry Wessler about being interested, or wanting to know when he was going to leave ARPA. How did you end up at ARPA in 1970?

CROCKER: I believe Barry decided that he was going to go to graduate school, and Cordell had a finite amount of time from his Army assignment. As I understand it, ARPA decided to look for a replacement for one or both of them. They looked into the community and talked to other people as well. I had gotten friendly with Cordell and spent some time with him occasionally. He intimated that I might be approached and suggested that I show up at a particular time. In August 1970 I visited Washington, stayed at his house, and went into the office with him the next day. Sure enough, Larry Roberts offered me a job. I thought about it carefully for awhile, and in September I sent a formal acceptance. From then it took 10 or 11 months. I started work in July of 1971. There was a security clearance process to go through, and I believe there was a hiring freeze at different points in that process. I was young enough that I didn't have a clear idea of how all the pieces worked. But I was working at UCLA and they were paying salary, and what I was doing was what they needed to have done. I wasn't in the office acting as Program Manager, but it was all sort of for a good cause. I was in limbo for this period of time, and it was pretty interesting. One of the effects of being at ARPA is because you're handing out money, you have a kind of instant popularity that can actually be a bit debilitating in a way. Everybody is happy to see you. Everybody has time for you. Everybody returns your phone calls, all the way up through the university presidents.

O'NEILL: So even without actually going there, just being announced or having it known somehow that you were going there

CROCKER: That's right.

O'NEILL: Were you still a student?

CROCKER: I was still a student and was falling get into the classic trap of getting a real job while being a student, and any number of people you meet who have interrupted their studies by taking a job still talk about going back to finish some day. I had already had a non-standard enough academic career that I knew there was a serious risk, that in going to work I had a good chance of not finishing school, and that would not be good. So I stayed enrolled -- paid fees and carried a registration card for the three years I was at ARPA. I had talked with Estrin, and I enrolled for independent thesis work, which was an ungraded subject carrying no credit, but I made a point of paying fees so that I didn't withdraw and later need to reenter or whatever.

O'NEILL: So you stayed in

CROCKER: I stayed in, yes. It was a bit of a problem, because the UCLA calendar was not on my mind, and the fees kept coming up at times when I wasn't thinking about them. More than once I had to call my brother up and say, "Could you go down and pay the bill, and I'll send you a check?" But I wanted to not sever that connection.

O'NEILL: And there was no problem with the fact that you were possibly, or potentially, giving money to UCLA?

CROCKER: Well, that was very easy. I just said I'd have nothing to do with UCLA in my capacity at ARPA. The ethical rules were very, very tight. It was clear that one had to conduct oneself very carefully. So it was evident that I ought to have one kind of relationship with UCLA, not two, and that was no problem. There were plenty of other people to do that kind of work. But just to give you some idea: when it came time to move to Washington, I got in my car and made a cross-country trip out of it. I had known, from graduate student conferences and from network work and other things, a good fraction of the community that we were supporting. Cordell had been the key organizer inside ARPA. I knew when I was coming that I would take over that work, so I inherited a completely

organized program. All the hard work of getting the goals set and of getting the funding organized and everything was all done. There were five main contracts and some small support contracts from specialists in the speech area, people weren't computer scientists natively but were speech people. One of them was a little group called Speech Communications Research Laboratory, SCRL, which at that time was in Santa Barbara. This was a very low-key operation -- modest budget, low salaries, not glitzy or high-rolling at all -- but high quality, a very, very fine reputation, smart people. I think because they were coming out of the speech environment instead of computer science, they hadn't had the status, the notoriety, that computer scientists had. These were people toiling away and doing first-rate academic work.

I stopped in to meet them, and we went out to lunch. I could tell the restaurant was a fancier place than they usually went to. As I recall, lunch was about three bucks a person, not a big deal. When it was time to pay the check, they wanted to pick up my check, and I said, "No, I can't do that." Then I realized that if they couldn't pay for my lunch, then it wasn't a business lunch that the company would reimburse them for. I took this in and said to myself, "This is very good. This is all friendly, and this sets precisely the right distance and serves the right purpose." So I paid my couple of bucks and they paid theirs, and I continued to think of them and still do as first-rate people. But things like that would show up over and over again. So, with that kind of mind-set, knowing that you're dealing with hundreds of thousands or millions of dollars of research contracts and, at the same time, being beholden as a student, there was not even a moment's hesitation that there had to be complete separation.

O'NEILL: Okay. When you got to ARPA, who was actually there? Roberts was still there, but Wessler had left?

CROCKER: Wessler had left. Cordell was almost gone, as I recall. And there were two others who were Air Force officers: Colonel Bruce Dolan and John Perry, and my recollection is that he was a Lieutenant Colonel and became a Colonel shortly after that. So in short order, the office went from being a Director and one Program Manager to one and a half, with Cordell overlapping with Barry Wessler at the end, to three -- those two others and me -- that was the

state of affairs for a bit. So the office was expanding. The other person who was in the office was Al Blue; did you talk to him?

O'NEILL: He was interviewed for the project.

CROCKER: Great guy. So that was it, the five of us. There was no Deputy Director, and there were a couple of secretaries.

O'NEILL: You mentioned that, going there, you knew you'd be taking over the speech understanding from Cordell Green, which was a program he had already developed.

CROCKER: Yes.

O'NEILL: Were you free to develop other programs? Were there other programs you knew you'd be taking over? What was your workload, besides the speech understanding?

CROCKER: Larry asked me what I was interested in. It was clear that the reason I had been invited to join the office was the work I had done on networking, and it was clear to me that the reason I wanted to go was because of the access to the artificial intelligence and other related kinds of work. I was perfectly happy doing networking stuff and had some affinity for it, but I really had strong feelings that I wasn't going to join the office to do networking and let somebody else do the AI stuff. Quite often people join ARPA, as they should, with an idea of developing a new project or having a program in mind that they want to breathe some life into. I came much more as a student, in my mind, and wanted to learn as much as possible, so I didn't come with the idea of restructuring anything or of creating anything so much as working in the structure that was set up. And my initial reaction, as I said, was that a lot of very fine work was going on and in all the right places, so this was a front-row seat on the best show in town. I was

delighted to do that. I didn't come with an agenda of things that I wanted to make happen. As far as I could tell, the right things were happening and my job was to cheer it on, facilitate it, and be supportive.

O'NEILL: How clearly defined for you was the role of Program Manager? Did you know what was expected?

CROCKER: No. I don't know how well-defined it was at all. Larry was in a line that had included Ivan Sutherland and Bob Taylor, and Licklider before them. And I guess that's it. I don't know much about how Lick had operated, but Ivan and Bob Taylor and Larry were very broad-based people who pretty much had everything organized in their heads and ran it as a one-man show. They had direct relationships with all the contractors, knew what the work was, and would go and sell it in the Pentagon. They didn't have a large staff helping out, so Barry Wessler was a big addition in Larry Roberts' staff, probably. So, to first order, Larry had a complete program and ran it himself, and my job and the others' jobs were to take on whatever was appropriate, but there wasn't any question that Larry had the overall picture. It was not as if he were an office administrator and each of us had territories and did as we pleased and he just made sure that we adhered to the basic rules. The technical objectives and the vision were completely his. So it wasn't clear at first what the division of labor ought to be, and I guess I viewed my job as understanding what the goal structure was, making myself useful, and doing what I could to be helpful. That meant looking in on the various programs.

O'NEILL: Site visits?

CROCKER: Site visits. To sort out the problems I could and take to him the ones that I couldn't. I remember that it was a pretty unstructured environment and knowing he was there was one of the anchors: if I made a mistake, people could come and appeal to him. It's scary out there if you're all alone, you have all that money and power, and you're trying to do something positive but it's easy to make mistakes -- mistakes in judgement of a technical nature or misunderstanding how the system works or giving somebody the wrong advice or whatever. I remember getting a strong sense of security from knowing there was recourse and back-up. In a small number of cases there were hard

decisions to make in that I wanted to shut off some work that I thought was genuinely unproductive. Of course, the people affected by such decisions think that you don't know what you're doing and that you misunderstand and, in my case, that "you're just a kid anyway," and so on one or two occasions there was an appeal back to Larry.

O'NEILL: So you were free to deal with the various contractors around the country and evaluate them on some level, anyway, and make recommendations to Roberts as to what should be done.

CROCKER: Yes. Yes, within limits. And the limits were never clear, but as I said it was a very small office. It never occurred to me to try to build a game plan and not tell him what I was doing. So if I thought there was an issue, I would tell them ahead of time.

I did get into an interesting trap once. I'd been a student at the AI lab at MIT, and Minsky had a variety of different mechanical arms, none of which worked very well. They were clever, and I had tried to program one of them as a graduate student. None of these were masterpieces of mechanical engineering. At Stanford, they finally built an arm that was a credible, that moved slickly, smoothly, and they had the motion equations, and so forth. But this was before that time. The Draper Lab was doing some work on arms, principally as part of the space program, and wanted to build an arm for us. I got a proposal from them and was trying to fund them to build an arm that I was going to give to Minsky sort of present it to him. Minsky wasn't too pleased with this, but I was not too pleased with the progress he was making on arms. This was an area where I thought I might be helpful.

The thing that caused great surprise was that I wrote up the funding memo requesting an ARPA Order and sent it to Larry for signature, and he bounced it back. It was the first time he had ever bounced anything back. I said, "Well, what's the problem here?" He said, "This says we're building robots. It turns out that we don't have, anywhere on record, that we're building robots." I said, "What do you mean?" We were building robots all over the place anybody could see we were building robots. I just never caught on, or bumped into the fact that, at that point, before AI had become a big deal, that Congress had no idea. So these were "intelligent systems," I think, was the name of the program. Under that name, there were various kinds of military justifications about future manned and

un-manned airplanes and command and control systems, and so forth. But robots actually never showed up in any of that paperwork, he informed me. It had never occurred to me that we were being cagey in any way, and so I agreed that, yes, that could be a problem if we now, all of a sudden, have on record that we were going to build an arm. Those are the kind of issues for which a boss who knows what is going on is a big help. Mainly we had a fairly close relationship. I had a set of things that I tended to: there was the AI program, there was automatic programming, speech understanding, and a portion of the network work. During this period of time, the network was growing, and it was a lot of work to add each node to the network. I would get involved because of the time that I had spent and the expertise. Bruce Dolan took care of all the mechanics of working with ordering lines, working with AT&T, but I continued to work with a lot of the new host sites and give them advice, get them connected, and so forth.

The first year I was at ARPA, I wound up going to Europe three times, and I found that to be very taxing. Not just the transatlantic travel, but having to deal with different cultures and trying to get organized. And a lot of it had to do with networking; I was invited to talk. When I got involved originally in networking, I got invitations to talk nearly everywhere. So in addition to travelling to different sites, there were different meetings. It was a hot topic; everybody wanted to hear about it. This continued at ARPA, and it threatened to continue internationally. And I thought, this could be, should be, a lot of fun, but if I do this I'm not going to get anything else done, and I need help. So I called up Vint Cerf, who at this point was at Stanford as a new professor, and I said, "Vint, we need an international network working group, and you're the Chair, and I'll ship you a few bucks, and you're in business."

O'NEILL: And from your perspective that was a way to keep you from having to go over to Europe to discuss these issues?

CROCKER: It wasn't just that I didn't want to go to Europe; it was a question of how to deal with the pressure for international networking and how to spin it off and get somebody else to take care of it. At ARPA, you have money and no time, and so a solution represents transforming money into somebody doing something. The more you can

get somebody else to things, the better off you are. You only have one source of leverage, and you just use it to the hilt.

O'NEILL: Let's see, when Roberts left . . .

CROCKER: He and I overlapped two years.

O'NEILL: Two years. Okay, so that was long enough for you to start feeling pretty comfortable at the office.

CROCKER: I was comfortable pretty quickly. I was comfortable, I would say, right away. I arrived on July 21, I think, and I remember asking what kind of rule he wanted in terms of approving travel, whether he wanted me to tell him ahead of time what my plans were, and he shrugged it off, saying do whatever makes sense. Ten days later, I was in Oklahoma City, at Tinker Air Force Base, trying to help them. They were attaching two machines to the ARPANET to do a comparison against AUTODIN I . . . one at Tinker in Oklahoma City and one at McClellan Air Force Base in the Sacramento area. I don't remember how the trip came up, but I had said I'd go. Either on that trip or another, I got Jon Postel and maybe somebody else to come join. It was clearly a worthwhile thing to do in terms of making the ARPANET look important in the Defense Department. I don't know who organized that but just pitch in and do something useful, and if they needed some help on what those protocols should look like, we'd go give them the help. So I was pretty comfortable. One would pick up the files and read what was going on. Most of the research places I knew something about. I'd been to Carnegie looking for schools, I'd been to Stanford a few times, SRI and Utah -- these were among the key places.

O'NEILL: So, even the areas that were outside of networking, you knew the people or had been to the places.

CROCKER: Yes, in fact, you see, nobody was in networking at first.



O'NEILL: But you knew these other people because you had visited while looking for grad schools and just through other contacts.

CROCKER: Yes.

O'NEILL: Let's go back to the other Program Managers when you first arrived and then we can spread this out through your time there, as well. What kind of relationship did you have, were Dolan and Perry doing the same kinds of things you were doing, or were they covering different areas? Were they more involved in [military] transfer programs?

CROCKER: We did have a rough division of labor. We didn't have any problem knowing who was in charge of any given contract; that was pretty clear. The areas that I was immediately in charge of were the AI and its cousins; there was intelligent systems and automatic programming, where there was a defined topic that had grown out of intelligent systems that was trying to look at a range of topics from advanced compilers to program synthesis of various flavors, and then there was the speech understanding work, which was, again, another very heavily AI-influenced program. There were also some aspects of networking, and I think that was the initial assignment for me.

At the same time, there had been a question of long-term changes in the climate. I don't actually know what the triggering questions were, but they were the same kinds of questions being asked now about possible long-term trouble. I don't remember whether it was triggered by nuclear testing issues originally or by other things, but a program was initiated to run climate models that was very heavily computation-based. These were related generally to weather models you start with the same basic physics, but you tuned the equation to be stable over integration of hundreds of years instead of over hundreds of hours: it's not so important whether or not it's going to rain tomorrow; it's much more important whether there will be any ozone left in a hundred years. Those kinds of programs use enormous amounts of computer time and are also very, very sensitive to the numerical analysis. This kind of work is not easy at any level, and there was no better home inside ARPA in those days than our office. We didn't

have the meteorology background, but we had access to and knowledge of the other kinds of activity. John Perry was a weatherman. He had a Ph.D. in meteorology from the University of Washington and had a career in Air Weather Service. He was an Air Force officer and a weatherman, so that was his program; there was no question about that. There were some thoughts about trying to start activity in automated manufacturing, and he spent some time trying to formulate some of those plans, as I recall. I don't remember precisely what else he did, but he was a smart person. Getting a Ph.D. in meteorology requires a fair amount of mathematics and other physical science subjects.

I think Dolan had a doctorate in electrical engineering from Stanford — another bright person; these were first-class Air Force officers. These were also the first military people I had ever spent any time with and I thought, "This is nothing like what I was reading." The Viet Nam war is in full gear, and the military was not looking very good. But this was certainly a different class of people in terms of education. They had a very reasonable, mature outlook. I developed a lot of respect for the military establishment. Students, who were having protests, may have had the right idea about some large issues, but they weren't very mature as people.

Dolan was doing a lot of work with the network, and we had programs in distributed operating systems, various kinds of distributed systems, that weren't perfectly organized in terms of focus. You get the work and then you staff those programs with the people necessary to get it done. They represented more grouping of existing work, with each effort reasonable in its own right and with some history for why it was being done. Then, for purposes of presentation in the budget and congressional testimony, they would be fit together. If you fit them together and packaged a program, you said, "Well, we're doing work in distributed systems that consists of these several kinds of things, and if you'll read carefully you can see each one of these things that you're doing — all of it seems reasonable as part of a larger program — represent these several different contracts."

O'NEILL: Okay. But it wasn't "this is what we want to accomplish . . ."

CROCKER: No, it wasn't a top-down, single-point

O'NEILL: Good, because some of these programs are hard to get a handle on.

CROCKER: That's right. And there's no criticism of that. If you have too strong of a point of view, if you choose only one way to structure things and say that we're going to start with a blank slate every single time we start a new program, you have chaos. You need a certain amount of continuity. The ideas that create new programs come out of the researchers in the field, principally.

O'NEILL: Okay. So, you would say that principally the ideas came out of the research community.

CROCKER: Yes. There's a question as to when you date the formation of an idea. Some are easier to date than others. I wasn't around at the formation of the speech understanding work, but it's my impression that at some point the idea of working in speech understanding took hold, and I have a very strong impression that it was discussed at a principal investigators' meeting in 1970 or perhaps a year earlier. I was invited as a Program Manager designee in January 1971. The Speech Understanding Program was already committed to in that the program was being formulated, Cordell was in the middle of writing it, and AI Newell had chaired a study committee to write the report. All that work was in progress. And I suspect that at the prior PI meeting, when the idea was in its formative stage, they were trying to formulate a program. So there's one that came out of the community, I suspect. Later, I actually tried to formulate a brand new program, called National Software Works, and the origin of that was the widespread feeling that software was in trouble and the question of what technology we had that could be helpful. The only technology that we had that was near application, not just an idea or basic research, was the network. So I focused on how could we use network technology to do something about the software problem and initiated a series of meetings to formulate a program resulting in the National Software Works, which resulted in a lot of interesting ideas and some interesting results but overall failed. So that was something that sprang from me at ARPA, although it involved interactions with lots of other people before it got going. In the Speech Understanding program, Raj

Reddy's work at Stanford, I think, was noteworthy, and he showed that you could get a certain distance that made people sit up and take notice. I think that was more a case of an idea springing from the community and then getting shaped by community interactions and from inside the office and resulting in a mission.

O'NEILL: Okay, office interactions. There are three Program Managers, and there's the Director, and there's Al Blue; are you getting together as a group and discussing these things? Is Roberts just going to each of you separately or, you know, interacting as you need to, or what?

CROCKER: We tended to eat lunch together daily, so we had plenty of time to interact. In those days, I don't believe that there was much joint work, although there was some. I got together with John Perry, and we went on some trips together where he needed networking help in his program. The same was true with Bruce Dolan. Networking really changed the character of the office, because now there were all these cross-currents to connect people together, and networking became a pervasive technology. So there were probably interactions for that reason, but it was not so common for us to get together and put joint programs together out of equal pieces. I don't really have a clear recollection of what the interaction was with Roberts. My view was that he had a complete picture of everything, and he delegated as much as he was comfortable with and as much as you could handle. In other words, a fair amount. There was so much to do that there wasn't a big issue there, at least from my point of view. Now, Bob Kahn arrived in November of 1972.

O'NEILL: That's exactly where I was going to lead this.

CROCKER: Bob came in a somewhat different way in that he had thoughts about what he wanted to accomplish, and I think he was keenly interested in automated manufacturing—factory automation, I think was the buzzword. That idea didn't get off the ground, because, from my perspective, there was difficulty in selling the idea at the higher levels of ARPA and the Pentagon. There may well have been some problems about getting coherent ideas of what

would you do if you were going to do this. So my guess is that it failed some number of sanity checks, technical and political, and so it didn't get going, and I think Bob was disappointed, at least for awhile.

O'NEILL: Yes, certainly he says he was. But that's why he had gone there. We didn't go into detail about what killed it, exactly, but it stopped.

CROCKER: Yes. It just didn't materialize. I wonder what Lukasik would have thought, would say about that, because it couldn't go anywhere unless Lukasik would know what was happening.

O'NEILL: Yes, but I didn't ask him about the automated manufacturing. Did you have a direct relationship with Lukasik? I mean, when he was Director of DARPA?

CROCKER: Yes. We overlapped considerably. Before I cover that, I was going to there was something else that was

O'NEILL: Oh, okay. Yes, we were talking about Bob Kahn.

CROCKER: There was a major demonstration of network technology down in Washington.

O'NEILL: Yes, the 1972 ICC??

CROCKER: Bob had been central to making the ICC happen, and he came immediately after that, that was the timing. He's an enormously talented and smart person, and I remember asking him when he came what he wanted to take off my plate and what he wanted to be copied on, in terms of electronic mail, which we were using heavily in those days. He said he was interested in everything, not to worry about flooding him, and so I made him a regular 'cc' on things.

One topic that had come up was narrow-band speech, speech compression, which has some relationship in technical terms to speech understanding but was focused on cramming speech, not on understanding it so much. I sat in on some meetings. A lot of money in the upper levels of the Defense Department had become freed up, because of something related to SALT talks. Somebody had decided that the money would be applied to speech compression and that an interagency group would be set up to formulate the program. We already had our own money, and we were pushing speech compression a bit, so we didn't need to vie for that money, but we did have an interest in participating in the working group so that it didn't go awry and so that we could be helpful to it. I remember sitting in on some meetings that involved Army, Air Force, NSA, and other personnel — DCA people, I think. A lot of our time was spent on how to distribute the money. The paperwork problem, getting contracts in force, and getting things approved all took a high percentage of our total time. The romance and the thrill of the job were the technical ideas, the people that you got to interact with, and the grand vision, but a lot of the day-to-day work was moving this slough of paper through the mill and sorting out these problems. The reason I was thinking of it is that in interacting on the speech compression work, there was money at the Air Force and, after some number of meetings, they all decided what they wanted to do and who was going to do it. Somebody wanted to have some work done at SRI, where we had a contract in place with the speech understanding work, and I wound up either assenting or volunteering to have that money come through our contract as an add-on. It was much, much easier to extend a contract, add on to it, and increase its scope than it was to create a new contract. Plus ARPA has contracting authority that is unique. This is a very big deal, if you haven't bumped into this. Every other agency, to my knowledge, generally writes its contracts with its own people.

O'NEILL: But ARPA used other contract people?

CROCKER: Right, and the "using" had a number of aspects to it. I suspect that it was set up partly to facilitate technology transfer, but it was also set up in a way that gave ARPA priority. To create a contract, you get your file ready for a contract you want to put in place. You identify the money, get the proposal, and go through whatever

hoops you have to. Then you turn it over to a contracting officer with a statement saying, "Here's my money and the proposal. I've done all the work I need to do. Go make a contract between the United States government and that institution." The process after that takes an extremely long time; it takes anywhere from three months, at the shortest, to nine months or more. I've never fully understood where all of that time goes, but it causes the primary objective to become how to maintain continuity and make sure there aren't gaps when people don't get paid, etc. For example, you must be getting some money from ARPA to do this, and the time at which that contract started must have been an issue, and it probably didn't happen on the day you turned in your proposal.

O'NEILL: Right. Before the money comes in there's that "blank" period.

CROCKER: Right. How long was that?

O'NEILL: I wasn't involved in the project at the time, so I don't know, but it was more than a few weeks.

CROCKER: Yes . . . probably a fairly small contract. Under \$10 million or something.

O'NEILL: Right.

CROCKER: So the contracting issue got to be quite messy, getting this little Air Force money to join our contract. We sent the money to the Army Research Office in Durham, which was the contracting shop that had the SRI contract, and they raised all kinds of issues: "Is this your money?" and "It will fund work too far into the future" and "We don't want to do it" and "We're too busy." My response was that it was an ARPA effort and they didn't have a choice about it. I thought all the legal issues were identified and resolved, and what finally happened was they found yet another reason that had not surfaced in the initial discussions, and two weeks before the end of the fiscal year, they bounced the money back and said, "We aren't going to do this." We had an emergency meeting at ARPA in June of 1973, as I recall. The newspapers were full of stories about the money trail and \$25,000 in cash in suitcases

being carried from Minnesota through Mexico, and meanwhile, we've got a few hundred thousand dollars that are about to evaporate -- go back to the U.S. Treasury, which is the same as burning it when you're inside of the agency.

We managed to send that money over to NASA to support the ILLIAC IV and to take ILLIAC IV money that was going to appear two weeks later in the new fiscal year and send it to the Army. The Army was willing to deal with next year's money because that came under different rules. The net effect was exactly what was supposed to happen: we got the right amount of money on the speech compression work, we got the right amount of money on the ILLIAC IV and, fortunately, I don't think my signature was in any critical place in all of that. But it really did color things. ILLIAC IV was what John Perry was doing. The ILLIAC IV was what we informally called ARPA's Viet Nam. It was a money sink. It took an enormous amount of money and kept drawing us in and drawing us in. That, I think, all by itself accounted for the bulk of John Perry's job.

O'NEILL: It would.

CROCKER: Big, big project. Lot of issues there. So one thread to follow here is the contracting issue, which is worth highlighting, because it actually colored everything. Contracting issues are a major aspect of how things get done or don't get done and how you choose to go about doing things. You think of having pure technical visions, but this other aspect is crucial. Another thing I can talk a bit about has to do with Bob Kahn. I got to talking about this speech compression work because that was a natural area for him to pick up on and, in fact, he did and Dolan turned it over to him. I knew I had to pinch-hit for it for a few months until he was on board, and then he picked it up. He was also working on the factory automation. I'm trying to think what else he directly worked on -- a lot of networking issues.

O'NEILL: Well, I was going to say he, at some point, gets into the Internetwork

CROCKER: Yes. He got into the Internetwork. There was also



O'NEILL: The packet radio

CROCKER: Well, yes. The packet radio and, in general, what we called the C2 program. DARPA is an agency that has 6.1 and 6.2 money those are meaningful terms?

O'NEILL: Yes.

CROCKER: At that time, our office was the only one, I believe, that had both 6.1 and 6.2 money. The other offices, the big offices with 6.2 money, were the Strategic Technology Office and Tactical Technology Office. They were oriented toward big military systems. Then there were some defense/science-oriented offices.

O'NEILL: Material Science?

CROCKER: Material Science was one of them, as was Human Resources Research Office, which was social science/psychology oriented. I think those were the other two at that time. They were basically small and had 6.1 money. We were the only office at the time that had both 6.1 and 6.2 money, and programmatically it showed up as two separate programs. So if you stripped away the org chart and the office structure and just looked at the budget, all you saw were budget items, line items, at a certain level. It just happened that the one 6.1 program and one 6.2 program were being managed by the same set of people, but programmatically they were two distinct programs. During those days, IPTO was very favored, and Lukasik and Roberts were very much on the same wavelength. Lukasik found opportunities to be helpful to Roberts and to help the program grow. I believe some time in the year that Kahn and Roberts were together, they created a third program element, which was another 6.2 budget, as a way of getting a next chunk of money that was substantial, rather than simply growing the two existing budgets. There were limits as to how large one could grow those without having the staffers in Congress, who would read the congressional testimony year after year, say, "But this program is only a \$10 million budget program. Are you asking for \$50 million? That doesn't make sense." (Those numbers aren't accurate, but the effect is right.)

O'NEILL: So, rather than growing, you want a whole new program element.

CROCKER: You grow the ones you can, but they really wanted to make a big step, and so the method of doing that was to create a brand new thrust, based on networking and its impact in military technology and so forth. Packet radio was one piece of that, and satellite technology was another piece as I recall. There were some other things. Bob Kahn worked very hard on that kind of thing.

O'NEILL: You mentioned that Lukasik got along well with Roberts.

CROCKER: Yes.

O'NEILL: Can you describe to me what happened when Roberts left? I mean, there was a period of time when there was no director, when Al Blue was the Acting Director.

CROCKER: It was chaos. It was a near disaster. Larry had been in place for, I think, six or seven years?

O'NEILL: He had come there in 1967 and left in 1973.

CROCKER: It was as if he had been there forever, and nobody was prepared for his leaving; there wasn't any machinery in place. It was very hard to recruit people out of the community, in fact. It wasn't easy in those days to go get somebody to decide to be a Program Manager. I was odd, in that sense. Most of the people who were quite good were interested in being professors, getting jobs at good institutions, and so forth. We did not have a Deputy Director, and I remember that Lukasik was harassing Roberts about that and urging him to get a Deputy Director, and I don't think Larry really wanted a Deputy Director. I mean, this is someone who had run the thing from scratch himself, and it was pretty good, from my point of view, working with him as a junior person, but I don't think he would

have been as comfortable with somebody who was close to being an equal and challenging his territory. So Larry's leaving left an enormous gaping hole and no clear successor. Now, Kahn, in my estimation, probably should have been put in place. However, Lukasik and Kahn did not get along all that well. Lukasik did not have a high degree of confidence in Bob. I'm not sure what the reasons were. Their styles were quite different. I don't know what the chemistry was, but it's clear that it was true, and Lukasik had a policy of having entrance and exit interviews with everybody who came on. Exit interviews are fun, because they're a free shot. I tried to figure out what I was going to say, and I spent my time telling Lukasik that I thought he was misusing Kahn, that Kahn was a major asset and that he ought to think very highly of him, and so forth. I don't think it had any real effect, but it was a chance to speak my mind. So it left an enormous vacancy. We tried to recruit another senior person. I don't remember how many different people we talked to, but we went through various names. In one case we made use of some recent conferencing technology. That was kind of fun. One person took the prospect quite seriously in some late stage of the discussion, and we wound up in a six-way network conference, using a system called Forum. The interviewee was at home on the West Coast. This was a Sunday. I was at home. I think Kahn was at home. Lukasik, Tachmindji, who was the Deputy Director of ARPA, and Roberts were all on-line at various places. This was a paragraph-oriented system; you'd type, and a whole paragraph would come out, identified by the speaker.

O'NEILL: More of a conferencing system.

CROCKER: Yes. It was a conferencing system. We had no voice contact, and our terminals in those days were SILENT 700s, principally, so we had paper scrolling out. The paper actually had the advantage that you could look back and see what was happening. It was really a neat experience because several threads were going on in the conversation, and you could keep track of them, and the bandwidth was actually higher than if we were all in one room where only one person at a time could talk. Here everybody could type whatever they wanted and the paragraphs would just come trooping out. You can read faster than you can type, so it was more exciting than you would think. This person was at home, in his den, with his wife looking over his shoulder, and asking a lot of questions. Tachmindji had not had any experience with this conferencing system, and there's a way to send

messages privately. I was giving him a moment or two of coaching about how to do something so he could get into it.

In the end, the interviewee declined, as did a number of other potential candidates. Licklider, who had started the office and had no more long-term career objectives, was an elder statesman. He agreed to come back for the good of the community, and he did. He got run over, pretty much, by everybody - by Lukasik and by the rest of the forces. His style was kind of relaxed and academic, and that might have worked in the 1960s, but the Office had evolved, and with the amount of pressure and the money, his style didn't work at all. Dave Russell, an Army colonel who had been a Program Manager in the Nuclear Monitoring Research Office, was installed as the Deputy Director under Licklider and essentially given responsibility for day-to-day operation of the Office.

O'NEILL: Okay.

CROCKER: Russell was a very mild-mannered sort of guy. On meeting him, you'd think there was nothing rough about him, but there's something taught in basic training about command presence and how to give orders to people who are marching. And in a room not much bigger than this, we had a little staff meeting with six or seven of us who were around, and he used that tone of voice to whip us into shape. And I thought, "This is interesting."

O'NEILL: So, this was even before he became the Director?

CROCKER: Yes.

O'NEILL: This was when he came into IPTO

CROCKER: Yes. This is not uncommon when you have an organizational problem: somebody is put in a pivotal position, even if there is a titular head in power, in whatever way is appropriate, just to make things happen. He

clearly had orders to be the Chief Operating Officer, to run the day-to-day operations, make sure the paperwork flowed. There were a lot of problems with the paperwork: if you get behind on the paperwork and the money doesn't get spent or the programs don't get justified properly, things start to unravel. There's a certain amount of factorylike operation you have to have.

O'NEILL: I was curious you mentioned about this conferencing system, with you and Kahn and Lukasik and the Deputy Director, Roberts.

CROCKER: And the candidate.

O'NEILL: and the candidate. Were Perry and Dolan gone, or were they not involved in this.

CROCKER: That's a good question. They weren't involved.

O'NEILL: Were they not involved generally with trying to find a replacement and, if not, why?

CROCKER: I don't remember the precise dates of when they were there. It could well be that one of them was gone.

O'NEILL: Okay.

CROCKER: Also, neither one of them was connected to the research community the way Kahn and I were. We had sort of grown up in it or lived in it a lot. Perry came from outside and was dealing with a certain set of computer applications but was not a native computer science researcher, and so the community he spent a lot of time with wasn't the community to generate a candidate. And Bruce Dolan was really much further away from it in many ways. He had been an intelligence officer, and he had a lot of skills, but they weren't the kind that were going to generate a strong candidate.

O'NEILL: Did you sit around and come up with a list of people, "you" being Kahn and Roberts.

CROCKER: Yes, at various times, but mainly as input to Roberts. There were multiple lists. There were times when we'd sit down and put names together and cross them off for one reason or another, and sometimes there would be someone we'd pursue. I think Keith Uncapher was a perennial candidate from the office's point of view. And he would refuse and refuse and refuse.

O'NEILL: Did you, personally, try to recruit anyone for Program Manager or Program Director?

CROCKER: I called up one person one day. I didn't know him, but he seemed plausible. Larry Roberts had said he wouldn't call him, so I volunteered. I very much liked this person, and I called him up. I knew him a little from previous interactions, and I called him up and said, "I want to work for you," knowing that he would take it to mean work for him at his place. But quickly he understood what I was trying to say, and he was amused, but it didn't have any further effect. He appreciated the sentiment, of course.

O'NEILL: How long did this whole process take? How long were you trying to find someone?

CROCKER: Well, it was Roberts was what happened here? Roberts announced in the beginning of the summer, June or something like that, and I guess he had his arm twisted to stay through the end of the summer 1973.

O'NEILL: Yes, because he left in September of 1973.

CROCKER: I think he intended to go faster than that. I think he had his arm twisted a little bit. It was the summer of 1973.

O'NEILL: Okay, and that's really the period of time when you were actively recruiting for a Director.

CROCKER: Yes.

O'NEILL: Okay. You mentioned starting up a new program, National Software Works, as an example.

CROCKER: Yes.

O'NEILL: What kind of criteria did you use to decide when to start something up?

CROCKER: Well, we certainly didn't have anything written down about what our criteria was for anything. I remember one day GAO was reported to be coming over to find out why ARPA was so successful, and that sent a shudder and a shiver through the place, as it always does when the auditors are going to come or the Inspector General's going to come or anybody's going to come and look at your files for any reason. It's not that you're doing anything wrong, but it's like the IRS auditing you because they're looking for an example of somebody who does it right. Somehow that doesn't make you comfortable. Sure enough, the GAO came over and looked in our files and said, "Gee, there's no evidence that you're successful here," which is just what you'd expect in an exercise like that.

There was a feeling that the office was growing and that the opportunity was there if you had a good idea. And there were plenty of examples, like the speech understanding question and the narrow band speech, which was in a different area, and it was clear that if you had a coherent idea there was an opportunity to go sell it and make something happen that was there. So, the criteria were vague in a way and, on the other hand, reasonably clear. There had to be a pretty clear idea of what the payoff could be. Suppose you did the work that you said you were going to do; how would it make a difference? You didn't need to have a spreadsheet that showed you were going to improve something by a certain amount. All you really needed was a reasonably coherent story.

O'NEILL: A difference to whom?

CROCKER: Difference well, to the client, and the client is the military.

O'NEILL: Okay.

CROCKER: You see, the Mansfield amendment was in place; everything had to be justified.

O'NEILL: So you had that in mind . . .

CROCKER: Oh, yes. Nowadays there's a lot of talk about a national ARPA and dual-use and so forth. And we were quite conscious at the time that we were creating technology that was generally useful. But you didn't really have a prayer of starting something up unless you could show that it was good for the military. On the other hand, technology is good for the military in very broad ways, and so it wasn't very hard to do that. But you had to have a reason.

Another Story. One day Lukasik got very angry. Lukasik had a temper; he was a flamboyant type of person. It's part of his charm, provided you don't get stung by it too often. He had a staff meeting. Larry Roberts went to the staff meeting, and my understanding is that he had come across a piece of work that ARPA had supported, and the credit line said the work was supported by ARPA. It was work done in graphics, and it was automation of Labanotation for dance. So you feed in dance notation, and out comes automated choreography. He had quite a fit over that, because he recognized that there was an enormous threat to DARPA if Congress, Proxmire for example, with his Golden Fleece award, got wind of it. It was a perfect candidate. What are we doing spending Defense dollars on automated choreography? What's the justification here? Word came down that he'd fire the next Program Manager responsible for a project like that. Now, I was responsible for the MIT, Stanford, and CMU AI labs among other places. I knew very well what went on inside those laboratories: if you called them up and said, "I'm coming by for a site visit,"



you'd see one story. If you'd drop in at 2 in the morning, just casually, you'd see the place going full blast with everyone just doing as they chose. I went to Stanford one day, and they were doing music synthesis. They had four speakers and automated synthesizing music, and they had a fly buzzing around that you could hear in four dimensions, and they were very proud of it. I asked how much computer time was going to this. And they said, "Well, it's using up about a quarter of our computing cycles," and they had accounting records to show that. Then in the daytime I'd come back and their management would tell me how they needed more computers, and I'd say, "Yes, I can see that you need more computers." So, it turned out this particular piece of work was not done in a place that I was responsible for. It came out of a graphics operation at the University of Utah, which was doing first-rate, advanced work in graphics and this was a graduate student using the facilities there. I thought "Gee, how did I miss this bullet?" I was the biggest target, given the set of projects that I had, and on that particular day I was lucky.

O'NEILL: I had asked, originally, about how you criteria

CROCKER: Oh, criteria, right. It was clear enough that you would not start a program whose output was going to be things like that. That was a loser from the start, in terms of justifying it and getting a program. On the other hand, you might be charmed by that kind of work, but you had better have some other things to show for it. It's a very rich set of opportunities, because the computing technology is good medicine for a lot of defense problems. It wasn't really hard. If you thought that something was technically interesting, there were quite often defense applications.

Let me now tell you about the NSW.

O'NEILL: We've looked into some of these programs, like National Software Works, and there's not very much written, and we can't find very many records, so anything you'd like to add ...

CROCKER: Right. Well, this was a program that only went a certain distance and finally stopped. As I said, the genesis was this notion that software was a big mess, as it has continued to be, and that single thought has spawned

a number of programs. There was the landmark study that Barry Boehm headed, CCIP '85, I think - and then there was the Monterey Software Engineering Conference that Jack Goldberg ran. There were a number of studies and then various DoD initiatives. The STARS effort, currently, is in some ways a direct descendant of the same impetus, and the ADA program, so there's a whole series of attempts to "fix the software problem" one way or another. And as I said before, I asked, "What do we have on the shelf that can be helpful?" I said, the network. Is there anything in that equation where the network is a good thing. We had a series of meetings in, I believe, 1972 sometime. People in the community, like Balzer and Cordell Green, tried to understand a bit of what the software problem was, in meetings with military people, and found it hard to find military people who knew what the software problem was or even knew what the military was doing. We found the research people in the military were, in some ways, no closer to the real military than we were. But out of three or four meetings, we formulated a rough plan. The idea, at the time, was to encourage the development of quality tools for software development by creating an incentive in the form of a marketplace. This marketplace was going to use network technology to improve access to the tools. This was before personal computers, before PDP-11s were prevalent, and so computers were big pieces of iron that sat in one place. The plan was to use the network as a highway to interactively get to these tools and be able to piece a software environment together out of tools that ran on different machines and used the network to move files back and forth so you could compile in one place and execute in another and things like that. It was totally flawed, as it turned out, because computers became relatively cheap, and you could buy a PDP-11 that had a richer and better software environment than any of the mainframes had, and that trend has continued. But it was ambitious - interestingly ambitious - in that we wanted to actually set it up as an automated marketplace so that money would change hands and people could have proprietary programs and bring them on line.

O'NEILL: So, as a programmer, I could have a package that would be available through this National Software Works, and someone who used it would then pay me?

CROCKER: Right. On a usage basis.

O'NEILL: Okay.

CROCKER: A meter would "click," either through invocation or however you wanted to do it. There would have been different pricing strategies. That was very exciting in that [it would now use the network in a way that was fundamentally different than if you didn't have a network.] You'd also be able to move files on demand from machine to machine.

The main contractor was Massachusetts Computer Associates (Compass). Steve Warshall was a key person. I liked his ideas a lot. He came to some of the initial meetings. Based on the ideas that Compass and others presented at those meetings, we developed the rough ideas.

I sketched the idea out for Lukasik, and Lukasik had one very firm reaction: he said he liked the idea, but that this was the kind of program that needed to be done jointly with the services, for tech transfer and other reasons, and that if I wanted to proceed with it then I needed to form a partnership with one or more of the services. Furthermore, he'd been burned in prior relationships with the services, and when he meant a joint program, he meant that their money came in tandem with ours and not, "We'll join you next year, if you get it started."

So I then set about trying to make friends with people in the Army, Navy and Air Force. The Army and Navy interactions did not go anywhere, in that I couldn't find anyone who got excited about it. In the Air Force, I got an audience with a two-star general, General Robbins, who was on the Air Staff, and he was in charge of computer matters within the Air Staff. Directly reporting to him were three operations: one was the Federal Computer Simulations Center, FedSim; another was the Air Force Data Services Center, which is the most visible Pentagon computing place, with a Multics system that's been installed there for years; and the third place was the Air Force Data Systems Design Center, in Gunter Air Force Base, in Montgomery, Alabama. They build the software running at all the air bases that does base processing, a lot of which is pretty mundane, such as personnel and payroll, housekeeping functions, and so forth. When I first met him, Robbins ushered me into his office and he started out

saying, "I've always admired ARPA," and I expected to hear that it was because of the great technical achievements or how smart we were, but he said, "You guys are so good at moving money and getting contracts." I said, "I see." So we sat down and we talked, and he said, "This is interesting." In fact, he had a couple of people in his office. I think his Deputy was Jim Burrows, who is now at NIST as the Director of the laboratory there, a key computer person. Robbins said, "Talk to my people, and if they like it, we'll put a deal together." He sent me to his people, both in the Data Services Center in the Pentagon and at Data Systems Design Center in Montgomery. The colonel in Montgomery I don't remember his name caught on instantly to the idea of a network. One of the things he saw was that they had a lot of Burrough's equipment all over their bases, and they would find out things about the Burroughs equipment that were not documented. He saw that if they had a network they'd be able to trade this information around quickly and keep up to date. I mean, bells went off his eyes lit up. So some junior personnel, young officers in the Pentagon and also in Montgomery, were assigned and we had a number of meetings. One of the people in the Pentagon was named Bill Carlson; I think he had a Masters from MIT and had Air Force obligations. He wound up coming to ARPA after he left the Air Force and running the Ada program. He became one of the key people inside the government in the Ada program.

O'NEILL: This was a Bill Carlson?

CROCKER: Bill Carlson. Now there are two different guys: William Carlson and David Carlstrom.

O'NEILL: Okay. That's why I was getting confused.

CROCKER: Right, and we had all kinds of trouble. But it's Carlson and Carlstrom. Dave Carlstrom was a more senior Air Force officer -- a major, I think, at the time -- and he came in days before Craig Fields. In fact they came in after Licklider was in place, I'm pretty sure. Anyway, we got some people from Montgomery and some from the Pentagon, and we had a number of meetings. We put together what was the National Software Works. The basic idea was just as I said. Then we had a number of technical issues regarding getting protocols in place that would make it possible

to have an arbitrary set of tools and make it possible to have enough of a common file system across different machines. This meant we had to define what the front end would look like, and all that. So I lined up a lot of technical work. I brought all that back, and Lukasik accompanied me over to meet with Robbins, and we shook hands on this deal. And this was in the spring, March or April of 1973, and Robbins said, "Okay, we'll do this and I have some year-end money, so I'd like to start off by contributing my share now." The problem was the year was going to start July 1st; they hadn't yet shifted into October 1st. "But I can't spend it fast enough, so I'm going to ship it over to you." Nowadays ARPA has contracting officers but not in those days, and there was no money whatever that was spent from inside ARPA; it always went elsewhere. The iron-clad rule was that you didn't take any money into ARPA. All we would do is ship it back out. Therefore, if you were going to do something that involved somebody else's money, you were to make a deal, tell them where the money was going, and have them ship it to the end-point directly. But he said he wanted to ship it to us so that it became ARPA money to get it off his books. Lukasik said, "Okay." And I said, "But sir, but sir" He said "We'll take it," and I said, "Yes, sir."

I came back to the office and walked over to the Program Management Office and said, "The Air Force is shipping three-quarters of a million dollars in here, and we need to get it on a contract by the end of the year." They said, "Absolutely not." And I said, "Lukasik said it's coming." I won't say what they said then, but we started, from scratch, with no proposals in hand, and the need to obligate three-quarters of a million dollars in a couple of months. Now, I've already emphasized that exercises like this are what dominated our life, and it didn't happen so easily. I had bumped into one place, one contracting officer, that asked for more work. This is, again, a singular event, this is not a thing that happens. They were down at Range Measurement Lab at Patrick Air Force Base in Florida, which we also had been involved with for other reasons.

I called up the contracting officer at Patrick and said, "I've got three-quarters of a million dollars, and I've got to create this program." I knew roughly what I wanted to do. I wanted to buy some IMPs and add to the network, and I wanted to get Compass started they had the best set of ideas that were central. I sketched out how much money I needed for each of these things and told him, and he said, "Fine, I'll get a pre-contract award audit started on

Compass right away." This was without having a single piece of paper in hand from anybody — no RFPs, nothing. I called Compass and said, "I need a proposal from you on your ideas, and here's what they'd better come in at." And on June 29th or June 30th, we had a signature on a contract, and that obligated the money, and we were in business.

That was one of my accomplishments for the year. I felt good about it. And it was purely administrative from the point after we had decided what to do, yet it was still non-trivial to make sure all that money came in. It came in two different pieces from different Air Force offices that this person had charge of and had shepherded it through. It turned out that when you take the money off the books in one place and put it on another, that triggers some reactions: it causes a re-coloring that the system knows full well about, and there's limited authority for doing that, and this action exceeded the authority. Therefore, they had to go back up through some layer in the Comptroller's Office to get additional authority. A lot of things are interconnected that you never know about until you trip them off. So that started the program, and I left shortly after that and Bill Carlson took it on, principally.

There were two big problems with that program. One was that the problems we were dealing with were considerably harder than I had gauged. For example, I knew that the existing protocols were not efficient enough and that we needed more efficient protocols for moving files and for engaging in general connections. I thought that we would be able to fold in a re-work of the whole protocol stack as a subtask of this effort. That I misgauged, entirely. We are still trying to get the protocols more efficient in a number of ways as the settings change. But Kahn and Cerf took on the infrastructure of the network and the coming need to interconnect networks in quite a different way, and that became a first class activity instead of being tagged onto something else. Then, of course, came the existence of smaller, cheaper computers that you could buy and have and that came with a full suite of software. See, UNIX had not yet taken over at that point; in fact, I spent a considerable amount of money developing operating systems for the PDP-11 — multiple operating systems; everybody was building operating systems for the PDP-11 prior to UNIX. One came out of the speech understanding work. One person at SCRL was doing a fine job of building a little high-speed, efficient operating system and we had the ILLIAC IV effort at Illinois. Another group was building an effort called ANTS, and one at SCRL was ELF, and I think there were one or two others.

O'NEILL: Okay, and those were all being funded through . . .

CROCKER: In one way or another, they were things that spring out. You see, one of the problems of running a sort of lush environment is that you get a lot of initiative, but you may get a lot of duplication, or things that cover similar, but not identical, things. At one point we had to put a team together to look at the relative merits of the different operating systems that these kids were inventing at every different place.

O'NEILL: Is that when they went with TENEX, or

CROCKER: No. TENEX was a PDP-10 operating system, so it was a different

O'NEILL: Oh, it was a 10

CROCKER: Yes. And that had come out

O'NEILL: That seemed to have become kind of a standard, at least in the community.

CROCKER: Yes, what happened was it became a standard in the community and DEC also picked it up as their operating system for their DEC-20. DEC had its own operating system called TOPS-10. The 10 and 20 are the same architecture, infrastructure, same instructions.

O'NEILL: Okay. And TENEX became TOPS-20?

CROCKER: TENEX became TOPS-20.

O'NEILL: Okay. But that is unrelated to all these other operating systems that you had . . .

CROCKER: Yes. PDP-11s were small machines and PDP-10s were big things, in rough terms, although about the size of one of a normal workstation now.

O'NEILL: Easily. You know, something that had occurred to me when you were talking about the National Software Works was your interaction with the Services. Was that a singular event? I mean, did you have very much contact with them?

CROCKER: I didn't. In that office, it was not common at the time to have a lot of interaction with the military, although on paper there should have been. Since that time there's been quite a bit. I mean, now we have, probably I don't know the correct proportions, but 50/50 military and civilian people as Program Managers and quite a few joint programs or programs set up to work military problems from a computer science point of view, or transfer things, or whatever.

O'NEILL: Testbeds?

CROCKER: Testbeds are a big thing. The Strategic Computing Initiative put in an order of magnitude more money. They changed the character I mean, this business of creating a third budget, a second 6.2 budget, was a big deal in its time, but it was a small thing compared to the \$100 million infusion of SCI. The first thing was just a training exercise for how to get it done big time. And I'm sure the character of the place changed a lot. The community enlarged and sub-divided so that now you have a large number of principal investigators. When you hold a Principal Investigators meeting, 300 people might come, and you might have only a fraction of the total environment. Now they've broken the office into two parts, and there are two separate offices. So any time you scale things up, the structure changes; it's inherent.



END OF INTERVIEW