An Interview with

SEVERO ORNSTEIN

OH 183

Conducted by Judy O'Neill

on

6 March 1990

Woodside, CA

Charles Babbage Institute
Center for the History of Information Processing
University of Minnesota, Minneapolis

Severo Ornstein Interview
6 March 1990

Abstract

Ornstein describes his experience at Lincoln Laboratory which included work on the SAGE, TX2 and LINC computers. He discusses his involvement with the LINC project, including its move to Washington University, and the later work there on DARPA/IPTO sponsored macromodule project. As the principal hardware designer of the Interface Message Processor (IMP) for the ARPANET, Ornstein describes the IMP design work at Bolt Baranek and Newman (BBN), the working environment of the group at BBN, his relationship with Lawrence Roberts, his interactions with Honeywell, and his work on the Pluribus multi-processor IMP. Ornstein also discusses the contributions of Wesley Clark and Norman Abramson, his involvement with the Computer Professionals for Social Responsibility, and his views on artificial intelligence and time-sharing. This interview was recorded as part of a research project on the influence of the Defense Advanced Research Projects Agency (DARPA) on the development of computer science in the United States.

SEVERO ORNSTEIN INTERVIEW

DATE: 6 March 1990 INTERVIEWER: Judy O'Neill

LOCATION: Woodside, CA

O'NEILL: We were just starting to talk, and you were mentioning the contributions of Norm Abramson and Wes Clark.

ORNSTEIN: Wes Clark is, I think, a more important figure in the history of computing than is generally recognized; there will be a limited number of people who will point to him. He was an extremely important figure at Lincoln Lab, and I met him there and worked with him for quite a few years. He was the designer of TX2, TX0, the LINC, and a number of other machines, and had a lot to do with the design of machines in the very early days of Lincoln Lab. He's a very retiring fellow and, I think, a real genius and had a lot of important early ideas. He is also very articulate. A little bit wacky, but a lovely person and a very good friend. In later years he hasn't stayed on the forefront of innovation. People have adopted things that he thought up originally, and I don't think he's gotten as much credit as he deserves. I recently read a history of the work at Xerox PARC, you may have seen the book, "Fumbling the Future." Xerox gets credit for at least one important idea which Wes had had quite a few years earlier. I wrote a letter to the authors of the book. By the way, I should finish that by saying that Wes not only did a lot of early machine design, but was also involved with the beginnings of the ARPA Network. In fact I think it was he who coined the name "IMP."

O'NEILL: That's what the reports I've read said.

ORNSTEIN: Right. I wasn't working with him at the time, although we had been good friends for years. I did not know of his involvement but later heard that he had done some consulting for the guys. Larry Roberts had been working on his degree at MIT when Wes was essentially technically in charge of the TX2 group at Lincoln. I was working in the group at that time. So people like Ivan Sutherland, Bob Taylor, and Larry Roberts will certainly know of Wes' work and have a real high regard for him.

I mentioned Norm Abramson because he also is someone who is not sufficiently credited. Norm devised a thing called the Aloha System which was really, in my estimation, a forerunner of the Ethernet. The Ethernet designers, of course, took that idea and ran with it hard and made many, many improvements. But the Aloha System was the first to embody the notion that one could simply radiate what one wanted to say to a set of receiving stations and that with a retransmission discipline the message would eventually get through despite collisions resulting from people wanting to "speak" at the same time. That general idea of sort of breaking in with your message into a medium was originally suggested, so far as I know, by Norm Abramson. That was well before the Ethernet came along. And the Ethernet did, of course, incorporate crucial improvements that made the idea work much better. Norm was interested in communication between the Hawaiian Islands, where the campuses of the university are on different islands. He proposed to do it by radio rather than on a wire - there were many, many differences. But he suggested the kernel idea and, I think, not many people know that.

O'NEILL: I wanted to find out a little bit more about your career and your background and training before ever getting involved in the ARPANET.

ORNSTEIN: Well, I'm an old man, and there were no academic courses in computing when I was a student. In fact the word "computing" was hardly in the language. I got into the computing field in 1954 ('53 or '54, I don't remember exactly) at the time when Lincoln was just being formed. I had been a geophysicist before that. My degree was in geophysics. A friend of mine who was working at the same research company, Gulf Research, where I was working came there having spent some time working on the Whirlwind computer at MIT. He also was a geophysicist but with an early background in computers. That was just at the time when core memories were first being tried out. So without going into detail, I eventually went to work at Lincoln as did he. That was the beginning of my career in the computer field.

O'NEILL: Who was this other person?

ORNSTEIN: His name is Howard Briscoe. Years later we both independently ended up Bolt, Beranek and Newman in

Cambridge, and as far as I know, he may still be there. I haven't kept track of him in recent years.

O'NEILL: Where was your degree from?

ORNSTEIN: Harvard. But it wasn't in computing, you understand. Nor was it an advanced degree. I don't have anything beyond a bachelor's degree, although I did some graduate work in math and physics at both Berkeley and Harvard. But my degree was in geology. After quitting graduate school, I worked as a geophysicist for a year or so before getting into the computer field. Later I ended up teaching courses in computer design at Northwestern University, at Harvard, and at Washington University in St. Louis. But I learned by doing and by working with people who knew.

The places I worked are fairly few. I worked at Lincoln Laboratory for about seven years. Then a group of us left and went down to MIT proper to form a group, under Wes Clark's leadership, that was interested in computing in medicine. That group designed the LINC computer (although it was really Wes' baby). The LINC was one of the very first mini-computers. It combined a number of technical features in a single machine. It had a display and a small tape unit. The tape units that later became DEC tape were really invented by Wes for the LINC. It was the first machine that you could take apart and put in the back of your car, carry somewhere else, put back together again, and it would work. That idea had never previously seemed conceivable. This was right around 1960, I don't know exactly which year.

O'NEILL: This was explicitly for hospital use?

ORNSTEIN: Well, LINC stood for Laboratory Instrument Computer. So it was really a laboratory instrument that was used in a wide variety of medical applications, and... well, that's another history in itself. After a couple of years at MIT, a group of us, essentially the same group, left Boston and went out to Washington University in St. Louis, where we spent several years. I left after three years because I really didn't like St. Louis. I went back to Boston and went to work for Bolt, Beranek and Newman. That was where I got into the network business.

O'NEILL: What were you doing at Washington University in St. Louis? What was the group doing?

ORNSTEIN: First, we were pursuing the LINC project. We had gotten involved with the National Institutes of Health. They funded a rather sizeable program of building multiple copies of these machines (LINCs) that were then placed with various medical researchers around the country. I think there were two batches of ten machines each that we built and installed. Doctors came to St. Louis and each brought a technician along. We gave them a course on how to program the machine, how to use and repair it, and then they took the machines off to their labs. Then we provided consulting at the various sites. I worked, for example, with a neurosurgeon and a cardiologist, as well as with a variety of other kinds of medical researchers, as did the others in our group. We went around the country giving advice, consulting, writing programs for these people, getting them started in the use of these machines. It was an experiment NIH conducted to see whether computers could be useful to medical research. It was a new idea at that point. And of course it turned out they were extremely useful.

When I went to BBN, I did a number of smallish things at first. I worked for Frank Heart for a while, and then I worked for the PDP-10 group, and I then I went back to work for Frank. I think I was casting about and I wasn't sure what I wanted to do. I had been teaching at Harvard at the same time, in parallel, with the work at BBN, and I was thinking perhaps of going to Harvard full time because Ivan Sutherland was there at the time and there was some appeal to that. It wasn't really until the RFP came from ARPA for the building of the network that I really got my teeth into something serious. I remember quite clearly when the RFP came. Frank handed it to me and said, "Why don't you take this home and have a look at it and see what you think?" So I took it home and looked it over and came back a day or two later and put it down on his desk and said "Well, sure, I suppose we could build that if you wanted to, but I can't see what one would want such a thing for." (Laugh) Prophetic words. In fact, I think that as the network has turned out to be used, it has had some surprises, even for the people who envisioned it in the first place. Of course it's easy to look back and say, "Oh I knew that all along." But I think there's a certain measure of surprise for practically everybody involved. Anyway, I certainly confess to being absolutely dead wrong about the thing at the time. (Laugh) But, nonetheless, we thought it over and decided it would be a fun thing to do - which for me was

always a primary consideration. Of course as we talked to Larry about it we understood a little better what he had in mind. I had known Larry, as I say, years before at Lincoln Laboratory. By this time he was head of the ARPA IPT office.

O'NEILL: Can we go back a little bit to this group at Lincoln. Can you tell me who was in the group? Were a number of people moving to these different locations, working on these projects?

ORNSTEIN: Well, group 62 at Lincoln Lab became known as "The TX2 group," and there were a whole lot of people there who later became well known in the computer field. It was thought of by most of the lab as the group that was most advanced as far as computers were concerned and that was doing the most fun kinds of work and had the most freedom. I didn't initially work in that group. In fact I worked with Frank Heart in an entirely different group and then joined the TX2 group when an opportunity to do so came up. Some of the other groups were more application oriented, whereas the TX2 group was a real research group. There were a number of bright people in that group, actually people whose names have since become well known. Larry Roberts and Ivan Sutherland, for example, were doing their graduate degrees at that time using the TX2 computer. Ivan did Sketchpad right there at that time. When I joined the group I started to learn about hardware and do hardware design for the first time. Before that I had been writing code. In later years I became known as a hardware guy and was never a software person in modern terms.

To answer your question more specifically, when I joined the group, Wes was beginning to think of the idea of the LINC, but he didn't have a very clear idea of it. He said "Why don't you go look around Lincoln and find all the places that are using specialized digital equipment that a small computer might be able to replace - if there were such a thing." I did that, and that was the beginning (around 1960) of the formation of the LINC group. Now the LINC was really largely Wes' conception. He is an inventive genius, and I remember many, many months during which he worked more or less on his own, occasionally coming back and talking to various of us in the group. Out of that was bred the first LINC machine. The group of people who worked on it was very excited by what was going on. It was a very small but cohesive group that grew, then, when the LINC became successful. That group of people, somewhat metamorphosed and enlarged, finally left Lincoln Laboratory. Projects at Lincoln were mostly Air Force projects.

The Air Force had one kind of overhead; NIH, which was interested in enlarging the LINC program, had a lower overhead and Lincoln decided that they didn't like that. So we decided to leave in order to pursue the LINC. We then negotiated with MIT proper. In fact, what happened was that a group of us moved down near the Institute under NIH sponsorship and made more LINCs happen. The program was sufficiently successful that our group and NIH jointly put together a proposal for a large inter-university laboratory to work on computers and biomedical research - which was to be headquartered at MIT but involved all the Ivy League colleges. It was an enormous proposal; I think it was 37 million dollars, which was real money in those days. When it became clear that the first several million of that was forthcoming and that this was a real thing, the administration of MIT got really interested. When the academic community understood what was up, they suddenly wanted to be involved. You see we were not in the academic fold, we were on the side in various research departments. There were complicated relationships there. Part of what had spawned the LINC was an interest in neurophysiology. Walter Rosenbeith, a well-known professor, wanted to oversee this new big thing that was going to happen, now that it was clear that it *was* going to happen. Eventually a confrontation took place, and Charles Townes, who was then the provost of MIT, ultimately decided that he had to side with his academic staff. So it was clear that if it was going to happen, it was going to have to be under the jurisdiction of some department. At Lincoln Laboratory we had discovered the disadvantages of being under other people's thumbs and since we had succeeded in creating this on our own, we felt we wanted to stay independent. So we decided to leave MIT. It was a very sad decision, and the center we had dreamed of never took place as a consequence of that. It set the field of computers in medicine back quite a few years, I think, in the process. Then we started travelling around the country, talking to a variety of universities, trying to figure out where to go. We finally ended up in St. Louis at Washington University, which has become, partly as a consequence of that, partly because of other people who were there, a major center for computing and medicine.

O'NEILL: Were any of the projects that you worked on, other than the LINC, real time computing systems?

ORNSTEIN: Oh sure. A great deal of the biomedical stuff was real time stuff, practically all of it. It wasn't post facto data processing at all. It was laboratory work, and a lot of it involved connection to instruments that were making real-time measurements - you know heartbeats, brainwaves, you name it. So yes, practically all of it was real time

stuff. And we'd had a lot of experience with real time stuff at MIT, too. One of the many features that made the LINC unique was that it had built-in, very simple to use analog-to- digital conversion channels so that it was very easy to turn it into an oscilloscope, essentially a capturing oscilloscope.

O'NEILL: Did you work on something called the Lincoln Experimental Terminal System?

ORNSTEIN: No, I didn't. I know vaguely about that, but no, I didn't have anything to do with that. What I did later - and I haven't mentioned - at Washington University after the LINC program became a going concern and didn't need our watching over, the next stage we thought was to build a more flexible set of building blocks for computers - pieces you could build computers out of or other specialized digital equipment to do special kinds of things that you could easily connect to a computer. So we worked on the design of what we called "macromodules." I had a lot to do with that and probably contributed more inventively in designing macromodules than I have in almost any other arena. I may have contributed more *work* in other areas, but the initial design of macromodules was primarily worked out by me and another fellow by the name of Mishell Stucki working together at Washington University. For various personal reasons, Wes wasn't in very good shape at that time. He would normally have been more the intellectual leader, but he was out of it for a while there and Mish and I simply plowed ahead on our own. Wes would check in periodically and look over what we had done and eventually, of course, he picked it up and ran with it. But Mish and I together, over the period of six months or a year, did a lot of the inventing. Later on others picked up the threads, found what was wrong with it, and so forth. But we put together a lot of basic ideas.

O'NEILL: That was an ARPA project as well, wasn't it?

ORNSTEIN: Oh, absolutely. Ivan Sutherland and Bob Taylor were both at the ARPA office at that time. Ivan came by after we had worked out a basic set of modules that you could really form a computer out of. I remember very clearly the day he and Bob Taylor visited. Ivan understood in fifteen minutes what it was we had done and after we described the various pieces to him, he sat down and, using them, put together a little computer. He said, "By golly guys, you've done it." I remember his moment of realization. So, yes, that was an ARPA review day; they came by to

see what we had been doing with their money.

O'NEILL: How long did you stay at BBN?

ORNSTEIN: I went there in 1967 and I left there to come out to PARC in early 1976. Then I was at PARC until 1983 when I retired. Not terribly long it seems. Did a lot in that time.

O'NEILL: Did you work with Bob Taylor at PARC?

ORNSTEIN: Yes. As you may know, at the time I came, there was complicated politics at PARC. Jerry Elkind was there. He was one of the people I had worked for previously at BBN before I settled down into the network project. When I came to PARC, Jerry was running the group I joined and Bob was his associate leader. Jerry was not universally liked as a technical supervisor. I think he didn't have the right touch. Although I like him personally, as a technical supervisor he wasn't terribly successful. A number of us suggested that the group would be happier under Taylor and that's how Taylor actually.... Well, it's a long complicated story. You can read about it in "Fumbling the Future."

O'NEILL: Were you familiar with or aware of the work that Roberts and Tom Marill were doing connecting SDC Q32 and the TX2?

ORNSTEIN: No, but I'm not surprised. That was probably at about the same time that I joined the TX2 group, I would guess. I don't remember exactly what year that was, but it must have been around 1959, 1958.

O'NEILL: I think that might have been a little bit later.

ORNSTEIN: Later?

O'NEILL: It might have been when you were already it St. Louis. I believe it might have been 1963 or 1964.

ORNSTEIN: No, I didn't know about it. Years before I'd worked on the Q7 some, [actually the forerunner XD-1] writing code for it. That's what I started out doing at Lincoln. As a matter of fact, interestingly enough, I haven't thought about it later, it was the communications part of the Sage system that I worked on under Charles Zraket - Zraket is now head of MITRE. He was the guy that I worked for at Lincoln. When MITRE was formed I was in a group that would naturally have become a part of MITRE. But in my usual rebellious way, I didn't want to do that. And I and only one other guy, I believe, from the set of groups that formed MITRE decided to stay behind at Lincoln because we were more interested in research and it looked as though MITRE was going to be building the air defense system and we weren't interested in that. So we stayed at Lincoln Lab and that's when I went to work for Frank Heart prior to working in the TX2 group.

O'NEILL: So before you started in Frank Heart's group you were working on Sage?

ORNSTEIN: Yes, I had been working on the Sage system. And the particular part of it that I was working on was what was called "cross-telling," which is the formal communication about aircraft, etc. between sub-sectors. I wrote the operational specifications for the kind of communication that needed to take place between two sub-sectors to pass both hostile and friendly aircraft over for tracking and so forth. And there were exciting one-kilobit lines between sub-sectors that carried all this information. That seemed like a flood of bits; what would one possibly send that could use a whole kilobit? So I figured out what to send.

O'NEILL: Was that all from the software end, were you doing programming at that point?

ORNSTEIN: Yes, I was actually writing code. Designing pieces of the system and also writing the code eventually.

O'NEILL: You mentioned Whirlwind... I'm getting confused, did you actually use Whirlwind?

ORNSTEIN: Well, no. The machines that I learned to program on were not where I was at at the time I was writing the code. When I was working as a geophysicist with this fellow Briscoe in Pittsburgh, on the side he was teaching me to program. He had a manual for Whirlwind, because he had worked on Whirlwind. The first programs I wrote were actually for the EDSAC machine because he gave me an EDSAC book to read, and I wrote some simple programs but I was never near the EDSAC machine obviously. I wrote more extensive programs for Whirlwind, but they were never run. This was just learning.

O'NEILL: Good, that clears that up. We started to talk a little about getting the RFQ from ARPA. Before you started working on that, were you working on medical projects at BBN?

ORNSTEIN: Not at BBN, no. I helped with the design of some analog-to-digital stuff for the 940. I worked with Chuck Seitz on that. I don't remember all the things I did. I went to BBN originally to work on various education projects with Feurzeig and Pappert at MIT.

O'NEILL: Oh. So you went back to BBN to do some education work.

ORNSTEIN: I worked on a variety of small projects, none of which were really panning out. I was also busy teaching at Harvard at the time - a course in logic design and introductory programming. Perhaps that is part of what kept me from getting really involved at BBN.

O'NEILL: You mentioned not having any knowledge of what was going on in the networking area prior to the ARPA.

ORNSTEIN: That's right. It was really brand new to me when Frank handed me the RFP. It was the first real exposure. I vaguely knew that people had been doing some interconnecting of computers, but nothing specific about networking.

O'NEILL: What was your working relation or situation with Bob Kahn? Was he involved at the same point with the RFQ?

ORNSTEIN: He was indeed. It was at that time that I met him. He was in Jerry Elkind's group at that time, a different group than I was in. I was working for Frank by then, and Bob was working for Jerry Elkind. I had previously worked for Elkind, but I didn't know Kahn at that time. I think Kahn had relatively recently come there from Bell Labs. Anyway, he was clearly, terrifically interested in the network project. The proposal that we wrote eventually at BBN was an enormous undertaking for the company, just the writing of the proposal. More dollars were spent preparing that proposal, more man hours charged to it, than I think had ever been done for any project, any proposal before. This was a really big, big thing, and a lot of effort got put into it. I don't know if it was decided that that was how it was going to be. It simply ended up that way - in part because there were quite a few people who were really, seriously interested. We all saw that it was a major undertaking and we knew that we were competing with large companies who would pour enormous resources into their proposals. If we were going to stand a chance of winning, ours was going to have to be the glistening, clear technical best bet because obviously a lot of these other companies had much larger reputations to offer. So we put an extraordinary amount of effort in the proposal. I was responsible specifically for the design of the interfaces to the IMP, all the special hardware. I designed those interfaces in considerable detail while we were writing the proposal, I would say I did 90% of the design then. I knew exactly how they would work - it was just a matter of laying in the gates after that. We chose to use a Honeywell machine and the man from Honeywell who was assigned to build those interfaces from my drawings didn't understand the drawings well and was not really careful. We ended up having to redo much of his work. I knew exactly what was wanted because, as I say, I had designed everything and had drawn fairly detailed pictures that went into the proposal describing how the interfaces were to work. Since we had a variety of host computers there had to be fairly general interfaces.

O'NEILL: Was that with knowledge of the Honeywell 516?

ORNSTEIN: Oh yes, absolutely. By the time I was designing the interfaces, the 516 had long been chosen and I

knew exactly how all the lines were going to connect to it. It was really a fairly detailed design, as I say, except for the actual laying down of the individual gates. But, yes, absolutely.

O'NEILL: How long did you actually work on the proposal? I think the proposal came out in June or July of '68.

ORNSTEIN: I really don't remember. I'd have to try and track it back, I don't know. But it feels like several months. It may have been more than a few. I just don't remember. It could have been up to six months, I would think.

O'NEILL: Do you remember who you felt you were competing with? Which other companies?

ORNSTEIN: No, I really don't any longer. I guess Western Union was perhaps one of them. You probably know who they were. We eventually got word, but not in the first place. We just knew that it was a big enough project that there would be a fair number of large companies competing. As time went on we knew more and more who they were, but I don't remember the exact unfolding of that picture.

O'NEILL: I was more interested in who you felt you were competing with - but just large companies is good enough.

ORNSTEIN: Well, large companies, Raytheon, Western Union, we thought all of the fairly obvious ones. We didn't think IBM, I guess we didn't think IBM would compete for that, but it was not clear.

O'NEILL: I haven't actually seen a lot of detail about which companies were competing.

TAPE 1/SIDE 2

ORNSTEIN: I think we knew more than I'm able to remember now. All I remember was being pleased and somewhat surprised, I guess everyone was a little surprised, that we actually finally were awarded the contract. Although by the time that we got down to the end, of course, we knew we were on an increasingly short list. By then we had done

so much work that we felt that if a good judgement was made, and we knew enough about Larry Roberts to know that he would exercise good judgement, we felt that we were really very hard to compete with from a technical point of view. There was a question, of course, as to whether BBN was able to mount such an effort. There was always that question, because it was a small company and the larger companies obviously could mount larger efforts. But by the end of the process, I think we felt we had a very good chance despite our size simply because we felt we had the best proposal out there. We had the whole system designed. We knew that other people generally didn't go to that extent.

You asked about Bob Kahn. Yes, Bob and I worked very closely together. He wanted to know everything. He's an omnivorous guy, and he knew very little about hardware design at the outset. He would take me aside and say, "Okay, explain exactly how it's going to work." He wanted to know all about it in great detail. He's a very thorough person, so he learned a lot. I did, as I say, the hardware design, but I was also heavily involved in the whole system design. We knew that the program, if we ever came to write it, was a really tight, real-time program. Both Frank and I knew a fellow named Bill Crowther, who had worked with us previously at Lincoln. He was a brilliant programmer and thoroughly understood machine language code, the kind of code that you have to whittle down, very well. So we hired him from Lincoln at that time. I don't remember it exactly. He had other commitments and he didn't come right away, but he helped us work on the proposal. So I spent time with him as well as Frank and Bob Kahn. I spent a lot of time, I remember nights till 3 and 4 in the morning, working with Bob Kahn in the back room of my house in Newton on the proposal - designing the system and figuring out how it was all going to work.

O'NEILL: So the "we" you used in response to the proposal, were you and Bob Kahn, and Frank Heart and then some consulting by William Crowther?

ORNSTEIN: Yes, but there were a lot of other people involved as well. Jerry Elkind, also, had input. There were other people, I think Danny Bobrow, I don't remember all of them. You know, everybody who was around was involved, and that's one of the things that made it so expensive. There were a lot of high powered people who were putting time and thought into the proposal, helping to hew it. I would be hard pressed to say, but I guess Frank

deserves credit for having done the most work on it. But there were an awful lot of other people who put in unbelievable hours - many of which they weren't paid for. As I say, Kahn and I were working till the wee hours of the morning very often, I remember clearly. I'm sure he'll remember that too.

O'NEILL: Was the incentive just that it was fun? That it was an interesting project?

ORNSTEIN: Yes. It was clear that it was a sizeable system and a fairly ambitious one, yet it was something that was clearly do-able. But, just about. It was clear that there were good ways and bad ways. Yes, it was a technically challenging job. For me, that was the primary thing. I don't think that I had time to think about the longer range possibilities of what it might mean.

O'NEILL: You mentioned that as you got closer to the proposal and realized that it was a very good proposal and the chances of getting it were fairly high, did you plan to actually work and implement it?

ORNSTEIN: Yes, absolutely. It was clear.

O'NEILL: You saw yourself as fitting into that implementation.

ORNSTEIN: Sure. And we saw Willy Crowther not only as working on the proposal, but also as the guy who would be primarily responsible for writing the code initially, as indeed he was.

O'NEILL: So you brought him over with that intent?

ORNSTEIN: Yes, absolutely.

O'NEILL: Once you got the contract and started working on it, how did that group function? You hired a few more hardware people, right?

ORNSTEIN: We did. Just before that I had finished the second session of the course that I was teaching at Harvard. I had had a bunch of bright students, and we hired a pile of them. We hired all the best students out of that course eventually. Some of them have become very successful. Ben Barker, for instance, was a guy who had been a student of mine and then later a T.A. We hired half a dozen or so of the students from that course over the next couple of years. Not all of them worked on the network project. I don't know whether you know John Robinson and a fellow by the name of McMillan. There were a fair number of people who got hired and who worked mostly on the network stuff.

O'NEILL: Did you maintain your position at Harvard during this time?

ORNSTEIN: No. After that there wasn't time. There was so much to do with respect to the network. And later on embellishments to the IMP and so forth. Still later I was in charge of the design of the big multi-processor we called the Plurubis (IMP).

O'NEILL: Oh, the Pluribus. What did you call it initially?

ORNSTEIN: HSMIMP. High Speed Modular IMP. That was just an internal name, that never was a public name. (Laugh)

O'NEILL: It is not quite as catchy as Pluribus. (Laugh)

ORNSTEIN: I got to choose the name Pluribus; I've always liked it.

O'NEILL: Can you describe what it was like once you got the contract? Was it still those kind of long hours?

ORNSTEIN: Yes. You could hardly tell the difference as far as I was concerned, except that it got to be more real.

Everybody started working like mad on the code and on getting Honeywell actually to start building the machines. I had a fairly heavy hand, literally, at some times in the building of the machine. As I say, Honeywell did not do well at it; they were behind schedule and mostly they sent us cabbages instead of computers. They hadn't understood the interfaces well and when the machine finally came, they didn't work. We had fierce times getting Honeywell to straighten the thing out. At one point, with Frank standing by with his jaw dropped down, I actually turned a truck around at the loading dock and told them to take the thing back. I caused a lot of fuss at Honeywell, but they finally shaped up. I had to be really quite nasty at times, and beat on the table, and shout and scream to get them to fix the troubles they were having. Eventually they shaped it up, but it took a long time, like the fourth or fifth machine. We were in a bind because we had a promised delivery schedule to DARPA and all of us working on the thing felt very keenly a sense of pride about wanting things to be on time. And they were. But a lot of blood was shed in the process to make it happen.

O'NEILL: What about the process of deciding to go with Honeywell in the first place, was there a lot of discussion about that?

ORNSTEIN: Oh, indeed there was. I don't even remember what the other contenders were. Frank will probably remember better than I. Certainly we considered several other ones. But I think that it was fairly clear that the Honeywell machine was the right one. I don't remember all the reasoning, but all the rationale that led us to the Honeywell machine was in our proposal.

O'NEILL: After the first three or four machines, was the working relationship with Honeywell on a fairly even keel?

ORNSTEIN: It eventually got to be. I always felt they overcharged fiercely, and we fought that as hard as we could, you know. But, yes, the relationship eventually straightened out. They knew how to copy things, once they got it right; it was just a matter of getting them to get it right. I would send back drawings with corrections and they would incorporate three quarters of the corrections and the other quarter were overlooked. We finally got them to straighten things out. After we got the first machine or first couple of machines installed, I finally took a break and

went off to Europe for a month and left Ben Barker in charge of the hardware. He was paralyzed with fear because he suddenly had a lot of responsibility thrust on him and he was not long out of school. But he was very good. By the time I came back, he was smiling and everything was coming along just fine. There was a lot of fun and a lot of excitement in shipping the first machines and putting the network together and seeing it all suddenly start working at several sites. Then, of course, followed all the business of starting to build up a sense of use for the network because, although the IMPs were in place, they weren't doing very much for quite a while until people finally started to use them, to develop protocols, and so forth.

O'NEILL: As I understood it, Honeywell did actually assign people to work with you. That didn't work out, is that what you are saying?

ORNSTEIN: Well, I guess there was miscommunication. The guy that they assigned, I think, just wasn't terribly bright. They were not used to doing some of the kinds of things we were doing. We diddled their machine in ways that nobody had diddled their machine before. For example, we discovered some design flaws in the machine. One of the reasons we chose the Honeywell 516 was that we thought it was a mature machine (we had a serial #500) that was not going to gives us grief. Well, we were wrong. We pushed it very hard; they had never had so much real-time traffic, interrupt and DMA traffic, coming into the machine before, and we uncovered a bug that they could hardly believe, a synchronizer problem. I don't know if you're familiar with synchronizer problems. Synchronizer problems are very, very subtle. We had to dig and dig and dig at them, and finally from their back room they produced a *really* smart enough guy - they do have a few - but, it was very hard to get this guy. We finally got him, I don't remember what his name was, but I do remember when he came and we sat down and I finally had someone I could talk to who would understand what I was talking about and believe me. It was a subtle problem. The program would run fine in the machine for days on end, and at the end of three or four days suddenly the machine would just die, inexplicably. It was a very, very low frequency failure, so infrequent that you could never look at it on a scope; you could only see the effects afterwards. In fact, we had to build special hardware that beat on the trouble spot many, many times faster than normal usage would. Then finally, with all the lights out in the room, we could faintly see on the tracing scope an occasional failure. That was when the Honeywell people finally became convinced that there was a real

problem. Fortunately it was a fixable problem as it turned out. We showed them the trivial fix they could make to the machine. It was a fix they had to make to all their 516 machines, which was a major undertaking. So, their people weren't absolutely top-notch people; they were okay. They were industrial-strength people, not research-strength people.

O'NEILL: Was there a lot of distinction between the hardware and the software side of this? Did you consider yourself as one group?

ORNSTEIN: We certainly considered ourselves one group. We were very close friends, the people who were in charge of doing the hardware and the software. That was one of the nice things about BBN, that there was no compartmentalization at all. It was certainly all one group, and pushing stuff back and forth across the hardware/software boundary was something that we did easily and cooperatively. There was never any question; we were like thinkers. The cohesiveness of that group at the time was enormous. I don't remember any tensions, technical tensions, whatsoever within our group. Bob Kahn, who wasn't politically in our group, got along relatively badly with some of our people at first because he really wasn't a computer person. He was an information theory guy from Bell Labs. Some people think he's never become a real computer person, which is true to some extent. He didn't have it in his early education. But he certainly tried very hard, I give him credit for that, and he learned a great deal. But he didn't understand, or if he did he didn't give a clear indication that he understood, a lot of the problems that we were dealing with and some of the things that he was suggesting were, we felt, off the wall, just wrong. People were, of course, busy and impatient and didn't want to take time to explain. So there was certainly that kind of thing. But not as far as the people who were actually doing the work. Bob wasn't doing the work. He was watching, thinking through a lot of the problems, and commenting. He did work on one very specific part of the hardware, the error detection logic, but not directly on the larger scale system problems, the flow control in the network, and things like that. Bob wanted to do a lot of simulation and to watch simulated network traffic on the screen. He wanted to see traffic moving and things like that. And we said, "Bob you'll never come to understand the problems looking at it that way." He had, what we thought, was largely a naive view of how to come to understand the problems that one would be confronted with. So there was some friction there. We generally did what we thought was best. The group

that was actually doing the work was a very cohesive group. There weren't, that I remember, any serious disagreements. I've been in groups where there were real tugs of war, but there wasn't in this case.

O'NEILL: So, this is Crowther, and Walden, and Barker, and you, and Heart. Is that pretty much it?

ORNSTEIN: Yes. Walden was a much more minor figure at the outset. Later he became a very important figure. In fact, he was at Norsk Data in Norway. I don't think he was part of the beginnings of the thing.

O'NEILL: I believe he went to Norway after some of this had started and then came back. But he had been involved earlier.

ORNSTEIN: I see. Is that what happened? He was a junior guy at that point and obviously has become a very senior guy since then. It was clear that he was very talented even early on and a hard worker as well. I mean, he's a really good guy. But I had forgotten his involvement at the early stages. He must have participated, but I don't remember him at that point. We were disappointed when he went off to Norway, I remember. But then when he came back, he rapidly became a senior member of the group.

O'NEILL: What kind of interaction were you having with ARPA at this time? Were you seeing Larry Roberts on a regular basis?

ORNSTEIN: Sure. I'm a bit of a maverick, I think, and I never take things as seriously as others. Frank took things very, very seriously and was very business-like about the whole thing. He had known Larry, also, of course, at Lincoln. But I was never as overawed, perhaps because I didn't understand, but I always thought of Larry as a friend and someone that I spoke my mind to quite straightforwardly - always. I tend to do that with most people. I remember during that era I was involved in anti-war activities, and I would kid Frank that in these meetings at the Pentagon I was going take my little "Resistance" pin and pin it on the colonel's jacket while he wasn't looking - things like that. I think Frank actually worried that maybe I would. In fact when we hired Will Crowther, I remember talking

to Will on the phone as we were arranging to go visit ARPA. Frank put me on the phone and said, "Tell Will not to wear his sneakers." Well, Willy wears sneakers - everywhere. The only day in his life, I think, that he wore shoes was the day he got married. So I got on the phone and I said, "Will, Frank says that you shouldn't wear your sneakers to this meeting." There was a silence at the other end, and then Will said, "Tell Frank they've seen my sneakers in JSAC meetings - it'll be okay." And of course, Will wore his sneakers and everything was fine. But Frank was more respectful to Larry, I think, than I was. I would have told Larry bullshit if it had been bullshit. I'm that way, pretty much, with everybody. But relations were good. Larry was smart, having decided that we were the best group to do the job; he was as helpful as he could be. I felt we were working *with* him on the thing. One of the reasons I have a lot of respect for the ARPA office is that up to the point at least when Bob Kahn was running it, I'd had an enormous respect for the people in charge of it because I thought they were *very* bright people. They were people you could talk to about things straightforwardly, and you didn't have to footsy around. This was not your average bureaucracy. These were people like ourselves - researchers. Larry would make suggestions and argue with us and very often persuade us. But we would sometimes persuade him too. I felt we had a working relationship with him rather than a subservient relationship. He was very, very helpful. Larry is extremely bright.

O'NEILL: You mentioned having meetings with colonels. Did you have a lot of exposure to the military during this time?

ORNSTEIN: No, no. There was some, I think. There were one or two meetings. You know, a conversation like that with Willy about the sneakers I'll remember until I die, even though it was only once. No, I don't think so; there was some, but not much, no. We were pretty much insulated, I think, from that.

O'NEILL: How about your interaction personally with the host commu nity, the host sites? Did you deal with them at all?

ORNSTEIN: Sure, some. Shipping the machines to the host sites involved some very funny experiences because we insisted that the machines be accompanied by a person. We didn't just put them on a plane. We actually sent a

person on the plane with them. That was something that the airlines hadn't experienced. Nobody was that untrusting. But we didn't trust *anybody* with *anything*. Eventually we gave that up. But for the first few machines, we rode with them on the plane. It was a damn good thing, too. I've seen all sorts of things happen to machines.

O'NEILL: You actually accompanied some of these machines?

ORNSTEIN: I didn't, but we had people who did. For the first few machines, when they arrived I flew out to the site, hooked them up to the phone lines, and got them going. It was senior technical people doing it. I did some of it; other people did some. So we got to know the people at the sites quite well. I've forgotten a lot of them, but that's how I met some people - Bill English, for example, at SRI. I met him there the first time because as I recall he was in charge of getting the IMP installed there. And people at Utah, I forget what the first sites were... UCS, UCLA was it? I've forgotten where the first machines were.

O'NEILL: UCLA, UC Santa Barbara, Utah... I can't think of the fourth one.

ORNSTEIN: BBN was probably one of them.

O'NEILL: I think BBN was the fifth node, as I recall. There were four right in the West Coast area. SRI, maybe.

ORNSTEIN: That's right. At UCLA I met Crocker for the first time.

O'NEILL: Was there discussion at that point about different approaches to various problems that came up? I guess what I want to know is did you have interaction with these people as new problems came up, or did you go off and solve them and then present them to these people?

ORNSTEIN: A mix of the two. Things had been so completely designed that by the time that we got the proposal and were awarded the contract, it was mostly just a matter of doing what we said we were going to do. There was

relatively little invention left at that point to get the IMPs connected together. A lot of invention came later on. For example, the involvement of satellites in the network; I wasn't much involved in that, but a lot of that design was done cooperatively with Larry Roberts. He was very interested in that. Norm Abramson ran a systems seminar in Hawaii every year during that period. We met there, gave papers, and discussed satellite involvement. That's where I met Norm Abramson. In fact one of the satellite systems was called Reservation Aloha, a token to Abramson's system.

[INTERRUPTION]

O'NEILL: Okay. As some of the problems or changes came up later, things like the terminal IMP, having multiple computers at a single host site, and having those computers further and further away from each other - how were these problems resolved? Was that a matter of community decisions?

ORNSTEIN: Yes, I think so. There were probably some pretty heavy discussions, although I don't remember any specifically. There was certainly debate about how to approach some of these problems. But I don't remember any real tug of war. Larry trusted us quite a lot, but he always exercised his own judgment about things. He listened, and if we had a convincing argument, he said that's fine. He was smart enough to spot oversights, and to the extent that he didn't, we made joint mistakes, you know. I'm sure that we didn't do everything perfectly, but I felt, as I say, that we were working together. That was true as we developed the notion of the terminal IMP. BBN did a lot of the work, obviously, but in some areas Larry had particular interest. He would spend time thinking about problems and was always very critical, asked hard questions, and tried to poke holes as much as possible. But that was how we all worked; we did that with one another. I don't remember any conflicts about how to go about things. We worked them out, and when the right answer came up, everybody could see that it *was* the right answer. If somebody could see something wrong with it, well, then we'd work some more on it. I don't think I'm suppressing things; I'm sure we had differences of opinion and argued for days on some matters, but I don't remember any serious disagreements.

O'NEILL: I'm trying to understand how the ARPA office influenced various developments. Whether they arbitrated

disputes, whether they just gave suggestions, were those suggestions very specific, whether they were just reviewing your work?

ORNSTEIN: Well, if Larry felt something was a key issue that a lot was going to hinge on, then he got right down into the nuts and bolts, into the details. That was true, for example, with the satellite systems, because he could see that that was really key, that a lot was going to hinge on some of the decisions being made there. Other matters that were less central, he would let others take care of. He was a very good judge of where to apply himself, and where to trust other people. Key points he looked at in great detail and argued and fought for things. Some things he saw that we didn't see. The way we worked with one another was also the way we worked with Larry. Now, I don't remember exactly when Bob Kahn took over the ARPA office. Do you know what year that was? I think it was at about the time that I left BBN.

O'NEILL: Well, it had to be after '72.

ORNSTEIN: All that work on the Pluribus, who was in charge of the ARPA office at that time? Well, I think Larry was still in charge through all of that. But I don't remember, he may not have been.

O'NEILL: During the time that Roberts was in charge, did you work with anyone else at ARPA? Did you work with Barry Wessler at all? I know the office wasn't very large at that point.

ORNSTEIN: Yes, a little bit, but Larry was really the figure that we trusted. I remember Wessler, but only vaguely. I didn't have a lot to do with him. As the program enlarged and as some of the problems in the software design, routing algorithms, and so forth got refined, I did not keep close track and, frankly, it's been so many years now that I've forgotten exactly at what level I was even involved in those things. I think, for me anyway, Larry was the primary figure that I worked with, to the extent that I worked with ARPA people.

O'NEILL: Did Roberts encourage you to disseminate information outside the immediate community, in terms of

giving conference papers, or writing articles, things like that?

ORNSTEIN: I think we all took it for granted that we would write papers when there was sufficient material. That was just part of what we did.

TAPE 2/SIDE 1

ORNSTEIN: There was some discussion about writing papers, but I don't remember there being any policy questions about it, it was just a question whether this was a suitable time, whether there was enough material, etc. There may have been some security question, but I doubt it, because the project was not security driven. Some of the later sub-projects were classified, but the main stream went into the open literature.

O'NEILL: Did you hear about or hear from people outside the immediate community of the people involved in the ARPANET, who disagreed either with the direction or any specifics of the projects, for instance at a conference, or whatever?

ORNSTEIN: No, not then. Later on, I've heard complaints about the nature of the IMP/Host interface, for example, the so-called "1822" host interface that I designed that people felt should've been done differently, many years later. But "1822" became a standard for a long time. As far as more general things are concerned, no. I don't remember any complaints. I think people were pretty impressed with how reliably the net worked. Of course, I'm sure there were people who said to us, why didn't you do this or that; it would've been a lot better, but nothing massive that I'm aware of. Again, Frank, being in charge, may have heard more complaints from outside than I would have. But, you know, I was pretty close to being in charge, too. I was basically the second in command for quite a while. In fact, Frank wanted me to become his second lieutenant, as the group was enlarging. It was clear the network project was growing; he wanted me to take over more responsibility, but I refused because I really wanted to stay with the bits. I'd always done that, and I did not want to become more of a manager than I was. I became enough of one anyway. No, I don't think there were serious complaints. I'm sure as the users came along, and as the next layer of protocol

got developed, there was discontent with what we had done underneath. And, of course, there *were* some limitations that were built in that we knew about and that eventually required major rewrites of the code expansion to more than 63 hosts and so forth, things of that kind. But those we had known about, and I think one of the key aspects of the successes that I've been associated with, has been the good sense to limit ambitions. I think man computer systems have failed because people saw too many possibilities and were unwilling to foreclose certain possibilities in order to get a bounded goal that you could get your arms around and get right, and then later on solve some of the ancillary problems. We did that very deliberately in the ARPANET design. We fixed certain things and of course with development software in the state that it was in at that time, such bounds were extremely important. Nowadays one has greater luxury because software tools are much more advanced than they were then. But using the kinds of software tools we were working with, simple assemblers and so forth, it was really important to keep the complexity down initially. We knew as things enlarged, if they did, that some rewriting would be necessary. But I think that one of the many reasons we were successful is that we had the good judgment to keep the thing simple. I think keeping things simple and under control is vital not only for those who are inventing, but also for those who are going to have to deal with a system later on. We did that. That was a problem with the Pluribus, where we were doing a more experimental thing that involved rewriting the IMP program for a multiprocessor. Will Crowther produced Pluribus code that nobody but he understood in the early days. That was exc usable, because that was not so much production line stuff; it was more of an experiment. In general, I think that an important characteristic of the group was that it understood where it was important to have flexibility and where it was important to keep it simple.

My relationship with the ARPA office throughout the years that I worked with it (as you see, both before and after the ARPANET) was always very good. I'm one of the founders, as you perhaps know, of Computer Professionals for Social Responsibility. I don't know if you know about that. We have certainly criticized some of the things that later on came out of the ARPA office. I and a couple of other people wrote a highly critical paper about The Strategic Computing Initiative. But I certainly made clear in that criticism, and have always made very clear, that I had nothing but the highest regard for that office for all of the years that I worked with them. I think that was one of the best kinds of government bureaucracy because there *wasn't* much bureaucracy there, at least in my time, and the kind of people who headed that office were top-notch researchers in their own right. I think that has been less true in recent

years. I don't know, but I understand that it's been much less true. I think that its success arose from the way it was run, as a sort of benign dictatorship. There was a lot of power in the hands of the person in charge of that office to make decisions, select projects for funds, and so forth. I think the state of computing in this country is where it is in no small part because those people were smart and they made the right decisions. A benign dictator is good; a bad dictator is bad. During that period I felt that that office funded the right stuff for the most part and that the dollars were very well spent. As I say, I think it's gone downhill in recent years, partly, I suppose, because of the changes in the way it's been run, new regulations, etc. But also because I think the quality of the leadership has degenerated somewhat. They were never well paid. I remember a meeting I went to with Larry Roberts. He was going somewhere afterwards and was driving me to the airport in his rented car. I don't remember where or when it was, Hawaii or New York, whatever. He had a Dollar rent-a-car; it was a beat-up little car. I said, "Larry, why did you rent this kind of turkey?" And he muttered something I didn't understand at the time about government rules and government expenses and so forth. I had always thought of him as passing out these millions of dollars, but it hadn't occurred to me that he was in fact living personally on quite a limited budget. So people like Larry sacrificed themselves for a while in order to get their hands on a big throttle that could work major influence on directions. But that they were making a personal sacrifice had not occurred to me until later when I realized that these guys were not able to do much technical work of their own except from 2 to 3 a.m. They had to watch over numerous programs, which I feel is a really crappy job for a real scientist. So I think of those people as having really given a piece of their lives to the field in a way that a lot of us didn't. I have a lot of respect for the office there - the way it was.

O'NEILL: You mentioned that early on you didn't really see that networking was going to have much impact. When did your opinion of it start changing?

ORNSTEIN: Well, I don't know that I could attach a date. I think it was, probably like most people, as I began to see it really starting to enter into my own life - in the message world, in particular. I began to realize that that created a whole new set of possibilities that had not been foreseen by me. (Laugh) If it had been understood by others, it certainly wasn't emphasized to me in the early days that this was going to be so important. I would have probably scoffed at that too, because it wouldn't have seemed as though it was worth all the fuss simply to get messages from

hither to yon. I think it's been a gradual process, the whole growth of the network business. In retrospect, it's pretty easy to see that it's an essential ingredient in the development of information technology. But at the outset I just got interested in building the thing and making it work, which is what I think computer people are often about. It was only after that, when I stood back and drew breath and saw people starting to develop systems that were actually using it, that I began to wake up to the possibilities. Then, of course, it's blossomed and it's continued to blossom. In retrospect such naivete looks laughable, of course, but it wasn't so apparent at the time.

O'NEILL: You mentioned that you felt Wes Clark is maybe not given credit for some of his ideas. Are there parts of the ARPANET story that you disagree with?

ORNSTEIN: No. You see, I don't know about that part of his career. I knew more before. I didn't know anything about his involvement in the network. It was prior to my own involvement, I think, that he had talked to Larry, et al.

O'NEILL: I meant other than Wes Clark. Are there other parts of the story that you feel are out of whack, or that you would emphasize differently?

ORNSTEIN: Well, I'm not sure, of course, what the story will say when it's done. I don't think that there is a lot of lore out there that is wrong or distorted historically. But I don't know much about what people think happened. I guess I won't see that until we see this history published. We were just in Hawaii a couple of weeks ago and I tried to look up Norm. I haven't seen him for years. As it happened, he was away. But I have felt all along that he has not been given sufficient credit for the initial Aloha System ideas. He's not much mentioned. I mean, you know, there's also Doug Engelbart, the fellow who invented the mouse. Everybody thinks Alan Kay invented it, but it was really Doug.

O'NEILL: Engelbart?

ORNSTEIN: Yes, Doug Engelbart. In recent years there has been a certain amount of recognition of his role. A lot of

ideas were Engelbart's. I think what's happened is that a number of people like Clark and Engelbart have had some important ideas at one time but then subsided from view. History writes over people pretty quickly. In Wes' case it's a great shame, I think, because he really was a significant contributor in many, many areas. He's a really bright guy, but now he's no longer in the forefront. When the guys at PARC invented what was called multitasking in the Alto, they were sitting and explaining it to Wes, who was a consultant to PARC at that time. They were very proud of this notion that they thought they had invented (multitasking) and Wes, stroking his chin, said to them, "Tell me in what way does this differ from the multiple program counters on TX2?" And they said, "What multiple program counters?" They had applied for a patent for this idea because it was a very good idea, but Wes explained to them that it was exactly the same idea, precisely, he had designed years before into the TX2. So they withdrew their patent application. I know that, they know it, Wes knows it, a lot of people know it. But in that book ("Fumbling the Future") it is never mentioned, and one of the major things that PARC is credited with in the book is the invention of the tasking hardware in the Alto. It was not that they stole it at all. They didn't know about it. But it had been invented almost 10 years before. They certainly reinvented it; I have no question about that. Their integrity is high. But I wrote a letter to the guys who wrote the book, saying there are a number of counts on which you have just not done your homework.

O'NEILL: Did your work at PARC build on your networking experience?

ORNSTEIN: Not much, some. Some, but not really. Hardware design is a pretty general technique. It's like writing software in that regard. I always preferred to involve myself in something new. No, I got involved in printing. The first project I did at PARC was the development of some specialized printing hardware for a laser printer. The first laser printers were of course done at PARC. A redesign redistributed the work between hardware and software, and I did have the hardware part. That was the first thing I did and had nothing at all to do with networking. It was just another hardware project. I don't even remember all the things I did. A major one, later, was the development of the DURADO computer, which another group in Xerox had been going to build, but which we ended up building ourselves for our own purposes. It was an extremely powerful personal computer. I and another guy were nominally in charge of the project. Do you have a list of people that you're contacting? Is the list set in concrete now?

O'NEILL: I wouldn't call it set in concrete, but the amount of time that we have to do the interviews is becoming more limited. We can talk about that.

ORNSTEIN: I just wonder whether there are people that I could suggest. I assume you're talking to Taylor or have?

O'NEILL: Taylor has been interviewed.

ORNSTEIN: Well, Taylor certainly is an important one. And obviously Roberts and Kahn. I'm just trying to think whether I might have any idea of any other people who would be knowledgeable - and I assume Steve Crocker.

O'NEILL: Well, the people I'm interviewing for the networking case study include Dave Walden, Will Crowther, Alex McKenzie, and Frank Heart. Howard Frank who worked at the Network Analysis Corporation. Leonard Kleinrock. I don't have a list here with me.

ORNSTEIN: Not Steve Crocker though, huh?

O'NEILL: Steve Crocker was not on the list. I attempted to contact Vinton Cerf who was at UCLA earlier, and it's not certain whether I'll be able to talk to him.

ORNSTEIN: None of those guys has had the good sense to retire.

O'NEILL: Is there a particular reason that you mentioned Steve Crocker?

ORNSTEIN: He was heavily involved in the network protocol development at one time. I don't know whether he has persisted. As you interviewed some of these other people, did they not also suggest names of people? How did you come by your list of interviewees?

O'NEILL: Yes. Very early on, people like Taylor and Roberts and Kahn were interviewed, and they had suggestions of various people to interview. In addition to that other names that keep coming up, for instance, people like Paul Baran. In the networking area, it makes sense to - Wes Clark is another name that was on the list as well. I don't have a specific interview set up with him.

ORNSTEIN: You'll enjoy meeting Wes.

O'NEILL: That's a case where there is some flexibility, and I will make more of an effort to do that. Frank Westervelt is another one who is on the list.

ORNSTEIN: I know the name; I don't know him. Well, Crocker was on committees of various kinds developing host protocols. You might ask someone else about him. I would rely quite a lot on Frank Heart. The trouble is that people do have friendships and points of view, so you want to sort of average these a little bit. If you start out with only the pointers that those first people you interviewed gave you, you really need to follow the references another layer deeper. Steve Crocker is certainly a name of somebody who at one time was an important network guy. I don't know what's happened to him since then.

O'NEILL: Yes, then in some cases, if you know multiple people are working together, you try to just hit one or two, see if there's a lot of differences and then decide whether you need to interview more.

ORNSTEIN: Will Crowther is not going to be terribly much help, by now. He tends not to remember things and he was never much politically involved. He always had his head right down in the bits.

O'NEILL: He said that he didn't remember a lot. But I'll in Boston anyway, so it's fairly easy.

ORNSTEIN: What's the schedule for the writing project?

O'NEILL: The contract is supposed to be completed, I believe, in May of 1991. A little more than a year from now. So the writing and whatnot should be done by the end of this year, more or less. There are multiple people working on the project as well, so we have various completion dates. When I was referring before to the story as it exists, there are some historical pieces - Larry Roberts has done a number of them.

ORNSTEIN: I haven't read those; I probably should.

O'NEILL: Then there was a conference last August, I believe, down at UCLA that talked about the beginning of the ARPANET which, unfortunately I wasn't working on the project at the time, so I didn't attend. A number of people gave their reminiscences.

ORNSTEIN: I'm surprised I wasn't invited to that. I was enough involved in it, certainly, at one time.

O'NEILL: I just assumed you knew about it.

ORNSTEIN: No, I didn't. I very likely wouldn't have gone anyway, but I'm surprised I didn't know about it.

O'NEILL: I'm not even sure who organized it.

ORNSTEIN: Sure. And I have indeed dropped out, as it were, from the computer field to a large extent. You know, people are to be forgiven for not including those who are out of it. The people who are in it are busy enough. For me computing was always intense, but a little bit of a hobby, as it were. I had multiple other interests - more so, I think, than many people. I've been pursuing those since leaving the computer field. Some of the smartest people in the field claim that the field is really somewhat shallow. We're not probing the secrets of the universe and the things that nature presents us with, but rather we're probing our own inventions to some extent and trying to get our thoughts organized. It's a synthetic rather than an analytic field, except for artificial intelligence which, as far as I'm concerned,

doesn't have its hooks into anything profound yet. The rest of computer science still is pretty much engineering, and in the long run I find and I think a lot of people have found it's not something that's intellectually consuming. I chose to get out of it when it looked like the next ten years was not going to be sufficiently different from the preceding ten years. Having gotten into it in the very early days, the creativity curve has been like this, and we're up here on the flat part now, where next year's machine will only be ten times faster than this year's machine. So what is the big deal in that? The ideas that were developed from the 1950s to the 1980s were tremendous. A lot of understanding about how things should be organized. I've never regretted getting out when I did in 1983 because there are plenty of other things to do.

O'NEILL: You mentioned your anti-war activities and founding of the Computer Professionals for Social Responsibility. Did it ever concern you, working under government contract, under the military?

ORNSTEIN: Sure. Sure, it did. But, you know, the network project seemed to me to be benign enough. I was able to say, look, the government shouldn't be funding things *in this way*. I mean the ARPA office, as I say, I believed was very well done, but I would have much preferred the funding of scientific research to come under an office of scientific research of some sort. The National Science Foundation, were it differently organized and differently run and a different kind of thing, should have been the place where scientific research was funded. It shouldn't have been always pursued under the guise of the military. Because the possibilities were much wider than just military applications, obviously. I felt that it was wrong that the country should be always having anything that was important to have a military flavor - otherwise you couldn't sell it to the taxpayers. So I thought that that was wrong. I still do. The particular thing that caused me, and a number of other people, to form Computer Professionals for Social Responsibility was that I was just terribly concerned that the way things were going in the world about that time (1980 - 1981) nobody seemed to have any grasp whatsoever what they were talking about. There was a lot of sabre-rattling going on, as you may remember, and I thought to myself, "They're going to kill me, and they're going to kill my kids, and they're going to kill my grandchildren, and they're going to cause havoc that will go on for thousands of years - without having any idea, these idiots." So I, along with everybody else, was really worried. I felt that there was a serious risk at that point that somebody was going to do something *really* stupid. A number of

us got together and talked about it. I was not inclined to form yet another organization because it just seemed like there would be overhead and waste motion. But I was eventually persuaded that the computer field was sufficiently closely connected into the arms race that it was important. And furthermore that computer people were thought of, to some extent, as wizards within the society, the way doctors and priests are, and that people might listen to us where they wouldn't listen to shoe salesman, you know, talking about it. So we formed our organization. Now, of course, that threat, as far as I'm concerned, is much diminished. It's turned out that there are many other social issues involved in how you go about using computers - crime detection, crime prevention, and so forth; there are all sorts of things. So although the organization has continued, I'm no longer directly involved in it. I was the founder; I was chairman for quite a few years. But now I'm just on an advisory board, and it functions on its own. I guess they're running almost a third of a million dollars now a year. Some of the money comes from foundations, a lot of it comes from individuals. But there have been a number of issues that are important to consider, and it's been a sort of social critic of the computer field and where it's going, saying we ought to be over here when we're over there. That's a useful function, it seems to me.

As far as the Vietnam war was concerned, I thought that it was a horrible mistake from day one, because I had followed the history from the time that the French were thrown out. I knew what was going on, saw what was likely to happen, and felt that we were just killing people and killing our own people needlessly and that what happened was going to happen sooner or later and the sooner the better - lives would be saved... like a lot of other people. In general it is unfortunate that matters have evolved in such a way that so much of what gets done has got to be done with a military cast. I think that that is out of style, really. It's a 19th century idea that we're hanging onto in the face of major changes in the rest of the world. We're somehow a couple of generations behind in our thinking, and we don't know what to do about it. Of course the military industrial complex is entrenched in the society - you know if you arbitrarily shut down a lot of these contractors there'd be serious dislocation of the economy. The thing that worries me is that we don't seem to be doing enough to work towards changing, and it's clear it's going to be more painful if we delay. So those are the kinds of things I'm working on now.

O'NEILL: I've asked all the specific questions that I wanted, unless there's anything else you wanted to add on the

ARPANET?

ORNSTEIN: Let's see, what were the other projects in addition to the network that you were working on?

O'NEILL: The other case studies that will be part of the history are time sharing, artificial intelligence, and probably graphics. They will be four specific case studies, and then there will be an overall history of the IPTO office using evidence presented in the case studies to make the points.

ORNSTEIN: As part of your study of time sharing, you should also talk to Wes. Because Wes was an important critic of time sharing - always.

O'NEILL: He was a critic of time sharing?

ORNSTEIN: Absolutely. As were a number of us who felt that it was a silly waste. We believed that in the long run what has actually happened was bound to happen, that it was going to be economically feasible to provide individual machines and avoid the overhead associated with trying to share a single giant machine. Wes was a very prophetic guy. Just after I had joined Lincoln Laboratory (that would be 1954-1955) Wes was talking about how it was not going to be too many years before we would have a computer that you could hold like in your hand. At that time computers were filling buildings larger than this one - a single computer. And he said, "Yeah, you'll just paint 'em on your desk just like that." I remember the gesture even; he was giving a talk about TX2 at the time. So a lot of us felt that time sharing was an enormous waste. As it's turned out, it's not important any longer. It's an important place in history, but the advent of very powerful personal machines these days - you go into a research center, you don't see much in the way of time-sharing in places these days. It's small machines - networked together, a very different emphasis.

TAPE 2/SIDE 2

ORNSTEIN: Wes was a vocal critic at the time, when he was a much more important figure. I was a minor figure compared to Wes at that time. You can get from him his feelings. He's very articulate and will remember, I'm sure.

O'NEILL: Did you use any of the early time-sharing systems?

ORNSTEIN: Some. Not very much.

O'NEILL: Was your criticism based on your usage of them, or just the whole idea?

ORNSTEIN: The whole idea. I read a lot about the systems, so I knew a good deal of what was going on. It seemed to me that a tremendous amount of effort was not directed at dealing with people's problems but with the problem of sharing resources. It derived from the notion that machines were always going to be big and expensive and that you therefore *had* to share them. We were the first group that started working on the notion that everybody would have his own computer and that you didn't have to worry so much about sharing - that sharing should come about by means of connecting individuals' machines. They had it backwards is what I'm saying. Instead of having this big humongous machine and figuring out how to parcel it out so that you didn't interfere with one another, our feeling was just that people should have individual machines and that you should be working on the problem of connecting them for those situations where they needed to be connected. That has turned out to be the correct view. History has shown that the economics have worked out so that now we can make extremely powerful, inexpensive, personal machines. The problems now are networking, communication, hooking these individual things together rather than parceling out a single big thing.

O'NEILL: When you say we, are you talking about Wes Clark and yourself?

ORNSTEIN: Yes, and the group of people involved.

O'NEILL: Generally the people working on LINC?

ORNSTEIN: Well, and at DEC there were people who were working on the PDP5 and the PDP8. These were small and were thought of as one-user machines. That was the beginning of the notion of one-user machines. Up until then, you either time-shared, in parallel, or you shared serially. TX2, for example, which was a powerful machine, in the early part of its life was shared but only in series. You signed up for a specific time, and when that time came you had the whole machine to yourself. It was only much later that it was time shared. I think we understood that TX2, while a powerful machine at the time, only barely represented what was going to be available to people individually in future years. After the LINC group left Lincoln Laboratory, time-sharing became a big shtick for a while because, of course, it was a while before the development of miniature stuff that made it economically possible for people to have powerful individual machines. For a while you could only access a really powerful machine as a shared facility. During those years TX2 became a timed-shared machine, ironically. But it was certainly never conceived of as that originally.

O'NEILL: Plus there was the data problem, too, getting large data bases and how to work all that.

ORNSTEIN: Sure, absolutely. People now have developed other means for doing that. But networking had to develop to the point where that became feasible. So there were lots of things in the way. But, as I say, there are two basic ways to go about sharing, and we were always pushing in the early days towards the way that I think has pretty much taken over the world. When I say we, there were gangs of people from Lincoln and I don't remember just who. Wes certainly, and people I worked with in the LINC group and at St. Louis, were always convinced that personal machines were the way to go because that was our religion at the time and eventually it became other people's religion too.

Networking, you say, and artificial intelligence. What to say about artificial intelligence? I have very strong opinions. Although I've never worked in the field, I've watched it closely. It was what originally got me into the computer business. I went to a talk - when I was working as a geophysicist, my friend Briscoe took me to hear a talk - by Wes Clark and Belmont Farley on neural nets. In 1952 or 1953 they had done some early stuff on Whirlwind.

There was a classic paper that they wrote then. Some of the ideas that were spawned there were ideas that had to do with things that have later come to be called artificial intelligence. The term, of course, didn't exist at that time. I have voiced some pretty strong negative opinions about the artificial intelligence field because I think that it has not always exhibited scientific and intellectual integrity. There I felt the ARPA office allowed itself to be bamboozled some of the time - to spend money in ways that I thought were wasteful. Artificial intelligence always was and remains the only interesting part of computer science as far as I am concerned. But I preferred to work on problems which had some chance of being solved in my lifetime. I felt that the serious, the interesting questions in artificial intelligence were not in that category. I still feel that way. I think that the phrase has been misused a great deal. There are some fundamental problems that have to do with what I think of as artificial intelligence, having to do with pattern recognition, that very little headway has been made on. And most of the things that are *called* artificial intelligence now are just standard programming practice, good clever programming techniques. Have you read Dave Parnes' critique of the computer part of the Strategic Defense Initiative? Do you know Dave Parnes? He's a first class software engineer.

O'NEILL: No. I'm not the one doing the artificial intelligence case study.

ORNSTEIN: Well, Parnes was a very vocal critic of the Strategic Defense Initiative. He wrote a series of critical short papers. He was originally on the committee that eventually became the Eastport committee. I don't know if you know about that. He was a member who decided to quit because he felt that it was a fraudulent effort, to justify building something that would never be trusted to work. He is one of the world's top software engineers. He's also very articulate, and he wrote a short, pithy series of papers that he handed in when he resigned from this job. The one on artificial intelligence is one of the best short critiques of the field that I have read. I would refer you to that because he articulates it far better than I. Mostly I think what I feel is that the field has never really dealt with the core problems. It doesn't know how. It has worked around the edges and called it artificial intelligence, but I don't think that it has made much real headway. I've seen so many promises made, so many people who should have more integrity, promising things that they must have known in their heart of hearts they would not be able to deliver, just to get funding. I feel that has really been very destructive. As a result the field has gotten a bad name for many

people. A few years ago, a big mania for "artificial intelligence" took over the world for a while. Everything became AI. I worked on a music program, for example, and people said, "Oh, that's artificial intelligence." Well, it had nothing whatsoever to do with artificial intelligence. Every piece of code you wrote became artificial intelligence for a while there. It's a shame because they have prostituted what is a really important and intriguing set of problems, it seems to me. So I feel that they have done a disservice to the field.

O'NEILL: Would you put expert systems in that same category?

ORNSTEIN: Yes. I don't think of expert systems as, for the most part, addressing real artificial intelligence problems.

O'NEILL: Can you give me the name again of the fellow who wrote the articles, the critiques?

ORNSTEIN: Dave Parnes. I just got a publication from him, so if you want, I could probably dig his address up for you.

O'NEILL: I was wondering if it had been published.

ORNSTEIN: Yes, it was published, but you know I don't really remember where. Let's see. Laura might have copy. I'll look it up just before you go. Parnes is quite an pithy young man. He wrote some very sharply worded criticisms of the artificial intelligence field that I think are justified. There's also another book you've probably seen, John Shore's *The Sachertorte Algorithm.* Do you know that book?

O'NEILL: No, I don't.

ORNSTEIN: Well, you should have that as bedtime reading. We'll get that reference for you as well. It's a book that starts out dealing with computer phobia, which is, after all, presumably a relatively elementary kind of thing. But in

the latter part of the book he really critiques the whole field to some extent as I recall. He's another excellent software engineer and also very articulate. The book is very well done; it's both accessible and amusingly written.

O'NEILL: I think those would be helpful.

ORNSTEIN: If you're interested in the history of the field in general, I think that both of those items are helpful. Is that enough? You got what you need?

O'NEILL: Yes, I think so. Thank you.

END OF INTERVIEW