

An Interview with
ROBERT E. KAHN
OH 158

Conducted by William Aspray

on

22 March 1989

Reston, VA

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

Robert E. Kahn Interview
22 March 1989

Abstract

Kahn briefly discusses his educational background and involvement with the development of ARPANET before focusing on the operations of the Information Processing Techniques Office (IPTO) of the Defense Advanced Research Projects Agency (DARPA). Kahn describes the development of computer networks with DARPA support and explains the process of contracting research at DARPA as well as the creation of DARPA budgets in detail. In this context he discusses the work of various DARPA and IPTO personnel including J. C. R. Licklider, Vinton Cerf, Larry Roberts, and George Heilmeier. This interview was recorded as part of a research project on the influence of the Defense Advanced Research Projects Agency (DAPRA) on the development of computer science in the United States.

Robert E. Kahn Interview
22 March 1989

Tape Index

Tape 1/Side 1.....	1
Side 2.....	12
Tape 2/Side 1.....	21
Side 2.....	31
Tape 3/Side 1.....	41
Side 2.....	50

ROBERT E. KAHN INTERVIEW

DATE: 22 March 1989

INTERVIEWER: William Aspray

LOCATION: Reston, VA

ASPRAY: Let's begin by taking a few minutes to talk about your own background before you came to DARPA. As I understand it you got a Ph.D. at Princeton in electrical engineering in 1964. What was your area of work? What was your specialty? What were your interests?

KAHN: The program at Princeton at that time was largely oriented toward applied mathematics and theory more than the practical side of engineering. My doctoral work there had to do with sampling and representation of signals. It was more an applied mathematics kind of thesis. From Princeton, where I graduated in 1964, I went up to MIT to join the faculty there. Actually, I had worked also at Bell Laboratories as a member of the technical staff, both just before and during the time that I was getting my doctorate degree.

ASPRAY: I see. The same kind of work there?

KAHN: At the laboratories I had actually started off in the Traffic Group. This was the group at Bell Laboratories that was responsible for the plant engineering of the Bell System -- a lot of interest and concern in things like queuing theory, analysis of switching performance, and global engineering of the system. This was the time when most people going through the labs were getting into more esoteric areas of communications: modulation theory, coding, color television, things like that. I opted to go into the telephone engineering part of the business, which was actually quite interesting, and heavily stocked with mathematicians I might add -- people who were doing work in combinatorics and queuing theory and the like, but an awful lot of seat of the pants engineering. It was headed by a fellow named Roger Wilkenson, a sort of Walter Cronkite of the Bell System at the time -- a grand old gentleman.

ASPRAY: I see.

KAHN: When I got to MIT I joined the Research Lab of Electronics. The particular group that I was in was headed at that time by Jack Wozencraft, who subsequently went over to Project MAC, and then, I believe, after a short tour at Lincoln Lab went out to the Naval Postgraduate School. I think he actually went back to MIT in between. I was at MIT for only a few years. I left in 1966 on a leave of absence, basically at the recommendation of Jack, who thought that it might be beneficial if I got a year or two of practical experience under my belt, since I was largely a mathematician. That group had quite a bit of fundamental grounding in practical realities of engineering. It was some of the best advice I ever got. I took a leave of absence in 1966 and went to Bolt, Beranek and Newman.

ASPRAY: How close were you to the people that were working on Project MAC? Were you participating in any way?

KAHN: Project MAC was in a different building, in a different part of the campus. It was headed at that time by Bob Fano, who had left the group that I had just joined to go over to Project MAC. So I knew quite a bit about Project MAC. I knew many of the people there, and in fact, I had subsequently worked with many of them at BB&N, but I had very little direct interaction with the Project MAC group at the time. My work at MIT was largely focused on more mathematical sides of communications. I was doing some continuation of my work on representation that I had started at Princeton. I was doing some work on rate distortion theory, which was another aspect of representation, and a variety of problems that were largely of theoretical interest to me, but weren't necessarily tied to any specific practical engineering problems. So the reason for leaving was really to get my hands a little more dirty with some of the practical problems of everyday engineering. I think I managed to do that at BB&N. In fact, when I got there I picked computer networks as a problem area to work on. This was a time when BB&N was largely a company that had made its reputation in the field of acoustical engineering. It had done a lot of the acoustic design for some of the major auditoriums around the country -- some notorious, I might add (laugh). It was a marvelous little firm at the time, a fraction of the size that it is today. It was largely a group of very professional people doing an extension of what MIT and Harvard were doing, except on a full-time basis, without the responsibility of dealing with students and teaching.

ASPRAY: It seemed to me that there were a number of people who went over from MIT to BBN. There was a pretty close connection.

KAHN: An awful lot of the key people in the field of computing and acoustics were involved there. I am not even sure I could give you a complete listing of all of the people, but there were very many. BB&N got the contract from DARPA around the time that I came, or possibly shortly before. I do not remember the exact timing of it; it could have been a year or two before -- to work on a variety of problems having to do with timesharing and its applications, artificial intelligence being one of the things that they were applying timesharing to. A lot of good people were there. It was really an exciting place to be. Some of the key people who were at BB&N at the time that I got there were people like Jerry Elkind, Jordan Baruch, Leo Beranek, Dan Bobrow, Warren Teitelman, Ross Quillian, Frank Heart. Over the years that I was there a number of other good people came in to join us, like Bert Sutherland.

ASPRAY: Were there any people in particular that you worked especially closely with?

KAHN: Well, I started out in the research group at BB&N. There were a number of different divisions. The group that I was in was headed by Jerry Elkind when I came. In fact, shortly after I came, BB&N got restructured into two separate sections, one of which was headed by Jerry Elkind, and the other of which was headed by Leo Beranek. Jerry not only ran half the company, he also ran one of the divisions under his half. So I reported to him at one of the divisions. Now, another division that reported to him was one that was run by Frank Heart. That was the group that Jordan Baruch had originally run, I believe, until Jordan left to set up some kind of a timesharing service jointly with General Electric to serve Massachusetts General Hospital. I think Frank inherited his whole group. It started out as a medical electronics or medical timesharing group, but it soon got into networking. When I had joined BB&N I began to work on issues of networking. It was the area that I had chosen to get into. And I was largely involved in design of various techniques for networking such as, error control techniques, flow control techniques, and the like. I was not familiar with ARPA at the time. I knew that ARPA had been involved in funding Project MAC, but it was not an organization I was directly involved with. I did not know the people there. But at one point, after I had started to work on networking, Jerry Elkind came up to me and said, "You know, it is really interesting; you really ought to let the people at ARPA be aware of it, because I think they would be interested." So I recall writing a letter to Larry Roberts, who had just gone to join ARPA.

ASPRAY: Did you know Roberts?

KAHN: I did not know Larry either. I still have the letter in my file somewhere. In retrospect, it is rather interesting, because I had no idea at the time who Larry was, or what he was interested in. Jerry just said, "Why don't you write him a letter and tell him what you're interested in?" So I wrote him this very short letter that basically said, "I have some very interesting ideas in networking and I thought these might be of interest to you." And if he had some time, I'd be happy to explain it to him. Well, of course, he had gone to ARPA with the idea of taking over their networking project, that I guess conceptually originated back in Licklider's time and had become a reality under Bob Taylor. Bob Taylor had brought Larry in to run the project for him. It was all news to me at the time. So I sent this letter out. Of course, I got a call back from Larry. He said, "Why don't you come down? Let's chat." So we did, and, of course, I found out at that point that he was actually interested in creating a real network. Having been a mathematician and theoretician, it really had not occurred to me at that point that I might ever get involved in something that could become real. It was a revelation of sorts to me. In any event, in mid-1968 or so DARPA actually issued an RFP for the ARPANET. This is a topic I know you didn't want to get into, but...

ASPRAY: Well, does it overlap?

KAHN: It certainly does.

ASPRAY: Please go on.

KAHN: They issued an RFP for the ARPANET, and Frank Heart, who was now running this other division, was very interested in getting involved in the project. It turned out that I had done more thinking in that area than anyone else at BBN. So Frank showed up in my office one day and said, "I understand from Jerry that you've been thinking about the networking area. Can we chat about it?" I said, "Yes, sure. Who are you?" He introduced himself and we had a very pleasant chat. When the RFP came out, I basically got very strongly involved in writing the proposal. One of the fellows that I worked with really closely on it was named Severo Ornstein, who is now retired -- he subsequently left BBN to go

out to PARC sometime in the 1970s. Severo was a pretty good systems guy and a hardware designer. He represented the hardware side of things to me. We put together a very well thought out proposal. It got sent in and BBN won. In fact, at the time the proposal was sent in I had no plans to continue actively involved in that project. My thinking at that time was more along the lines of finishing up my "practical experience" and going back to MIT. Writing something like this proposal seemed to me to be the epitome of practical experience. Then when we had won the contract I decided that, for a variety of reasons, it was better for me to get involved in the project directly, because there were an awful lot of issues that still needed to be resolved. There was a whole raft of system design issues that remained...Well, actually, I picked myself up and moved over to the engineering group that Frank was running. It was actually in a different building. So it was a physical move.

ASPRAY: What year was this?

KAHN: It was probably the beginning of 1969. And I stayed over there for about three years while we created the net. We went through several iterations of the system. When the net was up and pretty stable I then moved back into the research group. By that time Jerry Elkind had left and gone to MIT. Dan Bobrow had taken over running that division. I moved back into the research division very briefly. I stayed there for a period of maybe a year. During that period I spent most of my time actually planning a demonstration of the ARPANET -- a public demonstration of it in Washington, D.C., which we pulled off in October of 1972. Actually, that demonstration was what made the ARPANET real to others, because there was a lot of skepticism before: people could now see that packets switching would really work. It was almost like the train industry disbelieving that airplanes could really fly until they actually saw one in flight. It just did not quite seem like you could communicate by breaking messages into packets and shipping them; circuits seemed to be much more reliable to the untrained eye. Well, the demonstration was a major success. We had many different vendors contributing terminals. We had just about everybody involved in networking at the time there. Over a thousand attendees came through the conference. It was a hands-on, live demonstration. Many of the people who were involved are well-known leaders in the field today. It was a major event. It was a happening. You might be interested to know that I actually taped that session, so if you're interested for historical reasons I'd be happy to make the tapes available.

We went around and interviewed people as to what their reactions were. This was the first public viewing of computer networking.

ASPRAY: I think that we would be very interested in that.

KAHN: So I would be happy to give you that. I also had a photographer go around and snap photographs of people at the console so you could get a feeling for what it was like. I don't know why I did that. I have the archives on that. I also had a film prepared in which we interviewed some of the key people who were involved in networking at the time. I showed the film at that conference in 1972. I have a videotape of it, which I'd be happy to give you. It's a 16 millimeter film, which was transferred onto videotape.

ASPRAY: Yes, I would like to...

KAHN: It is a fascinating film. I think the film is as relevant today about the objectives of networking as it was back then. Interestingly enough, not that much has changed in the global sense. There are some technological differences. We now have workstations instead of mainframes. The computer industry has moved light years in that timeframe. We have local area nets where we didn't before. We have internetting where we didn't before. But, in terms of what people wanted to use machines for, a lot of that still remains a goal and objective today as it did back then. I think you will find that film interesting and if you like, I will just let you take the copy of it with you.

ASPRAY: Whatever you feel most comfortable with.

KAHN: Okay. So, let's see. I think just before that demonstration I actually left BB&N. We did the demo and then I joined DARPA. It had appended a D to its name by then, the silent D in front of its name back then. It became non-silent somewhere in the 1970s; I am not sure when. They never changed the name of the ARPANET. That always stayed ARPANET. They still issue funding on ARPA orders -- not DARPA orders. So there is some historical lingering to the

term. I joined DARPA at the end of 1972 (October 29, I think, is the actual date). But I don't think I actually showed up in the office until sometime in November because of the conference and wrapping up at BBN.

ASPRAY: Okay. What were your duties supposed to be when you showed up?

KAHN: Well, I ran a number of programs. My interests were fairly broad, and going into networking was, in some sense, a fairly significant change to me. Maybe it was not significant in some global scheme of things, but it felt like a big change. When I went to DARPA my plan was to start a program in flexible automation in manufacturing. I forget the name of the program. It was something like Manufacturing Automation, Factory Automation, Flexible Manufacturing, or some name like that. DARPA had either proposed it to the Congress, or were planning to propose to the Congress the startup of a hundred million dollar a year project. At least that's what I recall of it. When I joined DARPA, somehow that planned project had vanished within the first few months of my being there; what had been a planned program, or an approved program, whichever, suddenly disappeared. Larry Roberts, who was the head of the office at that point, had convinced me to go back and start working on some of the networking problems again, which I agreed to do.

ASPRAY: Reluctantly?

KAHN: Well, at that point in time I was looking to get into some new things for a variety of reasons, but I guess my feeling was that it was a big disappointment that that program didn't happen. On the other hand, I was now drawing on my strengths rather than branching out into a brand new area. So it may have been better for me professionally; I am not sure.

ASPRAY: Were you working within the IPT office at this time?

KAHN: Yes, right. I do not think I ever moved more than a hundred feet in all the time I was in ARPA, which was about 13 years.

ASPRAY: So this automation program was to have been something that was part of the program at the IPT office?

KAHN: That was my understanding. You would have to go back and talk to Larry Roberts about that because I was not privy to the details. I had just took it on faith. At any rate it did not happen. I am sure that it was fully their intent to start it. It was probably a congressional cut-back of some sort that happened. Those things happened all the time -- sometimes even in the middle of a program. Let's see, there is one other thing I should mention to you, which is that just before I had gone to DARPA, Steve Levy and I had spent quite a bit of time putting together a commercial organization to try and exploit packet switching technology. This was called TELENET -- it became a well-known commercial organization which I believe was bought out by GTE in 1978 or 1979, or sometime around there -- late 1970s. It is now owned by some other company, like United Telecom. It was a US Sprint company. Maybe it still is. Maybe they said Sprint -- I am not sure. The choice for me at the time was either staying with BB&N, becoming involved in running TELENET, or going to DARPA, and I elected to go to DARPA. It seemed to be more in line with what my long-term interests were, which was R&D, and helping to plan, formulate, and manage R&D programs. Larry Roberts left shortly after I got there to become the second president of TELENET -- Steve Levy being the first.

ASPRAY: Yes. What were your prospects at BBN? What did you think you would have been doing? BBN, that seemed to be an R&D opportunity also, just like DARPA would be.

KAHN: I do not think that I would have done that. I think my choice at that point would have been...My highest priority, of course, was to go to DARPA. My second priority probably might have been to go with TELENET. So my third priority would have been to go back to MIT.

ASPRAY: I see.

KAHN: No, I take that back. I think my second priority would have been to go back to MIT. I never seriously explored with them going back on the faculty at that point, because my first priority was to go to DARPA. But I think going back

to MIT would have been my second priority, and my third probably would have been staying to run TELENET, and the fourth probably would have been to stay at BB&N proper.

ASPRAY: I take it that by the time that you went to DARPA in 1972 you were rather thoroughly familiar with their operation.

KAHN: It is hard to really get familiar with their operations from outside. I do not think that there is any way that as a contractor, or an advisor, or an outside agent, you can really understand what goes on inside that organization. The dynamics of the place -- at least back then (maybe it's changed now; I would strongly doubt that) -- were so great that from day to day the changes were great. A lot of what it took to make things happen there really depended not only on understanding the mechanics of the organization -- how it worked, but on the politics, both with the organization, the agency as a whole, and with the Department of Defense. I do not think you can really get that observing from the outside. You might look at the White House and decide you really understand how the White House works, but you probably have an idea that somehow being inside gives you a completely different perspective on just how decisions are made, what kind of paperwork has to be prepared, who needs to see it, who can influence what, how are things carried forth, what the process is. All of those are things that you can get hints of, you can get glimmers of. You know that these kind of meetings take place, but you do not really know what happens in the meetings, you do not know what the personality give and take is like. You do not know how the forces are really exercised.

ASPRAY: Well, what about in another, more limited sense, in the sense of personally knowing the community of contractors, having some sense of the technical areas which the IPT office is funding? Having some sense of the workers and the project areas, you were somewhat familiar with that by that time, I assume. Or not?

KAHN: I was probably about as familiar with what the office was sponsoring as one could be outside of DARPA at that time. I am sure there were people who had more detailed knowledge of specific things, but in terms of an overall panoply of what was going on I think I had that, because of the fact that I was so deeply involved in this networking project. Remember, this was trying to link up the different ARPA IPT community participants at the time, which meant that I got

to know each of the different groups. We had to interface with them. We knew what their facilities were. To some extent we had to know what they were doing with those facilities. We knew the people there. One of the key contributions of the ARPANET itself was to take a set of researchers around the country who were largely (I'm not sure "disparate" is the right word to use here), but they were very separated. There was not a lot of communal activity among the agents.

ASPRAY: Independent agents, in a sense? Is that right?

KAHN: Yes, well, nothing is totally independent, but for example, I'm sure the graduate students at MIT did not know the graduate students at Stanford. My guess is that much of the faculty at MIT probably did not know much of the faculty at Stanford. That does not mean that they might not have recognized a name, or seen a paper, but the interchange was largely at a much more distant, arm's length level. What the ARPANET did was knit this community together more tightly than it had ever been knit together before. You now had the ability to actually interact with people, not on day one, but over the course of a few years, especially after electronic mail became available. ARPA had started the process of running a series of principal investigator meetings, sometime in the 1960s. I do not know what the origin of that was -- some combination of Licklider, Ivan Sutherland, and Bob Taylor, I presume. They would know for sure how that got started because they were the first three directors. I know that Bob Taylor ran PI meetings. He may very well have started them, but it is possible they occurred even earlier. That brought a small set of people together. By the time I became director we had principal investigator meetings that involved 40 to 60 people, as opposed to a much smaller set. And I was just at a DARPA principal investigators' meeting last November. They must have had several hundred people there. I am sure it was not more than several hundred, but it was several hundred. So it has really grown in size and scope over the years.

ASPRAY: And has a different group dynamic...

KAHN: Well, sure, any large organization versus a small one would. But those meetings helped to start to build a sense of community, and I think the network reinforced that at many different levels. There were also some efforts underway to

have graduate conferences where they brought key researchers from the graduate student ranks together to help facilitate a sense of community, but I don't believe that actually continued very long.

ASPRAY: Yes, Bob Taylor was telling me wistfully that they had been canceled after a few years and that he thought it was a bad thing to have canceled.

KAHN: He's probably right. So, just to close on this one point, I had a fairly good view of what was going on, both because of the network itself and because of the demonstration of the network, which involved me working with each of the sites to actually figure out what to demonstrate on their machines, who to work with, and the like. You have to recall that at the time we started work on networking there was not a lot of interaction between the computer and the communications fields. They were rather separate fields. The focus in the communications discipline was not on computerish things. That does not mean that there was not an isolated researcher or two thinking about that effect. Folks at Bell Labs had been thinking about various ring networks to hook together computers. There had been some isolated work in that area, but the ARPANET project was the first serious one to try and meld computers and communications in a serious way for networking purposes. There had been a little bit of work over in England by Donald Davies of the National Physical Laboratory, but what he did was hook a bunch of terminals up to one mini-computer, just to show you could type on one terminal and it would come out on the other through switching the packets out different lines. So he demonstrated the ability to switch, but he did not have anything that resembled a wide area network with algorithms, much congestion control to worry about, things like that. Much of the research work that was going on at that time was not yet in the mainstream of the industrial thought process. Timesharing really did not take hold until the late 1960s in the real sense. There had been quite a bit of exploratory work. Many of the facilities that were available around the country were special purpose, and they were either put together locally or they involved special-purpose operating systems designed by Berkeley or BB&N. These were not turn-key installations in most cases. Many of the projects were actually timesharing projects themselves, like the one at Project MAC and other locations.

TAPE 1/SIDE 2

KAHN: So in most of the cases the research groups themselves were so deeply involved in the work they were doing. Timesharing was a culturally different phenomenon. Most people had not interacted with machines in quite that dynamic a fashion. The basic use of machines was in batch processing mode. You would submit decks of cards; you got back responses the next day. The idea of interacting with machines was a new sociological phenomenon to a lot of people. The net effect was that although there was such intense interest in what people were doing, it was not that easy for someone to learn about what was going on unless he had some specific reason to be plugged into those groups. I was fortunate to be plugged into the communications field, so that I knew what was going on there. By virtue of being in a company that had been focusing on computing (i.e BB&N) although it was in acoustics as well, and having come out of the communications background, I'd had an intersection of knowledge in both of those areas that I was able to apply. There were not a lot of people that had that kind of background, although I think I was not unique at the time; I was just well-positioned.

ASPRAY: Do you think that was one of the attractions to your being invited to come to DARPA?

KAHN: Which?

ASPRAY: Having this overlapping interest in both communications and computing.

KAHN: Possibly, but I should not conjecture on that. You should go ask Larry Roberts who hired me to come in.

ASPRAY: Okay. Why don't you tell me about some of these special projects you ended up working at when you did arrive at DARPA?

KAHN: Well, they were largely projects that I formulated myself. One was a project that had to do with providing end-to-end security over networks. Another was a project that had to do with moving speech through the packet switched nets. Another one had to do with packet networks over satellites, using satellite channels to do packet switching, and building ground radio nets. None of these were new ideas to me. That is, there had been people working on aspects of

these. For example, there was an effort at the University of Hawaii that Norm Abramson, Frank Kuo, and Richard Binder were involved in. That was really the first project that I am aware of to demonstrate that you can actually send packets over radio channels. If you have ever been in Hawaii, you know that telephone lines are very noisy there (or were at that time). So the idea was to see if you could not somehow communicate to a central computer center using radio waves rather than noisy telephone lines. And they very successfully demonstrated that. In fact, they even had a few experiments to show that you could do relaying to other islands and further locations. So that was a start down the path of radio transmissions. They actually did some experiments using satellites to show that you could sometimes send something up by satellite and get it down. We did not spend a lot of time in that project worrying about sophisticated techniques for designing and managing these nets, because they did not have a whole lot of operational requirement, and they didn't have very wide area nets to deal with. In effect, one of the main attributes of their work was the operation of the set of protocols called "Aloha protocols", which are characterized by the absence of any non-trivial control mechanisms. They were the very first test of the idea that you could have a kind of chaotic network, with essentially no control at all. Every station did what it wanted, when it wanted. If you didn't have too much traffic it would still work. It would sort of be like having a big, open roadway and not too much traffic. You can drive blind, if you are pointed in the right direction, and you probably will not hit anybody. Of course, if the traffic levels build up, you start to get collisions, and then the two packets that have collided would be damaged, and you would have to retransmit them. They actually showed that as long as the traffic was less than about 18% of capacity you could still manage to get it through with reasonable delay. More than that, and you might never get it through, on the average. So they did quite a bit of early work on demonstrating feasibility. Of course, people had been doing network encryption for a long time, usually LINK encryption, rather than across the packet switched nets. The application to satellites was also something that had started in the Aloha Project. What DARPA succeeded in doing later was actually building a real operational system, quasi-operational, that worked on the INTELSAT IV satellite, which we used as the basis for interaction with a number of our European research counterparts. The most difficult part there was not really technological. There were some interesting technological issues of how you used such a system efficiently, and what the protocols should be, and exactly how you structure and architect the system. But a more difficult challenge in that project, as far as I was concerned, was making a system like that happen in the INTELSAT environment. Remember, there was no domestic satellite capability in the United States at that time. The only options we would have had were to go through INTELSAT or to use some military satellite system. We

chose to go with INTELSAT, although the military satellite system was an option that we were seriously exploring. But INTELSAT had no mechanism for supporting this kind of technology at the time. Their tariffs did not even permit it. We had to work through the Board of Governors at INTELSAT to get new tariffs approved; we had to get countries to agree to participate; plus we had to get the systems actually installed and working. So there was a -- I do not know what you want to call it -- a political/administrative challenge to actually make that all happen. That was my real contribution.

ASPRAY: How much of this work that you were doing in this period was research of your own, and how much of it was contracting outside, managing outside, getting others to participate?

KAHN: It was a mix. Let me just describe one or two more of the projects that I was doing in that time frame, and then I will come back and answer your question.

ASPRAY: Fine.

KAHN: The fourth project that I was involved in was one which had to do with Internet -- the ability to get machines on multiple nets, interoperating. You have to recall that in the time frame that we were operating back then the ARPANET was the only network around. Well, I was now in the process of creating satellite nets, and radio nets. Local area nets had not yet evolved. There were no such things as Ethernets, and while rings had been looked at in the lab, they had not been distributed and people did not have them. Dave Farber had been doing some experimental work at Irvine, and I mentioned the work at Bell Labs, but there were not really any local area nets available. There were computers around. The machines were all big batch or timesharing systems. My first application was with the packet radionet that we were creating. It was designed to work in a mobile environment, so all these radios were small and movable. This may have been the very first application of microcomputers in a military system. This was done at a time when the Intel 8008 had just become the state-of-the-art microprocessor, if you recall. National had come out with a machine called the IMP16. It was the first 16 bit version of a micro. These radios were also spread-spectrum radios. It was probably one of the first usages of surface acoustic wave technology for matched filtering. So we built in a little radio about that size [gestures

with his arms] -- what would have taken probably the better part of a room to build before, which meant it could not have been portable.

ASPRAY: Since we're on tape, was the size of the radio a foot by half a foot, or something like that?

KAHN: I would say it was about a cubic foot, and the base was about maybe seven inches square... eight inches square. It stood maybe a little over 12 inches high, maybe 14 inches high. I cannot remember exactly. But it was smaller than a breadbox, bigger than a cigar box -- somewhere in there. You could actually carry it. It was portable. You did not have to lug it, but it was not light. I do not remember the exact weight, but you could easily carry it. You would not want to strap it to your back or anything like that, and it needed power to drive it. But you could put it easily on a Jeep, or move it around -- carry it up to a mountaintop, things like that. We had the system and it could all work while it was mobile, so everything could be in motion. We later demonstrated it in a variety of configurations, including moving aircraft. [TAPE OFF] Let's see, we were talking about all these multiple networks that were in our "envision space" at the time. They did not really exist, but we knew packet radios would be out there. Since the only interesting computers at the time were big, multimillion dollar machines that needed dedicated rooms, air conditioning, and probably operators to run them, it wasn't clear how you could put them on Jeeps. So the question really was, how could we get access to computing to test with the packet radio net? The idea then was to connect the packet radio net to the ARPANET, which had probably 50 or 100 computers at the time and use those machines. It was an interesting challenge in its own right. We knew in the back of our mind that there would be other kinds of nets to connect. So we designed a mechanism for getting a machine on one net to work with a machine on another net in a general enough way that it, in effect, turned out to be the mechanism that was used on all of the local area nets in the ARPA community and throughout Defense when that happened. That's the so-called TCP/IP set of protocols, which you have probably heard about.

ASPRAY: Yes.

KAHN: I did that work with a colleague of mine named Vint Cerf, who now works with me here at CNRI. I'll give you a little background about Vint. When the ARPANET was first installed at the end of 1969, the first node was put into

UCLA on roughly the first of September 1969. It was sort of one a month, and by December of that year we had four of them installed. Shortly after the fourth node was installed, I had gone out to California with Dave Walden, who is now the president of BB&N Laboratories. Dave and I did all of the field testing and debugging of the network to bring it up. At that time Vint Cerf was just finishing his Ph.D. at UCLA -- or was working on his Ph.D. at UCLA, I guess. He had done a lot of the coding for the UCLA Sigma 7 computer, which we used as a vehicle for testing the network. Len Kleinrock had a contract to work on the network measurement center for the network, and Vint was one of the students working on this area. So he and I worked pretty closely on the network testing during that period, along with Dave Walden. Vint went to Stanford sometime, I guess it must have been 1973, after completing his doctorate at UCLA. He and I collaborated on the design of the Internet architecture. We actually wrote a paper on that back in 1973, which got published by the IEEE in 1974, that laid out the strategy. Then Vint and some of his students took the lead in actually taking these ideas and turning them into detailed specifications and real protocol implementations. He led the development effort. There were actually three parallel efforts--one by Vint and his students at Stanford, one by some folks at BB&N, and the third effort by some folks at University College, London. The three of them then interacted together. Vint was really playing a very major lead role in this whole thing. Then, of course, he subsequently left Stanford and came to work for DARPA in 1976, and in fact took over the Internet program from me at that time. I had started it and ran it until Vint came, and then he took it over after that.

Now, to get back to the original question you asked, "What was my role in it?" It varied. In the case of the packet satellite work, I played some technical role having to do with the overall architecture of the system and structuring it, putting in gateways, making sure it interfaced properly. But largely, the technical work on that was done by a team of contractors, including Irwin Jacobs who played a very strong role in leading that whole effort technically. At the time Irwin was president of a small company called Linkabit in California. He is now president of another small company called Qualcomm. I handled the political, legal, and administrative aspects of actually making it happen -- the liaison with the other countries: getting the approvals, working with INTELSAT, the tariff issues. These were not trivial. I wrote a paper that summarized that aspect of it. If you are interested, I can get you a paper on that side of it.

ASPRAY: I think we would like that.

KAHN: A very strong non-technical component to that program was required to make it a reality, because as I said before there was no domestic satellite capability. It was all mired in issues of politics. In the case of the end-end security program, that wasn't so much a problem that was technological as it was also political and administrative. It involved working with the security community to understand how to do this kind of thing: what the options were, what would be permissible, how we might approach it, getting agreement and approval to proceed. It took quite a bit of attention on my part. The technical part of it was turbulent, but there were mostly engineering problems that had to be solved. The Packet radio program was very different from the first two in the sense that I was really the chief architect of that program. At one point it involved close to 20 subcontractors working on different parts of the technology, the theory parts, implementing the radios, implementing some things called "stations", because we didn't have enough power in the microprocessors to do everything, so we had to delegate some of the functions to other, more powerful machines elsewhere. We had people building surface acoustic wave devices. We had more advanced versions of radios. We had a lot of different aspects -- field testing of technology. My role in that program was chief engineer, chief architect, and manager of all the different parties, making them all play together. It was a major technical involvement, as well as a programmatic and administrative one. I think the conception of the internetting effort is one that can be credited to Vint Cerf and myself. He and I really conceived of how to proceed with that, laid out the grand design; and then Vint worked with the community to develop and evolve it over time. It went through many, many iterations and revisions, because as you would try and implement something you would learn all the things that you did not think about before, what was right and what was wrong, and the like. He and I planned the original design, and he took it from there and just evolved it over the last decade. But that was a major intellectual contribution on our part. It was not so much a big management challenge, because when we started it out there were only a few places just trying to implement some things, but it turned into a major activity for defense because the original set of protocols that were implemented by machines, the host computers on the network, did not allow you to deal with multiple nets. I'll just give you a few examples. When the ARPANET was first created nobody envisioned that there would be other networks connected to it. We thought the things you would plug into it would be computers. Therefore, if something showed up to be delivered into the computer, there did not seem to be any reason to tell the computer who it was destined for. If I hand you a letter, there's no

necessary reason for me to tell you it is for you; the act of handing it to you means it is for you. What you want to know is who it came from. Well, if you now go into a network rather than a machine, there might be a hundred different computers on that other network. Which of those other computers is it now intended for? Well, there was no information that would let you figure out what to do with a packet once it went into another net, because all notion of who it was for had disappeared. It was there in the original net, but gone by the time it popped out. Likewise, when you sent something into the net at the source, there was no need to tell the network who it came from. Obviously, if you give me something; it comes from you. So those were some problems that had to be dealt with. The second thing was there was no concept of an Internet name space. You could identify a port on the ARPANET, but you couldn't say, "Send it to that net over there in Southern California, and, oh, by the way, to this machine on that particular net." There was no addressing notion for the Internet, which we had to create and then imbed in the protocols. Another thing that was really different was that the original feel of the ARPANET was like that of a reliable peripheral. That is, the computers treated the ARPANET as more like a line printer -- pick your favorite peripheral. The idea was, if you sent something to the network it would get where it was going with perfect reliability and the computer did not have to worry about it, just like if you sent something to the line printer you figure it was going to print it, unless your line printer is broken, in which case you fix it. But that is a rare event. Well, if you now start to think about going through a radio net, you might be talking to a radio terminal that enters a tunnel. You cannot get there right now, or suddenly there is interference in the airwaves. The packets do not arrive. So the notion of reliable transmission did not exist at all. There was nothing in the original protocol that would let you recover from that. If something got lost, it was broken; you had to go fix it. So we really needed a complete rethinking of the protocol suite which had to be reimplemented in the machines. It was a major effort to make that all happen. Of course, instead of dealing with a network that had a hundred computers, eventually we were dealing with probably tens of thousands of computers. So we had real problems about just managing the name space. If somebody said, "Send this to the following machine," and you knew the name of the machine, it didn't necessarily guarantee that the network could figure out how to get it there. We would need to know specifically what the internal mapping from name to address that was needed to rapidly deliver this message. So anyway, this was a major change that happened. The DOD adopted the Internet suite of protocols as a standard in 1980. And they made a conversion from the old suite of protocols which were called the NCP to TCP/IP--1 January 1983 was the cutover date. I managed the conversion and it was a very interesting exercise. It was almost as interesting an exercise as the original demonstration of the ARPANET

was in 1972 -- to get all the sites cut over, because you now have mismatches of protocols, and you had to coordinate hundreds of sites and make that all work, which we did. That aspect of it was also more of a management, administrative activity. But each of these things required a lot of detailed technical knowledge of what was going on.

ASPRAY: If you can, try to put this in a longer context: when Licklider was originally running the program, he had certain vision about what kinds of programs he was going to want done, but he did not really do any of the work himself. It was all contracted out. He did not have any technical officers there with him. His aide was a military person. Over time, it seems that people besides the IPT Director were brought in to do some technical work, or to manage technical projects -- I am thinking of Larry Roberts in particular. When you got there, you were clearly in this same kind of vein. You were coming from a technical background; you were doing some research as well as managing programmatic administrative work. Were there other people like you, and was this a longer tradition? Had there been a tradition by the time you arrived of having these kinds of technical people within the program working on special projects?

KAHN: Well, I would say that in some sense Licklider really had a good technical vision of where he wanted to go, and used that to guide what he selected in terms of funding. You would have to ask him -- but I am sure that it guided him in directing the program. Oftentimes the proposals that you got in were the result of face-to-face conversations with somebody who said, "Here's some ideas that I have. What do you think?" You might respond by saying, "Well, these are not relevant to our program, but I think these are really great." And, "Oh, by the way, have you considered applying this in the following way?" There are various ways in which, just through face-to-face iterations, people can get a sense of your own vision and propose. So I think it probably started with them. Ivan Sutherland, of course, was the world's expert in computer graphics when he came in to DARPA. Ivan did a lot of other things while he was there, including getting the ILLIAC Project started--the program launched, or at least conceived of. You ought to ask him. I am sure he played a major role in enhancing the development of computer graphics technology in this country while he was there. I do not think it was ever of the form where he planned what the contractors would do, but I am sure he played a major role in critiquing what they were doing, and picking the nuggets. I think Bob Taylor really diversified the program fairly broadly. I mean, Bob had a very panoramic view. I am rather impressed by his vision, and of course he showed a lot of

that at PARC when he was out there. He is a strong-minded fellow, but I think that has led him to really be very successful.

ASPRAY: He was the first one to admit, though, that in some sense he did not have the technical background when he came into that position, especially following Sutherland and Licklider, and that he felt he had to rely more heavily on technical help within the office.

KAHN: Yes, that is true, but if you look at the history of things that he has gotten behind, he just had incredibly good taste and intuition. It is not clear to me that somebody in that position needs the technical wherewithal to be successful. I mean, if you have got the technical wherewithal you can rely on your own instincts more than you can if you do not. But there are enough good people in the community that if you do not have them, if you are a good manager, you can get the right inputs. Bob leveraged that to the hilt. I think he was marvelous at that. He had extremely good instincts. He could tell whether something was worth doing. He may not have been the guy to do it, but he could make it happen through other people; that is the art of management. In fact, one of the things he did at DARPA was to bring Larry Roberts in to make the network project happen. That was one of the first projects that was actually managed, or started to be managed, out of the office. They had had projects before, but they were always projects that were delegated out, like Project MAC was delegated to MIT. So it was more like a contract.

ASPRAY: Right. I guess that is what I was trying to get at.

KAHN: ILLIAC IV was a big effort, but it was largely a project that was contracted to Illinois working with Burroughs. The Burroughs part was a subcontract, but they were all managed by that combine. The ARPANET was rather different. It was a major networking piece that was competed and awarded to BB&N. And it was more in the line of the ILLIAC IV and the Multics experiments up at MIT. But to make the network work, you had to get the whole community working together, so there was a sort of a project role that the ARPA had to play, to encourage the different contractors to become involved. Of course, I played a lot of that synthesizing role while I was at BB&N before I came to DARPA. The ARPANET demonstration in Washington helped a lot to bring the community together. But I think that may have been one of the first examples of something that was more like a project at the DARPA office level.

ASPRAY: Can you think of other examples of that style through the years?

KAHN: Well, I certainly ran a lot of them like that, and I think there were many after that. Before I got there, there were a number that were project-like. They were not exactly the same, but for example, there was a speech understanding project that Larry Roberts had started. I forget who he hired to run it. Perhaps it was Cordell Green initially. You could ask Larry about that. He had gotten a number of people from the community to help him put together a kind of description of what the speech understanding project should be. Allen Newell actually wrote a plan or report, I think, at some point. Whether it was the initial one or the following one, I don't recall. But Larry got some help from the community in shaping and forming that. He contracted with five different contractors to work on elements of the speech understanding program.

TAPE 2/SIDE 1

KAHN: That was an effort where they gave multiple contracts, and there was a strong sense of competition among the contractors, like there would be a winner and a cut some number of years ahead. In fact, I think they actually reduced it from five to two or three after two or three years. So it wasn't the same sense of 15 contractors all working jointly to achieve an objective, but rather, the set of them working in competition with each other. They made a lot of progress. There was some controversy as to that method of running programs, versus the cooperative method, but there was an awful lot of progress too. It was stopped around 1976 for a variety of reasons. It's clear that the internetting activity was run as a project -- managed out of the office -- because we were trying to get all of the different sites to work together. The packet radio program was run as a project out of the office. It probably was more my style rather than others to formulate projects where I could leverage what I wanted to do through many people in the community. Much of the networking activity that we funded that was project-oriented ran like that. There was quite a bit of conceptual work that was done that didn't really fit into that mold. People like Len Kleinrock, who proposed some theoretical work and would just go off and do it. I think it was more my style perhaps than others to actually formulate the project and run it, where

the focus of activity was on what we were trying to achieve, and where it wasn't done by one broad RFP that got somebody in the community to go do it, because we couldn't even articulate what we wanted to do well enough to write it down. I think that style persisted. Around 1975 or 1976 we began to get involved with testbed programs with the military services. There was one testbed with the Navy called ACCAT; we had several with the Packet Radio Technology; one with the Army at Fort Bragg; one with the Air Force out at the Strategic Air Command. We put the technology in place and worked with the military services to see what they would do with it, what the impact was on their capabilities and operation. Those testbeds involved multiple contractors working together. It was clearly not a matter of unsolicited proposals coming to us. It was more a matter of us formulating a project, and proceeding to carry it out. In the 1980 time frame, there were many other projects that started out that were really grand projects, like the Strategic Computing Program, which I started in roughly the 1983 time frame. I think legally it was fiscal year 1984 in congressional parlance. That was both a program and a project, and it had major stated goals, a plan. It was an embodiment of a lot of what had happened before in a much more structured and formal way. There was a director for that project, currently Steve Squires as well as a series of people before him. So that was probably a style change that stuck. The thing that is required to sustain that, of course, are really good technical people that can formulate those programs and be able to manage and carry them out. Strategic computing is almost an institutionalized program in a way. There was an original plan put forth that laid out a broad set of goals and was so umbrella-like that people could propose to carry out various pieces of it. Yet it was fairly targeted and focused in some sense. So what was the goal of it? It was more focused than a program in information processing, in which we said, "Well, push the technology forward. " Well, that's one level of abstraction higher than the strategic computing goal, which was to take artificial intelligence technology and create an industrial base for it, or to create multiprocessor systems or to create multiprocessor technologies which could run artificial intelligence programs a thousand times faster, or create a speech understanding capability which would run on that processor which would allow a pilot to speak to his "pilot's associate." I mean, there were a lot of very focused goals that were laid out there that typically did not exist in many of the earlier kinds of programs. But once you have got a goal like that, then managing a program like that with a number of distinct pieces, can be done by people who are technically competent in the areas. They do not necessarily have to create each of the subpieces for the whole program from scratch. I think that tradition has largely continued to today. Whether they can continue to attract really good technical people at the pay levels that the government is willing to agree to is a question. I

think they are underpaying those people by a considerable amount. But I am sure the Congress is fully aware of this situation.

ASPRAY: Yes. When you first came to DARPA, what other technical people were there? Who were they?

KAHN: I am probably going to do an injustice to someone or another that I can't remember. There's was a fellow named John Perry, who was an Air Force officer. He was running a program in climate dynamics, as I recall.

ASPRAY: Out of the IPT Office?

KAHN: Right, this was a program that had to do with weather simulation on the ILLIAC IV. I am probably not going to remember what everybody was doing exactly, or even who was there. But John was one I recall. There was a fellow named Bruce Dolan, who was another Air Force officer. He had been managing the ARPANET project. I never did find out what happened to Bruce. There was a fellow named Steve Crocker. Steve was a graduate student at UCLA. He had been the guy who had organized the community to work on network protocols -- the versions of protocols that were the original ones on the ARPANET, such as the NCP, that the TCP/IP protocol suite later replaced in 1983. He had been in the office. He later left to go back and get a Ph.D. at UCLA. There may have been one other program manager there. It beats me if I can remember though. Larry Roberts was running the office when I showed up. There was a fellow named Al Blue who was an administrative assistant. There were a number of secretaries.

ASPRAY: Okay. It had already grown by that time, compared to what it had been five years earlier in terms of the number of technical people. That's the impression that Taylor had given.

KAHN: That is probably right. I was never in a position to tell how technical John Perry and Bruce Dolan were. They were both Air Force officers. They were clearly working in technical areas, but whether they were just doing management, or actually contributing technically, I could not tell. Steve Crocker was clearly on a technical track. In fact, after he got his Ph.D. he was clearly very vigorously involved with the research community. But at the time he was still a

graduate student, so he was not a fully trained and seasoned technical performer. Of course Larry Roberts was. I don't know when he got his Ph.D., probably in the early 1960s at MIT. I think I was the only other Ph.D. No, John Perry might have had a Ph.D. But it was probably not in electrical engineering -- probably in something else. Even in the early days, when Bob Taylor was there, he had Larry on board as a staff member. And when Larry was there, he had a guy named Barry Wessler working for him. While Barry was a graduate student, he was on a technical track. So there was always some technical expertise around. I just think there was more of it then. In every stage of the game the people who were brought in on the staff were always selected for technical attributes at some level. That is, somebody who was not technical would never have been selected, no matter what, unless it was just for an administrative position, of which there were maybe one or two at any point in time -- a program analyst, or a budgeteer, or something like that. But the real work was done by the technical staff. There was not a uniformity of technical capability. But they were all technical to some extent, or they could not have managed the programs. Now the office budget level, at that time, I would guess was probably high twenty millions when I came in. When Licklider came in, I would guess the budget had gone up to about 40 million dollars.

ASPRAY: The second Licklider trip.

KAHN: The second Licklider trip. When Ivan was there, maybe it was 10 or 20 million -- grown from 10 or 20 million in the mid 1960s to maybe 40 million by 1974. Then it actually dropped down for a while. It was the Vietnam syndrome, and a variety of factors that cause the budget to be down. During Licklider's second tour I do not think he and George Heilmeier saw eye to eye on many things. Anyway, the combination of that caused the budget to go down. By the time Licklider left the budget probably dropped down to about 34 million, or something like that, and then eventually got back up to the 40 million level or so when I took over. The budget grew fairly quickly after that. When I left, the budget for information processing was probably around a quarter of billion dollars, much of it being in the new strategic computing program.

ASPRAY: What did the program look like when you first arrived? What were the major areas that were under investigation?

KAHN: I can not remember them all, but there was a program in speech understanding. There was a program in artificial intelligence research. It was fairly small -- three or four million dollars. There was a program called Automatic Programming. It was largely trying to deal with being able to have machines create programs from very high level descriptions. It didn't last very long. There was another program which we got partitioned in different ways. It included the Network Program known as ARPANET. There was the ILLIAC IV program, which still had some visibility. There was the Climate Dynamics Program, which was around for a while. There were a number of other, smaller efforts that were started, which I would somehow lump under the rubric of command-and-control initiatives. In fact, that was one of the areas that grew most significantly after I got there. The testbed efforts were put under that program. A lot of the new networking developments were put under that program. In fact, it later included things like packet radio, packet satellite, internetting -- things like that. That's my recollection of what the program was like at that time. Oh, there may have been a program on computer graphics. Oh, no, there was one more--image processing. It was not quite image understanding, which started up a few years later, but it had to do with the ability to use computers to process images to enhance them, deblur them, and shape and shadow them -- things like that.

ASPRAY: These are clearly lots of programs spread across a wide range of technical fields. An IPTO director cannot possibly have detailed knowledge of all of these areas. He can be expert in a few perhaps, but not in all of them. How was the technical management handled? Was it a division of labor, clearly, or did this have some bearing upon the choice of the program, or the choice of the person as a successor?

KAHN: To start off with, I think that it is clearly better when you have expertise right at the top. A good manager -- by good, I mean somebody who can decide who to listen to, and know that he needs some advice and input -- can function very effectively running that office. I think Bob Taylor demonstrated that. So it is not clear that an office director necessarily needs to be expert in anything technical. What he needs to have is good judgement and the ability to assimilate advice in those areas. That is the minimum. Now, more often than not, what you really need is good judgement, intuition, in deciding how to invest, provided that you can get some insight as to what makes more or less technical sense. Remember, you are dealing with inputs from the technical experts in the community. So you start off

with things that are usually on a pretty good track. The question is sorting out the few that may be slightly off track, or the ones that do not really align with your vision of where you would like to see things go.

ASPRAY: Yes. But that vision is partly one that is shaped by the technical...

KAHN: It could be, but it does not have to be. It depends on the level of definition of your goals and objectives. I mean, if your goal and objective is, let's say, to get people to be able to write programs easily for a machine, it's easy to tell whether you're there without knowing anything about the underlying technology. You can visualize whether it is there or not. If you have a good technical view, you may have some notion of exactly how you want to go about doing that, and you can guide it down that path; but if you do not, then you can rely on the inputs from the community. In some cases, you might actually get better inputs than you would have generated yourself. It depends on how good you are relative to the people in the community. There is a disadvantage of having too good a technical person as the head of an activity like that, because he may have some very fixed ideas about how to proceed, which are reasonable, and get so committed to them that he may become less open to different, even possibly better ideas that come from elsewhere. But in general, I would say that the office has been very fortunate in having people who have had both good judgement as well as good ideas.

Oh, one more point here having to do with the hiring of people. Generally, if you were going to run a program, let's say in image processing; when the office was small, the office director did a lot of that. But the actual work of moving a proposal through the system can entail a lot: you have got to work the proposal; you have got to read it; you have got to understand it; you have got to write the paperwork; you have got to staff it through the system, so you need somebody who is at least knowledgeable about the area. And generally, when staff is hired, you hire them based on the knowledge of what the program structure is likely to look like. If you are starting a new program in an area that needs somebody who is familiar with that area, you go look for somebody with those talents. For example, if you knew you were going to have a major program in software engineering, you would look for somebody with a strong background in software. If you knew you were doing something in image processing, you would look for somebody with a background in signal processing.

ASPRAY: So during your period as IPTO director, who and how many of these people did you hire? How many would be on at any one time?

KAHN: I don't remember how many I hired, but generally the staff was something in the order of eight to ten at any point in time, including myself. We had a number of other secretaries. So maybe the total office was 14 to 16 people.

ASPRAY: Eight to ten people of a... I don't mean to be pejorative to the others by saying "professional", but people who could carefully watch over a program.

KAHN: Well, when I was there, I would say that the staff that we had was extremely professional. I mean, they were people who were just unbelievably good at what they did. We had people like Vint Cerf. Vint, when he left, became a vice president for engineering at MCI. He is just one of the most gifted people around that I know. Duane Adams -- I don't know what his official title is now, but until they made the computer science department at Carnegie a separate school he was the associate head of the department for computer science. He has got the title of Associate Dean for Research now, or something like that. He is also very good. Larry Druffel is now the Director of the Software Engineering Institute at Carnegie Mellon University. Jack Hammett -- one of the very few Army Ph.D.s around. Bob Englemore, who became director of the Knowledge Systems Lab out at Stanford. Paul Loslaven, who played a very key role in running the Center for Integrated Systems out at Stanford. He was really a key technical researcher in VLSI. Barry Leiner, who is one of the people running the Research Institute for Advanced Computer Science out at NASA Ames, or the conjunction with NASA Ames. An awful lot of extremely good people. I mean, I have not given you a complete list of everybody that was there. These are all people who would be candidates to run the office. Ron Ohlander, who's now at USC -- was Acting Director there for a while -- did some pioneering Ph.D. work in image understanding, the first work in segmentation. A lot of very good people.

ASPRAY: Again, I know I am repeating myself -- if you take this list of people you just reviewed, how many of them would be in the office at any one given time? Four or five of them?

KAHN: I would have to go back through the list, but I would guess you have probably got seven or eight of those there at one point in time.

ASPRAY: Okay. Can you tell me the story about the move to your taking over as director. I know there is a period when Licklider comes in, it seems to me, to help out. And there is some drifting going on, it looks like to me as an outsider, but I do not really know this story at all. Can you tell me?

KAHN: Everybody sees these events from their own perspective. My guess is you would probably hear a dozen different points of view if you asked a dozen different people. Let's see, when Larry left, the director of the agency was Steven Lukasik. The deputy was Alex Tacmindji. Licklider came roughly in January of 1974. Larry left to take over Telenet at the end of September of 1973. They had a lot of trouble agreeing on who would replace Larry at that point. I do not know why that was, but there was just a lot of difficulty finding agreement on who a successor should be.

ASPRAY: Who would be responsible for making that decision?

KAHN: Well, it was a combination of people who were involved. Ultimately, it's the director of the agency, but he relies a lot on the current head of the office. I mean, Larry was such a strong leader that... [TAPE OFF] Okay, so there was some difficulty identifying who would take over. Partly, I think that was because Larry was such a strong director. It was not that easy to find someone who wanted to come in to replace him. He had been building some momentum in his own mind to want to leave to go take over Telenet, but then he had agreed to do that back in the spring, so he had been six months on the job planning to leave. I can't remember the exact timing, but I think it may very well could be that Licklider agreed to come take the job in September. As soon as Larry had learned that Licklider had agreed to do it, then he left. It might have been something like that. In any event, Licklider could not come on board until January. Larry was anxious to leave, so he left in September -- the end of September. Al Blue, who was the administrative assistant, essentially became the placeholder, acting director until Lick showed up. During that period, the office was essentially run by program managers with Al Blue knitting it together and having the signature authority for everything going on.

ASPRAY: Did it work okay for that period of time?

KAHN: I thought it did. Of course, it was like anything else. When you have got a director on board who is really calling the shots, then you get some direction from the top. Al was not trying to provide direction, because he was not a technical person. He was just trying to manage administratively. He did a very good job. When Lick came in, I think there were two issues there. One was that since he had created the whole thing in the first place, he still had a notion of what the office was like when he had originally started it, which was probably quite a bit different from the way it had evolved over the course of a decade or more.

ASPRAY: Can you identify some of the changes?

KAHN: No, I cannot, because I did not know what it was like originally. But I know that much of what Licklider had tried to do did not work as easily the second time as it would have worked for him undoubtedly the first time. For example, he tried to set up a program, I think it was called Management Information Science, or something like that -- a program to worry about managing information. This was at a time when the Congress was becoming increasingly adamant that defense work had to be defense relevant, and the work that he was proposing was fairly general, which was the theme of the office during its history. So now, it did not fly, and I am not sure I remember exactly why it did not fly, but it ended up getting recast as a command and control program, which was a perfectly good umbrella to do the same work, except that instead of sounding like it was management information it sounded like it was a real defense problem.

ASPRAY: Okay.

KAHN: Lukasik left in the end of 1974, and Alex Tacmindji took over briefly until George Heilmeier was named. George came in January of 1975. I don't think he and Licklider hit it off very well. Lick left in August of 1975. He was only there about 18 months. He might have only been there 18 months the first time. Maybe that is the length, in his view, of an ARPA director's stay. Dave Russell, who was an Army colonel, had gotten selected by George Heilmeier as head of the IPT office. I think George had talked to several other people about the possibility of taking that job. Two or three of them had declined to do it. He just selected somebody that he was comfortable with running the office. Dave

Russell was a very fine manager, and held the office together in a very effective way under George's tenure until Dave left in August of 1979, roughly. I took over the office at that time.

ASPRAY: Where is he now?

KAHN: Dave left to go to work for Satellite Business Systems, which was a joint venture at the time of Comsat, IBM, and perhaps AETNA. I think Comsat pulled out. Let's see, and then I think MCI bought it through some deal. Now he is officially an MCI employee. He is part of the old SBS operation.

ASPRAY: Located where, do you know?

KAHN: It is nearby, at Tyson's Corner. I could probably find you his telephone number, if you would like.

ASPRAY: Yes, I would like that. Were you, as far as you could tell, the obvious choice to take over the program?

KAHN: You mean when he left?

ASPRAY: When he left.

KAHN: I cannot comment on that. You would have to ask Bob Fossum. I think that George Heilmeier left in late 1977. Bob Fossum came right after that.

TAPE 2/SIDE 2

KAHN: My understanding was that Fossum had sent letters around to the military, academia and industry, asking for recommendations and out of that process I got asked to take the job. You would have to ask Fossum. I do not know whether I was the obvious candidate, or whatever. But Bob could probably tell you.

ASPRAY: When Russell was selected, was it a conscious decision to hire somebody from the military? I mean, nobody else had had that kind of background that had been director.

KAHN: George Heilmeier and I were both graduate students together at Princeton. So he and I knew each other going back into the early 1960s. We were good friends back then. I think he had some very strong feeling about how he wanted to see the office evolve. The way things ran, he pretty much left me alone to go run the programs that I was running. Oh, we arm wrestled from time to time over some technical points, but I always found the interactions with George to be straightforward, above board -- good, solid technical interchanges. I enjoyed them, actually. But I think he had some views about what he wanted out of the AI program that required him to take hands-on control of it. So, I do not think he would he would have felt very comfortable with anybody else running it that he felt did not share a common view with him of what he wanted to do. He and I never talked about that. So I do not know whether I had a common view with him or not. But I know that he was looking for some more specificity on what he wanted from the AI program than the people in the office were able to provide, or that the community was able to provide. He was looking for some kind of a roadmap. I think he found Dave Russell to be an ideal person for him to be able to work with, because he could essentially instruct Dave as to what he wanted to do. I think he and I were more technical peers at some level. I mean, this was more my field that it was his. He came out of physical sciences. I have a hunch that he felt it might not have worked as effectively. I do not know if I was ever a candidate. So that is sheer speculation, but in terms of supporting me and the programs that I ran George was terrific during that period.

ASPRAY: What was the attraction to you to take the directorship when it was offered?

KAHN: Well actually, at the time that I agreed to take it, I had only gone to DARPA with thoughts of staying there a few years and then leaving. I did not intend to make it a career. In fact, I am surprised, even at this stage, that I stayed 13 years. Most people do not stay that long. It is just that at every decision point along the way I had gotten my arm twisted. In fact, as I think I mentioned earlier, I was rather disillusioned when I first showed up, because I thought I was going to run a program in manufacturing. When that did not happen, I thought very long and hard about just going back

to something else at that point. But I decided to stay around, because I found it interesting. I thought that it was a place where one could have some impact and leverage, and it was an opportunity to learn about the government. At every point along the way, where I had said, "Well, I think now is the time to leave," I got my arm very strongly twisted to stay on. At the time I was offered the director's job, I told Bob Fossum, who was the director, in no uncertain terms that I thought that he should find somebody who had not been in DARPA for as long as I had, because I had been there at that point six or seven years, and I thought that was too long. He said, no. It's not relevant in my case. I shouldn't even think twice about it. He really wanted me to take the job. He pressed me very strongly to take it. I basically told him that I thought that the computing program at DARPA had been severely eroded over the past several years, and that the only way that it would be interesting for me to take it on is if he would work with me to grow the program at least back up to where it was before. He agreed to do that; which he did over the course of the next several years. It grew back from the 40-50 million range to close to 100 million by the time he had left.

ASPRAY: In its downslide, what had suffered? What programs had been cut back or cut out?

KAHN: It is hard to remember exactly. I know that the speech understanding program had been cut out in its entirety. But that had been replaced almost in its entirety by this Navy command and control testbed program. The ARPANET program had been phasing out, so that was just a natural erosion there. But on the other hand, the command and control program, which included that, was growing to replace it. I think the ARPANET was actually out of the command and control program. So while that was decreasing this other one was increasing to encompass it.

ASPRAY: The weather program?

KAHN: The weather program went away. The climate dynamics program was transferred to NSF, as I recall. I do not know that I can tell you exactly programmatically what decreased and what did not.

ASPRAY: I could probably find that out by putting the records together, although it's not clear since the records aren't very good.

KAHN: Let's see. The image processing program probably went away. It later got replaced by the image understanding program. So I think one was shrinking; the other was growing. There may have been a period where one had not grown enough to compensate for the other yet. The ILLIAC IV project was winding down, as I recall. Various other projects along the way had gotten canceled here and there. There was the program that Doug Engelbart had been running at SRI, the NLS work. Doug had been one of the real visionaries back in the 1960s. He still is visionary to this day, but his work just sometimes did not pass muster in DARPA at the time. So there were a number of efforts like that that were cut back. [TAPE STOPS] It's very hard to say exactly what was cut back. I mean, I remember little things here and there, but I think that while there was an awful lot of noise made about the AI program and its lack of a roadmap -- specifically articulated plans and goals, which was very hard to do at that time; it would have been nice if we could have done that -- the fact was that the AI program in its totality actually grew during that period of time. I think that is not well understood within the community. I don't recall whether the basic research part of it went up or down. The basic part of it might very well have stayed flat during that period, but there was certainly a very substantial growth...

ASPRAY: In 6.2?

KAHN: ... in the applications of AI technology. In fact, that was the very first forcing function that was ever applied to that community to produce something that would be useful. So I think that was a turning point in the whole evolution of artificial intelligence technology that came about because of George Heilmeier.

ASPRAY: Did you have a particular vision of areas that you wanted to start off or expand when you took over that you wanted this extra money for?

KAHN: Well, I came in with the following perspectives. One, at the minimum, we needed to get the basic research funding back to the level of where it was. Then, if we could do that I felt that taking the office job would have been worth it. That was essentially an agreement that I could have made with the director at that time and he agreed to that -- and worked with me over time. Bob Fossum and had a very good working relationship. But I also knew that just

working with the research community was not likely to make this technology available to defense unless we could get industry involved. So part of my thinking at that time was to find a way to somehow increase the program to get industry involved, because industry was much more expensive to fund than the universities. That was my long term goal going in.

One, let's get the basic research level restored, then let me go work on the problem of getting industry involved. If you do not have the basic research, then there is nothing to transfer; it is hard to get anything to industry unless you just want them to build big projects for you.

ASPRAY: Can you remind me what role that industry had had in IPTO work prior to this?

KAHN: Well, there had been a very significant role of industry in DARPA as a whole. If you went into the offices you would probably find the university component a very small fraction -- maybe only a few percent -- with the exception of IPTO and the materials office. The materials office budget probably had 25% funding of the universities. So they had a nontrivial basic research component. Again, I do not know the exact number. It might be 30%; it might be 20%, but it was somewhere in that range. The university support in IPTO at one point had been very large -- well over 50%. It was predominantly in the universities. In the mid-1970s it had started to switch as we got more testbed programs and more command and control kind of work. The industrial component had grown. The packet radio program, for example, was largely an industrially funded effort, although we were running it out of IPT itself. But still, I would say that the funding in the universities had gone down to well below half of what it had been a few years before, and maybe as little as a third, or even a quarter -- really, the university component was very severely cut back.

ASPRAY: Would you say that a significant reason was the social climate at the time; the Vietnam war, the Mansfield Amendment, and attitudes towards the Defense Department on the campuses? Did you have trouble with contractors for example?

KAHN: I am sure that was a part of it, but I do not know how big a part that was. I mean, it is nice to describe that in macroscopic terms as being the result of that, but I think that that would have more affected the overall shape and form of the DARPA budget. In fact, what happened was George Heilmeyer got very significant growth into the DARPA budget

starting around 1976 or 1977 through a program that he created called Experimental Evaluation of Major Innovative Technology -- EEMIT was some kind of demonstration to experimentally evaluate technologies that had been created, rather than just create them and put them on the shelf. That caused the DARPA budget to go up from the several hundred million dollar level to the six, seven hundred million dollar level within a few years -- a very steep ramp up in the 1970s. I think a lot of that continued into the 1980s, with the Reagan buildup. The rest of it was probably just personalities, the nature of the program, decisions of the ARPA management of where they wanted to put money. You know, much of it was just due to the people -- their decisions and interactions right within ARPA -- as with anything. By the time that I took over there was a significant industrial component, but it was not overtly involved in transferring any technology to industry. They were contractors that were there to run a testbed program for you. There was still a significant amount of work on networking at BB&N, or BBN was creating the technology work, but there was no overt linkage between a research effort that took place in a university and industry, which would take those ideas and go pursue them. Transfer was all happening rather by the invisible hand of the marketplace, or venture capital, or something. A guy would leave the university and go join a company. A company would then spin off. But DARPA was not taking any role in making that happen. What I felt was that there was a need for it to be a little more directive, because many of the programs which were creating new technology were not getting into the hands of organizations that could help meet the needs of needs of Defense. On the other hand, I did not think that I would be personally involved in that second wave. I thought that if I could just get the basic research level of funding back, maybe a successor could deal with that. Then when Bob Cooper took over as the head of DARPA -- I think it was in the summer of 1981 -- he was looking for major new things to do. We had already restored the basic level of computer science. He was asking, "What is it that we ought to do?" I was saying, "I think we ought to make a major push to do just this kind of thing. Let's get industry more involved, get universities cranked up a little more to work on transfer as well as the basic research. It's an opportunity to get industry involved." I laid out a strategy, and he made it his top priority, probably to the detriment of a lot of the other office programs, because the Congress didn't give us more money to do the program. Socially, you know... The budget profiles that we were working with allowed us to do this new program. We got all the new money, which did not make a lot of the other office directors happy. Bob Cooper attempted to knit the program together so it was a composite program which involved multiple offices all working together with different aspects of it. He set up the major management strategy. The information processing budget essentially went from around a hundred million a year to around 250 million

a year. Actually, you have to understand how you deal with budgets. In fiscal 1984, let's say, you're spending 1985 money; you're planning to spend 1986 money, and you're figuring out how you're going to budget 1987-88 money. So you're kind of handling multiple budgets in any point in time. But they're always in front of you. So whether you're actually spending the money or not are two separate issues here. Then we also had some interesting discussions, disagreements at times, about exactly how to manage the basic research part versus the applications part of the programming. People involved in applications were much more comfortable with timelines, and milestones, and details of exactly what you are going to do when. But a lot of those applications were based on technology that was coming out of the basic research program, for which we could not give such detailed timelines, so we could not see how to timeline the applications. So what happened was that we took the application programs and essentially separated them from the basic research so they could be managed with two different styles of management, because I wasn't that comfortable trying to manage the applications in the way that they were conventionally used to being managed, because there was too much basic research and series with it. It is one thing if you know you are going to buy this micro from Intel and plug it in here, and buy that memory chip and then hook them together. But it is another thing when you are going to create something here, and you are not quite sure what it is, or when it is going to be done, or what it is going to look like, or how it interfaces to anything. In fact, the applications got changed by virtue of this process, because it forced them to look at what they could get their hands on now, versus what might show up in the future.

ASPRAY: What part of your budget was devoted to applications?

KAHN: If I had to guess I would have said that in the out years it was about 30% maybe. In the initial years it was a little less. It might have been only 15% the first year or two. But anyway, that money was moved and managed separately. It was later brought back again into IPTO and they renamed the office ISTO. In between, in a year or two, they created a new office called the Engineering Applications Office where most of the strategic computing applications went, plus a few other things. But then they had two computer offices in DARPA, which was not an ideal situation. After I left, then, they essentially merged the two back together again and called it the Information, Science and Technology Office. I think, in fact, a lot of the applications... Some of them went away. Congress essentially took away

some of them. I think the Autonomous Vehicle Program stopped. You would have to ask some of the people there right now to find out.

ASPRAY: Can you recommend people to talk to?

KAHN: I would talk to Saul Amarel who is now at Rutgers. He was the office director after me. I'd talk to Clint Kelly, who ran the Engineering Applications Office. He is over here at Tyson's Corner and works for SAIC. You can talk to Craig Fields, who is the Deputy Director at DARPA, and who was Clint's deputy at the time. Let's see, I do not know how much Jack Schwartz would have known about that, but Steve Squires at DARPA might. You might ask Steve.

The other factor in the decrease was the Gramm-Rudman cuts which occurred in the 1985-86 time frame, which put some very severe constraints on DARPA. Part of the cutback was the fact that the money was split apart. I think there was something like 200 million left in IPTO, and there was probably 50 million in the EAO office. Then when Saul Amarel came in he encountered a number of budget cuts right off the bat. I think quite a bit of that was just due to the politics of the situation -- who got on the chopping block. It's hard to know exactly how to protect every penny. I think the big driver in this whole reduction, more than anything else, was Gramm-Rudman. The problem there was that the Gramm-Rudman reductions were lumped by categories. So the Army got hit, the Navy got hit, the Air Force got hit, and the defense agencies got hit. The problem was that the defense agencies included the SDI program, which was the biggest component of the defense agency's budget, I believe, and the president fenced that. He protected it. He said, if the agencies have to get hit -- by whatever the number was -- by 6%, you are not going to be able to take anything out of more than half of the program, and everything else is going to get hit by a much bigger percentage. I think the DARPA budget got hit by 15 or 20% -- some really big number -- that year. You could probably find out what the exact number was. The net effect of all these cuts was that the budget went all the way down from around 200 million to maybe 130, some number like that, when all was said and done. Now, I don't think that was all Gramm-Rudman. There may have been other competing demands. You have now got new management; you have got new ways of defending programs, new ways of articulating things. Management decisions change over time. I think the budget really did come down. It grew back up again in recent years. The last I had seen it was in the 170 million range, but that was many months ago. I

don't even remember what fiscal year that was, because I don't keep that data in my head any more. I do not know what it is right now.

ASPRAY: A technical question here -- technical in a sense: when we look through the budgets we see allocations for 6.1 funds and 6.2 funds. Does that realistically mirror the funds that were being allocated by the program office director to basic research versus engineering. Can we take those as a realistic basic for making...?

KAHN: No. In many cases the distinctions were irrelevant. I will give you an example. I was at a meeting in Seattle that the Computer Professional for Social Responsibility were holding. This was a group that had largely been very critical of Defense involvement in computing, and Star Wars, in particular. They were citing major changes that were happening, saying things along the lines that much more money is going into applications than into basic research, and it is terrible. I was saying, "I do not know that that is exactly the right conclusion to draw. It may be that there's a trend in that direction, but I do not think it's nearly so great. It is not like 25% is now in basic research, where 75% was before. I mean, if it is anything, it was 67% before and 66% now, or some small change." They gave me examples of why that was wrong, and showed me how the 6.1 number was this and now it dropped by this, and this went up. I said, "Well, let me give you an example," and I described the following situations: before we started Strategic Computing, we had something like 21 or 22 million in our AI program that we presented to the Congress. It was a 6.1 program line item. But when we started Strategic Computing, Congress knew that AI was a major component of that program. And that was a 6.2 program, because that's where the extra dollar flexibility was at DARPA. So we had a 6.2 program for Strategic Computing. And the Congress said, "We don't want two AI programs in two different places. Put all the AI stuff in one place." So in one year they made an accounting change that literally changed these dollars from 6.1 dollars to 6.2 dollars, but nothing else changed. It was the same contracts, it was the same line items. The only thing on the books that changed was the number that was associated with it for accounting purposes. So, you know, here was a case where the 6.1 funding dropped by 22 million. The 6.2 funding jumped by 22 million, and if you looked at the numbers you would have said, "Gee, the 6.1 funding was only..." Let's say the 6.1 funding was 50 million, and the 6.2 funding was 100 million (let's just pick two numbers). Well, so if the 50 million dropped down to 30 and the 100 jumped to 120, you would say one went up by 20% but you would say the other one went down by 40%. The fact is that there was no change at all. So I

think it's very hard to go by those numbers. You have to understand what is in the programs. The intent of 6.1 funding was to fund the more basic work. I think that as a general rule that is exactly what DARPA did with that money. I do not know of any cases where 6.1 money was actually spent on anything other than what it was intended for. 6.2 money was intended for exploratory development, but where is the boundary line between what is basic and what is exploratory? I think that in terms of computer science work it is a judgement call. It often looks much like a continuum. So the work that was really clearly hardcore exploratory work was in 6.2. If you were building a piece of equipment, it was almost always in 6.2. A lot of the work that was near the boundary could have been on either side with equally good justification, and often was. I guess the bottom line is that if you really wanted to figure out what the basic research program was, I would have said, "You should start with the 6.1 funding, and then ask, how much of the 6.2 is like the 6.1." So the basic research funding, or the stuff that was more basic than less basic was always lower bounded by the 6.1 funding. It could have been twice as much.

ASPRAY: Okay. Can you tell me about your program and how it changed over the years that you were office director?

KAHN: Well, let's see. There were really two major initiatives that I pushed as office director. There were many other smaller ones, but I remember, my two goals were, one, to restore the basic research funding and second of all, to see if we couldn't get some growth beyond that so I could somehow couple that to industrial involvement. The strategy in the second phase was the Strategic Computing Program. That was the mechanism. In the first phase, the justification, which was really a very strong one, was that the real future of the computer science community as a whole really depended on getting direct hands-on access to integrated circuit technology. Because, if you remember, there was a lot of work going on at the universities before, but it was all at the level of discreet modules -- and-gates, or-gates, and the like. Well clearly, you wouldn't ever build a computer today out of and-gates and or-gates. On the other hand, all of the efforts in integrated circuits back then were largely captive to the Intels, the Mostecs, the TIs, the Nationals, AT&T, IBM, etc. You really had to be inside those organizations to be able to get your hands on them. A key guy in a university could cut a special deal with somebody. Or if a guy left TI and went to work at Berkeley he could somehow cut a deal with his former employer and they would say, "Okay, you can do this, but you must sign this non-disclosure agreement. Maybe we will fab a chip for you in a year if you are really doing interesting stuff." There was no way that the community as a

whole could get their hands on it. The focus that I picked for the growth of the program was the VLSI technology, because I felt that that was very broadly based. It would apply to virtually the whole university community, and it was really a very good justification for increasing the program. If you go to the Congress and you just say, "You are not funding us enough. Why don't you give us twice as much money," it can't work politically. On the other hand, the community really did need the funds, and this was a good area to put it in -- an area that was not being addressed at all. That was where the growth occurred in the program for the most part. We created a VLSI design program, which caused the universities to learn how to do VLSI design, caused them to do design. We built a system called MOSIS for VLSI fast turnaround fabrication. A number of the key technical innovations were really developed by Carver Mead and some of his students down at Cal Tech. Then, working with Lynn Conway, they wrote and propagated a textbook on this subject which essentially told people how to do VLSI design in a way that was vendor independent, so you didn't have to sign proprietary agreements with the vendor to do VLSI design. That was the situation before: there was no way you could design for somebody's line without signing a nondisclosure agreement in order to learn about how to do design for their line.

TAPE 3/SIDE 1

KAHN: They came up with a methodology that was essentially drawn from looking at a number of different lines and understanding how to do the design that would likely work on all of them. That was the methodology we adopted. DARPA gave out a lot of contracts to different research groups around the country to explore some of their ideas. It was all idea-driven at the time. There were a lot of very good ideas that got explored -- the systolic array ideas at Carnegie, the scheme chip at MIT, the geometry engine and MIPS at Stanford, the RISC chip at Berkeley. Many interesting ideas like that were explored. I think it made a major difference to the computer community, because it gave them the opportunity to consider VLSI as an integral capability in their toolkit. So I think that was a major accomplishment. Hewlett-Packard was very helpful in making the first fab-line available for us to use in this process. The semiconductor industry was not initially very convinced this was a good thing for them because they had a captive hold on the market. But you know, the computer industry did not think that timesharing was necessarily good when we started either, because they were used to selling dedicated machines. Well, if you sold a machine that multiple people could use simultaneously

it might cut back on the sale of mini-computers, or whatever. I am sure packet switching probably raised similar concerns. I know packet satellite technology did. Instead of selling 'n' satellite channels they could only sell one, and it wasn't as good as selling 'n' in terms of revenue. But as soon as Hewlett-Packard agreed to provide fabrication services a number of the other companies agreed to provide them as well. Xerox PARC -- Conway's group at Xerox PARC, essentially collected all the designs and put together the tapes to make the masks, interfaced with Hewlett-Packard, and got the chips back. They ran that process a few times, and then I transferred that whole operation -- or rather I arranged that a similar thing be created at the University of Southern California, which they are still running today. It was not called MOSIS at PARC. It was just an activity over there. We commissioned a MOSIS capability at USC and it has been running there ever since. It has been very supportive. It is a piece of infrastructure that is necessary to make VLSI design easy almost everywhere else. So that was the strategy, and those were the two main efforts. We continued, you know, to...

ASPRAY: I'm slightly lost. One was...

KAHN: The first was this VLSI program -- had to do with VLSI design, architecture, and also the infrastructure to support it, and the funding of specific efforts along the way. Different, interesting ideas about VLSI. The second was a Strategic Computing Program, which we started in 1983 really.

ASPRAY: What programs were phased out in the new period?

KAHN: The Packet Radio program, which I had started, phased out during that period. Barry Leiner started a follow-on program on survivable radio nets. He was looking at the C³ aspects of survivability. Let's see. A number of the command and control programs -- I am sure that the ACCAT program phased out. This is the Navy command and control program that phased out during that period. We had started a program to develop some centers of excellence in computer science, and that program was phased out because of a decision to not work on centers but to work on specific ideas. I am trying to remember the programs that we were supporting. One has a tendency to remember the things he started up, the things that... Let's see, we had a military message experiment program out at SINCPAC. That ended

sometime just after I took over. I had started a program on distributed sensor nets, and we were just phasing out that program as I left. It either ended just before I left or just after, but it was within its last year. We started a program that was an exploratory program to look at the possibility of a satellite system consisting of multiple satellites. That was a defined two or three year program. That one ended, but then there was a follow-on program to that to create some experimental satellites. I can not remember the name of the follow-on, but I was not involved with that one.

ASPRAY: Has there been less of a communications emphasis under more recent office directors than under yours?

KAHN: I think that that is probably true. Because I had previously spent a lot of time working on networking (that's what I had been pushing for in the period that I was managing programs) I would not unfairly push networking as an office director. It would be too close to my own personal interests. Also, I felt that the level of investment was about right at that time. We did enter into some joint agreements with the National Science Foundation to work with them on networking, but I really put my emphasis on other areas that I had gotten to be involved in: some in multiprocessor systems, VLSI technology, parallel programming, artificial intelligence; those were really the main thrusts.

ASPRAY: Was there a graphics program at the time?

KAHN: No. I mean, there was some work in the VLSI program on chips for graphics. The geometry engine at Stanford was like that. The pixel planes program at North Carolina was like that. But those were not aimed at graphics per se, but innovative things you could do with VLSI. So the networking program did not go down, but it never really got the sustained growth that we got in the other areas. When I left, I think that DARPA had made a decision to cut back on the networking work, because I had always defended it fairly strongly while I was there, even though as office director I had not tried to grow it very significantly. It might have grown a little, but if it did, it grew at a much slower rate than all the other things did. When I had left, I think that the decision was probably that you've been working on networking long enough and now that Bob Kahn is gone, why don't we just cut back and take a fresh look. I think they actually cut back on the networking support for a while. But there has been a resurgence of interest in the government -- this whole issue of the National Research and Education Net, Fiber Optic Networking, and the like. There has been a resurgence of

interest within DARPA to get back into networking. I do not know to what extent that is a major change, or to what extent it is just a small reassertion of interest. We will have to see. My guess is that networking is just so fundamental to what that office does at DARPA that they can not ever really get out of it. The question is, how deeply do they want to be in it. That is a judgement call on the part of the management -- program managers, up to the office directors, up to the head of the agency.

ASPRAY: If we can turn for a few minutes to more strictly management questions, I am not sure I have the vocabulary in hand to deal with this all that well. But let me start at a place that I know I can get into it, and maybe you can guide the direction. Can you tell me about the budget process, both within the office itself, in terms of allocating money across the different projects, but also more specifically in competition with the other offices and the overall planning for the DARPA projects? What kind of process went on? What kinds of discussions took place?

KAHN: Let me draw you a picture. I do not know if you can hear this.

ASPRAY: I will move over a little closer.

KAHN: I will give you a picture of this to take with you. You have to remember this changes from year to year. So let's pick some numbers in here. Let's say this is fiscal year 1980, one, two, three, four quarters. This is fiscal year 1981; this is fiscal year 1982. And so forth. The fiscal year started. This would have been October 1 of 1979. So let's follow what tends to happen here. Let's see if I can remember all of this. Roughly in this time frame -- I am not going to give you exact dates, because this is all rough anyway. You get something called on guidance. This would be on guidance for the budget that was going to be submitted by the president here.

ASPRAY: At the beginning of fiscal 1981.

KAHN: At the beginning, yes. The budget that would go into Congress is roughly here -- January of 1981. That budget would be the fiscal 1982 budget.

ASPRAY: All right.

KAHN: This is the guidance for 1982. Remember all these numbers are fiscal numbers. This is basically guidance that says, roughly how much we think you ought to plan for -- guidance that comes back, I guess from OMB to the departments and agencies, and they distribute them. They take that guidance and each of the agencies then decides what to do with it. Say that guidance is, "Plan to go in at 200 million or 400 million this year?" The DARPA director would take that guidance (let's say it is more than 400 million) and he might say, "Okay, fine, I am going to hold back 50 million," or "I'll hold back 70 million for the moment." It depends on the strategy. Another director might say, "I will not hold back anything and I will just assume that there's ten percent in any office that I can count on." But let's say you hold that back. That means 350 million and let's say he has got 6 offices. He might apportion them and say to these offices, "You can have 70; you can have 110; you can have 40; you can have 15; you can have 130." So let's say he apportioned it to the 6 offices something like that. That is based on what they had last year. So let's say this office's budget was 400 the previous year. You might say, "Well, that's a small growth. So since these guys might have had 65 the preceding year, we will give them 70 this year. These guys might have had 100 the preceding year; we will drop them to 90, because we got rid of some things. These guys had 40, and we will bring them up to 70, based on ideas they have been talking to me about. So you give them some budget guidance.

ASPRAY: Over what period of time would the discussions have gone on with the office directors?

KAHN: I was going to say, let's say this guidance comes back from OMB to the department or agency. It somehow says, "For planning purposes this is what you should assume." And the agency director might say to his office directors, "Tell me what you would do if you've got that much money." So you go ahead and prepare for a budget review. The budget review typically occurs around here. So let me call this the... Maybe this is called the appointment process? I forget the name. In fact, we used to have three; there used to be a mid-term review, apportionment, and... This one was basically to figure out how to apportion the money. So what you do in this review is you put together a budget package, send it to

the head of the office and say here's how I plan to spend the money -- these 15 programs, and descriptions for each of them, how the money is going to be allocated, issues and resolutions, all that -- you go through this process. Then he comes back and he gets some guidance that says: Okay, this is what I told you originally. You're going to get 67; you're going to get 93; you're going to get 32; you're going to get 98; you're going to get 5; and you're going to get whatever the rest of it is. Or he may say, "Gee, I had really good ideas. Instead of being 70 here, you are going to get 100. So now you have 30 more. Whatever the amount. So now we go into the year. This is the guidance that they have been getting. All right, so now they have got to plan their program in some detail around this. Then, sometime around here is the FY 82 budget review. This is what is really going to go into the president's budget. That was all a preliminary process. That was a preliminary guidance. Now that we are nearer to the year and you get a better idea of what the real budget is going to be, and a better idea of what they're going to be doing with the equivalent of what they get for some big increase. It might have been, "Well, we want to build an electronic cathode, a magnetic gun that is going to shoot super ballistic at ranges of such and such." But when the details of the program are now presented, it does not look very feasible, or they do not need that much for this year; they need it two years from now. So this process can change it all again. Or they may get more guidance that says, "Look, we thought we were going to get this thing, and it looks like inflation is going to erode it and we are putting on a different number. So things can change. Anyway, this is the review that says, here is what we want to actually go into the president's budget with. And then, some time later, you get additional guidance back on that submission when the proposed programs are reviewed within Defense and DARPA. And here is a detailed write-up of our programs now. Somewhere back in here you prepare descriptive summaries to go to the Congress. These are the write-ups that actually show up in the president's budget book. Then the president's budget submission gets sent through the whole system and assembled into a composite whole. The president's budget goes to the Congress in January 1981. Then Congress deliberates on the submission. They're going to deliberate on that for much of the rest of the year. On the other hand, that money has got to be spent this year. There are three things you need to know about the spending process. There's the notion of committing money -- the commitment. That means basically signed out by the DARPA director. So why don't we just say "signed out". And then there is "obligation." That means "on contract." So if you sent money to the Navy and said, "Write a contract with MIT for such and such," it would have been committed, but it would still be inside the government. From then on the contractor would have been obligated. I believe DARPA is about to get its own ability to write contracts, if it doesn't have it already.

ASPRAY: I see.

KAHN: Then there is the expenditure. Now this is actually expenditure by the contractor.

ASPRAY: Yes.

KAHN: So you could have obligated a million dollars, but it might take them two years to spend it. The question is, when does this money flow? Well, I'm going to draw the commitment numbers for you now in red ink. Generally, you start committing dollars in the middle of the previous fiscal year for obligation after the new fiscal year begins. I would say that the dollars are usually somewhere between 40% and 70% committed by the beginning of the fiscal year. By the middle of the fiscal year, it gets up to something like 98% committed. Each of the offices would generally have committed 100%. There would be no reason for them to keep any. But the agency as whole would generally hold a little back, because you may get an expenditure for 1981 that you have to pay with 1981 appropriated funds, that you do not find out until fiscal 1985, because when they close out a contract and go back to do an audit, they'll find that the overhead rate was slightly different back here and you need to pay that extra money out of those dollars. So as you go back and do a history of DARPA, you find a small amount of the previous years money unexpended, but usually it is a very small percentage. But let's say that normally 100% of the money is obligated. The expenditure rate curve then shows how the contractor spent the funds over the term of the contracts that were written.

ASPRAY: And this was 40 to 70%, or something like that at this point -- committed at the beginning of the fiscal year?

KAHN: Yes, if I had to pick a number here I would say it is probably 60% expended at the end of the fiscal year.

ASPRAY: Sure.

KAHN: It's not zero, and it's not a hundred percent. And now I am going to show you, in black ink, the obligations. The funds cannot get on contract until the fiscal year begins. But, it might happen at the beginning again in a fiscal year that they might actually let many contracts instantaneously, and then the obligation rate would jump up from zero like this. Generally you do not get all the funds obligated by the end of the fiscal year. I believe the money can be obligated over two years if necessary. You might get up to 95% obligated, or something, by the end of the fiscal year with the remainder mostly obligated early in the next fiscal year. ARPA orders are issued to an agent of the government. Maybe to DARPA itself. Contracts are then written based on these ARPA orders. So somebody may receive an ARPA order here, but may not write the contract until out there. It is not shown on this obligation curve until it is on the contract. But then the expenditures might very well be spent like that. We have never had a good mechanism for tracking expenditures. But I would be very surprised if much more than half the money that was intended for that fiscal year was actually expended that year. So this is an expenditure curve. But in the middle of trying to spend your 1982 money -- remember you are doing this whole thing one year later -- you are starting to plan your 1983 budget. You get your guidance here, your budget reviews for 1983, and in the middle of here where you are still getting out ARPA orders for 1982, you are starting to do your 1983 commitments.

ASPRAY: So you are always going to have three years going on at the same time.

KAHN: Well, you could in fact have four, because you might have leftover money procured that did not quite get obligated in time, or committed on time. Or you may hold back money that an agent did not obligate and so you may be dealing with minus a year, or you may actually be dealing with plus another year -- so you are generally dealing with multiple years. But it is not only those three that are active at any one time. The other thing to keep in mind is what happens when... Here is a program manager, and let's say he gets a proposal from a contractor. The program manager then will generally send it up to the office director with paperwork requesting the proposal be funded. They may have some intermediate steps because they have assistant directors, and they've probably been staffed up. When I was there it was sort of one link, although we had a deputy director, so it sometimes went through three steps. But normally there was one tight chain here. From IPT it would then go to the program management office. Then it might go to the DARPA director or go back to IPT if any problems occurred with the paperwork. The program manager would create an unsigned

MRAO -- memorandum requesting an ARPA order for the director to sign. It would be sent up to the program management (the comptroller's shop). They would take that package and add a draft ARPA order. This package would go to the DARPA director's office. There might be a deputy director along the way who might actually just sign it and get it out. Then, when it is signed, they send it to an agent, like ONR, or RADC, or NASA. Some organization in the government who could write a contract. Now it could be DARPA itself. So they send the signed ARPA order. And the agent would have to staff it all the way through there. They have got about as many steps to take as DARPA did. The whole process presumably ends with a contract being signed. So that's how you close the book.

ASPRAY: How long would you say is a good average to pass through this path?

KAHN: It varies. Typically, a proposal might come in and the program manager might iterate it with the proposer. He could iterate for a year or two until he gets comfortable. Whenever he gets comfortable. Say he has a proposal that the proposer wrote just perfectly the first time. The time it might take him to get from submission to MRAO is a function of how important it is, of how high up it is on his priority list. A MRAO could be produced in one day if he wanted. Typically, this stage takes a few weeks to a few months. The office director will get to it anywhere from a fraction of a day to maybe a week. It is generally not much more than that much time for him to sign it out. The comptroller's office is usually pretty busy. This stage could take anywhere from, I would say, a few days, if it is crucial to get somebody to work on it to many weeks. It depends on how busy they are. It could be a few weeks typically. And over here in the director's office, this is very quick -- usually, I would say, they do not keep the paperwork for more than a few days at the most in the director's office. One to three days at the most at the directors office. They will either read it and sign it or dispatch it back. So, working it through the system, the biggest delay is upfront getting enough motivation, because this is the guy that's got to prove the case...

ASPRAY: The program manager.

KAHN: The program manager has got to build the case. [Referring to the charts] This guy will get it through fast. This may be a slow thing there. It depends on how bottle-necked they are. remember again, they get inputs from lots of different places. This is usually a rather big bottleneck here. This over here could take anywhere from... well, when I

was there anywhere from... you see, the actions can be prepositioned. I mean, the agent knew about this six months ahead of time. They knew the technical value of it when they saw the proposal. And they were all ready and everything had been staffed through. They could even go into negotiation with this guy, subject only to the money showing up. I've seen cases where the money has shown up and we have had a contract in place the same day.

ASPRAY: I see.

KAHN: But that's because they started the whole process nine months ahead of time, or something. But if these guys [Referring to charts] get an ARPA order just out of the blue -- it just shows up -- this could take anywhere from six to twelve months, if it is out of the blue. It could be anywhere from one day with prepositioning to let's say, a few months normally. That is, if they knew about it and were all prepared it would probably still have taken a few months to get it through the system. Today, they've got to advertise it for 45 days in such and such a journal. They have got to make sure if there is anybody who thinks they have got to go get competing bids then they've got to put out solicitations in advance to do it. Now they have to put out broad area announcements to get proposals, and compete through this whole process.

TAPE 3/SIDE 2

KAHN: So I think the contracting delays are much larger today, but you should not take my speculation on that. You ought to find out from the guy who runs this office. He would know the details. He should know all of it.

ASPRAY: Program Management Office.

KAHN: The guy who runs it currently is Jim Goodwin. You had better ask him. Today they have to put out either an RFP or a BAA. [Referring to charts.] After they do that you might get proposals over here. They might get the evaluation over here, and then they might get the ARPA orders out here. Those contracts in this range are signed here. Now the question is how long does this have to take? They could have started here and it would have been done maybe here. And so, what might go six months before now looks like a year and a half. Before, you might say, from the time

that the proposals come in until the contracts get signed, maybe that's only a year. Or if they can expedite this process, and get a contract or two to help them, maybe it is earlier. Oftentimes we had a proposal that was dated in January of one year that would not get on contract until July of the following year. Now, sometimes we were at fault. It just took a long time to get out of the system. Busy doing this stuff, not doing that stuff. But you can get some good statistics on just what the timing is from Jim Goodwin.

ASPRAY: Okay.

END OF INTERVIEW