An Interview with

FRANK HEART

OH 186

Conducted by Judy O'Neill

on

13 March 1990

Cambridge, MA

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

Frank Heart Interview
13 March 1990

Abstract

Following a brief overview of his fifteen years of experience at Lincoln Laboratory (including work on Whirlwind and SAGE), Heart describes his move to Bolt Beranek and Newman (BBN) and how he became involved with the ARPA network project. As the manager of the project at BBN for over ten years, Heart discusses his relationships with the group at BBN, DARPA and Lawrence Roberts, and the host community. Some of the problems encountered and surprises in the development of the network are addressed by Heart, as are the changes he has seen in DARPA over the years of his involvement with them. This interview was recorded as part of a research project on the influence of the Defense Advanced Research Projects Agency (DARPA) on the development of computer science in the United States.

FRANK HEART INTERVIEW

DATE: March 13, 1990 INTERVIEWER: Judy O'Neill

LOCATION: Cambridge, MA

O'NEILL: I'd like to start with a brief sketch of your career, your education, and your work experience before getting into the ARPANET.

HEART: Okay. I'm an MIT graduate from the 1950s with a Bachelor's and a Master's degree in electrical engineering. I went to Lincoln Lab and worked for fifteen years, and then came to BBN in December of 1966, and have been here ever since.

O'NEILL: Okay. You went directly to Lincoln after getting your master's degree?

HEART: Right. I worked at Whirlwind, and the Whirlwind became part of Lincoln Lab, and I went directly to Lincoln Lab.

O'NEILL: Can you describe some of the projects that you worked on at Lincoln?
In particular, the real-time computing projects?

HEART: Well, there was a wide variety of projects. I worked on all kinds of connections of computers to radars, and computers to seismic arrays, computers to computers, computers to radio telescopes, and pointing antennas, and various things of that kind. I worked on the Sage system for a long time. So I worked on an enormous number of different connections of computers to real-time devices.

O'NEILL: How did you first get involved in computer projects? In the 1950s, even at MIT, there weren't too many projects.

HEART: Whirlwind was active in the early 1950s. As a graduate student, I became a research assistant at

Whirlwind. So there were indeed computers. In fact, the first version of the Whirlwind computer I worked on had a total of sixteen registers all together, which you put things into with toggle switches.

O'NEILL: What kind of work did you do as a grad student at Whirlwind? Were you doing what we call applications? Were you programming?

HEART: Yes, programming. At that time Whirlwind was being considered for use in the Sage system, and Whirlwind was being modified to work as part of a thing called the Cape Cod system, which was a precursor to the Sage system. I worked on real-time control of airplanes with the Whirlwind computer.

O'NEILL: Were you running any of these programs, managing them?

HEART: Sure, at Lincoln, certainly I was, since I was at Lincoln for fifteen years. You know, I started out as a graduate student and then became a staff member, and by the time fifteen years was over I was running a group. So, sure, I was in charge of sizable efforts to build computer-based control systems and computer systems for antennas and radars over that period. I wasn't doing that when I first began, obviously.

O'NEILL: Did you recruit the other people that worked on some of your projects? In particular, the people who came over to BBN later who had worked with you at Lincoln.

HEART: Yes, some of them came to Lincoln to work with me in the beginning, and some number of those followed me to BBN when I moved to BBN - over a period of time. By the time I had moved from Lincoln to BBN, the MITRE Corporation had been formed and so some of the people who had worked with me at Lincoln had by that time gone to work at MITRE. So some of the people who followed me here followed from Lincoln and some from MITRE.

O'NEILL: When you were at Lincoln, did you have knowledge of, or interest in the work of Larry Roberts and Tom Marill when they connected the SDC Q32?

HEART: Well, sure. You know, Lincoln was a relatively small, closed community and certainly I knew both those people. I never worked directly with Larry Roberts. He was in a different group at Lincoln, but I knew him well, and I knew what was going on. And Tom Marill had originally, actually, worked for me at Lincoln at one point, before he went off to CCA. So that was also a small community. Tom came to BBN and then to CCA during the period that I was still at Lincoln. But I was not directly involved in those experiments that Larry and Tom did.

O'NEILL: Did you see any connection between the work you were doing and what they were doing?

HEART: Not at that time. As I say, those were experiments that I wasn't involved in.

O'NEILL: Why did you go to BBN?

HEART: I was attracted there by other people who wanted me to come. Actually, I wanted to come to BBN because I wanted to work on life sciences. I met some BBN people at a summer project that I went to at Woods Hole, and they began to think that I would be able to help them rescue a project that BBN had with the National Institutes of Health. I have a very hard time making major lifestyle decisions, and so it was very hard for me to move from Lincoln after fifteen years. But I managed to decide to move here. And I got here just in time to help officiate at the funeral of the project that they wanted to me to come run. So that didn't really go anywhere. But I actually had come to BBN to work on life science things.

O'NEILL: Were these other people, people who had worked at Lincoln, or just people you had met at this Woods Hole project?

HEART: No. BBN people whom I had met at Woods Hole.

O'NEILL: When you were at Lincoln did you have the opportunity to choose your projects? I guess I don't really

understand how Lincoln did its projects. Were there contracts from the Air Force?

HEART: Lincoln had at that time - it still has - primarily Air Force money. At that time certainly it had primarily Air Force money. And the projects it worked on were negotiated between Lincoln and the Air Force. The Sage system was a big deal. That was something that started when people were still at Whirlwind at MIT, and where a decision to move that - where it almost became a part of Lincoln Laboratory by a very complex process involving many important people. Getting the Sage system to be at Lincoln was a big negotiation between MIT and the government - not something young graduate students were involved in.

O'NEILL: Did you actually work on communication aspects of Sage?

HEART: Oh, yes.

O'NEILL: So you had what we would consider communications background?

HEART: Oh, yes. In fact the group that ended up at BBN probably knew more about connecting computers to communication lines and to real-time systems than any other group in the country. It was probably as knowledgeable as anybody in the country about how to do those things - because of the work at Lincoln. I'd like to believe that was why we won the contract, because we knew something.

O'NEILL: Okay, so you came to BBN and the project that you had originally thought you were going to work on was dead.

HEART: No, it wasn't dead. I was there for the year or two it took to die.

O'NEILL: Oh, okay. Did you start working at that point on the ARPANET?

HEART: No. That was in December 1966. The RFP for the ARPANET wasn't in the works until 1968. So there was a one and a half year period where I was working primarily on life science projects; the one that was dying and others that were starting. In fact, I continued to have responsibility for those life science activities even after the ARPANET project started.

O'NEILL: Well, if you presided over the NIH project for a year and a half, that would bring us to mid-1968.

HEART: Yes. My connection with the ARPANET did not really start at the very beginnings of the ARPANET insofar as it was being discussed in the government. In other words, there were government things going on that I was not particularly aware of. BBN had DARPA contracts, research contracts, and as such BBN people were involved with DARPA. But I wasn't. In other words, there were other people at BBN who had DARPA contracts of various kinds. Therefore, BBN as an entity, and some of those people, knew something about the early work that was going on in the government, thinking about the ARPA project. In particular, one of the people at BBN at that time, who was involved with the government earlier than I was on this, was Bob Kahn. Bob was involved in the piece of BBN that had DARPA contracts. So he was involved, at a relatively low level, but he was involved in some of the early, internal government discussions with the research community about the idea of the ARPANET. But that was quite a long while before there was ever an RFP.

O'NEILL: So the PI meetings in 1967 and those kinds of things, Bob Kahn was involved with but you were not.

HEART: I can't really tell you who went to those PI meetings. It may not have not been Bob. I don't really know that. I was not going to those meetings. The meetings in 1966 and 1967, I wasn't involved in. I don't really know who else was, or whether Bob was.

O'NEILL: So how did you get started: when the RFQ came out?

HEART: No, no, not quite. In the period of time when the RFP was being thought about by the government, Larry

Roberts was going around talking to various different organizations about their potential possible interest in bidding on this RFP. People were getting a certain amount of advance warning that there might be such an RFP. He was talking to AT&T, and he was talking to a lot of people. And one of the groups he talked to was BBN, and in particular, to me, because he knew that I had a set of people who knew something about this world. So I basically heard about the possibility of this RFP from Larry. But, you know, there was no real data. People were going around saying, "Gee, there's going to be this possible interest." He must have talked to fifty groups. (Laughter)

O'NEILL: Did Larry come and talk to you directly, or did he come to talk to a group of you people?

HEART: I guess I don't remember the exact details of that. I simply met him and heard about this. I don't remember the exact way I first heard about it. He was simply talking to a lot of people, and it was becoming bandied about that there was this possibility. BBN got very interested in this, and I got very interested in it. We made a decision to look into it. And once we started looking into it, we eventually got a copy of the RFP.

O'NEILL: What sparked your interest in it?

HEART: It was a potentially exciting new world. I don't think it was any one thing. It was just an obviously interesting area of R&D.

O'NEILL: Did you feel like it was an extension of what you'd been doing with the real-time systems?

HEART: Well, I think a better way to say that is just that we felt it was something we knew something about. I mean, it had aspects that we thought we knew more about than a lot of people. In particular, how to connect computers to communication lines. You know, in that period of time, the computer world was really somewhat divided into people who understood real-time systems and those who didn't. I mean, there were people who built computer systems using operating systems, and there were other people who built very finely tuned machine language programs for dealing with phone lines. And those two camps didn't interact a great deal. So the world was not full of people who

knew how to make computers run in real-time, and connect to real-time systems. We were not the only group, but it was a somewhat small universe. So it was an obviously interesting thing to a set of people who had been doing those kinds of things.

O'NEILL: On working to respond to the RFP, or RFQ, who actually was involved in doing that work? Was Bob Kahn working with you?

HEART: Yes, Bob Kahn was working for me at that time on this project. There were other people who were involved: Will Crowther, and Dave Walden, and somewhat later Hawley Rising, and Ben Barker, and others. But in the very early days, the people who had a very major role in the thing were Will Crowther, and Dave Walden, and Severo Ornstein.

O'NEILL: I thought Crowther came a little bit later.

HEART: No, he was there before the proposal was submitted.

O'NEILL: Did you actively try to recruit these people to come and work on this?

HEART: Well, no, it wasn't that way. Some of them were at BBN already. They were not recruited to work on this project. No. In general, the people who were here got involved in it because they were around me, and then others were added later in the process as it was going on. No, people weren't recruited to work on the project.

O'NEILL: Were they already working for you?

HEART: Dave was, and Severo was, yes. Crowther came, as I say, during the period the proposal was being written, as I recall - a little bit later in that phase. Crowther may have conceivably, again here's a place where memory fails, I don't really recall whether Crowther was influenced by the fact that this was going on or not. Maybe he said. I don't

remember the exact way that happened. You know there were people who wanted to come down here, potentially, and were obviously wondering what the hell we were doing. It may very well be that he was influenced by the fact that that effort was going on. That's possible.

O'NEILL: Can you describe for me your management style with these projects?

HEART: That's pretty hard. I think mostly I tend to believe important things get done by small groups of people who all know all about the whole project. That is, in those days all the software people knew something about hardware, and all the hardware people programmed. It wasn't a group of unconnected people. It was a set of people who all knew a lot about the whole project. I consider that pretty important in anything very big. So I suppose if you call it a management style, that would be something I'd state. I think also that they were a very, very unusually talented group. I think things tend to get done best by small groups of very, very good people - if you can possibly manage that. You can't always manage it. So if you again want to call it a management style, it is to get the very, very best people and in small numbers, so they can all know what they're all doing.

O'NEILL: Did you stay technically involved in everything?

HEART: Oh, sure.

O'NEILL: For the most part, when there was an issue to decide or some problem that needed to be solved, did you kind of pull everybody into a room and discuss it?

HEART: Well, yes. With a group that size there was very, very frequent interaction on problems, and people worked on them together. Although there were people who were obviously spending more of their time on hardware, and others who were spending more of their time on software. One of the critical things is the trade-off between those two things. In a real-time system you can do more hardware, or you can do less in hardware and replace it by doing more in software. So, that particular trade-off is really very, very critical if you're trying to do something in real-time.

So there was just a great deal of interaction of that kind among the people. There were only a half a dozen people. It literally was at the beginning. The proposal was written by less than a half a dozen people. The proposal effort itself was kind of interesting. I don't know if you want to spend more time on that.

O'NEILL: Yes.

HEART: It was an extremely difficult problem to produce a proposal like that in a very short time, because the government has the very natural tendency to spend an enormous amount of time getting an RFP out. Then they breathe a big sigh of relief, and they want the people to write the proposal in zero time. So there was very little time. And it was a very difficult problem because BBN at that time did not have any hardware capability. BBN had no construction. We had design capability, but we had no capability to build hardware at that time. So we had to find a sub-contractor who would not only supply the computer but who would be willing to do the special hardware pieces that were required. That was extremely tricky, because many manufacturers don't do that. They certainly don't do it rapidly. So we had the problem of finding a manufacturer who would be able to produce the computer we needed and at the same time be able to take on, on a very urgent basis, the construction of the special hardware to our designs. We were very fortunate that there was group at Honeywell. Honeywell had purchased a comp any called Computer Control Corporation, Three Cs. There was a group of people at Honeywell who were used to that kind of thing, who were not just building their standard machine but understood a bit about the special systems business. In fact, Honeywell had itself a tremendous problem because several of the proposal teams figured that out. Several of the prime teams figured that out, and they all descended upon Honeywell. Honeywell had to cope with the problem of several prime teams coming around simultaneously wanting instant response and wanting help in the bid. Fortunately we were very close to Honeywell. It was only a drive, rather than an airplane ride, and that worked out very successfully. But that was very difficult, and the people who had to put that bid together here were working very, very long hours. It was a very, very difficult bid to put together, primarily for that reason.

O'NEILL: Had you known the people at The Computer Control Corporation?

HEART: Yes, we had in fact. But that didn't help a great deal, because it still was a challenge for them to decide if they could throw themselves into this. The fact that one knew their names in a net didn't really help a heck of a lot. It was still a difficult thing to convince them of. They had to do some soul searching before putting the energy in.

O'NEILL: I was also thinking of your choosing them.

HEART: Yes. It was very hard. It was very difficult to look over everything that was available. Usually at any given instant in time, if you understand what you need in the way of a computer for a real-time problem, it's pretty clear after a while what the one or two choices are. And there really weren't that many. But it was still... the urgency made it very difficult. The fact that Severo was as good a hardware designer as he was and was able to work as well with... was able to sort of get in the trenches with the Honeywell people was a very critical aspect of being able to get the interfaces built.

O'NEILL: Were you expecting to get the bid? Can you give your level of surprise?

HEART: Oh. (Laughter) Well, you know, it was the kind of thing where it was extremely competitive, and there were other larger organizations. Most of the other organizations bidding were larger. So, no, we weren't at all sure we were going to get the bid. Quite the contrary. It was not a question of surprise, but it was just very competitive, and we didn't have any particular level of confidence at all.

O'NEILL: Did you pretty much know who your competitors were and what they were proposing?

HEART: Let's see. Probably not. We knew some of the teams, I'm sure. But I don't think we did know them all, no. The government was doing that bid quite carefully. It was a fairly arm's length kind of thing, and we didn't really know who all the other teams were, no. The government, after spending a lot of time thinking about it in their own councils, once they decided to put the RFP in the street, they played it quite straight. There was very little knowledge of what was going on available to the bidders. Unlike most of ARPA, you know, it was unusual. Because

in fact, in DARPA's experience, most of their contracting at that time was sole source contracting. So this, for them, was actually an unusual bid. Because they were urgently trying to play by the rules, and not play this from sole source at all. It was tough on them, too. They also had an unusual experience trying to do this.

O'NEILL: When you say them, you are referring to Larry Roberts and...

HEART: Larry Roberts and the set of people he was using within the government community to help him. There were a number of people in the research community participating with Larry and trying to make this happen - who had been involved for a couple of years in talking over the possibility of the network. Those were the people who participated in the evaluation.

O'NEILL: So you knew who was going to review the proposals.

HEART: We knew some of the people who had been involved for some time.

O'NEILL: Once you won the contract for the ARPANET, did the group change very much at that point?

HEART: Well, it increased slightly. Some other people who became very important were added. Let's see, Ben Barker was added to the group. There were two or three other people who were added to the group in that first year. It changed from a bid to a racing around to implement. Sure. It changed a lot. But it was a lot happier. While you're working maybe just as hard, at least you are being paid for it, and not thinking the effort is going down the tubes. It was different in that sense. It was actually a very, very unusual situation to have a contract of that urgency for as small an organization as we were. We had to cope with the phone company putting in lines. We had to cope with Honeywell building equipment. We had to cope with our own team building software. And it was really a very small group. It was a labor of love in those days, by a relatively small group.

O'NEILL: Did you try to get a larger group? Or did you want to keep it at that size?

HEART: No. Well, first of all, the bid was for a certain amount of money. And the bid was only going to support a group of a certain size. So a larger group wasn't an option, actually. But, also, I think that we felt strongly that only a very small group could do it. And if the group had to start interacting with paper, it would not have been feasible for a larger group, or a much larger group - in that time scale.

O'NEILL: How did you interact with ARPA?

HEART: Very closely. See, that was one of the things that was very nice. In fact, if you want to digress for just a minute into the reason I think the thing was as successful as it was, the primary reason was that there was close control from DARPA and a guy there with a great deal of freedom. If you had to give the single most important reason why it was as successful as it was, it was that Larry Roberts had a great deal of authority and freedom and was able to control not only the contractors who were working on it, like BBN, but also the users, since he was supplying all their money. In other words, all the sites at which the IMPs were installed were research sites being supported by DARPA. So he could get their cooperation by the simplest of techniques: he was supplying the money. The combination was quite critical because the sites had to do all the work too. They had to build special interfaces. The sites where the installations were made, those first four sites, had to build special hardware of their own, they had to build special software of their own, and they had to do it on a time scale that was very tight. And furthermore, they had trouble starting because BBN had to supply them specs for those interfaces before they could build them. So, not only were they under a tight time scale, they had to wait for a while until BBN could supply the specs for those interfaces. That, again, could not have been done if those sites had not also been under Larry's thumb, in a sense, and were trying to be responsive to DARPA.

O'NEILL: Were you ever consulted as to which sites should come on-line or who should be added, or was that all a DARPA decision?

HEART: No. That was a DARPA decision. That had to do with where their support was going, and those were

DARPA decisions.

O'NEILL: In general, how did decisions get made?

HEART: BBN was a team member, and so we were involved in all of the global decisions. But your degree of involvement varied with the decision, and certainly the question of what the sites were was in some sense pre-determined by where the DARPA support had been going or where there were people who could, in fact, take advantage of the thing. So certainly it wasn't something we would have a choice about. We were involved, but that kind of decision was a DARPA choice. Like many of the choices. You know, if you look at the RFP, one of the things you would realize is that while BBN did a tremendous amount of the implementation, a lot of the very key choices had been made by DARPA. They picked the baud rate of 50 kilobits, they picked the sites, they picked the issues about the check sums. A lot of work had been done in advance by DARPA.

O'NEILL: How did you interact with Larry Roberts? Did he come for site visits frequently? Did you go down there?

HEART: Sure. But you also picked up the telephone - all the time. It was a very, very close interaction.

O'NEILL: Was it say weekly, daily, monthly?

HEART: Not daily, no. I guess I really can't answer that. I mean, it varied a great deal. There was a great deal of contact. I'm sure somebody here talked to somebody at DARPA weekly. It was not an arm's length grant, where they gave us money and went away. There was a lot of interaction.

O'NEILL: Was there anyone else at ARPA or IPTO that you dealt with beside Roberts?

HEART: Well, over time a lot of people. Let's see, Barry Wessler was certainly a person who was there early on. There were a lot of other people over time at DARPA. And, you know, that changed from time to time. So over five

years there might have been ten different people involved. There were also people in the administrative capacities at DARPA. Al Blue, and others who moved paper and made things happen when you couldn't get them to happen. So a lot of people.

O'NEILL: How long did you personally stay involved in the ARPANET development?

HEART: I was involved for, you know, a decade.

O'NEILL: Okay, so you saw quite a few things.

HEART: Oh, yes. I was involved for ten or twelve years - directly. It came alive very quickly, but it took a long while before it was a utility. And so I was involved for quite a long time directly.

O'NEILL: Can you describe for me what you mean by a utility?

HEART: Yes, sure. A utility is something people depend upon. Like the electricity, or the phones, or the lights, or the railroads, or the airplanes. Yes, it was a utility. That's the thing that was the amazing surprise. It was started as an experiment to connect four sites, and it became a utility much, much faster than anybody would have guessed. People began to depend upon it. And that was a problem, because that meant when you changed it, or it had problems, they all got mad. So that was a two-edged sword. But it was also very exciting.

O'NEILL: Can you identify when that change took place? Was it after the experimental system was in place?

HEART: Well, I certainly can't pick a date. I think that as the system began to go into more sites, people began to assume it was going to keep running. And especially as electronic mail became an important component of the use of the system, people began to assume it was going to exist. Then it became very important to them. But, no, I can't pick a date.

O'NEILL: I've gotten different reactions from people, some people feel that they knew early on that the network was going to be a big deal and change the way people worked. Others felt that it was just an interesting technical problem at the time.

HEART: There was a third class who believed it would never work. Seriously.

O'NEILL: None of the people I've talked to, but okay.

HEART: Right. But there were others. Because you know the notion of packet switching - some people thought the routing stuff would never work. Some people felt it would just plain not work. It would not behave, and it would be unusable.

O'NEILL: Did you have direct discussions with these people saying this?

HEART: Sometimes, sure. We thought they were wrong. We thought it was going to work. (Laughter)

O'NEILL: Can you give me some idea, not necessarily names, but who these people were. Was it the academic community, was it the telephone people?

HEART: I don't know. There were people all over the communications community. Certainly the big carriers who had turned down the opportunity to work on it and decided against it, and some academics. There were a variety of people who thought it wouldn't work, because the technical problems were not small. It had to send stuff reliably, it had to have retransmission that worked, it had to work through phone line errors, and had to work through phone line downages and outages, and so on. And it had to solve the routing problems, so the stuff didn't keep circulating and got to the place where it wanted to go. There were people who thought we wouldn't be able to do that. It was a technical opinion. And there was no proof. You couldn't go to some book and find the answer; it was an experiment.

We were quite confident, but there were people who didn't believe that.

TAPE 1/SIDE 2

O'NEILL: Obviously you weren't in camp three, but of the other two, did you have a feel for the broader applications of the network?

HEART: We thought it was important from pretty early. Yes.

O'NEILL: What was your relationship as manager of the project to the host community?

HEART: Well, we were a member of the ARPA community, and so we interacted very directly with the host community that was working with the net. We were members of the committees that tried to work on the host protocols. We were members of the committees that went to the principal investigator meetings. We were members of the community. In the host community, in particular, we were providing them specs and we were providing help. When they tried to build interfaces, they required help occasionally. We had people whose job it was to try to help them get connected. And we would help them with hardware and software problems. We would help them with testing problems. Because it was in our interest to have them have a happy connection. So we put considerable energy into working with the host community - not in applications, but on the connection, the hardware and software of the connection. There is a quite well-known document called 1822, which was a specification for the hardware interface. We had people who understood that completely and could help somebody work their way through it when they needed to. So there was a lot of interaction. And we participated in the host protocol developments. But that was up a level; in other words, we were directly concerned with the connection of the hardware and the software at the channel level. The host protocol committees, the Network Working Group as it's been called, we participated in as members, and we played a large role in it, but that wasn't our primary function. We had no special relationship there, we were simply one of a number of people who contributed and worked on it - because it was interesting and important and DARPA wanted us to. But it wasn't that we had a primary role there.

O'NEILL: How about with other contractors, did you interact with them? I'm thinking of people like the Network Analysis Corporation.

HEART: Oh, sure. We had very close direct contacts with the other groups that were working on it. Sometimes we'd argue with them, but, yes very close contact. And in particular, the Network Measurement Center, Lenny Kleinrock's thing, had to depend upon measurement facilities that were built into the software of the IMPs. So their statistics had to be captured by dint of software we wrote, producing those raw data. There was a lot of contact over how that all worked.

O'NEILL: Were there discussions with these groups, either the host groups or when the hosts came together as a group, or with the other contractors, where DARPA had to arbitrate a disagreement? Or were things pretty much worked out within the contractor community?

HEART: Well, arbitrate is a strong word. DARPA was a player in all of those discussions. The people at DARPA were strong technical people in their own right. I mean, they were not just managers. Larry Roberts was a technical person, who had a personal technical interest in things. The other people who came into that office, Steve Crocker and others, were people with strong technical capabilities, and they participated as players in those discussions. So they were not passive managers. Now, arbitrate, that's... Certainly if anybody had to arbitrate they did. I don't know that I would view most of those discussions as arbitrations as much as they were simply participation with a louder voice than some of the other people.

O'NEILL: Okay.

HEART: I mean they were technical people in their own right. They weren't just arbitrating between competing visions of other people. In some cases they were leading. In the host protocol area, the DARPA people played a very central role. Crocker and Cerf and others were involved, you know, very directly in the Network Working

Groups, as kind of group leaders at some points. Have you collected the written material that exists?

O'NEILL: I've collected some. I'm not sure what you are referring to.

HEART: Do you have the recent book that was issued about the Network Working Group with Crocker's introduction?

O'NEILL: No.

HEART: Well, it has stuff... You might want to ask either DARPA or Steve Crocker for this. But this is probably just some background you might like. It has an introduction - the rest of it you wouldn't be interested in - but the introduction you might find interesting. So you might want to try to get hold of the introduction to this.

O'NEILL: This is the *RFC Reference Guide* by Joyce Reynolds and Jonathan Postel. Okay, great.

HEART: How about the report that was written some years ago about DARPA itself?

O'NEILL: The Barber study?

HEART: You have that?

O'NEILL: Yes. I do have access to that. Also, the Completion Report that was done here at BBN. Then Alex McKenzie had some documents he's letting me look through. So it's starting to build up. How did Bob Kahn fit into the group implementing the ARPANET?

HEART: Well, Bob was interested in several different problems. He had worked originally, before the RFP had ever been issued, on phone line error problems. He was working on error control problems. So one of the original

interests he had was in that. He also was very interested in the routing problem and worked on the routing. So the two things that I remember were those two. Now he worked on a lot of other things, simply as a member of the group. Since it was a very small group, everybody participated in everything. But if you had to pick the two areas that come to my mind as the ones he paid attention to at the very beginning, it was the phone line error problem and the routing problem. Bob, of course, left BBN at some point to go to DARPA. So at that point he began to have a very different role in the thing. Bob was also a very, very hard-working guy and wrote very well. So a lot of the documentation Bob played a large role in. He was a good codifier. He played a significant role in getting stuff on paper.

O'NEILL: Did you work with him once he had gone to ARPA?

HEART: Oh, sure, yes. In a sense we had to. He became the client. So sure.

O'NEILL: Talking about documentation, did ARPA actively encourage your group or you to disseminate information?

HEART: Oh, sure. In fact, you probably have seen the summary of papers from the two joint computer conferences. Two sets of papers were written. Yes, there was encouragement to participate in producing those papers and to publish stuff, absolutely.

O'NEILL: How did the collaboration process work on those articles? Did you as manager of the group, say, have to coordinate these articles or suggest who should do what?

HEART: Well, sure. As I say, since it was a quite small group, everybody had to work on it. Sure, I was helping to figure out who was going to do what, and trying to make sure they did. But I didn't have any larger role than others in trying to produce them. As I say, Bob and Severo played a pretty large role in getting those things actually on paper.

O'NEILL: Did ARPA review them before they were published or look at them, that you recall?

HEART: I guess I don't remember the answer to that one. My guess is that when those papers got written, they probably went out to a lot of people at once, and there wasn't any special DARPA review. But they may have, I really don't remember.

O'NEILL: During this time, as manager of the group, did you feel the need to shield your technical people from some of the political aspects of working on a government project?

HEART: No. Because, in that sense, it was a very benign government project. We had very close connections with the client and very friendly relations to the client. There wasn't much to be shielded. Unlike some of the more recent history where there's a lot of shielding to be done, that wasn't really a problem.

O'NEILL: That's one of those areas where the technical people certainly felt that there wasn't, but as manager you might have felt differently.

HEART: Well, there certainly may have been things that I did and didn't bother other people with, and if you call that shielding, fine. But it wasn't that there were really serious political problems that I was shielding anybody from.

O'NEILL: Did you have any exposure to the military at this time?

HEART: Sure. In fact some of the sites, after a little while, became military sites. So we, as the network operator, were involved with those sites and had to work with them. Again, they were fairly benign sites at the beginning, I mean places that were within the R&D community - RADC and other places that had a research bent. It wasn't until much later that groups came on the network that really had no interest in the R&D, just wanted the network as a service. The early sites, that were military sites, were people that had research interests. So they weren't really substantially different from the other sites. It was a little different, of course, but it wasn't substantially different. By

the time the network got to the point where it was involving other sites that were much more interested just in having the thing as a service utility, by that time DCA had gotten involved and a very different set of people were involved. The military did not represent much of an influence on the early ARPANET - is I guess the way to put it. To the extent that there was interest, they wanted to be part of the academic research community rather than our serving the military community.

O'NEILL: Okay. So they were trying to fit into what you were doing.

HEART: Exactly. In the beginning. Now it wasn't true much later when the network got much bigger. Then military sites would join in because they wanted to use it, and then they had different goals, and they felt differently. But in the early days, that wasn't true.

O'NEILL: And you were still involved with the network during switchover to the Internet?

HEART: Yes. Oh, yes. Well, for the first dozen years, 1969 to sort of the beginning of the 1980s.

O'NEILL: Can you describe your reaction to the growth of the network and the Internet?

HEART: I think that as the ARPANET became a big deal, and as the whole world began to adopt this technology, and as there began to be other networks all over the world, we found that very exciting. And we became a tourist stop for everybody and their brother to come and see the thing, and that was always nice. It became a focal point for many, many other people who wanted to understand what was going on, wanted to come by. As that began to grow, more networks grew in the United States as well as elsewhere. So it was a fairly natural evolution. It was all part of a very happy experience to watch it becoming popular. It made people feel that they had been very successful.

O'NEILL: Did you actively try to work with some of these other groups? You mentioned they were stopping by.

HEART: Yes. BBN did try to. We tried to do various different things. In fact, the eventual decision of BBN to form Telenet was based on our wanting to try to leverage this experience. I don't know how much of this story you've heard, but at one point there were some people working on the network project who decided that BBN was not being sufficiently aggressive in exploiting it, and they left. They formed a company called PSI maybe, something like that. I've forgotten the initials. If anything were to goad BBN to do something, it was the fact that the people left to go do it themselves. So shortly thereafter BBN, in fact, formed Telenet. Telenet was eventually sold to GTE, but that was a long time later. BBN hired Larry Roberts to be president of Telenet. So BBN tried various different ways to exploit that technology, and is doing so to this very day. A major division of BBN, BBN Communications, is trying to sell packet-switch networks around the world. So certainly it was a big deal for BBN, and many different things at BBN grew out of the work. I would take the position that a non-trivial chunk of BBN's growth over those years was due, in no small measure, to the ARPANET contract. So, yes, it was a very important thing for BBN.

O'NEILL: Did you consider going to Telenet, or doing something like that, striking out on your own?

HEART: Not particularly, no. It was nice here.

O'NEILL: As I understand it, early on you had connections with the people over in England, Donald Davies and the National Physical Lab?

HEART: Yes. They had done work early on. They had built a small network that was co-timely with the early ARPA network. They were one of the groups that was, in fact, suffering from not understanding real-time systems so terribly well. They, in fact, were one of the groups that had gotten trapped into trying to use the operating systems in their computers and to connect their computers together with operating systems involved, rather than building tailor-made machine language code. So they were working under a fairly serious disadvantage early on, and Crowther's hand-crafted machine code gave our system a large performance advantage over their early systems. But they were a smart group, and they did some very good things. But there unfortunately wasn't a lot of cross-learning in the very early days. It might have been better had we known about them earlier and spent more time there. We

weren't able to take advantage of really very much they were doing. Although we did interact with them a lot, and the connection got to be quite friendly after a while.

O'NEILL: Were you aware of the work of Paul Baran at the Rand Corporation?

HEART: Oh, sure. Paul Baran's work had been one of the things the DARPA community people all knew about. So that work was, in fact, in the minds of the ARPA guys who worked on it at first. And we certainly all knew about that work and had looked at it, yes.

O'NEILL: Do you know how you knew about it? I mean, were you sent copies, or was it just in the community?

HEART: Again, I can't remember that either. It was in the community, and certainly the DARPA people all knew about it. So sure, that work was available to people. Again, it really did not make a huge difference in what one did, because we were building a new system. And there was nothing in the previous work that said how you should do check sums, or how you should do retransmission, or how you should do a lot of the things that had to be done. It mostly provided a backdrop of information about the world and about some of those ideas. It wasn't helpful in the small. It didn't help the guy writing the code write his code.

O'NEILL: Okay. You mentioned E-mail before, that has been described as a surprise to a lot of people. Would you include yourself in that group?

HEART: Well, the surprise was the extent to which it became important. When the network was originally built, Larry probably had - if you had to list his goals, you can look at the DARPA order, but if you had to list his goals - he certainly had high in his set of goals the idea that different host sites would cooperatively use software at the other sites. There's a guy at host one, instead of having to reproduce the software on his computer, he could use the software over on somebody else's computer with the software in his computer. And that goal, has, to this day, never been fully accomplished. That goal still to this very day has not been really accomplished to the degree that it was

hoped for in its early days. So that turned out not to be the main thing that was created by the ARPANET. Other things, people interactively working together, people being able to communicate with each other more easily, people being able to have joint projects and just being able to use software they both had developed over the Net, became much more important. In that sense, when the mail was being developed, nobody thought at the beginning it was going to be the smash hit that it was. People liked it, they thought it was nice, but nobody imagined that it was going to be the explosion of excitement and interest that it became. So that was a surprise to everybody, that it was as big a hit. And it became a major network use; it became perhaps the single largest network use, finally.

O'NEILL: Did you use the mail yourself?

HEART: Oh, sure. People internally used it a lot. As I say, it wasn't that people didn't know it was a nice thing. I think they just didn't have any idea of the social impact it would have and the degree to which it would acquire usage. The people who were in the community were using it right from the beginning, a lot. But they were using it without really realizing that the whole world was about to want it.

O'NEILL: You mentioned that the sharing of resources on the level that ARPA envisioned it, in your opinion, didn't take place, and then the surprise of the mail having more of an impact than anyone would have suspected. Are those pretty much the differences from expectations that you would list? Are there other things that you are surprised about how they turned out?

HEART: Yes, there were other things that were surprises. One interesting one was that when the network was first envisaged, the idea was to have an IMP and a host. And one of the first things that happened once the project started was that within a month or a mo nth and a half it became perfectly clear that all the host sites had more than one host they wanted to connect to the IMPs. So that was a big surprise. (And that wasn't as important as the next thing I'm going to mention.) But the notion that people wanted to connect more than one host was a big surprise to our group. Because that was just a big change in what we had to do. But then it began to be realized that the IMPs could be used to connect the hosts together at a site. That became very important to many people, because for the

first time there was a protocol and a specification for how you connect a host to an IMP. But once you do that, you can connect two hosts at the same site together through the IMP - something which had been very difficult for people to do, especially with heterogeneous computers. So the notion of using the IMP as a local connection was quite a surprise, to the extent that it became just common and had not been envisaged. It wasn't as big as the mail surprise, but it was a surprise.

There were other surprises. I don't know that I can necessarily come up with them right now, but there were. I think that it was a surprise how well one could debug phone lines. In other words, the phone company had never been able to tell when a phone line was about to fail. Their technology for dealing with phone lines was when someone called up and said, "I can't talk over the phone," they would send someone out to figure out what was wrong with the phone line. The IMPS watched the phone lines all the time, all the time, and they could tell when a line was degrading, not just when it was failing. So there were amusing instances when somebody here would call up the phone company office in California, and tell them that the phone line between Los Angeles and San Francisco was about to break. And the phone company guy, after first thinking we were calling as a joke, would then say, "How could you possibly know that in Boston?" A lot of that went on. So there were surprises in how well the IMPs were able to do phone line monitoring, and debugging was an interesting thing, which I don't think the... The carrier certainly adopted those techniques afterward.

O'NEILL: Were they fairly open to all of this?

HEART: Oh, yes. But it was just a source of, at first, some amusement and then some amazement that the thing was such a very good phone line monitor. And that proceeded even into other things. When the network had satellite links it was able to tell when the satellite power at some site was slowly drifting. And it was able to monitor the performance of phone lines. Another kind of surprise was the difficulty of the whole problem of monitoring. You know, when you are at a terminal and you punch your key, you want something to happen. When you're connected at your terminal to a phone line in your site to some computer and it's connected in turn to the net, the net's connected across the country to something in California, which is in turn connected to some other computer, but

you, as the guy who hit that key, do not care about that; you just want something to happen. So when it doesn't work, you don't know what the fifty things are that might go wrong that could possibly be affecting the fact that you hit that key. And there had to be people who could debug that problem. You had to be able to call somebody and say, "I just hit the carriage return and nothing happened." Well now you had the problem of trying to figure out where in the country that trouble was, whether it was a distant host, or whether it was the host connection, or whether it was an IMP at the far end, or whether it was in a phone line. That was a very difficult problem, and it's still a difficult problem. People still spend tremendous amounts of money and energy on network monitoring systems, and they generally do it very badly. It's a very difficult problem. That was something which has represented a tremendous amount of energy over the years in trying to build systems that could monitor and debug the performance of networks and find difficulties when they occur and isolate them. For example, we at one time had an autodialer. And every night that autodialer would dial all over the country to every input terminal IMP port to see if you could, in fact, get into the IMP from there in case the equipment at the input dial end was broken. Because, otherwise, you would not know whether it was broken because it wasn't in the main phone lines. So we actually went to all kinds of lengths... I think people certainly did not anticipate at the beginning the amount of energy that was going to have to be spent on debugging and network analysis and trying to monitor the networks. That was a very hard problem.

O'NEILL: Did you set up the Network Control Center to deal with that problem?

HEART: Yes. It got set up many, many times over the decade, and it changed greatly over the period of time as one realized increasingly how much energy was required. Those programs had to get better, and people had to learn new tricks. There had to be new software written. That was a big problem. It still is a big problem. To this very day, even BBN doesn't do so well at that in its private network business. It's just a very difficult problem. And it's been even more a problem in the Internet. When you think of going through five or six networks to get to something, and now you hit your key and nothing happens. Now there's not a central authority anymore. At least in the ARPANET days, you knew where to look. I mean there was a BBN who had a contract to run the network. So at least in that period people knew where to call. We didn't always like getting the calls, but they knew where to call. When

something breaks in the Internet, it's often not known where to call.

O'NEILL: Was the network control center considered part of your contract?

HEART: It was part of the contract.

O'NEILL: Okay. So that was envisioned early on in the process.

HEART: Yes. It was envisioned, but it was not envisioned how tough, and expensive and difficult it was going to be. People knew they had to do it, but they didn't have as much of a feeling for the degree of problems. Tremendous efforts went into that. I mean, the early IMPs could cross-patch their lines remotely. We could, from here, turn around the output line of an IMP in California to test whether the bits were going back in properly, to try to distinguish between an IMP failure and a host interface failure. Because when those things broke, how could you tell? You couldn't try to get help from the sites. The IMPs were unmanned, and the host sites didn't know much about them, and they didn't want to play with them too much. They didn't have people who knew much about them. So that was all being done remotely. That part of the problem was and is very difficult. It was a surprise how tough it was to agree on the host protocols. I mean the network's utility was delayed for at least two years because of a misapprehension of how hard it was going to be to get the host protocol suites in place. Given that the network works, now the host software has to work. It's like picking up the phone and calling France - if you don't speak French you've got a little problem. So even if you get the connection to the two telephones, if you don't speak French you don't communicate very well when you call a random person in France. So, in the same way, even though we managed to put the network together, if the host sites don't have their code in their host computers properly working... I think it was misperceived how long and hard that would be to do. The Network Working Group was organized, people worked on it, but there was a lot of argument, a lot of discussion. It took much longer than people thought it was going to take, even with ARPA being on top of it and able to attempt to urge cooperation.

O'NEILL: BBN did play a role in the Network Working Group?

HEART: Oh, yes. But we were one of many there. We weren't running that.

TAPE 2/SIDE 1

O'NEILL: After your work with ARPA in the ARPANET, were you involved in other DARPA or government projects?

HEART: Yes. BBN has had a continuing relationship with DARPA, and I'm running a division that has 300 people in it, and so it has a lot of DARPA projects to this very day. It's had a lot of DARPA projects over the years, so, yes.

O'NEILL: Have any of them been anything like the ARPANET in terms of how they are run? You mentioned that Roberts had a lot of authority and freedom.

HEART: Well, I think probably the way to describe that is that it has gotten worse steadily. And then with the advent of the Competition in Contracting Act it got worse with a big step function. So it's been worse, much worse.

O'NEILL: Can you explain the Competition in Contracting Act?

HEART: Well, I think that, as I think I said earlier, most DARPA contracts, with the exception of the ARPANET contract at that time, were sole source contracts, where DARPA had a history of finding good people in good places and giving them support to work on things. That's not possible in the current environment in anything like the same way. I mean, these days almost all DARPA contracts are competitive, rather than sole source, because Congress has mandated that that shall be the case. And so these days often contracts go to the low bidder rather than to the - I'm exaggerating slightly for emphasis - but I think it's just a different environment. Also DARPA is much bigger. The people are not always of the same quality, in my opinion, as a Larry Roberts was. Certainly there are good people, but they are not that good usually. The government environment itself is not as pleasant, so DARPA has had

increasing trouble over the years attracting very good people. It's had a harder time attracting good people. It's not as good a place in a lot of ways. And isn't able to do as good work, in my opinion, because of those changes. I think Congress has shot itself in the foot in those regards. Congress tries to fix things that they think are global problems. And whenever you fix a global problem, you're going to manage to foul up some piece of the problem where you didn't focus on, and to have that change take place in DARPA was very foolish for the country. It has cost the country very severely. So that's one major difference.

I think, in addition, DARPA has gotten bigger, and the budget has gotten bigger, and the problems have gotten different, so it is not anything like the collegial... For example, for a number of years there were principal investigator meetings at DARPA that were really a lot of fun. They were some of the best meetings one ever went to. If you talk about all the meetings people go to in their lives, of professional groups or various society meetings, probably the meetings that DARPA held in those early years for principal investigators were among the most interesting, useful meetings that ever took place in the technical community. They don't happen anymore. There are no such things anymore of the same degree. They are now much bigger, and they're in some small area, there's not the same... it's just very different. So DARPA has become a different place. It's still a very good place, still probably one of the best places, but it's just become different. But BBN has had a long series of DARPA contracts, and we have many network contracts still - very many, many network contracts. I think DARPA still is one of the best research support places the country has got. People talk about trying to emulate it with a civilian DARPA these days, and that may happen. But there certainly is not the same degree of freedom there that there was once.

O'NEILL: Do you find that they try to keep as close tabs on what is going on?

HEART: It varies a lot from program to program. Some of the program managers there do try very hard to do that, others don't. But sure that probably hasn't changed as much. Once they get something going, if it is a pretty good guy involved, he does stay involved pretty closely.

O'NEILL: As you know there's been an interest in the history of the ARPANET with the twenty-year anniversary just

passing. Are there parts of the story as it's being told, the ARPANET story, that you would disagree with or would emphasize differently than what you have seen?

HEART: Well, probably everybody looks at it a little differently. But I'm not sure what you're thinking about.

O'NEILL: Let me see if I can think of a good example. Well, some people would emphasize the availability of mini-computers as leading significantly to the ability to do the network at that particular time.

HEART: Probably it's the case that a lot of things follow economics. Certainly, if one could not have gotten small computers for something like the prices they were selling for, one couldn't have done this the same way. I mean, DARPA couldn't have afforded to pay ten times as much at that point in time. Certainly I think I would agree that it was pretty important that there were a few sources of computers for prices that were in those ranges. If it had been ten times as much, it wouldn't have happened. Economics does drive what you can do. If you couldn't have gotten fifty-kilobit phone lines for the prices that they were then available for, but had they been a lot more expensive, DARPA might not have tried. Sure, economics affects things. But I don't think there are any differing opinions about that, particularly.

I suppose one thing I'd say is that I'm sure different groups view the importance of their own contribution with a slightly different coefficient than other groups may view them. Some people probably think that it was absolutely critical for the Network Analysis Corporation to design networks - absolutely crucial. I would probably not think it was absolutely crucial, because, in fact, most of the plans for what you connected weren't based on logic. They were based on who had contracts, and who was ready to do the installation, and where there was a user with money, and various other things. So I would probably have a different coefficient on the value of some given little piece of activity. But I don't think those are big differences. They're probably small differences. You know the people who were involved in this all were having a very good time. The ARPANET was a big thing in most of their lives. So most of the groups of the host sites, or the Network Analysis Corporation, or at Kleinrock's Network Measurement Center, or here, or others, were all having a very good time. They were all really having the time of their lives. While

they might have some differences, the overriding single bit was they were all having a great time. And they all thought it was very exciting.

O'NEILL: Do you think they were having a great time because it was a fun, interesting project, or was it something else?

HEART: I think they thought it was important, and it was a fun, interesting project.

O'NEILL: So you think there was an appreciation for the importance of what you were doing?

HEART: Oh, sure. Sure, I think so. Sure.

O'NEILL: That usually adds incentive to work long hours. Okay, that covers the questions I had, unless you have any general comments you would like to add to the record.

HEART: As I say, the only thing I would emphasize is that if people try to look at all the projects that this country does in the computer world, and all the failures there have been, all the large military computer systems that have fallen on their face, or had a factor of ten overruns in time or money... If you look at all those things, and then you ask, "Why do these few make it and be successful, and why do all those not be successful?" That's an important thing, if one could somehow get that lesson across to Congress and others, it would be kind of nice. It's very hard apparently because they keep doing it the wrong way, mostly. What happens is that people who know nothing about the technology write long lists of requirements. Then they go out to industry and ask industry to bid on these silly requirements, and they get back silly bids. Then they pick the low bidder of the silly bids, which often is the one who understands the technology least; and then ten years later they have a disaster. Instead of people who understand the technology very, very well in the government who are able to drive it to the groups that are the right groups, even if it isn't the low-price bid, and who take close control of it, who don't have an enormous list of specs, and then other people who never wrote the specs, but who have a participatory proto-typing kind of relationship

with the thing where it grows. If one could get across that lesson, it would be kind of nice. But I don't hold much hope for that.

O'NEILL: Can you give me any examples of what you consider big failures, using this kind of method?

HEART: Well, gee, there are probably... No, I probably can't do that, just because I wouldn't want to be quoted. But, I mean, they're legion. The number of defense systems where they've tried to build computer systems, and they've been horrible, there are probably lists twenty or so long. I probably wouldn't want to be quoted on naming them, but it's many, many, many subsystems that have been ten years late at ten times the cost - things like that. The ARPANET went from a contract award to an installed operating system in nine months. It went from a contract award to equipment being delivered, on site and running, in nine months.

O'NEILL: That was slightly ahead of schedule, wasn't it?

HEART: Just a teensy bit.

O'NEILL: It was supposed to be in by the end of the year, or something?

HEART: No, it wasn't that much. It was just a little bit. It was very close, very close.

O'NEILL: Financially, did the project stay within the budget?

HEART: Oh, yes.

O'NEILL: I hadn't ever heard that it didn't, so I assumed that it had.

HEART: It did.

O'NEILL: Okay. Thank you very much.

END OF INTERVIEW