An Interview with

ALEXANDER MCKENZIE

OH 185

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

Alexander McKenzie Interview
13 March 1990

Abstract

Following a brief overview of his background, McKenzie discusses his connection with the ARPANET project at Bolt Beranek and Newman (BBN), including an indirect influence on hardware selection, his early role as a project generalist responsible for handling user questions and giving presentations about the ARPANET, and the running of the Network Control Center (NCC). McKenzie addresses why the NCC was set up, how it expanded, his view of the computer utility concept, and his interactions with the IPT Office, the other members of the group at BBN, and the rest of the community. The interview ends with an evaluation of the impact of the ARPANET. 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.

ALEXANDER MCKENZIE INTERVIEW

DATE: 13 March 1990 INTERVIEWER: Judy O'Neill

LOCATION: Cambridge, MA

O'NEILL: Let's start with a brief sketch of your career, your educational background, and what you did before starting on the ARPANET at BBN.

MCKENZIE: Okay. I received my undergraduate education at Stevens Institute of Technology. I received a Bachelor of Science in mathematics in 1962. I went to Stanford University as a graduate student, originally in the mathematics department as a Ph.D. candidate. But I quickly switched into the computer science department, which at that time offered a terminal master's degree. I received my master's degree in, I think it was April, of 1964. But, actually, I finished my work there in December of 1963. My first job after graduation was with Honeywell in New England. I wanted to come to New England, so I took a job there. I worked in the Fortran compiler department of Honeywell Information Systems until February of 1967. I started at BBN in February of 1967. The first few years I was here I asked for, and was mostly assigned to, projects that would give me a good opportunity to learn about innovative things that were going on at BBN. My very first job was the documentation of the time-sharing system that BBN had built for the National Institutes of Health and the American Hospital Association, so I could learn about time sharing. Then I participated in a number of consulting, and study, and design projects until early 1970.

In the beginning of 1970, I took a six-month leave of absence from BBN, and my wife and I went to Europe and traveled around - camping - for six months. When I came back in November of 1970, ARPA was putting a lot of pressure on the sites at which ARPANET nodes were located to begin using them. I was not involved at all with the initial deployment of the network in 1969. And I, obviously, was not involved with the network in 1970 until around November. But when I came back, as I said, I think ARPA was putting a lot of pressure on the sites to begin using the ARPANET, and so they were beginning to take it seriously, and they had a lot of questions. In the meantime, it was still in kind of a shaky state. The development team was working very hard to try to make it more solid, and get the bugs out, and so forth and so on. And they were being interrupted a lot by these site representatives who were beginning to take it seriously and wanted a lot of questions answered - and they were typically the same questions

from each site. So Frank Heart asked me if I wouldn't spend a few weeks learning glib answers to questions about the ARPANET; understanding who did what so I could go talk to the right people if there were more detailed questions that came up. And then act as the interface between the development team and the community, giving the information to the community that it needed, and shielding the developers from constant repetitions of the same questions, so they could keep up with their development job. As a result, I became BBN's first generalist in the ARPANET. I had never actually built any single piece of it, but I knew at least something about every aspect: from the hardware design, to the interface design, to the software design, what the limitations were in the through-puts, and all that sort of thing. So I ended up being involved in the thing called the ARPANET Network Working Group, which did the design of the host level protocols, and going to a lot of meetings. And eventually, a few years later, after 1972 when the ARPANET had its public demonstration and the whole world suddenly began to believe that packet switching was real, and organizations all over the world wanted to learn a little bit more about it, I was in the enviable position, for somebody who likes to travel, of being the generalist who could go and talk about sort of any aspect of the network. So I had a very nice, pleasant existence for a couple of years, traveling around and lecturing, and doing some consulting to other organizations that wanted networking, and so forth. It was a very pleasant position to have been in, to have gotten tapped as the generalist.

In the early seventies, around 1971 or early 1972, I took responsibility for running the Network Control Center, as it was called at that time, here at BBN, which was responsible for trouble shooting, and diagnosis and repair of the ARPANET. I had that job until, I think, about mid 1976, when the ARPANET was formally transitioned to the Defense Communications Agency for operation. Roughly at that point, I stopped being responsible for the Network Control Center, which is now called the Network Operation Center, and went off into other projects that were not really directly related to the ARPANET.

O'NEILL: I have questions about a lot of what you just said. Let's start back at Stanford. Who were you working with, and what kinds of projects were you involved in in the computer science department?

MCKENZIE: The department was just forming at that time. The head of the department was George Forsythe, who

was a numerical analyst. I was not so interested in numerical analysis. John McCarthy had just come from MIT, and he was teaching courses in artificial intelligence and LISP. The thesis topic that I picked for my master's thesis was work on the Burroughs 5000 compiler. The Burroughs 5000 was a machine that was really engineered to execute ALGOL, and its compiler was written in ALGOL. ALGOL was basically the only language that was supported by the machine. It was a very elegant machine. (When I got out of graduate school, I seriously considered going to work for Burroughs just because I liked them - they seemed to be, for a computer company, very innovative. But I wanted to come to New England more, so I ended up taking a job with Honeywell instead. I didn't know what DEC was; at that time it was too little. They made the PDP-1, but that wasn't so well known or popular.) I did that work under a man named Harold VanZoren who had come from Carnegie Mellon a couple of years before. I don't think that he ever became a full professor there. There was, at that time, really nothing that I'm aware of in any university dealing with computer communication. Maybe some little bit of stuff was going on at MIT, but I'm sure nothing else was going on anywhere. So my background there really had nothing to do with what I ended up doing with the ARPANET.

O'NEILL: So you went to Honeywell for a few years, working on Fortran, and then to BBN. Why did you come to BBN?

MCKENZIE: I wanted to be in a smaller company that was doing more innovative things than what was going on at Honeywell. Honeywell tried to find its market niche in building faster, cheaper clones of IBM equipment, by taking advantage of the technology progress since the IBM stuff was developed. It was interesting for a while to be working there, but after a while it didn't look like I was going to learn very much. I wanted to be somewhere that was more on the cutting edge. BBN seemed pretty good.

O'NEILL: You mentioned that when you first went to BBN you had the intent of learning about the various areas that they were involved in. Did you learn about the ARPANET involvement at that time?

MCKENZIE: BBN really was not involved in the ARPANET in 1967. The RFP was written in 1968, and the contract was awarded in the beginning of 1969. As I said, most of the first year I was here I spent learning about and

documenting the so-called hospital computer system, a time sharing system that BBN had developed with NIH and American Hospital Association funding. It was an operational time sharing system for a specific purpose. It was one of the first in the country. So it was a good opportunity to learn about how time sharing worked. Time sharing was still pretty new at that time, and it looked like the wave of the future, and there was a lot of communication in it, and real time system design that just didn't exist in the things I had been doing before. I don't remember the chronology of other things that I did immediately after finishing that documentation project. But one of the things that I did was to work with Dave Walden, who was a relatively new employee at BBN also, on the specification of a mini-computer based data management system, or database system. This was an idea that we had that mini-computers were going to be coming along and were actually powerful enough to do real work. And we went and looked at a lot of different mini-computers that were just coming to market to try and pick one as the platform on which one could build a data management system. We ended up, after this study, picking the Honeywell 516 computer, primarily because it offered the widest range of I/O and peripheral equipment from a single vendor, of all the different mini-computers. Others were faster, or had more interrupt registers, or better precision arithmetic, or different other advantages. But the advantage to the Honeywell 516 was that it had a complete line of disk drives, and tape drives, and terminal multiplexers for terminal access, and all that sort of thing. So it looked like, if we were going to actually build such a system, it would be the most trouble-free platform. Ironically, that study was a very strong input into the selection of the Honeywell 516 as the first IMP, even though the first IMP didn't need or want mass storage or terminal multiplexers or tape drives or any of those things. (Laughter) But we had looked it over, Dave and I, and it looked like a pretty good platform for some other purpose. And when we had to make a quick decision - when Dave was involved in writing the proposal for the ARPANET - he and Frank took a cursory look to see if there was any disadvantage to the Honeywell machine, and noticed the proximity of Honeywell, and the fact that it could be ruggedized, those were factors, too. But I think a really important factor in the selection of that machine was that we had picked it for a completely different reason, for a completely different project. I was involved, I guess, in the early days of the ARPANET from that point of view.

O'NEILL: Had you been involved with the 516 when you were at Honeywell?

MCKENZIE: No. Actually, I don't think Honeywell bought the company that made the 516 until just about the time when I left. It might have been even after I left - a company call Three Cs, I don't remember what the three 'C's stood for anymore.

O'NEILL: In November of 1970, when you first started formally getting involved with the ARPANET, how did you go about learning about it? I have heard that there were some reports that had been written, the IMP Guy Notes, or something like that?

MCKENZIE: The IMP Guys Notes, guys, yes.

O'NEILL: Why don't you tell me how you learned about it.

MCKENZIE: I think mostly it was just by talking to the developers. Walden, [Will] Crowther, Severo Ornstein, Bob Kahn, Frank Heart, a guy named Jim Geisman, Truitt Thatch, I think was still here. A number of people who were involved in the early development and have since left. But probably mostly it was Walden, Crowther, Ornstein, and Kahn.

O'NEILL: You mentioned going out and talking to people, being a kind of general consultant for their networking interests. Were they interested in hooking into the ARPANET? Or were they interested in doing networking in their companies?

MCKENZIE: Different things in different time periods. In the first few years of my association with the ARPANET, beginning late in 1970 and going through most of 1972, most of the world seemed to believe that packet switching couldn't work. So the only people I was talking to, at that time, were the universities primarily, and some non-profits, that ARPA was telling, "You have to join the ARPANET. Whether you like it or not, you have to do this." That might be a little over stated, but I think there was pretty general reluctance most places to go spend scarce resources on connecting up to this new network concept, because I think mostly, people believed that it wouldn't work. There

were a few enthusiasts at many different places, but by and large the attitude was "Maybe somebody else could do this first. Then, if it's not a failure like we think it's going to be, then maybe we could do it." I think at UCLA, where Kleinrock, and Crocker, and Postel, and Cerf, and Ari Ollikainen, and a number of other influential people were, that wasn't the attitude. They were enthusiastic. But I don't think there was much enthusiasm anywhere else. So in those early years, what I was doing was talking to the people, and trying to help them figure out how to get their computers connected up to the ARPANET, because ARPA had told them that they had to do it. And being involved in the Network Working Group where we were trying to hammer out the ideas of what are now called protocols, and protocol layers, and protocol stacks and all that sort thing. We were all feeling our way because there wasn't any body of current expertise or knowledge or anything. Steve Crocker, especially, took a major leading role in that whole development area. There were also a number of other people who were authors of the early RFCs, the first 50 RFCs or so. Some of the names I can remember are a guy named Kalin from Lincoln Labs, Eric Harslem and John Heafner who were at Rand, there was somebody at University at Utah [Barry Wessler]... Mike Padlipsky from MIT (from the Multics project). We all had long, vigorous arguments about what the right layering principles were. We had a concept that layering had to be done, but exactly what the right way to do it was all pretty unclear to us.

O'NEILL: Was BBN sponsoring these trips that you were taking, or was it ARPA?

MCKENZIE: It was paid for by ARPA as part of our support activity for the ARPA Network.

O'NEILL: To your knowledge, had anyone else been doing that before you became involved as a generalist? Were the developers doing it on an ad-hoc basis?

MCKENZIE: I think Kahn did it a little bit, and I think Crowther did it a little bit. Maybe Walden also. But I think that the first concept of the network didn't really go beyond the idea, very much (maybe in Larry Roberts' mind it did). But in BBN's mind our job was to deliver some hardware boxes and make them work. The idea of the Network Working Group, I think, was Roberts' idea to say, "We don't want to put all our eggs in the one basket of one contractor, to tell everybody how to do everything. Everybody has to be involved because there are good ideas everywhere, because

ARPA wants people to feel like they have a stake in the outcome." For a lot of reasons both technical and political, ARPA had the idea that even if BBN was developing the hardware of the network, the IMPs, that the host organizations which were ARPA research sites, in order to ever get them to take the network seriously, and really connect to it, and really begin using it, and really begin finding applications, that they had to be involved too. So ARPA was funding people from each of these sites to get together to work on these issues. I think I was probably the first person from BBN to do it in a sort of formal, recognized, that's-your-job kind of way. But there probably were meetings during 1970, before November, that somebody went to. Since I wasn't here, I'm not really sure exactly how that took place.

O'NEILL: When you were talking to these people, they were not too enthusiastic about joining into the ARPANET. How did they communicate that to you? Was that just the general feeling you had? Did they actually say to you, "Well, we're not really very interested in this, but we have to do it."?

MCKENZIE: No. No one ever said that. That is the feeling I had. In a sense it's sort of the not-invented-here feeling that one can often pick up when talking about any project to someone who wasn't really involved in it. In the first year, it was mostly BBN who was involved in building the ARPANET because the nodes had to be designed and built and installed before anybody could really begin to do anything, and other people had their own agendas. In a sense, one of ARPA's motivations for building the network was so that they could share computer resources. Instead of buying everybody a PDP-10, they could buy half that number of PDP-10s, and that was not exactly highest on the priority and interest list of the sites who were going to end up not getting the PDP-10s - to share them with some other faraway people. High on their interest list was having their own. So I certainly don't want to imply that anybody was out actively trying to sabotage the idea of building the network. And, probably, at every site there were one or two enthusiastic people. But mostly, for most of the people at any given site, it was at best neutral and at worst maybe a little antithetical to their own interests and aspirations. Furthermore, a lot of people believed it just wouldn't work, that it couldn't be done, or it couldn't be done anywhere near on time - that it would be five years late and then wouldn't work very well. So there was a big acceptance barrier to overcome.

O'NEILL: When you were getting involved in the networking area, were you familiar with any previous work?

MCKENZIE: No, absolutely not. I knew nothing about communications per se in November of 1970 when Frank asked me to do this - it was outside my experience.

O'NEILL: You mentioned that BBN sort of had the approach of, "We have hardware to deliver, and make it work." Do you remember how you felt about networking at the time? Some people have mentioned that they did have an idea that it was going to really change things and change the way people worked. Other people have said it was just a fun, interesting project, or whatever, but didn't really think about any bigger implications.

MCKENZIE: I think I was in the second category. I didn't have a vision of how this was going to change the world. I came back from a long vacation, I was asked to do this as my next job, and I did it. As I was doing it, it was fun and interesting, but I don't think that I really had an image. I've never been very good at looking into crystal balls. I'm not good at it now, and I don't think I was good at it then either. So I didn't really have any idea of how this project would go on. The project, after all, was originally to build a four-node network that might ultimately expand - if it was wildly successful - to 19 nodes. That was the contract. Even the people who say that they could see how it was going to change the world, I'm not sure if they imagined that there would be thousands of packet switches around the world in twenty years. Maybe they did. I didn't get that feeling, but maybe they did. I didn't though, absolutely. (Laughter)

O'NEILL: Do you recall the approximate size of the network when you started working on it?

MCKENZIE: I think it was in the vicinity of ten nodes.

O'NEILL: After your initial phase of involvement as a consultant and whatnot, you mentioned that you started working more on the operations. When you were working with host sites, what were some of their complaints about actually using the network, as opposed to earlier when they were hooking into it?

MCKENZIE: I don't honestly remember any serious set of complaints. It was intended that this be a research project. That's why ARPA funded it, as a research project in communication. And, also, to support other people's research in distributed systems, and things like that. Once a set of host protocols were defined to allow connections, such as the NCP (Network Control Program), and Telnet (the remote logging protocol), and file-transfer (I guess that came at about the same time, for transferring files around), and there were a few implementation of those, it was remarkable how quickly all of the sites really began to want to view the network as a utility rather than as a research project. So a lot of the energy that I spent, and other people that I worked with spent, in the Network Control Center was trying to make the network run, if not with the reliability of the power company or the phone company, at least more reliably than anybody's computer center, because the biggest set of complaints were if it went down, if there were problems.

There is a different kind of complaint that I guess is important and worth mentioning. BBN, I guess in 1972, deployed the first Terminal IMP, or TIP, the thing that provided direct access for dial-in terminals, or direct-wire terminals. (In those days, a really fast terminal meant one that went up to 300 baud.) BBN tried to take the view that we were the TIP contractor, and our operational responsibility was to make the IMPs and the TIPs work properly, and that is where our responsibility ended. And then if some site had a modem bank connected up to the TIP, that it was the site's responsibility to make the modems work. We didn't have anything to do with them, we didn't buy them, we didn't install them, we didn't have any authority over them. But what would happen was that ARPA, for example, situated a TIP at the MITRE Corporation in Washington, and paid MITRE to get a bunch of modems to connect up to it, and then ARPA told all of its small contractors in the Washington, DC area, "If you want access to the ARPANET resources, call these telephone numbers for these modems," and they did the same thing at NASA Ames in northern California, in the San Francisco area. The users didn't know anything about networking particularly, or whether a problem was a modem problem, or a network problem, or a TIP software problem, or a host problem. They only knew they were supposed to be doing their work by dialing this number and then typing some magical stuff and pretty soon they'd be connected to a time-sharing system where they had an account. And if they weren't connected to it, then it was a problem, and they wanted one person to call up. They didn't want to be given the run around. So a lot of energy, ultimately, went into the Network Control Center assuming responsibility, or doing its best to try to

assume responsibility, for monitoring things over which we really had no authority. We added to the instrumentation in the Network Control Center, which started out being instrumentation...

TAPE 1/SIDE 2

MCKENZIE: ... to communicate with the IMPs, with the packet switching nodes, to find out how they were doing. Through that we were able to see how the circuits that connected them, the lines, were doing. Because if two IMPs on a line reported that the line was down, then we could fiddle around and try to figure out why and make a call to the appropriate telephone company to get it fixed. But we began adding instrumentation and tools and equipment in the NCC to monitor critical hosts, the ones where a lot of people from other sites had their accounts.

For example, USC ISI ran a complex of time sharing systems on PDP-10 TENEX systems. And ARPA gave accounts on those time-sharing systems to many small contractors, so the NCC began monitoring the status of those hosts, whether they were up or down, and to the extent that it was possible whether their protocol machinery was running properly, and so forth, so that we could call them up and say, "Hey, it looks like you're having a network problem; let's work together." Even though we had no authority over ISI, we couldn't tell them to do anything, but we could cooperate, and the users could feel like they had one point of contact to call up and ask about troubles. If ISI, for example, was having hardware troubles, we tried to get them to call the NCC - they usually did, they were better than many - to give us status information about when they expected to be back up, so that the users anywhere could call the NCC and we would be able to say that there was a problem. Similarly, with the modems connected to the Terminal IMPs, the TIPs, eventually we got lists of all the telephone numbers for all these modems that ARPA was giving out the telephone numbers for, and we got a WATS line and an autodialer, and we had a program to call up every modem in the middle of the night, and keep trying it until it made five tries or three tries or something, or until it got through. Then in the morning, we printed out a list of all the modems that we weren't able to access. Then the NCC operators would call those modems up to see if they got the carrier tone. (Because the autodialer was pretty dumb, we had to have a level of manual trouble-shooting after that.) Then we would call up the sites for modems that seemed to be broken, and say "Would you please get this fixed, or take it out of service, or busy it out (essentially take the phone

off the hook) so that users won't get stuck there and not be able to get access to the network." And we had to do that mostly because the users really needed a single point of contact. It's absolutely no use to some geophysicist in Washington who's trying to do something in the DARPA seismic monitoring program, to say, "Well, you can call the Network Control Center, if that doesn't work call MITRE and ask about their modems, and if that doesn't work call ISI, and ask about their computer." Nobody would work that way. I wouldn't work that way myself. So a lot of the activity in the NCC over those few years, from roughly 1972 or 1973 until 1976 when I left, was spent expanding our coverage. Even though our authority didn't expand, our responsibility expanded quite a lot.

O'NEILL: How large was the NCC at this point?

MCKENZIE: Well, we ran seven days a week, 24 hour a day coverage. We did that, I think, beginning in 1972 or so. By 1976 we had, I think, three-person coverage during the period say 10:00 a.m. to 5:00 p.m. EST, and two-person coverage from 7:00 a.m. to 10:00, and from 5:00 p.m. until 8:00 or 9:00 to cover the West Coast prime-time, and one person coverage in the middle of the night and on Saturdays [and Sundays]. So I can't remember anymore what that turned into, but that was something like eight full-time operators, or maybe it was six. Reporting basically to me as part of the NCC operation after 1973 or so were the network operational programmers; people who were just doing bug fixes, and trouble shooting, and fire fighting, and continued improvements in the software but not inventing something brand new. There were probably three or four programmers - full-time people. That seems about right, three or four people or that equivalent, doing those kind of activities. I think by 1976 the network was somewhere in the vicinity of 80 to 100 nodes, and probably between 20 and 30 TIPS, the rest of them were IMPs. I think at that time we had Honeywell 516 IMPs, Honeywell 316 IMPs and TIPs, and a few Pluribus IMPs. (The Pluribus was the multi-processor that we developed as a second generation switching engine for the ARPA network.) So there were a lot of different kinds of hardware, and the programs on the different kinds of hardware were not completely compatible, so we needed programmers who were experts in each one and we had to have a lot of operator coverage and it was a pretty busy place most of the time.

O'NEILL: Were these programmers writing tools as well?

MCKENZIE: Yes, when they had time.

O'NEILL: How was the Network Control Center set up? Who's initiative was it, as it expanded? Did ARPA have direct influence on what was done, or was BBN taking care of it?

MCKENZIE: Oh, ARPA had direct influence on anything they wanted to have direct influence on. (Laughter) There's no doubt about that. In the very earliest days of the network, the Network Control Center was just a Teletype connected to the IMP at BBN. Well, in the very earliest days I guess you couldn't say that there was any operations. It was telephone coordination with the sites, because the first four sites were on the West Coast, and Utah. BBN's IMP was IMP number five, and when that was installed then there was a way of getting reports from the other IMPs back here, and there was a logging Teletype that was connected to the IMP and people, I guess, informally went and looked at it every now and then. If somebody called in from a site, and said "Hey there seems to be a problem," then somebody would explicitly go and look at the Teletype and try to figure out what was going on. I don't remember off hand (although there are papers written about it, and they're probably in this pile), exactly when we began to have a staff of operators and people who were explicitly responsible for watching the Teletype. When it got too painful to watch the Teletype, some of the programmers hacked together programs to run first on our PDP-1, and later on one of the PDP-10 TENEX systems, to provide computer assistance for sorting the logs and filtering out extraneous information and storing it on disk instead of on the paper scrolling out of the terminals, and so forth. But that all evolved kind of gradually. By the time that I took over responsibility, which I think was in mid 1972, there were operators. I don't think that we were running 24 hours a day then, but there were operators and then it grew. I think mostly the operation grew at BBN initiative, not ARPA initiative - probably under ARPA voltage - saying "It's not going well enough. You have to do better." But not really telling us exactly how to do it.

O'NEILL: Was there a separate contract for the Network Control Center?

MCKENZIE: No. It was all part of the development contract, really.

O'NEILL: So it was just one of the duties of BBN as part of the development.

MCKENZIE: I'm not positive of that answer. That's my recollection, but it might've been different. We've gone through phases with DARPA of having every little task have a separate contract, to mushing everything together into one so-called umbrella contract that covers a lot different tasks. I don't exactly remember what it was like in 1972, 1973.

O'NEILL: That's fine. You mentioned the view of the people using the network of it being a utility, as reliable as any other utility. Did you view it as that, too? Were you trying to run it like a power utility, for instance?

MCKENZIE: I was, yes. I don't think we were ever nearly as reliable. We felt good if we had the IMPs up 98 or 99 percent of the time. That would be an abysmal record for a power utility. Yes, I think that I was given the job of running the Network Control Center because I kept advocating that we ought to be trying to run it as a utility. I probably started advocating that before it was appropriate, but once it became appropriate (and that did eventually happen) probably that's why I got picked to do that.

O'NEILL: What was your idea of a utility?

MCKENZIE: Well, I think basically the idea that it should be... that we should take it very seriously if anything was broken. That we should put more emphasis on the operational aspects and less emphasis on... If one had to make a choice at any given instant which should be the subordinate goal, doing network research, or providing network connectivity to people who were doing other kind of research - I guess my view was that we should subordinate the goal of doing network research to the other goal. That was not a universal viewpoint, by anybody, even by DARPA. We reserved a couple of hours, I think it was on Tuesday mornings, I think it was 6:00 to 8:00 east coast time, to do network software testing, and we reminded everybody every week that that time shouldn't be viewed as reliable time; still people didn't like that. But we needed to have some amount of time, but we tried to really make sure that the rest

of the time we weren't doing experimental things, and we were taking it seriously when something broke instead of treating it casually, I guess. Again, the expansion out to doing monitoring, and taking responsibility for monitoring things that were not really not under our authority, like the modems and the hosts, was, I think, another aspect of trying to make it more utility-like and provide a central point of information and complaint for users, and that kind of thing.

O'NEILL: That sounds like the difference between a development shop versus a production shop - was it that kind of idea, where you are offering a production service?

MCKENZIE: Yes, I guess that's fair. I mean, it was always a mixed goal. It was a mixed goal from the very idea of creating the ARPANET up until long after the operation was turned over to DCA. At the beginning, the emphasis had to be on doing network research because nobody had ever done this before. After the network was turned over to DCA, their emphasis had to mostly be on running an operational facility because that's what they are supposed to do. During the intervening six years, it gradually slid across the spectrum from one to the other. So I don't think that there was ever a point, while DARPA was running the network - or even during the first few years that DCA was running it - when one could say, "Research isn't appropriate in this environment." It was still appropriate to be doing research, but that got more and more controlled and in some sense subordinated, pushed off into the inconvenient times (the middle of the night, the early hours of the morning) in order to provide support for the other research programs that were not research in networking, but research in something else that DARPA was funding people to carry out using the network as a tool. Eventually when it became real clear it was operational they turned it over to DCA, and then it was mostly operational. Although when DCA had assumed control, there was still research that needed to go on. Eventually they split the network into the ARPANET and the MILNET. The MILNET was supposed to be really operational, and if new software had to be tested out it was supposed to be tested out in the ARPANET because that was more expendable - allowed to be a little less reliable.

O'NEILL: One more question about the utility service idea. I think of utilities, like power utilities, as being accessible to anybody.

MCKENZIE: Well, we were only running the ARPANET for ARPA. They could make it accessible to whomever fit their policy guidelines or whoever they wanted to let on. When I say utility, I'm not trying to talk about BBN trying to commercialize the ARPANET. We commercialized the technology by establishing Telenet. When I talk about running it as a utility, I mean running it as a utility for the people that ARPA said were allowed to be using it but supporting them. ARPA was really sponsoring a lot of research, distributed research, where, in order for the researchers to make progress, they had to have access to this communication tool. It was not only our goal, it was ARPA's goal to not have the people come back at the end of their one or two year grant period, and say "Well, you know, we really couldn't make any progress on our research," whatever it was, let's say natural language understanding, "because we needed to communicate between the east and the west coasts to make any progress and we didn't have any facilities for doing that because you told us we had to use the ARPANET and it was no damn good." ARPA was pretty clear to BBN that they didn't want to hear that kind of thing. So we tried in the NCC to make it so they didn't hear it.

O'NEILL: During these years (my timing might be off) but didn't the network expand internationally at that point? Did you have international sites?

MCKENZIE: We were a little bit international from pretty early on. If I dig out a set of maps, I have a pretty complete set of maps.... Maybe I should let you borrow those, too.... But we had a link from Washington to University College in London, and from there to Kjeller, Norway. That site at Norway was a site where... I don't know the details... where seismic information was collected, at least, if there weren't seismometers or seismographs or whatever they are there, they were scattered around in Norway. They were getting the seismic information to a place in Washington that was sponsored by DARPA. This was part of, basically, the nuclear monitoring program. The ARPANET was viewed as a way of sharing or speeding up the transmission of data or sharing a line. I can't remember exactly what the pre-existing situation was, whether the seismic site in Norway was sending tapes by airmail with seismic information to Washington; or whether they had a leased line that came directly here. But they didn't need all of a leased line, and so if the ARPANET provided that line, then they could send all their data over it. And in addition to that, some

researchers that DARPA wanted to collaborate with in Oslo and especially in London could share the use of that transatlantic circuit. So the ARPANET went there. We also had a satellite link to the University of Hawaii. That's not really international, but it's out of the continental U. S. That's really all that there was internationally until after I left the Network Control Center, I believe. DARPA sponsored, in the late 1970s, the development of a project called SATNET that BBN was the contractor for, which provided a shared access satellite channel that reached to Norway, London, somewhere in Germany, and the United States, for a while. Eventually it expanded to Italy, but I think that was actually in the mid 1980s. But that post-dated DARPA's involvement in the ARPANET per se.

The switch from the Network Control Program, the original host protocol (these days, in ISO terminology, the level 4 protocol) to TCP/IP took place in the early 1980s. Although I think the basic ideas of TCP were in an article that Cerf and Kahn published in 1974, it did not really become the official protocol of the ARPANET, or of the INTERNET, or of DARPA sponsored networking, however you want to say that, until the early 1980s. And I don't even think there were experimental implementations before around 1978. The details are in that paper that Walden and I wrote. Up until the time when TCP/IP replaced NCP, there really had to be only one network. There wasn't the concept of going from one network to another. Up until that time, in the early 1980s, there were hosts on the ARPANET; there weren't IP gateways as we have today. And up until that happened, I don't think there was substantial international expansion. So I think the international expansion really dates to very late 1970s, or early 1980s.

O'NEILL: Did the few international sites that you had while you were running the NCC cause any particular problems? Were they any different?

MCKENZIE: Yes. The line from the U. S. to London ran at 9.6 kilobits per second, as opposed to 56 kilobits, which was the service that we had everywhere else in the network. That always meant some extra tinkering with the parameters of things to make it work properly. But I don't remember any really severe problems. Also, it was a stub circuit, so called. The whole idea of the ARPANET in reliability terms was that there would be two ways to get from anywhere to anywhere else, at least two ways. So, no single failure of a node or a circuit would isolate, would break any communication (except communication that terminated at the failed node, of course). Both the overseas sites in

Europe and the connection to Hawaii were single tail circuits. So if they broke, then that site was isolated. Now, in most of the ARPANET when a circuit began to give errors, when there began to be errors on a circuit, we wanted to declare the circuit to be unusable and get it fixed as rapidly as possible before the percentage of errors got too high. Let's say if there were 10% errors on a circuit we wanted to declare it unusable. Otherwise, information would be getting queued up for that circuit because stuff would be lost in the errors and have to be retransmitted, so we couldn't get so much through-put, and there would be more delay. So the algorithms in the packet switches were biased towards rapidly declaring a noisy circuit to be unusable and only slowly accepting it back. That meant that the singly connected paths to Europe and Hawaii, since we wanted the IMPs to be running the algorithms everywhere. Another part of the goal of the ARPANET was that everything was the same, so nodes didn't really need to have special information about their special circumstances; they would all follow the same algorithms. That meant that a line that had 12% losses to Hawaii kept being declared down and breaking the communication. And since we were using a host-to-host protocol that threw all state information away once there wasn't a connection, then that meant it was very hard for the Hawaii people and the London people and the Norway people at times when the lines began to get noisy, it was very hard for them to do any work. Ultimately we had to adjust those parameters, and take special note of those things. But that was a peculiar problem - not so much because they were European or Hawaiian, but because there was only one way to get to them. That interacted badly with the algorithms that we were using everywhere else.

O'NEILL: You mentioned Telnet. Were you involved in that at all?

MCKENZIE: Telnet the protocol, or Telenet the company?

O'NEILL: Telenet.

MCKENZIE: Telenet the company. I was not particularly involved with that, no. BBN, as you probably know, established Telenet as a wholly owned subsidiary. BBN got venture capital investment that diluted BBN's ownership, but I think BBN retained more than 50% ownership in Telenet, although I'm not positive of that. BBN

recruited Larry Roberts from DARPA to be the president of Telenet. He wanted it, basically, to be independent of BBN so he could pursue his own development. They set up their offices in Washington rather than here. Some BBN programmers who'd been working on the ARPANET went to work for Telenet, or transferred to Telenet, as it were, for a period of time. I think some of them stayed there, and some of them came back here. Eventually, after a few years, Telenet was sold to GTE which had enough money to invest in it, to really make it fly. It could afford bigger operating losses for the first few years than BBN was willing to afford, so they could build up the network into a real national presence, which is what they needed ultimately to get subscribers. I, personally, had nothing to do with the Telenet venture.

TAPE 2/SIDE 1

O'NEILL: I wanted to ask you about how you interacted with the developers who were working on the ARPANET. You mentioned that you got most of your information initially from them.

MCKENZIE: The developers at BBN?

O'NEILL: At BBN, yes.

MCKENZIE: I think that we were all one team, really. Our offices were adjacent, and we talked all the time.

O'NEILL: So you were part of that group?

MCKENZIE: Yes, we talked all the time, yes. Sometimes when they wanted to do more researchy things, I had to hold them back, and they were undoubtedly irritated with me. Sometimes we had to work out compromises between changes to the software that were interesting for research and changes to the software for more reliability. But, by and large, we were all compatriots. They didn't work for me, and I didn't work for them. We all brought our strong opinions to the table and worked out compromises.

O'NEILL: Were you involved in whatever sorts of design reviews or whatever they would be doing as well? Did you interact only in issues of the Control Center?

MCKENZIE: We were all one team. Yes, I was involved. I was involved with things that they wanted to do, and they were involved with things that I was interested in proposing. There wasn't any real split between the developers and the operations people. There just were not enough of us to make that necessary. We all worked for Frank Heart whose management style, at that time at least, was to be involved in everything that was going on and have the final say on anything that was happening in his department. So we couldn't drift very far apart. (Laughter) He wouldn't let us.

O'NEILL: How large was the development group by 1970? I know initially there was a handful, four or five people.

MCKENZIE: By 1970?

O'NEILL: Well, that's when you first started getting involved.

MCKENZIE: I don't know. We were six people, or something like that. Six or seven or eight, not too many. In fact, I don't thing the project was ever, up until it became DCA's job, I don't think there were ever - leaving the NCC operators out - I think that there were probably never more than a dozen people or so - maybe that might even be high - working on the project at any one time.

O'NEILL: At this time was Bob Kahn considered part of this group? I know that for a while he worked in another group. What was your relationship with Bob Kahn?

MCKENZIE: Bob was part of the group when I became involved with it. My perception was that Bob really was quite instrumental in having BBN bid on the original proposal for the ARPANET, and that he was always pretty

upset that having convinced BBN that BBN ought to bid on the project, that he was not put in charge of it. He never, in my opinion, while he was here really accepted the idea that he shouldn't be in charge of it, or that he wasn't in charge of it. So there was often a lot of tension between Bob and Frank Heart over project direction and what we should be doing exactly; in management issues there was a fair amount of tension. Frank was in charge, and Bob wasn't. I wasn't in any way in the middle of that, but I think that the tension that I sensed between Bob and Frank meant that any of us who did what our boss, Frank, told us to do were maybe a little bit estranged from a really close relationship with Bob.

O'NEILL: Did you consult with Bob on technical issues?

MCKENZIE: Oh, Bob was involved in all the meetings, and so forth, and so on. It wasn't like he was relegated to the outer darkness. He was a participant, too. He had strong opinions just like everybody else. And this is a group of people with very strong opinions, this set of people: Frank Heart, Dave Walden, Severo Ornstein, Will Crowther, Bob Kahn, myself. We all had very strong opinions, and none of us were at all bashful about letting them be known and lobbying for what we thought was the right thing to do. Bob was part of that just like everybody else.

O'NEILL: Did you have much interaction with Bob Kahn once he went to ARPA?

MCKENZIE: No. I did not. Again, it was my sense that because of the tension between Bob and Frank that I mentioned earlier and because... There were two groups at BBN at that time that were involved in computer-related things. Frank was in charge of one of those groups, the Computer Systems Division, I guess it was called. And the other was the Computer Science Division and just at the time that Bob went to ARPA, I can't remember who was in charge of that, it might have Jerry Elkind, or Dan Bobrow, or Ray Nickerson; it was one of those three people. But the key communications guy in that group was Jerry Birchfiel who is still at BBN. My sense when Bob went to Washington was that Bob steered all the DARPA new activity that got done at BBN to Jerry and his group and that he steered it away from Frank and his group. So beginning at that time that Bob went to DARPA, I think the amount of interesting R&D work that came to BBN didn't changed, but where it went changed. And as a consequence, I

really had very little interaction with Bob.

O'NEILL: Who at DARPA did you have the most interaction with?

MCKENZIE: At the beginning, Larry Roberts. The program managers for the ARPANET work, the people who were essentially my boss at ARPA, were Larry Roberts, followed by Steve Crocker, followed by Craig Fields, followed by Steve Walker.

O'NEILL: Was there a significant change in how you interacted with ARPA based on these individuals or for other reasons? Can you characterize the change over the time that you were working with them?

MCKENZIE: Each of those people had a different personal style. My biggest memories of changes are of personal style, not of what ARPA was interested in doing. I guess Larry Roberts and Steve Crocker were both much more interested in the research aspects of networking. Craig Fields was the first of the ARPANET program managers who began the trend, I guess it would be fair to say, towards ARPA being more interested in the operational issues. Craig probably was split 50/50 in his mind between operational importance and research importance. Steve Walker, I think by the time he became the program manager of the project, the plans were already being developed to transition it to DCA. His focus on the ARPANET was much more on its utility aspect, its operational aspects, rather than its research aspects.

O'NEILL: How did the people at ARPA let you know what they wanted? Did they call you on the phone and say, "We want this or that"? Did they work through Frank Heart? How did that work?

MCKENZIE: Well, there were a variety of ways, of course. The operational aspects, running the NCC, was only one part of what we were doing for DARPA. So their interactions with whoever was running the NCC, in each case, was only a part of their interactions with BBN. In terms of interacting specifically with NCC concerns, I think mostly each of those people dealt with me directly, mostly by electronic mail. With Larry, I think it was almost exclusively by

electronic mail. With Crocker, it was probably 80% electronic mail. Probably the same with Craig Fields. Steve Walker was maybe down to 50% electronic mail, 50% phone conversations. But I don't know, these are recollections from fifteen, sixteen years ago. They are probably not too accurate.

O'NEILL: Let me just follow up on that anyway. At what level of detail were they giving direction, or discussing things?

MCKENZIE: There were three big classes of things that we did in the NCC. One, just dealing with short-term operational difficulties: firefighting, fixing things when they broke. The second class of things was short-term planning. ARPA decided that a new packet switch, a new IMP, is going to be installed at Carnegie-Mellon University; how is that going to go into the network, on what date are the telephone lines going to be available, on what date is the packet switch going to be available, how many interfaces have to be supplied with it - delivery and installation and de-installation kinds of issues. The third kind of thing was the IMP program should be changed in order to provide such and such a feature, do more of this or that or something else. On the first front, the fire-fighting, if something was really broken, if some researcher was complaining to DARPA that they couldn't get anything done, then ARPA would issue ultimatums. Ultimatums is a little too strong, but not too much too strong. (Laughter) You know, "Get this thing fixed!" or "Tell us what's going on here! I hear this has been bad for a week. What are you doing about it, and what's the prognosis? And when are we going to have it ready?" and "So and so is complaining that he can't get his work done, so what are you going to do about it?" Sometimes they would want very detailed reports of everything we were doing to deal with the problems; sometimes they would just want reassurance that we were paying attention, we knew what the problem was, and we had a plan. It probably depended on how loudly the person complained to them. (Laughter) On the planning stuff, ARPA had a contract with the Network Analysis Corporation to tell them what the topology should be for new network nodes. And sometimes they just told us that, and that's what we did. Sometimes we argued back on some grounds - we argued back... I argued back. Sometimes I would say, "That doesn't seem very practical, because it requires us to change the configuration of the IMPs, the number of interfaces that they have, or something, too much" or "we don't have that many interfaces in stock, and if we pick this other plan, then it would be easier to do." So there would sometimes be a

lot of detailed discussion about that. Sometimes not. On the third thing, on software directions, I think probably it was rare that there was anything really detailed that we exchanged information about. There it was really more broad brush kind of direction.

O'NEILL: Did you deal directly with any of the other contractors? Like the Network Analysis Corporation, or was that just through DARPA?

MCKENZIE: It was really a pretty small community of people. We were all invested in the success of the thing. We all talked to each other. Maybe we didn't talk enough, but we certainly didn't... It was a hierarchical arrangement in the sense that we all took our direction from ARPA; and if we didn't like what some other contractor was doing, we would go appeal to ARPA to tell them to not do it. And sometimes they would do that, and sometimes they wouldn't. But almost always based on a real analysis of the issues, not arbitrariness. But we also all talked to each other. We talked to the Network Information Center at SRI, we talked to the Network Measurement Center at UCLA a lot, we talked to Network Analysis a fair amount also. I won't say that we all saw eye to eye or were really good friends with each other at every instant in time, but we communicated a lot - all these groups, I think.

O'NEILL: Were there set meetings that ARPA called with contractors, or was this interaction pretty much informal and as issues came up?

MCKENZIE: Well, there were meetings of the Network Working Group, continuously, every few mo nths for a long time, into the mid 1970s. Exactly when those things got so big and unwieldy as to stop having them, I don't remember - maybe after Crocker's tenure as the program manager. I don't think we had Network Working Group meetings under Fields. All of us would be at those meetings, so we would get business done - other business that wasn't protocol work done on the side, other kinds of coordination and information exchange. ARPA, for a long time, had Principal Investigators meetings. I think it was at Lake Arrowhead in California. I never went there, but Frank Heart went there, and Len Kleinrock [Network Measurement Center] went there, and Howie Frank [Network Analysis Corporation], and Doug Engelbart [Network Information Center]. I don't know how long those kept going. I have the

feeling that was through 1972 or 1973. So at a pretty high level, the important people from each of these sites got a chance to bend each other's ears and offer complaints or congratulations, and you know, stay in touch. My interactions with the other centers were pretty much informal: e-mail, a fair amount; telephone, sometimes; occasional visits, usually in conjunction with some other trip, to some other place, for some other purpose.

O'NEILL: Were the meetings of the Network Working Group sponsored through ARPA?

MCKENZIE: Yes.

O'NEILL: Did they decide where they would be, and when? Or was it more a matter of them providing the money?

MCKENZIE: I don't really know. Steve Crocker was sort of the first chairman, and I think he stayed chairman through Roberts' tenure at DARPA and his own, maybe. Whoever was in charge at ARPA didn't actually pick the time and place. At least they probably were perceived by all of us to have veto power over the time and place. The first Network Working Group meetings happened a lot in conjunction with the AFIPS National Computer Conferences, the Spring Joint and the Fall Joint. There were sub group meetings of people working on let's say, the Telnet protocol or the file-transfer protocol at the site of one of the active people in the group. In general, everything happened either on the east coast or the west coast because almost all of the ARPANET sites were either in California or along the east coast. Typically the meetings would ping pong back and forth - first let's say, at MIT, and next at UCLA, next at MITRE, Washington, next at SRI, just to spread the travel load around in a more or less equitable manner. If there was some group that because of the interest of its participants had 80% of its members on the west coast, then it would probably have 80 or 90% of its meetings on the west coast.

O'NEILL: I was trying to get a feel for how much the Network Working Group was internally directed and how much was ARPA directed, in terms of what they were doing.

MCKENZIE: Well, ARPA certainly set the goals. I can certainly remember when Larry Roberts came to some

Network Working Group meeting, where we were haggling about the details in the Network Control Protocol, the NCP protocol, and he said, basically, "You guys have spent enough time arguing about this. I want implementations in three more months. We can't stand to have any more delay, too many things are being held up." So whether we had implementations in three more months, or whether it was really five, I don't know. But the message was pretty clear. It was not "Do it this way" but "Come to a decision, damn it, you've been arguing long enough. All the ideas are out on the table; now you've just got to decide." So in that sense, ARPA provided goal direction, time-table direction, things like that. The program managers at ARPA didn't really have the time to tell people how to do it. They really much more aimed at funding people who were smart enough and self-motivated enough to recognize what the problems were and go solve them. There was, in my experience, zero or close to zero micro-management of any ARPA program - either at BBN or anywhere else that I ever heard of. ARPA's objectives were to find people that they thought were sufficiently smart and sufficiently motivated and give them a ball and let them run with it.

O'NEILL: Were you encouraged by ARPA to disseminate information to write articles, to attend conferences?

MCKENZIE: Absolutely. I was a relatively a low-level person at that time. So I never talked to anyone at ARPA from a contractual point of view about how much they were allowed to encourage us. But I always had the sense that they wanted people who had learned stuff about communication to go spread that information to the rest of the country and the rest of the world. We certainly reported in our reports to them that so and so on this contract traveled to this conference and gave this paper or met with these people, and so far as I know, there was never any question that that was the right thing to do. Maybe there was some specific instance some time as to whether it was appropriate. But as a general rule, not only at BBN but at all of the other centers that were involved, I really think that ARPA tried very hard to get everything published. They wanted us to put all of our quarterly technical reports from BBN into the National Technical Information Service. And at some point the National Technical Information Service balked and said, "We don't want these quarterly technical reports. That's not what this is for. This is a repository for final reports," or something, and ARPA said, "No, you've got to take them because there are lots of people who are interested in this. And the research is being reported in these quarterly reports, and lots of people want to know about it, so please take them." And to BBN they said, "Keep on sending them," and we did. (Laughter) So I think

they tried very hard, all of those people that I mentioned whom I interacted with from Roberts through Walker, to get the information out to the public.

O'NEILL: So these were explicit, or was it implicit?

MCKENZIE: Some of it was explicit, and a lot of it was implicit. It was encouragement by example, too. Roberts, I think, was the chair of the first session of five papers at the 1970, I guess it was the Spring Joint Computer conference. I think Kahn was the chair of one on packet radio a few years later. And there was one in 1972 that I think was chaired by somebody else from ARPA, or ARPA picked the person or nominated him or something like that. So it was a lot of leadership by example in that domain. DARPA people were trying to get the word out, asking contractors to come with them to meetings with other parts of the government, for example, and other universities, and so forth.

O'NEILL: Can you describe a little bit of the collaboration process that went on in writing some of these articles? How did the articles actually get written? They all seem to have multiple authors, and it's not clear to me how they got produced.

MCKENZIE: I think mostly they got divided up into sections, and the different people wrote... that's one step ahead of myself. Some of the people are good writers and some of the people are not-so-good writers, of course, as in any sample of the engineering population. There were a lot of people who were good writers on this project: Frank, Walden, myself, Kahn. I don't remember whether Severo Ornstein was a good writer or not, Crowther's not such a great producer of text. But I think, generally, everybody wrote sections, and then they all got pasted together, basically. We used electronic document preparation systems. The information got passed around in electronic mail. Everybody read everybody else's stuff, and criticized it, and made corrections. Then somebody - either an editorial assistant, a staff member whose only job was to do documentation kind of things, or somebody from the group - it might've been me, it might've been Dave Walden, might have been John McQuillan, who was also a very good writer - would take all those comments and put together sort of a final draft and pass it around for everybody to look at one

more time. There would be a few corrections, and it would be done. So I think in those papers where there are lots of authors, that probably the work was pretty much done by those people. It wasn't just, you know... There aren't authors on those papers, even the multi-authored papers, that are just there for the honor of being authors. Everybody worked hard on it, mostly by dividing up into sections, I think.

O'NEILL: Was that pretty much the way the completion report was started as well?

MCKENZIE: Yes. Walden had something, I had something, Frank had something. I don't remember if anybody else worked on the completion report or not. Probably McQuillan...

O'NEILL: He is listed.

MCKENZIE: We divided it up into sections. My section was, looking at these notes, mostly about applications and protocols.

O'NEILL: At that time, did you think that a book would be published about the experience on the ARPANET? The big three-volume report? I had the impression, or maybe it even says in there, that that was going to be published.

MCKENZIE: I think different people had different ideas. BBN had a contract from... The completion report is a thing that has to be produced by ARPA. It is ARPA's report to its own managers. As I understand it, it can't be done by a contractor, that's not something ARPA is allowed to sub-contract out. We had a contract to gather information from which they would write the report. We, of course, were interested in saying lots about this wonderful thing that we built. And we gathered a lot of information, and we were all pretty wordy, Walden and Heart and I, and McQuillan. And we produced this thousand-page thing, whatever it was, order of magnitude of a thousand pages. We sent it to ARPA and said here's the information we gathered. (Laughter)

TAPE 2/SIDE 2

MCKENZIE: I know that I thought that this was just going to be a report. A few people would want copies, and most of the world wouldn't care. I know that Frank Heart felt that the ARPANET was just as important to the country as the Manhattan Project. And since there were books written about the Manhattan Project that appeared on the *New York Times* best seller list, there ought to be a book about the ARPANET development that ought to appear on the *New York Times* best seller list, and this was going to be the basis for that book. I never particularly shared that enthusiasm of Frank's, but that was his goal. For Walden and McQuillan, I'm not sure what they thought. Frank and I were probably at the two extreme ends of opinion about this. (Laughter) Whether they were at one of the ends, too, or in the middle, I don't know.

O'NEILL: In the article that you and Dave Walden have written for the *Encyclopedia of Telecommunications*, you talk about the impact of the ARPANET on local area networks. In particular, you mention a fellow named Robert Metcalfe, and also a group at MIT that was working on Token Ring Nets that had evolved from the ARPANET. Can you give me a little more detail on that? Did you specifically work with Robert Metcalfe?

MCKENZIE: Sure, yes. Metcalfe was involved in the Network Working Group from early in its history, maybe not from the very beginning, but from early in its history. He wrote his doctoral thesis about packet communication. He was really heavily involved in watching the ARPANET grow, and watching its protocols develop, and seeing what the strong and weak points were, and all that sort of thing. So he was clearly influenced, as he thought about communication, by the ARPANET experience. The group at MIT included Ken Pogran who is here at BBN, and Dave Clark who is still at MIT and pretty important, and I don't know who else was involved - probably a lot of people. All were involved in watching the ARPANET grow and in communications research related to the ARPANET. They had the idea of Token Ring Networks, which ended up being the foundation for the company which would be called Proteon. It's not, of course, that the ARPANET was the only place where people were thinking about rings and bus networks, and so forth. There was a lot of other work, also, going on: the Cambridge ring in England; there was a ring network that Farber was involved with at University of California, Irvine. There were other local area network approaches and all of us, I think, in the community... People who were in the community who were at the level of

Dave Clark, or Bob Metcalfe, or myself, we all really knew each other and what everybody was doing, and so forth. There were international workshops that ARPANET people went to to talk about their ideas, and other people went to. I remember in particular there was a NATO workshop on networking in Brighton, England. I don't remember when it was, it must have been 1974, 1975, where we were talking about the ARPANET, and somebody gave a presentation on the Cambridge ring, probably Metcalfe may have been talking about what turned into the Ethernet or maybe not; I may have that wrong. So there was a lot of exchange of ideas. It wasn't really a very big community of people who were involved in these things, and I think mostly we all knew each other to some extent.

O'NEILL: Was Metcalfe at a host site?

MCKENZIE: Yes. He was at MIT, I think. Was he at MIT or Harvard? I don't remember where his thesis is from, maybe he worked at MIT. Let me look.

O'NEILL: So Metcalfe was at MIT, his thesis was done at MIT.

MCKENZIE: Yes, eventually he went to Xerox Palo Alto Research Center, and they were connected to the ARPANET also. They were doing some DARPA research. He was there and they were connected to the network at the time that he really developed the Ethernet.

O'NEILL: Other than these people that you've mentioned, do you see any technical carry-over from the work on the ARPANET with the local area networks and the Ethernet?

MCKENZIE: I think that the whole idea of protocol structuring in layers, which permeates all of the data networking stuff that people are doing now, was developed in the Network Working Group. Maybe it was developed by Roberts and fed as the ground-level ideas into the Network Working Group. But certainly it was refined there in its first refinement. In terms of technology - chips, or electricity, physical layer kinds of issues - no, I don't see much connection between the ARPANET and local area networks. I think the main contribution, in a sense, was the proof

that the idea that you could have a network of dissimilar computers - not in the same room - doing useful work together. Or even maybe, just the concept that you could have dissimilar computers doing work together is the strongest contribution. Let me say it a different way. In the early days of the ARPANET, the ARPANET *was* a local area network. When we installed the IMP at MIT, MIT connected to that IMP, and then through it to each other, the MULTICS system, the ITS (Incompatible Time-Sharing) system, and a few other computers. Although they were physically adjacent, essentially, they had never been connected together before. Nobody had the concept to see that they could be connected together, and that it would be useful, potentially, to connect them together - because they all worked differently. The IMP provided a common interface, a common kind of interface. If they all matched that interface, then they'd all be connected. If they could talk over the ARPANET to California, they could also talk to the guy next door. So, in the early days, at a few sites at least - and MIT was a striking example, I think Lincoln Labs was another striking example, and BBN was probably another one, and UCLA where they had a Sigma 7 time-sharing system and an IBM 360 91 was another example - we were, we the ARPANET, were the local area network. Not very fast, and not very efficient, by today's standards. But the contribution to local area networking was more the idea that it was possible, and might be useful, rather than the technology. The technology doesn't really carry forward very well at all.

O'NEILL: We have covered what I wanted to ask you about. I want to open it up and ask if there's anything you want to add, any general comments you would like to make about the ARPANET, or your involvement in it, or ARPA in general.

MCKENZIE: I think that the example of the way ARPA worked in those days, in the 1970s - not only in communication but in a lot of other programs - was pretty effective in terms of cost and return on investment. That the model of some really bright people in the ARPA office, with some particular goal in mind, something they thought would help the country, or the DoD, being given the authority to go out and find the smartest people they could find who were interested in working on that problem, and then giving them relatively free rein to do some research or some development, produced pretty remarkable results. And I know that there has been fraud and abuse in government and in contracting and so forth, but it seems to me that the kind of rules and regulations that there are

now, that are attempting to prevent that, really make it very difficult for the government to get the same kind of power out of its research dollars these days as it was able to then. I know it's hard to find a balance between accountability and free rein, and these days the government approach seems to be more on the side of accountability and less on the side of free rein. But I think that a lot is being lost. I remember a discussion that I had with Jon Postel back in, I don't know, the early 1970s. We were talking about travel to a Network Working Group meeting. And he was saying that he really would like it - maybe it was a subgroup meeting - if it would be some [one] place rather than some other place, because the government accountants were getting university accountants to not want them to spend so much on travel. And his comment was, "I know that in my contract there's only a fixed amount of money. I want to use that money the most effective way possible. I don't know why they think that, when my salary is paid out of this, that I'm going to waste money on unnecessary travel. I want the money for my salary; I want to keep having a job, but sometimes the travel is necessary. They don't seem to understand that that's the case. They think it's just a boondoggle." I think there's a lot to that attitude. That, yes, we spent a lot of time in Network Working Group meetings hammering out ideas, but something really good came out of it. Yes, BBN wasn't required to file monthly fund expenditures reports and justify every decision on the basis of cost effectiveness. We could pick the Honeywell 516 because it was OK - it was an OK processor for this. It would do the job, and it wasn't, maybe, the most cost effective and we didn't have a subcontract bidding plan in effect to give everybody the opportunity to bid, and thereby put another three months of delay into the schedule - we could just decide. I think that was pretty valuable. If you're talking about multi-million dollar projects, it's probably more important to pay attention to those things than it is in tens-of-thousand or hundreds-of-thousand dollar projects. But, even there, there's probably things where better results would come to the government if they could find a way to shift back a little bit more towards accountability for the results, rather than accountability for the pennies. I think that DARPA in the 1970s did a really good job for the country in that way. It was a joy to be associated with the ARPANET project. It was fun. It was challenging. And I think it was good for the country. It's not so easy to find that mix now, and I think regulation is a big part of it. Thank you for the opportunity to have this soapbox. (Laughter)

O'NEILL: Well, thank you.

END OF INTERVIEW