

An Interview with

VINTON CERF

OH 191

Conducted by Judy O'Neill

on

24 April 1990

Reston, VA

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

Copyright, Charles Babbage Institute

Vinton Cerf Interview
24 April 1990

Abstract

Following a brief overview of his background, Cerf describes his involvement with the ARPA network, including his work for the Network Measurement Center while a graduate student at UCLA, and his relationships with Bolt Beranek and Newman, Robert Kahn, Lawrence Roberts, and the Network Working Group. Other topics include: various influences on the development of the TCP/IP protocol, IPTO funding while he was at Stanford University, his decision in 1976 to become a program manager for networking projects at IPTO, and the military use of IPTO networking projects. 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.

VINTON CERF INTERVIEW

DATE: 24 April 1990

INTERVIEWER: Judy O'Neill

LOCATION: Reston, VA

O'NEILL: I would like to start off by talking about your background. I know that you were at Stanford and then at UCLA where you got your Ph.D. Why did you go to UCLA?

CERF: Why did I go to UCLA? I did my undergraduate work at Stanford, and the recommendation at Stanford was that you should go someplace else to do graduate work just to be exposed to a different set of faculty. As it happened I decided that I wanted to work for a while after I finished my undergraduate degree at Stanford before I would go on to do any graduate work. Part of the reason was that I was a little tired of school, and part of it was that I wasn't quite sure what I was really going to wind up doing. I had been a math major, but I had gotten very involved in computing early on in my undergraduate career. So I wound up working for IBM in Los Angeles. Originally I was supposed to go to help install the first 360/91 - this would have been just in 1965, middle of 1965 - to help install the 360/91 at Los Alamos. I was all ready to pack up and go to Albuquerque when they ran into a problem with a time-sharing system that they were running out of the Los Angeles Data Center in downtown L.A., called Quiktran. It was probably one of the first interactive interpretive FORTRAN time-sharing systems. People could actually dial up and use - IBM 1050 remote terminals - it wasn't just teletypes. So I wound up being the systems engineer for this time-sharing system, running on an old IBM 7044, the predecessor to the 7094. I did that for a couple of years, and it whetted my appetite for operating systems because I had to dive into the middle of it to understand how to make changes to the system. Here I was, a spanking brand new systems engineer, all of twenty-two or so, and I actually got to make some changes to the time-sharing system and to the operating system. I wrote them all up and sent them off for approval, and lo and behold, they said, "Yes, it is a good idea. Go ahead and do it."

O'NEILL: Was this sent to IBM?

CERF: This was IBM. So I did that for a couple of years, and then I realized that I didn't know enough, I didn't have enough theory to really do a credible job in computer science. I said, "Well, it is time to go back to school." By this

time I had met my wife in Los Angeles, so we kind of settled there. By a series of very fortuitous circumstances - an old high school friend of mine, a guy named Steve Crocker, was already doing his graduate work at UCLA - in fact, he had done his undergraduate work there. So Steve set me up with his thesis advisor, a guy named Gerald Estrin. Jerry took me on as a research student at UCLA doing work on what was then called the Snuper Computer project. This was a DARPA-funded thing, and it was a very unfortunately named project, I suppose, because this was in the late 1960s and the Vietnam War was happening and people became more and more sensitive to what you could do with computers. The Snuper Computer was simply a machine that was capable of monitoring the performance of another machine in a relatively non-invasive fashion. In particular they had a 360/75 installed on the campus in the campus computing network. We got a Sigma-7 computer from what was then Scientific Data Systems and later was purchased by Xerox and turned into Xerox Data Systems. And we used a bunch of high impedance probes to go right into the display lights of the 360/75 so that we could monitor things like the program counter, which I/O channels were in operation, and was the system in I/O wait or was it running?

O'NEILL: So you were actually monitoring the display panel?

CERF: Yes, we actually monitored the 360/75 display panel! So I was writing software for the Sigma-7 system to capture data about what was going on in the IBM machine. Somewhere along the line DARPA decided it didn't want to pursue that project anymore, and as that phased out, the ARPANET project popped up and UCLA turned out to be the first node of the ARPANET. It was responsible for doing network measurement. This was under the auspices of Leonard Kleinrock who even at that time was one of the foremost queuing theorists in the country. He had a very special interest in store-and-forward networking because he had written his dissertation on the subject and published a book. It was something like, I'm not going to generate the title correctly, but it was something like "Network Queuing, Stochastic Flow and Delay," something like that. [*Communication Nets: Stochastic Message Flow and Delay.*] It was published in 1966 by Dover Press when he was still at MIT. So Steve and I, although we continued to be guided by Jerry Estrin as far as our dissertation work was concerned, actually took research assistantships under Len Kleinrock's sponsorship and began working on the ARPANET project.

O'NEILL: So Steve was still there, he hadn't gone off to DARPA?

CERF: Steve was still there, he hadn't left yet. He didn't actually leave, I don't think, until at least 1970. At least that - oh, it must have been beyond that. Let's see, Steve was still at UCLA up to 1972 even, because he was part of the big project to demonstrate the ARPANET in October of 1972. So he was still involved and still at UCLA. He must have left shortly thereafter though. I'm not real clear on the exact dates, but I remember a lot of things happened at the end of 1972 where people just shifted gears. I left UCLA and went to teach at Stanford in October of 1972. Bob Kahn left Bolt, Beranek & Newman, after having taken a year of his life to get this big demo of the ARPANET done, and came to DARPA. And Steve Crocker, I think, must have left UCLA somewhere in the same general timeframe and gone to DARPA. I don't recall exactly how long he stayed, but he left before I came to DARPA, so he must have left in 1975 or something like that. Because I didn't come to DARPA until fairly late in 1976 - something like October of 1976.

O'NEILL: Just to back up a little, when you were at UCLA you didn't have any involvement in any sort of networking until this Snuper Project was canceled, and you then got an assistantship with Leonard Kleinrock and started into the ARPANET?

CERF: That's right. In fact it was really almost amusing in one respect. It wasn't any fun to have the Snuper Computer Project canceled, but what was amusing is that during that time - this would have been around 1968, early 1968 - Steve Crocker and I were consulting for a company called Jacobi Systems, which was a firm in the Santa Monica area. They basically did a lot of systems work for the Navy. My recollection is that they were doing automatic testing of electronic equipment, of computer based equipment, and they would write programs that would step these things through their cycles. They used UNIVAC machines. Well, somewhere along the line they decided to branch off into another area and build small-scale time-sharing systems. And they wanted to have some interactive time-sharing facilities built for them. Steve and I got interested because one of the principals in that company knew Steve from earlier periods when Steve had worked with him on the campus. So Steve got us both involved in designing and building a small dual processor time-sharing system made out of Interdata machines - that company doesn't even exist anymore, I don't think. We used an Interdata-3 and an Interdata-4, one to do I/O and one

to be the actual time-sharing engine. Steve designed the time-sharing engine, I did the editors and I think he wound up also doing a big chunk of the FORTRAN interpreter. So we actually got a system up and running. Why is this relevant? It's relevant because at the time that DARPA sent the RFP out for the ARPANET in June of 1968, Jacobi Systems was one of the bidders. So I built a GPSS simulation of the packet network, and we used that as part of the input to the proposal that we sent. Well, we didn't win. I think we were close to it, I don't remember whether we were in the last four or the last six or something. There were something like twelve or fifteen bidders, and we sort of lasted for a while, but compared to Bolt, Beranek & Newman we probably weren't quite as credible. The nice thing was that even though Jacobi Systems didn't win the bid, Steve and I got directly involved in the ARPANET because Len Kleinrock had gone on with his proposal to build the Network Measurement Center. So we got involved anyway, it didn't matter which way we went.

O'NEILL: How large of a company was Jacobi Systems?

CERF: My memory is a little vague, but I think that it couldn't have been more than fifteen to twenty people. In its heyday it must surely have supported about that many. Not a gigantic operation by any means, but they had a big Univac 1108 sitting right there on the premises. I remember being impressed by a thing called a Fastrand that had, my recollection is, 100 megabytes of storage, which was a lot in those days, on this enormous spinning cylinder. It looked like a stainless steel roller that you would use to cold roll steel, and when the thing fired up, you sort of expected the whole building to twist. It was just an enormous piece of equipment, a huge heavy thing. They had to put reinforcing steel plates on the floor to spread the load and everything else. You would never want to use it on a ship - the first rough sea and the thing would blast right through the wall of the vessel. So by late 1968, or maybe even August or something of 1968, I think the awards were out, and it couldn't have been very much after that because they delivered in less than a year. The first Interface Message Processor (IMP) was delivered around Labor day - a couple of days before Labor Day in 1969.

O'NEILL: The contract was not actually out until the end of 1968.

CERF: December, okay. So it was like a nine-month gestation, so to speak. In any event Steve and I and several other people at UCLA were very involved in the development of the first protocols that were used to run on the computers that were on the outside of the ARPANET itself.

O'NEILL: Now at this time at UCLA you were going to connect the Sigma-7, right?

CERF: That's right. In fact, we had the first node on the net. A guy named Michael Wingfield built the first - other than BBN - 1822 interface (BBN document number 1822) to link the Sigma-7 up to the ARPANET IMP. So I can recall... in fact, that is how I met Bob Kahn. Because he was at UCLA at that time - this was 1969, now, after September, probably in October or thereabouts. Bob came out from BBN to begin making measurements of the behavior of this packet-switching network because they had not done very much in the way of test and experimentation. They had a few IMPs at BBN, but they hadn't hooked them all up over long-haul lines. So by the time Bob came out, I think there were four nodes on the net, maybe only three. Ours was first, SRI International was second, Santa Barbara was third, and Utah was fourth. I don't have any more of my old data, but I know there were at least three nodes, and maybe Utah was up by the time Bob came out.

My job was to write the software that was to drive traffic into the network and capture information about how many packets were lost, and what the delays were for getting echoes back, what the total bandwidth was that we were able to achieve, and whether or not multiple sources of traffic on multiple logical links increased the total bandwidth or not. We wanted to be able to watch the congestion control mechanism that would decide to switch routes. For example, if you went from UCLA up to SRI there was one direct link, and then there was another one that went by way of UC Santa Barbara - it was a triangle. An interesting question is, if you launched a lot of traffic from UCLA towards SRI International, at what point would the system decide that the queues going straight up to SRI were too long to sustain any further traffic and would decide to switch over and send the stuff by way of Santa Barbara. We could watch it, like a firehose, going back and forth. We actually generated more traffic between UCLA and SRI than could have been carried on a single channel as a result. Because the single channel was a 50 kilobit circuit of which you could get about 40 kilobits of useful traffic if you drove it as hard as you could. But by switching back and forth

between the two lines, I think we actually generated in excess of forty kilobits. I don't remember any more whether it was - it may have been a fairly dramatic number. It may have been as much as seventy or so kilobits of total bandwidth because the traffic was literally being bifurcated. Of course, that strategy was probably okay for such a simple topology, but when things got more complicated it wasn't clear that bifurcating traffic and filling up two queues necessarily gave you that kind of result. The sum of the two might not actually be realized as end to end traffic. It depended on the topology at the other end.

O'NEILL: I'd like to understand more about the environment at UCLA and how you and Steve Crocker and...

CERF: ... Jon Postel, and later Steve Crocker's brother David was involved in this thing too, among others. Bob Braden, Robert Braden is another one of the old-timers in this bunch. Incredibly, they are all still involved in networking. We were permanently infected by packet-switching. (Laugh)

O'NEILL: So the group consisted of about ten graduate students?

CERF: Well, let's see. There were fewer than that, really. At UCLA, the sort of UCLA Mafia sort of consisted of Steve Crocker, who was sort of the ringleader; Jon Postel, who ultimately became the editor of the RFC series and has been ever since; Bob Braden was involved because he was at the computer center. He was the guy in charge of the 360/75 at the time, and it later became a 360/91. But he wrote the first NCPs - that was the host protocol at the time. He wrote the first NCPs for the IBM machine all in assembly language, God help us. There was a guy named Bill Naylor, William Naylor, who I think came a little later on in the picture; a guy named Louis Nelson who had a lot to do with the development of the time-sharing system, in sort of maintaining the environment of the Sigma-7 computer; a guy named Charles Kline who was a programmer, an undergraduate, in fact, who did a lot on the operating system of the Sigma-7. We did a homegrown experimental time-sharing system at the time. Another one at UCLA who worked on the ARPANET effort was Holger Opderbeck.

O'NEILL: Was DARPA funding the time-sharing work as well?

CERF: You know, I honestly can't quite recall how that all worked out. I think the answer was that we had a choice of running a crummy batch operating system or writing our own time-sharing system, and we decided that we would write our own time-sharing system. We needed it really. We needed multi-process capability in order to run the measurements because we needed multiple processes driving the system.

O'NEILL: So SDS only provided a batch operating system?

CERF: Well, they had a... what did they call it? There was a very basic operating system that was provided to us and it wasn't really an interactive time-sharing system. It's funny how hazy that is to me, I mean I was down in the guts of that operating system at the time. What did we call it? I even remember modifying the thing to put up a header when it came up... something about the zipadeedoodah release. I made a bunch of modifications to that system in order to run the network measurement system. Because we ran it at first just as a single-thread job. But it didn't give us anything like the kind of flexibility that you get from a full-fledged time-sharing system with virtual memory. The thing was configured to use virtual memory, but it didn't use it very effectively. So what are you going to do? You have got a bunch of computer science graduate students who really want to attack some hard problems. In fact, there was a basic time-sharing capability that was very primitive, and we ran that way for a long time. As I recall it took us... the machine came in 1967 in the summer, and we were probably still running, boy the timing is really not clear. But I think we were still running the basic operating system that I was maintaining until at least 1970. We really didn't come up solid on the time-sharing system, we called it the SEX, Sigma EXperimental, Operating System. Of course, everybody loved the title of the user's manual - SEX User's Guide. Who did I leave out? There were some other younger people who were doing some basic programming jobs for us, but they weren't so terribly involved in the networking aspect of it.

O'NEILL: How involved was Kleinrock in determining what you were doing? I don't quite understand how that worked.

CERF: Len was actually fairly relaxed about what was going on. He made us do project reviews from time to time so he understood what was going on. He had a lot to say about the kinds of measurement functions that were needed. In fact he had several graduate students who were doing both queuing analysis and modeling and direct measurement in the network [Len Kleinrock's early students on queuing analysis of networks included Gary Fultz and Jerry Cole]. I spent quite a bit of time working with those graduate students to modify the software, the measurement system, to run experiments, and capture data. So he was very involved in that. The part that we were more left on our own about was the specifics of the operating system design and the evolution of the protocols for the ARPANET itself. In fact, if you've talked to Steve Crocker... you should, because he can relate to you firsthand his feeling that we were stepping in to try to devise this set of protocols expecting... you know, we were just rank amateurs, and we were expecting that some authority would finally come along and say, "Here's how we are going to do it." And nobody ever came along, so we were sort of tentatively feeling our way into how we could go about getting the software up and running. In the longer term, Larry Roberts was very insistent that this intrepid bunch of graduate students - not just at UCLA but at other sites like MIT, and Utah, and SRI, and UC Santa Barbara - get their rear ends in gear and actually make decisions about the protocols and get them instantiated and get them running in all the operating systems. In some instances Larry objected to the designs that we had done and made us go back and start over again. So in terms of the protocol work I would say that Len was probably less directly involved in that than he was in the measurement work, which of course was what he proposed to do anyway. And Larry Roberts had rather more to say about the protocols. Bob Kahn was the one who was most deeply involved in the guts of the network itself, from our perspective. He was the guy who would come out and camp out for weeks at a time to understand the behavior of the system and why did it do X, Y and Z. He would force it into certain modes which he predicted would fail, and other people wouldn't believe it would fail this way, and he kept saying, "But it will" and nobody would listen until he came out to UCLA and he and I just blasted the thing, and it would break. And he'd be happy because he proved that certain things wouldn't work right.

O'NEILL: In fact, the reason for asking about Kleinrock and how involved he was to try to get at how involved you actually were with DARPA and with Larry Roberts. It sounds like there was quite a bit of personal interaction.

CERF: Oh, always. A great deal of it.

O'NEILL: So you weren't going through the PI or anything like that?

CERF: No. The style at DARPA has always been that the program managers are quite intimately involved in what's going on. As an example, a practice that was not followed regularly after his departure, but while Barry Wessler was at ARPA he would put on an annual retreat of two or three days' length for graduate students. It was a remarkably good idea. First of all he would be there, personally involving himself in listening to what graduate students were reporting. He would interact with them on matters of substance. But more important, from my point of view, the graduate students themselves got a chance to know each other and learn what they were up to, and what the various groups that were supported by DARPA were up to. This was supposed to mirror the annual Principal Investigators meeting that was also an exciting kind of thing to go to. I can remember going to the graduate student retreat in, I think it was Allerton, Pennsylvania. Steve would remember because I know he was there too. And we met a lot of very smart people, many of whom ultimately became very active in the ARPANET project in one way or another, whose careers... Well, for example, people like Bob Metcalfe, who ultimately went off to invent ETHERNET and had started 3 COM and all that stuff, was one of the graduate students there. Postel and Crocker...

O'NEILL: So you remember those meetings running for at least the time that Barry Wessler was there. That would be... I'm not sure exactly how many years that was.

CERF: Well, it seems to me that Barry came almost concurrently with Larry. And Larry got his arm twisted to come in 1966, or 1967. Is that right?

O'NEILL: It was December of 1966 that he actually went to ARPA.

CERF: Because he had already done some packet switching experiments between Lincoln Labs and SDC. There was a Q-32 to TX-0, or TX-2 I can't remember which. It was a point-to-point link where they tried some things out. Then

Bob Taylor twisted the then director of Lincoln Labs arm to make Larry come to DARPA.

O'NEILL: That wasn't a packet-switching experiment, but it was a networking experiment?

CERF: Not in the sense that... there weren't separate nodes, but it was still packetized data going back and forth. I mean it's not a continuous transfer. They had to send blocks... it was a block exchange and they had acknowledgements and things like that. So the feasibility of inter-computer exchange in a kind of block like mode was being looked at in 1966. Of course, they were totally oblivious, I think, at the time, to Paul Baran's distributed communications report, which at the time was partly classified. It was done at Rand Corporation in 1964. Actually he did the work in 1962 and the final reports came out in 1964.

O'NEILL: As a graduate student involved in networking - were you aware of Baran's work? Had you looked at that?

CERF: No, absolutely not. And what is ironic is that Paul was one of my thesis advisor's graduate students. He was Jerry Estrin's graduate student, but he left before he finished his degree. He didn't get a Ph.D., and he went off to Rand and did all kinds of interesting and inventive things. He reported, as I recall, to a guy named Keith Uncapher, who is now a vice president at NRI. (Laugh) The world is so tangled together it is hard to believe. I was unaware at the time of that work. So my involvement in networking was literally a hands-on experience. We didn't know what computer networks were really, except that there was this thing that BBN had invented called an IMP. Somehow or other we were supposed to try to use it to make computers communicate with each other, make processes talk to each other. So our principal theme was this network was designed to support interprocess communication, and our job was to figure out what kinds of things would these processes want to communicate. Out of that experience came things like file transfer, electronic mail, and remote log in, ("Telnet"). And those three services were sort of the mainstay in most applications up until recently. I would say in the last decade the proliferation of local area networks and the ease of broadcast on them has led us down another interesting path of distributed processing and things like shared file systems and the like, where multiple process communication has been the norm. In a long haul store-and-forward network, it was uncommon to have more than point-to-point logical connections going on; you know, client-

server mode operation. But in the local area nets you get a much more distributed style. Now there is starting to be an appetite to merge the two, and multi-cast facilities are starting to be looked at in a wide area context. So we will probably see some evolution of wide area services that look a lot like local area net services, over time.

O'NEILL: When you were working with Bob Kahn early on, you were writing... maybe you just need to explain to me again. Were you putting in the hooks in the operating system to get the statistics that he...

CERF: I understand. What I did was to write an application program that was capable of configuring itself so that it could start processes running in the machine. It would generate data to be sent out onto the network. I would instrument those processes. They would record how much data they had sent, whether they got acknowledgement back, and the like. Then I would generate statistics that would explain, from the outside, what things looked like. In addition to that, it turned out that the IMPs themselves were instrumented. So I also wrote code that allowed me to control what information would be captured from the packet switches themselves. So there were all kinds of hooks that BBN put in to make it possible for an outside entity to go and pull out snapshots of memory, collect specific statistic packets and so on. They would accumulate statistical information in a table and send a packet to a designated target, usually my Sigma-7, to capture, record, and then post-analyze. So there were two different aspects to the instrumentation work.

TAPE 1/SIDE 2

CERF: What I was saying is that I was capturing end to end statistics as seen from the processes that were trying to use the network, and also capturing statistics from the insides of the packet switches to see what they were seeing. They would see how many packets were queued up on a given line, what the error statistics were, retransmissions on a given line, what the variations were in packet lengths. Of course, I could choose to configure the traffic, the rate at which I wanted to produce packets, how long they were, did they have variable length, did they have fixed length, how many processes would be running concurrently, trying to push traffic out the door, how would they compete with each other, would they both get an equal amount of bandwidth, would three suffer and one win - those were all

questions that we were trying to answer about the way the system would work. If I did anything to modify our own operating system on the Sigma machine it was only to get down and add drivers and the like that could go and talk to the packet-switch through Mike Wingfield's 1822 interface.

O'NEILL: When you dealt directly with DARPA, with Larry Roberts, you mentioned that he expressed disapproval at some of the designs for the host-to-host protocols. How did he do that? Was this a meeting of the Network Working Group and he would come and say, "I don't like this" or...?

CERF: Well, there were Network Working Group meetings that were held not exactly regularly, but fairly often. Sometimes they would happen in conjunction with a conference like the Spring Joint Computer Conference or the Fall Joint Computer Conference. Most of the time I think that Steve Crocker and Larry were the two that interacted most directly. Steve was sort of the putative head of the Network Working Group, and he would report progress to Larry and if Larry was unhappy with something, he would holler at Steve. So a good thing for you to do to get a clearer picture of that whole business would be to talk to both Steve and Larry. My perceptions are that most of the time Steve would take the heat, and he would come back and we would all caucus and figure out how we were going to fix it.

O'NEILL: So you got information through Steve?

CERF: Mostly, yes. Some of us had direct interactions with Larry, especially once the electronic mail system was up and running. But that sort of didn't happen until after we had gotten the basic protocols written in the first place. So once the Telnets, FTPs, and NCPs were written, it was then that e-mail came along and everybody sent messages to everybody else. But before that mostly it was, I think, direct interaction between Steve and Larry.

O'NEILL: Do you remember when you started using electronic mail?

CERF: It was pretty early on, and it is a little bit hazy now, but my recollection is that Larry may actually may have

been the first one to write a reasonable program to parse e-mail. He wrote a TECO program that would... let's see, the way I remember it now is that we used to send e-mail on the Telnet channel of an FTP connection. And the message would be appended to a file with a particular name, like mail.txt or something like that. It would be appended in a certain agreed way so that you could write a program, in this case a TECO macro, that would search down through the text file and pull up just the message itself. There was a header on the front of each message that said how many characters long the message was and you could skip from header to header and it would list the titles and the to and the from fields and the like. So that would have been... that was certainly demonstrable by the October 1972 timeframe. I think the first real e-mail floated around somewhere as early as 1970. But it was not uniform everywhere. Probably the thing that would help remind me most accurately of when we were stable would be to see when RFC733 was written. Because 733 was the original document describing what the format of electronic mail is for the ARPANET. And the author of that was Steve Crocker's brother: David Crocker. There may have been some other co-authors. I think we have a copy of the RFC index here if you want to look at that for some reason.

O'NEILL: I can get hold of that, too. I am not familiar with TECO, can you tell me what that is?

CERF: TECO stood for Text Editor and COrrector. It was actually quite a complicated programming language. It was sort of a programmer's text editor. It had one letter commands. You could type this alphabet soup that would cause the system to move the cursor from here to here, copy this set of characters that moved from this place to this place down somewhere else, count so many lines down, inject this, delete that. It was just a list of commands that you could type, and they were all one or two characters long with a couple of parameters. In fact some of us still use it - he does. You still use TECO from time to time, don't you Bob? [Bob Kahn]

KAHN: Uh, no. There is no machine that I know of that I can get to anymore that still runs it.

CERF: You mean once you got off of ISI-A TOPS-20 systems, you didn't use TECO anymore?

KAHN: I loved TECO, it was and still is my favorite editor of all time.

CERF: Was it a LISP program, or was it just an interpreter basically?

KAHN: An interpreter.

O'NEILL: Was it available on multiple machines at the time?

CERF: Well, it came with TENEX as I recall.

KAHN: Yes, Dan Murphy built it...

CERF: ... at BBN. What, in the late 1960s?

KAHN: It may have been an outgrowth of earlier stuff at Project MAC, I'm not sure. It was just the simplest, most straightforward editor you could imagine. It's like dealing with anything else. If you think about where the electrons go, it is easy to reason from first principles. Most people don't like to do that.

CERF: (Laugh) If you could remember things like "semicolon y" meant "yank a file in from the file system."

[INTERRUPTION]

CERF: We were just talking about TECO. I don't know whether you got Bob Kahn's comments or not.

O'NEILL: I think they will be picked up. I have one other question on UCLA. I have seen some references to something called the UCLA Computer Club. Was that a formal entity, or was that just informally what you called yourselves?

CERF: Well, let me see. I don't actually know very much about the UCLA Computer Club. My name is on a list which occasionally causes newsletters or other announcements of ten-year or twenty-year reunions to happen. I was never a member of the Computer Club, so far as I know. Steve was a very active member. It was an informal group that got enough support, I guess, from the faculty that they had a room of their own. They actually had a room of their own. But I don't think that it was ever chartered; it wasn't incorporated or anything like that. The students were as fascinated as you might imagine with computers, because in those days in the early 1960s they were very new to the campus. UCLA had been especially pioneering in that respect, because it had a Bureau of Standards Western Automatic Computer installed. And that would have been in the late 1950s. Then it got a 7090 around 1960 which was a very big, very fast machine at that time. So it was very much at the forefront of numerical processing and things of that sort. So the computer club was mostly students, some of whom were not even UCLA students, and some were high school students. Crocker may have become involved in the computer club while he was still a Van Nuys high school student.

O'NEILL: Did you work with other people from BBN or from the other contractors?

CERF: Well, it depends on what timeframe you want to talk about.

O'NEILL: I am still talking about the period while you were at UCLA.

CERF: Well, the answer is yes. I had interactions with some other BBN people like William Crowther, for example. He was a guy who worked with Bob and wrote some of the papers on the anomalies in the algorithms of the ARPANET. But I think my real interactions with BBN didn't escalate until after I left UCLA and went to Stanford. That can't be completely right, because BBN was clearly very involved in the operation of the network. It had to have been the case that for me to do any network measurement work I had to be interacting with people besides Bob. But my recollections about which people I dealt with at the time is hazy. There were a set of people there, they called themselves the "IMP guys" - people like Dave Walden, and Alex McKenzie. I don't even remember now whether Bob Bressler was a part of that that early in the game. He was an MIT student at one point then went to BBN and became

very involved in the ARPANET. So I did work with other people, I mean it was inescapable. I had to have worked with them, but my recollections are most firmly of working with Bob. [On further reflection, I worked a lot with Alex McKenzie who was the first Network Control Center manager.]

O'NEILL: In 1972 you went to Stanford. Were you also personally involved in the demo during that time?

CERF: Oh, yes. I had finished my dissertation work at UCLA in March of 1972, and I literally devoted almost all of my time from March until October working on and helping to organize the demo of the ARPANET at the Washington Hilton Hotel. That was where a team of us got to... Again, we took advantage of our work together in the Network Working Group. Here is where people like Jack Haverty, who is now at BBN, and Alex McKenzie, and Bob Metcalfe, Jon Postel. I don't recall... Steve Crocker must have done something, I can't recall exactly what he was doing at the time. Maybe he wasn't - no, he must have been involved in it. Of course Bob managed this whole project, it was his baby. Gee, he actually has a movie. You may have seen it, or there is a videotape.

O'NEILL: I'm hoping to get a copy of it. I haven't seen it yet.

CERF: He also had a whole bunch of slides at one time that he showed a couple of years ago at Interop '88, or something like that. He might actually have lists of who was there participating and packing tools. A lot of the community was there to demonstrate what was available.

O'NEILL: What did you do during that time from March through October?

CERF: I remember doing things like maintaining the mailing lists, making sure everybody knew who was responsible for what, trying to organize what the demos were, who was going to do what demos. I remember that Metcalfe led the team that wrote the manuals for how to use the system - literally - sit down, type the following things, and this will happen. I remember Bob Metcalfe was the one that did that or arranged for it. So my job was strongly on the coordinating side as I recall. And I can remember being up pacing the floor, walking around trying to make sure there

weren't any loose ends dangling in the application model, trying to make sure that everybody knew what they were going to demonstrate, and when. A lot of other people were doing that sort of thing, too. I don't mean to mistakenly argue that I was at some central point of it all. But I certainly was busy, along with fifty other people it seems like, trying to get that whole thing functioning well.

O'NEILL: No, I actually was more interested in what someone who wasn't the central person was doing. I was more interested in what kinds of tasks needed to be done. It seemed like an incredible amount of work.

CERF: A lot of coordination and a lot of reminding people. I remember putting telephone numbers and distribution e-mail address lists together to make sure people could stay coordinated with each other about who was going to be responsible for manning their booths, and when, and how often would we do demonstrations of things, and stuff like that.

O'NEILL: Were you doing this while at..

CERF: I was still at UCLA at the time.

O'NEILL: Were you out in Washington temporarily?

CERF: No. I didn't actually move... In fact, I didn't actually get to Washington until just a few days before the actual conference.

O'NEILL: During the conference itself, what kind of reception did you feel there was to the demonstrations?

CERF: (Laugh) I thought there were two kinds, maybe three kinds of reactions. The first group of people were the diehard circuit switching people from the telephone industry who didn't believe it could possibly work. And they were stunned - because it worked. (Laugh) "It works? It couldn't possibly work!" Then there were the people who

didn't know anything about computer communications at all to speak of and were sort of overwhelmed by the whole thing. Then there were people who had been exposed to the stuff in one form or another and were just as excited as little kids, because all these neat things were going on. My God, some of the demonstrations involved Chinese character displays. There was a gal from MITRE, as I recall, who did that work - Susan Poh. There was a multi-processor air traffic control simulation system called McRoss from BBN that was being demonstrated. We demonstrated our tools for network measurement. People would be doing things and we would show displays of where all the traffic was going and which IMPs were talking to each other and things like that. There were simple demonstrations of e-mail, and text editing, and file transfers, and remote log ins, and things like that. There was access to the online system at SRI International. You could get into the databases there and get documents, RFCs and things like that. It was just a remarkable panoply of online services, all in that one room with about fifty different terminals. It wasn't even uniform. I mean, as far as the equipment was concerned it was a deliberate attempt to interface as many different kinds of terminal equipment as possible. Of course, this drove the poor guys at BBN nuts because they had had to program this terminal IMP, this TIP, to accept interconnections from all these different machines, and they were all RS232 but they didn't all necessarily interpret characters in the same way.

O'NEILL: So there was a lot of interest then and enthusiasm?

CERF: I think so. I would consider it a watershed event in the sense that it made packet switching real to people other than the ones who were involved in designing and building the ARPANET. It really marked a major change in the attitude towards the reality of packet switching. At that same meeting, the ICC 72, is when something called the International Network Working Group was formed, INWG. Some portion of that group spun off a little cadre of four people who ultimately led the charge at CCITT to create X25 - Barry Wessler, Dave Horton in Canada, Remi Despres in France, and another guy whose name I'm not going to remember, it's John something and over time he became one of the major CCITT players for the UK. This INWG group... Actually Steve Crocker did it to me. This must have been very close to the time he was transitioning into DARPA, because he somehow stuck me with the job of chairing that group. So I chaired INWG until I left Stanford, from 1972 to 1976, and ultimately we affiliated INWG with IFIP. And it became IFIP 6.1. There is this Technical Committee 6 for data communications in IFIP, and I engineered along

with the then TC6 chairman a formal affiliation of INWG as IFIP 6.1. Now there's IFIP 6.2, .3, .4, .5 and .6 - I think, at least 6.5 - for various aspects of data communication.

O'NEILL: Okay. How big was this group? There were the four you mentioned.

CERF: The INWG group was about 25 people or so, which ultimately grew to, you know, a mailing list of several hundred. BBN was instrumental in maintaining the distribution lists and pumping the information out. I mean for years Alex McKenzie was the secretary of that group. He probably has every INWG document ever published.

O'NEILL: So in 1972 you took off for Stanford.

CERF: Oh, I should point out that the INWG group was international in scope. It involved people from England and Sweden and France and Germany and Italy. Really very eclectic, strongly academic. It was not the same people who necessarily would go to ISO meetings or CCITT meetings, with some exceptions. They still exist today. I mean, that group is still quite active and still part of IFIP.

O'NEILL: Through that committee did you have interaction with the people at the National Physical Laboratory in the UK?

CERF: Yes, absolutely. In fact, Roger Scantlebury was one of the major players. And Donald Davies who ran, at least he was superintendent of the information systems division or something like that. I absolutely had a lot of interaction with NPL at the time. They in fact came to the ICC 72 and they had been coming to previous meetings of what is now called Datacomm. Its first incarnation was a long title having to do with the analysis and optimization of computer communication networks, or something like that. This started in late 1969, I think, was when the first meeting happened in Pine Hill, Georgia. I didn't go to that one, but I went to the next one that was at Stanford, I think. That's where I met Scantlebury, I believe, for the first time. Then I had a lot more interaction with him. I would come to the UK fairly regularly, partly for IFIP or INWG reasons and also because University College London had become

a node of the ARPANET in the... when did that happen? That must have happened fairly early on - 1973. Because I remember... it must have been very early on. My first kid was born in September of 1973 and it interfered with my going to a meeting in England. That meeting in England would have been in September and that was when the ARPANET node was first demonstrated by Bob Metcalfe from Sussex, University of Sussex. But it was actually demonstrating a node that had been placed at University College of London in downtown London. The guy who has been associated with that site for years is Peter Kirstein. Peter is another person to talk to about early history because he's always stood... In fact he'll be at my house this weekend. He's coming to an IETF meeting, if that means anything to you, in Pittsburgh next week. Anyway, the INWG group was very international, and as the ARPANET branched out and touched parts of Europe, there was a mixing together of INWG activity and Network Working Group activity.

O'NEILL: Was the INWG group responsible for your ideas on the INTERNET?

CERF: Some of it. In fact the guy that... several people had a lot of influence on how the design went. Bob and I spent a lot of time working through various concepts and we wrote that paper in 1974. But I had also a lot of exposure to Hubert Zimmerman and to Louis Pouzin, both of whom had been doing experiments at INRIA, it was called IRIA at the time, on packet switching. They had developed a system they called CYCLADES, and the underlying network was called CIGALE. It was a pure datagram network. CIGALE means grasshopper in French, and the packets would hop through the net. Anyway, Pouzin's ideas on windowing techniques were very appealing to me, and I incorporated them into the initial TCP design. A guy named Gerard Lelann was at IRIA working with Pouzin and came to my lab at Stanford for a year and had a lot to do with the early discussions of what the TCP would look like. So did Bob Metcalfe, it turns out. Metcalfe was at Xerox at the time and in June of 1973 we began working together, Lelann, Metcalfe, and I, on the design of the host-to-host protocol for INTERNET. Eventually Metcalfe got impatient with the rate at which things were going. I was trying to get a large number of people to agree on a set of protocols, and every time you brought in a new player we had to go through the argument again. Meanwhile Metcalfe had five or six guys over at Xerox trying to get the local area nets running. Finally they said they didn't want to wait until this process of agreement and consensus finally concluded, so they went off on a slightly different

tack and invented XNS that took some different choices than the TCP did. And they got it up and running before ours, in fact. Of course in the long run we've... They kept it secret, and that was a mistake, I guess, now looking back. If they hadn't kept it secret, we might all be using XNS instead of TCP. But as it stood, TCP turned out to be the open protocol that everybody had a finger in at one time or another. That is just how it all worked out.

O'NEILL: How did you get interested in the problem initially?

CERF: Well, when I came to Stanford I was busy teaching operating systems and things of that sort. But Bob Kahn, by this time, had started both the packet satellite and the packet radio programs. And specifically he had the problem, "How am I going to make a protocol work through a noisy packet radio network where I don't have telephone lines, and I don't have point-to-point links, and the connectivity keeps changing, and the probability is high that if you radiate something it won't make it to the destination?" The NCP protocols were not ready for that. They assumed that the underlying network would maintain sequencing of packets. Even though the ARPANET was considered kind of a datagram-like system, because you put a label on the front and say here deliver this. Underneath, inside the IMPs, there was a virtual circuit being built, and things were delivered in sequence. And if they weren't in sequence there was something wrong. So the end-to-end protocol (NCP) assumed that. Well, that wasn't going to work in the packet radio network. So when Bob said, "A - how do I get it to work in packet radio nets? And B - how do I get access to the computer resources on the ARPANET from the packet radio net? What do I stick in between?" From that whole set of problems arose the notion of gateways between the networks and the concept of a much more robust end-to-end protocol.

O'NEILL: So it is really an internal problem then, as opposed to people wanting to connect multiple networks?

CERF: Well, no. I wouldn't say that exactly because the packet radio was an external network and the question was "How do I make computers in both the packet radio net and the ARPANET communicate with each other?" Where the ARPANET side might be reliable, but the packet radio side is not - and I need to have a common protocol in between. Well, I don't have to. I mean, that was one strategy which is what we chose eventually as an end-to-end

solution. So it was a problem of interconnecting two different kinds of packet networks, and eventually three because of the packet satellite system, and later four with the ETHERNET showing up as another kind of network that had different internal characteristics. But somehow machines had to communicate across all of these different kinds of nets successfully. We wanted to have a common protocol and a common address base so that you couldn't tell, to first order, that you were actually talking through all these different kinds of nets. That was the principal target of the INTERNET protocols.

O'NEILL: Okay, back to Stanford. What kind of projects were you working on while at Stanford, other than teaching? Were you receiving DARPA money, working on DARPA projects?

CERF: In fact, just before Steve Crocker left DARPA he helped to put together the arrangement for me to have funds to start working on this INTERNET project. Bob Kahn had started the project; I guess Crocker may even have been - I don't remember whether he was ever formally charged with managing it. I think Bob probably was, but somehow Steve was in there. I think he introduced me to Craig Fields. My first introduction to somebody at DARPA other than Bob Kahn and Steve Crocker was Craig. So it was fairly early on, I think by 1973, I was under contract to carry out the INTERNET research work. It probably would have been late in 1973.

O'NEILL: So before the INTERNET work then you hadn't been doing any DARPA projects?

CERF: Well, before I had a contract I was actually starting to work on the problem, because even in June, or even before that, the problem had been posed, and I was beginning to do some design work, and trying to think my way through the problem. But I didn't have any formal support for it until later.

O'NEILL: When you say the problem had been posed, you mean that Bob Kahn sat down with you and said...

CERF: Bob Kahn raised the problem. He said, "Here's the problem, what do we do?"

O'NEILL: It wasn't any more formal than that?

CERF: No. No, it was just "Here's an interesting problem."

O'NEILL: While you were at Stanford did you attempt to have students supported through DARPA as well?

CERF: Absolutely. I had a whole team of graduate students, some of whose names are now fairly familiar in the industry. Judy Estrin, who is Gerald Estrin's daughter, was one of my master's degree students. Of course, she went on to found Bridge Communications, and now she's running, or she's executive VP of the company that makes those little NCD X-displays. A guy named Richard Karp, another old high school friend of mine, decided to go back to graduate school and I took him on as a research assistant. He wrote the first TCP in BCPL on the PDP-11/20 at Stanford. He went on to get a Ph.D. in theorem proving and now is president of a company called ISDN Technologies out on the west coast in Palo Alto. Another person is Yogan Dalal who was a graduate student at Stanford and was deeply involved in the design of the TCP, the first go around, and also did a lot of work on the first documents that came out in 1974. He is now vice-president of software engineering at Claris Corporation, a spin-off from Apple. A guy named Carl Sunshine, who has written several books on the subject of internetting, did his dissertation work in my group, and is now running a lab at Aerospace Corporation. He took a job that Steve Crocker vacated in order to go to work for Trusted Information Systems. (Laugh) I mean, you'll never be able to disentangle this group. Let's see. Then I had visitors who were there, not graduate students.

TAPE 2/SIDE 1

CERF: I already mentioned that Gerard Lelann was with us for a year. There was another guy from Norway, Dag Belsnes. He did some really interesting work on pure datagram protocols and how you get reliable connection initiation. In fact, he managed to prove that a three-way handshake was not enough and that you actually needed a five-way handshake to make sure that everything was right. And we decided that was overkill and accepted the limitations the three-way handshake imposed on us. There were some others. A guy named James Mathis who went

on to work at SRI International on the packet radio project and now is at Apple. He built the first TCP for an Apple system. Also Darryl Rubin, now a vice-president at Microsoft, and Ronald Crane who is a key person at Apple.

O'NEILL: Were all of these students supported out of ARPA funds?

CERF: Yes.

O'NEILL: You didn't have to look for other funding sources?

CERF: Well, there was some other funding that I had available, too. There was something called the Joint Services Electronics Program, or JSEP, and there were small amounts of funds available for the faculty in the digital systems lab where I was located. But almost all of my funding came really from DARPA.

O'NEILL: I would like to talk about your move to DARPA in 1976. Why did you go there?

CERF: Actually, I had been asked to go more than once. I had my arm twisted one time. At the time I think Steve Crocker was still there, and Bob Kahn of course had been there for a while too, and they both said that they really thought I would be more helpful if I was at DARPA rather than at Stanford. The first time that they offered this, it was in the dead of winter, and I flew back here and they had the worst snowstorm they had ever had in a decade. And I turned around and went home. (Laugh) I said, "I'm not interested in living in this." I think that must have been in 1974 or 1975. A year later the issue came up again. This time I was more interested, in part because I was finding my time very fragmented at the University. A lot of people needed to see me, I had seventy students in the master's degree program and about a dozen graduate students in Ph.D. programs. I was finding it hard to get any research done because I was so busy doing that, or writing proposals, or writing reports telling people what I would do if I wasn't spending all my time telling them what I was going to do. It started to appeal to me that maybe I could get more done, if I had more resources available. At DARPA I would be able to call on the resources of some of the best places in the country to actually do substantive work. [INTERRUPTION] So it became more appealing. And

my wife finally said that she was willing to go to Washington because she was curious. She'd grown up in Kansas and lived in Los Angeles and San Francisco. It was really her interest and enthusiasm for moving to Washington that sort of tipped me over the edge. I was very reluctant to go. Not so much because of the geography, that was an element, wondering what it would be like living here in the winter. I grew up in Los Angeles. But I also was a little worried about taking the job because I thought, "Gee, it's very visible. If you screw up, you will screw up in a very visible way and everybody will know." "Ah, that guy really loused it up." That gave me a few sleepless nights. But I had the same trouble deciding to go to Stanford when I was offered a job right out of graduate school at UCLA. I thought, "I don't have anything to teach anybody. What am I going to say?" That made me nervous. So I wasn't sure about going to DARPA, but I finally got my arm twisted and they said, "Well, you'll never know unless you go." So I said, "Well, all right, I'll try."

O'NEILL: Was it primarily or exclusively to work on networking projects, or were there other kinds of things they wanted you or you wanted to get involved in?

CERF: No, I was almost totally enmeshed in the projects involved in packet satellite, packet radio, internetting and network security.

O'NEILL: Were these projects set up at this point, and you were managing them, or were you free to initiate projects?

CERF: Bob had got them all started really. I mean he had started the packet satellite program, he had started the packet radio program, he had started the internetting program that I was mostly involved in at Stanford. I had actually already had some exposure to the network security program. Another person, Steve Walker who is the president of Trusted Information Systems now, had started that program or at least picked it up. I think Bob was involved as well. But I had been doing some consulting out of Stanford on the network security project. So by the time I got to DARPA I was actually very familiar with all of the projects that had been started. It was nice. I didn't have to go and initiate any of them, I just had to grab hold of them and try to steer them. That made it easy - that made it easier. I shouldn't say it was easy.

O'NEILL: What did you consider as your mission when you went to DARPA?

CERF: What was my primary mission?

O'NEILL: To just continue what was going on?

CERF: The single-minded goal was to get the Internet system up and running. My principal concern was to get the set of protocols established, implemented, tested, and deployed that would let us demonstrate that multiple packet networks could indeed... hosts on multiple packet nets could reliably communicate with each other. There were lots of ramifications for the military. For example, we absolutely wanted to bring data communications to the field, which is what the packet radio project and the packet satellite projects were about; how to reach wide areas, how to reach people on the oceans. Can't do it by dragging fiber, can't do it very well with terrestrial store-and-forward radio because line-of-site doesn't work very well on a wide ocean. So you need satellites for that, or HF but we wanted high bandwidth. And the best way to get high bandwidth, at that time anyway, was still satellite systems. So the whole effort was very strongly motivated by bringing computers into the field in the military and then making it possible for them to communicate with each other in the field and to assets that were in the rear of the theatre of operations. So all of the demonstrations that we did had military counterparts. The students who were working on it didn't necessarily have that model; they didn't have to. But we did one demonstration of the INTERNET in 1977. One of the first demos of a multi-net system operating where we had a vehicle going up and down the Bay Shore Freeway with a packet radio in it, ran that through several hops around the Bay Area where we had repeaters on mountain tops, ran that through a gateway across the ARPANET, across the packet satellite net to UC London, across a dedicated satellite link from Norway that we got to via the packet satellite net, back to the ARPANET, then all the way down to USC - across the terrestrial ARPANET. So what we were simulating was a situation where somebody was in a mobile unit in the field, let's say in Europe, in the middle of some kind of action trying to communicate through a satellite network to the United States, and then going across the US to get to some strategic computing asset that was in the United States. So we simulated that sort of thing, at least as far as the

communication system was concerned. And there were a number of such simulations or demonstrations like that, some of which were extremely ambitious. They involved the Strategic Air Command at one point where we put airborne packet radios in the field communicating with each other and to the ground using the airborne systems to sew together fragments of Internet that had been segregated by a simulated nuclear attack. There were all kinds of challenges for this technology to overcome that were military in nature, that were problems that were caused by very hostile environments. Now as it has turned out, the robustness in the system has been helpful in the civilian sector, too. They may not be as dramatic but a cable cut through an optical fiber line is just as devastating as nuking some central office somewhere, as far as communications is concerned.

O'NEILL: Did the military suggest these kinds of scenarios, or did you present to them "This would happen, if..."

CERF: We worked together with various of the military commands. One of the things that DARPA had to work the hardest at was to involve the operational command in the examination of new technology. So the people who ran the Information Processing Techniques Office, a fellow named Dave Russell who preceded Bob Kahn, for example, spent a lot of time visiting various of the Army commands, briefing them about packet radio and networking technology and trying to stimulate their interest and willingness to allow us to deploy in the field this kind of equipment for their use in exercises. Ultimately we did that. We deployed a whole bunch of packet radio gear and computer terminals and small processors to Fort Bragg with the 18th Airborne Corps and for several years did a whole bunch of field exercises. We also deployed them to the Strategic Air Command in Omaha, Nebraska and did a series of exercises with them. In some cases, the outcome of the applications that we used were so good that they became part of the normal everyday operation. A small example of that was in the case of the Fort Bragg 18th Airborne Corps. SRI International developed an application that would let you automatically figure out what the load plan should be for moving, let's say, a heavy mechanized division. Maybe it's got trucks and tanks and maybe some helicopters and vehicles and the like. The question is "How do you assemble all that equipment; in what order do you place it, on which aircraft, where do you put the people, what about all of their supplies, and everything else? How do you do the loadmaster job? Which aircraft do you load up and when, when do they take off, when do they dump, how do they dump? Do they do low altitude parachute extraction, do they dump from high altitude? What's the story? Are

they going to land under fire, or not?" All those things affect the load plan. It used to take days to figure this out. Guys would run around with little stubby pencils and paper trying to figure it out. Halfway through somebody would say "Our C130 isn't available, or the C5 or C7," I never get these numbers right, "isn't available." And you would have to redo the plan. This was cleverly calculated with an application developed by SRI International. A guy named Mike Frankel led that effort. The program could recompute the load plan in midload in thirty seconds and pump out the new manifest. And it got everything down to the inch. "You put your vehicle in this way, tie it at that place in the aircraft with this orientation. Here's the order in which you do it." So they would sit there on the tarmac with a packet radio and a terminal and a small computer and enter all the data and the computer someplace else in the ARPANET would actually run the load plan and fire the results back in thirty seconds, and they would pump it all out. If something went wrong in the middle, they would start it all over - not all over. They would start in the middle they would say, "I have loaded this far, and I got all this crap in these airplanes, now what do I do?" They actually made that their regular load planning aid, as far as I know, rather than doing it by hand anymore. We ran one test where the computer generated the load plan and all these master load planners went through and did their best and they didn't do as well as the computer program. That was neat.

O'NEILL: Why did you leave DARPA?

CERF: I had started to think about college costs for my kids (who were ages nine and four at the time). The best estimate was \$25,000 per year *per kid*, or at least \$200,000. I could not see how to do this on my DARPA salary. As I pondered this problem, I had an invitation to meet a guy named Bob Harchasik who had recently joined MCI after serving as president of TYMNET for several years. Bob wanted an engineer to design and manage the construction of a "digital post office." There was a very substantial amount of resource available from MCI to do the job. This was a challenge that would use all my DARPA-acquired skills and know-how. What emerged was MCI Mail - an effort which has borne fruit and is still very much in use today. (Indeed, since I joined CNRI, I've arranged to link MCI Mail to the Internet!)

O'NEILL: To finish up, because we are just about out of time, do you have any further or general comments about

DARPA or the ARPANET, its impact, etc.?

CERF: I think that there isn't any doubt that the investment that DARPA made in the ARPANET put packet switching technology on the map. It added whole new communications technology that didn't exist before DARPA invested in it. It convinced people it was real and has spawned a phenomenal explosion in new kinds of computer communications techniques. The ARPANET sort of triggered interests at the University of Hawaii in what they then called the Aloha Net. And the Aloha Net was what caused Metcalfe to invent ETHERNET. Because he wanted to do the same thing on a cable, or he said "Why not do the same thing on a cable?" The INTERNET, which was spawned out of this conglomeration of different packet technologies that DARPA initiated, has already had a pretty dramatic impact on both the military and the commercial world as far as I can tell. You can't pick up a trade press article anymore without discovering that somebody is doing something with TCP/IP, almost in spite of the fact that there has been this major effort to develop international standards through the international standards organization, the OSI protocol, which eventually will get there. It's just that they are taking a lot of time. So I don't think that we would have seen the kinds of applications that are emerging now and available to us had ARPA not made such a major investment in packet switching. Of course, it was forced to do it, not because it thought it was a neat idea, but because any other way of getting computers to communicate was too expensive. The real reason that packet switching happened is that when DARPA said, "We have thirty sites; we want to connect them all together and they are all time-shared, how do we make that work?" AT&T said, "Well why don't you connect all n machines to all other n machines." That was n squared over two, roughly, lines. That was expensive because it went all the way across the country. Even ARPA couldn't afford that. So then they said, "Why don't you build a circuit switch system?" And the answer there was, "It is too slow." Because you have twenty-five people on one machine and they want to connect to twenty-five other machines, you would have to have twenty-five different circuits all going at the same time. So that didn't work. Finally DARPA just threw up its hands and said, "This can't be the right thing. There must be another way to do it." And invented packet switching. So there. (Laugh)

O'NEILL: Okay. Well, thank you very much.

END OF INTERVIEW