Interview of Steve Crocker

Interviewed by:
James L. Pelkey

Recorded April 26, 1988
Washington, DC

CHM Reference number: X5671.2010

James Pelkey: How did you get involved in all of this?

Steve Crocker: Thanks for the leading question, but let me talk about a couple of other things for a second. You mentioned, talking to ARPA, and then the next thing you said that Frank Heart was on your list to talk about ARPA. When you say the word ARPA to me, I have a very clear distinction in my mind between the agency, which existed before and has continued on and is much broader than computer science --

Pelkey: I'm talking about ARPA now.

Crocker: I know you're talking about ARPA now because that's a common thing -- versus the Arpanet itself, and what I wanted to say, I don't care if it's on the transcript or not, is that the agency itself had a much stronger role than might be apparent even after you look in and see the developments, and the top management of the agency, the directors during the period that you are interested in, embodied a considerable amount of vision, and in some ways, people like me and Steve Walker and even Larry Roberts were agents of that vision. Although we had a great deal of thought and we had a great deal of freedom to be helpful, there is even a larger point of view that was in progress, a larger vision, and in terms of tracking when were the concepts known and how did they unwind themselves, how did they spin out to have the effect, it might be useful to talk to some of those people.

Pelkey: Like Licklider and Taylor.

Crocker: Well, Licklider and Taylor were, at various times, directors of the office within ARPA that managed that technology, which was the IPTO. It's now ISTO, but it was IPTO then, Information Processing Techniques Office. That is one office within what was ARPA in '72. It went through a minor change and became DARPA. The directors of the agency had to deal with the issue of how much resources go into computer science versus going to strategic weapons or tactical weapons. How do you defend a nascent program in a budget environment where you have to justify the budget to Congress and through the Defense Department and the Office of Management and Budget Cycle? Where do you lay down the major commitments? What do you protect, what do you surface? Other things like this. Eb Rechtin --

Pelkey: How do you spell his name?

Crocker: His first name is Eberhardt, I think, and his last name is Rechtin, R E C H T I N -- was the director of ARPA until sometime in '71, just before I got there, and his deputy took over, Steve Lukasik, and was director until late '74. In terms of the critical developments of the Arpanet and the genesis of it and the nurturing of it and so forth, those two clearly understood what was going on, clearly were supportive, were absolutely critical in the process of nurturing that development, protecting it from a budget point of view, and all of those sort of things. Rather than being mere paper pushers who took all things and said: "Well, if you can do it, fine, here's a few bucks to do it," had a very strong sense of what it was all about, and probably things couldn't have gone so well if they hadn't been so supportive. They're both in Los Angeles, actually. Eb just retired as president of The Aerospace Corporation and is now associated in some capacity with USC, and he went on -- this is interesting, actually -- he went on from

DARPA to become an assistant Secretary of Defense for Telecommunications, the only that ever existed. That was a director slot that got upgraded to an assistant secretary slot. He was it, then it got downgraded or folded in with something later, but he is clearly a guy who has been involved in communications in a broad sense. Steve Lukasik is a vice president at the Northrup Corporation, and his office is in Century City. It's not that hard to track them down.

Pelkey: Could you help me with Rechtin's number?

Crocker: Sure.

Pelkey: It's interesting; this is the first time that their names have come up.

Crocker: Is that right?

Pelkey: Absolutely.

Crocker: Larry didn't mention them?

Pelkey: No.

Crocker: Well, I can't say what's on Larry's mind. I don't know what's on Larry's mind.

Pelkey: I understand, but this is very helpful, because I have wondered about that.

Crocker: Now, they will have a different point of view. They have seen it from there.

Pelkey: Strategically, I would like to hear their point of view.

Crocker: They're fundamental, I think. I wanted to toss that in. Now Frank Heart, of course, is central to the Arpanet, and others --

Pelkey: Subnet.

Crocker: The subnet in particular, and others at Bolt, Beranek and Newman. An interesting guy, if your going to go to BBN is Dave Walden.

Pelkey: You're the second person. I've seen his name, obviously, many times, but --

Crocker: Yeah, now Dave Walden is a guy who came in at the bottom as a junior guy and is now quite a senior person. The IMP code: the code that was in the IMPs themselves. Now the IMPs were small machines. Honeywell 516s originally and then 316s, which were essentially equivalent, and they had 16K of 16 bit words of memory for data and programs and everything with no external storage, no disk or anything. That tells you something about the amount of code, not all that much, written in assembly language. The code was written by essentially two people: Will Crowther with Dave Walden assisting him, so Dave was sort of the junior member of that team. Dave is now president of BBN Laboratories, or he was last I looked a few months

ago, but anyway, BBN has grown from a company that was, in those days, I once looked, you can get real facts from them, but as I recall, around $20 million, to a big company now, hundreds of millions of dollars, and it was one company and now it's a collection of companies, and Dave runs the BBN Labs, which is in some sense the flagship, the repository of the research side of things. So he's got a clear view from the whole slice of things. In fact, Frank Heart now works for him -- interesting reversal there. So I wanted to add a couple of people to your list. There are plenty of other people; I don't mean to leave anybody out.

Pelkey: No, no, but I do want to make sure I touch the critical people.

Crocker: Now I'll tell you my story, as they say. There were actually two fairly different things that came together. This is -- I'll allow myself to engage in some personal kinds of things. I had spent a lot of time as an undergraduate, and before that a little bit of programming and drifting into systems kind of programming; operating systems and that kind of thing, but I had had -- my main objective was to do things that in those days were considered part of artificial intelligence, and in some ways are still part of it; formal proofs of correctness and automatic development of software, things like that. It's an accident, in terms of my own agenda, that I got involved in the communications side of things. I actually made an explicit decision not to let it overtake me, and perhaps that was the wrong decision, but I spent a lot of time involved, and people like Vint Cerf, who was a very close friend of mine, got dragged in right behind and kept going, and Metcalfe too, in some respects. They've done pretty well; they've done pretty well for the field and they've done pretty well in all respects for themselves. I had been at MIT for a short period of time studying artificial intelligence and I had come back to UCLA, first just for a summer, and then decided to stay on.

Pelkey: What years were these?

Crocker: I came back from MIT to UCLA at the end of May1968. I had grown up in Los Angeles, I had gone to UCLA as an undergraduate, and I had come back for the summer as a summer job, and then I had decided that I was much more comfortable at UCLA and stayed there. During that summer, all of the following things happened: ARPA had an RFP on the street -- well, it was preparing an RFP, and then it had the RFP on the street during that period of time. My strongest involvement with that RFP was that, while I was working at UCLA, I had gotten involved with a part time consulting job with a small company nearby that turned out to be one of the bidders on the subnet, and so I helped write a section in a proposal, and I was concerned, actually, about whether the fact that UCLA was going to be one of the sites was at all a conflict of interest with helping -- I was young; I didn't know exactly where the rules were, but I had some sense that there might be an issue there -- and I went and asked Barry Wessler, who was program manager at ARPA and was -- I'll tell you how that conversation came in -- whether or not there was any issue, and he had a very interesting response which was: "Not only is there no conflict of interest," and in fact the rules are pretty clear, the selection people, which were the government, were the only ones that were protected, and the fact that I might get involved with using the Arpanet had no relationship to anything -- "not only was there no conflict of interest, but it would be a bad thing if I felt there was any conflict of interest, because that would weaken the opportunities for the company that I would be helping, or that would weaken my options," and in some sense that, quite opposite from some of the strong pressures that, particularly post-

Watergate, of avoiding appearances of conflict of interest. If you wanted to avoid any inhibitions that one might feel -- people a lot now-a-days avoid doing things because they are afraid that it might be viewed as conflict of interest, he wanted to avoid the opposite, that people might feel inhibited and not participate, and that would interfere with ones freedom of movement on a personal level, and would also weaken the responses to the RFP -- I always thought that was a very enlightened and sage attitude.

Pelkey: This was the September time frame?

Crocker: Well, I didn't meet him until early July, so it was later in the summer sometime, or maybe the fall. Somewhere there, I don't remember the precise time. Now, ARPA was pushing --

Pelkey: Excuse me, what were you doing at UCLA?

Crocker: I was a graduate student. I had come back from MIT and I was working -- at that time --

Pelkey: In a PhD program?

Crocker: Yeah, it was --

Pelkey: Were you associated with Kleinrock at this point?

Crocker: Well, ok that's right, the whole situation. There was the Computer Science Department, there was a contract under Jerry Estrin, not under Kleinrock at that instant, and I had gone to work for Jerry Estrin. He was doing measurement of operating system type of measurements, and in particular, he was trying to build a computer that would plug into another computer and do measurements of it, and he had procured a Sigma 7, which seemed like the right choice at the time, and needed software developed for it. I had worked for Estrin toward the end of my undergraduate career, and actually, that's an interesting aside. In the '65 time frame, ARPA had supported a project at UCLA, this is now going back a few years, which Taylor and, actually, Sutherland had started it up and then Taylor took over. The sequence of directors in the office was, I think, Licklider, Sutherland, Taylor, Roberts and then Dave Russell -- I'm sorry, Licklider came back for a brief period of time and then Dave Russell and then Bob Kahn and so forth. Networking had been a theme in the DARPA office, in IPTO. There had been early computer network linking experiments between Lincoln and SDC, which I only later heard anything about. They had started up a project at UCLA, which, it's pretty amusing in a way, to look back and see what there parameters were. UCLA had three 7094s, an unusual situation, in some ways a very rich environment. They were run by three entirely different organizations and had many political issues and financial issues and technical issues. They were basically three separate fiefdoms. One was run by the medical school, and was supported by NIH under its computers in medical research area, and was a quite well known and vigorous and successful computing facility. For example, the BIMED series of statistical programs were built in that environment, I think pretty much. You might crosscheck me because I'm giving you un-researched history as I saw what was going on. Another was Western Data Processing Center,

which was one of two facilities that IBM had set up in the mid '50s to promote graduate business schools' involvement in computing. They had established two regional centers; one at UCLA and one at MIT, paid for the machine, put the machine there, provided considerable amount of help, and the idea was that the machine was to be available at UCLA to Graduate Schools of Business in the 13 western states. They had also installed their own office of IBM employees that were effectively part of the IBM Scientific Laboratory system, as best I could tell, and again, there are other people that can give you the real history. In fact the people who were intimately involved in that are still at the Los Angeles Scientific Center up the street. It's very easy to find them, for example. That might be another lead. They had a 7094 installed, and they had managed -- there was a period of time when, when the campus wanted to go buy its own -- what I've missed in those two is what did the real university use? What did the central facility use? When it came time for the university to buy a machine, IBM managed, in some controversial moves, actually, to have it be the 7094, and there was a short period of time when one machine was used in three shifts; one by Western Data Processing Center, which was an organ of the business school, one by IBM with its own corporate employees, and one by the university in a group that had originally built the SWAC, the Standards Western Automatic Computer, a 1949, 1950 era machine, and turned into the Numerical Analysis Research Group, and was sort of an unabashed set of machine language programmers that knew an awful lot about computing down at the bit level. The Western Data Processing Center crowd was characterized by two-minute FORTRAN jobs and IBM was characterized by wearing suits all the time, if you like, and they ran three different operating systems and they did a complete flush of the machine, it was like periods processing. There were some tense periods where the staff really didn't get along with each other to the point of barring key personnel who were central to one shift not being permitted in the building during the other shifts; some fairly tense periods of time. Well, eventually, that machine moved out and became the central machine for the campus. Another one was brought in to WDPC, so there were three machines: WDPC's, the central machine for the campus which was called CCN, Campus Computer Network, but the network was more in hopes rather than in fact at the time, so it was the central machine for the campus, and then the Health Sciences Computing Facility, HSCF was their initials, and on top of all of this, ARPA supported a linking experiment which was going to connect the three 7094s in this political and technical environment. The distances involved -- UCLA is a fair sized campus, I don't have a clear idea, but it's maybe half a mile from the northern end of this to the southern end with a stop-off slightly off target to the middle one -- they were going to use 360 Mod 40s, one of the first set of 360s that came out, relatively modest sized machines, and those 360s were going to do two functions that we would now describe as two entirely separate things. They were going to act, in effect, as IMPs, and they were also going to act as front-end workstations. Now 7094s, you must remember, are batch processing machines. They weren't timeshared. One of them was used for timesharing at MIT in the CTSS system, but that was a one of a kind, special purpose machine. It wasn't what the machine was basically used for. These 360s were going to have a small timesharing system whose outlines were not as ambitious, but are like UNIX, in a sense. They were intended to be interactive, terminal-based systems where users would compose programs and would edit them and maybe would run small things, and then would submit them to the batch processors. You would be able to move files from the 7094 to another 7094, or back into the 360s. The 360s had file systems on them and so forth. Then the 360s themselves would be linked in a triangular network with cable laid on the campus, coax of some sort. It never happened, so I don't know what the --

Pelkey: And the 7094s were linked to each other as well?

Crocker: The 7094s were linked only to the 40s. There were going to be three 40s -- three 360/40s -- one at each site, and those were to be timesharing. That's why I say they were the IMPs in essence and they were --

Pelkey: What year was this?

Crocker: '65?

Pelkey: Well, that's pretty progressive for then.

Crocker: Well, it was an ARPA thing. That's what it's all about. Now that project was flawed in a couple of ways. The staff that carried out the project was taken entirely from WDPC. That project was formed by taking [Western Data Processing Center] WDPC staff people and dividing them and forming another project, but it was housed at WDPC, the principal investigator was Clay Sprowls, who was a professor and was head of the WDPC and he was the titular head, but had not very much to do with technical aspects of the project, so it was never clear what the cooperation was going to be from the other sites. Now, for various reasons, that project collapsed. Some pressure was put on it to meet certain goals, and there was a question as to whether or not it wanted to do that, and Sprowls fired the top half of that staff in January of '66, I think, and -- that's approximately right, and the whole project kind of died towards springtime. It limped along for a little while. The residue of that money wound up in Estrin's hands, which was a reasonable thing, but that was down in the engineering school. It was entirely separate. And in the summer of '66, I wound up working for Estrin. It wasn't exactly handed over, but there was money there and he said come around and work when that project died, so I did, and we had a good time. I worked for him in the summer and fall, just before I went off to MIT. So it was not unnatural to come back and work for him a year and a half later in the summer of '68. By this time, I had spent the whole time at MIT.

Pelkey: Were any of those 360/40s ever --

Crocker: One 360/40 was delivered. We designed and partially implemented a complete timesharing system from the ground up. It had no relationship to the existing software. As a technical matter, we did something that was fairly ambitious. The 360s have base registers, and a very short address, 12 bits, imbedded in the instruction, and then it references a base register and separately an index register. This was pre-paging, before virtual memory, and we imposed restrictions on the code that could be run in a way that made it effectively re-locatable code. In other words, we said all the code had to be arbitrarily re-locatable if we update the user visible registers, three or four registers that you could use as base registers, so the style in which you wrote programs had to conform to having that code picked up and moved arbitrarily with the operating system adjusting those registers, and that would have made it possible to run it just like an ordinary timesharing machine with swap-able processes and so forth, and we did all that. It was neat. It meant that users couldn't write programs unless they conformed to that and have them do anything. There were protections, I guess, so the programs couldn't go out of bounds. I

can't remember precisely, but mapping registers, page tables, didn't come in until the model 67. Now, of course, they're standard on everything. Oh, yes, Multics had them -- there were a series of papers in the fall joint computer conference in Las Vegas in 1965, and those papers provided a great deal of information to this project, and our design was far less ambitious than Multics, but the 940 had come out, the Berkeley Genie project, which had taken an SDS 930 and added paging hardware to it and made it a timesharing system, and the Atlas Project at Manchester. So these things were known, and we had done a crude form of adapting those and adhering to the rules without making hardware modifications to the 40 and trying to pull off this small coup. It was a neat little thing. I was the most junior member of the project and designed a little editor that would -- there were also issues of protection, of how many people could be referencing a file and what happens if somebody's changing it while you're accessing it, things like that, single-writer, multiple-reader kind of problems that we looked at, and efficiency issues about how often do you have to update or can you keep and alter file, so a lot of exciting things, but the project died. It was too much pressure in that environment, and I'm not sure exactly, but Sprowls pulled the plug. He said he didn't want to have to run a big time project, at least that's how I understand it. Now I end up meeting Estrin and working for him, and we worked on instrumentation of programs, and entirely different kind of subject. How often does a program execute various portions of itself, things like that? Then I went off to MIT and then when I came back, which was initially just for the summer and then decided to stay -- oh yes, there's one other interesting thing. Vint Cerf and I had gone to high school together and we were bosom buddies. We were very close. We would sleep at each others house and we would knock around.

Pelkey: Is this true?

Crocker: Oh, yes. We were later -- I stayed at UCLA through the end of '66, then I went off to MIT at the beginning of February ‘67. By that time Vint was married. He had gone off to Stanford. He was a much better student. He had gone off to Stanford and graduated from Stanford and come back and was working for a year for IBM before going off to graduate school. The draft was an important thing in those days. Vint was entirely exempt from the draft because he is very hard of hearing, but I wasn't. I had limped through school. I had taken longer time to go through, so I didn't have the freedom to drop out of school, per se. Even that wasn't very much help, so he was going to go to graduate school. I had been looking at graduate schools and one school that I had initially focused on was Irvine and I had gone down and -- but eventually somebody got hold of me and said: "Look more broadly," and I wound up at MIT, but I had some knowledge and some affection for the department down at Irvine. Vint had applied down there, and for reasons that I think were just the usual kind of complete miscue; they rejected his application, which was just a plain mistake. They didn't understand who was applying, or it got misrouted, or whatever, and he had recently gotten married, and his wife was not eager to move far away. So I introduced him to Estrin. I basically marched him into Estrin's office and told him that there was a major find and that he should grab this guy quick. Estrin's a bright guy. He had no trouble – (I make it sound like I was in charge of everything). So as I went off to MIT, Vint was entering graduate school at UCLA, and when I was looking for a summer job in Los Angeles for the summer of '68, Vint engineered a job for me, and I came back and spent the summer, and then I wound up staying. We shared an office through graduate school and had a great time. I was the best man at his wedding, and later he was at mine. There are a few more things which I will describe. Ok, so I left behind Vint, then I came back for the summer. Estrin

had focused on this instrumentation process, and he had continuation of the DARPA contract, which had now transformed its goals into looking at measurement stuff. Kleinrock, I'm not sure exactly what his role was. He was in the department. I met him that summer, but he was not in charge of a lot of things. Almost completely separate, another interesting thing was happening in the ARPA scene. In those days, ARPA supported a relatively modest number of substantial contracts, at MIT, at Stanford, at Carnegie, at SRI, and at UCLA and so forth, and the heads of those projects were called together for an annual meeting, whose primary purpose was, as best I could tell, was to rap, just to chat and to see if any great, brave new ideas would get surfaced. So it was the annual Principal Investigators' meeting. Newell or Simon would come from Carnegie and Minsky from MIT and Fano or somebody from Project MAC, and McCarthy from Stanford --

Pelkey: Culler?

Crocker: Culler, right, the Culler-Fried System, which was interesting. (Daniel) Slotnick was pushing on the ILIAC IV ideas, and he said: "This is so much fun that we ought to clone this down to the graduate student level, and get the graduate students together," and that was done. So in early July of '68, the first graduate students' conference came together.

Pelkey: Oh, ok.

Crocker: Now, I had left MIT for the summer only. I had not made a decision to leave permanently, but I was in Los Angeles, so I was selected as one of the people coming from MIT --

Tape Side Ends

Crocker: So I got selected from MIT and Vint got selected from UCLA, and about 30 people -- somewhere I have the list, it's probably at home, maybe I'll find it and send it to you -- an interesting collection of people. It was hosted by the University of Illinois because Slotnick had made the suggestion; so he got stuck with the arrangements, although he wasn't there. There were no senior people there. Barry Wessler came from ARPA. I never had met him or heard of him before, but he worked for Larry Roberts as program manager, and he came out. Let's see, from memory, some of the people -- well, two guys showed up from Utah. One was Alan Kay, and the other was a guy named John Warnock, who is currently founder and CEO of Adobe. There was a guy name Bob Balzer who came from RAND. He was a senior research guy at ISI. A guy name Marv Schaefer, who is my colleague, another vice president of TIS (Trusted Information Systems). He came from SDC. Cordell Green came from SRI where he was a graduate student at Stanford, but he was working at SRI. He's an important AI guy. Two of use were selected from the Artificial Intelligence lab at MIT, the other was a guy named Pat Winston who is now the director. It was an interesting group. In fact, the story I like to tell is that I made an initial assumption going in that this would be fun to spend a couple of days, but who are these guys, after all. Maybe 10% of them will be really interesting, ok. I land at the airport and six of us wind up on a Sunday evening with a little time to kill before the bus is coming and we go out to dinner. One of them was Alan Kay, who was spellbinding, you know, 20 years ago. At the end of three days of cooking with these people, and we just exchanged ideas and talked about

what we were doing and what we were interested in, and it was great fun, I went through the list and just checked off the people that had made a strong impression on me and then counted it and it came out to be a full half. I thought: "Well that's pretty good." A short time later, a couple of years later, I'm at ARPA, and I inherit these conferences. I move them out to Pacific Grove, to Asilomar, and had great fun. We had the same format. We spend the first night with everybody quickly introducing themselves and then choosing some topics to rap about, and I would sometimes circulate through the group. Somebody would be talking, and various students would be listening, and I would circulate, and one kid covers up what he's doing, and he doesn't want me to see. What is he doing? He's doing the same damned thing. He's checking off, he's putting his impressions down, ok. Another arrogant youngster from MIT, as it turned out, but by now I'm three years older and he doesn't want me to see what he's doing.

Pelkey: So you have this conference in the fall of '68?

Crocker: No, this was July of '68, and Barry Wessler kept trying to get us interested as a group in this network, and the group's reaction was not very positive. It wasn't negative. There was no protest, but he wanted us to think about the issues of file sharing and some of those kinds of issues. We were all into AI and graphics systems and other topics that were hot at the time, and there wasn't a lot of response. It just sort of fell flat. That was one interesting benchmark.

Pelkey: How did the students get selected; by the universities? Was there anything special about that, other than that they were competent, motivated people?

Crocker: They were selected by the projects, as opposed to the universities, so in the very big environments, they were part of those projects, and some of them, like Balzer already had his degree, but he was the closest thing that RAND had to a graduate student, in a sense. So they were the people who were visible in those projects who were about the right age. There wasn't any selection process or anything. It was just: "Send a couple guys." So, the one thing that was true; as that that community was not keenly involved in networking, per se, and not really responsive --didn't quite see it. On the other hand, ARPA was pushing it and it was a serious vision. Now, I don't want to characterize everybody with that. There were probably a few, but my impression was that was not a common thing; now at Lincoln, maybe, although I don't know. BBN had not been selected yet. The RFP was either not on the street, or maybe it was on the street but it wasn't --

Pelkey: Decided.

Crocker: But certainly BBN was not involved. The other thing, and again this is entirely personal, I had had some view of where the interesting research in the world was going on, and place that I wanted to follow and people that I wanted to follow, and I was impressed by the fact that, at this graduate students' conference, almost a perfect fit with my model of the world was there. That is, all the places that I thought were doing interesting work had people there represented, and very little else. It never occurred to me that there was a reason until that time, that somebody else had -- that there was any organization in the world, that anybody had the right idea. I was just looking at it from a student point of view, and I thought that was remarkable, that there was somebody, particularly in Washington, where 'who knows what they

do in Washington,' had any idea of where talent was and what the right problems were, and I went up to Barry Wessler and asked him when he was leaving his job, clearly implying that I was ready to take it over as soon as he was ready to leave, no embarrassment at all. He looked at me --

Pelkey: You were a shy young man.

Crocker: Yeah, real shy. He looked at me like 'where'd you come from, kid?' I don't think anything particular happened, but I had basically announced myself, and my interest was driven by ARPA's direct involvement in artificial intelligence -- which was the hot topic, publicly, at the time; it was ridiculed, in a way, but not by me -- and by other technologies that I thought were critical, like timesharing systems, and just by it being on the leading edge of all these things, and the network was not -- Barry mentioned it, but that was not part of anything, from where I was sitting at the time. Meanwhile, ARPA was proceeding with its plan to bring the network into existence, and had decided that the first four nodes -- the program was structured that they were going to buy four nodes, install them, and at least the structure of the program was that they would then make a decision about whether that was successful, and go on. I think that decision was more in name than in practice, because the distance -- the time delay between the fourth node and the fifth node was pretty small, and that decision looked like it was made sort of: "Oh, yeah, we'll make that decision, but meanwhile, go right on with things."

Pelkey: One question. My understanding is that during '67 and certainly in early '68, ARPA was out at conferences and talking about this project, trying to get people interested. Larry Roberts describes that if anybody would listen to him, he would try to get people interested in this project.

Crocker: Well, Roberts came to ARPA, as I understand it, to do this, to make this happen. He had been one of the principals in getting that Lincoln/SDC experiment going, and other things. Bob Taylor was running the office, and Larry came in as a Special Assistant to the Director of ARPA, and later took over the office, but that was his mandate --

Pelkey: So even with this PR and ARPA going out there, this student group, you were just --

Crocker: Well, we never heard about it as a group. Now, Paul Baran was writing his papers, and one of the interesting questions that comes up: "Those papers tell you everything you need to know. Why do you have to actually spend the money to build it, because there's nothing to be learned? The answer is all there," and as is completely clear, one of the problems in technology transfer is the demonstration in order to engage in consciousness-raising. You can write the most elegant paper and have it completely clear --

Pelkey: Scenarios is also the same --

Crocker: And a good paper will affect a few key people, but a demonstration will change the world, so there has to be that step. So four sites were selected: UCLA, SRI, UCSB, and the University of Utah. Now, I don't know the precise reason why those were the four sites, but it was clear that they were there. There wasn't any choice, at least for the first four. The simple

statement that I heard was geographic coherence, since they were all on the west coast, that would reduce the cost of laying those lines, and why west coast versus east coast, maybe because the political environment was simpler, although I don't -- that was more by implication. Well, you see, the places on the east coast, take MIT as an example: very, very intense environment, lots of strong players, all of whom have very definite ideas of what to do. Even though you've got the money, you don't just sort of walk in there and say: "Look, your guys' research is going to turn in this direction and we're going to impose this on you." You might get a very strong reaction, and there's no point in using force, that doesn't get the job done at all.

Pelkey: It actually would have been interesting if they had gone to the Shannon group at MIT --

Crocker: See, Larry, among others, had been at MIT. I don't know what his relationship is with Shannon, but I'm sure he had a very clear idea of whether that was a good or bad idea, and I wouldn't be surprised if he thought it would be a very bad idea. As you point out with the little diagram about computing and communications, this revolution came out of the computing environment, not out of the communications environment. Going to AT&T was not the right idea. In fact, sort of out of sequence here, one of the anecdotes that emerged from after the first couple of years, a big point in developing the Arpanet was that there were multiple paths. Every point was connected in at least two distinct ways to every other point, so that if a line goes down, it didn't force separation. Now, the maintenance problem is that it turned out that these lines were sometimes perfect and sometimes broken and sometimes in a flaky state. When they were in a flaky state, you would call -- BBN was running the network operations center, Network Monitoring Center it was called, and they would call -- Network Control Center actually is what it was called, NCC. They've gone through some name changes. The IMPs would, in addition to sending messages to each other, would also keep track of the number of retransmissions necessary, the number of broken things. Background processes would accumulate the statistics and send those off to the Network Control Center, so statistics on these lines were kept quite accurately and quite up to date, minute by minute practically, so that when a line broke, the Network Control Center would know it right away, and when an IMP broke, it was visible as the equivalent of all the lines around it breaking, because that's the only information you had, and there was a two minute delay in the Network Control Center program, at a certain point, where first one line would get announced as broken at the scene by the other IMP that it's connected to, and then one after another, these lines would be marked down, and then the Network Control Center, sort of like calling an election, would two minutes later declare that the IMP must be broken as an inference on the fact that all the lines were broken. You could watch in a display that they had at one point of these lines breaking, you'd know that it was about to tell you that the IMP was broken. They had quite up to date information, very detailed, so that led to some very interesting interaction with the phone company. There might be a line across New York state that is flaky, and they would call up AT&T Long Lines and say: "We would like attention at line number such and such," and that would have to go from AT&T Long Lines to, maybe, the New York State Telephone Company, if it was entirely within their jurisdiction, or whatever, and then BBN might call them up and say: "You better wait up, we've got a problem down in Florida," and AT&T would say: "Well, what's Florida got to do with anything. This is a New York problem." We'd say: "No, no, that's part of our cross-country system. Wait until we get that line fixed, because we don't want to take this line down there." So, after a while AT&T began to catch on and it began to make an impression in there, and eventually they had a line from the

Network Control Center, a teletype line going down to some AT&T office in New Jersey -- this is stuff that you can get much more accurate and much more detailed stuff from the BBN crowd -- and they began to pay attention, but there were very humorous anecdotes. A General Telephone guy would call up about, say, a line in California, and would be told to call the 617 phone number to get more details if he sent out a repair call. So he'd call back and get the Network Control Center and he's tell them the circuit number and say: "I'm supposed to call you, I don't know why. Are you having trouble with this?" Then the Network Control Center people would say: "Well, it seemed to be down for a minute at 12:15, and then it came up and it was good for five minutes, then at 12:17:30 it was down for 10 seconds, and then it came back up," and they'd read off this list of things, and this guy is out there with his probe out in Los Angeles, or somewhere in California saying: "How do you know these things? I'm sitting right here and I can't tell anything about how good this line is, what are you doing?" They had all these self-measurement things going on like that. So that was great.

Pelkey: That's fantastic.

Crocker: And I don't know what the evolution was inside AT&T. That's another interesting question about the story that might be told, about what happened inside AT&T and how far it penetrated and so forth, but AT&T after a while began to respect this and be cooperative, and I think there's some story there. Now, as I say, that was a bit out of sequence.

Pelkey: Coming back to, now you've left the conference. You had announced yourself to Barry --

Crocker: Announced myself to Barry. That was sort of out of sequence --

Pelkey: But it was important, I suspect, to subsequent events --

Crocker: It was because it made something happen, in some ways, for the wrong reasons; one of these classic right thing for the wrong reason kind of things. So, that was just a little chunk of the summer. It was fun, around, just before the Fourth of July. Meanwhile, ARPA commissioned a set of meetings of representatives of the four sites, and a fellow working at SRI at the time, Elmer Shapiro, was given the task, and he called together representatives from the four sites. Although I was ostensibly only around for the summer, a task I had taken on was to see if I could clone the MIT operating system onto the Sigma 7, and get an operating system going for the Sigma 7 that would be appropriate for doing the kind of work that Estrin needed, which was the timesharing system. The Sigma 7 software that came from what was SDS and became XDS was essentially batch oriented with inadequate timesharing system overlay. We needed better than that, so --

Pelkey: This was CTSS?

Crocker: No, this was the AI lab's timesharing system, which was called, in their style, the Incompatible Timesharing System, ITS, as a deliberate jibe about the other half of what was going on there, and it was a very simple, very elegant, stripped down operating system that was very efficient and ran on PDP6 and later PDP10 hardware, and I had thought, somewhat

ambitiously, that I could take the central ideas in that and sort of map that design over to the Sigma 7 and we'd have a small coup. So we set out to do that. Anyway, I went to this meeting that Shapiro called.

Pelkey: Do you recall who else was there?

Crocker: Well, I recall a few of them. Jeff Rulifson was there from – you know what, this meeting might have taken place even before the graduate student meeting, because I think -- yes, I think this meeting took place before. Steve Carr came from Utah. I can't remember which of the people came from Santa Barbara. It might have been Ron Stoughton? It might have been somebody else? There were a couple of people involved, and there may have been a couple of others, but those were -- certainly Rulifson and Carr were key people. I don't remember whether Vint came from UCLA with me or not. Who else? There weren't more than about six of us. It was a compact meeting. Elmer called the meeting, but he didn't have enough to say. We asked some question and he said: "Yeah, those are the questions, in effect," and we were expecting the answers. We were expecting what I would now characterize as what the protocol was going to be. What were we going to do with our machines that we were going to connect up to the network? I remember Rulifson and I sort of caught each other's eye, that we were much more kindred spirits and on the same wavelength on these things, and not much came out of the meeting except that we had met each other, and decided that it would be fun to talk more about these things, and that we had some things that we ought to talk about. We didn't have a charter. There wasn't any agenda. There wasn't any mandate to go forward. There wasn't any organization imposed. There wasn't even any authority to do anything, but the question had been raised about -- well, we had been told, basically, that these IMPs were coming and we would get connected, and that was all the structure there was. So we kind of created the opportunity, partly because it was just fun, to have an excuse to get together again, and we did that. That became the genesis of the Network Working Group, and we found reasons to meet at each of our different sites. It was the original boondoggling mechanism. We wound up going to Utah. We wound up going to SRI and to Santa Barbara and UCLA. We sort of rotated hosting these things. I remember it must have happened before, because when I went to the graduate students' meeting, I was surprised that Steve Carr was not one of the two people, because I thought, at least from the perspective I had of Utah, he was certainly one of the guys making things happen. Well, it turns out Warnock and Kay are pretty good guys too, and that was fine, but I remember what the sequence must have been, because I was curious that Steve didn't come. I kept expecting that we were just sort of educating ourselves; and that eventually, somebody official would show up. BBN got the contract --

Pelkey: So you had had numerous meetings by this point?

Crocker: Well, a few, not that many, but a few.

Pelkey: So now it's Christmas of '68 --

Crocker: Christmas of '68, and they get awarded the contract, and roughly around Valentine's Day we had our first meeting with BBN. A few of us went back. I remember I got snowed in.

Pelkey: And the few of you that went back were you that --

Crocker: Well, we were the people working -- we were the same people, basically. I don't remember precisely who came, but it was that de facto Network Working Group, who had continued to have a few meetings and worry about host-to-host protocols. We came not as a group, but we came as representatives from the sites that were going to be connected. So we came as the customers, in a way, and listened to what BBN was going to provide. It turned out that there was, originally, an attempt by BBN to provide host- specific software in the IMPs, some tailoring, so the battery of exactly what the hosts had to do versus what the IMPs were going to do didn't clarify until later, and BBN retracted back to a fixed set of services and a serial line interface that they had defined, which has become known as 1822 Interface; 1822 was the report number that first laid this out.

Pelkey: Do you recall when that came out?

Crocker: Well, it must have been in '69 some time. It should be easy to find out. It's still called 1822, I think, so -- the date is some current date, but it's not so hard to find the original one. That was really our very first exposure to the BBN crowd there. One of the things that we had thought about was that the host protocol ought to be very suspicious about the network, and it ought to have check sums, and we spend a little time trying to devise a check sum, not that was the strongest thing from a cryptographic point of view because we thought that the reliability issues would be pretty straight forward, it either was right or it wasn't and we just needed something simple, but we tried to imagine what would catch the kind of errors that might happen, like packets being reassembled in the wrong order, so we wanted to make sure that, if the IMP swapped two packets, it turns out to be not the kind of thing that ever happens, but we knew that things as long as eight packets were being sent as a unit across the host-IMPs interface, the IMPs would break it down into single packets and ship those packets, so we postulated the idea that they might be reassembled in the wrong order, so our check sum had a little shift in it to make sure that it would show up as a change if the order was shifted. We also wanted something that was not computationally too onerous, because the traditional kinds of check sums that are very good from a check sum point of view are fairly hard to compute in a general-purpose computer, very easy if you get down to the hardware level and create shift registers and so forth. So we put that into our initial thinking. BBN had a very sharp reaction, Frank Heart in particular, that it was a waste of time, that it was computationally a tremendous waste. We went back and forth a little bit, and I remember in that meeting, I asked him: "How reliable is the host-IMP interface?" He was arguing how reliable the IMPs were, and I couldn't argue with him because he could make any claim he wanted, and also he clearly knew what he was doing. He had a lot of resources and there was no reason to disbelieve him. So I remember tackling what seemed to be the weak point that he couldn't absolutely guarantee were these host-IMP interfaces, and he said: "As reliable as your accumulator." That's what he said, and as it turned out they are not unreliable in general, but there were some very, very treacherous, intermittent errors across those interfaces on some hosts, and it would have been helpful to have check sums, but we didn't put them in. We got argued out of it, and the original host-to-host protocol, the NCP, had actually a considerable amount of concern -- I don't know how successful it was -- but a lot of the design reflected efficiency concerns, to be able to keep the overhead very low. One of the observations about the transition from the NCP protocol to the -- NCP was

originally not the name of the protocol. It was called, simply, host-to-host protocol. NCP was the name of the program that had to be written inside the operating system, and was the Network Control Program, and the protocol's name was "host-to-host protocol" which is not a name at all, it's just a sort of generic identifier, but later NCP became the designator of that protocol, in contrast to its replacement, TCP. TCP includes check summing, as a lesson learned, and consequently the overhead is higher, and it's proper to pay that price.

Pelkey: So you had your meeting and you argued these points and --

Crocker: Yeah we lost; we got argued out of it. We could have done either one. It was our protocol, in some sense.

Pelkey: Was the concept of this specification that became 1822, was that discussed at that meeting?

Crocker: I don't recall whether it was out or not. I remember, we were concerned about how many messages we could have in flight and how many connections we could have open at one time, in effect, and how the IMPs were going to manage that, and they seem to have thought their way only part way through those issues, but they evolved very rapidly. Bob Kahn was working for Frank and had done the primary design issues of capacities and loading and the central conceptual issues. Will Crowther was the implementer of the IMP software.

Pelkey: Now at this point in time, was it clear to you at UCLA and the four IMP sites that you had to write the software that sat in the hosts to interface to these IMPs?

Crocker: Sure.

Pelkey: So it was absolutely -- and the four of you, presumably, were also charged with the mission to make sure of the -- in fact, to talk to each other.

Crocker: Well, we just adopted it ourselves. Nobody came and -- in fact, that's absolutely crucial to the way in which the Network Working Group functioned for a period of time. As I said, we got together. Somebody had to pay our bills, so our bosses -- I don't remember where the shift took place from Estrin to Kleinrock, but relatively soon in there, but whether is was Estrin or whether it was Kleinrock, they were willing to let me travel, and the other guys seemed to be able to travel, not so much at first, but -- So it wasn't an invisible operation, but, at least from where I was sitting, there wasn't a major hierarchy and briefings and nobody was setting milestones saying: "When are you going to have this design done?" Nobody was coming from Washington and holding contract reviews or preliminary design reviews, critical design reviews.

Pelkey: It must have seemed kind of bizarre to you guys.

Crocker: No, what was bizarre was later when I found out that that's the way the real world worked, but it did seem -- there was an issue that was bizarre. It did seem kind of surrealistic. We kept waiting for the professionals. We were sure that there was somebody who knew the subject. First of all, we didn't know much about this, and second of all, who was cast, you know,

sort of adults, we were just kids, but surely somewhere there was a group, and if, it wasn't BBN, then it would be some other group at DARPA. So we viewed ourselves as kind of educating ourselves, and thinking about this, but we didn't view it as our territory, at least I didn't, but we didn't know what was going to happen. Nothing seemed to be going on, so this was kind of a very informal group. So we met a few times, and in the spring of '69, just after BBN was in the process, we began to meet with them. We had a meeting in Utah and decided that it was time to write down the thoughts that we had. Now some of our thoughts were much higher-level thoughts, having to do with applications and how to use computers for some complicated applications. One of the thoughts that we had was that the kind of computer that you could envision would involve two smart machines: a front end machine local to you, and a back end machine elsewhere on the network; and that the division of labor between what would be done locally and what would be done at a distance might be very application dependent. Rather than being a fixed interface – as is represented, for example, in Telnet where you have a virtual terminal but nothing much smarter, or a file transfer protocol -- we envisioned some fairly high powered graphics applications, for example, where much of the computing might be done locally for graphics display, but the database might be elsewhere, for example; and rather than defining a fixed interface that all application programs adhered to, we envisioned a programmable interface in which, as you set up the connection and as you initiated the conversation, the application program would download to you at the site you're at, a program that you would then fire up, and it would have this generic interaction -- It would have this specific, non-generic interaction. For example, if it was a text editing program, and you had a graphics display, say local editing, erasing a word or erasing a character, might be performed locally, and until you had settled on that, it wouldn't be transmitted back. It wouldn't echo every character, or if it was a more complex CAD type program, some of those computations might be done locally and some might be done - - and so we spent some time, we recognized all our machines were different. Heterogeneity was -- we had a Sigma 7 that was different, a 940 and a PDP10 at Utah, I think, and Santa Barbara had a 360/75, all those kinds of things. So, we started to go down the path of defining a new language that would be compiled locally and run, but which could write programs in, and special purpose for the point of writing these front end applications programs. Some of the early RFCs talked about this language, that was known variously as DEL, for Decode Encode Language, or NIL, for Network Interchange Language, so the very early RFCs had some of that in them. Only later, when we had specs from BBN, did we get down to the business of defining a low level host-to-host protocol, and try to knock that off as a task. Anyway, In March of . . .

Tape Side Ends

Crocker: So in March of '69, we decided that we had been meeting too many times and that we needed to save some of our thoughts and write them down. So we decided to generate a set of notes. Again, the context was that these were things coming out of these meetings, but we didn't believe that they had any status. Nobody had come from Washington and said: "You guys are real." We hadn't been officially tasked, or anything like that. We were just operating, basically, almost defensively as representatives of the sites that were going to have to deal with this network that was going to be imposed on us. So we were more of a user committee than we were a part of the team producing anything. I remember being very, very concerned that I didn't want our actions to be viewed as intruding into anybody's territory or encroaching on anybody's

official charter, and I worried over the words of that first memo far more than was necessary, of course, but I think I spent some very late hours writing and rewriting the words for the original thing, and one of the things I did is I wanted these things to drip with humbleness -- not our style -- so I said: "These are not design notes. These are not official of any kind. These are requests for comments, just a way of engaging people in conversation," so that's where the term came from. These were requests for comments, and I said: "You could write anything, and they have no status, and you shouldn't reference them," and that became a point of contention because they began to have quite a bit of information in them, and I remember there was some heartburn about whether these were published or not, in some sense. It was helpful that they weren't, because the Lincoln Laboratory, for example, had a hard time -- they would have to go through a clearance process, and we somehow managed to avoid, for the most part, that, because of the low status. Dave Walden gave me trouble about it later, about making these publishable, and I think we should have done more than we did at the time, but the original thought was these were just no-status notes. You could write anything you wanted, as short or as long, and send them out. There was a mailing list, and then we sent out -- people came onto the mailing list, and the mailing list grew and grew, and several of the yearly RFCs were just new mailing lists of everybody you had to send -- You were responsible for sending one to everybody on the mailing list, and then we imposed a rule that only one person from each site should be on the mailing list, and they would take care of secondary distribution and so forth. I think I handled the numbers. The only coordination was you had to get a new number; people would call me up and I'd give them a number. Now Postel, I think, or somebody, assigns those number. Postel worked in the group at UCLA, so we had -- our primary job was building this operating system --

Pelkey: Who was in the group at UCLA?

Crocker: Vint and I and Jon Postel, a guy named Charles Klein and a couple of programmers, and Mike Wingfield was a graduate student. He had ROTC commitments and he went off to the Air Force as soon as he finished. He built the first host-IMP interface, and I think now he's at MITRE. He went to BBN for a while after being in Rome to do his Air Force commitment, and now he's at MITRE as a senior guy in networking.

Pelkey: Because everybody had to build that -- connected the host to the IMP.

Crocker: That's right. That was a big, big problem, because if you had a lot of hardware expertise and you were building a lot of hardware and you had control of your machine, well fine. But if you had a standard piece of equipment from some vendor and the research you were doing didn't entail getting into its guts on a hardware level, you were faced with a pretty serious problem, because there was no standard interface. RS-232 was about the only thing you could count on connecting anything, and BBN didn't supply an RS-232 based interface, because they wanted to go at 100,000 bits per second instead of 1200 or 2400 or 9600 baud, or whatever. So everybody was faced with that problem, and of course there were no vendors supplying standard parts at that time. We went to SDS and asked them for an estimate, and they wanted $19,000, which in those days was a lot of money, and they wanted about six months, and we didn't have that much time. Wingfield was a very clean, effective hardware designer, and Estrin went to him or maybe Kleinrock, somebody went to him. He thought he could do it for about $4,000 in six weeks.

Pelkey: And he was a graduate student?

Crocker: These were things done in the shop, and so he was tasked with doing that, and in five an a half weeks and under budget, he had a box about a foot cube, maybe a little bigger, and elegant. It wasn't compact in the way you could do it now, but a super clean design, very simple design. There wasn't much that had to be done. The main problem was knowing what the interface had to be, so you had to get into your host computer, and he showed me - - He had a whole bunch of colored buttons on it for test setups and he showed me one day that you could single-cycle the thing and march these bits through and you could watch the bit march through the registers and single-cycle, and it worked fine for a while. He went off to a conference briefly, for a couple of days, and sure enough it broke during that day that he was gone. So I went over to the box, and said: "Well, I remember. He showed me this once," and I started to push buttons, and I watched the bit fall off at a certain point between, I think, bit 24 and bit 25 in some shift register. There was a bit serial interface and it had to assemble the word and then pass the word into the computer, coming from the IMP, and then going to the IMP it had to take a word in parallel out of the host and then ship it bit serial, ok? So you could see it sort of assemble these bits and then it would do the shift and so on. So I pushed these -- I single cycled the way I vaguely remembered him doing it and I watched these bits fall off, and I said: "Ok, so it's broken. I can see that it's broken," and I didn't want to try to get in and fix it. It turns out, if I had opened up the cover, I would have seen a wire come loose at precisely the right place and I could -- I told him this is what I had seen, and -- elegant job, I mean total visibility in what was happening. So that was the first -- we wound up being site one, UCLA was, for some reason, was number one, so that Sigma 7 was host one on the network.

Pelkey: So when did you hook your Sigma 7 to the --

Crocker: Oh, that was a bad day. This was the summer of '69, and site one was scheduled for the first of September; September 1 for the first one, October 1 for the second one, November 1 for the third one and December 1 for the fourth one. I think they were -- one a month was the basic plan. Well, we were having a little trouble with the software. No big surprise, getting our software ready, and I knew that BBN was struggling with some subtle timing problem in the IMPs and they were real worried about whether they were going to make their deadline. I looked at the calendar and I said: "Oh, well, September 1 is Labor Day. There's not problem, because that's a holiday, that's Monday." So I reasoned like this, I said: "First of all, BBN is running late. Second of all, September 1st is a holiday, so it'll ship sometime after that. They'll ship it out here and Honeywell will come in and try to get the machine working, and then BBN will send out a team and see if they can get their software working on top of that, so I got an easy two weeks minimum." BBN fixed their problem, looked at the calendar and said: "We're not going to miss the first date. Air ship the thing," and it wound up on our loading dock on Saturday, the end of August, and on the same day somebody showed up and wheeled it in and plugged it in and the program kept running from where they had unplugged it, and they had a working IMP. So all the reasons why I thought I had two weeks collapsed to about minus two days, and there it was. So that was September '69.

Pelkey: Can't win them all.

Crocker: Wingfield's hardware was working, and only my software --

Pelkey: You and Jon and Vint --

Crocker: Who all did we have writing the software? We did a lot of the design, and I don't remember who coded it. Good question. Meanwhile -- that was about six months later -- in the spring we started this sequence of notes, the Requests for Comments, and they were very unofficial sequence. In September the IMP came and in October SRI got theirs, and I think November 21st, Larry Roberts and Barry Wessler came out, and we ran a demo, an ad hoc demo, in which we had the Sigma 7 at UCLA talking to the 940, I think it was, at SRI, and we were logged in using them, so it was kind of a Telnet-like thing, and Larry was pleased. That was good. Now, we had been working on the design of the protocols, and we had -- there wasn't much organization to all this. It was sort of ad hoc, catch as catch can. We'd start with some very high level ideas, and then when the hardware interface was defined we knew where the bottom level was and we had to get there and we had to do something about it. The notion of layering was relatively clear, but we didn't know quite how to deal with it, and there were also mechanical problems like coordination and synchronization of processes, things like that. I remember going off and trying to find out what the literature was on inter-process communication. Well, there's no problem, we'll just use standard inter-process communication mechanisms that must exist. When I was at MIT I know that the Multics guys solved all these problems and had elegant solutions for everything, so I went and read the critical theses at MIT that I had only glanced at before, maybe Saltzer's thesis or whatever, and they talked about a single bit or two bits that you use for synchronization, a three state process. I said: "Oh, I know all that stuff. We use that on a single machine, but it doesn't work across a network because you don't have any shared memory, and so forth," so that was no help. So we were rolling our own, and we kind of had in mind that we wanted -- we didn't want to orient it toward terminal to timesharing system kinds of things, or file transfer or batch process. We viewed all those as kind of special. We wanted a generic underpinning, but we didn't know quite how to get there. We were struggling with it, so the first designs we came up with, which were, by this time, under some time pressure to have something, we came up with a specialized design that would allow you to log in from one place to another, but that was it, pretty much. We presented this to Larry in December of '69 at Alta, in Utah. There were two meetings. That's right, there were two meetings at Alta. There was a research group that was concerned with Petri Nets, which was a technical thing that was a hot topic in those days, and somebody had called a little workshop and held it in a ski lodge up in Alta. It was nice. That butted up a weekend, and right after that was a review of where we were with the network, and Larry had called that, so it was a clear opportunity for spending the weekend skiing there. I drove my Porsche up, and my Porsche broke and had to be in the shop. Vint and Jon Postel came up for the second meeting. I had gone up for the whole thing, and my style in those days was to believe that airplanes were supposed to wait for me instead of the other way, and they told me afterwards that they had said, flying up, that if I had to fly back with them, they were going to have to run to make the airplane. Now I had driven up and I expected to drive back. My car broke, I flew back with them, and sure enough, we were running down to get the airplane. So Larry listened to this and said: "No, this is terrible. This is -- how am I going to send electronic mail, for example." We said: "We kind of understood mail as a higher level process that needed to make use of these general

mechanisms," and we said: "Well, we haven't got there," and he wouldn't let it go through. He said: "It's not an adequate first implementation. Go back and design it again," so we went back and pushed on inter-process communication as the fundamental underpinnings, and --

Pelkey: Now, this is still just the four universities that were involved in this?

Crocker: By that time -- let's see, it's hard to remember. We could probably go dredge them all out. I think the group was growing by that time. For example, I had met a fellow --

Pelkey: Because other sites were going to know that they were going to come up?

Crocker: Yeah. That's why I said there wasn't any real gap between the fourth and the fifth, sixth, seventh, and so on, maybe a two month hiatus or whatever, but other places knew that it was coming fast, and so --

Pelkey: They were interested in the same thing.

Crocker: -- BBN, Lincoln Labs, Harvard, MIT, I suspect. Paul Rovner, for example, who had been at the graduate student conference, started to show up at some of these meetings and interact. He was a Harvard student, but he was full time at Lincoln, and he was working on graphics programs, so he got involved in some of these discussions. He showed up at the first of those meetings in Utah, I remember, because he wound up later introducing me to my wife, but that's another whole set of stories.

Pelkey: So Larry --

Crocker: So he sent us back to the drawing board, basically, and we then focused on getting out a basic host-to-host protocol and a Telnet protocol on top of that, and a file transfer protocol in parallel to the Telnet protocol, and visions of a remote job entry protocol, which actually did come into existence, and mail protocols on top of file transfer protocols, and this whole hierarchy of things.

Pelkey: Now was there also the concept at that point of time, in terms of Larry was not only -- there was a way of managing these different projects that were going on, or was Larry -- that you realized that you had these higher order things and you also had to have these lower layer things --

Crocker: See, I wouldn't have said it was a management issue. That's not a response I would have given. I would have said it was more on an engineering/design kind of problem, of how do you partition the design in a way that gives you the maximum flexibility, a modular set of building blocks. For example, one of the things that was very awkward was up until that time, if you looked at the design of an operating system in a computer, the view implicit in the design of the operating system was that it was the center of the world, and anything that was attached to the computer was a slave device. You could attach tapes and disks and card readers and printers and so forth, but the initiative was in the computer, and it would say when to talk or when to listen to those things. You try to connect two such devices together, two computers together, and

they only know how to talk to each other as if one is the master. Well, if each is the master then the other must be the slave. We knew that we wanted a much more flexible vision in which the initiative could be at either site, and there would be a coordinated or cooperative model of communication, but the applications that first come to mind, being able to remotely log in to another thing or to move a file or do remote job entry, all are back in the master/slave model. We wanted those to be special cases, rather than being the only thing you could work with, so we knew that if we put those in as your basic model, that more ambitious things would always be fighting against those kind of things, the business of being able to download a program, the business of being able to have one machine act as a surrogate for another. Say you had a collection of machines around the network and you wanted one of them to handle the log-in process and then hand you over to another machine, those ideas -- they seemed obvious to me, that you wanted just the maximum flexibility. The roles should be whatever you wanted to, rather than being committed ahead of time, to be able to configure it. So a general purpose inter-process communication facility was definitely needed, and then you'd build things up on that. How many layers was not a question that we thought hard about. What's now viewed as seven layers was not -- we didn't strive for that or envision it, it was just however many you needed, presumably not so many. In some sense, we viewed it as two or three, if we had tried to count, but having a number wasn't important. The fact that things were built on top of each other, that you had the least commitment at a certain level and then you could build on top of that, and then when a consumer would come along with an application to do and he had to build code, he would tap in at the right level. He would make use of it, whatever the level was that was most appropriate, and if he had to dig down below, extra cost for him but he wasn't locked out, down to whatever layer. So that was the motivation, and I guess you can view -- it's a very generic kind of motivation. You can view it as a management motivation, but I would say it was more a clean design and modular building block kind of thing. Then all of the advanced ideas about how you get computers to cooperate on various tasks were -- the fundamental elements were there in place, and you could build on top of that. Those were the days when, instead of viewing the network as an electronic mail system, which was kind of an afterthought in a way, there were all these visions of shared databases and load balancing, or jobs would be shifted from one machine to another to balance the load.

Pelkey: E-mail, while it was at some level defined, at least in Larry's mind, for most other people, mail was an afterthought.

Crocker: Yeah, and frivolous.

Pelkey: And it turned out to be --

Crocker: Nearly everything. If you look at the amount of traffic that's sent and you try to ascribe it to E-mail versus other things, E-mail is the vast -- there are some, actually, I now see some more subtle reasons why that's true that have to do with security and privacy. E-mail turns out to be the only mechanism that allows communication, private communication between parties that are not previously connected to each other. If I want to move a file from one machine to another, I've got to have an account, I've got to have a password, or it's got to be wide open and nothing -- same for Telnet. Mail is not that way. Mail is - - there was implicitly a security mechanism, not one that you couldn't crack, but things went from person to person, and

not to everybody else. I mean, if I send you some mail, it's not the norm that everybody can read it, and it's not the norm that I have to know anything about you ahead of time. I can send it to you cold and you can send it to me and so forth, so there's some level of what we now call authentication.

Pelkey: Now, you somehow became a senior person at UCLA.

Crocker: Well, the group wasn't all that big.

Pelkey: Right, but you were the leader.

Crocker: Vint was much more disciplined and was getting on with his own research and getting through school, and I sort of fell into this business of getting involved with the network in a more direct way.

Pelkey: How did Jon Postel get responsible for the RFCs in the early stages?

Crocker: Well, he wasn't at the early stages at all. I was.

Pelkey: When did he become responsible?

Crocker: What happened here. Vint graduated. Postel graduated. Let me untangle the sequence. In summer of '71, I got invited to go to ARPA. No, I'm sorry, it wasn't that late. It was summer of '70. There was a year that transpired, almost a whole year from the time that I was invited until the time that I went, so things actually happened in fairly quick order. The Network Working Group became defined as an entity, in some sense. We met at the Spring Joint Computer Conference in Atlantic City and presented a paper on the host-to-host protocol in the spring of '70. So, the first IMP went in in '69, and this is now the following spring, so not a very long time. We wrote a paper. There was a session, actually, a whole session unveiling the Arpanet, and there were a couple of papers from BBN. Larry Roberts gave the lead paper, and I gave what amounted to the runt end, the last paper in the sequence, about these host-to-host protocols. I remember Frank Heart put up some diagram that had sort of the spectrum of the problem, and the host-to-host protocols were a thin little sliver on one end and the data communication, the long distant telephone lines were the other thin sliver, and the IMPs solved this problem that was really right in the middle. I remember looking at it and: "Well, that's your point of view," and from my point of view, all that lower level stuff was the bottom layer, and then we had this vast territory to try to untangle. We had a big Network Working Group meeting in conjunction with that, so that was spring of '70. By this time the group had grown to a number of sites. We had a meeting at UCLA. I'm trying to remember when this meeting was at UCLA. It was earlier and it was a bit of a disaster. People from MIT came for the first time, and we had this host-to-host protocol that we were telling them about. There was a feature in the host-to-host protocol that was aimed at being able to shift a connection from one machine to another unilaterally. Well, it went like this: Let's suppose that A and B are talking, there's a connection established between A and B, and let's suppose that B is also, at the same time, got a connection established with C. Now, the program that's running on B could go into a mode in which it reads from A and copies it out to C and vice versa, and so it simply acts as a conduit, copying from A

to C and from C to A. Now, if it just did that, it's a bit inefficient in that, geographically, it's a distributed data copying process. We envisioned that there might be a reason why this happened a lot, in particular, suppose you wanted one machine to act as a log-in process and have a database of who the appropriate users were, and then deal that person off, for load sharing purposes or whatever, off to another machine. So it wanted to say: "Ok, I've authenticated who you are. Send all of your further messages over to this other port, and you'll get all of your responses back." Well, it could simply do it itself, but that's a high overhead sort of thing, so we built in, as a primitive function down at the host-to-host protocol, that B could identify two connections and tie them together by telling A that, in the future, it was to continue its conversation with C, and vice versa, and not only just with the host in general, but with the specific port, or specific socket as we called it. So it would introduce A and C together, and it would do this after the connections were established, and sort of at an arbitrary time. That posed some fairly hairy synchronization problems, because B would have to send a message to A, saying: "Hold up a minute, you're about to be introduced to C." Then it would send a message to C saying: "Hold up a minute, you're about to be introduced to A." The only problem is that C might be connected to D, and about to introduce B and D together, and when you've thought your way through that problem, now consider the problem that D might be connected to A and about to introduce C and A together. So there could be a ring, and worse yet, A could be about to introduce D and B together, so they might all drop out. So I had what I thought was kind of a nice little solution, that a message would go from B to A and a message would go from B to C, and when it got there, if they were not simultaneously in the process of trying to drop out of the loop, they would freeze and reply positively, and they would get introduced this way, and if they had simultaneously initiated the same thing so that there were messages in flight that crossed, then C would know that it was in that state, and it would take the message that it got from B and pass it on to D, saying: "I know you want to introduce me to A, but you should really talk to D," and that message would keep going, and it would finally hit the end, in which case it would bounce back, and wherever it would find the end of the chain, and that would bounce back and it would bounce so that A and D would find each other. Well, if there was a loop, in fact, then it would go all the way around and it would come around the wrong way, and so B would see that this message that it sent out in one direction was coming back the other direction, and it would know, therefore, that it had a loop, and all the other sites would have done the same thing concurrently, and they would all arrive independently at the same knowledge, and they would realize that they could all drop out, that nothing further was going to happen. They would just close off their connections. There was a lot of controversy over this, because it was somewhat more complicated than the other parts of the protocol. There was quite some resistance, and the Multics people, surprisingly, were the most resistant. They said we couldn't implement things like this, saying that many of the ideas in the protocol I had thought were Multics-like ideas, where we had sort of taken Multics as the role model for what a real multi-purpose computer, multi-user modern computer was, and had adopted its inter-process communication ideas as best we could. The people sitting on IBM 360s had had no facilities at all for inter-process communication had no problem, the Multics people said: "No, no, we can't implement anything as complicated as that," and it turned out that there were pragmatic issues about who was allowed to program at which different level, and the people that were assigned to this task were locked out of the lower levels, and we were sort of envisioning the whole abstract thing, and so it was educational. Eventually, there was enough controversy over -- that portion of the protocol was actually overruled -- Barry Wessler came and visited me and said: "This is not going well,

drop that portion of the protocol," and gave me an order, basically. On August 5th, 1970, I remember quite well, so we dropped it, and the function that was going to be performed by that got put in at the next layer up in a somewhat different way. I would have viewed it as clumsy at the time. There's ways in which what I was proposing was very clumsy, so those are the kind of issues that are never perfectly resolved, but an interesting piece of humor. I got into a conversation with Will Crowther who was the chief IMP software designer about all of this, and he said: "Why don't you just, with respect to the issue of detecting the next host might be trying to do the same thing at the same time, why don't you just have the hosts send their message, and if it gets to the next host and the next host is already trying to reconnect -- we called it the reconnection protocol -- then it would just reply negatively, or it would just forget about it, and if you didn't get a reply, you would just try again, and eventually, somebody would succeed," and I was astounded. I was just shocked, not in the negative sense, but just couldn't imagine where an idea like that would come from, because it was statistically based, and it wasn't the way I thought about things. I thought about things as interlocking puzzles, and you get all the pieces to fit exactly right, and the idea of a statistical process, which is pretty easy to see that it had a high chance of converging very rapidly, so it wasn't a problem, but it just wasn't the way I thought about software at all. I couldn't imagine where he could think of an idea like that, because it wasn't the space of ideas, it was more a limitation on my part. All of a sudden, I had a piece of insight. I said: "What did you major in in college?" Sure enough, he said: "Physics." See, I was a math major, and physics was really not a thing I focused on very much, but it was a viewpoint of how the world works, and I viewed it more in algebraic terms, and he viewed it in more -- Heisenberg was a real idea to him. So I was surprised. Then later --

Pelkey: How did you implement it?

Crocker: Well, we threw the whole thing out from that point of view.

Pelkey: You know, that surprises me that you make that statement, because I'm led to believe that Kleinrock was a statistician's statistician at some level.

Crocker: Kleinrock --

Pelkey: At this point in time you're already working (unintelligible)

Crocker: Absolutely, except we weren't talking to each other. It wasn't a matter -- I say that and it sounds like it's a hostile thing. These were detached processes. Kleinrock was doing fantastic stuff in queuing theory and analysis of loads on networks and he had a stellar set of graduate students doing theoretical kinds of things, and I think he was having a pretty dramatic affect on other aspects of the network design, and he was certainly supportive of what we were doing. My travel, for example, shot upward dramatically, and he was supportive of that, and he -- Roberts supplied more money, so there wasn't any tension, but there was, in effect, two completely disjoint efforts. He did not look very closely at what we were doing. He never engaged in scientific or intellectual discussion about the protocol development, nor was I burdened with having to learn the wonderful things that he was doing, which probably would have been helpful to me. He did look in at the operating system work that we were developing once and he asked a question having to do with capacity, how many terminals we could support,

and I actually didn't understand that you could ever run out of capacity, in a way, so these were not tightly organized activities.

Pelkey: And you're now, at this point in time, you're traveling to many universities --

Crocker: Yeah, going around a lot. In fact --

Pelkey: You're really building up your frequent traveler miles at this point.

Crocker: There weren't such in those days, although United had a deal where you could join their airline club if you could document 100,000 miles, free. Different rules in those days, so I had no trouble filling that out. I was invited to speak in Waterloo and in New York and in various places. I was having a great time.

Pelkey: So the summer of '70, you de facto have become the head.

Crocker: That's all that happened. Somewhere in May, I think the issue came up and sort of put to Larry in some meeting that there was a group wanting to know how official all this was, and he said: "Well, what's your problem. Your in charge," but it wasn't as if it was something that had to be decided, that's all there was.

Pelkey: Now, Larry must have been putting pressure on you. Was it Larry or Barry who was saying: "Come on, what the hell's happening?"

Crocker: Well, primarily Larry, but it wasn't very frequent. He was a relatively taciturn guy.

Pelkey: Once a quarter?

Crocker: Well, not that much. Two or three times, maybe, during the whole process, that I recall. He was -- first of all, I was a mere graduate student. He was probably talking much more to -- he had a lot of other things going on. In August, the rule came down to drop this reconnection thing. Later in August, I was in Washington, and I got invited to come to the ARPA office and it had been cleared to several people. Cordell Green was, by this time, in the ARPA office. He had ROTC obligations and they managed to steer him into the ARPA office. I remember I stayed at his house, and he kind of alluded that good things would happen. I would come to the office and I might get offered a job. He didn't want to make a -- Larry offered to come work sometime in August, and then -- I made a big deal of thinking hard about it, and my view was I was interrupting my graduate studies and, from my point of view, it was part of my studies. This was an educational opportunity more than anything else, and I said yes, but the paperwork, security clearances and hiring freezes and all kinds of things, delayed it, so I didn't show up at ARPA until '71, July of '71. By this time, I had spent pretty much the whole year running the network working group, helping define the protocols. The group had grown to hundreds of people by this time, a small number of hundreds of people. Meetings got larger and larger -- what other critical things happened there.

Pelkey: When did you meet Metcalfe?

Crocker: Well Metcalfe was working at MIT.

Pelkey: He was building the IMP interface.

Crocker: He was building the IMP interface, and he became one of the people that came to these meetings. Then what happened was that it became clear that there was an education problem, that many sites needed help getting on the Arpanet, and there was a learning curve issue, and also there was a big difference in experience. Some groups had people that could pick up the ideas very simply, very easily, and go right to them, people like Wingfield that could go build an interface in spare time. Others really needed a fair amount of help. So, a few people, Postel, Metcalfe, Rich Kalin from -- he was at Lincoln at the time, I think was a Harvard student -- and I can't recall, but about a half a dozen people I designated as "facilitators". Now, they also seemed to have enough freedom in their own environments so that they could participate in this.

Pelkey: When did you decide to do that?

Crocker: Well, somewhere in that general 1970-ish time frame.

Pelkey: Was this independent of the scenarios?

Crocker: Yes.

Pelkey: Had the scenarios been defined?

Crocker: I'm not 100% sure that I am totally in sync with you about the scenarios.

Pelkey: (unintelligible) this conference that --

Crocker: The ICCC in '72. Totally different. The scenarios were an artifact of that conference, so I'll get to that --

Pelkey: But Bob Kahn --

Crocker: Bob Kahn did it.

Pelkey: There was a lot of pulling in this to get -- some people perceived that there had to be something to get all these guys to see this --

Crocker: We had a prior form of this, where we had a big Network Working Group meeting at MIT in September or October of '71, I think, and the game was: everybody was to try to log into everybody else's host, so we had a big matrix, and in the course one day or two days, people were going to try to initiate a Telnet connection from their host to the other host, in both directions, because they are asymmetric when you get to that level, and SDC distinguished themselves by being completely out of it and having nothing ready, and everything else kind of worked. The matrix was relatively well filled in, and SDC's absence became notable only

because there was a big hairy empty row and an empty column, but an awful lot worked at that time. The ICCC was more of a -- that was essentially an internal, inside the community push to get everything up and running, and at that time, you could sort of say that the network was there. Now, somewhere along the way, the TIPS, now known as TACS, were defined, and BBN built them, and they embodied this host-to-host protocol inside of them, and the Telnet protocol as what they were doing. I'm talking too much.

Pelkey: No.

Crocker: Bob Kahn was invited to come down to ARPA, and he didn't come until -- I showed up in July of '71. He showed up in November of '72, so a year and a quarter later. ICCC, International Computer and Communication Conference, happened in October, as I recall.

Pelkey: Right, of '72.

Crocker: He had a major role in organizing this big, very flashy demonstration of the Arpanet, and it consisted of having a TIP rolled into the ballroom floor of I think it was the Washington Hilton, and vendors were invited to bring all their different kinds of terminals. They were all connected, and sites on the Arpanet were connected -- this thing would be connected into the Arpanet, so AT&T had to supply the long distance lines. The lead time on 50 kilobit lines in those days was pegged at nine months, sometimes you get them a little sooner, but long lead time items. AT&T was very cooperative and managed to get that put in, and they took this on and so forth. I didn't have a great deal to do with it. I was fully busy with other things, and Bob was pushing it. I went to a few meetings, and an awful lot of the people I knew got dragged into it, but by that time I was involved with a lot of other aspects inside the office. It was a huge success, and the scenarios were, now that you've got everything connected up -- you've got a terminal connected to a TIP, which is connected to the Arpanet, which then can go connect to all these different hosts, what are you going to show the people? So a series of demos, the scenarios as you call them, were put together, and I guess there was a -- I vaguely recall a notebook so that people could walk in and --

Pelkey: I'm led to believe that all of October November of '71, when you went through this exercise --

Crocker: That was a Telnet bake-off --

Pelkey: Telenet bake-off.

Crocker: Telnet versus Telenet are two confusing terms. Telnet's the protocol, Telenet was the company that was formed.

Pelkey: Telnet bake-off that -- in '72 all the IMPs were talking to each other, sites were talking to each other and using the system.

Crocker: I don't know what the traffic level was. I don't know how often things were working, but a substantial amount was in place by that time. There was at least a core of sites that were on the network all the time and working, and things were in place.

Pelkey: If I could, unfortunately, I have about 20 minutes left, but I am interested in the Metcalfe incident.

Crocker: Oh, ok. I'll tell you that -- let me just get myself out of this. When I got to ARPA I got swallowed up into a whole set of things, and one of the things that was happening was that, building a network became a national objective for a number of countries: Canada, UK, France, among others. The Arpanet demonstration was spectacular, in the sense of causing changes throughout the world.

Pelkey: This is the '72 showing.

Crocker: Well, I don't think of that particular showing, but -- in fact, the events I'm describing happened even before that, in '70 and so forth. It is literally true that building a network and having something comparable to the Arpanet in some fashion was a national priority in several different countries, so I naturally find myself traveling even more, and I wound up going to Europe a couple of times; three times in the course of one year, I think. I had not traveled much internationally. I found it overwhelming, and time consuming and debilitating, and I really went to ARPA, as I had said originally, because I was interested in the artificial intelligence and stuff like that, and there were a lot of exciting projects. Speech understanding was just starting up, and the network was just a thing I was doing on the side because it was there and I needed to be helpful. It is the reason they invited me to come, but it isn't the reason why I went. So Vint had finished his thesis, had gotten invited to teach at Stanford, so he went to Stanford, and I turned around shortly after I got to ARPA, realized that this travel was killing me, and called them up and said: "Look, Vint, I've been running the Network Working Group. Why don't you be the International Network Working Group," and said: "You do it," because he had spent his junior year abroad in Germany and he was an extremely literate, capable guy in multiple languages, and I just knew he was far, far better at handling all that stuff than I, so I just turned around and said: "You do it. Here's a few bucks." I liberated a few bucks from the Stanford contract, and away he went. He went tooting off on that and did a spectacular job, and I concentrated primarily on the other topics, and handled a bunch of networking stuff because it was useful to do. Many stories I could tell, but I'll cut that.

So the Metcalfe incident. I had -- Norm Abramson came into the office one day -- I hadn't known about the ALOHAnet stuff, and he described it, and again, I guess it's the same reaction -- I hadn't put the two together before this -- but the same reaction I had when I listened to Crowther's. He described this computer and the terminals and outbound from the computer to the terminals, radio messages addressed to a particular terminal would get heard by that terminal as well as everybody else, but they would know who it's for and they would print it. Inbound, the terminals would just fire off messages whenever they wanted, and of course they would collide in the air and disrupt each other, but if they didn't get a positive acknowledgment they'd retransmit, and there were some statistics that go with that, and you only gave up a certain fraction of the bandwidth for doing that, and I was just astounded. The boldness, the sheer

brashness of that idea -- and of course it would work, once you say it. I remember trying to absorb it and understand where the thing was that would sort out who would get and get in what order, which is typically a multiplexer, and I said something like: "The medium is the multiplexer in this case," sort of a McLuhanesque statement, so that was picked up and quoted a couple of times. So I thought that was pretty exciting, and again, that was outside of the space of ideas that I would traditionally invent for myself.

Pelkey: Who was it that came in?

Crocker: Norm Abramson, who was a principal investigator and had been around, he and Frank Kuo were --

Pelkey: And that was an ARPA funded contract?

Crocker: That was an ARPA funded project – there were great things happening a lot of places. That must have been later. They weren't represented at this graduate students' meeting in '68, for example, so this was a couple of years later. So Metcalfe and I, by this time, had gotten to be friends, and would pal around whenever we could, and he had been at MIT working, and had been a Harvard student, and he went from there to take a job at PARC. He came out and visited me one day when he was at PARC. In fact, I'm trying to remember -- there was an interesting day when he came out and visited Bob Kahn. Bob Kahn had a house for or five doors away from me, around the corner, and Vint Cerf came out and visited me on the same day that Metcalfe was coming to Kahn's. It just happened. They arrived about an hour apart, but they were on the same airplane, and it turned out that Vint was traveling university mode and flew coach and traveled with buses and so forth; Metcalfe traveled in the front of the plane and rented a car the minute he hit the ground and showed up just lickety-split like that. It was just very interesting, the difference. Metcalfe was visiting, and I guess I wasn't keenly aware of it at the time, but he claims to have gotten the key idea for the Ethernet, which is again this contention-based protocol. It's a symmetric version instead of an asymmetric version, where all the hosts were both senders and receivers, this is all well known now, and went off to PARC. PARC was a very rich environment; a lot of very bright people, and he was able to put the project together and get people to cooperate, demonstrate the thing, and change the world.

Pelkey: He recalls that night staying at your place, and this Abramson article was on your coffee table or something.

Crocker: Maybe so. Maybe I was just excited. Maybe because Abramson -- I don't actually recall.

Pelkey: Do you recall him saying anything to you?

Crocker: No. I don't remember -- I remember later that he said that was the day it was invented. As I say I got to be host of the father of the Ethernet -- a very tangential role.

Pelkey: From ARPA, you had launched, now, Vint as coordinator for the International Working Group, and you went off and stayed at ARPA and did your thing, and you started pursuing your own interests at this point in time, which was artificial intelligence -

Crocker: Program verification, in particular. I stayed at ARPA for three years, and then I came out to ISI and -- you see, I had been a little tardy about finishing a dissertation, and so I wanted to do that. All the time I was at ARPA I carried a registration card for UCLA and paid fees and I wanted to distinguish carefully in my mind from those people who had dropped out of school and said: "Well, I can always go back," and never did. So I viewed it as a hiatus in my education, or sort of a part of it, and in fact, a kind of special experience.

Pelkey: So your involvement in the data communications business, although you stayed involved because you were part of that crowd --

Crocker: It took a long time to lose the label and -- Now, I went to Aerospace and set up a computer science laboratory, and, among other . . .

Tape Side Ends

Pelkey: . . . The modems and multiplexers, those guys really didn't pick up much on all this. Why do you think that was?

Crocker: I don't know. Why didn't they pick up on it? I think the commercial people worry about where the market is and how to position whatever it is they're going to build so that the customers are going to buy it, a very rational concern, and you can't go faster than a certain amount because customers aren't ready to go more than a certain amount. So it's a kind of evolutionary process, and in order to cause a revolution, you've got to have, effectively, complete solutions. If you're going to design something like an Ethernet, you need the computers that can talk to it and the interfaces and the software and the whole thing. It takes a long time. The modem people were selling modems and a big deal to them was to go from 300 to 1200 baud and from 12 to 24, 9600 and so forth. They couldn't redefine the bandwidth of the lines. The Arpanet came along and said: "We want 50 kilobit bandwidth and we'll spend what it takes to get that much." It was a huge, huge amount of money, but it wasn't the norm. It won't be the norm actually, until 56-kilobit service digital is available cheaply. We had to -- setting up this office, we connected into the Arpanet; I assumed we'd get a 50-kilobit connection over to the IMP that we were assigned to which is at UCLA. Called up GTE to get the data service and they wanted a fortune for a digital link from here to there, and the Arpanet is so heavily loaded now that, in fact, you don't get the throughput that justifies that, so with great pain in my heart that I would downshift to anything as simpleminded as 9600, but I asked how much for that. Thirty dollars a month for a 9600 baud line, leased line, from here to UCLA, and it was like $1500 a month for -- maybe a little less, it was a factor of more than 30 for bandwidth that was, at best, only five times more, so it was grossly disproportionate, and still, and this is long time later, now twenty years later, and still don't have that kind of service. So, I don't know, for the marketplace, the thing's got to move in sync. It can't get too far ahead, or the infrastructure that supports the product isn't there.

Pelkey: One of the other strange things is that University of Utah had very little impact on the communications and networking. Some people believe it's because they kind of imported the software from -- I guess they got Tenex from BBN.

Crocker: This is a strange question in a way -- why would you expect them to have?

Pelkey: Well, UCLA was so active. UCLA was clearly the most active of the first four nodes.

Crocker: Well, I think that's the wrong question. Let me propose to you the following: I think you're putting two things together that don't necessarily belong together. The first four sites were selected primarily, I would suggest, and again this is a thing you'll get different opinions about because I didn't participate in selecting, but I didn't see anything that motivated the first four sites except geographic proximity and being major research laboratories in the ARPA supported community, which meant that all the political issues about whether they would participate and who would pay for it and all that was subordinated; it didn't come up. That's the end of that story. Now, when you look at the four in particular, you say: "What do they do in real life?" Getting connected to the Arpanet was a side issue, ok? UCLA had Kleinrock doing interesting things in networking, SRI, it was Englebart's group doing augmentation of human intellect, doing some very, very important things in man-machine communication stuff, which led to the creation of the mouse, Santa Barbara was Culler-Fried, basically, and doing things that were, from the interactive systems point of view, were important, but they were also doing signal processing and some speech work and things like that; and Utah was doing very important things in graphics with Evans and Sutherland and so forth. Well, of those, the ones that were directly applicable to the Arpanet were Kleinrock's stuff in the one hand and the SRI Englebart stuff, but there wasn't any reason to expect, I wouldn't think -I think you should come at it the other way; you can name almost any of the major institutions and they were having an effect on whatever it is they were doing, and they were incidentally getting connected to the Arpanet. In fact, Tom Stockham was at Utah and did spectacular work in speech and digital restoration of Caruso records, all kinds of things, and I remember, Bob Kahn and I made a sojourn; we made a trip across country in the fall of '70. We wound up with BBN trying to write up these memos of the places we had visited in early November of '70, and I remember sitting with Stockham. He knew, by that time, that I was going to go to ARPA, and he put to me -- he said: "What are you going to do when you're forced to make a choice between networking research and artificial intelligence?" He was basically asking: "Are you going to let this networking stuff interfere with real work?" He was testing my values and so forth, and I assured him that I agreed, that artificial intelligence was far, far more important than pushing a few bits around, and although I didn't think the question was going to come up in that form. I didn't think it was zero sum game, and I guess he was satisfied. It didn't matter so much, but even then, the community's attention was not turned entirely toward networking as the only thing. There was a broad spectrum of very important research going underway, then and now, and the networking, the problem of managing a success and keeping it from turning into a disaster. So I think it's an unfair question, in a way, to ask: "Why didn't Utah do more networking?" You could ask for any institution that's doing very well in some area: "Why aren't they doing -- "

Pelkey: My question wasn't to imply that they weren't doing -- my focus is on communications, and --

Crocker: The fact that they were in the first four wouldn't think -- it would be an entre if they chose to do that, or if they had the people that -- but, now Steve Carr was heavily involved, and I haven't kept up with him, but he went into doing some of that kind of stuff, and a few -- there was a pocket of activity, but it's not a necessary, or even necessarily desirable -- you certainly don't want everybody to have that happen to them, so it's relatively important, actually, that there were places that were practically unaffected, that became consumers. What they provided to the world was of a totally different nature.

Pelkey: Another kind of a question. I have a view that two important things went on in the communications domain: one is ARPA went from ARPA to DARPA with the Mansfield Amendment of '71ish?

Crocker: Well that wasn't the -- the name change was, as we say, a hack. ARPA had been part of the office of the Secretary of Defense -- this is an arcane piece of bureaucratic lore. The office of Secretary of Defense not only includes the Secretary of Defense and a very substantial number of people -- I think a few thousand people that sort of work for him in some sense -- while their was a manpower push, didn't have anything to do with the Mansfield Amendment. Mansfield Amendment had to do with relevance, that was a different axe to grind, but there was just the traditional hiring freezes, trying to trim the budget and so forth, and ARPA was operating as a piece of the extended office of the Secretary of Defense just like you have the Office of the White House, or the Executive Office, which has got two buildings full of people, all sort of besides the president who nominally runs it all, and so it was decided at some level, way outside of anything that involved us directly, that ARPA would be moved out of the office of Secretary of Defense and be called a Defense Agency, and that would put it in the same status as the Defense Mapping Agency and the Defense Nuclear Agency and some others, and I'm not precisely sure about the nuances of all of that, but on paper it meant that the director of the agency reported to the Secretary of Defense, instead of up through some other channel. In practice, the reporting chain was right through the same people, but it's a rule that if you're a defense agency, your name begins 'D' for Defense, so ARPA became DARPA. Lukasik was the Director of the Agency, and then the change happened on the first of July 1972, and he got us all together and he said that he expected the effect to be 5% measured in some dimension. He was trying to say that it was pretty nominal, and among other things, it meant that he could sign his own travel orders and some simple-minded things.

Pelkey: How about that? My understanding was that there was a change because of the Mansfield Amendment relative to the charter of investing of ARPA/DARPA, that went from more fundamental research, which was clearly the case in the '60s, to much more defense related, more immediate project oriented than fundamental research.

Crocker: Others can tell you. I was pretty junior and protected a bit from that. The Mansfield Amendment said that agencies like the Defense Department, I guess in particular, had to be mission oriented. The work had to be related to what they were doing, as opposed to -- you couldn't justify the research just because it was good research. That was the NSF's job. It certainly was another wicket to be careful about, just like Environmental Impact Statements later

and all sorts of things, but it didn't really change anything very -- I guess the word change is what's hard to --

Pelkey: It wasn't a sudden change, but over the '70s, the issue became increasingly -- it kept becoming more and more important how to get around this definition.

Crocker: I guess it had an affect at two levels. First of all, any specific effort, you had to be careful that there was a documented relationship at some level; maybe a sentence, maybe a paragraph, or maybe something larger. The other is that, in terms of gathering a budget together and building a research program, the kind of work that had to go with it and the things that would sell versus -- tended in the direction you're saying, to be more application oriented. But there was another affect that was much more dramatic happening at the same time, and that is success. The ARPA office was a small office that was originally a so-called 6.1 office; it had basic research funds. The Defense budget has several major line items. Line six, at the top level, is research, development, test and evaluation, and under that, 6.1 is basic research, or just research, and 6.2 is development, and 6.3 is test and 6.4 is evaluation, and the amount of money available in those four categories increases substantially as you move from one category to another. Now, IPTO had been a 6.1 office, and one of the ways of increasing the amount of money that you could spend on computer science issues was to create a 6.2 budget, which is an applied budget, and that's what happened. So there were -- as the office grew, a much more substantial fraction of its budget came from funds that were earmarked, independent of the Mansfield Amendment, just part of the normal way of the world, to be applications oriented, and therefore less researchy, and the Arpanet certainly was a big consumer of money, and therefore had to consume a lot of 6.2 funds.

Pelkey: Now, on the other hand, divestiture came, and the charter of Bell Labs is going to change too. It's going to be much more commercially driven than being a research center for the United States.

Crocker: Perhaps.

Pelkey: So the places where we as a country are investing in fundamental research, be it in specific professors or universities working in small groups or on larger managed projects which is more what the ARPA mode was versus the NSF mode, and if Bell Labs is going to become much more commercially oriented, IBM spends a little on research but, the industrial sector doesn't spend much on research. If, in fact, research 20 years later leads to economic products and so on, the United States competing in the economic order of decades from now, our ability to commit and manage research strikes me as a very important issue, in terms of future economic growth -- not this next decade, necessarily, but in the future. If we're going to compare to Japan or Europe, there's seemingly much more commitment -- I don't know the dollar amount -- to fundamental research as a basis on which they can start to catch up with the advantage that we currently hold.

Crocker: This country -- schizophrenic is to strong a word, but it certainly has at least two points of view about things. When it gets scared that the research base and academic base and intellectual base isn't strong enough, then money gets pumped into funding these things, but

fundamentally, I think the United States views itself as overwhelmed with short term problems; an economic system that is in debt and getting worst, employment problems, inflation, and things like that, and losing its competitive edge in the marketplace, and driven by short term issues. Where is it going to find the money to spend on research, just as where is it going to find the money to spend on social programs or anything else. I don't see -- the pendulum has been swinging from the '60s pretty much uniformly to viewing research as too expensive, sort of a luxury item, and maybe I overstate the case and in some ways I have not been tracking research budgets carefully. It's not quite as dreadful as some other areas where people see their whole programs go down the drain, but I don't see any significant renaissance where we have a flowering of basic research driven from the top level. It certainly hasn't been on the agenda for this administration, and I don't think it's going to be on the agenda for any of the next few. Bob Kahn's initiative, his National Research Initiative, is a very ambitious attempt to try to attack that directly, and it's such a different kind of thing than the country has ever seen before, I can't say whether it will be successful or not, but certainly that embodies the kind of frustration -- is a response to the frustration that you can see when you see the major successes and how important they are, and you don't see the mechanisms in place to support that. That's one of the things that has made ARPA so special, is that it's survived, and even ARPA's gotten -- it's a very, very good place, but if you compare it to where it was 20 years, it has a whole different feel; much more bureaucracy, much more careful.

Pelkey: Steve, thank you very much for your time. It's been a real pleasure.

END OF THE INTERVIEW

