



30th Anniversary of Internetworking using TCP: Session 2: the roles of BBN and NDRE

Participants:
Barbara Denny
Paal Spilling
Virginia Strazisar Travers

Moderated by:
Vint Cerf and Marc Weber

Recorded November 8, 2007
Mountain View, CA

CHM Reference number: X4307.2007

© 2007 Computer History Museum

Marc Weber: Barbara Denny, Paal Spilling and Virginia Strazisar Travers and we're talking about their role in the three network transmissions in 1977 and also in internetting in and around that historic event. So You want to just start with Ginny, talking about when did your involvement start?

Virginia (Ginny) Strazisar Travers: My involvement started in April of '75 when I came to BBN and I pretty quickly got involved in doing the work on the gateways, which were the machines that connected the networks. That was really what became the internet, was the ability to connect multiple networks. We had an experimental network called the Resource Computer Network at BBN and we had a gateway there between that and the Arpanet. And then I installed gateways between the Arpanet and the Packet Radio Net and between the Arpanet and the Atlantic Satellite Network.

Vint Cerf: Ginny, the voice you hear off camera is Vint Cerf. When you were doing the work on the gateway, Bob Kahn remembers that gateway code was in the station originally and then somehow migrated out. When you got to BBN were you working on a completely distinct device that was acting as a gateway?

Travers: Yeah the-- for the gateways between different networks were generally done in a completely separate device and it was a PDP-11. But for the Packet Radio Network, we combined the Packet Radio station code and the gateway code into a single PDP-11. So that was the distinction. But the other gateways, the ones between the satellite network and the Arpanet were always separate PDP-11s. So they communicated on one side with like the Arpanet and the other side with the satellite network and directed the packet between the two networks.

Paal Spilling: But that wasn't one of the original design, at least the original idea to have a combined SATNET and second node and a gateway. But then BBN decided to split.

Travers: That could have been the case. I mean even within BBN, there was one group of people who were doing things like the Imps and the Tips and the satellite Imps. And there was another group of people which I was one that was doing you know, the internet protocol work and the TCP work and the gateway work. So I can believe that there would have been a split there.

Cerf: Well in fact, it's Vint again. My recollection is that before Bob insisted on-- Bob Kahn insisted on separating the satellite network from the interior of the Arpanet, there was a very strong preference in one part of BBN to integrate everything together into one common set of devices.

Spilling: That was my intention.

Cerf: But that would hide all of the interfaces in effect and the only thing you would ever see would be the external interface of what in effect is the Arpanet. And Bob wanted very much to have an externalized and open interface that other people could build to, that's where Internet is different. And so I could easily see he had a potential conflict there but in the end, the people who were paying for the project got to decide what the architecture was.

Travers: I mean I very distinctly remember it as being separate because I did the software and somebody else was doing the hardware end of it. And so we were trying to coordinate these trips so that

like for example, we got the PDP-11 installed in London in December '76 and then I came and did the software. And then what happened when I went to Oslo was, there was a-- I think there was a TCP/IP meeting there and then they were supposed to install the hardware and then I was supposed to install the software. And I think that there was some glitch somewhere along the way with the hardware, because I ended up staying with you for a few days and then I went off and toured Norway and then I came back and installed the software. So I have fond memories of visiting with Paal.

Weber: And you-- give a little bit of your background. You were just out of college when you took the job Ginny?

Travers: I went to Case Western Missouri and I graduated with a bachelors degree in '74 and then I went to work for MITRE for 10 months. So in that job I had a little bit of exposure to the network world. But it really sort of kicked in gear when I came into BBN in April of '75. The thing was that Case was one of the first-- not one of the first but one of the early nodes on the Arpanet. So that by the time I got to BBN, I already knew people like you know, Jerry Burchfiel or Ray Tomlinson or that sort of if nothing else from all of the reports in that because we were running a TENEX system on the Arpanet at Case.

Weber: So and you went around the world installing these gateways in which period of '76?

Travers: I started, let's see the best of my recollection, I installed the work here probably in the summer of '76, installed the one in London in December of '76, installed the one in Norway in, would have been summer of '77. So that would have paved the way for the three-way demonstration in the fall of '77. And then that was the last of them that I installed. I think that subsequently there were gateways that went into Italy and Germany. But I became involved in other projects and so those were the ones that I had a part in.

Weber: And so Paal, tell a little bit about the-- when Ginny arrived, what was going on in your institution and how did she come to stay with you?

Spilling: Maybe I should give a little background from--

Weber: Definitely.

Spilling: -On myself first. I'm educated as a nuclear physicist and did my PhD in experimental nuclear physics in the-- in Holland. I started to work for NDRE in early '74. Oh, sorry, in '72. And then after a few years I was asked to join the Internet (or you can say the computer communication) project which NDRE would like to participate in. And so that was my sort of starting with computer communications without any background at all. I've just seen a computer. I hadn't done any programming except a little bit around modifications to a program I was using to analyze nuclear physics experiments. So it was a very hard time to catch up with all these ingenious guys. But I found the group which we worked with, very stimulating. I mean it was sort of kicking my back the whole time to try to be in synchrony with the rest of the group. Okay, I think that's sufficient for my background. Then in end of '74 I was asked to join this-- the packet thing project and started to work full time from the beginning of '75. But then that was partly on catching up with TCP/IP and next to participate in the packet switch satellite network development. So I had to devote time to both of these. And fortunately I was in London in October '75 for two months and overlooked the initial TCP/IP test between Stanford University and University College in London. They

had two separate implementations. And but I think they had a little smaller difficulties in getting TCP to establish a connection. And one problem I remember was that Stanford had a different way of calculating the checksum than UCL. So in the beginning all the packets were just thrown away because the checksum didn't fit. But that was easily solved and then they were able to establish their connection. So with that was wonderful to witness.

Cerf: Paal, I'm wondering, it's Vint again. My recollection of some of this is like yours. That the checksum calculations weren't always the same-- implemented the same way. When Jon Postel was doing bakeoffs where we would get, you know, a dozen different implementations all trying to talk to each, we first started by turning off the checksums to see how many could actually communicate. Then we turned on the checksums and all but two of them didn't talk to each other anymore.

<Laughter>

So we asked the two guys who could communicate, exactly what algorithm they had used and whatever it was we adopted it. So everybody else had the same thing. My recollection of those tests with UC London, is that we were getting a grand total of two or three packets per second out of the TCP.

Spilling: Oh that could be.

Cerf: It was a fairly placid kind of exchange.

Spilling: Oh we had to remember the line --- well, maybe they had separate satellite connection at that time between UCL or Goonhilly and Etam. But I can't remember the capacity of that leased line.

Cerf: It was 9.6--

Spilling: Nine point six.

Cerf: So that line drawn coming from Kjeller down to UC London over the satellite.

Spilling: That was a little later.

Cerf: Sorry?

Spilling: That was a little later. They have at least line from UCL to Goonhilly to Etam and then into the Arpanet.

Cerf: That was a 48 kilobit line. There were two lines. There was a circuit that went up to Norway at 9.6 and there was a 48 kilobit line that went to Goonhilly Downs.

Spilling: Sure.

Cerf: And I remember at the time that that was considered blazingly fast for British Telecom to supply. But before we had the satellite connection and we were just doing the TCP testing in 1975, we didn't have the-- the Satnet was not part of that test.

Spilling: No.

Cerf: It was going through a satellite link that was internal to the Arpanet.

Spilling: To the Arpanet, yeah.

Cerf: Down to Tanum and into wherever it landed. To Kjeller and then the land line at 9.6 down to UCL. So that data rate was fairly modest during that time frame. I have a question about the three node, three network paths Ginny.

Travers: Yeah.

Cerf: Even though you had installed gateways, I think at all three of the satellite downlink sites, at Etam, at Goonhilly Downs and at Kjeller. My recollection of the three network test is that that the only gateways what were actually used for that test, were the ones at Etam and the ones, or connecting the Etam, and the ones at Goonhilly. I don't think any traffic existed into Norway. Do you remember-- since the packets never actually left the--

Travers: I think the point was like to show it going all the way through everything and sort of coming back. So it was like some kind of a round trip and like you said last night, we may have engineered the packet algorithms and routing algorithms in the gateways in order to achieve it.

Cerf: I'm pretty sure that it was engineered to run through only two Gateways, because if it had to go through Gateways then they would have had another satellite up-and-down hop. And I don't think we did that.

Travers: I mean, the other thing is that this experiment was one of a whole slew of experiments that we were doing. Somebody asked last night like, well did you keep the whole thing up? Well we kept it up as much as we could. I mean the Packet Radio Network was different in the sense that that was like only up and running when they were out running around with the van and they had all the radios powered up. But Arpa Atlantic Satellite Network that was supposed to be providing real connectivity for people to be working on it.

Cerf: Right.

Travers: So we tried to keep that part up as much as possible.

Barbara Denny: Yes so, I guess I'll start with a little background about myself. I was like Ginny. I just had a bachelors degree and BBN was my first full time job. I went to Carnegie Mellon but I had exposure to the Arpanet at that point in time and it actually was instrumental in some ways, me going to BBN. I had an account on cmu-a and I wanted to keep it. So that was one of my criteria for looking for a job. Who

am I talking to that has the network? And of course I talked to BBN and they came on campus, and I went up there and I started working on Packet Radio in the fall of '81. But at that point, I had to use the TIP to still go back to CMU because the Packet Radio Net and our test bed really wasn't connected all the time. And so if I wanted to send mail or email my friends, I would have to go back to the TIP then access cmu-a and send it as BD0N, that was my user name. So that was a lot of fun--

Cerf: What was your responsibility at BBN and the Packet Radio program?

Denny: Okay, yeah. Well I worked for Jil Westcott, and Mike Beeler and Charlie Lynn were part of that group at that point; and I was responsible for making the Packet Radio protocols work at that point in time. One of the first papers I got was Ginny's "How to Build a Gateway"¹ and it [the gateway] was in there but it never caused me any problems. <laughter> So I was like okay. And they were just making the transition from the single station control [controller] of the Packet Radio Net to the multi-station. So I got Greg Lauer's code and just running experiments and dealing with SRI and supporting it. And at that point they had started, I believe, the test beds in Fort Bragg. And so one of my jobs was to make sure we kept the testbed users happy, because we really, as you know, wanted to get the military excited about the technology.

Cerf: Exactly.

Denny: And so actually a lot of times I would call SRI because in those days, BBN only had a small set of radios. And if I wanted to do any more complicated experiments, SRI had more extensive test bed. And I'd call them and say, please configure the net this way. I want to run an experiment. We had the crossnet debugger, and that would allow me to work at BBN and get fun stuff and dump the information I would need from all the different tables and try to figure out what was going wrong, what was going right, and fix things when I found them.

Cerf: You know, I haven't heard the term crossnet debugger in a long time.

Denny: Yeah, yeah it was really XNET—

Cerf: That was a very important tool.

Denny: Yes it was.

Cerf: It let you look inside what was 3,000 miles away to find out what it was doing.

Denny: Yeah. So that's all that I did on my first project at ... when Packet Radio was still in the first infancy and then the subsequent work I went to SRI and continued on with Packet Radio and by that point, the routers really were separating out and I worked on the LSI-11 router for the [re]constitution protocols.

¹ Note of the editor: This paper was written by Lixia Zhang.

Cerf: Did we actually, I don't recall that we ever called them routers. I think we still called them gateways.

Denny: No, they were always gateways, right.

Cerf: It didn't turn into a router until Cisco Systems started calling them that.

Denny: Yeah, my recollection is, yeah, it was originally the gateway and then what happened ... it might have been Cisco ... I also heard, like, application [level] things were called transformation gateways. And we're going, "wait here's a gateway---which gateway are we really talking about?" And so at that point there was a transition from being known as a gateway to a router. But to be honest, I don't remember exactly when that flop occurred. But it definitely was a change in terminology.

Weber: And when did you go to SRI?

Denny: I started at SRI in '83.

Weber: Eighty-three.

Denny: At that point I was working on a router in Ada²

Cerf: Good lord.

Denny: Yeah. Ada was popular and actually that's what was my program when I was at CMU. I was working with the grad student department developing the software development environment for Ada. So I already had the Ada skills. So I got asked to try to develop one of the first routers in Ada. So we spent some time doing that. And the technical problem they were looking at was how to do emission control. Unfortunately that project was ... due to a funding cut, we never got to finish the implementation of the planned technology to try to handle emission control, which is when you cannot ... you can only receive, you cannot transmit. And so I had to make that all work with the TCP/IP protocol suite.

Cerf: Yeah this was a big issue because the navy was particularly concerned about ships at sea, submarines especially, not being exposed by radiating, because they might be detected by satellites. And so when you are in EMCON [emission control], that meant you didn't radiate at all. But that made it really hard to make TCP work because there was this thing called ACKing and it can't acknowledge the receipt of the traffic. So you have to assume that it got it or you have to transmit multiple times. I often wondered whether we got anything to work in that domain, because if we did, it was after I left ARPA.

Denny: Oh. Yeah no. Unfortunately, like I said, our funding got cut because there was a transition in the Navy department and they decided they didn't want any [funding] of the Gateway anymore for that problem. So we never really got to finish that technology underneath it. But a subsequent fun one, which unfortunately was not till the mid '80s, if you'd like to know what that was --- the reconstitution protocols. And that's when we tried to do a mobile IP scenario and used the aircraft to connect.

² Note of the editor: Ada was a new programming language.

Cerf: That was like 1982 I remember.

Denny: Yeah, it started and experiments went all the way through until like '86. I was just talking to Ginny. Things stick out in your mind. For me it was the fact that I was in Omaha for an experiment and then the disaster with the shuttle happened. And I was like oh my God as I was sitting at the terminal trying to get things to work at Omaha.

Cerf: Wow. That would have been in '86, late '86.

Denny: Yeah and the Looking Glass [Airborne Military Command Post] was still up, because I remember once we accidentally hailed them trying to get our plane. Luckily we had a military rep there who was quick to go over and explain what was going on. And sorry, we didn't mean to hail the Looking Glass when we were trying to a dry run.

Cerf: Several of you were looking at reports from SRI.

Denny: Yeah.

Cerf: Earlier. Do you recall whether any of those included the reconstitution experiments at all? Because I've been trying to nail down the date that that first--

Denny: Experiment?

Cerf: Demonstration of reconstitution was done and I keep thinking it was in '82, certainly before I left ARPA and which would have been around October. But I don't have a very clear recollection of that.

Denny: Yeah ... Don [Nielson] could be able to help you along with Jim Mathis. I basically took over the reconstitution protocol from Mark Lewis who had done the first set of demos and I was involved more in the later set at demos. So I could check. I might have documentation at home. [I'll] look in the garage³.

Cerf: Now there're two that are--

Spilling: I'd be interested.

Cerf: There are two other people at BBN that-- whose names haven't come up. Radia Perlman is one and Julie Sussman was the other.

Denny: Yeah.

³ Note added by the editor while reviewing the transcript with Cerf: Michael Frankel managed the reconstitution experiment for SRI.

Note added by the editor while reviewing the transcript with Denny: When I joined the project, Boyd Fair was the person doing the management.

Cerf: One of the two had some involvement in figuring out what the re-constitution rallying protocols had to look like. Was that Radia by any chance?

Denny: Boy.

Travers: It could have been. It's really hard to say.

Denny: Yeah because--

Travers: We have different people working on different things at different times.

Denny: Different times, yeah.

Travers: Julie worked on the earlier Packet Radio Network stuff and then Radia came in a few years later. And I'm not sure which people, you know, worked on which things. I mean there was a group of us and a lot of us shared doing different things.

Denny: I'm not aware of the early BBN effort in that respect. I only know the SRI contract, because I was at SRI at that point. And at that point, it wasn't like a joint contract like a lot of packet radio is. We were really pretty much working [independently.] I think part of the funding was DARPA and part was also out of Rome Labs. And as far as what Radia did, I don't know. She had left BBN just prior to my joining. I did meet her at a party. But I didn't work with her, and Julie we overlapped at BBN and I did talk to her. But again, her involvement in packet radio at that point had wound down a little bit too because I never dealt with her directly when I was doing the code.

Cerf: I doubt that it would be Julie who did that work because the recon protocols were up at the IP layer.

Denny: Right. That's correct.

Cerf: Whereas, all the other stuff would have been more at the packet radio level.

Denny: Yeah so ... so as far as I know, a lot of the key contributions were made by Jim Mathis and Zaw-Sing Su doing the design of how, you know ... realizing you need to split the TCP binding problem between the address and the transport; and how that was causing us grief.

Cerf: Still a problem today.

Denny: Yeah. And in fact, one of the hardest problems for me was trying to figure that out was due to the IMP. The IMP buffers weren't that long and I was trying to make the radio realize that "oh, you really changed networks because it broke." Well I was having trouble where it was the IMP that was backing up; and I couldn't get it to recognize correctly that it had been moved, versus just congestion problems at the net. So here we were having this struggle late one night trying to get ready for a demo; and not have... you know, having trouble figuring out exactly, you know, what's the right number of things to

realize you've dropped before you really have moved. And it was touch and go there for a while. But we made it, which is good⁴

Travers: I mean I recall in the early days, we started to split apart the meetings. And so you would have several days with their TCP meeting and then there would be a day that was an IP meeting and then for a while we were into this awful thing where we threw a Gateway meeting in on a Saturday or something like that.

<laughter>

But that's sort of how I recall things started to split out and so you'd have some people covering one set of meetings and other people coming in for other sets and we started to split out the functionality.

Weber: And this is within BBN then?

Travers: Oh, within the whole group. There tended to be people then who started to specialize more in the TCP issues and- and I forget, there were a few jokes going around or maybe Vint remembers some of them about how could you tell when you are in a TCP meeting versus how could you tell you were at a IP meeting. As I recall the reason you knew you were at an IP meeting if we were arguing about the length of addresses. And we argued about that for years.

Cerf: Well it was basically one year. My recollection is all of 1976 was spent arguing over how big should the address space be for IP. And in 1977 I said finally, okay, we're done when nobody can come up with any definite conclusion. We're picking 32 bits and we're going to get on with it.

Travers: Yeah I mean and the thing people probably don't understand in retrospect, is why we were having these arguments and what the parameters were. And so the question was, how many computers are you going to have to address. People don't realize this was before the personal computer. You know, so like BBN was probably you know, like SRI or some the other organizations. I mean we had a huge number of computers which was you know, in the tens maybe you know. So and then when you started talking about network addressing, I distinctly remember sitting there saying, well it took us six months to get the last two connected. So we're going to need two to the 16th here? I don't think in my lifetime. And the reason we were arguing about it, the other side of it was, we had these like 9.6 kilobit lines that we were using to interconnect things. And then you start running telnet protocols over. And people are saying wait a minute, we're trying to send one character and we've got how many bites of addressing with it? And you know, I mean now that's not a concern at all. But back then, it was a major concern, because you were really slowing things down by throwing all this addressing information on top of it. That seems kind of- of a strange argument these days.

Weber: What was the atmosphere like at BBN in those days?

⁴ Note added by Cerf: The idea of this effort was to partition the ARPANET and use the packet radio network to glue it back together using the IP routing system. To make this work, you had to recognize that a core network of the Internet had partitioned and one needed to create two new subnetworks with associated and distinct Internet network IDs and route OUT of one part of the ARPANET and back into the other by going through the third, airborne and ground packet radio network.

Travers: <laughs> In what way?

Cerf: Depressed.

Weber: Was it exciting? Was it--

Travers: BBN's always been pretty easygoing you know and it was a fairly small group of people and I think it was exciting across the board in the sense that you know, we were doing something new that the establishment basically said couldn't be done. I mean I remember sitting in, in meetings with representatives from the PTTs and saying you know, this packet stuff is never going to fly. We know how to do communications. I distinctly remember one meeting in which they said, you can't call what you're doing a gateway because we have defined what a gateway is. And we know and you don't. And you know, we were like, so what? I mean I think there was sort of that, that sort of atmosphere all the way from the program was up to people like Vint and Bob who was saying you know, "No you can't tell us what we can't do. You know, this was really-- it was a technology experiment and it's like we're going to see what we can do with the technology. And if it's outside the way people do things now, if it's outside the standards, we don't care. We're going to see what we can do with the technology." And that's really what pushed the whole thing forward.

Weber: Paal did you have any dealings with or was the Norwegian PTT involved at all? Not really?

Spilling: No they-- my boss Ing Rolune, he was able to persuade them to join but not as in-- he was not able to force them to be an active member. They volunteered to be an observer but they provided, free of charge, the uplink to the satellite. They provide two lines from Kjeller to the attendant ground station in Sweden, one 48 kilobit line and another 9.6. And then I think London came up with, you can say, a 9.6 line from UCL to Kjeller. But we weren't-- and I think that Peter Kirstein had the same problem in England to really get the PTT involved. They at least in Norway they were just willing to observe to see all the stupid things we were doing. But not participate, just comment and observe. And I think when the-- at least my impression was that, say, the standardization bodies and the ARPA community, they were just like dogs and cats. They didn't like each other. They didn't understand anything of-- at least the standardization world didn't accept the way ARPA, say, the communication paradigm. And so they didn't believe in this and that was my problem, when I quit NDRE after having spent--

Weber: When did you quit?

Spilling: In August '82 after spending one and a half years with SRI, coming back-- then I attempted to stay longer. But the big boss at NDRE said, well you have-- now at the NDRE, you have to come back. And when I come back, I was just sitting in my corner. Nobody would be very interested in what I had been doing. And then after a few months or a few, a year or so, then I went across the fence to the telecom administration, the research department of the NTA. And they had more money but less interest. So I got money to establish my own ... after a short while ... establish my own Internet, connecting the university in Oslo, Bergen and Trondheim with our facilities at Kjeller and being able to communicate into the- the U.S. part of the Arpanet. That--

Cerf: And the money came from the Norwegian Telecommunication Administration. They did give that.

Spilling: Yeah, because they had money enough. So if I was wasting some money on this, it was okay. And- and when it was not possible to get anyone, except one person, out of 150, 200 employees to participate in the TCP stuff. And that guy he implemented gateways for the innovation part of the Internet, making use of the bridge telecom, that bridge communications device. But that was all and after a while it all died out. They would all use the Internet facility, but not participate because they didn't believe in it.

Weber: So in '77 at NDRE, they were going along with it but.

Spilling: Yeah.

Weber: And what was your official role at that time?

Spilling: Well it was to partake in the- the Satnet development. That was I mean, the main, at least one major goal and the other one was to participate in packet speech experiments.

Weber: But what was your position within the NDRE? What department were you in? What was your--

Spilling: Department of Electronics.

Weber: And your.?

Spilling: And just a research fellow.

Weber: But obviously the institution did participate up to the high level but only as a kind of--

Spilling: Well I think it was- it was permitted ... I mean the group of Yngvar Lundh was permitted to participate. There was no money and I think the NDRE institution itself was not interested. And I think in the long run or at least the impression I had was that the technology was at least ten years ahead of say the acceptance. I think that was the same in England, maybe all over the place. But at least in Norway and in Europe, for sure it was at least ten years ahead. So when the TCP/IP stuff was sort of standardized in '82, it took another ten years before and I think that prompted you to the Web browsers before it was accepted as the communications media.

Cerf: I can confirm that because it was literally just after the Internet rollout in January of '83, the U.S. government and DoD being included, decided that the OSI protocols would be the preferred target. They had five years of documentation from 1978 to '83. It was in-- the International Standards Organization had standing as an international standards group. Now and this other stuff, this TCP/IP stuff was just a DARPA experiment. So for the ensuing, almost exactly 10 years, the U.S. government adopted something called the government OSI profile or GOSIP which was published by NIST, the National Institute for Standards and Technology. And finally when the Internet Society got started in January of '92 when I served as its first president, not long thereafter I sent a letter to the head of NIST, I forget who that was at the time, suggesting that it was time for them to carry out an evaluation of TCP/IP versus OSI, TP in particular, and to determine whether or not the two were sufficiently equivalent in functionality that it would be acceptable to adopt the use of TCP/IP as an alternative. And they appointed a blue ribbon

committee. And after a year of discussion, they came back saying, yeah you know, the two are pretty similar. The only difference is we had by that time, how many years of experience and implementations of TCP and zero years of experience with the OSI. So Paal is exactly right, that it took ten years of fighting to finally overcome that particular challenge.

Spilling: And I think it's sort of a paradox, the way ITU and the standard bodies are working. They work in periods of four years. And at least the communications activity in ISO was very academic. I believe that. And so then we're working say 4 years to make a standard or standards and there were like IBM and certainly many other manufacturers participating. And they all got, depending on how strong they were, got their options into the specifications. But then ... and then it was standardized and then it should be implemented. And then they had two devise workshops to figure out which of the options should we make use of. So they were spending one or two or maybe three years on selecting the options. And what it turned out was that the options that I selected, made the ISO the TP for the transport protocol. And they were part equivalent to the TCP. Just an enormous waste of time.

Weber: Well just tell a little bit about writing the gateways and traveling, seeing the different sites you were involved in it.

Travers: Well you know, that was really the interesting part of it and I guess that's why, that's what all my memories are around. I mean it was a very different world in terms of what computing facilities we had and such and we were using paper tapes and we were writing an assembly language and things like that. You know, we'd get the PDP-11 running in the lab and you know, in many ways you know, people talk about how much fun it was. And in many ways, it's a lot like what we do now. I mean you got a big demonstration coming for DARPA and we had meetings typically every three months and they went around from you know, it was BBN, Washington, ISI, SRI and then thanks to Vint and Bob, we slotted in the European trips so that we always had like you know, right before Christmas in London or it was a summer in Norway. They were pretty good in getting those right. But you know, it was a lot of working late at night you know. I can still remember Vint hanging over my shoulder, "Are we gonna get this running? Are we gonna get this running?" I mean it was like totally not obvious and especially network experiments there are so many different pieces involved and all those people have to be up at the same time in order for anything to go through. And then you say, "oh finally we got it, you know. That packet went all the way through." And it's like yaa you have any proof? You know, you have any snazzy user interface? The answer's like no.

<Laughter>

I mean network experiments were just like the most like boring, uninteresting things, you know. But we're all excited because we got all the pieces working. And you know, I got to meet people from different parts of the world because we were together every three months at meetings and it was a real young crowd. You know, there were a lot of people who were working fresh out of school, people who you know, had just come to DARPA, people who had just come into some other research institutes. So we had a lot of fun together. It was great to be able to go and see people in their home countries, to spend some time with Paal, to visit in London. So that, you know, that was all really exciting. And those are, I think, the fond memories we have is much more so than the packet running through the network. It's the people we got to know.

Denny: Yeah.

Weber: Had you been to Europe before that?

Travers: No. No so you know, this was really exciting--

Weber: So what was your impression arriving in Oslo?

Travers: <Laughs> It was a great time. And I spent time traveling the country. I was very much into the outdoors. I actually made arrangements to go hiking with the sort of the Norwegian National Hiking Association. And you know, back in those days, it wasn't like you jumped on the web and you found out all the information. I mean this was like multiple phone calls back and forth and you know, making the transactions to set this up and then I'm going to meet the group in a Norwegian train station and go hiking for a week. It was like pretty out there. And you know, in between luckily I had opportunity to stay with Paal which helped, because I mean, I didn't know the language or anything. But Norway is just a beautiful country and I just regret that I haven't gotten back there yet. But it's on my list.

Spilling: We enjoyed having you there.

Weber: So she was this young, almost student age fellow.

Travers: Yeah, yeah and Paal was living I guess out in what would be the suburbs of Oslo but more like out in the country. I remember going to the farm next door and Paal and his wife showed me around and showed me all the barn animals and stuff. It was just like, it was really a lot of fun. I mean, people socialize together a lot more and you know, really got to know each other as friends as well as co-workers and that was great.

Weber: You were some of the first people to really use email. You were among the groups that were the first to use it as a real community tool. I mean, do you feel that-- was that different than today, the ways that you all interacted electronically? Tell us about that.

Travers: I think the thing that was different was, you know, the state of communications back then. You know, that was back when you called home and reversed the charges and somebody kept their eye on the cost because it was dollars per minute for phone calls. International phone calls were even worse, you know. What do you mean you're having problems with the Gateway so you were on the phone to Oslo for an hour? You know, you just, you didn't do that. So yeah, we were some of the first people who benefited from that. I mean I can remember if you had the network connections up in that and by that time the Arpanet for example, was pretty solid. So you had good network connections to California. It was like way easier to type at people and it was just the TELNET protocols underneath, you know. But that supported being able to type messages at each other. Well that was years and years and years before IM but we were using that as communications because the alternative was picking up the phone and those phone calls were expensive.

Weber: Because--

Cerf: I'm sorry, there was the time zone issue too, trying to schedule time zones and maintain those conference calls was hard, you know, with anywhere from three to five hours of time zone difference or

more. Eight hours from the West Coast in the U.S. to parts of Europe. So email had overcome a lot of those barriers just by allowing you to communicate whenever it was convenient for you.

Weber: And I realized that this whole group that had never really come together as a group and we've got what, 20 out of 35, you all, it was by benefit of electronic communications that you were a group. I mean at the time you wouldn't have known each other otherwise. I mean how much do you, when you met, when you went around and met each other, did you feel that you already had a relationship based on email and messaging?

Travers: But we, we were also in meetings like every three months.

Cerf: Yeah, I think you're misunderstanding something here. We had face to face meetings on a very frequent basis-

Spilling: Sub groups.

Weber: Okay.

Cerf: Every three months and so-

Weber: But not everybody was it?

Cerf: Well there were at least 70 to 100 people who would come to those meetings. So that was a fairly major event.

Weber: Okay.

Spilling: And I would like to comment on that. In, at least my experience, in using email in the old days is that if you haven't a face to face relationship with the people you're corresponding to, an email is not, it's not binding. I mean in order to establish a good working relationship, you need person to person communications. That has to be established prior to using email. And back in those days I had two facilities with the services I was using. One was the talk, which enabled me to communicate with my family home every time I was out to these meetings. And another function which I found nice once in a while was the, WHOIS. You could ask the computer who is on, logged on to the computer? And then I could disturb Vint for example or some of the other guys. So that was my main use in those days.

Weber: Do you have any questions.

Cerf: I don't think I have any other specific question to ask. We covered quite a bit of territory in this discussion. I think it's important to know that the tests didn't end with the 1977 three network demonstration. That was just the first shot, in some sense, of getting multiple nets demonstrated. In fact, looking at some of the reports that Don Nielson dragged out, looks like there were five or six network tests eventually, which were done including packet radios. So of course and then finally the Internet just took off and there were lots and lots of networks. But getting this to work the first time for three was, and a significant milestone I might add. And these three folks were part of all that story.

Weber: Well thank you.

Spilling: Okay.

Travers: Thank you.

Denny: Thanks. <laughs>

END OF THE INTERVIEW