An Interview with

DAVID L. MILLS

OH 403

Conducted by Andrew L. Russell

on

26 February 2004

Newark, Delaware

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

David L. Mills Interview

26 February 2004

Oral History 403

Abstract

Internet pioneer David L. Mills discusses his career working with computers at the intersections of academia, government, and private industry. Mills earned his Ph.D. in Computer and Communication Sciences at the University of Michigan in 1971, and then worked at the University of Maryland (1972-77), COMSAT Corporation (1977-82), and Linkabit Corporation (1982-1986) before joining the faculty of the University of Delaware in 1986. Dr. Mills invented the Network Time Protocol, chaired the Internet Architecture Task Force, and made many other significant contributions to the development of packet-switched networks and the Internet. Major topics and themes of the interview include: the evolution of the Internet standards process, the social and cultural aspects of the ARPANET and Internet research communities, international collaborations and tensions within computer networking communities in the 1970s and 1980s, the Government Open Systems Interconnection Profile (GOSIP), and his interactions with colleagues including Vinton Cerf, David D. Clark, Jon Postel, Peter Kirstein, and David Farber.

Russell: This is Andy Russell, and today is February 26, 2004. I am at the University of Delaware with Dr. David Mills. Dr. Mills, can I have your permission to record and use this interview for my graduate seminar in oral history?

Mills: Absolutely.

Russell: If I want to use the contents of this interview for other purposes at a later time I will request your explicit permission.[1]

Mills: I am legally a public person, an employee of the state. The Freedom of Information Act applies.

Russell: Good. I have a list of questions here and I am sure other things will come up as we get through it. My first question is, did computer science exist as such when you entered the field, when you started working with computers?

Mills: No. I started, I think what happened in 1960 I started grad school. And after getting a master's in EE [Electrical Engineering], I found computers much more interesting. And so I started over again, which eventually became computer science. When I received the Ph.D. in 1971, the field still hadn't really existed. They had some programs, I had the second dissertation at Michigan in what was then called computer science, so it was quite new.

Russell: It seems that your first experiences with computers were with "Big Iron" mainframes.

[1] Written authorization was given for public posting by the Charles Babbage Institute.

Mills: Oh yes.

Russell: I'm curious, can you explain some of the differences between IBM-style mainframes and the ARPANET/ Internet-style network architecture?

Mills: There was a fairly interesting cultural state of mind in the middle 1960s. What started it was the IBM 360. These were large computers. You devoted a whole room to them, you had huge air conditioners. The computers by today's standards could be implemented on microchips. But back then it was "big iron." The philosophy was, the computer cost so much that you had to have one of them for the university, and you couldn't afford to decentralize, because that would strain dollars from the central facility. So, what happened with DARPA [Defense Advanced Research Projects Agency] was that we began to believe that it wasn't just big iron time-shared computers, it was decentralization. It started out with graphics displays. Obviously, centralized graphics displays were silly. So they moved them out into smart, relatively small computers with graphics displays on them. These were small computers. All of the sudden we had the idea of the minicomputer, exemplified by the PDP-8. And so the development of operating systems for them. And the whole idea that computing could be done in a decentralized way, and it was cost-efficient.

Russell: This was the mid-1960s.

Mills: Yes.

Russell: So then the ARPANET and networks, computer networks as we know them today, came from the idea that computers were distributed.

Mills: It started out with a service-center mentality, which is to say ARPA, it was called that at the time,[2] had centralized service machines on the West coast, and in New England. The early ARPANET, the earliest ARPANET, existed primarily to get access to those machines. It was only as it quickly developed, and especially with electronic mail, that the idea that these machines would be used to interchange ideas and develop programs and things like that, and this was now in the middle 70s, I think is the right way to do it, and the attachment to the ARPANET was called an IMP, an Interface Message Processor. Some of these IMPs had modem lines attached to them, and they were called TIPs, for Terminal Interface Processor. And the idea was that [unintelligible] the terminal authenticate yourself to that TIP, using passwords, and then you could access any computer on the ARPANET. But this idea shifted pretty quickly in a variety of connecting minicomputers. The ARPANET was a connection where you had your mailboxes on the local machine, and the network would be used to exchange files, data, mail between these computers. But the service center idea continued to flourish at this time. It was only until about the middle 1980s that the service aspect of ARPANET slowly disappeared along with the ARPANET itself.

Russell: Your biography explains some of your work at COMSAT [Communications Satellite Corporation] between 1977 and 1982, including work on IP [Internet Protocol], TCP [Transmission Control Protocol], Telnet, FTP [File Transfer Protocol], and so forth. Can you

[2] In 1958 President Dwight D. Eisenhower created the Advanced Research Projects Agency, ARPA. This agency was renamed "Defense Advanced Research Projects Agency," or DARPA, in 1972, back to "ARPA" in 1993, and then back to DARPA in 1996.

describe the work environment at COMSAT and your interactions with other DARPA contractors?

Mills: Well, COMSAT was then the only interface between domestic telephone service and international satellite. The international portion was from INTELSAT [International Telecommunications Satellite Organization], but Congress decreed COMSAT was the only interface between the two organizations. And, thus, COMSAT's business was telecom. So it was also in the business of developing satellite modems and earth stations. Much of that stuff was being paid for by DARPA, and then the IP project came along. Part of the IP project started at what was called the Atlantic SATNET, which was a network of several countries, five in all, of earth stations operated mostly for the domestic telephone organizations and COMSAT. And so COMSAT had an [unintelligible] taurus antenna, which was *mine*, and access lines, and that's where I got my start. But the interface, you see the difference between the research community and the service community. The research community was kind of orthogonal to COMSAT's laboratory objectives, and there was always a bit of friction between what we're doing here and whose money was funding. I managed to be on the side of that equation for the research, and so it didn't bother me that much. But there was some friction.

Russell: And I assume that Bell, or AT&T, was the service community, the telephone provider.

Mills: That's true, but MCI had started. MCI had started as a building full of lawyers with an antenna on the roof. And I say that seriously. I think it's fair to say back then that there was more a war of turf than you see today. Turf-building was MCI's—and others—were building

turf. AT&T was the dominant supplier. What this meant was, that when you got a public tariff earth station with the access by [unintelligible] telecom provider, then they had to be completely separate. So at the COMSAT earth stations, there were equipment rooms at the station which contained telco equipment, and even the station engineer didn't have keys to those rooms, it was that separate. I had a little theory about this. That when you interconnect two carriers, who charge tariffs, that the tariff they charge for the interconnect between the two is driven to the minimum to the point where the only connection is a screw-type terminal strip, and the value of that strip is zero. You arrange the whole system so that ends up to be the case.

Russell: During this period you developed some relationships with members of the DARPA research community.

Mills: Yes.

Russell: Can you tell me how you all worked together?

Mills: It actually was a really interesting and useful collaboration. But you have to understand the culture. This was the culture of grad students, with this fantastic toy. We had a 15 million dollar toy—it was 15 million dollars a year. And it was a sandbox. And we essentially were not told what to do. We just were told, "Do good deeds." But the good deeds were things like develop electronic mail, and protocols. I guess the really important comment is a cultural point. How do we develop habits of working with each other? Since we were very seldom in the same place at the same time. We weren't physically in the same room very often. So we worked

together with electronic mail, and moving media back and forth. We developed teleconferencing, which we did a lot. So I think it might be kind of a cultural experiment, a sandbox where you can see how to develop projects like this, and the fact that we had these young grad students, working with a diffusive kind of objectives, we kind of had to work our own objectives. We still had to interface with the military, because after all DARPA is a military funded organization, so we had to be relevant to that community. That was interesting because we did various kinds of demonstrations, I call them packet-poppers. Somehow I got to be kind of the focus of those things and the experiments I did. A lot of the things I did was making networks work, so one of my favorite packet poppers was some Admiral at sea, in Monterrey Bay, wanted to talk to his counterpart at USECOM [United States European Command] in Stuttgart, Germany. And this involved a satellite hop, involved a packet radio hop, it traversed several networks—and he wanted to talk in real time. Just actually being able to do that was a good demonstration for the military. Another was packet radio. To get the whole idea of packet radio. And one packet-popper to do that involved me sitting in the back of a Land Rover bouncing around the hills of Malvern, England, reading my mail in the U.S. And dodging sheep on the back roads. That was fun.

Russell: So mostly, you said over electronic mail, and rarely did you meet face to face. On what occasions did you meet in person?

Mills: We probably met 3 or 4 times a year, but, because it was very much an international organization, we met in Germany, in Norway, in England, and in Italy.

Russell: And this was on an ARPA travel budget?

Mills: Yes. For a while, I was kind of commuting across the Atlantic. It got to the point where I got to know the cabin crew of Pan Am 107.

Russell: So there must have been a lot of interchange of ideas across the Atlantic—this wasn't strictly an American…

Mills: No, no, not… There were a lot of very valuable contributions. The equivalent of NASA in Germany, which is called DFVLR [German Test and Research Institute for Aviation and Space Flight], they imported or had developed a good deal of satellite technology. COMSAT labs had also, the Norwegians had developed a good deal of technology, so it was very much a collaborative effort.

Russell: Did everyone speak English?

Mills: Oh yeah.

Russell: And did you speak any German?

Mills: I spoke a bit of German.

Russell: From DARPA's perspective, do you think that researchers in the private sector or

government were treated differently than researchers at universities?

Mills: That's an interesting question, did DARPA treat them separately? Well they treated them more or less the same. However, the commercial contributors had agendas—serious agendas. I mean, why get involved with this unless you can make money? The fact is that there were two ways of doing this. One is to think of DARPA as subsidizing corporate research. The second thing is think of corporate research being able to get ideas from DARPA, and then go off privately and make money. Now it turns out there are two kinds of contracts. One contract is a cost plus fixed fee, which is what I worked on. That's research. The end is research products. The other—if you participate in that kind of contract in some government agencies, you cannot participate in production. So there was an incentive to do the research for the research's sake, but there was an incentive not to do the research if you want to produce something for pay. Since I worked both for universities and for corporate, I saw those motivations, loud and clear.

Russell: It seems from your setting that you preferred the research setting to the corporate setting.

Mills: I had a case at COMSAT where I told my boss, you know. My boss had said that "there's not enough money for us to complete the work, and we're going to have to use the money that we would give to you to somebody else." So I told Vint Cerf from DARPA that, and Vint was quite unhappy. So he went to my boss, which he later reported to me, and he said to my boss, "What is it that you're not going to do now in order to leave enough money for Mills to do his work?" Well, there are two questions. It made me happy, and it made my boss very angry. A

year later when I left the company he took me aside and he said, “Mills, I should have fired you a year ago.”

Russell: Your mention of Vint Cerf brings me right to my next question. I’m interested in learning more about leadership within the DARPA community—especially the characteristics of the leaders. Not just Cerf, but other people, how did leaders emerge within the community? Were they the people who controlled the money, or the most technically astute people, and you mentioned—it’s widely reported that there were lots of different colorful personalities within the group. I’m trying to sort out who rose to the top, who had control.

Mills: Well, the short answer to that is to say that agendas like the Internet are started by someone. They don’t happen by themselves, somebody gets some kind of vision, and I give that vision to Vint Cerf and to Bob Kahn. But a lot of that—you see the tip of the iceberg, but you don’t see the geese paddling under the water. Vint and Bob had to convince the director of the agency, the director of DARPA, to spend 15 million dollars a year on this hokey thing. I think the leadership qualities shows in the fact that they were able to get the money. I think technically, that those two were very highly regarded technically. You know how the DARPA program managers work—they’re hired for two year periods, and then they’re expected to leave. Okay, the problem is, how do you have an agenda that lasts for 25 years? In a climate where the principal investigator leaves after two years, it just doesn’t happen. Now Bob was there for some years, he finally quit to start the Corporation for National Research, CNRI, and Vint stayed there quite a bit longer, but that program wouldn’t exist, it would not have flourished, if those guys couldn’t be around for a while.

Russell: And someone like Jon Postel was known for being very good technically and working very well with people. He seems to have been respected by everyone.

Mills: I can't—if somebody says, "What is the position description for Jon Postel?" you could not write it. It's very hard to describe what he did. I knew the guy well. I always regarded him as the glue. That's just it—he was the Internet glue. He was the RFC Editor, he had his own policies for doing such things as assigning numbers. Somebody has to assign numbers. So he was the numbers czar. That's how we'd proceed. So, if you got a protocol number, Jon assigned it. If you got a network number, Jon assigned it. And so everything—things like that—everything—documentation series. Now there's an issue there. Most projects are documented by some kind of ongoing serial publication. Even before Jon was involved, they started the Internet—well it wasn't even the Internet—it was the ARPANET Request for Comments [RFCs]. And the whole idea of the request for comments, as a public document where an unfinished idea was exposed, sometimes a stupid idea was exposed, sometimes *purposely* a stupid idea was exposed to get something going, I did a few myself. But the idea of documenting in a widely-circulated serial publication of ideas was something that hadn't been done before. We did have the Internet Engineering Notes series, that were supposed to be engineering notes. But the RFCs are slightly different, they were works in progress. Later on the RFCs became standards vehicles, and not so much works in progress, and what developed was the Internet Engineering Task Force's Internet Drafts. And they are definitely works in progress.

Russell: More on the personalities of this community—did personal quirks or work habits or personal conflicts influence the technical work of the group? A lot of people were trying to work together, and I wonder to what extent social dynamics influenced what you were making.

Mills: You know, for a 25 year experience, I can't recall many cases where a personality got in the way. There were characters. Jon was a character. He was a very private guy. And so we just kind of—I just kind of steered clear of him, you know I just agreed with what he did, and I socialized with him, but he wasn't someone I would schmooze with. Dave Clark—a good example. He was an extremely bright technical person. He was assigned kind of informally I guess the Internet Architect. That was his title. And Jon was RFC Editor. Who else was involved? Well, there were several people who came and went. I'm trying to think of the names that persisted for a while. One character was Peter Kirstein. Peter was a professor at University College London. I value Peter because he was the social and cultural interface between the British government—and all of its quirks—and the rest of us. Peter always managed to get his way. He was politically very astute. I think Paul Spilling from NTRE, which is the Norwegian equivalent of Bellcore. Back in the days when SATNET was established, the technical aspect was interesting and hard to know. But the real aspect was to get the governments to agree to share their domestic earth stations. I think Bob Kahn was the one to blame for that one. I saw the documentation exchanged between the national administrations and Bob Kahn, COMSAT had it. And I looked it over and figured, wow, this was hard to do! He had to essentially beguile the technical administrations who were not friendly to this—the rank and file of the various countries were not friendly to the big, bad U.S. doing this DARPA stuff on their antennas—that was interesting. I think—there are 2 aspects. Vint was very easy to work with. He's a nice

fellow, and he's receptive to ideas. He's a facilitator, probably more so than Bob Kahn. Dave is a technical guy. His style at meetings, the meetings kind of crackled, because new ideas would fly back and forth, and frequently a bunch of us wanted to talk at the same time. He'd go like this [points]. And then he'd say, there's this issue, and you'd raise your hand. And he'd say, "You, go." I called it eyeball meetings. Floor control by eyeball.

Russell: It sounds like an academic seminar.

Mills: Oh yeah. Jon was a hippie. He had a long beard, wore sandals. Once we had a meeting at my employers, and everybody came into this room. My boss was a very straightforward guy, he took no informality, you had to be in coat and tie. So all the guys came in, and then Jon came slapping in with his sandals. My boss was going like this [makes a face], you know? I said, not to worry.

Russell: Were there many other hippies in the group?

Mills: I can't remember, I didn't notice. I mean, nobody noticed things like that.

Russell: I want to move forward chronologically into Internet standards institutions, mostly in the 1980s. My first question is about the ICCB, Internet Configuration Control Board.

Mills: That name is revealing. Because it says we're not controlling anything. We didn't control anything. We couldn't call it something that the government could construe as managing. We

didn't manage the Internet. We were just a committee down here that solves problems—configuration problems. We called it the ICCB, later it became the IAB. But notice that—Internet Activities Board. What does that mean? It doesn't mean anything. We proceeded—first of all, there are several things to mention. One of them is that we were building standards, and we were doing so in a private way. ANSI and ITU considered us irrelevant.[3] So the idea was that, if we really wanted to standardize something, it wasn't going to be with the international agencies, we had to do it ourselves. I think since then by the way a lot of progress has been made, to interface with ANSI and ITU. But back then we had to do it ourselves. So the whole idea is, what is a standard? Now, in conventional use, the standards are developed by a standards body which is responsive to the industry, government, and various interests, and so they make this framework in a standard. But, you can't implement that standard—a standard is a set of ideas and principles. It's not something which is interoperable, you have to decide if a subset of that standard is going to be interoperable. So you have to kind of, you know, tweak it, so that everybody agrees on the exact same bits, and then you of course farm it out to industry and somebody makes it. But the Internet didn't work that way, because the standards kept changing, so we developed something and then changed it. So standards evolved, they didn't fall into place. And what's more, the Internet standards tended to be those written for implementers. International standards were written as documents to be *obeyed*. They did very little implementation—you obey this standard, we don't care how you do it. The Internet's kind of backward, in a sense, we built something, see how well it worked, changed it if it didn't work, and the result you'd document and say hey, implement it this way. Well, this is not good practice, frankly. But that's how we did it. Take the example of TCP. TCP, RFC 793, is an implementers' document. I read the document—I helped write it—I read the document and I

[3] ANSI is the American National Standards Institute, and ITU is the International Telecommunications Union.

implemented it. Because I knew what it was telling me to do—"here's the idea behind TCP." Well, DARPA had contracted with a company, I think it was SDC—System Development Corporation, which may not exist anymore. They contracted with them and they said, make a formal standard. They hired me as a consultant, so I know what they did. And they wrote a book. It's a pretty good book. It was a formal standard of what TCP implementers should do. The book was 300 pages. And I don't think anyone in the government ever heard of that book. I know it existed because I helped write it. So far as the actual document, it's probably sitting in a vault someplace in Defense Communications Agency. So there's a little difference there. Having written standards myself, and trying to bridge the gap between the two—I wrote RFC 1305, the Network Time Protocol [NTP], and what I learned, it took me a year. It was a very difficult document to write, because I had to rationalize the code, which I and others had written, with the formal standard, and change them around so that each of them accurately reflected the other. I'm doing the same thing now with NTP version 4, and again it's a labor intensive—tremendous labor-intensive project.[4] So there are three views of standards. There's the OSI[5] view, in which you have a cookbook, a recipe book, I'm sorry. That's a good way to put it. The ISO has a recipe book, and the Internet community has a cookbook. And something in between is what I tried to do.

Russell: I noticed that IP is "protocol," TCP is "protocol," network time "protocol," you didn't call them standards, you called them protocols.

[4] David L. Mills (1992), "Network Time Protocol (Version 3): Specification, Implementation and Analysis," RFC 1305, http://tools.ietf.org/html/rfc1305; D. Mills, J. Martin, J. Burbank, and W. Kasch (2010), "Network Time Protocol Version 4: Protocol and Algorithms Specification," RFC 5905, http://tools.ietf.org/html/rfc5905.
[5] OSI stands for Open System Interconnection, a model for network interconnection created by ISO, the International Organization for Standardization.

Mills: Well, it's the Internet Protocol standard. The document might have the descriptive on it, but it actually represents a standard.

Russell: Is there a difference between a protocol and a standard?

Mills: I haven't considered one, I mean, if I write something that says, "This is the Network Time Protocol specification," which is in matter of fact my title, then I assume it's treated as a standard, when it's nominated as such. But in the current procedures, you have a long trail to go through to get it with a standards logo on it. NTP 1305, for instance, is a Draft Standard, it never reached the [Internet] Standard status. It never needed to, because nobody else bothered to implement it. But the real bottom line is that the IETF [Internet Engineering Task Force] RFC style of standard is very religiously—it's a religious requirement. And it's required to be in what I call Postel format, which is plain ASCII in a certain indentation and font, and everything including figures has to be done with ASCII text. But I refuse to do it, and I told the IETF, I'm not going to do this. And so, you'll have the document. It'll be a proposal, but it'll never be an RFC, because I'm not going to go through 150 pages of detailed diagrams and flowcharts—I'm not going to do that in ASCII text. So that's where it stands.

Russell: Going back to the ICCB, I'm wondering how you would coordinate research agendas. I assume you had many different people working on many different things.

Mills: There were 10 task forces. There were chairmen of the task forces. As chairmen, we'd need to come to the ICCB. I was first chairman of the Gateway Algorithms and Data Structures

Task Force [GADS], because I was interested in routing gateways and things like that. And I was chairman of that for the entire lifetime of that task force. There were several other task forces. There was the Security Task Force—which was Steve Kent's, I felt that was a very good task force. There was a tactical task force, what can you do about the dismounted cavalry using IP? Well, that task force never did anything. The others varied. Later on, we took the idea of GADS and we discovered something interesting, that the ICCB was doing engineering. And we thought our charter was to do research. So we really needed to separate the agendas of the ICCB and the focused engineering group, which became the IETF. What had happened was, my task force members in GADS were split. And one guy who worked for me, I nominated as the chairman of this new group, which was the IETF. And that's how the IETF was born.

Russell: And that was Phill Gross?

Mills: Yes.

Russell: Back to the transition between ICCB and IAB, what do you remember about the transition between those bodies? Barry Leiner was the one in charge, is that right?

Mills: Yes, again, you see, the program managers evolved. Each program manager had a slightly different view. There was the feeling at the time, it actually was a really crazy time, the Internet was about to be a victim of success, in other words, life had changed. It was no longer really a research environment, now we had to make some money out of it.

Russell: About what time was this? Around the mid-1980s?

Mills: [Groans] As of ‘86, I think, was the time when it really became important to separate the ARPANET—and the Arpanauts, I like that word, Arpanaut—the haves, from the have-nots. Because it was very clear that the ARPANET itself, which had already been cloned to MILNET, that the ARPANET itself had already served its purpose. And now it was time to move on to a world of Internets. And the question was, how do you organize things in some kind of social fabric. And our core vision at the time was, what is this thing that we call the Internet? What do we do with it? Is it a telecom thing? Is it a research thing? Is it a public thing? We didn’t have a very clear vision—we had no idea that the Internet today—this is just mind-boggling. But we did feel that there was an agenda to convert this research idea into something which was more publicly accessible. The vehicle for doing that was perceived to be the NSF [National Science Foundation]. So NSF started a research program in which they built this phase 1 network of five interconnected minicomputers, which happened to be driven by software, but at the time, the gateway manufacturers—Cisco didn’t exist. Len Bosack did, and I knew him well, and I can tell you about the history about how that got started. At the time, the notion of something you could buy that was the Internet—we weren’t quite there yet. You couldn’t go out and buy an Internet. You could buy a UNIX system. You could put Berkeley software on it with Berkeley networking software. And we did—for academics, that worked just fine. But the box hadn’t opened yet. Now, at this time, the position of the ICCB hadn’t changed. If you consider a research and engineering organization, we were just a bunch of researchers. And a whole new dimension had to go out, in terms of politics, transition to public use, lobbying the government, setting up the NS servers, setting up the ISPs, setting up interchange points, east and west coast,

all that kind of stuff. And most of us weren't interested in that. So we kind of drifted away.

Russell: What was INARC—the Internet Architecture Task Force?

Mills: When GADS was split into IETF and INARC—I was chairman of INARC for many years—looking back at that time, I'm not sure what my charter was, I'm not sure what others expected me to do. "Do good works." I think what kind of cemented my position then was the fact that NSF had just adopted my software, and now the genie was out of the box. So the issue you had to do technically was how do you take all these crazy universities, with all of their crazy networking systems, formerly connected to the ARPANET, and now they have to look out the back door and be connected by this dinky 56 kilobit NSFNET. That had a lifetime of 2 years, and it was grossly overloaded, and I learned an awful lot, and that was a very useful experience for me, probably not less so for some of the researchers, but after 2 years… [end of tape]

Mills: We were talking about the transition to public use.

Russell: Right, and INARC and the NSFNET.

Mills: What I perceived as my own business was being able to see the technology be extended to a much much larger body of use. So I worried about the growing number of networks. I was very near-sighted, in the sense that I looked at my software having to handle 63,000 networks like we had today. There's no way it could do this. I was stupid. I didn't look at the moving technology ball of computers are getting cheaper and larger, and who bothers? 63,000 items in a

directory are trivial these days. But I didn't see that at the time. So I worried a lot, fiddled around a lot, and I felt that there were 2 overriding issues. One of them was getting networks very large, and doing things necessary, with whatever infrastructure we had, to get large. And the second thing was getting people connected. And I perceived at the time to get people connected—because there were no ISPs—everything was dialup. You had to go to modem. I perceived the need to have a good reliable, secure dialup mechanism. And that was done by Van Jacobson, that was called PPP. That was what I was interested in because that was important. As I look back at that time, I was pretty stupid. [laughs]

Russell: So you welcomed broader participation, including commercial participation?

Mills: Oh yeah.

Russell: Today, some people are worried that commercial participation is going to change the network, change the network architecture…

Mills: It already has!

Russell: … disrupt the values.

Mills: It already has.

Russell: How so?

Mills: Well, if you consider who uses the network now, who gains directly from it, well, of course, I use the network that somebody else provisions, to talk to some commercial service—like Google—to use their cycles, which are paid for by advertising. Now, as I said, 20 years ago, I never imagined that. But that's what happens, right? You don't pay to use the network. Danny Cohen's got this wonderful idea of talking about service. He says, it's like Disneyland—you pay for a ride, you pay for a day, or parents pay. Right? We have some combination of all three right now.

Russell: We touched on this a little bit—INARC and IETF split out of GADS in 1986. I was wondering if that had any significant effects. IETF seemed to be more into engineering and INARC stayed more research. Is that correct?

Mills: Yes, exactly right.

Russell: The Internet Research Steering Group [IRSG] came to be in 1989. I'm curious what that was and what your role in it was.

Mills: As one of the task forces, I had a seat at the IRSG. But I perceived that the IRSG was more of a policy resources type group. Frankly, I haven't seen a lot of output from that organization. But I'm not a member, so what do I know?

Russell: It still exists today?

Mills: So far as I know, yeah, but the only task force that's had a longitudinal presence is Steve Kent's Security Task Force. It's always been there and may always be there. That portion of the Internet landscape is really one that's been stable for a long time. The engineering principles have changed. I tell my class this—if I were to implement TCP today, it wouldn't look like it does now. It would be really different. And that's because back in the ARPA days, the resources were buffer space in the gateways and live bandwidth. And you had to tune everything in response to that particular model and that's how TCP was invented. TCP essentially is a resource that you can hold back for a while—that is, to transmit buffers—and you could drop them over the fence, one after another, because it's delay, and you have to be very careful of the network in between, not to overrun it. So that's why TCP is so carefully done. But if we did that today—*bleh*—I've got 100 megabits to that computer, anywhere I want to go. It's the flow of mice in a big pipe, instead of the roar of elephants down a tiny pipe. Quote me on that.

Russell: What was the End-to-End Task Force? Why did it change its name to the End-to-End Research Group?

Mills: Part of that was the fact that Bob Braden runs the group. Bob had a pretty—I wouldn't say narrow, but a pretty focused intent. I perceived the INARC group as being general network research. I'm concerned about what's between the TCPs. He's concerned about the TCPs. But our various missions kind of wandered back and forth. And in fact, Bob was more patient. I was a professor, I had to do professor things. And I couldn't go to all the meetings he was holding. I was a member of that group for some years. I suppose, on paper, I probably still am. But I just

stopped going to meetings, because we weren't talking about stuff which was in my particular research area.

Russell: So there were people who were members of both groups?

Mills: Yes, in fact I can say that differentiated services, between Van Jacobson and Dave Clark—I view that as a network issue. Bob views that as his issue. We had no turf discussions between us, but I felt that, well, this is something I normally would consider part of the network. If Bob takes it over in his space, and runs the group, fine, I'll just go to his meetings. Which is what I did.

Russell: It doesn't sound like a scenario of turf battles…

Mills: No, no.

Russell: … even though the names were different.

Mills: Not at all. "You want to handle this? I'll handle this." One thing perhaps that—a couple things I really think that did show up on the radar, one of them is IPv6. Now that caused a whole [unintelligible].

Russell: And about when did that show up?

Mills: Boy—life is just one series of flashbacks.

Russell: That's well documented, I can look it up.

Mills: That was Bob Hinden. I went to some of the meetings then, but I perceived that there was an interesting discussion going on when we realized… I think it was the first time, with IPv6, that we had to worry about how the user viewed the Internet as opposed to the view from inside the shop. We all—myself included—were looking at the Internet as an engineering adventure. We wanted to make the Internet work well. But we didn't realize that there were some users out there who didn't care about this. They cared about what they could do, and some of the things they wanted to do were kind of orthogonal. And one of them was, billions and billions and billions. Carl Sagan comes to mind—"billions and billions and billions."

Russell: It seems like this effort to transition to IPv6 has been going on for some time now, and there's no clear end in sight.

Mills: I have to say, I have some pretty strong thoughts about this. When we designed IP, I remember a meeting we had at University College London, we were trying to decide how wide the address field was. And at that time, the IP address field was broken into an 8 bit network identifier and a 24 bit host identifier. There were no classes at that time. And we tried to—we threw ideas back and forth about how big the address field should be. And remember when this was—at the time, we were graduating from a typical 16 bit minicomputer into a 32 bit minicomputer. And so we thought, let's be modern, we'll use 32 bits. We had to put the IP

header 32 bits all the way through, and that's because we had a whole bunch of different machines. The most common intersection of all the address renderers was 32 bits. Of course, you have to bear in mind, that some machines had 16 bits and some had 36. And putting a 32 bit checksum on a 36 bit machine is a labor of love. We had some machines where a character was 5 bits, and a word of 36 bits included 5, no 7—I'm doing this wrong—5 7-bit characters. And you're trying to transfer words from that machine to some other machine where you're transferring characters and all of the sudden you have a lot of backing and filling to do. So we thought 32 bits was right—so everything was 32 bits. All of the sudden now we're dealing with 128, and we're doing it all over again. Also at the time we felt that since the IP header is 160 bits, and the TCP headers are the same, so we had 320 bits of overhead all the time—we were worried about overhead. But now, you know, we're not so worried. So a huge header in IPv6 is acceptable.

Russell: So it's a combination of trying to be efficient, but also…

Mills: … the technology of the time.

Russell: And also, how can you scale it looking forward?

Mills: We never thought we'd get anywhere near to filling up the 32-bit address space. That's 4 billion, and that was what we thought at the time, 4 billion is way more than enough, and we thought of technologists talking to technologists. We didn't think about people talking to people. How many Internet addresses do you have in your house? More than one, I bet. OK? I must

have 100. Literally 100. When I first asked for a network number from Jon Postel, Jon had, we had just decided to form class A, B, and C, and my net address was net 29. That's what it was. And Jon said, will you change over to the new system, and I said, fine, I'll be the first one. So I said, "Give me 128, something, one." He said, "I assign the numbers, not you." He said that to me. He said, "I'll give you 128.4." Which is what he did. Now, at the same time I asked for and got 128.5 for Ford Motor Company—I had correspondents there—and 128.8 for the University of Maryland. So I thought, boy this is neat. I've got my own network number. I used it for 25 years—when I came here, I brought it with me. And so on this campus, 128.175 is coming to campus, 128.4 is me, my stuff. It came to a head not long ago when the campus infrastructure wants to take over the department infrastructure. And I lobbied strongly against that. So I said, "Let's use 128.4 as the department," and they did. Not before I got a call from someone. They said, "That's a valuable network number you have—will you sell it?" And I said, "I don't think it's mine to sell." And I thought about it. Does the 128.4 belong to DARPA? No. Does it belong to the university? No. I guess it belongs to me. I guess I could sell it. But I want it. And I asked the IT department here, how much is it worth? He said, "About a hundred thousand dollars." Bleh! I said, "I'm not going to sell it." So I told my daughter about that, she's a network technician for Happy Harry's Drug Store Chain, she said, "Will you put it in the will? Leave it to me?" [laughs] I said, "I can't." So that's how we got 128.4. That's why 4's a small number. But not as small as I asked for. I had autonomous system number 2 was assigned to me. The network system I had at the time was good, it was called DC network. DC stands for either District of Columbia, Drastically Corrupted, or Distributed Computers, take your pick. When I did that, of course, when I developed the gateway-to-gateway protocol, I used that experiment, and everybody knew it. When I came up here, I didn't need it, I let it sit there for a

while. Sometime after I went to look for it, to see, to make sure it was still assigned to me, somebody else had it. How did somebody else get it? It was assigned to me. And as I discovered, they were using the term "DCNet" in their products. I hadn't been watching. So I got a call from someone recently that says, "Hey, what about this autonomous system number, are you still using it?" I says, "I can't—somebody stole it." He got very concerned. He said, "They can't do this, you know, that's something which the IANA won't allow." And then he went back to IANA to get them kicked off. I haven't heard back from him yet. Still pending.

Russell: Were you active in the IETF? Did you start any working groups or were you active in the hierarchy of the IETF?

Mills: I let Phill run his shop, and since I had my own meetings to go to… I went to a couple of IETF meetings, but, they meet so often in so many different places I just couldn't do it. When something came up that affected me, I went to a meeting. I perceived—by the way—I think that is a real success story. It's kind of like ANSI. Only ANSI with legs. ANSI has no legs, but IETF has legs. People go there, and they discuss things, and products come out of it, like IPv6. I think it's a real success story. If you can do something like IPv6 from the standards standpoint, and as a result of this work, my computer has IPv6 on it, that's really interesting.

Russell: It's a different scenario than ISO…

Mills: Far different. The players are different, the agendas are different, the results are different. The IETF I think is a really really good success story.

Russell: In 1992 Dave Clark tried to soothe a rift within the Internet community, and told the IETF plenary: "We reject: kings, presidents, and voting. We believe in: rough consensus and running code." This has turned into a motto of sorts for Internet standards…

Mills: Say it again?

Russell: "We reject: kings, presidents, and voting. We believe in: rough consensus and running code."

Mills: That's true. Definitely true.

Russell: Can you elaborate on that?

Mills: Well, I wasn't that close, I think, to Dave's agenda at that time, but I think the philosophy is exactly right. It hits the nail on the head. That's the difference between the Vint Cerfs, the Dave Clarks, me now, and corporate research and development. There's a lot to be gained by letting the researcher have his head. It's like my job—nobody tells me what to do.

Russell: It seems a summary of the technical philosophy as well as a way of working.

Mills: Yes—I would guess that the occasion that Dave would make that remark would be at the IAB, and with reference to the IETF and the informal style that we had and enjoyed so much.

Now, there's an issue there that the IAB became polarized and became politicized, as it should have been. They were the right guys, and there were a certain number of dust-ups with the federal government, and industry, and who owns what, in terms of names—owning names is a very valuable resource and economic advantage. So, there was a whole part of the community, and there was some bad press back then too, and a certain amount of king-making. But I'm not the reference for that.

Russell: This was in response to a crisis over the address space, and the IAB wanted to move to incorporate CLNP to try and solve the address space problem, and the IETF didn't like that so much.[6]

Mills: There's a bit of history that might not be really relevant—well, I know it's relevant. Are you going to ask me about the GOSIP[7]?

Russell: I am, I'm coming right up to that.

Mills: Do you want to save this discussion until then, or do you want to do it now?

Russell: Let's go now, that would be fine.

Mills: I know something about that, because I was very concerned about that, by being a…

[6] CLNP is ConnectionLess Network Protocol, developed by OSI as part of the ISO reference model.
[7] GOSIP is Government Open Systems Interconnection Profile.

Russell: Maybe you can start by describing GOSIP?

Mills: Well what happened was the Priests of the East and the Priests of the West. The Priests of the West were the Darpanauts; the Priests of the East were the European telephone consortium. What is the acronym for that—CEPT. That's the Council for European Posts and Telegraph. They were an organization of countries, of legislative branches of countries. And these countries needed to protect domestic telephone suppliers. So their mission was to capture the market for their telephone companies—their government telephone companies. So they were politicians, and their vision was to capture business. They were the Priests of the East. And their management, the people below the CEPT, the people who actually were doing things, were marching with those orders. They wanted to standardize things so that they could talk among the telephone companies. They didn't really care about the users. And the Priests of the West worked from the bottom up, you know, we wanted to talk to people. But the fact was that the federal government in this country perceived the CEPT model as being something we should adopt here, not because we had to protect our suppliers, it was just the nature of the ISO, and the UN, cooperation, this kind of stuff. A very strong component of this country advocated that. I don't think it was industry—it was personal. [unintelligible] turning the Internauts into using the European-derived standards, just for the sake of standardization. Like I said before, though, because the European standardization style was as Dave exactly described—he described the kings, and the fiats, and voting, and stuff like that, that's how ISO works. So, the National Academy convened a task force. And the task force was given three agendas—well, it attempted to make three outcomes. Outcome number one, adopt ISO, set fire to TCP. Outcome number 2 was adopt TCP, and let progress ensue. Number three was to let TCP/IP to exist for a few years,

and then get rid of it. That was their charge. I was never on that committee, but I watched them very closely. [unintelligible] was on it, and so was the guy who used to have this office, Dave Farber.

Russell: What committee was that?

Mills: This was a National Academy committee.[8] I think it was from the computer science directorate. So they met and they decided outcome number three. They said, this has to do with, GOSIP was Government Open Systems. And what the output of the committee was direction for the government that they would fund suppliers, government contracts would fund OSI and the budget for these agencies—like the Army, Navy, whatnot—their budget would have to include OSI. As an optional matter, they could include TCP/IP, but only if it was going to last for a short while. And that was the output of that committee, and there were people who believed that, believed it strongly. I got involved with it many times as consultant to the Navy, the Army, DCA, and they would say to me, "Geez, how am I going to do this? I've got TCP/IP running, now I've got to get OSI running." And I essentially told them, I said, "I do not have good news for you. You can't buy it. I can't go to Digital Equipment and buy OSI." I can buy Digital networking, which clever enough had the abbreviation Digital Network Architecture—DNA. Just a comment on the margin: DNA was the best conceived engineering product I've ever seen, in terms of networks. It was *designed*. There were architectural books on it. There were implementations that were close to provably correct, and it was well done network engineering. In the IP we had, it was not good engineering. It wasn't provably correct. But Digital's was; of

[8] Computer Science and Telecommunications Board, *Transport Protocols for Department of Defense Data Networks: Report to the Department of Defense and the National Bureau of Standards* (Washington, DC: National Academies Press, 1985).

course, they are no more.

Russell: What happened to DNA?

Mills: DNA I think is still used in the DNS system. If you have a DEC system that runs on a VAX, and it uses VMS software, it probably still does have DNA running.

Russell: But never implemented as broadly as TCP?

Mills: No, no. So getting with GOSIP. What happened was, I went to the Army, I think, is the one that really got me. I said, you know, "I haven't got any hope for you." And they said, "We're going to have to continue to use TCP/IP, because we can't get OSI." At about the same time…

Russell: This is about the late 80s?

Mills: Roughly, yes. About the same time an interesting thing happened. In the INARC, I had raised an item to say why can't we talk both protocols? If we could push the CNLS—ISO 1594 I think it is—if we can—that didn't sound right, whatever it is, it's the datagram. ISO had to be dragged kicking and screaming into connection-less transport. They hated it. They wanted X.25 underneath TP class 4, TP class 2, and then above that, you know, they wanted rigid—and they hated TCP/IP. So I think the ISO kind of gave in, and said, well, we're going to have to have a datagram transport. And they did it, and it wasn't too bad, either. And so the question I had for

the group was, why can't we run both? The packet comes in, the packet goes out, you're just looking at headers, and what header does this go. And that idea was—I don't think that idea was unique to me. So, when NSF was expanding, and they adopted IBM as a supplier, and as that evolved, IBM put connection-less transport, and it seems to have worked. It worked—it probably, I'm wondering if today, how much of the Internet would right now will in fact switch this thing? I wouldn't be surprised if Cisco software would work. So it was the kind of thing that was a freebie. It's like virtual tunneling, only it's tunneling with the network layer. So that's what came out of that, and I don't even know if anybody ever used ISO on the NSF backbone. I don't know. So what happened—two things happened. One of them was that, once upon a time there was something called JAN standards. Joint Army-Navy standards. This came out of World War II. JAN standards were around for a long time. Vint managed to get 792 and 793 declared as JAN standards—now that was a labor of love. But the military actually did that, so they are military standards. Now at the same time, however, the military, about the time of GOSIP, had this tradition, this agenda shoved down their throats—which eventually got you Halliburton, by the way—the idea was to get the government out. It was called COTS—Commercial Off The Shelf, COTS. This became a mantra. Get rid of the standardization process. One of the objectives of the Army at the time, use COTS, and avoid standards. Don't use JAN standards. I remember having a discussion with a bunch of guys on this, and I said, "You want to use TCP/IP because it's not standard, right?" They said, "Yes." And I said, "Well, look again, it's a JAN standard. [laughs] Standard 78." I got a big laugh. It is.

Russell: You mentioned earlier that you believed that GOSIP didn't catch on due to economic issues.

Mills: I think the killer was COTS. The agencies perceived, use commercial products. Well you can't buy OSI, but we can buy Unix—we can buy TCP.

Russell: So OSI had to be implemented in something for them to buy?

Mills: Yes. Now, what I think is very interesting is the Europeans. I was in a meeting in Trieste, Italy, with Vint Cerf, as a matter of fact. We were both there. We were Priests of the West, very definitely, and everybody else in the room—to a man—were all Priests of the East. And the moderator of the conference was from CERN, Switzerland. And he had a rather large network there, which was not, of course, TCP/IP, and he was like Nikita Khrushchev, he pounds the table and said, "There will never be TCP/IP in CERN!" He said that. I heard him. And well, the situation is quite different now.[9] There is however an interesting exception to that. The European Space Agency and NASA had gotten together a set of standards for spacecraft. And they did this in order to inter-work at the spacecraft level between the USA satellites and rovers on Mars, and the NASA community. So at the lower levels of the protocol suite, we already had these protocols established, you can't change it. There was a significant number of people at NASA-Goddard who wanted to abandon that and use IP/TCP instead. I'm watching this thing develop.

[9] CERN stands for Conseil Européen pour la Recherche Nucléaire. In spite of the CERN moderator's conviction, TCP/IP networks were soon deployed at CERN. Ironically, the most familiar Internet application—the World Wide Web—was designed between 1988 and 1991 by Tim Berners-Lee, at CERN, to run over TCP/IP. In his autobiography and in other historical accounts, Berners-Lee recalls how he admired the design of TCP/IP, and was surprised at CERN's "lateness in adopting the Internet." Tim Berners-Lee, *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by its Inventor* (San Francisco: HarperSanFrancisco, 1999), 18.

Russell: Vint Cerf is involved in that.

Mills: Yes. That's the Mars Internet. As far as I know, I had the only official contract to work the Mars Internet. My mission was to do NTP timekeeping. It was fun.

Russell: I want to get back to NTP, but a final question on these Priests of the East. The history of the Internet documents who the Priests of the West were pretty well, but I don't know if I would be able to name any of the Priests of the East. They seem like faceless European bureaucrats from the way you're describing them.

Mills: European Union. That's the way they work. They're faceless.

Russell: So they weren't people or personalities that you worked with in the same way that you worked with Postel, or Cerf, or Kahn?

Mills: I have a couple of comments on that. I believe the father of ISO—the name has to come to mind, he wrote the first IEEE article on the whole principle of layering.

Russell: Zimmerman?

Mills: He was one of them… Pouzin. What's his first name—his last name was Pouzin.

Russell: Right, Louis Pouzin.

Mills: I met him once. He wrote the first IEEE article, which I read, and I read this thing again. And if I look back on it, that was a very seminal article. The whole notion of how to address services and build architecture, it's a wonderful article. And the model is very good. We don't exactly follow the model, but the principles are there. So he's a Priest of the East.

Russell: When Vint Cerf talks about TCP, he credits Pouzin's ideas as having an impact on him.

Mills: Who else can I think of, technically? Others that I know of, and these include people that I've met and worked with in the DARPA community, including Peter Kirstein and Paul Spilling at NTRE [Norwegian Telecommunications Research Establishment], and do I have some from some other place—yes, the guys at DFVLR, I'll think of their names in a minute—these guys were more facilitators. They were Internet researchers, but they had to be two-faced, they had to answer to their bosses who wanted OSI. Especially in the case of Peter Kirstein, he had a really strong advocacy to adopt the Rainbow Books, which would have been the UK's contribution intended for the European Union.[10] I don't think we ever got there. I implemented the grey book, which was the file transfer. And, it worked. But, I used it before there was FTP. I implemented FTP but nobody else did. I talked to myself. I could talk with—there was some other name to it, NVP or something like that, in which I could talk to my friends overseas, so I used it for a short time. I think a number of books are long since gone.

Russell: What about people who promoted or championed GOSIP? Marshall Rose is one

[10] The "rainbow books" were a series of guidelines for computer security that were designated by a different cover color (Grey, Tan, Bright Blue, Neon Green, etc.), to address different aspects of trusted systems.

name…

Mills: Marshall is the one that comes to mind. He was chief, cook, and bottle-washer. Marshall implemented some things. Have you talked to him?

Russell: No.

Mills: He's a good person to talk to. He was our man at OSI, and yet he came to the meetings. He didn't try to proselytize. He was trying to learn from, and contribute to.

Russell: I've spoken to Mike Padlipsky a couple of times,

Mills: The Internet—what did we call him? The Internet Seer? He has a very colorful vocabulary.

Russell: He does, very poetic at times.

Mills: I just got a message from him about that. Without looking at the author I knew it was Mike Padlipsky. I have a certain reputation myself of speaking what people call Millspeak. I love the English language, I love to kind of torque it in various ways. And I think you will know, if you read my messages, you'll know it's me. I had to submit a paper once that went to the editor, came back from the reviewers, and I changed it and submitted it for re-review, and the reviewer came back and said, you took the Millspeak out of it!

Russell: Is that a bad thing or a good thing?

Mills: I don't know, it's just in my character. Ginny Strazisar, who was a researcher at BBN many years ago once said at a party, she once said, I don't understand a thing Mills says.

Russell: It might have something to do with this office, I think Dave Farber is famous for Farberisms.

Mills: Yes, yes.

Russell: Maybe it's the water.

Mills: He's the first one to say that the fundamental resource limitation in the future will be the space on your belt. [end of tape]

Mills: I first met him [Farber] when I was just a young a graduate student, and maintained contact all the way through. He wanted me to do something—I was a young grad student, and I liked to build things, and he wanted me to—shoebox, it had to do with a shoebox. He wanted a computer in a shoebox.

Russell: In a shoebox? Literally in a shoebox?

Mills: Literally, yeah. This was in 1962 or 3. [laughs] He got his wish, eventually.

Russell: You would be fabulously wealthy if you did it at that time.

Mills: At that time, sure.

Russell: Do you have any further thoughts on GOSIP or the politics behind it?

Mills: Now remember, this is one man's view, and I wasn't part of the halls of power in any way with GOSIP. All I saw was what I saw as a consultant to the agencies. What I saw was fear in the agencies. They didn't want to go to GOSIP, it was a nightmare for them. They didn't have a network that could talk with GOSIP, they didn't have computers that could talk GOSIP, but they did have TCP/IP. Again, the strategy at the time was if somebody gave them some money to buy a computer, and it came with GOSIP strings, they would ask for a waiver, and then they'd keep asking for waivers. Eventually the pressure just went away. There was a real fear at the time, when I came here, there was a real fear at the time that we'd wind up with a profusion of different systems on campus. And the people—Dave Farber, in particular, and the then-chair of the department, Pete Warter, and I felt very strongly that we wanted TCP/IP to be on this campus. So this department was TCP/IP, but of course there wasn't any networking anyplace else. So if you wanted a good network you came to us. And what we had eventually spread to the rest of campus, but not without a fight from IBM. The campus administrative tools were all IBM, they ran on IBM VMS machines. The IBM VMS culture was totally different than the TCP/IP culture. I don't really know if we actually even have survived that division yet. We

might still have an IBM VMS system on campus.

Russell: What do you mean by the culture is different?

Mills: [pauses] I think the way to describe that is to say that the IBM culture is batch-oriented. Truly. They never were able to wrap their minds around the notion of interactive time-sharing, and I've got proof of that. In the middle 1960s, Michigan got an IBM 360 model 67—that was an experimental machine, there weren't many of them around. They were dual-processor, configurable systems, very expensive and very large. And I got to use it. Myself. *Only me*. But the idea was that we had to—the whole notion of supplying service to a campus population was, you've got a terminal, you dial a telephone number, and you're dealing with yourself and sharing some files. The way you talk to the system is interactive terminals. IBM is very different. In IBM, the mode of operation was you control your program with punch cards. You hand them to a faceless person behind a window. They take that deck of cards, they put it on a small machine, and they let it tape. They put a whole bunch of decks of cards on tape. And they take that tape off the little machine, and you put it on a big machine, and run one job after another. These jobs produced printed output on a separate tape, and you take that tape off and put it in the small machine, it's got some printers on it, and you print everything and you get back the printed output and the cards to the center, he comes and picks it up. The idea was that your job turnaround was a day. So whatever you, however fast you could work was limited by the turnaround of that deck of cards. And so you learned to think that way. But look at how different that culture was. From submitting cards and getting them back a day later to typing a line with the keyboard. It was a totally different culture. So IBM—that culture of the deck of

cards lives on today, only instead of the deck of cards, you have job control language. It's just cards in a file somewhere. I could be selling our administrative department short when I say this, because I have no evidence they ever changed, except for one thing, in that everything I do now involves the Web. We don't use paper in this university anymore. That causes some real problems, not the least of which for me, because I have limited eyesight, and being able to deal with the little tiny print on some of these websites, I just can't deal with it. And I get angry and bad-mouth Web designers.

Russell: That's another standards issue…

Mills: Oh yes, the W3C [World Wide Web Consortium].

[brief discussion on Mills' website, browser compatibility, web design and coding]

Mills: Meanwhile I got angry about cookies. Don't put cookies on my machine unless I give you permission. No, we won't put any cookies on your machine, but it's like the IRS, when they want your social security number. You don't have to give us your social security number, it's required by law that we can't require you to do that. But if you want to file your taxes, you have to do that.

Russell: Do you watch the W3C, or do you keep tabs on them?

Mills: I don't keep tabs on them, but I did visit the website, and visited the—what is it, the 504,

the Senate document. There's proposed legislation that would require certain things of Web services. It's been a long time since I've seen that document, but the whole agenda of saying, if you provide a Web, then our legislature will require you to obey these following constraints. Just as if you had a wheelchair ramp. And I'd watch that issue, and have quoted that to university administration, saying that "if you build a website, be sure that it conforms to this." I got a very nice response from the Vice President for Information Technology, he said it's a good idea and he'll do it. But there's so many Web designers, that have their own little agendas, and they use their tools, and their comment back is saying, "I use this tool, I have to do what it tells me." Which I come back with, "Look, you can do this in the source." I happen to use GoLive, and I stay very close to the source. I don't do anything fancy, and I don't want anything fancy!

Russell: I have some more questions about DARPA support. You mentioned that there was a culture where DARPA would give you money and do what you wanted.

Mills: You still had to make a proposal. In the beginning, when I first started this business, I'd go to someone like Vint Cerf, and said, "I'm going to do this," and he said, "Fine," and gave me the money. You can't do that anymore. You have to go to BAA [Broad Agency Announcement], where they list the announcements. You have to respond to intricately-worded invitations for proposals. You have to do pre-proposals, and you have to do white papers, you have to do proposals. There's a great deal more paperwork now. I don't—an interesting question is, could the Internet have matured at all under the present-day regulations? I don't know. It'd be much more—we didn't have the Freedom of Information Act. We didn't have the rule that says that public meetings of government agencies have to be open to the public. You

have to have an agenda, there's a name for that regulation. But, I was a member of the NSF Network Technical Advisory Group, the NTAG. It was very informal, I'd go to Washington now and then, and go to the fourth floor of the, opposite the Executive Office Building there the meeting, no problem at all. Come and go. That's how I got the software through the NSF. And you'd put the back window open and the White House, you could see the missiles on the roof—yes, missiles on the roof, you can see them—they exist. [laughs] My eyeballs saw them. But the whole point was very informal. Well, the NSF had to stop doing that. They had to discontinue this kind of thing, because everything we did had to be, have an agenda, public agenda published six months in advance, we said we can't work this way.

Russell: ARPA, especially in the early days, seemed to be designed to counter that, to be light on its feet and move quickly.

Mills: Yes, it was. I'm not sure I could have done the work I did if I had to respond to a Broad Agency Announcement. I mean, it really was the case where they just kind of said, "Do what you think you need to do."

Russell: What about oversight on the other end? Not the application process, but once you were done with something. Was there anyone watching to make sure it had military use? Or potential military uses?

Mills: I don't think there was anyone—there were certainly people watching out for what the director of DARPA, for instance. He wanted to see a program, he had to sell the program to

Congress. The program managers had to sell the program to him. And then the individual guys like me had to sell to the program managers. There was always that oversight. DARPA had—a very good idea—they had a meeting of Principal Investigators every year. We would go and we would tell each other what we were doing. As much as I hated it, it was a pretty good exercise. You were talking to other people in the room, with technical backgrounds, who were going to critique what you were doing, and just like any symposium, you had to sell your product. So that oversight I think was very good. Not only did we find out what other people were doing, but we could also pass judgment on it.

Russell: That was a sort of peer evaluation?

Mills: Peer. Peer-to-peer, that's exactly right. But isn't that how science works? There's not a lot of people who will pass on my research that I do for a living. My peers will do that. But there's nobody on above who says, "Bless you my son, you've done good."

Russell: Do you know if any of the protocols you developed through DARPA funding were used in military engagements? Or did you ever have any qualms about receiving money, or think of the money you received as coming from the military?

Mills: I don't think that this was the responsible science, the—what is it, the Association [Union] of Concerned Scientists? Whatever it is. That was mostly in response to things like the atomic bomb. I don't think any of us thought of anything we were doing as being a military controlled or used product or project. Quite the contrary, we had to defend the fact that, what's the military

going to do with all this stuff? There was a good deal, in fact, the fact that Bob and Vint sold RFCs 792 and 3 to the military as JAN standards, the fact they were able to do that is a real reflection of their political astuteness. But I don't think there was any feeling that what we were doing was aiding comfort to the enemy, or immoral, or anything like that. We felt quite the opposite.

Russell: The historians of physics make a big deal out of that, but I don't see that much of it in computer science.

Mills: I was once, as a freshman in college, I went to a talk given by Edward Teller, and at the time, of course he was working on the bomb, he was trying to promote the idea of using atomic bombs in peacetime, for construction purposes. And what he had proposed, and I was there when he said it, he said, he proposed to use atomic bombs to dig another Panama Canal. The words in his mouth. I went away from the meeting saying, ah… [laughs]

Russell: You came from academia, and then worked for COMSAT, which might be described as government, and then went to Linkabit. Can you tell me a little bit about Linkabit and about working on the Internet in the private sector.

Mills: Yes, I came from many years of academics, I was an academic character. In fact, I was denied tenure at Maryland, so I had to find a real job. I went to COMSAT, and what attracted me to the place was Harry Van Trees, who was a well-known scientist, wrote a wonderful book on estimation theory, and he was a senior vice-president for research. I thought, hey, I'm

moving up in the world. I was the only guy on the floor who was not from MIT. [laughs] Harry was a graduate of MIT, by the way, who's now at George Mason University. And so, of course he had program managers for him, a guy named Estil Hoversten who was responsible for the COMSAT participation in the DARPA project, and life moved on. Eventually Harry decided to form his own company, and since he knew the guys at Linkabit, and since Linkabit had an absolutely sterling reputation in everybody's mind, I wrote—boy how names go—Ira Jacobs—was probably the most revered communication scientist in the country. And so, they had started Linkabit, and they wanted to expand to the East Coast, and gain government exposure. Harry was under-secretary of the Air Force, and so he was in the contracting process, and he knew the military was about to contract for a very expensive satellite system, and he wanted that system. So he formed Linkabit East. And the first guy he hired was me. So we went out and I hired some people—Phill Gross among them—and everybody that worked for me at COMSAT came with me to Linkabit. Well, to make a long story short…

Russell: How many people?

Mills: I think I brought two with me. So this was a small company, and I had to find places to live, and so did Harry, and we tried to start the company essentially. It was an adventure in its own right, but the work that I had was DARPA. I didn't want to do anything else, but Harry wanted that contract. Well, he lost the contract. So he had to live on what everybody else in Washington has to live, a bunch of beltway bandits. So we were lobbying a lot to get business. And, to make a long story short, the way you live as a beltway bandit is, you're selling yourself, you're selling man hours. That's what you're doing. It took five years for Harry and I to realize

that you make money in this business by selling man hours cheap. So how do you get cheap man hours? Well there are a couple of tricks you can do. One of them is that you entice the young engineers to work overtime. To do that, you buy them dinner. Dinner is a expense account, that's an overhead. The cost of the dinner pays the man hours that you get out of the guy. [laughs] These are dirty tricks, you know. But that's how you get the cost of a man hour down. The other way you get it down is that people like me are really expensive. Take my salary, multiply by three. That's what we charge the government. We gotta add up all the loads. So you can't use very much of me on a cheap contract. So you write me in at five percent. Now, I'm not useful at all, except as a resource. The poor young guy who has the contract doesn't know what to do, and I've got to tell him, do it this way. He needs guidance, he needs me to be there more than five percent. I'm spread over twenty contracts. And so I worked pretty hard too. And the poor guys, they don't get enough of me to do a good job, and they don't know what to do. So the result is a contractor like us did lousy jobs. I wasn't proud of what we did, other than the DARPA contract. I had a contract with the CIA that I hated. I hate the CIA. [laughs] Over at the CIA building, the spook mentality is just maddening. It's not secrecy—nothing to do with it. The CIA and the NSA look inward so much, their cultures are so inward-directed, they just don't work well outside their cultures. If one of the spooks decides to tire of the life and get a job outside, and I give you the retired general is a good example too. They look for a job outside and they come up against the mentality of the beltway bandit, and they just disintegrate. They just can't deal with it. So I did that for a total, I guess, of ten years. And I got pretty tired of it. I just felt, this is not something I want to do for a living. Happy as a clam when I came here.

Russell: Then you returned to academia.

Mills: I came back in as a full professor. I left as a untenured assistant professor, and came back ten years later as a full professor. I never went though the review process.

Russell: And Dave Farber was here when you were here?

Mills: I think he's one of the reasons I was here. He was able to make the case.

Russell: And you worked together for a few years?

Mills: Off and on, yeah. We had different research projects, but we socialized and we were much I think the same kind of personality. He's a very twinkle-eyed personality… Dave and Peter Kirstein are both world travelers. They both like to wander around and spread good works.

Russell: This seems to have something to do with the culture of the Internet, this international culture we were talking about before.

Mills: We like our foreign friends. I love my foreign friends. I think, perhaps, because it's nice to be kind of worldly.

Russell: Throughout the nineties, and it's especially true today, many people who have been involved in Internet standards are worried that the IETF and the Internet standards process will

be, or already is, captured by the vendors. Do you share this concern?

Mills: That's not necessarily a bad thing. Now, if you want to say—who are the stakeholders? Good question, who are the stakeholders? The stakeholders are really the vendors. Am I a stakeholder? I'm just a researcher. No, I don't see anything in the IETF that needs my input. It's a standards organization. Well, that's not quite true, I have several agendas of my own. One of them is network security. Because my research idea is network security. So I pitched the ideas to the IETF, and I did. So if you are going to do security in this particular way—this is Network Time Protocol—then here's my suggestion for how to do it. Because I have a research interest. Now, of course you need to say, what vendor would want to own that issue. Maybe—well, I don't know. I'd be delighted if some vendor were to take it up. So, I think the stakeholders are the vendors. And the ISPs.

Russell: Are you still active at all in the Internet standards groups associated with the Internet Society? IAB, or IETF, or INARC?

Mills: Only—well, of course, INARC is no more. The IETF will last forever, as it should. The postel.org I hope will last forever. That, by the way, I'm sure you've seen that.

Russell: Yes.

Mills: I'm a little bit uneasy about that, because there are so many things that I've forgotten, and I keep thinking, gee, if I only wrote that stuff down, if I only kept a notebook. But I didn't do it.

What I did do was every Friday, every Friday I had to do a corporate report for Linkabit, as everybody did. You took the previous week's report, and change a couple of words. The monthly report for Jon, I did that for years. I put that on floppy disk—8 inch, double side, double density floppy disk. Nothing around to read it anymore.

Russell: You still have the disks?

Mills: I still have the disks, but I have no way to read them.

Russell: This is a problem that historians are increasingly paying attention to. People in the past had lab notebooks…

Mills: Yeah… Mea culpa.

Russell: What about email archives? Same problem?

Mills: Well, I have a zip disk with 10 years of, let me see. Running from—I think it goes back to certainly the time I came here, and before. Because I used to keep mail on a machine at ISI, on the west coast. ISID was my mail. I managed to archive much of that stuff, but I don't think I—that never made it to floppy disk. The archives I have on zip disk I think are the first ten years I was here, which would mean from 1986.

Russell: And that can be read by…

Mills: It's on a zip disk. If you wish…

Russell: It's plain text?

Mills: It's plain text. If you wish, you're welcome to it. There's an awful lot of stuff there that doesn't pertain to anything that you care about…

Russell: I can imagine. That's a related problem with digital materials, if they're there, there's so much of them that it's hard…

Mills: Just please don't make me go through and remove all the personal stuff.

Russell: Can [you] describe your work creating NTP [Network Time Protocol]? I was hoping you could talk about how you came to be interested in time and why it's important for the functioning of networks.

Mills: Looking back at motivation, I have to say that what attracted me to it was very interesting cultural issue. I've always been very much concerned about what it means to me, what kind of animal am I, how am I trained? I think you have one of those [holds out a business card], look at the title.

Russell: Professor.

Mills: Yeah. Of?

Russell: Electrical and Computer Engineering and Computer and Information Sciences.

Mills: A lot of words there, eh? A lot of words. That's me. I have feet in both cultures, I've been trained in both cultures, I've taught in both cultures, so if anybody, I think, has feeling about what it means to be a computer scientist or electrical engineer, you know, *moi*. Now, I've said to my students many times, "What do you consider yourself, a computer scientist or engineer?" And they all say, "mumble mumble mumble," and I say, "No, see, it's your toolbox. It's how you're trained." And very few engineers today are trained in computer science, and vice-versa, although we do take each others' courses. Maybe there's some crazy professors like me that really, that's the die-hard, are research-prepared to be multicultural if you wish. Now, consider computer time. It has firm roots in both cultures. It is a computer protocol of interest to the computer science community. But if you dig down inside, there is a great deal of control theory in there. There's computer engineering and computer science. So NTP is a wonderful place for me to use my background. And, I have to say, not many other people who can speak my language. Now that's what computer time is. It's a way of distinguishing my own career and saying, "This is me, and you're not." And I have found—essentially I've found nobody that has that kind of kit. I'd dearly like to talk to somebody like me. Very few.

Russell: It seems you've been exposed to such a wide amount of people.

Mills: Yeah, and I've had a few differences of opinion. Primarily because I think of myself as a researcher that designs the protocols and algorithms. I understand the technical aspects of it. The production aspects, the—release engineering has become a new field now. Release engineering. That's what they're called. People call themselves release engineers.

Russell: I've never heard of it.

Mills: Something leaves me. And now you've got to take that product and make it work in a large variety of different computers, operating systems, and cultures. And making that happen is a legitimate part of engineering. You have to figure out what are the tools to maintain this thing? How do I keep it going? How do I distribute changes? How do I update the firmware? There's a whole body of work there, and the tools to configure for different operating systems are draconian. There's this thing called autoconfigure. If you do a package today, and you want it to run on Windows machines and Unix machines, and you expect to need it, then you have to have provisions in the code to cater to different operating systems. And this autoconfigure thing figures out what the environment is, and tells you… I don't understand it! It's beyond me, I don't have credentials in that field. I have, for instance, the documentation, it's done on the web. I get fierce objections to that. They wanted to do man [manual] pages. I said, "I don't do man pages." Why not? Because my style is html. And they say, no we can't accept that, we have to put a converter in there, and then the actual distribution will put man pages. I get angry with them, and say no, you can't do that. Don't do that. I get serious objections, philosophical points like that. But that—I can see, you see, that stuff, that's not me. That's somebody else. So in recent contexts, I've just surrendered. Here's the body of code, you do with it what you wish.

But it's used in such a wide variety of computers, everybody's routers are running this stuff. I think I told you about the time when I showed up in Netgear, and Netgear had a product. That product was a home router which had NTP built into it. It was non-configurable. You didn't even know about it. It just ran. And, if it configured, it would go out and talk to one time server at the University of Wisconsin. And, if defectively configured, Wisconsin would suddenly reply back, and the reply would get thrown away by the firewall in this user-unfriendly network router. It wouldn't too bad except when this didn't happen, that router would try every second continuously. And even that would be survivable if there weren't too many of them. Netgear had sold 750,000 of them. So the net result was that the ISPs in Wisconsin drowned—denial of service. And that wasn't too bad until I checked into it. They asked me to be on the committee, I was. And so, we contacted Netgear, and Netgear had no idea what I'm talking about. I said, "Your product is overwhelming the network. It's defectively engineered." And the answer I got back was, they didn't engineer it, they outsourced it. They outsourced the design, they outsourced the manufacture, they had a contract to manufacture 43,000 units a week, I'm sorry, a month. 43,000 a month. And there were many orders in the pipe. I said, you must change it, you must change it right now. Well, with 100,000 units behind, they can't change those. Change the firmware. "Well we can't do that." "Sure you can, you put it on your website, say download the right firmware." Well they didn't write the firmware—didn't know how to do it. Well, I got really angry, because my name was getting mentioned in sort of—I'm the single person choking the Internet. So I got really unhappy and I rewrote RFC 2030 that describes SNTP. And so I dropped that on the RFC-Editor's desk, and I ran across the Postel formatting problem again. So, I said, for God's sake, this is an informational RFC, get it out there so people like Netgear can get it to their suppliers and they won't make these bad mistakes. Well I'm in

the IETF now, and they want to make it a standard, so now I've got to go back and re-edit it again for the standards process. [end of tape]

Mills: … the Internet-Drafts process again, and we'll fight all over again about the Postel format.

Russell: What would you prefer to the Postel format?

Mills: PDF. Give me a stylesheet—PDF.

Russell: I've seen this on the IETF list, you're not the only one bringing this up.

Mills: I know, I know. The issue was in RFC 1305 I wrote in late 1992, and I did it in PostScript, this was prior to PDF. And, it's about the same time that the IAB had just decided, as a matter of principle, that everything had to be in ASCII. But I had already had it submitted, it was already considered for a standard. So I said, look, grandfather this thing. I'm not going to write this the way you want it. Just grandfather it. But they refused to do it. So, as of right now, it is still a Draft Standard. Literally, it's a Draft Standard. What can I say? So—the RFC Editor, Bob Braden is the RFC Editor. He well knows my position. But I don't think he's the obstruction. It's the IAB. And the IETF.

Russell: I wonder if you had any other thoughts to add. The area I'm most interested in is some of the organizational and cultural issues, especially in the 70s and the 80s, before the broad commercialization of the Internet.

Mills: Let me see if I can think of any organizations that DARPA came up against or worked with. The computer science community and the UUCP [Unix-to-Unix Copy Program] mail, the whole culture that was kind of in parallel, applications—they eventually did the news system, people like that. They were very much involved, I think Dave Farber perhaps represents that community better than anybody else. They had their own culture, they were, of necessity, forced to use X.25. They probably were the best people to cement the idea in our community that X.25 would never work. Because the ISPs, well, it wasn't an ISP then, it was something else, but they felt that if the customer talks to us, they talk to us in virtual channels—X.25. But the resources were so miniscule that a virtual channel was an expensive resource. And you couldn't get it to go very fast. So you get multiple virtual channels. So what they did was to make a driver that would use all the virtual channels available for transport, and TCP/IP over that. That was a coup. They actually managed to get 50 kilobits of throughput on the line. They were pretty good at having—they were forced to use public facilities. Their technical contribution was just that. There were, I think a very good thing was working with the national research organizations, especially in the aerospace area, working with DFVLR and NTRE, and CNHCE in Italy, and exposing them to the DARPA mission, and also the reverse was really very useful as well.

Russell: What about Asian colleagues? Or colleagues outside of Western Europe?

Mills: They were dark. No participation whatsoever. Later on the Japanese got involved with IPv6, I think, with KAME, and they did a good job on that.

Russell: What was KAME?

Mills: I don't know what KAME stands for, but the Japanese—the IPv6 implementation for Unix. We've had several research projects, the most, I think, important one besides DARPA was NSFNET, which as a vehicle in and of itself, doesn't exist anymore. And also, the CAIRN [Collaborative Advanced Interagency Research Network] initiative, where we all got one and a half megabit lines to our offices at a time when they were very rare. And we had about fourteen different organizations, including the ISPs, including the telephone companies. And so this got us into organizations like Sun [Microsystems], computer manufacturers, people who don't ordinarily play in this game. That was DARTnet [Defense Advanced Research Technology Network], sorry, and then DARTnet evolved into CAIRN, and CAIRN was the same thing, supercharged. So they were working out how to use 1.544 megabits and how to put gateways between that and connections to lower-speed lines, for instance. In that sense, we had wide participation in those particular projects. There's also Internet2 and things like that, but I think you know about that. I think the CAIRN website still exists. C-A-I-R-N, like a rock.

Russell: What accounts for the success of IPv4?

Mills: IPv4 or IPv6?

Russell: IPv4. The Internet, as we know it.

Mills: OK.

Russell: And part of it is what you said, that vendors were already implementing it in equipment. I'm trying to get a handle on the length of time between work in the IETF or the Internet standards groups and implementation by the vendors; in other words, how a standard went from the research community to being implemented.

Mills: Well, In terms of the original IP, the only vendor that survived was Unix and Berkeley. That's really why the success of the Internet right now, I think you have to say that there is a specific target that you'd point to and say, this is why it exists. Unix and Berkeley implementation of TCP/IP. That happened purposely—when I was in the INARC, one of the big problems I had was, Mike Karels from Berkeley was just ungovernable. He was a real problem. It got to the point where DARPA—I forget who it was—said that I have talks with Mike Karels. Mike Karels was part of the INARC. I gave him real attitude. "Why did you do this?" But whatever he did, it survives today. But the fact that this Unix system could be plucked—for the academic community, it was free. And until, relatively recently when the copyright ran out, the AT&T licensing ran out, until that time it was expensive as hell for commercial use. But after that time, and especially since Linux and FreeBSD have shown up, man, well there's no choice, it's *free*. Of course, you find it everywhere, you find the protocol stack—the Tahoe protocol stack—is everywhere.

Russell: The Tahoe protocol stack—is that the Berkeley…

Mills: That's Berkeley, yeah. It went from Reno to Tahoe, and points West. What's beyond

Tahoe, but you're still in Nevada. Donner Pass? [laughs] I used to live in California. Keep the TCP out of California. Go the other way. So I think that's really the fair thing to say, TCP/IP and Berkeley and Unix is what drives the bullet.

Russell: And Bill Joy was one of the people who was doing that?

Mills: At Berkeley? Yes.

Russell: So then this proliferation of TCP/IP products came from that?

Mills: From that *one* implementation. It's like the mother of creation. That's where it came from. And it's been copied and improved ever since. See, at the time, Bolt Beranek Newman had a TCP/IP, they had built one. And people had built them for other machines, like the TOPS-20 and Multics. But the ones that survived were the ones that were in the PDP-8, because the PDP-8 was cheap and small. So there was a version that was—several versions were done for the PDP-11. And eventually DARPA told Mike Karels, take that stuff from BBN, and put that in the Berkeley Unix, and distribute that. Mike Karels said he would not do that. Instead, he re-engineered the Berkeley TCP along the lines that BBN had, and produced that. I guess I can't fault him too much for that, but as I say, he didn't respond to any directions. I caught him only once. In Berkeley TCP 4.2, Mike had thought that the world exists on an Ethernet. The retransmission time-out is initialized at 200 milliseconds. It caused real grief when he put it on the network. [laughs] He eventually agreed to make it larger. [end of interview]