Oral History of Gordon Bell, part 4

Interviewed by:
Gardner Hendrie

Recorded October 19, 2006
San Francisco, CA

CHM Reference number: X3202.2006

Hendrie: We have with us today Gordon Bell in a continuing series of interviews covering his long and very eventful career in the digital world. I think we left off last time with the beginnings of the VAX project and the collecting of a team. Could you start there and talk about once a decision has been made to proceed to design this system, how the team was assembled and who were the leaders and what were the first things that needed to be done?

Bell: Okay. The first meeting now looks to be April 1st, 1975. I have on my website a report by what was called the VAX A group. So the VAX A group was the small team that was with the lead architect, Bill Strecker, and we talked about what the machine is going to be and all of the issues and made various decisions, and that was the team that then spread out and actually implemented it afterward. There were six people. I led the group. It included Bill Strecker who I claim is the architect of VAX per se. He had been studying this kind of an architecture while part of the research group. Steve Rothman who was a hardware implementer and was part of the 11/70 group, and he communicated with the hardware group to actually implement the first machine. Dave Cutler who was responsible for VMS and was going to be the implementer of that. Richie Lary who was I'd say probably the most versatile generalist in all of DEC and still remains that way. And then Tom Hastings who was a very meticulous scribe and actually took all the notes of everything we did. And this was described in a memo. I wrote the various design decisions down and why we ended up making the ultimate in the CISC architectures, that is it had everything in it. It had a PDP-11. It had a COBOL instruction set so that you basically could do virtually all of the things that COBOL wanted kind of in one instruction -- it turned out to be a micro instruction or a subroutine, queuing so that you could add to queues. It had stacks and things that were remnant of the PDP-11. So that was the basis. We met for about three months, worked through all of the issues e.g. pages sizes and what was going to be in the architecture, and every couple of weeks, we allowed I think one or two people from marketing to come and talk to us and then we sent them out for questions. I don't remember any of the questions that we had for them. I don't think there were very many. And then we had one customer input and that was Ken Thompson from Bell Labs, and we talked about what VAX was going to be and is that going to be fine for UNIX. So that was the process. Meanwhile, Steve Rothman had communicated with most likely Bob Stewart who worked for Bill Demmer, and they actually implemented a very general machine that is one of the prototypes sort of hanging around. By the time the architecture was done I think we had a working machine in September. It was a microcoded machine -- it was basically a programming task and the people could then start thinking about using it or whatever for benchmarks or testing codes.

Hendrie: They had built a general purpose machine, microcode sort of machine that could act as a simulator, knowing that it might serve this purpose or was it for some completely other reason?

Bell: I don't remember exactly how or why we had ended up with it, but we had this machine basically that was done. We had a bunch of technologies around that had been done for PDP-11's, and so a lot of those were put on steroids to do the 11/780. There was a bus that ended up being the bus that we used for the 780 and it was a very general, or well, it had been certainly modified and extended and everything when we finally built it. It was the first machine we ever had where we basically loaded the microcode. IBM had been doing that for mainframes but we had a floppy I think in the base of it to load the microcode for us. Then we proceeded on to implementation and I remember going out and, quote, selling a few. The first couple, I had made a presentation at Carnegie. A professor by the name of [John A.] Pople who ultimately won the Nobel Prize for computational chemistry -- he was sort of the inventor of it -- was going to get one of the first ones and he actually ended up getting number one or number two. A couple of the early machines went there and I think that was one of my sales calls when it was introduced in late'77. I

don't remember when they actually delivered it -- I think it was later, in early '78. And then I went to Japan and gave talks about it at the various universities and then to Australia in the summer of '78, and by now it was clear that the machine was going to be successful. And by the way, in the process one of the other things that I did was to write the ads for it since I had always complained bitterly about advertising. What we ended up doing was a very simple idea. It was a testimonial ad program and it got to be quite a thing to get to be asked to give a testimony -- users loved to be selected for our "ask any user" ads.

Hendrie: Pushing you to buy a machine to do it. Right?

Bell: Yeah, the campaign was called "Ask Any User" and the users basically featured various aspects of the machine. It had new arithmetic and it was a nice virtual memory and a good sized memory and so at that point in time it was a very hot machine and it sort of became the standard for computation. In fact, I still to some extent live with the legacy as all the supercomputing guys always accused me of killing supercomputing, of being the one person that really fouled up and destroyed supercomputing because then everybody had their own computers and now they didn't need the centralized Crays and things like the large comp centers.

Hendrie: Those were relegated to only the most difficult, largest jobs.

Bell: Anyway, so that was really kind of a byproduct of all of the VAX, and when it was clear it was going to be a success and it was worth betting everything, I had gone to this trip to the Far East including Japan which was my first trip to Japan. Actually, that was in the summer of '78. I came back from that and spent 3 weeks diving in Tahiti. So we had enough experience with it, the machine was going well and so basically at that point I proposed and wrote the VAX Strategy which essentially bet it all on VAX and that we would build a number of machines and that we would stop PDP-11 stuff. I didn't say anything about the PDP-10 at that point. We stopped building PDP-10's later and everybody accused me of killing the PDP-10, but in fact in reality what happened is no one there was able to build PDP-10's. Alan Kotok had stopped working on PDP-10's and no one knew how to build a PDP-10, so we had had an engineering project and we had to abandon it because the guy doing it had built 360's and when he found out the PDP-10 was nothing like a 360 then all the pipelining techniques wouldn't work at all and so basically we stopped the project.

Hendrie: What was that machine called internally?

Bell: It was called Jupiter.

Hendrie: Who was the project engineer?

Bell: Bill McBride was head of the whole project, and I don't remember the engineer that was doing it but he was hired because he had done such a good job on the 360. But the team hadn't really built a 10, so basically all that learning had gone.

Hendrie: And it didn't help, required a new set of rules and . . .

Bell: Right. All of the idiosyncrasies and what you had to do because the 10 was great in that everybody loved the byte pointer in there but that means virtually all the instructions are sequential. It had a huge

number of bytes pointed at things and that meant that fundamentally every instruction required an indirect address and so we had to do something to that which was not built into the hardware.

As an afterthought, 08/08/24, a company called Systems Concepts in San Francisco had built a great PDP-6 using TTL. Shottky and I tried to get it licensed in late 1982. Stu Nelson was the designer and Mike Levitt the president. But when I had my heart attack, communication stopped with them because I was the only one they trusted. The deal would have been worth 10 million or so to them with absolutely no work except to sign the license. According to Wikipedia their name was changed to the SC Group and they sold a few computers.

Hendrie: I'd like to go back a little bit and maybe talk about some of the issues that came up from the time in '75 when you sort of got this project rolling that you had to deal with.

Bell: Right. Well, I think the big thing was getting to the point in '75 because there was a long period of gnashing of many teeth, of do we extend the PDP-6 or the PDP-10 or do we extend the PDP-11? And by now the PDP-11 was going like gangbusters, the 11/70 was out there, the rule was 8-bit and the 10 was doing well too. It was clear we could build smaller machines and that in a sense it didn't make very much difference. But on the other hand we had so much tied up in the inertia, velocity, momentum. Everything was going to the . . .

Hendrie: Of the 16-bit line.

Bell: Of the 16-bit line.

Hendrie: And 36-bits, was it going to look good on your image even though it did the job.

Bell: Right, and so that was part of the big decision and that was done prior to deciding that we were going to go that way when we started up in April '75. In some sense, I don't think there was much of a decision about the approach we took. In retrospect several years ago Cutler asked me, he said, "Well, wait a minute, why did we go that way versus a RISC?" And I said it was unfortunate because at that time it turned out all the technology was just at the wrong point on the curve. We'd write fast, we'd read write memory and couldn't build a large cache that was the basis of what RISC needed, would have been predicated on so there was a big thing about that. The microcode memory worked very well. There was an issue of compatibility that came up and so those were a lot of the concerns about the decision to go the way we went. Then once we made that decision there was the VAX A Bluebook which has the instruction set in it. In fact, I was rereading it the other day because I gave a talk about the history of minicomputers through DEC to a group at Berkeley, UW, and UCSD as a course. I reread that part in the VAX A Bluebook to decide how we had made those decisions. But at that point once you went down the microprogramming path, you basically put everything in, anything that was going to be used in a sense more than once you sort of put in. We had a polynomial instruction in the architecture too.

Hendrie: I didn't realize that.

Bell: Right. No big deal. You just put the tables there. We had a very elaborate "call" instruction and I don't remember what the ultimate impact of the effect of a call was. We didn't use it, we weren't using it

that much or it was too specific, but certainly we had everything that would. . . You sort of put everything on the stack and called the procedure.

Hendrie: Was there any provision for the customers doing any of the microprogramming?

Bell: No, we didn't want to go down that path at all, absolutely. We wanted to stay away from user microprogramming because of the difficulty of maintenance and . . .

Hendrie: And support.

Bell: Support and compatibility going forward. The biggest thing we had was insistence that we were going to maintain compatibility because there were many different models for the PDP-11 but they were all fundamentally incompatible, and it was I'd say a tribute to programming that in fact we could run across all of these machines. There was an article that Ron Brender wrote in the software group called "Turning Cousins into Brothers and Sisters" or something like that, but basically all the early floating point or the arithmetic had different options and there were a whole bunch of different options and different levels of memory management and it evolved over time. Everybody was sure "we can't stand doing this again".

Hendrie: You can't do that . . .

Bell: Never can do that again. So at that point we really worked on compatibility. That was a key part of it, that the models had to be compatible as we went forward.

Hendrie: What was next after the original VAX? What was the next machine that was started?

Bell: I'd say . . .

Hendrie: Taking the next step . . .

Bell: The next step was this big event in the summer of '78 when I was on a diving sailboat for three weeks and it was really kind of an "ah ha" and saying look, we've got to build a range of machines but unlike the 360. And so I got that idea from the 360. I had written about that at Carnegie when I wrote "Computer Structures" with Allen Newell. There were two things that were different. In the case of the 360, all the 360's were fundamentally mainframes. They were all operated the same way. They had different operating systems. In the case of VAX, they all had the same operating system, they had the same environment, and the thing that differentiated them was in fact where they lived in the hierarchy, that is, we had predicated clusters so our mainframe was in fact what has now become the way you build mainframes, scalable clusters. So we pioneered clusters by aggregating 780's and then bigger machines together saying we would never be able to afford building the big machines that IBM was building and the way to solve that problem was simply the cluster things or MPs or cluster, whatever made sense at that level as opposed to continuing to push as much into technology. Now that lesson got forgotten later on and that was one of the bad parts of what happened to DEC. The 9000 was clearly one reason for failure because pouring $3-6 billion into a product that didn't make it was quite a bit of money. I don't know how much it cost, but it was . . .

Hendrie: Very costly.

Bell: Very costly and could have been something else and it also established a kind of psychology of "Gee, we're going to use that machine to go off and head into IBM", which was crazy. So anyway there was the hierarchy which was at the top, it was a cluster, and then the sort of distributed many kind of thing that we all knew and loved, and then at the lowest level then we said everything's going to be work stations and PCs and that those things are going to be the desktop and those things had to all run compatibly and so you could choose anywhere you wanted to compute and that was the basic idea. When I came back from the summer trip to Tahiti, I started that crusade to go install that and that meant getting rid of a bunch of 11's and other projects and focusing on VAX and then starting the VLSI programs to get MicroVAX.

Hendrie: You had some concept of how you could execute a really small machine.

Bell: Yeah, and those things all took off. Initially, there was a huge amount of uproar about that. We stewed over that.

Hendrie: You mean in terms of killing off the 11?

Bell: Yeah, and deciding that we were going to bet it all. Initially Ken was fairly non decisive and aloof and observant. One day he'd ridicule it, then he would listen and then ridicule, and eventually I think he actually bought in to it. And the culmination of that was I went to the board, I think in November or December of 78 or early 79, and I described what we were going to do and that became known as the VAX strategy. It was a diagram in the shape of an 'E'. At the top part of the 'E' was a cluster with a cluster interconnect and then all of these levels and things were connected together by an 'I' which was a network interconnect. And fortunately in the '78 time frame we were able to make the connection with Xerox, and Bob Metcalfe for the Ethernet.

Hendrie: Tell us a little bit about that story. What happened there?

Bell: Well, first off, the strategy was in a sense predicated on an 'I for Interconnect,' of having some way of doing that, and we had two or three different ideas about connecting everything using a local area network.

Hendrie: You had some idea of a way to do it but you hadn't thought of . . .

Bell: And at that time our friend Apollo wasn't started so we had two or three different network interconnect structures there that we were looking at. We had some rings and we had other stuff. We didn't know a lot about the Ethernet at that point although that had come out I guess in '76 or so and that's what PARC was using to connect their workstations. What happened was Bob Metcalfe joined us as a consultant. He had left PARC and said, "Gee, I'm here, what I want to do is propose that there be a standard here and I want to try to broker a deal here with Intel." Or I mean with Xerox. And so that was kind of the beginning of the DEC, Intel, Xerox cabal that in fact ended up proposing Ethernet. And I can probably find it, but I don't recall exactly the first meeting of that but it was with Phil Kaufman, who was later the president of . . .

Hendrie: Of Silicon Compilers.

Bell: Silicon Compilers. He was at Intel and he and I got together. We met using the ATT's PMS picture phone meeting service in San Francisco and Boston so we got together virtually. I think we signed up for two hours of time at a couple hundred bucks an hour and we had a meeting and basically agreed to be partners in terms of developing the Ethernet standard. And I think it happened kind of concurrently that we had approached Xerox for the patent and to be part of that and they were the group brought in and David Liddle headed it. I recall the meeting and it was in the Parker Street building and our guys had been meeting with their guys. I don't remember exactly who they were but I remember coming to the meeting and at the end of the meeting I said, "Okay, where are we?" And they said, "Well, we've got to get your management to do something and then we got to get our. . ." And I said, "Okay, let's make it very simple." Here's our management asking for this patent so I said, "What do you want the letter to say?" And I sat down, wrote the letter on a random word processor right there, signed it, and they walked away with the letter of understanding. They went to the . . .

Hendrie: Went to the Xerox managers . . .

Bell: . . . and they said okay, these guys want it and so it was a patent.

Hendrie: You wanted to take out a royalty free license . . .

Bell: They basically agreed that this be a standard and that was the beginning of Ethernet.

Hendrie: Had Metcalfe left?

Bell: Yeah, he had left.

Hendrie: He had started 3Com?

Bell: No, he hadn't started 3Com yet. That actually happened after I believe. . . He might have started it, but he was there as kind of a free agent advising us. He was a consultant so we actually I think paid him for consulting on what a network would look like. Anyway, that all got put together, and the culmination of that was the 1982 -- I think February '82 -- announcement and that was Liddle and I and Bob Noyce announcing the Ethernet at the World Trade Center in New York and we talked about this being the standard, this is how we see computers being connected together. And then I remember one of the people from Business Week was there saying, "Well, Wang's got their Wangnet and why aren't people going to use that network because it uses the existing cable TV?" And I said, "Well, first off there's not much cable TV in corporations, and even though people know how to install cable, we don't really see these things as being compatible at the time or necessary." And then I said, "Oh, and besides, the biggest problem with cable TV is this is sort of a direct coupled coding as opposed to analog coding, and even though you've got multiple channels on cable and people were starting to fool with cable TV for networking, the biggest problem is think of it like your house or your city distribution. You've got water, sewage and gas coming into your house and in principle you can put them all on one pipe but sorting them out is the difficult thing." Anyway, so I sort of blew her off. But anyway, that was the announcement and then we made the announcement in London and we also went to Holland.

Hendrie: What were the responsibilitiesof the partners in this?

Bell: Initially, the way the thing was structured there was a working group of I think eight that were designing it, and in fact one of the things that was contentious was I'd get this call from Phil Kaufman periodically and he'd say Olivetti wants to be a partner and if you don't want them to be a partner they won't play. And I said look, we've got eight people, they know more about networking and I don't know that Olivetti knows anything about networking and no. That was basically a nice way of telling them to pound sand. There'll be plenty of time to comment on it after we get it proposed. I said I don't even go to the meetings. And then one Saturday night I remember I had gotten a call from Rob Wilmot who was the chairman and the CEO at ICL so he was the head of the U.K. company, and he said, "Hey, this is exactly what's needed, we really need it, we believe that the industry needs it. I will bring you all the Europe ECMA. I'll bring all the European computer manufacturers into it." And I said great. And he said I want to take a look at it and I said come on over. And so he brought over a bunch of his designers and we ended up at a meeting at dinner. The final dinner was at Colonial Inn in Concord, and so that brought them in. And then ultimately after it was announced and the products were out there, IBM sort of came out with their wiring plan for rings, but their LAN was years after. Basically, they didn't have a product at all. The token ring was many years after Ethernet was announced and what they did announce was their usual way -- that IBM will take care of you here, just install these wires, and it was sort of this bundle of wires and it had a couple of cables, a fiber and this and that and everything.

Hendrie: It was a hodgepodge trying to appear to be part of this.

Bell: Yeah, and so anyway install now and pay later or play later.

Hendrie: You said you didn't go to the meetings. Who were the technical representatives from DEC?

Bell: Dave Rogers I think was the technical lead, and I don't remember who all of them were but they were the best people that the world had at that point in local networking because we had the Xerox guys who had had experience with their earlier net, the 3-megabit net, and so we had that working experience and then we had some good circuits people. It might have been actually Tony Lauk who was one of the DECnet architects, so we had that kind of architectural experience.

Hendrie: And the Intel people were going to implement . . .

Bell: Yeah, they came in. I think maybe there were a couple people from Intel and it wasn't fully balanced on three, three, three or something like that. But in fact they produced the Ethernet Bluebook basically that then was commented on. And I think HP had probably made the biggest contribution having to do with some grounding and stuff that made it work better. But anyway, that was the origin of the Ethernet. My original slides for the product introduction at the World Trade Center are on my website. The slides had a couple of great quotes: "Ethernet is the UART of the 80s" and "the network becomes the computer" that SUN adopted. Several folks later told me that they thought I was crazy making the claim, but now after the fact they acknowledge that I was totally correct.

Hendrie: Going back to the original VAX development, were there any things that came up in the development that proved to be particularly difficult hurdles to overcome or represented challenges that you hadn't anticipated were going to occur during building this quite large machine?

Bell: When we first built the 780 I would say the answer is absolutely not. We'd kind of done a lot of it before. It was the first really serious microprogrammed machine that we had built. It was bigger but I don't believe there were any real crisis. There probably were some in terms of, oh, we don't have any space to put in these instructions. I recall Richie Lary worked on it from time to time to get things to fit so there might have been a crisis there but that was probably it. We had simulated and understood the whole architect, the cache structure that we had, and it was nice -- it really worked well. The decision to use the PDP-11 and have that be our build in an RSX-11 compatible machine was absolutely great because in fact it allowed us to run a huge amount of the software. All of the software was running on day one on PDP-11's or the compilers and things like that when we announced so we could come in with the software running and that was a big deal. So in a way it all went very smoothly. The crisis really on VAX occurred when we implemented the ECL machine, or the using of ECL gate arrays machines. We had two projects, the 750 and then the 8600 which is Venus, and in the case of the 8600 the project just totally got out of control and was an absolute mess that cost DEC dearly in many ways.

Hendrie: This was the ECL project?

Bell: This was the ECL project and in a sense it came out in, oh, I don't know, '85 or so. Vinod Khosla, the founder of Sun, says that if it hadn't been late Sun would never have been able to get started. What that did was because everybody was waiting for a faster 780. They couldn't get it and so Sun could now take all that money that was allocated for 780's, the 8600s, and buy workstations with them and that's what they did. And so that was kind of a byproduct of a project getting really screwed up. In fact, the 8600 was just an unmitigated disaster in terms of a project and I went out there to Marlboro and was out there for six months working on it.

Hendrie: It was done out in . . .

Bell: In Marlboro, and in fact it was under the same guy managing the ECL PDP-10. Bill McBride wasn't the project engineer but he had another counterpart who was leading the 10. Ulf Fagerquist of course was the boss for both of them. Then there was George Hoff who was the engineering manager for the 8600 or for Venus. And fundamentally it was kind of a classic story of what can happen to an organization. When I went out there it was sort of clear what happened because I sat in on one of Ulf's staff meetings and there were all these people around who had nothing to do with the project. Here's George and I think McBride and then everybody else was sort of the finance guy, the human resources guy, and this guy, the facilities, all the 10 or 11 people in a room, and so Ulf was spending his time on all of this crap and had no knowledge of the projects. From this experience I made a rule: "If I ever know more about any project than anybody in the management chain then let's get rid of the intermediate managers." This was a pretty hard thing to test because I generally knew more about every project in the company than somewhere in the chain. So you'd go up and down the chain and occasionally people would know something about it but in general the managers didn't know what was going on.

Hendrie: At what stage in this project did it start to unravel? Was this in the paper design stage?

Bell: Yes, it was in paper design. It was a little bit past the paper design. The team hadn't quite finished all of the design. It was manifesting itself as, "Gee, we've got to take a risk on getting Motorola, buying chips from Motorola for the gate array, so we've got a bunch of essentially manufacturing gate slots that are lined up and now we're having trouble and do we take a risk on those things building the chips?"

So that was kind of one of the warning signs that popped up, and the design isn't quite done so there were a bunch of things like that. We had a huge design review -- I know it was a Friday the 13th, I don't remember which Friday. All I remember was that it was that bad, it was in a very large room, a cafeteria, at Marlboro and everybody was sitting there, and Bill Demmer was there and he was from Littleton and George Hoff talked about the design. George said, "Well, we've got these problems and we're not quite done with this or that" and was asking whether we should get these chips fabricated or not, and then lots of things like, "I'm not sure about what Motorola is doing and I don't trust those guys" and all this stuff. And so there were all these bad signs from Motorola and so it's their problem. "No, I don't know what the gates are going to look like or the power", some damn thing like that, and so it was the end, 4 or 5 o'clock, and I said does anybody have a feeling about this? Well, this is going to be tight to make the schedule. All the people in the room thought it was going to be tight, but doable. And I said, "This is the worst project I have ever seen, this is a disaster, a f***ing disaster, this machine will never in a million years work, there is no way that this thing will ever work, you have no verification of it, you're not simulating it at all, there's no way, you're not even doing any code reviews, there's no good debugging philosophy to it where you can do that, you've got some of these little things, you're not even done with the thing!" And then I went in on Monday morning to the operations committee and I said, "I got news and this is really, really bad." And Ken said, "Well, you'd better leave right now." and I said, "I think you're right." And so basically Mary Jane and I went to Marlboro and spent six months and started working on the project. We had to start a simulator from scratch. We had some simulation efforts but we couldn't buy one. We had that. There were a whole bunch of things that were started to get that. Then I remember Ulf and I went to Motorola and looked at where they were and I came back and I said it's not Motorola's problem, they're going to be just fine, don't worry about them. Shortly thereafter, I said to Alan Kotok, "Alan, you've got to come in here and look at the project".

Hendrie: What was Alan doing?

Bell: Oh, Alan was doing telephones probably. He loved to work on the phone system or whatever.

Hendrie: He'd gone off to do something more interesting. He was tired of doing 10's. You weren't going to build any more 10's so . . .

Bell: Yeah, we weren't doing any 10's and so he was having fun with a phone system or some damn thing like that, and it was always something to consult on and so I enlisted him to come back and said you've got to play the referee architect. Venus was one of those projects that in the first week individuals or groups were telling me they're really the architect driving this whole design thing, and then the second week everybody was telling me no, they had nothing to do with this.

Hendrie: That is funny.

Bell: Yeah, it was pathetic. In fact I called it the Soul of the New Machine 2, or Soul 2. I think it was called after the book about DG. Not even the DG team did this as poorly. It was of course late and was really a financial disaster. I can't remember the number, but DEC was losing 2 or 3 million a day on not having it, so you knew what the effect was going to be. And I said if we had had it we would be delivering this amount.

Hendrie: How late did it end up being by the time it got out?

Bell: I think it was probably three years late. It was a factor of 2, yeah, at least a factor of 2 schedule slip.

Hendrie: Huge amounts of . . .

Bell: Yeah, I said we're not going to build, we cannot build this machine until you can show that it's going to work and because you've got these gate arrays and they're not probe-able or fixable.

Hendrie: You can't stick a probe in them.

Bell: Yeah. Well, they had some fake gate arrays that they could probe and look and see what was happening but you couldn't build a whole machine out of them. Or if you did, you had to plan that's how you're going to build it and we didn't have a plan that said we're going to build an emulating kind of machine and that wasn't it or whatever. And so there was that, and I'd say that was probably a big effect on DEC, and then in a way to survive we developed a simulator.

Hendrie: It did eventually work, right?

Bell: Yeah, it eventually worked. DEC sold a lot of them and it was successful. Meanwhile, I had gone to Littleton and got them to start the 8800. I sort of pissed them off because I went out there and they were farting around with 8800 with a TI technology called I2L [Integrated Injection Logic] that never hit the market. I had them switch to ECL, really as a backup to the 8600. In retrospect, they might have better jumped to CMOS. That way we would have converted some great designers earlier and changed their careers.

Hendrie: What were they theoretically doing?

Bell: This was the 780 team so they knew how to build computers. The other guys, well, really they probably didn't even know how to build a computer but they were building something out of I2L [Integrated Injection Logic]. So I was watching all that and this is TI and they had a similar kind of thing and this I2L is getting slower all the time, it's not clear TI's going to make this or not, whether they're going to manufacture it or not, and they were off working on that. I said, "Forget it, you guys, don't build that, this is not going to work, it's just too hard, you're not going to get the speed with the TTL clock that you have to get to have a 780 replacement." And so they switched. The two came out just about the same time as the 8600, maybe a little bit later, but that was kind of a backup because I wasn't sure that the 8600 was ever going to work. In fact, I don't know at the time whether I would have said we don't need both of them, but the 8800 was a superior machine. It was a little bit slower but it was a dual processor and so it was just a nice, clean machine and I trusted the guys who had done it before. And then meanwhile over in Hudson people were saying, "Yeah, we're moving along, we'll eventually get there with CMOS." But they were years away, three or four years away from getting enough high speed, so we were stuck for awhile. But anyway those were some of the things.

Hendrie: Some of the things that proved to be . . .

Bell: Right, it was that transition. In a funny way, making the transition to VLSI, going through gate arrays. because Dave Cain had made the gate array transition -- and you probably knew Dave as the founder of a company that used the 68K to make a lab computer. He had made the 750, a beautiful

machine, a small machine almost as fast as a 780 but very reliable and it was a gate array machine, and then there was a 730 which is an easy machine but a lower speed and very small.

Hendrie: Talk about the issues with building the MicroVAX portion of the line.

Bell: Well, that was . . .

Hendrie: You had one start that didn't go quite as far as you thought.

Bell: There were a couple of machines. One was the first MicroVAX. It was actually done in Seattle. Dave Cutler did that. Dave of course is a whole story, many tapes in himself, but don't think I said anything at all about Dave Cutler.

Hendrie: No. You didn't say anything at all about Cutler. I was going to roll back and get the software story in a little bit.

Bell: Right. Well, certainly Dave had finished the operating system and drove that and basically wrote a large part of it and had a small team that did that. And then after that he went over and wanted to do a PL/1 compiler because he said, "Gee, I've never done a compiler."

Hendrie: I was going to say . . .

Bell: But I needed a compiler. He was in New Hampshire at the time and so he had half a dozen guys and they built the PL/I compiler because they wanted to start with something really of a challenge. He had a very nice group, a very tiny, little group. The part that was kind of interesting about it is that he had architected it such that the backend used a common run time and so ultimately that became the backend for all of DEC's software, so it was languages.

Hendrie: All of DEC's compilers.

Bell: All of the compilers, and so he worked on the run time and the backend and that was for Pascal and I think C and PL/I. I think Fortran ultimately moved over to it but that was another machine, another language that it had initially had. So Dave had worked on that and then that came out, and I don't remember when it came out. It was probably '80, '81, something like that. He then basically came to me in 81 I believe and said, "Look, I've had it, this is a totally bureaucratic place, can't get anything done, all this bullshit, and so I'm leaving." I asked what he was doing. And he said, "Well, there's a guy at Livermore that wants me to build a language and we're going to build some compilers," or something like that. And I said, "Okay, well, look. We'll fund you, we'll invest in your company. Why don't we set it up and we'll meet the guy." Whoever it was was a typical Livermore physicist who knows everything. So they came out and we had dinner someplace in New Hampshire -- Dave and I and probably Larry. I don't know whether Portner was part of that. There were probably four or five of us at dinner. This guy had such a tremendous ego and he sort of said when we asked what his company was going to be doing, "Well, my programmer here will do this," and so he just kind of sloughed Dave off as oh yeah, he's . . .

Hendrie: As his programmer.

Bell: His programmer that's going to implement his vision of this great piece of code. So after this Dave said, "Well, I've decided I'm not going to do that." I said okay, that's all right, I can understand your frustration because we had just moved the disk guys to Colorado and I said I think it's not good to not have everybody in Maynard anyway. So take anybody you want, go anywhere you want and do anything you want to do. And so he and Roger Heinen and a couple of other guys that he had worked with went around the country looking at sites. They looked at the valley and ended up going to Seattle, and then he said, "Okay, I'm going to build a MicroVAX. I've always wanted to do some hardware, I'll build that." Because the MicroVAX wasn't going to be ready in time because it was a VLSI chip and the chip had a long gestation time and it wasn't ready. And so he said he'll take a more conservative route, we'll do whatever, and it turned out that he ended up buying the first chip from Silicon Compilers which was a custom data path, a VAX data path. So it was a 32-bits, in essence a fancy ALU, but in fact it had the registers and everything else with it and the control and so you just add the microprogramming and you've got a VAX. So Dave pulled some people that could manufacture stuff and that team actually ended up putting it into manufacturing and I don't know how many were sold.

Hendrie: In Seattle.

Bell: Yeah, in Seattle and that was the first MicroVAX. It was small. It was a workstation sized machine and he also said, "Gee, by the way, I'll throw in another operating system just for free." So basically he put a run time, a real time VAXElan, which had a lot of stuff in it too. It was probably the only operating system that was written totally in a provable code. It was kind of a very fancy typed check Pascal that was used for real time control. You downloaded the VMS images to it and so it was totally compatible with everything, with all the other machines.

Hendrie: He was going back to his roots to a real time operating system.

Bell: Yeah, and so it was probably the world's best real time operating system bar none. So anyway, that was Dave.

Hendrie: He kept wanting to do something he hadn't done before.

Bell: Yeah, absolutely. And the irony is he's still doing it. He of course did NT and that was the basis of all Microsoft operating systems. He was the guy who actually did the 64-bit NT. And Microsoft getting that out, which is no menial feat. Now he is working on a new project that I think will, quote, save us again. And there's never anyone who's been as productive or prolific and really just produces this sort of beautiful work as Dave. There's just nobody like him. We've got all these guys that have done really one nice thing but no one who just continues to put out really solid, beautiful software.

Hendrie: Elegant products.

Bell: Yeah, and so he's doing it again which I'm delighted about.

Hendrie: That'll be very interesting.

Bell: Yeah, and it is. I'm excited about it.

Hendrie: We did the other day Dave Cutler's story. Talk a little bit about if there were any issues in the software development of the original VAX. We talked about the hardware and some of those things and yet it was either the compilers or the operating environment.

Bell: I think it went very smoothly. It wasn't your kind of typical thing where you're starting something new. The labs [Bell Labs] had gotten an early VAX and I think that they had ported a UNIX System 3 or 5, whatever, AT&T. Then we had a bunch of VAXs that went to Berkeley that they built the Berkeley Unix on. Coming back I think about these folks that keep reappearing. Dit Morse who is the only guy I'd ever fired when he was heading the PDP-6 timesharing system in 1963.

Hendrie: You told us quite a few tapes ago about that.

Bell: Yeah. Anyway, Larry Portner came to me in the late 70s and said "Hey, I'm going to hire Dit." and I said "Dit, I haven't seen him in years, since 1965 or so!" And he said, "Yeah, we need somebody who's really been working on file systems." And so he basically came in and did consulting on the file system. I think ultimately Larry also fired him, but he did spec the file system. There were others, and it was like these different people, Bill McBride. Whether these things get seen or not, it doesn't matter but it's kind of interesting that there's a little management philosophy that shows incredible tolerance. So he in turn really screwed up the PDP-10, and then I think Ulf got rid of Bill McBride because the PDP-10 was equally screwed up. And Bill Demmer came to me and he said "Hey, I need a project engineer and do you mind if I hire Bill McBride?" And I said, "Well, I wanted him fired but if you're going to do it and hire him, that's your problem."

Hendrie: It's on your watch.

Bell: And then three months later well, Bill left or he fired him. I remember Gordon Moore telling me how they had interacted with people at IBM. He said the problem with IBM was they're just absolutely brutal in the way they manage. He said at Intel if we have a guy and he screws up then we give him another job and about the third time we say well, maybe it's the guy. But at IBM the guy screws up he's out of there, no second chance. So we give him about three chances. And so I always remember the story of that -- oh, yeah, we don't want to fire the guy, we want to give him another chance. But anyway so Dit had done the file system design and then was let go.

And people come to me from time to me and they'd say, "Hey, there's this great language on VMS." So many people! Anyone who has used VMS just says, "Oh, god, why isn't that out there?" But it was the DCL, the DEC command language. I was never a fan of command language so I didn't even know. I knew the 360 one from CMU and it was the world's ugliest language, just ugly. But Dave had a beautiful command language that would drive everything and that's what a lot of people wrote applications in, in this command language, and it was a very nice thing that he had done.

Hendrie: Can we go back to the MicroVAX and take the next step in the MicroVAX story?

Bell: Well, what was going on in MicroVAX was it was being designed in Hudson. DEC was on a treadmill to continual movement to I'd say master VLSI. I don't recall that there were really any disasters per se, but it was much more of just the sheer size of the jobs at the time and the state of the tools. This

team is probably the most talented in the industry and those guys are spread throughout the industry at AMD (Rich Witek), Intel that took over, and companies they started like Dan Dobberpuhl.

Hendrie: It was just a horrendous amount of work.

Bell: Right. And I think Carol Peters was one of the people that had worked a lot on MicroVAX. The final MicroVAX II, which is actually the one- or two-chip MicroVAX, was subsequent. It came out after MicroVAX I, the Cutler machine using the Silicon Compiler.

Hendrie: At what level was the technology of this? Was this N-Channel Silicon gate or was this CMOS?

Bell: At that point it was CMOS.

Hendrie: It had gotten to the point where you were doing CMOS.

Bell: Yeah. I am not sure if there were any intermediate N channel stuff or not.

Hendrie: So it was a dedicated CMOS.

Bell: Yeah.

Hendrie: Full custom CMOS chips.

Bell: Yeah, absolutely. I had come back to DEC in 1972 to get into that whole business, and then we started out with the LSI-11 which was a deal that we had made with Western Digital and then gradually acquired the technology and the learning. Carver [Mead] had helped a lot on all of that and we got the people together to just understand it. Alpha was probably one of the best, hottest machines around. Also they developed the StrongArm for embedded computing that was ultimately sold to Intel. The basic problem was that the company spent like a semiconductor company but didn't have the sales and marketing to go along with the many great products it produced. Some of the people that were involved in that, and then also I guess they did the StrongArm chip, and so a lot of the people that are key in CMOS really came out of the Hudson environment. Dobberpuhl, for example, and the rest of them, so I think it was a super group. I'd say probably the only bad thing that happened -- and I can't really reconstruct the thoughts and the history of it -- was undoubtedly too much do it yourself. We can't really afford to be investing in those very large plants for our own stuff. So not making the right deals in manufacturing and capital equipment and all of those kinds of decisions. They just didn't feel right certainly.

Hendrie: Did you get involved at all in any of the early DEC workstation/PC ventures?

Bell: Yeah, those at that time were disasters in terms of looking at would I do anything different or not, and I think given the way that it all played out, probably the big error there happened around the time I left in '83. In a sense the whole story was written about that time and that was the standard had been established and that was all right when I left. I'd say the whole book on the history of PCs had been opened. DEC introduced three separate personal computers at a time when the world had clearly

standardized on the IBM PC standard. This was really the root of the problem. The PDP-11 might have been okay, but it was held as proprietary.

Hendrie: If you were a good prognosticator, you could see what was going to come.

Bell: Yeah, and you didn't have to be very good either because everybody was out there. So I'm not going to fault us for doing the three because there was, for example, the DECmate, and that was a PDP-8 and that was a fairly simple thing because of the tiny, little engineering group. They were very productive. They had customers, they had legacy, they had software and so we weren't investing a lot. It was used both for small business and word processing.

Hendrie: A company project.

Bell: And in fact I tend to look at it as the DECmate was a wonderful typewriter and it competed with Wang and so I tend to think of DECmate as just our typewriter or our terminal or whatever. It was very nice and so I wouldn't have changed that. The PDP-11 was our bet on what was going to be the work station PC, and in fact as it came out it was absolutely superior to the PC, there's no question about that, but the whole issue of standards really overwrote everything else. Well, standards and then also the other thing was, geez, we didn't even allow people to connect to it. So DEC, who had been the sort of champion of having I/O and buses because they had a very nice, little bus structure -- I don't know whether it was hot pluggable or not -- and we had both RT-11 and a version of RSX running on it. But it didn't have the panache of what IBM had. The 11, called the PRO, was held proprietary so there was no way to compete with an open standard like the PC. The Intel-Microsoft PC was the Rainbow.

Hendrie: And all of the other vendors.

Bell: Right. I don't know how long it was before it was very clear that that's the way the industry's going to structure, but it wasn't very long after the PC. And DEC never got it. It took years. I know they didn't get it in by '86 or so. I was at NSF and Ken sent a PC down for me and it had everything but software and I said, "Where's the software?" He said, "Well, we got to have special software because it doesn't run IBM software, doesn't run Microsoft software." So they never got it or they didn't get it at that time. They ultimately got it but it was too late. DEC was not used to following standards it didn't create.

Hendrie: Far too late.

Bell: Yeah, and again it's one of those things. DEC could have been a player in the PC like HP or anybody else that was starting up because it was a volume kind of issue and DEC was okay in manufacturing. DEC could get itself together to do volume when it had to and it would have been a great challenge, but when they finally did do it they were always behind. And so I'd say certainly the PDP-11 was a waste and it didn't have the right stuff, but in this case it was really the PC coming in as proprietary when the world wanted standard.

Hendrie: You were just talking about Barry James Folsum, Mary Jane Forbes, and Don Folsum.

Bell: Right. So we had the three machines, the PDP-11 was the Pro. During the project, it had been called a KO for knockout or Ken Olsen machine. And then Barry James had been the advocate of

bringing in initially, I think, an 8008 version, the predecessor to the PC. No, I guess it was the Z80 Zilog version based, the 8-bit machine and then making the transition to the 8088 Intel 16-bit architecture. So he was driving that standard.

Hendrie: Was this the Rainbow?

Bell: That was Rainbow. Yeah, that was the first machine. The compatibility issue I think, in terms of really not making that a fully compatible machine, was certainly the major and fatal flaw in terms of what DEC thought about. Because it was so clear once other people were manufacturing it. Certainly that that was standard and everybody was building to that standard. The paradox of that, by the way, is that everybody says, "Oh, if IBM hadn't made that a standard, then they would have had the whole market and blah, blah, blah." And I said, "You don't get it. The fact that it was a standard, otherwise it would have been like the UNIX market, the fragmented market." The guy who actually made it a standard was Wintel and that was the standard. I know at the time or shortly after the announcement, Bill sort of says, "Gee, we can't afford to make an operating system for you, one that's unique. You're going to have to be brought in. It's got to be compatible with what we're doing." That whole issue of platform standards and compatibility is what made the whole PC go, of course.

Hendrie: And that was so counter-cultural to a company that prided itself in creativity.

Bell: It was counter-culture and, in fact, almost so counter-culture to the industry. In spite of itself it ending up with a standard, to a certain extent with Linux or with Red Hat, so that there's a company there and the companies seem to be maintaining some compatibility. My own belief is that they will ultimately fragment, like in a way that's not unlike what happened to UNIX, but perhaps not as bad. But certainly to get the kind of volume you have to have a standard and that's what didn't happen before Linus.

Hendrie: Was that still going on when you were there? Let's get back to your career.

Bell: Yeah. I left after my heart attack in the summer of '83. I'd say at that point that summer or shortly thereafter it was pretty clear to me that the PC was the standard. Clearly DEC hadn't gotten the message to that effect in '86 or so, because when I was at NSF Ken sent me a fine DEC PC. In fact, it was nicer than a PC, but it had no software for it, so what the heck? It didn't matter. I think we have all these different things that were important things to the change in the course of the company -- in '83, not making that switch, certainly by '84 someone saying we're going to be the best PC manufacturer out there. If you're going to do that or you could have gotten out of it presumably, too. But DEC's aspiration was to certainly be kind of a full service company and so you basically have to be compatible at that point. You have to offer those things. And DEC had that knowledge, because it had owned the whole terminal market, printing and CRTs. We still see some of them. Those things are still around, the VT-100 and friends.

Hendrie: You mentioned your heart attack. Would it be appropriate to start off with what led up to that and what happened?

Bell: You never know what leads up to a heart attack. A cardiologist that I talked to within the last 10 years sort of said, "Ah, stress has nothing to do with it." But basically, my view is that your psyche and stress is a major contributor, as is your heredity and whether you smoke and stuff like that. Well I didn't

have those kinds of problems. But I had a double bypass in March of '83. I had some sort of angina attacks before that. I had been logging those. I had two or three events that were kind of blackout kinds of things before I had the actual one that was fatal.

Hendrie: You knew there was trouble somewhere in River City.

Bell: Yeah, and then a full-blown heart attack. We were skiing in Snowmass [Colorado], it was early in the morning and I basically passed out. I was with the DEC ski group and one of the wives, Kay Marcus, heard me fall or something and then Bob Puffer, who worked for me at the time, basically gave me CPR and saved me. I was then in a coma for about a week in Grand Junction and then I was flown directly to the Denver hospital when I was ready, and then the next day I had a double bypass. I took a stress test about two weeks ago. I had a minimal heart attack when I was ill. I was biking in France and I blew out one of the bypasses. So I have been living with one pipe there. I took a stress test a couple of weeks ago and my cardiologist just said, "Gee, I think maybe you ought to go in and do something." So actually, I got a stint a week ago. So now this one bypass is working better. I will probably want to see if we can get a little more blood flowing there eventually. They're willing to take another pass at it in another year or so. GB: In July 2007, I had another double bypass to restore the old flows.

Hendrie: That was an era when not everybody knew CPR. You were lucky.

Bell: Oh yeah. Bob was certainly very simply my savior. I owe my life to him. Bob was a volunteer fireman, among other things.

Hendrie: So he had learned these things?

Bell: Learned all of that and gave me CPR. When I was at DEC, Ken and I used to have very strong hard arguments. Certainly, those were quite stressful to us both. After that, I basically said, "Well, I've done about as much as I can do. The VAX strategy is in place. It's clear that all the products are in the pipeline. We could make semiconductors, we had the clusters, they were just being announced. I want a less stressful life." It was that kind of a thing. I think for anybody who can or has a heart attack probably really has got to think about some change in their behavior. I certainly changed my diet a bit, but not drastically. It wasn't a hard problem. Now I use chemicals to get it down. My cholesterol at the time was 220 or something, which was normal for that period. Now its 120, but with Lipitor and its friends.

Hendrie: So you decide this was really the impetus to . . .

Bell: To change. Yeah, because I wrote the note when I left DEC. I came back from the hospital and I sat in on some very contentious meeting about something and I thought, "You know, this is just too hard." To go ahead and make that decision that I would have gotten it made in a particular way, but it would have been a lot of work and a lot of angst. I think it actually had to do with workstations. It was a workstation kind of meeting. What's DEC going to do in workstations? It was one of those things of "You know, it's going to be a fight just to get it done and to stay on path." One of the things that happened after I left was much more diversion of let's appeal to the market. We've got to listen to the market. We need a workstation. Okay, well the MicroVAX hadn't been there yet. On the other hand, everybody was selling 68000 workstations. Now we're going to make 68000 UNIX workstations. Well, why are we going to do that? We've got the MicroVAX. It was right at that threshold. What happened is a bunch of different

machines got introduced, and it was at that point, I'd say the peak of DEC. So I'd say from '83 on, if you look at the revenue, it was virtually straight up until '88 when DEC had its cap. I don't remember what the market cap was. It was a very significant market cap. I guess it was number two when I actually left. But it maintained number two. At that point in time, at the peak, it really had introduced all of these platforms. So there were a ton of platforms. It had the MIPS as a UNIX workstation, but after Sun had kind of established itself as the UNIX workstation company. I think there might have been a couple of versions of UNIX for the MIPS and then there were a couple of versions for a PC. So DEC had a lot of different platforms. It had like seven platforms that had to be maintained as programming environments and support.

Hendrie: Almost like the PDP-11 days of all those programming environments.

Bell: Yeah. And in this case, you were buying the programming environment from an outside source. I think DEC had maintained a couple of them, but it then really just couldn't maintain those machines. That was a huge cost of DEC's demise.

Hendrie: That made it really very difficult and more and more costly to move forward without the commensurate volumes.

Bell: Yeah, right. You diffuse the market. You can't afford investments in all of them and also people don't know what to buy.

Hendrie: Exactly.

Bell: Because you have to make a choice and be ensured that there's going to be the software for it, which they couldn't say.

Hendrie: From your own personal point of view, you decided that you probably didn't want to stay around DEC and decided to leave. What did you do?

Bell: Henry Burkhardt came to me and said, "Hey, I'm starting this company with Ken Fisher. Do you want to join us?" Basically, I talked with the guys and it was an insane business plan. But my attitude was hey, I'm leaving, here's a train. I'll just get on this train and if I don't like the train, I'll get off. What the hell. No big deal. That's basically the attitude I had. I had joined the board of Silicon Compilers at the time, I think. So I said, "Well, there's that life. I can just kind of do that kind of thing as a free agent. But here's the start-up. Let me see what that's all about." So we formed Encore. It didn't take me very long to get absolutely educated about start-ups, having never been part of one since the founding of DEC.

Hendrie: What went on there? What did you do? Or was it really Henry's baby?

Bell: As I said, it was kind of an insane business plan.

Hendrie: Why was it insane? Talk about that.

Bell: The Encore story. I have a little bit of that in High Tech Ventures. John McNamara who wrote the book with me said, "We'll get sued. You've got to take it all out". So we took it all out. It didn't add that much to the book. But the plan was we were going to buy a computer company whose name I forget. I recall they were in Florida. I don't remember which one. It was kind of an old-in-the-tooth business and a 16-bit architecture, but they had a little bit of a niche. We were going to acquire them. In order for that to happen -- because you can't quite do it that way -- you needed another company in addition to what we had. So we had bought another software company that was a friend of Henry's called Formation, which was a DG software company.

Hendrie: I know about that.

Bell: In North Carolina, which we ultimately sold to Sperry.

Hendrie: Mr. Gilmore.

Bell: Gilmore, yes exactly. And it was a very nice product and it, in a sense, saved Sperry and made them viable for a little longer. They're still viable, but we met with the board member. The plan was, okay this team of three and a couple more. We had a CFO from Prime and Carl Wassman, who was a financial type.

I will always remember the day Encore "came out". Probably by far, Encore's "coming out" was the most embarrassing day of my whole life. It was pathetic. Here Ken Fisher was, this marketing guy, this sales genius that had made Prime and Prime had been straight up. Well Prime had gone up. We introduced VAX and Prime went straight down. We just totally cut them off at the knees. He was fired from Prime and so, of course, like any good sales guy, he's pissed. "How could they do this to me? I took the orders to make Prime, never mind that Apollo came from Prime. I'll get even." We formed the company and lo and behold, he then hired the whole damned top-level sales management team about six or seven sales district managers and VPs, all of that. He said, "Well, we've got to announce the company." We couldn't say, "Well, our plan is to buy this other company or to acquire this other company or whatever." So we did this in New York. This was at some damned financial channel -- it was put out on the the financial TV network. The most embarrassing day of my life was sitting up there with 12 salesmen -- Ken and Henry and 12 salesmen, a line of these guys who I had never met and talking about, "Oh yeah, we're starting this great new computer company." It's like, "Oh, shit. Is this thing going to work? What is going on? This is not the way you do things."

Hendrie: This isn't how you do it.

Bell: This isn't how you do it.

Hendrie: You build this product and then you tell everybody.

Bell: Right. That's why I say it was just absolutely terrible. Right after the announcement of our coming out party, we went to meet with one or two of the board members of our potential mergee. One of them was Morgenthaler from Cleveland, the father of Gary, the Silicon Valley VC, who was on the board of the company to be joined with. We talked about acquiring the company. We had a slight problem in difference. They had a book value of like $100 million or so. We wanted an equal book value. For five

guys or ten guys, we said, "Our book value is like $100 million. We're going to do an equal merge. No, no, maybe we'll give you $5 million for that and then you're going to come in and run the company." It took about a minute or two into this conversation for me to get it. I don't think I understand these financial transactions, but oh, I think I do understand! It took a couple of minutes. Oh, this is a dream, it is not going to work at all.

Hendrie: These people have a balance sheet and they can add it up to $100 million. We just say $100 million.

Bell: Yeah, that's what our balance sheet is. Gee, our stockholders could sue us if we did this deal for half of the company. So we went home and that was Encore I. Ultimately, I call these Encore I, II, III, or IV or whatever. I think we ended up with Encore IV or something like that.

Hendrie: So this was Encore I.

Bell: Encore I. Then Ed Fredkin came in and said, "Hey, you guys. Let's make PCs and sell them." Okay, I guess we could do that. So we did and Ed was sort of beating us up about trying to do that. Nobody could get enthusiastic about making PCs and selling them. Henry and I were like, "Okay, well let's see. How do we do that?" We couldn't get with the program. Then I think there was another one in there we explored for a while. But what we ultimately ended up with was a series of startups because after all somewhere in the process we ended up with 50 million in the bank. We acquired a DEC engineering team to make a multiprocessor. We had Formation. There was a group in Littleton that we hired and we made the Multimax, which was one of the early multiple microprocessors I defined as a "multi". It was right after Sequent and the other company who was first was called Sequoia. Anyway, it was one of the first multiple processors.

Hendrie: We were making multiple processors, but fault-tolerance was the story.

Bell: For fault tolerance, you had a clear focus. Anyway, ours was a general purpose machine. We used the National chip, we were a competitor with Sequent. Charlie Rupp came with us and he built a very nice, absolutely gorgeous terminal which was a 19" terminal that had multiple windows that you could connect to. So we'd go to these sales guys and say, "Well, you have to establish a channel for these. You have to make a market." They said, "We're mini-computer salesmen. We don't know how to make a market. Or sell terminals." So anyway, we didn't get very far with that. But it was a very interesting and quite a nice terminal. So the terminal was just sold by ourselves with a few on the side. The company ultimately sold a number of Multimaxes.

Hendrie: You built them and they worked?

Bell: We built them and they worked and there was a little bit of flack on that, and again, I went back and started doing more engineering. Somebody said they were trying to help marketing and whatever. As the design progressed, it got behind and whether or when it would work is unclear. Somebody said, "You know, if this thing doesn't work, it's your fault." I said, "What? Okay. You mean I'm the chief technical officer, so man, I've got to go work on this product." In fact, I was talking to somebody the other day of how the Mac saved my life. I had gotten back in with the engineering of it and the project simply wasn't

clear. The project was out of control. These were young engineers and they were enthusiastic and bright and all that and lacked discipline and management.

Hendrie: And hardworking.

Bell: And hardworking. But "What are we going to do next?" was how they were approaching the engineering. So basically, on a Saturday I went to New Hampshire and bought a Mac, because they had just started shipping MacProject. I came back, sat down one Saturday morning and started putting the project into MacProject. Then I came back. I bought my Mac on Monday and then I said, "Okay, we've got to get serious." When we finished, it had 500 blocks and we had posted it on a large whiteboard and that became the project management tool. But then we did the standard project management stuff of "Okay, every Monday morning we're going to have an update of the project: what did you do last week and what are you going to do next week and are there any changes to the schedule?"

Hendrie: Where are we?

Bell: "What are you we going to do?" So that was getting Encore going. We got a DARPA contract, which was based on a design because I was interested in building scalable machines. We proposed to DARPA that we'd build a 1,000-node multiprocessor. We basically took the Multimax, which is a 20-processor. Then you could put 20 of those together and put them in a hierarchy and then use the hierarchy to manage the coherent memory across all of the machines. So that was basically the proposal for what became KSR.

Hendrie: I was going to say, I'm hearing KSR in the story a little bit here.

Bell: Exactly.

Hendrie: Okay, memory coherence.

Bell: That was at a time when I believed in a religion where memory had to be coherent. I gave up on that religion in 1994.

Hendrie: But this isn't 1994.

Bell: This isn't 1994 yet.

Hendrie: This is 1984 or 1985.

Bell: Yes. We made the deal, I think, in 1985. We got a DARPA contract. I made the DARPA proposal. DARPA got granted the proposal. It was accepted after I left. I left in this case very early, because what happened was Henry and Ken got into an argument about the company and how it was going to be run and things like that.

Hendrie: Oh, really?

Bell: So fundamentally, the board fired Henry and so I said, "If Henry goes, I go."

Hendrie: And that was that.

Bell: And that was that. That was simple. So I left. It's too bad. I think that the company could have done all right, but in fact, it was just poorly managed. By the way, Julius Marcus for marketing came over to help. And Bob Puffer for manufacturing so it could have worked, but Ken valued friends over competence. His friends were just incompetent.

Hendrie: You had really good people.

Bell: We had really good people. And the problem was Ken didn't know what good people were. After I was there not very long I sort of said, "Well, he's maybe five milli-Olson's as a manager." Even though I may or may not have agreed with the way Ken Olsen managed, he did manage versus leave them alone. "Ken Fisher, what's your management philosophy?" "Hire good people." "Then what?" Encore survived much longer than I ever expected.

Hendrie: Didn't it merge with a minicomputer company?

Bell: SEL. It was a different one. It wasn't the one we ended up with. It merged with them, and then I'd say a few years ago the assets were sold to Sun which actually ended up with things like the memory coherent protocols and things where some patents were involved. Those were sold, too, as the motivator.

Hendrie: I see. I didn't realize that.

Bell: Sun paid $150 million a few years ago, I believe.

Hendrie: So Ken Fisher is probably doing fine. As soon as he bought the company in Florida, I said, "I know what's going to happen. He's going to move the company to Florida, so he can live in Florida."

Bell: That's right. Anyway, that was Encore.

Hendrie: That was that. So you quit?

Bell: Yes.

Hendrie: Now what are you going to do?

Bell: So in January 1986 I didn't have a job. I had been in California. I was on the Technology Advisory Board and I was a board member at Silicon Compilers. So I went out there to the valley periodically. One of the times I was out there, I had met Ben Wegbright and Steve Blank and maybe a couple other guys. By the way, one of the things we also did in Encore was encouraging John Hennessey. "Hey John, how about becoming part of Encore and we'll fund you to do the MIPS chip."

Hendrie: Oh, my goodness!

Bell: So John still credits me with getting him to start MIPS.

Hendrie: Really? Did that actually happen?

Bell: John got it going. He went and just did it.

Hendrie: He got funded.

Bell: Not me. Mashey [sp?] and I don't remember who actually funded MIPS.

Hendrie: It got funded by a couple of venture capitalists out in California.

Bell: Exactly. John Mashey and John Moussouris I believe were founders.

Hendrie: I was going to say John Moussouris was the other founder.

Bell: Who owes me $100 bucks, by the way, for a technology bet. But anyway, that was another story. John got the company going.

Hendrie: Okay. Let's take a break. I'm running out of tape.

Hendrie: Why don't we take a pause in your career because I think you're going to move to the West Coast and talk a little bit about your involvement or what happened in the genesis of the Computer Museum. Maybe tell that story a little bit at least from your point of view.

Bell: Right. You saw the e-mail that I sent out in the last couple of days. Given Bill Gate is going to come into town, I just sort of wanted to find out when our first connection with Bill was and it was after I had left when the museum and Bill had connected. That was the main thing. Well, anyway, going back to the genesis of the Computer Museum itself was the first really written public item was a little brochure, a six page fold out brochure, of a thing called computer generations. I had developed a theory of how computers formed based on technology generation and that these generate different kinds of machines. I had worked that kind of as a hobby way back to 1620 or so with mechanical devices and probably you can even go back to ropes and things like that, stones and ropes for records. But anyway in various times I had gotten interested in the problem of collecting artifacts. Ken Olsen was also particularly interested in that too. Ken wanted it from a perspective of the machines that he had worked on, namely Whirlwind and TX-0. So in the early '70s he had acquired the Whirlwind and the TX-0. A guy by the name of Bill Wolf had acquired the Whirlwind in the late '60s or so principally for the software. But of course there's not a chance in hell you could ever make that machine work again or afford to power it.

Hendrie: He put it in a building in Concord.

Bell: Yes, but I don't think he ever got it working. There were just too many cables and just too expensive and the power bill would kill it. It was at the wrong time because transistors were coming in. Ken was able to buy the machine along with Bob Everett who headed Mitre and was responsible for Whirlwind. It was put in a storage in Concord in cabinets, in storage crates, and then in fact I think also it ended up in some trailer in one of the DEC parking lots by the pond at one point in time, but that was kind of one of the geneses. We had also acquired the TX-0 through a collection of circuitous transactions because it was government surplus and initially obtained by a junk collector who wanted the gold. That was another thing. That was Ken's kind of first transistorized machine.

Hendrie: I didn't realize that he had also . . .

Bell: Yeah. Wes Clark was the architect of the TX-0 and the TX-2 and then Ken was responsible for the circuitry and some other things too. So anyway, that's how we got those artifacts. When I came back to DEC in '72, I had started to acquire various artifacts and thought, yeah, it's a pretty interesting time. Let's really collect those logic and other items, and I started collecting things in my office. I think at that point in time I had a piece of the DEUCE and I don't remember how I'd gotten it. It was the early vacuum tube machine that I had worked on in Australia I discussed earlier. And then in '75 was the first kind of record of a museum and we called it the museum project. I don't think we can find those memos, but by the way we might be able to now that Ken. . . although it's not clear how Ken's files are or whether that's searchable, whether we can find them through his files or not. Surely Ken's archives have items on the formation of the museum at DEC. Unfortunately, mine don't.

Hendrie: At Gordon College.

Bell: Yeah, about the museum. I think this would be a good point to go and try to ask them for anything that had to do with the museum.

Another origin came about when I used to go to Washington periodically for the National Academy of Engineering meetings. The first time I went to the Smithsonian I was really pissed that they had an absolutely terrible exhibit. They had some old computer, a Burroughs that I had never heard of, in one corner and they had some junk lying around. I tried to find out who was responsible and a lady by the name of Uta Mertzbach was responsible. She was a mathematician who was a Scheutz (engine) Scholar. I never could get anywhere with them and so I said, "Gee, we got to collect all this stuff before it disappears." So I basically said to Ken, "Well, we've got to do this, we've got to have a museum to collect all those things", and I had given some talks on my generational theory and then made the whole brochure. A lot of that was kind of still riding on the inertia that came out of the book that I had written with Newell on computer structures. And so we ended up taking the closet in the main Mill building 12-1 lobby and putting a glass front on it and putting all of our artifacts from my office in there.

Bell: And so that was the beginning of the museum. It was called the museum project. Then it moved to the Marlboro lobby. I don't remember what the date was. We had the whole lobby. It was a beautiful building. The RCA computer company built this Madison Avenue building in Marlboro that at the time was in the middle of nowhere.

Bell: And so we took all of that space and it became the Digital Computer Museum. It was really strictly a play on words as the Digital Computer Museum and it operated there. Of course it backfired and we had to call it the Computer Museum. Gwen at that point decided she wanted to run the museum and took

it over and took that as her project and did all the exhibits. Well, we used DEC industrial designers to do design stuff there and then we pulled everything together and meanwhile we'd been given a piece of the TI ASC. We had been given other machines and other artifacts and this was while it was still a part of DEC. And then Gwen was able to make it a public museum as part of DEC in sort of the '81, '82 time frame, which was no mean feat because of the IRS kinds of requirements of sort of hands off and all of the other stuff that you have to do and how you feed money to it and who's on the board and stuff like that. And then we got board members and what have you, renamed it to be The Computer Museum and then it moved to Boston.

Hendrie: That had happened . . .

Bell: In '83 or so when I was leaving DEC, it was moving to Boston. Anyway, that had been a deal that Gwen had cooked up with Mike Spock who was the head of the Children's Museum. He came to her and said, "Gee, the other tenant as part of the Children's Museum (which was an auto museum) is folding."

Hendrie: That's Larz Anderson.

Bell: The Larz Anderson Auto Museum. Yeah.

Hendrie: The museum that had come in from out in Brookline.

Bell: Anyway, so we got the space and that was the transition about the time I left, and then that was a potential threat too because I was no longer part of DEC. But Ken and everybody maintained their responsibility to the museum and it continued. DEC was the dominant provider in Boston and sponsor.

Hendrie: I wanted to make sure we covered that. Now you've left and you've finished Encore and what year are we? Do we know?

Bell: Yeah. This was '85, December of '85 or January '86. I decided at that point I was going to go live on the West Coast. I had been out there on a couple boards. So the MIPS had been happening, and one of the guys out there I said I saw was Ben Wegbeit with Steve Blank. And so we said, "Well, let's make work stations based on the MIPS or build a simple computer" and basically that was the genesis of that startup that was undoubtedly the most fun I've had as a startup.

Hendrie: That was how that . . .

Bell: . . . that started. And then Allen Michaels got involved because these guys were all the founders of Convergent. The company got together in January of '86 and it was called The Dana Group because it met on Dana Street in Palo Alto where Ben Wegbreit lived. It was called Dana Computer and it remained Dana for a while until the Dana Corporation . . .

Hendrie: Came in and said . . .

Bell: . . . came in and said, "Hey, you guys, get the hell off of our name!" So it became Ardent, and notice that Ardent starts with an 'A' and it's Allen looking through the dictionary for a name.

Hendrie: I see.

Bell: What are we going to be and it was literally we have to change the name tomorrow morning. We even had a contest. Steve Blank wanted to call it Defiant.

Hendrie: It was based on MIPS. When did you meet Hennessy?

Bell: I have known Hennessy going back to while he was doing his work at Stanford, so I had known him from DEC and Stanford Days. I knew what the project was and that's when I said the world needs a chip like that because we're stuck in Intel hell.

Hendrie: Intel hell and Motorola.

Bell: Yeah, right. Motorola was really doing that but at that point the RISC thing was just ripe. It was just right at the right time and you make the compiler and then you had the right caching structure and everything worked exactly right.

Hendrie: Exactly, and everybody followed and nobody succeeded.

Bell: Right. Yeah. In fact, the irony is I wrote a paper for some consulting company that sold position papers, and I think that was maybe '92 or so, but I think I'd given myself 10 or 12 years that there are six platform companies and at most there can only be three and so HP went down first.

Hendrie: Now you get to guess which ones.

Bell: Yeah, right. HP and DEC and then Sun hung on. The first two were clear. It was Intel and IBM and then the question was is somebody else going to be a survivor too.

Hendrie: Tell me a little bit. Did you recruit any people from your previous history at DEC into Dana?

Bell: Well, I didn't stay very long at Dana and I don't remember the exact chronology. I think at the same time Erich Bloch had contacted me to go to NSF. And the Dana guys were doing just fine. I was helping recruit there and doing architecture and I didn't want to manage the project and so Erich had recruited me to start up the National Science Foundation Computing Directorate. I did that starting in the spring of '86. So I went to NSF then in early 1986 and that was a startup, too, as I took all of the funding areas from other parts of the NSF and formed a directorate called CISE for Computing and Information Science and Engineering. The 2008 budget is maybe a half a billion dollars of computing research money, maybe more. It's one of the largest directorates, but I pulled together some of the hardware engineering, all the software engineering, and the supercomputing centers that were just forming that had network as part of them, pulled all of those together and formed a directorate, AI and I think Library Science. There's a bunch of different areas, and that was the beginning of the computing directorate. That was really a fun time. I enjoyed it. There were alternate days I felt like somebody could actually run the country, but it was not nontrivial but I got an idea of how things worked and what the executive branch did. And Reagan was the president then, and the problem with him and I'd say the problem with most of the presidents is that they don't have any understanding about science or engineering or where does all this money come from

and certainly science figures in to all of it. There are a lot of hugely contentious issues and he's not willing to hear conflict. The last thing somebody wants to have working for them is to have this person who's got all these problems and certainly that's an area where you've got all of these cross agency problems which are hugely complex. Virtually all of the Presidents' science advisers have gotten fired, or neutered.

Hendrie: Because . . .

Bell: Because there's something that has to be dealt with.

Hendrie: Science is not politically correct.

Bell: Yeah, and that they demote it or certainly they're way down in the pecking order of what to do because all the power resides at the cabinet level. It's like Rumsfeld having to argue with the Department of Energy. So here you've got these two very powerful kinds of guys going after each other and then you throw in these other little bureaucracies like NASA and the National Science Foundation. The wonderful thing about it is that NSF is a thing called an independent agency like the Veterans Administration so it just kind of sits over there. It doesn't report to the executive branch and so it's good and bad. It kind of gets money from the committees per se and so you've got a huge political problem there because it's the President's budget but yet in this funny way he's not quite . . .

Hendrie: . . . a powerful advocate.

Bell: Right. There's no line management going back to him. In fact, I just had a discussion with one of my friends today about that. I'm on the Department of Energy Committee to look at supercomputing. Well, it turns out another old friend from the Computer Museum is the Secretary of Energy, Sam Bodman.

Hendrie: Yes, Linda's former husband.

Bell: Well, anyway, Sam Bodman has obviously perhaps given money to the Republicans or to Bush, in addition to being competent and knowledgeable about energy as a former MIT Chemical Engineering Professor. So the Department of Energy is reeling in money and everything else for overall support. And so our little committee is looking at supercomputers and what are you going to do with all of those machines, and meanwhile they've just gotten their charter that basically looks to me to be the same as the National Science Foundation's in spite of the department's charter to focus on energy. So they've got so much money that they're going to have computing for whatever you want. It's no longer a mission program but it's, "Hey, if it's just support of good science we'll give you computing."

Hendrie: That's pretty interesting.

Bell: Very interesting and so there's competition now at the federal level for who funds what science, as if we didn't have enough intrigue before.

Hendrie: Has the NSF fundamentally taken the place of DARPA? DARPA really did this role of funding computer research.

Bell: Right, and in fact what has happened is that over the last few years NSF had to absorb that not because it wanted to but more because DARPA has really abdicated anything but mission funding. And in this regard I think there are all the people in computing. . .The universities are not getting money from DARPA now and principally because that's been a decision by the director of DARPA. It's a tragedy, a real tragedy for the country. This is bad, this is really bad. The current director of DARPA, a guy by the name of Tony Tether, fundamentally is giving money to the Beltway Bandits undoubtedly for PowerPoint slides. Over time DARPA was having to become more political vis-à-vis the universities and so it had gone through a period of a very small number of very large projects, small number of schools getting funding to being more egalitarian. It looked more like NSF in funding too many, in my view, too many projects, not enough . . .

Hendrie: Not enough in-depth projects.

Bell: And then now it's that there's none at all, so that's kind of where things are with NSF and how it's gone and I think we're in a pretty tragic sad state at this point in time with DARPA, and now we may be in an equally bad state with the Department of Energy taking over more science per se from NSF.

Hendrie: When you got to NSF you had a charter. Erich was the director to form a computing division. What did you do?

Bell: This was a sort of true tree graft of pulling together a bunch of programs from other divisions. So I basically rearranged an organization, and I'd say the part that was probably the most contentious was what are you going to work on? I outlined a program basically in parallelism, that in fact our biggest problem with computing was going to be how to exploit these machines that were going to be coming into existence over the next couple of decades. And so that was one of the cores of the thing if I go back and look at the statements I have made. The other thing that happened was what to do with the supercomputing thing because we had just. . .That goes back to the VAX because the computing guys had sort of said well, I had wiped all these centers out, now what we've got to do is re-create centers because we're not getting enough computers, we can't run big jobs on our VAX. By the way, it might have been totally different had the 8600 come out. My belief is it would have or maybe not, because in fact what happened was the Cray in '76 or so and so the Cray-1 and Cray XMP and stuff others like that were coming. Those were very powerful machines that are different animals.

Hendrie: They did have a big leg up on big . . .

Bell: Yeah. We simply couldn't do large . . .

Hendrie: . . . scientific problems.

Bell: You couldn't do those. So what happened was when I got there the division director was out opening centers and getting more solicitations for centers, and I looked at the demand for time and then I said, "Wait a minute, we don't need any more centers, let's get these running." So I ended up saying, "Well, I think we needed about four centers, not six centers."

Hendrie: At this time there were . . .

Bell: There were six.

Hendrie: And he wanted to open more.

Bell: He wanted more and then he also wanted centers because to get the program going in the very beginning he went out and bought computer time at some of the other centers and so these guys wanted to be centers too that would get NSF support. There were a few million going to Minnesota and other places and now we had real centers being formed, one at San Diego, one at Pittsburgh, one at Cornell, one in Illinois, and the von Neumann Center was opening up using the CDC Star. They weren't able to deliver it and they wanted to switch to Cray. And I said no, they got funding based on the fact that they had a different computer and now that the different computer folded I said well, let's fold the center. And so basically after I left the center ultimately was closed down and then Cornell was an IBM center. They were funded significantly by IBM and they had the 3090 vector. IBM had put in several of big vector 3090s, vector machines, there. And so one of the things I did was decide that well, we're going to run Unix on all of these machines so that we can now have some application compatibility.

Hendrie: Affordability of some of these applications.

Bell: Absolutely, the big problem here was apps. I remember writing a position paper to Larry Smarr at Illinois and to Sid Karin at San Diego asking them to switch to UNIX. They had Crays and they were running one of the Livermore time sharing systems or the DOE's, and I said, "Hey, you guys, you're going to run Unix, Cray- CTSS, the Cray timesharing, the Unix timesharing system." Ultimately they did but there was a lot of gnashing of teeth.

Hendrie: And screaming.

Bell: "You're telling us what system to run?" And I said, yeah.

Hendrie: You don't have to do it but the funding is likely to dry out.

Bell: Yeah, I invoked the "golden rule". However, that is small compared to the decision to take networking away from the centers and create a separate part of CISE to handle NSFnet. They had the networking. They were each building a network infrastructure for the U.S. to support just their own users. As you can imagine, we were creating 5 independent star networks. Why did we need networking? Well, it's for our supercomputers. That's the way it was sold more than any other way because we almost understood the application. Yeah, we got a network and here's the Illinois network and . . .

Hendrie: And it starts at the centers.

Bell: Yeah, and so I formed a network division and then that was the division that actually worked on the NREN Plan. NREN was the plan for internet. That plan was used to get funding and we funded regional networks from that group to tie in to the backbone. With a backbone network we got all of the other agencies involved in doing all that. This was in response to the Gore Bill [Supercomputer Network Study Act of 1986]. Gore, NSF, and his staff had gotten Gore to write a bill that sort of said the NSF should write a plan for . . .

Hendrie: National infrastructure.

Bell: Yes.

Hendrie: Or the information superhighway.

Bell: I led the cross-agency group from DOE, DOD, NIH, NASA, and NIST that put that plan together. We had a meeting in 1987 in San Diego. I remember getting up the last day of the meeting in the morning and saying here's the plan. There were about 300 people there. We had had sort of subcommittees and everybody was talking about different things and I had listened for several days, went in to listen to all the technology and other ideas for research. Well, what's happening, what's the state of art, what is it going to be, and it was clear that Optical was nowhere in the immediate timeframe. So I drew a diagram that had these five waves of bandwidth versus time going out to 2000 or 2005 and then I said this is the plan, this is what we're going to do.

Hendrie: This is what we're going to do.

Bell: Yeah, and note there's no research in networking per se until we can go to Optical networks. This is an engineering problem, we've got to now get this managed and we started building the network out and then ultimately there'll be fiber that will kick in around early 1998, when we started using higher speeds, but until then it's just network engineering. The first thing was phase zero stabilizing what we have because we don't got nothing now and then put in T1 and make that work and then we go to T3 and then we go into optical. So that set of curves and those dates turned out to be almost precisely what's ended up happening. I don't remember how we made the deal, but Bob Kahn was there and I said how about becoming the funding instrument for all of this because NSF can't deal with all of the funding and making this happen, just go out and get it done. And the first deal was that Michigan took over the management at the network center for the first wave, and I think they still do most of that and they work with IBM who supplied the initial routers and gateways and was responsible for the network.

Hendrie: For the basic backbone and this is the national backbone.

Bell: The backbone's been replaced but with higher speed links. Now NREN 2 is out and so we are starting down the path that seems to be back to switched optical links. I led the group that developed a three volume plan or rather the summary plus background of two volumes, well, the three volumes. The plan is of what we're going to do and why we're going to do it is about a 15-page report that went to the President's science advisor, Bill Graham. I tried to get him to write a really strong computing infrastructure charter, but he didn't want any part of that. He was a very weak guy because he had to interface with Reagan or someone on his staff. The way it worked was there was a cross agency group called FCCST, Federal Coordinating Committee for Science and Technology, for IT and so I led the networking part. There was somebody from DARPA leading the people infrastructure and their high performance work, Jim Decker led the supercomputing, and I put the networking part together and so we had a great plan. I walked in with my friends and said, "I want to respond to the Gore Bill. Here, the Gore Bill's due. We're going to send this off and have it done." And they said, "Like hell you are, you can't do that, it's too good, no, you'll make us look weak and disorganized" like they were. And so we waited another year to create the high performance initiative, a little blue book that came out. In fact, Steve Squires drove the production of it. He came in from the DARPA because they were working on high performance computers and also CS infrastructure.

Hendrie: Was that the principle activity or most significant activity that you were involved with at NSF?

Bell: Yeah, I think the thing that we did was really establishing networking, getting the supercomputing program on the right foot, even though it meant there was some blood flowing, and switching, going to a more universal standardized environment. And certainly the NREN response to the Gore bill was really important and getting that program off and putting the parts together was really great. Also, the overall program and organization I set up, along with the goals for CISE that included a focus on parallelism were important.

As a by-product, I established the Gordon Bell Prize for parallelism in 1987 that is over 20 years old and currently funded at $10K/year. I have funded the ACM to run it for another 20 or 30 years, too. The prize is administered by the Supercomputing conference and papers are presented there. In 2008, the parallelism is about 200K to deliver almost one half petaflops. In 1987, that was one half a gigaflops with 1K-fold parallelism. So I consider the prize to be very useful that came from being there.

At that same time, I was commuting to the west coast because the Dana group was building its graphics computer. I would go out every month or so on a weekend or to a board meeting and interact with those guys. And then I don't remember when it happened but, at one point, Allen . . .

Hendrie: Allen Michaels?

Bell: Allen Michaels called me and said, "I fired Ben. You've got to come here and run this project."

Hendrie: Oh, my goodness.

Bell: I said, "Oh, wait a minute."

Hendrie: Why did you do that?

Bell: "Why did you do that?" And he said, "Well, he's not getting it done" and then he said, "What do I do?" And now let's just talk about the Ardent group, it was first Dana, then Ardent, and the best group of engineers I have ever met in one small . . .

Hendrie: In one spot.

Bell: In one spot. I mean, we built the supercomputer, a parallel vectorizing FORTRAN compiler, blazingly fast graphics, the whole thing with this team of about 50 engineers.

Hendrie: How many?

Bell: 50.

Hendrie: 50, okay.

Bell: Everyone had done spectacular things both before and afterward, you know? For example, Cleve Moore Moller was the founder of MATLAB. Okay, "What do I do?" And I said, "Okay, Allen, I'll come out in a couple weeks but here's what you do. You take all those guys and you put them in a room and you don't let them come out until they have a schedule that they all agree to." It's that simple. Once they have the schedule, have them meet weekly and ask everyone what they did the last week and what they would do next week. There were about a half dozen groups: the computer hardware with about four complex boards with very complex custom chips for the vector processor, memory, i/o, and graphics, graphics software, the O/S, and compiler.

Hendrie: Yes, okay.

Bell: And so fundamentally that's what he did. That's what happened. When I came out, I think the schedule was being formed at the time, and then I interacted a little bit as they were doing it and ultimately got a schedule and I introduced the term schedule fantasy factor and said, "We have to manage to the schedule and that's what's going to happen." There SFF was about 1.8, i.e. it was taking almost twice as long to do things as they predicted. What happened was classic with I'd say a super bright engineering manager. Ben had projected his schedule onto everybody somewhat by intimidation.

Hendrie: They project their own ability.

Bell: It was a way of extracting commitment. And so what Ben was doing was he had, like, ten guys reporting to him and he managed them one on one and he was the gateway and he clearly could do that. But he did it sort of by intimidation. He'd bring them individually into his office and said, "Now, you got to do this, this, and this" and, "We need this by this and that, will you do it?" and "No, I can't have it done," "You got to have it done," he asked, "Next?" There wasn't co-ordination of the parts except through Ben. Worse yet, with that kind of management, you don't have any coordination and then also you had no commitment.

Hendrie: You have no commitment. You have unrealistic schedules. So you got rid of that, and it worked.

Bell: Yeah. Back to basic engineering management, just like I did at Encore when I had to take that one over.

Hendrie: Now, tell me, what was the fantasy schedule thing?

Bell: I don't remember when I (or if it is something someone else came up with) had come up with the name schedule fantasy factor. You take any future milestone time, you look at what the commitment time is for that milestone, and then divide the actual by the committed and anything greater than one is a fantasy. I think their schedule fantasy factor, at that time was 1.8, so it was taking them 1.8 times longer to do something than when they said it was going to take. After they made their own schedule together they got it down to 1.0, a rarity for engineering.

Hendrie: Exactly.

Bell: And so . . .

Hendrie: There's every reason to predict that it's the same factor forever.

Bell: It's a constant for a group. Yeah, that's exactly how optimistic people are. You have to be realistic.

Hendrie: It's got to be realistic.

Bell: And so they got it right down to one. Actually, we were on time and, by the way, any time anybody comes in with less than one, I mean, to me they're heroes. The only time I ever saw that was the Network Appliances startup team got a product out in 0.8^{th} the time they committed.

Hendrie: Wow.

Bell: And so at the time there was a question about funding. The archangel who was leading the group who was dealing with a lot of the funding wanted to hold them up because they hadn't achieved their funding milestones, they didn't have the VC yet and so some money was due and I caved. I said, "Hey, you know, let's not hold them up" and he wanted to use that to extract more flesh and so I said, "No, come on, let's get on with it." And so all the angels were really pissed at me because here's a guy that caves on funding.

Hendrie: Right. But they'd done just a spectacular job.

Bell: Yeah, exactly. They should be rewarded. I Said, "I'll never be a VC."

Hendrie: Okay.

Bell: So I stayed at Ardent and got it to ship.

Hendrie: About how long was that?

Bell: We shipped in '88 which is still amazing for the project, the three-year project, and then we merged with Bill Poduska's company.

Hendrie: Yeah, it was Stellar.

Bell: Stellar and Ardent to form Stardent.

Hendrie: Yes. And they both thought they were their own worst enemies, and if the merged they'd be okay, only to discover Silicon Graphics was actually the enemy.

Bell: Yeah. It was pretty bad but, anyway, I still get Christmas cards from Kubota, which is the Japanese company that funded both of them. I left and then did angel funding and consulting until '95.

Hendrie: So you left, maybe this is '88?

Bell: '88 or '89, yeah. I was enjoying life and just floating around the valley and was on Suhas Patil's board at Cirrus Logic and other things including consulting at SUN and Intel. I tried to get Intel to build a vector processor. The net result of this is during the period 1983 till 2007, I have invested time and/or money in over 100 companies or about four per year.

Hendrie: Wow.

Bell: Yes, from '83 'til now. And then one of the other things that happened during that period while I had worked with Heidi Mason was the creation of the Bell Mason Model and Diagnostic for startups. Heidi had been the editor of a book that I had written at DEC. She was out here doing marketing -- she had started a marketing company and then we started collaborating on consulting. She had me help look at some companies regarding technology. That's when we created the Bell Mason Diagnostic that we have licensed to Coopers and Lybrand and others. The idea for it came to me while going back and forth to Washington while I was an NSF in 1988. The idea is that there are 12 dimensions for measuring a startup company. You plot the state of the company on a 12 dimensional kiviat diagram. Excel calls this a radar chart or polar chart.

Hendrie: Well, tell us more about that. What is that diagram?

Bell: The whole idea of evaluating startups and projects had its genesis back at DEC and in the early '80s.The popular idea that was being pushed in the '80s was the idea of expert systems and heuristics. I wrote a little paper called "Heuristics for Building Great Products" and said management and business is just a set of rules. So I cooked up about 40 rules and segmented them into about seven dimensions. You had product, market, team. After I left DEC, I got involved in all the startups. I like startups because of their simplicity. When I left DEC, I basically said, "I'm not going to ever go back to a large company." You know, these startups are so simple because they're just projects. They are not companies! Think of it that way. So, by this time I was commuting, and Heidi was talking to me about all the marketing problems that were going on and the companies she was advising -- it's their plan, it's the team, it's the sales, etc. We were doing all this stuff, and then I had enough knowledge too, because I was on the Silicon Compilers Board. John Doerr had founded that company and was the chairman. And so I had a pretty good idea of how startups really worked. And then I started writing about it and I cooked up the idea of just 12 dimensions. I liked that number, and then the idea of heuristics was how you look at and measure them based on things that they should be doing. So we had a grading system, we had the dimensions, we had what I thought was the orthogonality. I haven't changed my mind at all from all that. And so I started getting more experience with it, started writing. I was fascinated by this the idea of stages. I found out the stages didn't mean anything. I read all the academic papers and the stages only meant we ran out of money. We need more money, that's what the stages were.

Hendrie: Yes, those are the stages. Stage 1 is how long the first money will last.

Bell: Yeah. And I then used or rather took the software engineering model to say these are really the stages of the product, these are events that you can measure. So Heidi and I cooked up the diagnostics, which has about 1,000 questions over the 12 dimensions. There are actually four stages but then sub-stages and thus specific times at which you can really measure companies. And then I wrote the book, it came out in '91, and, again my old friend John McNamara came to my rescue. I started writing the book and I had a publisher -- actually Addison-Wesley -- and they tried to get me in a co-author. But all the co-authors were doing was re-editing or rewriting stuff I had written, and John looked at it and said, "No, its

written wrong. Just turn it 90 degrees. A chapter per dimension as opposed to a chapter per unit of time." And so I changed all of that and then I wrote the first draft and then John rewrote it and then we ended up with a really good editor at the end, too, and she did a lot.

Hendrie: Good.

Bell: I like the way it reads.

Hendrie: Good. Hadn't John written, way back when, a book on communication?

Bell: Yeah, John had written a book on communications at Digital and he liked to write. In 1978, Craig Mudge had convinced me to write a book about the DEC computers. And so I started the book with Craig and found out that Craig needed more help to get it done. He wasn't writing at the velocity that I write. And the book would never get done or would certainly take a long time.

Hendrie: He'd never get his parts done?

Bell: So Heidi and I called John and I said, "John, you got to help on this book" and so John came and helped get the book Computer Engineering done.

Hendrie: Very good. Excellent.

Bell: Yeah, so the key is knowing who to ask.

Hendrie: Yes, exactly.

Bell: When you get into a bind.

Hendrie: Okay. We got off on that book but you're eventually going to get to Microsoft? Have you gotten there? Are there any other stories on the way?

Bell: Yeah. I think, you know, there are 90 stories on the startups, and what I did on those, by the way, I ought to do something with them. I have about four hours of some very pithy remarks on some of those because I basically took all of the shirts and cups and coats and sat there before a video camera and rambled on, "Here, see this is the story of this $185,000 coffee cup and what I learned."

Hendrie: Yes.

Bell: Yeah. I did the startup and consulting gig until '95 and then I went to Microsoft. But in that time, in '94, I ended up for whatever reason meeting Jim. I had never met Jim Gray before, and so Jim and I met in '94. He came to my house and we talked about what computers were going to be like and we decided we had exactly the same religion about the way computers were going to be built and he called them the smoking hairy golf ball or bricks and I had just written a paper on networks and nodes which we changed

to SNAP, which is Scalables, Networks and Platforms, and the only thing we got wrong was that we thought ATM was going to be the interconnect.

Hendrie: The connecting fabric.

Bell: The connecting fabric so that we could deal with the scaling across the geography, but ATM didn't happen. But everything else was the same and that scalables were going to take over all these other classes because I had refined my computer class formation theory. All different little classes and niches are going to be wiped out and replaced by scalable PCs. They are going to wipe out workstations, PCs are going to go right up through to minis and mainframes and eventually supers. They'll be the dominant structure and a scalable network. Then Jim wrote that up and he gave a talk about it in Berkeley in three McKay lectures. That was kind of the beginning of our belief.

In '91, I had met Nathan Myhrvold -- I think he called me, probably Cutler gave him my name. I had convinced him to give a talk in Geneva when I was giving a talk on startups because I was giving a talk about my book. We had two days on where computing was going. So I got Nathan to come and speak. We dined in the Two Star Geneva restaurant. Nathan said, "Okay, now come and help me start research at Microsoft."

Hendrie: Ah, okay.

Bell: And so I went to Microsoft and Nathan and I noodled about how it was going to start. I don't remember that he ever asked me to run it or not but at least I know I said, "Let's find somebody really good to run Microsoft research." And I've still got my notebook that has the list of people, qualifications, and our grades as to their suitability.

Hendrie: Oh, very good.

Bell: And the pros and cons of who we were going to hire.

Hendrie: Yes. Very good.

Bell: And then Nathan and I then went off and went to Bell Labs, met some people there, we knew various other people and then we went to CMU and Rick Rashid was at the top of our list. And so we went to Pittsburgh and met Rick who was much better than what we had thought because he had kids and the kids were using computers and he had appreciation of all the subtleties of consumer electronics. The family came to Redmond and then said, "No." He met Bill and said, "I'm doing just fine. I'm doing the Mach operating system and I'll just keep doing this." And so I told Nathan, "First off, he can't say no. What is this?" And so I wrote a four-page letter to him and his family saying you can't say no, don't you get it? This is the best job in the world. You can always be a professor. You'll be an even more valuable professor if you do this. But here, come and do this. This will offer you an incredible amount of financial freedom and, if you ever want to do anything, you can.

Hendrie: You can do it, yes.

Bell: You can do it. I started the technology advisory board for MSR and we got the same guys who are on it now -- Hennessey was on it, Andy Van Dam, Raj Reddy, Ed Lazowska. Hennessey resigned ultimately, and then Richard Newton who is now deceased was a member.

I've been happy going to the TAB meetings and the like, and then Jim and I got together in '94, worked for awhile and Jim had been consulting and then he called me and said, "I can't deal with this. I just need a place to work and a project. This is just not any good for me." So I said, "Well, okay," so I emailed Microsoft and I said, "Hey, we got to hire Jim and set up a lab." And so they took him immediately. It turned out he had been talking with them for a year or so, so probably all I did was to crystallize the charter.

Hendrie: It didn't take very long.

Bell: It took, yeah, about a nanosecond to hire Jim, and then Jim calls me, he says, "Okay, you've got to be part of the lab. You need a job, you know, your life is just floating around doing whatever. You need discipline". So he basically convinced me to go to Microsoft. In August '95, I joined.

Hendrie: So you've been there.

Bell: Yeah, I've been there ever since.

Hendrie: Very good.

Bell: And then his first question was, "What are you going to do?" You know, he hired me and then the next day he said, "Okay, what's your project?" And so I said, "Telepresence," because I'm going to stay here and I want to work at home. And I started out in that and then I, fortunately, hired Jim Gemmell, which was great, and Jim worked on various parts of telepresence and scalable multicast. One of the things was that I had us acquire a company doing telepresence. We gave up on it in '97/'98 because our protocol stack is too slow, the machines aren't right yet and it's going to be a long time before he gets there. There's just too many problems at this time. Interestingly, it has taken about 10 years and Skype is probably the most important catalyst.

And then I fell into the current project in '99 of deciding to scan all my bits and that got me onto this path of cyberization of the world. I think I've almost got everything that I've been working on, capturing all the personal bits I can find since 99, and building Bush's Memex that we call MyLifeBits..

Hendrie: Excellent. Do you have any comments to make about all the things you've done, you've accomplished, you're sort of the proudest of?

Bell: At one level, I go to the museum affair that we did like Fellows the other night and I thought, "Gee, that was really good. It's an institution." And, in fact, it's beyond a project, it may live forever and it's a great thing. So the museum is very high up in that list in spite of all the contortions that it went through to get here. It's something I feel proud about.

Hendrie: Yes, exactly.

Bell (6/28/08): It is critical to add here that if Gwen hadn't decided to do the museum at DEC, none of this would have happened the way it has. The Computer Museum in Boston was critical for collection and to get some passionate people involved that have remained as such including yourself who has been on the board from almost the beginning, Dave House, Grant, and John Shoch, and undoubtedly others I forgot. Ironically none of the Boston trustees have remained engaged.

Bell: Because, in fact, sometimes you kind of look at this thing and say, "Wow, that doesn't seem like that long." But on the other hand, it's over 30 years ago that was rooted in 1975. At one point I thought, well, if we could have a permanent building by '95 or so, something like that, that would be fabulous. I guess I'd say it really exceeded my wildest expectations. I think it's a tribute to just a few like you and Len Shustek and then the way we actually were able to get the money and the building at the right time, that timing was just spectacular because that gave at least five years of boost and probably ten. Having Bill contribute is a big deal, too, to get his support.

Hendrie: It certainly is.

Bell: Because that's all time, I mean, that's all trading money for time.

Hendrie: Exactly.

Bell: So the museum's got to be way up there. I always feel bittersweet about DEC because of wanting it to live forever. And that it should have lived.

Hendrie: It should still be here?

Bell: Yeah. And it sort of makes me very angry. Actually I've gotten over it mostly for example, you know, Dave Cutler had a similar feeling. So you get angry at these egos or incompetence that had it happen that way.

Hendrie: Yes.

Bell: But you certainly see DEC as a great time, that it was a great experience and you have to live with yourself, you have to take a philosophy away from that of, well, okay, at least a lot of good people were trained. I guess I tend to look at everything as kind of projects in a way and some are more important and lasting than others. That's basically how I conduct myself -- as an itinerant project engineer.

Hendrie: Yeah. Which ones . . .

Bell: I like writing. I'm in a quandary now trying to decide what to do next. When I was doing speech, research I quit it because it was a 20-year project. Well, it's a 40, 50-year project. Memex or MyLifeBits is sort a of brain augmentation or memory augmentation device, recording and augmentation and storing and all of that at a personal level, a person's life events and knowledge so I can't even see an end. There's no end to the potential for it.

Hendrie: Yeah.

Bell: So in a way I feel good about finding a project that has no end. On the other hand, it's a little frustrating as an engineer.

Hendrie: It isn't something you've done before.

Bell: Right.

Hendrie: All my projects have had ends.

Bell: We can do this in two years, three years, six years, let's go do it. I've loved all the projects. I mean, it has been frustrating at times but that's the way all projects are or should be, at least. Otherwise, it's not a project.

Hendrie: You probably didn't do something, would take a big enough step.

Bell: Right. Exactly.

Hendrie: It wasn't hard. Alright.

Bell: But, anyway.

Hendrie: Good. Well, thank you very much, Gordon, for your patience.

Bell: I agree. This is amazing, that we got to convergence.

Hendrie: We converged. Good. Thank you very much.

Bell: I was worried there about an hour ago.

