

INT: When did you fall into this role?

BS: The Mill?

INT: No, with Digital.

BS: With Digital? That would be just over 20 years ago. One of my professors at school suggested that I check this place out when I was looking for a job. So I gave him a call, and they said something along the lines of, okay, we'll think about it or something or other. We'll set something up. And then my wife and I went down to the Cape for a week of vacation, and we were passing back through our apartment in Cambridge, and on our way from somewhere else, Maine, I guess it was. And while we were there changing our clothes, the phone rang and it was Digital at personnel saying, why aren't you out here for your interview. So I explained that nobody had told me about it, and came rushing out here, and had an interview. They told me to look for the old mill in Maynard? So I was sort of thinking, hmm, there must be one of these quaint old grist mills with this gigantic modern corporation behind it, right? We went looking around Maynard, and finally found this gigantic brick monstrosity and concluded that it was Digital. And eventually found the right door and got interviewed.

This would be '71, yeah.

INT: What did you find?

BS: The 1145 group. Dick Clayton, Bruce _____,
Ben Hughes, all sorts of people. A few of them still
here.

INT: Well, let's see, that was eventually in Plant 2
(?).

BS: Yeah, the first office I had was probably about ten
offices toward the parking lot from here. My first
office was directly underneath the loading dock on the
next floor up. So whenever they dropped a barrel on the
floor above me, I'd jump out of my shoes, and the dust
would rain on my head.

INT: They offered you a job immediately to do what --
what were your expectations?

BS: Well, my expectations of course, coming out of
school and into a real job were a little bit weird. But
I sort of knew I wanted to get into the processor group.
That sounded like the most elegant or sophisticated or
something. So the 1145 group seemed pretty interesting
to me. And they made me sort of assistant engineer on

the memory management _____. They also gave me the responsibility for a special purpose debugging, hardware debugging console that they had started developing. And it was unusual. It was the typical sort of thing. We did address break points and looks for certain types of data on the bus and that sort of thing. And it looked much too complicated to fit in the space on the backpoint that they had allocated for it. So the first thing I did was go to talk to the software people to find if they really needed all this function that they had asked for. And I asked who I should talk to. And they said, well, you need to go talk to Dave Cutler. You might find him a bit unusual, they said. So I went off to see Dave Cutler and explained what I was thinking about. And he made some comment along the lines of something like, You fucking asshole hardware engineers, why are you even bothering to do it if you're ^{not} going to do it right, and on in the same vein. And that was my introduction to Dave. So I went back and figured out how to cram an interface module into the space on the back plane, and put the rest of it in a separate box. And eventually they cancelled the project, because they figured nobody was going to buy the thing. There were only I think two of them ever made. They worked, but they weren't worth finishing.

"Mr. Abusive"

11/46

INT: That was an extension of the first PDP-11.

BS: Yeah, it was sort of quite an extension, but it was an extension. I guess Digital had discovered compatibility at this point. So after the 1120, they made some other PDP-11s that were pretty similar in terms of construction set. Many of them had run from a lot of almost the same software as the original 11, and the 1145 was one of them. Much faster, and it had a floating point.

INT: What was the reputation at that time of the 11 group?

BS: Within the company?

INT: And outside the company?

BS: Well ... Well, you're going to get kind of a warped viewpoint, because my outside view was only from college, and my inside view was of course limited by being a new person. But from outside, about all I'd seen of the 11 group was the PDP-1120 handbook. And it looked like a sort of unusual, innovative and interesting sort of machine, but I didn't know much of anything about the people who designed it. From within the company, I guess there was beginning to be some rivalry with the PDP-10 group even at that early stage.

INT: They were still in the Mill right?

BS: Yes the 10 group were still in the Mill. They were working on the KI-10 at that point. And there were already starting to be snide cracks made back and forth about the ... The 10 group apparently felt that the 11 group didn't pay attention to details, didn't do a thorough job on technology, had a machine that was too small for any practical purpose and so on. And the 11 group sort of felt that the 10 group wasted money, built gigantic boxes that didn't go any faster than our 11s and so on. So it was already starting by then.

INT: But do you think the 11 was a radically different machine at that point in terms of the instruction set.

BS: Yeah, the instruction set was certainly radically different. In retrospect, one might even say downright weird. I don't know that the design style was all that radical at first. The 1120 was done before I got here, and basically a small to medium scale integrated circuits and ordinary state machines. Much like some of the previous machines I guess.

INT: Just existing technology.

BS: You should talk to an 1120 person to find out

whether they thought it was existing technology. Looking at it from my point of view it didn't seem all that new and exciting. Which probably made a lot of sense. If you're trying to get a project done _____ risk at a time _____, often more than enough. The 1145 was a little more adventurous in terms of technology. I think we were sort of the first large project to use TI's new _____ series TTL. That was an adventure in itself. Probably one of the first in Digital to use _____ boards. The first at Digital to build a memory out of MOS memory chips. A somewhat adventurous technologies. I guess we were also the ones at Digital to discover synchronizing failures, things likethat.

INT: At that point you had to run everything off of one clock.

BS: Yeah, the 1145 was intended to be a big synchronous machine, and it was a much faster synchronous machine than we tried to build before. So everything was getting a little bit pushed.

INT: What happened after the 45 blew (?) out?

BS: They put me to work on multiport memory project. It was originally going to be used for connecting sever 1145s together for a multiprocessing system. We got as

far as the I guess you'd call it prototype. We had a bunch of noise problems. Which we didn't really have the experience to understand at that point. We probably would've solved it eventually. But the thing that really killed the project was that eventually somebody realized they actually didn't have any software to use on this multiprocessor, and that nobody was about to write any. So at that point they decided it wasn't worth anything.

INT: Well the software group wasn't as big, or was it more contained in the separate product engineering group?

BS: I don't know. My view of software at that time was fairly confused. I remember that there were several different operating systems being worked on for the 11, and they didn't exactly seem to report to any of us. They must have reported to somebody else. I do remember Dave Stone being one of the lead software managers back then. But exactly who decided there wasn't going to be any of their software or why I don't know.

INT: Was there close contact between software engineering and hardware engineering?

BS: There was a lot of contact with the diagnostic software engineering, but I don't remember an awful lot with the operating system and the power engineers. Maybe

I just wasn't in the right place to see it, but I don't remember very much contact. I do remember a few instances of the software people were apparently very involved in a couple strange instructions involving semaphores, and there was a peculiar instruction for stack marking (?) which they helped invent. It later turned out not to be particularly useful. What cooperation there was didn't seem awfully _____ at that point.

INT: Then what?

BS: Then I was doing a brief stint doing environmental testing for a 16K sense core memory, a new core memory module that _____ was developing. This sort of grew out of multiport memory which used the module. But now the basic memory using that module is getting ready to be tested in the environmental chamber. So I got assigned that job. We'd heat it and cool it, humidify and dry it out, shake it and do various things like that. I still have memories of trying to debug one problem that only showed up at 90 degrees and 90 percent humidity, sitting inside the environmental chamber with an oscilloscope and having to run outside every so often to cool off. We eventually did fix the problem, but it's not a job I'd like to have permanently. And then I got involved with the 1170. John Levy who is not longer here

had invited Bill Strecker in to talk to our group. Bill was I guess fairly new then. And he'd come from RCA and he gave us a talk about cash (? --cache?) memory systems. And this sort of struck something in my memory. I remembered reading an earlier memo by Don Grenada about putting a cash memory on a machine that was going to be the 1145 for never got built. And I started thinking about putting a cash memory on the 1145. And I talked to Bill and he did a bunch of simulations, and Steve Rothman and I got around to writing down some proposals, plans and stuff. And eventually that grew into the 1170.

INT: What was the concept of the 1170?

BS: The original concept was let's put a cash memory on the 1145 so we can get good performance without the unreasonably expensive semiconductor memory that we were using at the time. The 1145 was designed to run at 300 nanoseconds using bipolar main memory, or 450 nanoseconds with _____ memory, or something like 900 nanoseconds with core main memory. So you had the interesting speed potential of running at 300 nanoseconds, but nobody could afford very much bipolar memory. In fact, you couldn't put very much on it. You could only put in four boards worth. Most people used it with core memory, so a lot of its speed potential was being wasted. I was thinking of getting the speed of bipolar on the cost of core. And it

grew from there. We also ended up expanding the address space, which wasn't something I had originally been thinking about, but apparently was a good idea. I made a few other changes. And the thing worked out pretty well.

INT: That would be '73 or 4.

BS: Somewhere in there.

INT: How big a team was working on it. Was it a growth of the 45 group?

BS: It was sort of, let's see, the original 45 group had partly grown and partly split up. The 45 was a fairly big project and I guess some people during the course of the project either got burned out or decided they wanted to try a different career or something like that. Because a number of people went off to other places at the end of that project. The people that were left sort of continued along. I was part of that group. And the 1170 team sort of grew out of that in turn. Once we got the project approved it was, let's see _____ and I were working on the cash, Steve Rothman was sort of supervising the thing. Steve Jenkins was doing the _____ adapter, and the bus at that time. Sharon Smith was doing some work on the thing, lot of fine memory management I'd like to think. [INT UNCLEAR] No, not

very originally designed _____. I was assistant junior engineer then. My brain is not recalling everything after all this time. _____ was involved in that too. It was a fairly small group by today's standards. I guess life was simpler back then. In some sense of course it was a lot simpler. Nowadays we think of 15,000 gate arrays as being no big deal. And we try and get them done with one or two engineers. The whole 1145 probably didn't have many more than 15,000 gates in it. And so I guess life was simpler back then.

INT: But you didn't have any design array outpools (?) to do.

BS: There weren't very many design array outpools. There was actually a layout. There was a wire listing system, which ran on a PDP-8, and it involved going over to one of these old teletypes that produced paper tapes, and carrying them down to a PDP-8 and running it through a PDP-8, and eventually it migrated to a PDP-10, and we didn't have to deal with the paper tape which was a big advance. And the wire list at least let you check your layout when it came back and fix all the bugs that were introduced because most of the layout was done in the early days by hand. Eventually it started getting some automatic layout systems. But there weren't really very many other tools. We weren't using any of the simulators

or _____, or schematic drawing packages, *the layout was done* done by people who took drafting in high school. That aspect of the job has changed quite a bit.

INT: In what quality? Qualitatively is it different designing a single _____ with some of the layout tools and simulators and modelling tools, but is it qualitatively different in terms of what the hard things are?

BS: I think the design tasks are much as they've always been. It's just a question of whether you have the computer do the bookkeeping for you essentially. You still have to figure out what it is you want to design. You still have to figure out the stacks, in other words in detail what it is you want to design. The computer doesn't really do much for you there. The only thing that it helps you there, is it's easier -- in the old days you'd write your specs out and then a secretary would type it up. But now you can use your editor and type it up yourself. That's about the only difference there. You still have to think about how you're going to do your data path and control and how you're going to partition the various pieces. And once you've gotten that far, then the computer actually can start to help you. Because what you used to do is get out your pencil and paper and drafting template and draw gates on paper,

and then you decided you didn't like them, you'd erase them with an eraser. And nowadays you can get out your computer and your mouse and draw your gates on the screen, and if you don't like them you can erase them by pushing a button. But the drawing advantage of putting it on the computer isn't so much the advantage of drawing it. The advantage is that you've got your hard design in the computer and you can do other things with it. So now the real fun begins, because instead of having to figure out by hand which are the longest planning (?) paths in your design, and adding up by hand and keep track of when you change something, now the computer can do that for you in probably more detail than you actually want to know about, fairly quickly. And instead of you having to rely 100 percent on thinking through every possible case in your logic, you can have a simulator that can try out cases at random for you of special patterns you can think of and give you a somewhat better chance of your logic working when you actually build it. And instead of having to give your schematics to another person and have them draw lines on a circuit board or a chip, the computer will now draw the lines for you. And then you actually build the thing and get it back in the lab, and now you're back to pretty much the same thing we used to do. You have to sit there and try a bunch of things, and if it breaks, you have to get out the _____ to find out what happened. So the computer has done a lot of

bookkeeping for you, which has the advantage of letting you be more thorough and making sure you cover all the bases instead of just the ones you think about. Sometimes those things can be overkill____. You cover more bases than you actually need to, get confused on such a volume of trivial data that you have trouble wading through it all, you become obsessed with the minutiae of the design instead of getting it completed. It's not all sweetness and light. But it would probably be very difficult to handle a large complex design of the type we do today without the computer to do that.

INT: DEC has sometimes had the reputation of overdesigned products, making them maybe more elegant than is sometimes necessary.

BS: Yup. Yeah, I can understand that. It does seem to be part of the hardware engineering culture around here to make the thing work theoretically in the worst possible case. And probably you could sell a lot of machines without doing that. It's fairly deeply built into the culture here. I don't know exactly where that started, but it certainly started before I got here. The PDP-10 group even when I first arrived was very much that way.

INT: They seemed to be very proud of the elegance of the

design of the PDP-6, which didn't really work in practice, but much of it was carried over to the 10.

BS: Yeah. I don't know whether it has to do with the educational background of some of the earlier people or not. Perhaps coming from universities, one is used to solving problems and getting the right answer instead of just something that works. So perhaps that has something to do with it too. But it predates me.

INT: Can you think of any other elements of the engineering culture that were unique to Digital or at least characteristic?

BS: Well, I'd have a little trouble with that because this is pretty much the only place I've worked except for summer jobs. One thing, I had a summer job with an RCA defense division, defense and aerospace _____ group. And they had unionized technicians, which meant that engineers were not allowed to touch soldering irons, set up oscilloscopes or anything like that. So I was very relieved to see that Digital didn't have this problem. That was kind of a strange environment. I guess another interesting thing about Digital, it might have something to do with the early people, it might have something to do with the Mill. But Digital engineers and presumably other people seemed to take some sort of perverse delight

in putting up with conditions that nobody else would. I don't know if you've been out to places in California and seen how they live out there. But you just couldn't hire people from the rest of the country to live in the kind of environment that we have here. I mean, I like the Mill as an historical artifact. But ... We do put up with an awful lot for no real obvious benefit.

[LAUGHTER]

INT: When groups started moving out of the Mill for space reasons, in the mid '70s, did that split up the engineering groups and make them more remote, less integrated?

BS: Oh, absolutely. Actually it started before that. Do you remember when the PDP-10 group moved out to Marlboro, that was pretty early, by my standards. And that pretty much took them out of any contact with the rest of the world [OBSCURED BY LAUGHTER]. I mean, yes, we all worked for the same company, and every so often we'd attempt some meetings and so on, but it really --

[END OF CASSETTE SIDE 1]

BS: -- it only increased the rivalry between the groups I think.

INT: Did the rivalry grow through the mid '70s?

BS: Yeah, to some extent. It really hit when the VAX came in. The 1170 was pretty upsetting to the PDP-10 group, I guess, because we sold a lot of them and it was pretty fast and it was popular. But the 10 group was doing a good job with the KL and all that, and that was fine. The thing that really caused a lot of intergroup trouble I guess was the VAX. On the first, on the one hand, there was sort of the philosophical insult that it was based on the 11 rather than being based on the 10.

INT: It sort of proved that 11 was king.

BS: Yeah, right. People had figured out that there was a lot bigger base to build on in terms of customers, not necessarily in terms of the machine elegance or anything like that, but just in terms of selling the blasted thing, people had figured out you'd sell more of them if you made it sort of like an 11 instead of sort of like a 10. And then the fact that the machine essentially was a replacement for the 10, by any great practical purpose. Essentially the same root (?) size, pretty much just as fast, capable of having a large memory, a real full scale operating system. It was real obvious that the handwriting was on the wall for the 10 group, and that I think -- I'm not sure that was handled

in the best possible way. So then there was some bad blood there. And then to cap it off, the Venus project was started up in Tewksbury by the 11, 780 leftover group. And Gordon moved that project to Marlboro, and this left neither group happy, right? The 11 group didn't like it because their project had been stolen from them basically. It didn't seem to fit in with the Digital value system very well. And the 10 group didn't like it because they were having this strange architecture crammed down their throat. A number of them sort of left the group, or I guess some may even have left the company over this. And a few 11 people were shipped down there to work on this with them, in positions of responsibility, which probably didn't sit well with them either. So there was a lot of friction around that one. [BOTH TALK AT ONCE] And once it got down there, it got completely screwed up. This is not necessarily to say that it wouldn't have got screwed up if it had been left in Tewksbury. I'm not trying to wear a halo here about this. But, you know, life is not a controlled experiment and it did go to Marlboro and it did get screwed up down there, and that then made everybody feel even worse. So that episode I think contributed a lot to the troubles we've had between the groups this past few years. Yeah, I can remember when they moved Gus out of the Mill. We were pretty cynical about it. At that point it was the 11780 group that was

being moved. And they sort of waited until we had finished the project so we didn't have any way to get in the way, and then they told us your jobs are being moved to Tewksbury, would you like to go along? Basically in those words, that's almost a direct quote. And this did not make any friends with anybody.

INT: To the shopping center.

BS: Yeah, that joint was the pits. The Mill has its problems. But all in all, Tewksbury, at least for me, I found was a much worse environment. It's isolated. It's, the neighborhood is not wonderful. And the building is a big grey featureless monolith with no windows. It's just as awful. They should have turned it into a warehouse, not an engineering site. A complete disaster.

INT: Now that was from the 1170. After the 1170, you bounced up here.

BS: No, that was after the 780.

INT: So you were involved in the 780.

BS: Yeah. We went off on some tangent. After the 70 finished up, what we actually did, a bunch of us, was we

went to work on something called the Unicorn. Which was going to be small PDP-10. And it was going to be a small PDP-10 with a unibus, sort of an 18 bit unibus, we were going to add a couple bits, we had a whole plan for how we were going to do this. And we, we had a fair sized group working on that. Don (?) Hughes, Dave Rodgers, a whole bunch of people were there, Steve Jenkins, _____, all sorts of people. We got as far as getting boards into layout. But the Achilles heel of that project was that obviously we needed the PDP-10 software. The PDP-10 software group was in Marlboro, and they were damned if they were going to build any software -- I mean, so, obviously they didn't go to management and say we'll be damned if we're going to build any software for people in the Mill. What they did was they, oh, it would take 47 billion man-years to port drive this to the unibus and yatta, on and on and on. So eventually, the project got cancelled because we couldn't get any software for it.

INT: Well, the operating system wasn't that powerful either.

BS: Well, yeah. It was the basic PDP-10 operating system. Whether the project would've actually been a great success, I don't know. I mean it was a PDP-10, and PDP-10 by then were, you know, nobody quite realized it but they were pretty passe by then, so maybe the project

was better off cancelled. On the other hand, a year or two later, people in Marlboro went into the backroom and built a small PDP-10 with a unibus coming out of it, and didn't seem to have any trouble boarding the software. What'd they call it, the K-8, KS-10, I forget. [BOTH TALK AT ONCE] Yeah, the little one. Anyway, I don't think it sold all that wonderfully. So probably it's just as well the project got cancelled. And then we sort of waited around for people to finish the VAX architecture. That was --

INT: Was central engineering formed yet? Gordon's facility.

BS: I guess I don't actually know. I guess I don't know that.

INT: In those days, how did people select different projects or how were they selected for different projects?

BS: Things were a lot more project oriented back then and less group oriented. When a project finished sort of, the engineers were finished and they went off looking around for things to do by word of mouth, they'd hear about it, they'd talk to their buddies, or their buddies would hear that they were finishing and called them.

Managers sort of knew who the engineers were that they wanted, and called them up. People sort of spread around to other projects. Some people would stay and do the follow-on engineering for the old project and wait for a new thing to start up, and some would go off to other things. And it worked okay. There was a limited number of engineers. Most of the projects were all in the Mill.

— Marlboro quickly became an isolated culture, but most of the rest of the stuff was in the Mill, and you could swap from hardware engineering to memory to disk back to CPU again without any particular trouble.

INT: It sounds very nice.

BS: Yeah, I actually preferred it that way. Maybe it's nostalgia, but I actually preferred that system. But then as more and more groups split and moved off to remote locations, it became less and less practical to trade people from one group to another both in the geographic sense and the fact that you didn't know anybody from the other projects, you didn't know the managers, you didn't know what projects were going on. There really wasn't, was not great communication. You didn't know what everybody's reputation was. It used to be no problems in the old days. So the thing gradually became sort of more group oriented. And this may also have been in some sense what the managers wanted, right?

INT: Empire building?

BS: Yeah, to put it one way. Engineers tend to think of the product they're working on, or the project. And they tend to work in sort of episodes. Managers tend to, I mean empire building is one way to say it, but they tend to think more in terms of careers, and larger responsibilities, and they tend to view finishing something as a setback, right?

INT: Speaking of which, up to this point were there rewards when the project went well or was cancelled, or was a failure?

BS: Well, yes of course. Let's see. If the project was successful, then the reputation of the engineers who worked on it was enhanced and they'd get better offers for the next project.

INT: For example, the 1170.

BS: If the project was unsuccessful, then everybody assumed that the engineers had somehow mucked it up and they'd get less interesting offers the next time. I guess there was probably stock options in those days. I was pretty junior at first, so I didn't get any for the

first several years I was there. I think I got 100 shares for the 1170. But I guess there were some stock option things there, and I sort assume -- well, I don't know if I should say I assume ... The natural assumption would be that the size of your raise had some bearing on how well you'd done or something. I'm not quite convinced that's actually true. I think it had something to do with who your manager was or something like that.

↳ or what group

INT: But when the culture is more organized in groups, then like long term --

BS: It tends a little more toward long term stuff. It tended until this recent disruption of the fact that we're losing money hand over fist came along, it tended towards a little more stability. The raises seemed to get apportioned out more evenly. I mean it was just a little less discrimination, shall we say, in the process.

INT: Then the 780 came along.

BS: Yeah, eventually they finished spearing (?) on to the VAX architecture. I don't know. It's hard to say how things would've been different, if we had designed the VAX for speed instead of to make it look like an 11. Because ever since that architecture, we've been fighting with it. The architecture is just a pig. But anyway

they finished the architecture --

INT: The price for early success.

BS: Yeah. And they put together the group to do the 11780 hardware and they put together a software group to do the operating system. And they put all this under Bill Demmer, sort of as a unified management, Manhattan Project type of affair. And it worked. The stuff got done in a reasonably timely manner. It worked out pretty well. It was basically a pretty well run project.

INT: What was your role in it?

BS: I was in charge of the cash memory system for the processor. So I was sort of one of the five-man sort of executive committee for the hardware design. Steve Jenkins was doing the data path, Al _____ was doing the construction unit, Jed Larnik was in charge of the microcode, _____ did the memory. Bert _____ was on the unibus option. It was a pretty good sized project. Projects had been gradually creeping up in size over the years. This was getting to be a decent sized project. And we eventually made that thing work. I remember Roger Gord, who was the software engineering manager, the day after the thing was announced, he called up one of his friends at Prime Computer Company, poking fun at the guy.

And the guy at Prime said, Oh that? Oh we're not worried about that machine. And Roger Gord just laughed and laughed. But that one sold pretty well. And _____ to Tewksbury.

INT: Then what happened?

BS: The CI. The CI bus. Yeah, and the ~~Ether~~net. All that stuff. I was sort of on the periphery of the ~~Ether~~net project. Yeah, I mean, I talked to them and worked out a little bit of stuff, but I wasn't a major player in that one. And I was sort of on the periphery of the BI project. The CI project though I was more directly involved in. And then, ~~Tandem~~ (?) had just sort of started making their mark about, and everybody decided, aha, we ought to have one of these too. And so we started thinking about how to do it. And we concluded that we ought to have separate computers, not a multiprocessing shared memory system, in order to, so that we have a better chance of surviving operating systems crashes and things. That was good. And that in turn led to the idea of having a bus which would connect these computers together so that they could be separate yet still communicate at high speed. And then somehow the thought got in, well, oh, well, they all ought to be able to get at the storage, so let's put the storage on this bus too. And Don Venada sort of contributed the

idea of making it a serial bus instead of a multi-wire parallel bus. And it all started to go together. Who was there? Strecker and Barry Rubinson and myself, Dick Hudspeth and some other people started thinking about what a controller for this bus would look like. And we made it excessively complicated. And Venada and his people and I and some others started thinking about how to do the serial bus protocol on the actual wire. And that actually came out fairly well considering the circumstances. And eventually we got that thing to work.

INT: It's fabulous where the ideas came from. This one idea leading to the next.

BS: None of this stuff sort of springs full grown from the brow of Zeus. There are a bunch of ideas sort of floating around the system and some of them sort of come together and seem to fit and you assemble something out of them. I'm sure there's even more people involved in the CI, for instance, that contributed various pieces for this thing. And it turns out the way it turns out. It's not like it's some perfect --

INT: But the roots were in redundancy as opposed to parallel processing.

BS: Yes, the roots of the project were very definitely

in redundancy and keeping things from failing. And then essentially it's the same old story. Somebody finally figured out that they weren't actually going to build this software. It's the story of my career, right? There's this gigantic project supposedly building this fault tolerant system and it just sort of like deteriorated or fell apart or got too complicated for people to think about or something. But we didn't do it anyway. And then ... But fortunately I guess you could say, Dick Hudspeth and some of the BMS people did put together the system that would become Clusters. Which was again a success for Digital, although not exactly what we'd originally intended. In some sense the 1170 was a success that we hadn't quite intended. We'd originally assumed that it would be used mostly for technical computation. But it turns out its real success was banks and things like that. I guess there's something to be said for putting products out there and finding out if people can use them, rather than attempting to decide in advance with some omniscient vision exactly what the market's going to do with your stuff. Or for that matter cancelling projects just because original idea But yeah.

INT: There were almost no efforts in the commercial sector, except for Tracks, and that blew out.

BS: Yeah. I mean a lot of this stuff we sort of deliberately tried to do for the commercial sector and then failure after failure failure after failure.

INT: And then the accidental stuff. I talked with Skip Walters, and he seemed so, we'd pitch them to MI Streckers, Dupont Milliken, all those people down there. We can do this, we can create an office system. We're demoing it. Meanwhile their CEOs were up here in the Mill, visiting with Ken, and Ken says with no intention whatsoever of putting any commercial stuff in the VAX, just forget it, it's just a technical machine. And then they came back and said, What?

BS: Yeah, All-In-One, that's another sort of weird success.

INT: What a brilliant way to burn up a myth(?).

BS: Yeah. If we hadn't sold so many of those things, they would've lined the entire group up against the wall and shot them. Technologically, engineering-wise, elegance-wise, the thing is just a disaster from day one, right? But it turns out to do what some bunch of people out in the world who I just cannot understand want done. VTX is another one. I don't know whether we sell a lot of that. [BOTH TALK AT ONCE] Complete disaster.

INT: But KP(?) just smiled all the way to the bank.
When we were doing it we saw how much VAX power you were going to have to use to run it, and how much disk, and we thought that was pretty good. So much for elegance.

BS: Yeah, I don't know.

INT: So you didn't stay on the CI project, after that was done.

BS: No. We, let's see what happened there? [BOTH TALK AT ONCE] Rick Casabona and Steve Jenkins sort of took over the CI project from me. I had been experimenting with management at that point in my career. And it took only a brief experiment to convince both me and everybody around me that this was not a good idea. So they took that project over. And what happened then? I guess at that point or sometime thereafter we started thinking about Venus.

INT: Which was currently underwater in Marlboro.

BS: No, this was before that. This was when we first started thinking about doing a high-speed pipeline machine design. Judd Lenick was very involved in the early days, and myself, and Dave Rodgers, and Steve and a

bunch of other people. We started thinking about it. We had a spec written for the thing. And then Gordon Bell stole the project and moved it down to Marlboro. Some of them went with it, like Judd went down with it. And eventually the situation down there destroyed him or he destroyed himself or something. He went down to the West Coast for a while, but eventually left the company. Al Lenius left the company, and became a manager at Prime, and I understand he's been very successful. Dave Rodgers, he left and became a vice president of engineering at Sequence (?) and he's pretty successful. A lot of them have gone various places.

INT: I see them going to management. Is there really a dual track in DEC? You can keep engineering and not sucker ...

BS: It's a funny situation. Yes, of course, right? There is a dual track. There are people who keep getting promoted up the scale. The pay does get better as you get further up the scale, and there is a dual track up to the level of whatever it is, you know, senior corporate building engineer, which is close to God or something. And you can keep doing engineering sort of all up that scale. The company is not, at least as far as I can see, is not as engineering driven as it used to be. So the senior engineering people don't have as much influence on

what's going to happen as perhaps they used to. The company, there is a dual -- there's a lot of, but it's fairly managed by the managers. I guess the other piece of this thing, there is a dual career ladder. The numbers are very different. If you look at the number of people in each of the various grades, the engineering ladder looks pretty much like a pyramid, there's fewer people as you go up. You sort of expect this, right? but in fact there's about five times fewer at each level as you go up, the last time I looked. On the management side, it looks a little different. It turns out that there's a big bulge around manager and senior manager there, and that if you just sort of add up the number of managers, there's a lot more of them than there are senior engineering, even in engineering. You can sort of understand that you need a lot of managers to run 100,000 people. But it's not quite so clear that you need these many managers to run the engineering piece of 100,000 people. So, if they're looking for somewhere to get the next 8,000 people lay off, they might look at that bulge. But it's a natural thing for managers, right? Managers would like to believe that what they do is important, right? Like anybody would. So since managers control sort of the pay and procedures and privileges, they will tend to give it to people that they think are doing important things, which sure enough are managers, and it all gets fairly incestuous. Enough of that digression.

Everybody's heard it before.

INT: What happened after Venus?

[END OF CASSETTE SIDE 2]

BS: Right at this instant I'm not quite remembering what happened after Venus immediately anyway. Eventually Venus got itself all snarled up. And Gordon turned appointed Alan Kotok to go in and straighten it out. Have you met Alan?

INT: Oh yes. Fabulous.

BS: Alan is very good at a lot of things, but being high-energy, managerial and decisive is not necessarily one of them. So Gordon eventually I think got impatient with Alan --

INT: And threw Glorioso.

BS: Well, first he threw me in.

INT: Oh really?

BS: Yeah. I got a call from Gordon sort of one night the day before we were supposed to -- my wife and I were

packing to go down and visit her parents down in New Jersey somewhere and I got this call and it was suggested that I should go down and help out Alan. So Alan and I Suzanna Nathan, a personnel person, now working over in _____'s area, sort of deposited (?) ourselves and interviewed everybody in the world and started raising hell, and sort of reorganized the project a bit with Alan as the manager. And then I left. I got out of there. The handwriting was on the wall and it was in big neon letters too. And eventually Gordon sort of realized that Alan was not really what you want to manage a situation like that, and tossed in Glorioso. Actually, Steve Rothman was in there at some point, wasn't he? I'm confused. but at some point Glorioso got tossed in there, and eventually the thing did get down~~ed~~.

INT: I talked with Glorioso before he left. Everyone said that it was really screwed up, just to throw it out. So he said he talked with all the engineers and said, it looks like you people are good engineers --

BS: Yeah, the people are good engineers.

INT: And then he looked at what was going on and according to him, he said the problem was that the proposed delivery dates were just not realistic, and that they were starting to screw things up the more they tried

the unrealistic things, just causing more trouble for the next deadlines.

BS: Oh yeah.

INT: So he said, push it out, put control back, and do a real quality job on it. So eventually ...

BS: Yeah, that's sort of a common problem at DEC, and I don't know, maybe everywhere for all I know. We tend to get ourselves in this vicious circle regarding schedules. Suppose you knew everything about how to do a project. You were just -- you had it, right? And the only things you couldn't quite predict were a few things whether your lead engineer is going to get hit by a meteor, you know, there's always some random variables right? Suppose you drew the bell-shaped distribution of what date this thing was going to be finished on. Suppose you knew enough to do that. First off, you don't know enough to do that. But suppose you could, and you looked at that, and said, now what shall I propose as that date. A rational human being might pick the middle of this distribution, and say, okay, I've got a 50-50 chance of being early and a 50-50 chance of being late. You can at least come up with some theoretical reason for doing that, right? That would be one approach. A cautious person might pick somewhere fairly late in the distribution. And say I'm

pretty sure to meet my schedule. What we do at DEC is we pick somewhere way on the early edge of the distribution. Because we know that if we don't pick that real early date, is going to take the job away from you and give it to somebody who will predict that date. We also do it sort of to encourage the engineers. We want to put pressure on them, make them think harder or something, I don't know. But basically the project will get stolen from you if you don't predict an early date. I mean, the 1160 project for instance, that happened. And of course, if you predict an early date, you're pretty much guaranteed to miss it, right? We've been going on for 20 years, but DEC hasn't quite internalized this fact yet, that if you predict an early date, you're not going to make it. DEC has never quite cogitated this completely. So of course we tend to miss our dates, the management, especially the bean counter portion of it, gets all upset about unpredictability. People get under horrible pressure to meet these ridiculous dates and screw up, which gets them further behind, this never time to do it right, all the time to do it over. There's probably other extremes you could go to, but DEC has clearly decided which extreme it wants to be on. And this is another piece of DEC culture just like sort of overdesigning stuff. So while Venus was screwing around, we started the, well actually, we talk about we. What was to become the Nautilus project was started. And

that's a project with a really long gestation period.

INT: Why?

BS: The goals, technology and everything changed as it went along. It was originally started as a follow on to the 11750. And so there was an extensive study to determine what the right technology and architecture and so on to do a follow on to the 750 was, and they settled on a version of I-Squared (?) L Project, the stuff that IBM Germany invented and got used in watches for a while. It's an interesting logic family. It's a bipolar transistor but fairly high density for bipolar but not very fast. And various companies were predicting that they were going to come out with speedier versions of this, and _____ it'd be great. So the project went along -- and also this is a side note. There was a new connector figured out. I mean, we couldn't use any existing connector for this, it had to be a new one. So they specified a brand new connector that you could plug in either from the top or the side, it had a zero insertion force, nifty connector, okay? At some point management realized that this project is not moving along as fast as they wanted, so they more or less threw out the management and engineering team that was working on it and gave the carcass of the project to us. Don McGinnis, Steve Jenkins, myself, a bunch of other people.

And we started looking at this thing. And we started making it move a little bit. We were never very fond of this logic family or the connector. But we started making it move, but meanwhile Gordon had sort of realized that Venus was in trouble, and then he eventually realized it was really in trouble. And so one day he came up and said, you know, we gotta get something out pretty soon now. How'd you like to take the micro-VAX MOS chip design and lay it out in some fast TTL logic, and basically we said, no we would not, thank you. And he didn't like that answer. So we went away, spent a week or two proving that this was a bad idea. Who knows whether it really was or not, but we proved that it was. And then he said, Well, okay, how'd you like to translate your current design into ECL. And we said, Hmmm, well, that sounds a little better -- we'll go prove that's a bad idea too. So what we did was we proved that we should really do a slightly modified version of ECL. And so what is it, August '82 I guess it was, Gordon said, okay you go do that instead of what you're doing now. And we said, but what about, you know, phase one and phase two and all those phase reviews stuff. And he said, this is the phase review, just do it. So we started converting all our designs to ECL and produced what eventually became the 8800. I guess the other decree he came up with was that it had to be a dual processor. We had originally been planning a single

processor. And we would've been perfectly happy to make a single processor in that tool. And that became the 8800. And it did use -- remember the connector? Well, nobody ever really used the side entry version of it. The top entry version really doesn't work worth a damn. The thing has been the bane of our existence ever since we've started that project. And we're still using it on the BI and SMI and everything else, but it's still there. So eventually we finished off Nautilus and got that out the door. It was August something or other, I forget the year exactly.

INT: _____ disappeared then? Disappeared in April '83?

BS: Yeah, I guess it must have been '82. He disappeared in '83. I don't know what the story is behind all of that. Have you talked to Bill Hefner? He says his major ambition is to be the first DEC Vice-President to retire from old age. As far as I know there are none in line ahead of him. Something else has happened to all of them. Bill may make it. We shipped the 8800. Ken wanted a quad processor version of it. That was sort of his instigation. He came over to the power on ceremony, and said, why don't you make a quad of this. Okay. We're still not sure it was ever a great idea. I don't think we sold all that many. I think it was too late in

the life --

INT: I've heard a few people buyed it. Ultimately expandable.

BS: But the 8800 probably should have been the last ECL machine. I have a great respect and admiration for ECL. I think it's a wonderful logic family, great stuff and much more fun than this CMOS garbage. But the 8800 maybe should have been the last one. Anyway, we tried to do a follow on called Argonaut. And that go cancelled, one could say that it was because it wasn't a good idea, or one could say that it was because Aquarius was going to come along any day now, so we didn't need the Argonaut.

INT: Real soon.

BS: Yeah, real soon now. I think in retrospect the company made a fairly good-sized mistake by cancelling that project. It was not going to be a really great project, but I think they could have sold enough of them to keep us from getting into quite so much financial trouble. And we tried to do the Raven or _____ project, which was going to put the entire thing, an entire processor on one gate array. But the bipolar stuff wasn't quite dense enough to pull that off, and of course Aquarius was going to come along any day now and solve

this problem, so we didn't need that one either. And then I moved to work stations. A new adventure for me.

INT: The work station group grew out of the PDP-11 micro-11 people that were ... but then there was in the early '80s they were building work station projects in group, apparently.

BS: Yeah, Digital really messed up getting its way into work stations. You should talk to some of the old timers here, John Kirk and so on, to get a view of that. But Digital really messed up getting into work stations. Do you remember we sent Cutler out to Seattle to build the micro-VAX 1 and a bunch of stuff. If we could have gotten Cutler to build work stations, instead of trying to build the replacement for VMS, the ultimate VL operating system, there probably wouldn't be any UNIX revolution [mumbles]. But we didn't do that. Instead we let Cutler screw around with other stuff, and eventually run him out of the company, or let him run himself out of the company, depending on your point of view. And we fragmented our effort in work stations. We let a whole bunch of people try all sorts of crazy things. We had lunatic management for the longest time. It's almost like we were looking for ways to do this wrong. You couldn't even have thought of better ways --

INT: There didn't seem to be any strong vision all the way through it as to what the point of this was going to be.

BS: Digital had gotten very ... Remember Jack Shields? One of the other vice presidents who didn't retire? In this case, probably well deserved. He had been moving for years and years toward a vision of Digital as competing with IBM on their own turf so to speak. I mean, completely stupid in retrospect. But he wanted to have a large sales force that relates to the customer and that is suave and sophisticated like all IBMs. He wanted to have a large service organization. He wanted to do lots of stuff for the customer to make them feel wanted. He wanted to be just like IBM, this is in terms of the field organization, and in terms of the products we could put them, you know, it would be 100 percent reliable and wonderful and he didn't care how much it cost, just like IBM, right? And this whole sort of stuff had pervaded the company pretty well by this point, unfortunately. And so people had sort of lost track of the idea of building a personal, small fast machine that would just sit on your desk and get some work done. And instead they tended to think of grand visions of entire companies networked together, and oversized machines in oversized boxes and high overhead structures in the field -- the whole complex of stuff that went with that. And that, I

don't think Digital could have put together a decent work station group or a decent work station software package, I think it was just impossible. The culture had gotten too screwed up to do that. And I might even argue that we still can't. We seem to have at least gotten the idea now but it's a little late. We made a lot of mistakes around that. That and the whole Prism, Hudson _____, complete fiasco. We asked the wrong questions, we got the wrong answers. We went around asking customers how much of a cost performance advantage would it take to interest you in converting away from VMS and VAX and into Risk and Unix. They said they wouldn't do it for 2 to 1, they probably would for 4 to 1. So we went out and gave them a 4 to 1 excuse for converting, right? We made machines that had low performance, we delivered them late, we made sure that they were very expensive, and we forced our customers just right out of our camp. Someday this will be a classic in a marketing textbook -- how not to do it.

INT: This is really the outgrowth of the 11 group.

BS: Well, sort of. There were a bunch of 11 groups, remember. There was the 1120 group, which sort of went off and did the 1105 and the 1140 and the 1160 and then sort of fell apart. There was the 1145 group which sort of _____ [backfired?] in some sense. In some ways this

group is sort of, I don't know, you might think of it as an outgrowth of the PDP-8 group almost. But you should probably get that straight from Kirk or somebody who's been around here for a while. I'm sort of the new kid on the block and I shouldn't be talking about their history. Those pieces of the various 11 groups and the 8 groups are _____.

INT: Is there still a strong network of engineers? I mean physical network that helps tie the place [UNCLEAR].

BS: Yeah, the electronic mail has really changed things quite a bit. There still is a network of engineers. Despite everything. It's not quite as all inclusive as it once was. The people out in Colorado for instance are kind of isolated from the people in CPU for instance, not like the old days, when you could go to their reviews, and they could go to yours. The West Coast work station group is sort of by choice a little isolated from the East Coast. There's other things like this. But there still does exist a -- an informal underground network of engineers, and people still do change from one project to another, so cross-pollination (?) and get to know people that way, that sort of thing. At one point there were attempts to make it a little more formal. Gordon was doing sort of engineering seminars and coconsulting engineer dinners and stuff like that in an attempt to

foster that cohesiveness. But when he left, essentially that had fallen apart pretty much.

INT: And now we have business units.

BS: Yeah, now we have ABUs, IBUs, PBUs, CBUs, anything with a U on it. A hundred of them or so, I don't know how many.

[END OF CASSETTE SIDE]

