

[NOTE: SUPNIK IS TERRIBLY UNDERMIKED.]

BS: I started at Digital in July of 1977. I had been working in a small company called Applied Data Research on and off for the previous 13 years, either full time or part time when I was in school, in various positions, _____ in Massachusetts. I'd been involved with DEC's equipment from the very beginning. Applied Data was one of the pioneers of embeddedness _____. We did some of the first analytical implementation systems. When I was still in college I had an absolutely marvelous experience of the first automated television station came out, KMOX-TV out in St. Louis. I had six weeks on site while they were bringing the new control system over. I had the run of this brand new television facility. The greatest toy a 19 year old ever had. And it was absolutely marvelous, despite the fact that when I got there, I found I couldn't rent a car because I was too young. Here I was the company representative on site of Applied Data Research, and I couldn't get to the customer because I was underage. So I had intensive involvement with PDP-8s and PDP-11s, primarily on the real-time, scientific side. And in 1977, when I wanted to get involved in different, broader areas of computing, it was very natural to think of Digital as the place I wanted to go. So I had interviewed at DEC, and I had already two job offers, one from the EMS group and one from the

software machines group. And I was on my way to some interview, and in the middle, when I got lost in the basement. Things were pretty casual in those days. Go find your next appointment, it's over there somewhere. I was wandering around the basement when I ran into a friend of mine, Barry Rubinson. Oh, you're here interviewing? You really ought to go to talk to the people in engineering. They need some software expertise. You could help them out a lot. So indeed I went to talk to Mike Riggels, who ran the storage advanced development group. We struck a deal, which was I would bring software expertise to this problem of making intelligent disk controller, a very novel concept in '77. And they would tolerate and try to cure my massive ignorance of how hardware worked, since I hadn't a clue. I was a self-taught programmer with a little college background, but _____. So starting off in the basement, _____, and then moving on to the lofty heights of 3-6 (?), [REST UNCLEAR]. We worked on what became eventually the HSC and the _____. _____ the architectural problems of a high-performance, intelligent disk _____. And just about the time the advanced development was finished the whole group was imported (?) to Colorado. Sort of a first-rate expansion out of the _____, Colorado Springs. And I wasn't really going to do that. I was going to stay in living in the number one target for nuclear attack in the United

States. So, I set out to find something else to do.

BL: Could I go back to the disk and ask you to do -- I think you brought up a theme which was that it was a very hardware oriented engineering group.

BS: That's right.

BL: And it's interesting to see where the software expertise if any, ever started entering into the mainstream. To what extent was there the software group, to what extent did you deal with the other software groups in the Mill, which were apparently really small.

BL: Oh, we were totally abandoned, renegades, both within hardware engineering and within software engineering. Barry Rubinson came up with the key group (?). _____ and some other people, _____ business group. And we literally wrote everything ourselves. We wrote the whole software environment for the HSC, which was pretty complicated, since it was a multi-transfer (?) control system. And we developed a driver-work (?) interface into the operating system, _____ test bands (?). And this was entirely appropriate for an advanced development. There really wasn't a model yet of how this was all going to fit together as a system. And it did by pressing Mike

Riggels part to assert that the application of microprocessor technology and the intelligence of the disks, plus the _____ would result in other products. And _____ come together in a systems sense, was the great, a great leap forward in our product strategy _____ orchestrated over the 1979 _____, which most people associated with the company conversion of the bands (?). There were actually several other steps. One of them was converting on Internet strategy. Where the CI, the BI, Internet were all really one collected, adopted the company's interconnect architecture for the early and mid '80s.

BL: _____ the Unibus concept.

BS: [UNCLEAR] So the energy, the interaction in this advanced development team was really interesting. It was a clash of cultures between the software people and the hardware was just phenomenal. The software people, particularly software people at that time, coming out of the renegade tradition of software engineering, _____ . One of the things that's characteristic of software people who were trained in the late '60s and early '70s is that they keep irregular hours. They hate coming in in the morning, stay late, and _____. And the reason was that computers were scarce and one got machine time at three the morning, if

they got it at all. The rest of the time it was used by _____. Well, the hardware engineers, they came in at 8:15, and they went home at 5, and trying, everything was by the numbers, _____, and the software was messier, _____. And rooming (?) them together in a cultural sense was a major accomplishment. And then getting them to contribute jointly to an architecture was next, was even a step beyond that. Because initially it was, gee, I [UNCLEAR PHRASE], work like this. Okay, well here's a set of hardware things that will do what _____. The very first model of the HSC advanced development prototype, which was called _____, _____ was a complete failure. The whole model we had for how to put together an intelligent _____ system didn't work. We made a couple of basic mistakes around how often the intelligence got involved, what kinds of operations had to occur, _____. But from that learning came a model of how to properly structure _____, and that became the basis of the HSC. And at that time, because the hardware and software groups were off in their special spheres, saying, _____. All right, we now understand the dynamics. What can hardware do, what can software do. Which has the primacy in their areas? Software supplies policy, hardware supplies _____.

But for me all this progress occurred under the shadow of

the breakup of _____. [REST UNCLEAR] So I personally had become fairly intrigued with this synthesizing role between hardware, software, architecture, systems, and customer requirements. Sensing that there was a job to be done there that didn't fit any of the conventional disciplines. And I had this feeling to reinforce, I've got to talk with the personnel person, the mandatory chat about _____. And they said, there's very clear choices. You can be hardware, software, you can be management, you can be technical. I want to be a hardware-software-systems-technical-manager. And what turned out to be marvelous about Digital is, that I was able to create the role of precisely seeing my needs and my talents. And the company was willing to get what it did. [UNCLEAR PHRASE] So, when I was looking around, this was in late 1978 now, _____, again I set up this, some people I knew, one of my earliest _____, before the early '70s, I got hired as a consultant by Digital, and answered the question whether the second, third and fourth models of the PDP-11 family were _____. The 1120 had come out and at that point the 11 was, and we'd built the 1105, the 1135, and the 1145, and kind of --

BL: [UNCLEAR]

BS: Yeah, and there are some slight differences, and there were no tools for testing whether they conformed to the architecture stuff. So they brought me in and said, go figure it out. And I was able to show that in fact _____ . The 1105 and the 1145 were 11s, and the 1135 wasn't. Actually it was the 1140. And it was due to microcode bugs that hadn't been detected by diagnostics. _____ . So I knew the people who did small systems _____ the problem. And the one that led _____ in, when I started looking for something to do, I connected with them, and again the company was kind of a quid pro quo with us. We're in this small systems business which is going in the direction _____ . We've got this playpen CLSI (?) project, and the PDP-11s, and we needed someone to analyze what's going to happen and what we ought to do, and _____ . So I went to work for a small systems group, which was at that time under Dick Clayton (?). And spent about six months trying to analyze trends in small computers, trends in technology, what our customers wanted. And by the spring of 1979, had a proposal that said we really need to aim to build a microprocessor-based PDP-11 that's as fast as the biggest one we'd ever built, which was the 70, that's the future of the product line. You'll never need to build a really big PDP-11 out of other kinds of technology again, that's what the VAX is for. And in the usual tradition of DEC, the proposal was _____ , very

nicely. Now it's yours. Congratulations, go ahead, the LSI project is going to go off and implement this.

_____ called the J-11 _____. [UNCLEAR PHRASE] The last time we ever went outside the company to collaborate on a microprocessor project. At the time, there was a sort of prevailing industry wisdom that you used C-MOS to make _____ chips. And so, starting with a small set of collaborators who had worked on the PDP-11, _____, back into the learning by doing mode. Great fun. And Portner was in Florida. I learned more about the _____ in Florida, than I ever learned in a lifetime. _____ restaurants and other.

BL: Now you microcoded it.

BS: No, I actually didn't. For the microcoding, I'd worked on the microcoding in the _____ model. And I found it sort of interesting, but for the J-11 I was very much the project manager of the group. And there was a microprogrammer was _____. And saw that group for, into 1981, and then, again, I was sort of overtaken by events. All these recommendations for the J-11 was _____ parallel sort of recommendation that had been put together that Digital had to take semiconductor technology a lot more seriously for its competitive future. This work was spearheaded by, I think, it was _____. And that was really Scorpio. And they

recommended a major programming investment that means that Digital would now be self-sufficient in semiconductor technology _____, particularly for microprocessors. It was a four-point program. Start developing basic technology ourselves. Up to that point, it was largely _____. Develop a design capability for really complex projects. Tools. _____ fabrication. And lastly design the box (?). So it was altogether a major uplift (?). [UNCLEAR PHRASE] And it was based on the perception that it wasn't to ride (?) much more than _____. And down the road, semiconductors and microprocessors were going to take apart all _____. [UNCLEAR PHRASE] And the program was put into effect in 1981. So this was going along, and there was this gradual racheting (?) _____. It was impeded by the fact that our semiconductor effort was divided among _____ organizations. To have semiconductor manufacturing, which was _____ 1974, I think, when the vendor of the original microprocessor, the LSI-11, went Chapter 11 right in the middle of the _____. And Digital went out and bought a semiconductor plant in Western Massachusetts, from MOS Tech, and turned it into an LSI-11 factory in four months. [UNCLEAR PHRASE] There was a group that was focused on the designs of microprocessors, which was in the small systems area. And lastly there was the Scorpio team focused on trying to make the great leap. And the people who over that

responsibility, Joe _____, manufacturing, _____,
[UNCLEAR PHRASE]. They got together on their own in
early 1981 and said, this is crazy. We have to have just
one engineering (?) team (?). And they ran a model
processing organizational merger. Very quiet, but very
professionally, _____. [UNCLEAR PHRASE] And in
the course of this they went out and sort of did a
classic assessment of material, _____, filling
out a questionnaire. And on my questionnaire, it said,
what do you do, what are your skills, what do you want to
do, what are long term aspirations. And I wrote down
that _____ advanced development group, just concentrate
_____. Maybe I should do that again some
day. [UNCLEAR PHRASE] I must have been the only person
who circled the box, because when they organized, they
said, well, we're going to have an advanced development
group, a microprocessing group, a _____ group, _____ and
resources. And Bob, you're it. So in April of 1981,
when semiconductor engineering group was _____,
I was told on literally a week's notice that I was about
to take over the advanced development function of the
_____ advanced chip design, advanced method of tools
design and process technology. [UNCLEAR PHRASE] And I
was going to go in managing a team of 12 people who doing
all these pieces of the J-11, _____. This was
quite a leap. Thanks, _____ was organized.
[UNCLEAR PHRASE] And we were off there, and now, this

was just yet another instance of being placed in a sort of learn or die situation. The people under me in the process technology group were, they were experts. [BOTH TALK AT ONCE] It was really pretty interesting. When I first met the head of that group, a marvelous, marvelous woman named Ruth Rawa, became a semiconducting engineer in the 1950s, when women in engineering were _____, and she stuck with it ever since.

BL: And she's still _____?

BS: She retired about 18 months ago. Rawa. I believe she's still in Chelmsford. And certainly the semiconductor manufacturing people would know. [UNCLEAR PHRASE] Not even that. And she comes in and she sort of looks me over for a while, and then she said, You know, Bob, when I started in semiconductors you were in kindergarten. You're absolutely right. Well, it was that job then that actually, finally brought me to _____ what I wanted to _____.

Because now, there were twin challenges. One challenge was to demonstrate that I had _____, to [UNCLEAR PHRASE]. And the second was, back then, _____ that our technology _____. Two things came out of that. One of which was the company's fundamental semiconductor strategy. And the second was the company's microprocessing _____. And they were really

related. But they had what I would call, I don't know, _____ . And for those who would like to see history as an inevitable march of social, political, economic forces, there is no evidence to support it in my experience. Let's take the _____. As a matter of the _____ got set up. We had recruited a lot of expertise from RCA, because RCA was gradually dropping out of the _____, semiconductor companies. They were well _____ in the area. And RCA had a long tradition of C-MOS development. So in the fact of the competing claims for C-MOS and _____, we had some people who were predisposed towards C-MOS. And that together with the J-11 experience, which had been sort of really one of the first efforts that was on a high-performance _____ processor, converged to a conclusion that C-MOS was the wave of the future.

BL: And yet no one else made that.

BS: No one else made that leap of the time, right. Just as in 1980, _____ concluded that multiple levels of metal interconnect was _____. Jeff Kalb (?) was brought in from _____ to become the head of the LSI group in 1982. [UNCLEAR PHRASE] Basically overseeing the merger of the organization _____. But Jeff was a specialist in the _____. And so he walked in a Valley Talk (?) guy, looked at what we were doing, and

said, you're crazy. Multiple levels of metal, this is berserk (?). The way you're doing it will never work. So in fact the very first management challenge I had within the process group, was to defend our process strategy from my _____. And eventually, after some really tough-minded, technically _____, expertise, and other ... [UNCLEAR PHRASE]. And a real technical specialist cutting across five levels of management if necessary together we put all four, in a couple of days, we came out of it with _____. Multiple levels of metal were really necessary, and that one way, one extra feature we had brought in, _____. So we got onto multiple levels of metal early and we avoided a significant manufacturing problem. But it was process strategy. Our meetings and our extrapolating processes let us, _____, that C-MOS was going to drive out M-MOS, and far more radical, that MOS in general was going to drive out VCL (?). [UNCLEAR PHRASES] Now this was, mind you, when Digital was just starting down the road investing in _____. [UNCLEAR PHRASE] And I think the only reason it's worked out is at that time, nobody took semiconducting seriously. This was all an internal program.

BL: [UNCLEAR]

BS: If we had been taken more seriously, we probably

couldn't have come to such a stark conclusion. Because the pressure to compromise would have been really overwhelming. But being a playpen, _____.
[UNCLEAR PHRASES] So we got out of the Echo business, and really focused down on C-MOS, and lay out a process strategy from the very --

[END OF CASSETTE SIDE 1. QUALITY OF SOUND IMPROVES TREMENDOUSLY ON SIDE 2.]

BS: -- process strategy that was very clean and very simple and very focused, which was necessary given how relatively little Digital could invest against _____ or Monroe (?). We said, we're going to do C-MOS, we're going to do multiple layers of metal, and we're going to focus everything on building very fast microprocessors. And we haven't deviated from that in ten years.

BL: Still there now, huh?

BS: Still there now. And we've built -- we're in the development of our sixth generation of internal process technology. We do one about every two years. The other side of the equation was the micro processor strategy. The effort we were going to do of VAX was very much in the model of, the VAX is what it is and okay, we got to build a VAX. It turns out that to do that we had to

cover 8 chips. VAXs were very big and they had lots of micro code. So we were not really building a low end system in silicon technology, we were building something like a mid-range system. It turned out we were pretty much going to make a smaller machine lit by the power of the 780, it was like a washing machine and it wasn't very inexpensive(?). And we really hadn't made the conceptual breakthrough about firming the VAX to the real low end of the business and make it a mass market barometer. Micro processors were these toys. They were used by IBM PC, what does that have to do with VAX? Interestingly enough the idea for a micro VAX originated outside of Digital. This is one of our guilty secrets. We were approached in the fall of 1981 by Xylog(?). Xylog(?) was one of the successful vendors of 16 bit micro processors but they had concluded that, against the established Intel or Motorola competition, they couldn't introduce a 32 bit architecture of their own. They needed to adopt more of the ones that were there, what better one than the VAX? Alright we'll talk. [UNINTELLIGIBLE]. . . And in they came with this revolutionary suggestion, - of course we're going to put it all on one chip, we're just going to drop the parts of the instruction set that are too big. Life goes on. - What do you mean you're going to drop the parts of the instructional set? - Oh yeah, we'll just emulate them with the remaining parts of the instruction set. So, off we go. We were cogitating this

for the rest of 1984. The conclusion was this was really a nifty idea; to get a VAX that was in the same cost and form factor of these micro processors, would permit us to drive the VAX architecture as a standard throughout the industry. So, guilty secret number two is, - we never intended to do the program ourselves. Our intent was to license a semi-conductor company to build a good one, not a small one like Xylog but a big one, to build VAXs and make it a mass market _____. In the early winter of 1982 we were going around to semi-conductors offering them, begging them to get VAX licenses. And the crazy crazies all refused. They all said that their own proprietary micro processors were too important to give up.

BL: Based on the Intel model?

BS: That's right. So Motorola and National turned down licenses to the VAX architecture because it would have meant they had to stop on the eight thousand twenty and for the National thirty-two thousand. We came down to a decision point in May of 1982, where, either we were going to have give up on the idea and base our low end products on one of these other micros. And we were actually looking at whether to port VMS towards these other micros or we were going to have to do it internally. And internal notion was never on the part and wasn't part of the strategy. The strategy was to

make it a standard but we internally had nothing to do with it. I was the technical consultant to this whole process, I was part of the usual Xylog(?) evaluation team and I was working with Remata(?) who was spearheading the marketing side of this. I sat in on all these negotiations. I listened to all these excuses and saw the technical estimates that said we could take the ____ 6. I'm sitting there scratching my head saying, - this isn't right. I could do this. I know what building blocks we have and our capabilities. Starting about March of that year, I just started drawing sketches on pieces of paper for how to put a chip together that would meet the goals. I started researching VAX microcode to figure out whether it would fit and I developed a model that said we could go over to our big ____ and basically steal some interesting pieces and ____ some interesting pieces and call them together _____. So I started proposing this in an informally way and Roy kept saying, - no, this doesn't need to go _____ and finally came in to May of 1982 and the business team said, - we're going to adopt an industry _____. They sent the usual memo saying, - we've reached the conclusion and the answer is. And I fired off - at that time I was just an engineering manager I wasn't even a consulting engineer - saying bullshit. We can do this ourselves. This is not rocket science. Within two days, again, Jeff Kalb(?) and Steve and worked with Jack and

Ken. They had an answer. It was - we'll do it ourselves. We are going to make a VAX. You're in charge. Now, I was still running this advanced development group. I had process technology and CAD technology and _____ advanced design team. We're _____ . I became the project manager. I recruited Digital's finest. I used to say that Dan Dyrofoam(?) was the best _____ designer on the East Coast but I'm now convinced he's the best _____ designer in the country. He was there. He had just returned to Digital after having been an independent consultant for a couple of years. There he was fortuitously. It wasn't planned. He happened to be there once we were starting this up. And he volunteered. Rich Widick, who had been software engineer in the RISC group and in our CAD group, walked in and said, - _____ I think I could help here and he turned out to be the architect, the person who actually put the chip together at the lowest level and has done the architecture, our leading edge designs. I looked for the guiding hand, the great invisible force and I don't see it. Very casually, it could have gone the other way. The team was put together. The other thing that was really revolutionary, the first time we were going to try and do a single chip processor and the other thing is, we committed ourselves to an insane implementation schedule. We said we're going to get the design done in fifteen months, from July of '82 to October of '83. It turns out

it was 18. But it was still significantly less than anything that had been attempted before.

BL: Do you remember why you chose such an aggressive number?

BS: Because we had to. Again, we were working from models of, - what's the competition doing? And we were projecting out where Intel and Motorola would be, when they would get the 32 _____. The goal was to be a vendor in the chip market. So, this is what our competition was saying. Now, they missed by years. We missed by three months. It was strictly market driven and didn't know any better. Here's the ____ serendipity again. Originally when we put together this architecture for the one chip we sub-set the VAX in a couple of ways; we cut out a bunch of instructions, we also simplified the memory management architecture in such a way as we thought that VMS would still run but it would not be suitable for running really large scale systems because it was intended to be a low end chip. While this effort was going on, the other person involved in defining the sub-set, in getting this whole _____ was Dave Cutler and Dick Husfetter(?).

Dave and Dick represented VMS. Just after we cut the definition loose early in 1982 and as we were talking to chip vendors, Dave went out to Seattle, started DEC

Alaska(?). He said, - I'm going to do a cursor system, fast time to market technology, and I'm going to write a real time operating system called the _____. Well as Dave got into the problems _____ software developer, he started pushing back on some of the sub-setting, particularly the sub-setting in _____ management. At first we restored one feature and then, it would have to have been relatively early in 1983, restored the whole thing, the whole area of _____ architecture. And at that point, we didn't know it but we had just again, taken a major _____ because now there was nothing it couldn't run because it could emulate its missing instructions. It was now suddenly a full VAX. We still called it a micro VAX but there was nothing missing from a customer's point of view, there was just speed issue, performance issues. That, as it turns out, would really put the cat with the pigeons because at about that same time in 1983 Jesse Lipkin who was in the PDP-11 business got very excited about this. I can now make a machine that looks just like my small 11s except I'll put a VAX in it. What? We can't do that. This little box, it's only two feet high and nine inches wide, it's got a queue bus in it? Yeah, got a queue bus in it. _____ Kathy Morse got very excited about the VMS team, fortunately Dick Husfetter did the same _____ and _____ but Kathy stepped right in, picked up the whole. It was all crazy. We're going to run this thing on a tiny, little box. At

that time, the smallest machine we made which was the 11-730 was the size of a dish washer. We're going to run queue bus peripherals, I mean those things are good for _____.

BL: Aside from that the entire plan was to put it on BI.

BS: Yes, that was the official company strategy.

BL: The company has this strategy, - if this chip comes out it will support a BI based system.

BS: So for a long time, nobody took it seriously, which is fortunately. Jeff Kalb(?) was a stalwart supporter and everybody left us alone. We missed the October '83 design completion date. We had been too optimistic particularly about the integration phase of the design. It was really too bad because we had chosen a date to finish the design that was coincident with Jeff Kalb's(?) birthday. We missed it and it took three months. We were devastated at the time and yet retrospectively, it's still one of the most incredible design _____.

The design was finished at the end of January 1984. When we fabricated it, we had some very severe initial debug problems because the design was so tuned we were trying to balance space and speed. One of the circuits we had pushed too hard, it was just too risky a design. It

didn't work over a broad range of process parameters. We were able to debug the chip in pieces but it took us until the summer of 1984 before we first got a part that was fully functional. That was July. By August we were running VMS. Coincidentally, about that same time, official interest in the project in the company went to zero. Beginning with fiscal '84, the forecast went to zero. The marketing people said, - no, we'll never want a small VAX. The owners of the established systems said, it couldn't possibly perform the way we said it would because it has to be a 780 to do that kind of performance. We were about to announce the 8600 which was a mere two years late at that point. That was sort of a dear point, I've gotten used to this since then but this was the first time this had happened to me, this phenomenon that six months before it's shipped, the system _____. I'm told that six months before the 780 was shipped there was actually an official motion at the executive committee to cancel it in the interest in Digital's future. That could be hearsay.

BL: Then you started debugging it while queue bus idea and you're stuck in the middle.

BS: Once the system started coming together, everything turned around. As we started having these demonstrable tiny systems, as we started measuring them and show that

they really do deliver 780 like performance in many ways, interest really began to grow. But the other thing that happened was, - people came to the realization that we built a VAX, the chip is a FAX, it embodies the essence of Digital's hardware technology and the door is slammed shut on the licensing. It was about that time that the door was slammed shut on the licensing, that _____ BI _____. The company was turning away from its OEM model doing business in a big way. To this day, that's the part of the program I most regret the loss of. The fall of '84 was another transition also for me which was, after three years of running sem-conductors, Steve Ticher wanted to try something else. [UNINTELLIGIBLE] system to work stations where he started up our Palo Alto venture [UNINTELLIGIBLE]. And I moved over from _____ the advanced development program to running the micro processor group and I took the microVAX with me. I have had more fun doing this than I could ever have imagined. Project _____ written micro code. I did write _____ . From soup to nuts, I had finally learned that this technology is fun, it's marvelous. We were inventing and innovating and solving problems every day.

BL: I love the way you throw way, - [INAUDIBLE AND UNINTELLIGIBLE]. That seems to be the one reason why you're so well respected in the group. Everyone sees

this manager and yet he does the most difficult and technical work at the same time.

BS: It's not the most difficult piece.

BL: Nobody knows how.

BS: Yes, we're a dying band of brothers, we, micro coders in Digital. Me and Richie Larry are, Aileen Sandberg, not many left.

BL: Strange, you have the reputation of being the top micro coder ever.

BS: I do a lot of it. If you consider that I have written micro code for every single of our VAX microprocessors. In the case of micro VAX, C-VAX and N-VAX. Micro VAX involved C-VAX and N-VAX I wrote the first drafts for C-VAX and eventually did essentially all of it. And Roger and I did a lot of contribution. Even from the 11 the multi chip VAX, I did some of the final auditing and performance too. I like it It's fun. Anyone who is trained in programming the PDP-8 will appreciate their _____.

BL: Why is that?BS: It's banging(?) and it's optimization at really this fine grade level. Riche Larry is probably the best PDP-8 programmer of all time

had a very simple theorem, PDP-8 had a basic unit of addressing called, a page which had 128 instructions and Richie Larry's theorem was that any 129 word PDP-8 program can be re-written in the 128 words. And this feeling is instantly recursive. And that's what micro code is about, is a fixed resource, 1600 words in the case of micro Vax, got to fit in the entire VAX instruction set minus the 30 we left out. Previously the smallest implementation had been four or five thousand words. It's a challenge.

BL: How do you go about doing it? BS: In some sense the principals are well established, the VAX instruction set doesn't permit you much flexibility. Essentially for me it was a lot of _____, _____
_____. That is I would get a draft down, simple straightforward implementations wouldn't fit, they were too big. in micro VAX I got to about 75 or 80 percent of it and I was out of space and so then it would just be, - alright, are there better algorithms? Are there better consolidations? What are the time-space trade-offs? It was a relentless _____ process. I was making something like 20 or 30 edit passes, three micro codes a day. A lot of the things on it _____ week-ends and stuff like that. The other thing is, it relates to one of the things that, by the strategic level which is the ability to take up massive data and massive details that

fundamentally _____ and correlate them. So for me the way the micro code lives in my head is as a series of pieces from which there are, at least initially, no formal structural relationships on which I'm going to impose structural relationships, interdependency, order. What starts out as a nice, clean unoptimized code ends up as nice, optimized, small code. It was just a lot of fun. The other thing was, is also a principal of good engineering is, I'm very pragmatic about micro coding which means I have absolutely no shame about looking at other examples and borrowing what looks good. So I had religiously studied the 780 and 750 micro code and the development of micro _____, the Nautilus family just as the micro VAX chip, it's fundamental architecture came from the multi chip of the 11____ and many of its actual silicone sections come from that. The difference is, it's knit together as one _____ whereas the other is knit together as seven or eight. And the difference in application was between having to build a modest size system and to build a true micro processor. In substance, the rest is history. Jesse's micro VAX was the most successful start-up of a business in Digital's history. In fact the micro VAX business is a bigger and faster start-up than SUN. It started in 1985 and in 1991 it is a bigger business than SUN, - I'm sorry, 1982. It certainly came off the blocks, incredible task and completely reoriented the company's view of the VAX. The

VAX was a computer run system. The company was geared up to sell the '86 _____, 375 thousand dollar CPU. Now we trot in this 19 thousand 9 hundred and 95 complete system. And slowed those powered(?) people. VAX power to the people. So the process strategy was one piece and the micro processor paradigm was the second piece and _____ together was a coordinated process product strategy that first got formulated in 1982 when micro VAX was launched. What this did was, and I think it really was the first of the companies to look out, in this case five years, and say, the growth of the technology and the growth of our understanding about designing good chips, makes possible the following capabilities: we shall define products at specific points and we will use this as a planning matrix for driving not only our process and chip designs but for our system partners to drive their system design. The first one looked out from '82 to '86, it was micro VAX, C-VAX and module(?). We updated it in 1985 in Rogell(?) _____ index. I think they did another pass after I left _____ third overlapping spectrum. But it has set the systems development agenda for Digital Equipment Corporation, for everything below the very highest machine. First of all the machine the micro VAX were based on the micro processors. In 1985 as we were beginning to work on the next one, C-VAX and showing people the model of a continuous technology evolution at this phenomenal rate of 50 percent performance

improvement. Brian Elison(?) in the mid range got hooked on the idea, this makes multi processing affordable. He laid out a system concept for a multi processor, an affordable multi processor. Then he and I working together, worked through the principal of what does it mean to make the system immune to the changes in technology? The notion of divorcing the system from a specific processor, making it possible to replace the processor, so-called _____. That resulted in the Calypso platform on the 6000 Series, which is the most successful individual computer series in Digital's history. _____ fifth generation of upgrades to that, had the fastest start in revenues of any product in Digital's history. By moving from the advanced development group to the micro processing group at that time in '84 I got the pleasure again of going from the formulation side of strategy to the implementation side of strategy. From '84 to '88 I was the manager of the micro processor group and I just had a ball. That has got to have been the best working experience at Digital because here was an enormously large scale program, not only in terms of technology and products, but in terms of fundamentally altering Digital's view about how computers were built.

BL: I think you're going to have to throw me out now. You've given me so many perspectives and observations. One thing I don't want to ...

[END OF SIDE 2].

[END OF INTERVIEW WITH BOB SUPNIK].