



Oral History of Robert Supnik

Interviewed by:
Gardner Hendrie

Recorded May 2, 2017
Carlisle, MA

CHM Reference number: X8215.2017

© 2017 Computer History Museum

Hendrie: We have with us today, Bob Supnik, who has very generously agreed to do an Oral History for the Computer History Museum. So thank you very much, Bob!

Supnik: You're welcome! It's always a pleasure, because the History Museum is absolutely one of my favorite institutions on the planet!

Hendrie: Aw, now that's really-- now those are very kind words. I think maybe what I'd like to start out with is if you could talk a little bit about your family background. Your mother, your father, where you grew up, any siblings you had, just to understand what your environment was like when you were really young.

Supnik: Okay, I was born into an upper middle class family that lived in Long Island. I was actually born in Brooklyn, but they moved out to the suburbs of Rockville Center when I was two months old. My father was in the corrugated box business, a family business started by my grandfather. And my mother was, you know, a homemaker, housewife, socialite, whatever. In 1955, we moved to upstate New York to Syracuse, so my father could run a branch box plant up there. And he decided that he wanted to go independent, have his own business. So he bought that box plant out from the family and became the President and CEO. But it wasn't economic scale, and it wasn't really, really set up to be an independent plant. Plus we were heading at the time into the very severe recession of 1959, which was one of the most severe post-war recessions. And the family went broke, completely broke. We lost everything. By 1961, my parents had lost all their savings, they had lost the house, and they were now faced with how were they going to put three children through college. I have two older sisters, Lee, who is six years older than I am; and Gale, who is two years older than I am. Well, they had a couple of advantages, one is my mother had a background in history, and my father was a lawyer by training. So they both went back to school. My mother got a Master's degree in Education and became a high school teacher. And my father, after struggling to reestablish himself -- because he was by then in his late 40s-- became a bankruptcy lawyer, pretty appropriate given that we had been through bankruptcy. And so we three siblings did go to college. My older sister, Lee, went to Cornell with some scholarship help. My sister Gale went to the University of State of New York at Buffalo, very low tuition at that time. And I went to MIT on a National Merit Scholarship. So we were able to squeeze through despite the family's low income.

Hendrie: How old were you when the bankruptcy event occurred?

Supnik: Let's see, I was 12. And I remember it very distinctly, because we had had quite a nice house. I had had a very elaborate Lionel train setup. And one of the things we did to raise money was to hold a yard sale, and I sold off all my trains. So you know, that made an impression on a 13-years-old at the time. And I think it was actually a very good thing, because it taught me a lesson about being resilient and bouncing back when things turn sour. My mother taught until she retired. My father worked as a lawyer till he was 76. The fact that we moved from a fairly nice house into a duplex. They were not fazed by that,

and they saw to it that we three children stayed on a path of being educated and moving forward. That's kind of the hidden advantages, I think, of a background like mine. I can say, "Well, I started from nothing! My family had no money," etcetera. What I had was a strong propensity to be educated, a strong drive from my parents to be someone who read and who thought and who tried to be creative. And so, the notion of being quote, "self-made," unquote, doesn't appeal to me, because I fully acknowledge all the help I had from my parents, my sisters, my relatives, institutions, and so on.

Hendrie: Okay. So you went through the public-- you kids were-- the kids were all educated in public schools.

Supnik: In public schools in Syracuse.

Hendrie: Did you have any teachers that sort of influenced you back in that period? Did you like Science and Math? What were your most-- what were the subjects you liked the best, and did you have any particularly influential teachers?

Supnik: I was definitely a science geek, a math geek. But also a reading and literature geek. So probably the-- in terms of influence I would point to a Latin teacher that I had in high school, who had-- was very charismatic. Mrs. Meechum [ph?]. And what she gave me was the structure of understanding language, which has enabled me to be articulate and also quite a good editor to this day. The downside is, having learned Latin, which is a language you only read, I have never to this day managed to master the speaking of a foreign language. So you know, it had its price. The other one was probably the tenth grade Math teacher I had for geometry, Mrs. Blanchard, who really understood that there wasn't enough in the standard curriculum for me and worked hard to set up alternatives. And the school had been very supportive. I had taken 11th grade Algebra when I was in 9th grade as an experimental program, much to my sister's chagrin who was also in the 11th grade. And Mrs. Blanchard and the Math Department at my high school really saw to it that there was enrichment available to the point where in my senior year, having run through the high school curriculum, I, and the couple of others who had gone through the program, were packed off to Syracuse University to take a year of calculus. And as it turned out, a semester of Computer Programming.

Hendrie: Ohhhhh.

Supnik: Ohhhh, yeah.

Hendrie: Oh, yes!

Supnik: And in general, I found that the teachers I had in high school in the Science and Math curriculum, well, whether it was Mr. Spadafora for Physics, or the biology teacher, really knew how to awaken a love of learning, you know, right in the public schools. So that sort of carried me forward in terms of where I would go educationally, because I would go major in math when I went to MIT. But the career, the career comes out of family. My older sister, Lee, was going to Cornell, and she was being courted by an electrical engineering student named Dave Waks. And so Dave would come to visit, and here's, you know, pesky younger brother hanging around, getting in the way. So. Dave either humored me or wanted to work with me, but he started telling me about his computers, because he was a computer person. This is back in 1961/'62 [actually, 1959/1960]. And he would show me manuals and show me pictures. One summer when I went to visit them after they got married down in New York. He took me into where he was working at the Courant Institute, and he showed me the CDC 6600 being constructed. Because in that day and age, Seymour Cray's machines were literally debugged on site wire-by-wire. You know, they would measure it, they would say, "Oh, that's a nanosecond loss," snip, rewrap. And I was just gobsmacked! I mean, there was a fascination and a romance with them that sort of began to obscure everything else.

So, in my last semester in high school, I was able to take a Fortran programming course at Syracuse University. And I had access to their state of the art 1620, where everything was on card decks, and they had set up the computer room so that the-- they had a really nice wheeled chair that you rolled across to get the card deck, you rolled back and you just, you know, pushed yourself across the room to get these masses of cards where they needed to go. And as I remember, what I did is I put the accounting for our class treasury-- I was the treasurer-- I put that on the computer as my first program. And that was probably more dangerous than useful. I'm not sure that the books ever did balance, but I was kind of committed. So when I graduated from high school, Dave was by then working at a company called Applied Data Research, in Princeton, New Jersey, which was one of the very first independent software vendors. They probably would have called themselves a software services shop, but then they started making things they could sell over and over again for the IBM 360. So I worked there as an intern at a time when nobody knew what college interns should be doing. So the first summer, I drew flow charts, hung curtains, ran errands for the office manager, who was a formidable person. And at the end of the summer, as my reward, I was allowed to go into the IBM data center in Philadelphia and write and run an actual Fortran program, which drove the huge Calcomp plotter. And I think I made one of those circles with a million cross-points that give you a very elaborate star. So then I was at MIT. MIT had no Computer Science program. And--

Hendrie: Well, excuse me, can we roll back?

Supnik: Sure!

Hendrie: I want to understand what places you thought of possibly going to college.

Supnik: Ah!

Hendrie: What were your choices and how did you end up at MIT particularly?

Supnik: Oh, okay. Well, it was clear I was going to go to someplace with a decent technology background. So I applied to, let's see, I applied to MIT; I applied to RPI; I applied to Syracuse University as a safety school. And in particular, I had won a full scholarship there for a contest on citizenship when I was a senior. I applied to my father's alma mater, Lafayette; and at my mother's insistence, I applied to Harvard, although I had no intention of ever going there. <laughter> And I got in everywhere, except Harvard. And I went to MIT, because it had the reputation for science and math. And I never visited there before I went. I had never seen it! I had no idea what it was like! I didn't do a campus tour. I just read the literature and said, "All right. That sounds okay."

Hendrie: The best-- you know, that was the best school you got into!

Supnik: And we have a visitor here, Ms. Pumpkin [a cat]. She will stay out frame, I hope. She may contribute vocally, if I don't pay some attention to her. So because I'd won a National Merit Scholarship, I was able to afford the tuition. All I had to pay was room and board. I mean, tuition was the astounding sum of \$1,700 in those days.

Hendrie: What year is this?

Supnik: It's 1964.

Hendrie: 1964, all right.

Supnik: 1964. Went up to \$1,900 the next year, and there were riots! "\$1,900! Too damn much!" was the slogan. I believe it is now \$50,000? Maybe 60?

Hendrie: I'm not sure.

Supnik: Not sure. So--

Hendrie: So now you get to MIT, and are you still interested in math as your major?

Supnik: Absolutely! I think I'm going to be a mathematician.

Hendrie: Okay.

Supnik: And that the computers are just something to do for fun, because they're interesting.

Hendrie: Now, also at MIT, they have a Railroad Club that-- and you had mentioned that you had had-- been interested in model railroads at one point when you were younger. Did that tempt you at all?

Supnik: I never got there. For some reason, I never knew about it until I had graduated! <laughter> I have no idea why. I mean, it was very prominent, well-known, and I started interacting with DEC while I was still in college, and of course, Alan Kotok was the guru of the Model Railroad Club. But I never got there.

Hendrie: Okay.

Supnik: The way MIT was structured, at that time, was the first freshman year, there were required courses in math, physics, chemistry, and humanities, which was considered a full load. So to do something with computers, I took the introductory computer course as an overload. And so that first semester I sort of don't remember breathing much.

Hendrie: Okay.

Supnik: For one thing, a full year of Syracuse University Calculus had caused me to place out of the first semester at MIT of Calculus, and I shouldn't have, because MIT covered a full year of Syracuse University Calculus in about six weeks. <laughter> So I was-- in second semester calculus, I was really in trouble! I don't recall doing almost anything for fun that first year!

Hendrie: Now what computer did you work on?

Supnik: Well, the introductory course, which was called 6.41-- if I remember correctly-- taught you a higher level language called MAD, the Michigan Algorithm Decoder, which was a sort of a bastardized version of ALGOL, but considerably simpler than Fortran. And they had constructed a-- the MIT CS Department had constructed a very fast turnaround compiler, load-and-go system for the MAD language. And they-- that's what you did your problems on. So they had a series of programming problems, you punched them on cards, you ran them in a batch, and you got your corrections.

Hendrie: Do you remember what machine they--

Supnik: [An IBM] 7094.

Hendrie: 7094.

Supnik: It was the 7094.

Hendrie: Yep, okay.

Supnik: In the Computer Center.

Hendrie: All right, good! And then the second semester I took 6.47, which was the second course, and then the following year when the chemistry requirement was no longer there and I could drop down to a reasonable load, I took the penultimate course, the third one and last one that they had, 6.251, Introduction to Systems Programming, which taught you how to do basic assembly language programming on the 7094 around a model assembler, that was very simple, that you put ornamentation on. That would be the end of my formal computer science training. <laughter> Because--

Hendrie: That's all there was!

Supnik: That's all there was! And so I banged around while I was studying math, and I really accelerated through the math curriculum. I finished the first year of graduate Algebra in my third year. I banged around in whatever department had a computer course to offer. I took one in the Economics Department. I took one in the Naval Architecture Department. They were mostly just programming courses that gave you the opportunity to use their machines. One problem I returned to repeatedly was trying to make a game theory solver. I had read a book called "The Complete Strategist." And I'd been fascinated by game theory. And at the time the book was written, the deterministic matrix inversion method hadn't been invented. So I would write Fortran programs with the various techniques they offered to see if I could replicate the solutions they had in the book. And eventually, of course, it was cracked by better mathematics and I wrote a program to do that. I think I did most of that work on a 1620 in the Economics Department. And what was amazing about this is the machines were not in high demand. The Computer Science [actually, Electrical Engineering – there was no CS department] machine was always very heavily used, and they had just begun to create the time sharing system called CTSS, and that was always in use. But the machines in the Humanities [actually, Economics – no computers in Humanities] Departments, it was very easy to get time on. But in the meantime, I was working summers at Applied Data Research. And the second summer I was there, having raised my pay from \$1.05, which was the minimum wage, to \$1.10 an hour, they decided I was as useful as anyone else, and so they put me to programming the front end of a program they were trying to invent called Autoflow.

Hendrie: Ohhhh, very interesting. Yes.

Supnik: This is the days when flowcharts were everything. So they were trying to create a program that would make it possible to automatically generate flowcharts as documentation from annotated comments, or what we would call structured comments.

Hendrie: Mm hm.

Supnik: And I was working on a machine called the RCA Spectra 70/20, which was RCA's knockoff of a IBM 360/30, but further cost reduced, and with, you know, a very, very bizarre instruction set that was not compatible. So I worked on that for most of the summer, and I got it running. I remember that machine because it had the first really high-speed card reader I had ever seen. You know, one of these things where it blew air through the cards, and then sucked them in with vacuum. And they would always come out slightly curled and bent at the edges. And after a while, if you didn't reproduce the deck, that curling would become an airfoil, and the cards would go shooting off the end of the machine out into hyperspace! <laughter> So it was great fun! Well, they were so astonished that I'd actually accomplished this, and before the end of the summer-- because I was very fluent with assembly language, I really liked doing it-- that they said, "Why don't you try and move this to the 360. There's this new thing from IBM called the 360." So I got sent in, I was 18-years-old at this point, to the IBM data center in Philadelphia with a purchase order and instructions to port my system over to the 360. So in walks this 18-year-old kid, you know, in a white shirt and a tie, because it's IBM, and they're looking at me like, "Oh, what are you doing here?!" Well, the 360 at that point had no real operating system. The assembler was a card deck. The loader was a card deck. Everything was card decks again. So I ported it over by the end of the summer, and the President of the company was so impressed that he threw a party for me. And he got me so drunk that I was ill for three days! Now remember, it was, back then, 18 was the legal drinking age.

Hendrie: Yeah, wasn't anything illegal that you did.

Supnik: Wasn't anything illegal, and my brother-in-law, Dave, was there to drive me safely back to a place where I would-- could recuperate. But whoo! That was an experience. So in my junior year at MIT, was when I first encountered DEC. Dave wanted to move Applied Data, or at least part of Applied Data, into the business of what we would call embedded systems, or OEM systems, where you would use a minicomputer attached to something, and it would control it. So in order to get the capital together, he persuaded the government to fund a research project by a computer scientist named Anatol W. Holt, fondly known as Tolly, on Petri nets. And Petri nets were the kind of thing that would appeal to DARPA, and so DARPA gave some money. And the way this was going to be done was on a graphics processor attached to a PDP-8, called a 338, a beautiful machine [designed] by a guy named Dave Brown. You know, the PDP-8 was this big. The graphics processor was *that* big. And was far more intelligent. But there was no programming system for the PDP-8, except paper tape. So Dave conceived the idea that we would use--

Hendrie: This is your brother-in-law.

Supnik: My brother-in-law.

Hendrie: Yeah, okay.

Supnik: Dave Waks, that we would use a PDP-7, which had a very primitive mass storage operating system called DECSys, basically just enough to run a compiler or an assembler and load a program, as a front-end for this PDP-8. We would do all the program development on the PDP-7. He would write a simulator for the PDP-8 on the 7 to do additional debugging, and then you would load code from the 7 to the 8 over an inter-processor link, which had never existed before either! So, you know, DEC was game for anything in those days, and their Computer Special Systems Division agreed to put this configuration together, provided that we from Applied Data took responsibility for writing diagnostics and proving that it worked. So in the spring of 1966, the configuration was about ready. So Dave came up from New Jersey, and he picked me up, and we went out to Maynard. I was out from Cambridge to work for a couple of days on the factory floor of Digital Equipment Corporation. Now in those days, DEC was like one floor, maybe two floors of what is now Building 5 of the Old Mill. And it had not been renovated in any way, so the floors were slippery with all the sheep lanolin, and the brick was grimy. There was hardly any security. We, you know, walked in, we got met by our salesman, he said, "There's your machine. Here's some badges so you can get access whenever you need it, and go for it." And together, we worked out a diagnostic to test this inter-processor link, which turned out to be the prototype of every inter-processor link DEC ever built after that, and--

Hendrie: Who designed the inter-processor link for you? Somebody at DEC?

Supnik: Yeah, DEC's Computer Special Systems Division.

Hendrie: Ah, okay, but you don't know who?

Supnik: I don't remember who.

Hendrie: Yeah, right.

Supnik: But it became the basis of things like the DR-11 parallel interface family and everything else. So that summer, the machine was delivered to Applied Data, and we realized almost immediately, you know, Dave realized almost immediately that it was lacking the most fundamental tools you needed to do this work. So, for example, we had bought a card reader because clearly with a single user machine, there

wouldn't be enough time for people to bang on the teletype at 10 characters per second to edit the programs. So the idea was to do initial program editing on the card reader, read it in, and then just do updates online. Or interactively rather. There was no software to do that. There was no software that would read a card file, a card deck and write it to DECTape, which was the storage medium of the day, all 145 kilowords [KW]. So that turned out to be my first programming assignment on this project. I wrote the card to DECTape routine with all the necessary character translations and bizarre stuff like that. And other people, because it was a very young crew, wrote basic DECTape routines and a cross-assembler, and Dave wrote the simulator. And by the end of the summer we had the programming system together that would permit the PDP-8 to be programmed on the 7, partially debugged on the 7, and loaded from the 8. And by this point, I was enamored enough of the programming that I was going to Princeton at every school holiday and staying with him and working at Applied Data. And MIT had long school holidays, because it was six weeks between semesters.

Hendrie: Yeah.

Supnik: And so I was becoming more and more enamored of computers, but progressively less enamored of math, because I discovered a fundamental thing about myself, sort of after I'd run all the way through the first year graduate curriculum, which was, I didn't have the kind of mind for abstract mathematics. I couldn't see things. I could work theorems as a logic problem, you know, "This follows from that, that follows from that. You could apply this axiom or that one." So it was a logic puzzle, work it through. But if you really want to do advanced mathematics, you need to be able to sort of visualize the structures in multiple dimensions. You know, not only in three dimensions, but above three dimensions. I couldn't see it. So at the end of my third year at MIT, having got enough credits for the degree, I finally decided to do something else. And there still being no computer curriculum, I decided to do something as radically different from math as I could, and I went into history. Now my mother was a Social Studies teacher, so I had a background in it.

Hendrie: And you said your mother had a history degree.

Supnik: Yes!

Hendrie: Yes!

Supnik: You know, the apple does not fall far from the tree, as they say. So I switched over to the Humanities Department at MIT and enrolled what was called the Double Major program, which means I would stay a fifth year, get two degrees, two bachelor's degrees. And I studied history for the last two years as well as--

Hendrie: Yeah, okay.

Supnik: So I was working on that, and I wrote an undergraduate thesis about <sighs>-- I want to say World War II diplomacy over the Polish problem. Yes.

Hendrie: Ah, okay! I was going to say what was your area of study that you specialized in?

Supnik: It was 20th Century European History, and what I didn't realize about it was it was going to make me terribly depressed! It was a terribly depressing subject!

Hendrie: Yes.

Supnik: But I applied for graduate school. I had every intention of getting a PhD in History, and I got into Brandeis. But in the meantime, I was working every summer in computers, and now, every vacation on computers at Applied Data, doing more and more interesting projects. So, in 1968, for example, we did one of the very first embedded computers to automate a television station. So, in the old days, the way television stations worked is there was a big control panel with faders, and buttons, and sources, and things. And there was a script. And somebody who was really good at this would be manipulating the sources to get things on the air at the precise time. And that included rolling the film and tape machines the precise amount of seconds in advance to get them up to speed. Well, a company called AMP, AMP in Pennsylvania, wanted to break in to this industry. So, they created a switcher (as it was called) that could be controlled by a computer. And they teamed up with Applied Data to create this system. And our experimental first customer was KMOX-TV in St. Louis, which was building a brand new facility. So, we created-- Dave created the hardware. He was an electrical engineer, so he designed the interface. This included a CRT display that was just purely a tube that was being driven directly out of the PDP-8. The PDP-8 literally drew every pixel as its null job. And I designed and wrote the software. I wrote-- created a small real-time operating system and all the tasks and so forth to keep this all running with a program or schedule that was input on cards and ran in this 8K word PDP-8. Yeah.

Hendrie: Wow.

Supnik: So, I wrote it-- most of it in one summer. Someone else filled in some of the gaps during the following year. And when I came back the following summer, it was time to install it. So, there was going to be a six-week shakedown period, while-- to get it running before the station actually was moved over. So, I flew out to St. Louis and got to the airport and showed them my travel paper. Here, cat, come here <whistles> Yeah.

Hendrie: Thank you.

Supnik: That's good. He's afraid of loud noises. So, you can just scare him away by clapping your hands or something. Got to St. Louis airport with the salesman, showed them my reservation for the rental car, and they said, "We can't rent to you." I said, "Why not?" He said, "You're twenty-one. We don't rent to anybody under twenty-five." This was like Hertz or one of the big ones. So, great, now what do I do? I'm going to be here for six weeks with no transportation. So, the salesman, bless his heart, said, "I'll rent it. You drive it. And don't tell anyone. And don't get in an accident." So, I had six weeks with this television station basically to myself and the setup crew. Best toy a boy ever had, let me tell you. And we had-- they got it shaken down. What was interesting is that because the PDP-8 had no mass storage on it, everything, every fix I made to the program had to be done by hand as a patch, hand assembled, patched in. And then occasionally I would fly back to Princeton, reassemble the whole program, bring it back. Well, got it running before I went off to school that year. That was also the summer I got married. So, it was quite an exciting period and just one of the many interesting and strange projects that Applied Data did trying to put computers with things. One of my colleagues used a PDP-15 to automate a planetarium. Later, he and I would use a PDP-11 to automate a pallet system in a warehouse. It was all brand new. It was--

Hendrie: This was an era when I remember those machines. We called them real-time control computers.

Supnik: Yeah.

Hendrie: Before anybody invented the term minicomputer. That's what you were doing.

Supnik: So, I was at Brandeis still working at Applied Data. But now I was working more with the branch of Applied Data known as Mass Computer Associates in Wakefield. And they were primarily a government contracting house. They did research projects for DARPA and so on, and so very different work. There was a project writing a compiler for the DDP-516, or 716, at a time when Bill Poduska was still involved. There was a task to interface a CDC 6600 to the ARPANET via a PDP-11. That was exciting. And-- but what I was finding is I was more and more engaged with the computers and less and less engaged with the history. The computers made me excited. The history made me depressed. And finally, in the middle-- the beginning of '72, I decided I've really got this the wrong way around. And computers should be the career. And history should be the hobby. So, that's what I did. I went to Applied Data and said could I have a full-time job in Boston. And-- cat. I got my consolation Masters and dropped out of history. Although, I still read it extensively, and-- of course, my interest in computer history is--

Hendrie: Yes, the combination of the two could be pretty interesting.

Supnik: It's palpable. The-- so, Applied Data set us up in a tiny little office in Boston in what was the old Somerset Hotel. It had just been renovated as office buildings [office space]. And there were like five, and

eventually seven, of us doing embedded systems and simulation work. The PDP-7 had been outgrown. So, Dave decided to take advantage of the-- DEC's first timesharing system, the PDP-10, and write a simulator for the 10 that could be used by multiple people at the same time called MIMIC. And MIMIC was designed by myself, Len Fehskens [recently deceased], and Michael McCarthy, with Mike doing the I/O system. I did the command interpreter. And Len wrote some of the key simulators. I wrote some of the other simulators. And Applied Data used it as an internal tool, of course, for doing its embedded computer work. But we also sold it to other people doing embedded work via service bureaus. So, it got put up on TYMSHARE, and Compushare, and a number of other PDP-10 timesharing bureaus of the day. And eventually this led us to cross paths with DEC again, because DEC wanted to use it as a design tool, something we had never considered. They had seen that Len Fehskens had written this brilliantly accurate 1145 simulator, I mean to the cycle, to the nanosecond.

Hendrie: Wow.

Supnik: It could run any diagnostic. And they said, "You know, I guess this has the structure we need for doing register transfer level simulation. We think we need this for this project called the LSI-11." So, they bought MIMIC. And I think Len left Applied Data to work for them and created a very detailed LSI-11 micromachine simulator that was used to write the microcode for the project. So, that was another interaction with DEC. And at about the same time, I got a consulting engagement with DEC to figure out whether the second set of PDP-11s after the 1120 were really PDP-11s or not.

Hendrie: Ooh.

Supnik: Because--

Hendrie: That's pretty interesting.

Supnik: Yeah. because they had no tools, no architectural verifier, nothing. They had no idea whether they were compatible or not at the very detailed level that would be required. So, they brought me in. They said, "Here. Here are these three machines. Go break them." So, okay, "Give me all your software." So, I got all their software. And I ran it. I ran variations of it. And I finally proved that the 11/05 and the 11/45 were 11s. But at the time, the 11/40 was not. It had some obscure bug. And the way I found this was by firing up the timeshare basic system called RSTS [V4]. And RSTS had a spawning feature where you could spawn a shell, and then spawn a shell, and spawn a shell, spawn a shell. And you just increased the stress on the machines the more layers you added. And eventually, it just-- on the 40, it just crashed. The others, it just tied itself in knots. But the 40 actually crashed. So, that's how-- and then I had to backtrack what had gone wrong. That was a-- that was really fun. And I met a lot of interesting people at DEC that way, like Steve Teicher, who was in charge of the 11/05 project, and Dick Clayton who was running the 11/45. So, I got to know some--

Hendrie: So, what was the-- what was-- where was the fault in the 11/40? Do you remember?

Supnik: I'm not sure I can remember after all these years.

Hendrie: Okay, just curious.

Supnik: I'm pretty sure it was a condition code not being set properly on some instruction. That pattern-- Yeah, okay. So, it wasn't a hard fix. It was just a microcode fix. But at this point, the main project that I was working on for Applied Data was to automate a mass spectrometer for a company called AEI Ltd. in the UK. And they had some new very fancy mass spectrometers with the ability to use two magnetic beams instead of one so that you could have a reference sample in one and the actual sample in the other and not have to cross contaminate them, very--

Hendrie: Very interesting, yes.

Supnik: Very exciting. And so, they wanted a mass storage-based system that could handle these machines and handle fairly capacious amounts of data. So, we went looking for a real-time operating system with disk capabilities and multitasking because they needed-- they wanted to have the ability to have reports generated in real time. So, you're collecting data. You're analyzing it. You're printing it. You have to have multiple tasks.

Hendrie: Yes.

Supnik: And DG [Data General] had RDOS, which pretty much fit the bill. And DEC had nothing.

Hendrie: Really?

Supnik: Had nothing. That year, in '72, RSX-11M was still two years away. And RT-11 had not grown any sort of real-time scheduling capability.

Hendrie: Okay.

Supnik: So, I mean I went to talk to them and talked to the software developers at DEC. And the software developer I talked to was a gentleman named Dave Cutler.

Hendrie: Oh, yes.

Supnik: Yeah, this will be the first of several interactions with Dave that, how shall we say, do not go well.

Hendrie: Lots of interactions with Dave that don't go well.

Supnik: Well, we were pretty skeptical of what he was talking about. He was trying to describe M and what he had in mind for M. But because it wasn't there, and we needed to get started, we ultimately had to tell him it wouldn't do. And certainly, the older operating systems like 11A, B, and C, or D just didn't work. They were either huge or slow. So, we went off on our DG equipment and created this very large-scale software system. One of the things that I really liked about it was that the front end was literally built. You put in a descriptor of the capabilities you wanted. I want this report, and that report, and that report, but not that report. I'm going to use this kind of machine [mass spectrometer]. And it built an executable front end from a library for you to run.

Hendrie: Yeah, so it wasn't a-- it wasn't just a huge kluge that you had to have the whole thing in there. You could-- fundamentally, it was a self-configuring to just what you wanted to do.

Supnik: Right and we had to do that because we didn't have enough core.

Hendrie: Exactly.

Supnik: We had the--

Hendrie: And time did not allow you to be rolling things in and off back and forth off the drum

Supnik: Right, not off that magnificently slow Diablo-33 with its 2.5 megabytes of capacity.

Hendrie: Right.

Supnik: Or maybe it was 5 megabytes by then. So, we had that project. We had the planetarium project. We had the pallet project. We had lots of interesting embedded contracts. We also had some real trouble because, at some point, the city of Boston had discovered there was this computer programming shop in Boston proper. And at that point, the notion of being kind to technology firms hadn't crossed their mind. So, they came in and said, "Oh, look at all these computers you have. That's personal property. We're going to property tax you on them." And we were moved out of Boston within two weeks. We were gone. We moved over to Cambridge to a kind of a slummy office upstairs from a restaurant in Central Square at a time when Central Square was a pretty groddy place. And my wife and I were living just off Central Square in a triple decker. So, it was very convenient. I walked to work and continued to do that until I got

mugged, at which point, we picked up stakes and moved to Arlington. But around '75/'76 Applied Data, which was beginning to suffer a bit financially from competition, said, "You know, it doesn't make sense to have an office in Wakefield and an office in Cambridge. You guys are moving to Wakefield." So, we moved to Wakefield, set up our little embedded computing shop on the first floor. But it wasn't really a good match between the management structure of Mass Computer Associates, which was very much a research and government contracting firm, and us. And I didn't get along very well with the guy who ran it.

Hendrie: <Laughs>

Supnik: And at some point, we decided it was better if I was elsewhere. So, in the spring of '77, I decided to go change jobs. Now, by this point, I had developed quite a few contacts at Digital.

Hendrie: You knew a lot of people there. Yeah because you'd-- yeah, they'd done some special systems for you. And the--

Supnik: I had been consulting for them on a couple of occasions.

Hendrie: Yeah.

Supnik: And in fact, in 1975, a fellow named Barry Rubinson tried to hire me into the PDP-15 group. Fortunately-- I sort of was not ready then. And the PDP-15 group was like disbanded within six months. So, it was a good thing. So, when I put in a resume to DEC as a programmer, as a software engineer and by now a section manager, so potentially a supervisor, and so I got the usual-- what I would call all the usual interviews. I got an interview with the CAD group. I got an interview with the VMS group and an interview with another software engineering group. So, I went in to talk to people. And-- including Dave Cutler, again. That went much better this time. But while I was being interviewed-- DEC was very informal in those days. And he was, "Well, you know you're here. Go find your way over there. And it's down the stairway and cut across these three buildings and go up to the clock tower. And if you're lucky you'll find it." Well, I was-- I went down. I went too far. I ended up in the basement of the mill. And I was wandering around hopelessly lost. I mean I went by some ancient power-generating equipment that was still there.

Hendrie: From its mill days.

Supnik: From its mill days.

Hendrie: Yeah.

Supnik: And then Barry Rubinson found me, this fellow that had tried to hire me into the PDP-15 group, stumbled across me and said, "Why are you here?" "Well, I'm interviewing." He said, "Oh, if you're here interviewing, you really should go talk to the disk engineering group. They need people with software experience because they're trying to do some things with software now." And he shanghaied me out of the interview loop right over to the office of the vice president of advanced development for storage, Mike Riggle.

Hendrie: Okay.

Supnik: And Mike described his vision for intelligent controllers for storage. And because it was going to be software-- have an extensive software component, he needed to build a programming team. And he needed somebody who knew how that worked, particularly around real-time control systems. And I'd never studied hardware in my life. I had no idea what any of this was. And I said, "Sure, sounds great. I'm in."

Hendrie: Yeah, well you'll figured out-- you figured out anything you needed to know about the hardware you were working with when you did these real-time systems.

Supnik: Yeah, of course. And the-- so, I went to work in mass storage engineering, the first software engineer, other than Barry, they had ever seen. And immediately, I understood there was going to be some degree of culture clash because software engineers, in those days, worked whatever hours they had to [in order] to get machine time. When we were-- had only the PDP-7 back at Applied Data, it got scheduled to three in the morning because there was only the one terminal to use. And there were a lot of people queuing for it. So, the hardware engineers, on the other hand, I thought they had joined a union. They worked eight to five. And man, they came in before me. And they were gone by the time I was sort of warming up. What's going on here? But they were extremely generous with their time and their knowledge. I particularly remember a gentleman named Pete McClean, never ever gave me a hard time when I would ask the most fundamental stupid questions about how a disk worked or what the issues were in latency, and seek time, and algorithms, and so forth. And so, we put together a prototype of what was called the New Disk Subsystem. And it was going to have a PDP-11 as plain controller because we knew that PDP-11 chips were coming and some hardware to talk-- I think we used just the standard RK-11 interface to talk to the PDP-11. And then we used a parallel interface with DMA to talk to the quote host.

Hendrie: Yeah.

Supnik: And Barry and I wrote the software together and started to measure it. And it was a catastrophe. It didn't perform at all. And the reason it didn't perform is that the PDP-11 was getting involved with every buffer transfer. It had to—it got interrupted when data came in off the disk. Then it had to program the

DMA channel. And then it had to be interrupted when the transfer was finished. And it all made sense in theory. In particular, the idea was to have the intelligent controller take over the error correcting process so that we could use error correcting algorithms that were too complicated and sometimes too lengthy for a driver. We're moving from the old forty-four-bit code that could correct up to, I think, an eleven-bit burst to a hundred and forty-four-bit Reed-Solomon code, which was much more powerful. And-- but it didn't work. So, I'm sitting there wondering what to do when I was introduced to a gentleman named Bill Roberts.

Hendrie: Okay.

Supnik: Bill Roberts had been at Western Digital and was therefore one of the principles in the LSI-11 project. And I was friends with the people in the small computer group who had done the LSI-11 because that was Steve Teicher's group. And Bill got invited over. And I described the problem to him. And he said, "Well, of course. You can't have a general-purpose computer in the data path. You have to separate the control path and the data path if this is ever going to work". And so, he started thinking about the problem. And he proposed that he could build a single board controller that used a timed shared 2901, which was a bit slice machine microcontroller, to run the two data paths. And it would have enough intelligence left over to play control processor, enough spare cycles. A really, really, really elegant design, I was just about gobsommed. And so, we hired Bill to create this thing. In the meantime, DEC, in its first major, as we would say, overseas expansion-- it was out of state, decided the whole storage group was going to move to Colorado Springs. Well, we had just bought this house in Carlisle. And there was like no way on God's Earth that I was going to Colorado Springs. So, I was now under a time limit. So, I wrote up my reports on the New Disk Subsystem, saying it didn't work. And--

Hendrie: And why.

Supnik: And why. And then when we got Bill's model running, it did work. It actually did deliver the performance and still was able to do the error correction invisibly behind the back of the host processor.

Hendrie: Okay.

Supnik: And it could-- we could put seek optimization and command queues and things that weren't really well known to be in disk controllers in those days.

Hendrie: Okay, it had enough cycles to do that. Yeah, well it was a very fast chip set.

Supnik: Very fast chip set, and it had a generous amount of EPROM on it for this purpose. But we had made a slight and fundamental error. We had not secured the exclusive intellectual property rights. Either

DEC was naïve in those days, or Bill had offered us a bargain if we had joint rights. So, DEC took that and built what became the UDA50 intelligent disk controller. And the design also went into what became the HSC50, where the PDP-11 was there, but only as a policy processor. Bill, on the other hand, took that design and founded Emulex and started producing knockoffs of every DEC controller on the planet by microprogramming that same basic design to talk to industry standard disks, which DEC's controllers wouldn't do. So--

Hendrie: And of course--

Supnik: Yeah, and he made a killing.

Hendrie: I mean every OEM is going to say forget the DEC controller and disk.

Supnik: Yes, right. Go buy raw CDC disks and mark them up, and there it goes. So, it was a-- from a business point of view, it did not work out as well as it might have. But boy was it a-- it was a cool design. I mean it was just a wonderful design. So, I was looking around for work and talking to people I knew at DEC and wondering whether I should go back to software after all this time, a year away. And when I went to talk to Steve Teicher's people, they said, "Oh, well if you don't have anything to do, why don't you come work for us. We're trying to design some more chips. We could use some help." [Like I said, this is the vocal one. The other one is--][another interruption by the cats] So, I joined them as a systems analyst. Basically, I was going to look at the problem of computers-- of what strategy this nascent chip design group should follow moving forward. [Don't be friendly. Cat, come over here]. And at the end of six months, I produced a report. And it said look, we should build one more PDP-11. We'd [already] built two chip sets. Let's build one more that's as good as the best PDP-11 that has ever been built, the 11/70. And let's switch our focus to designing an LSI VAX, because the VAX was out as of 1978. And it was clear that was the future. And Steve Teicher said, "Oh, that's a good idea. You're in charge of the 11. Go build it." So, I now found myself, with no prior experience in chip design, CPU design, or anything else--

Hendrie: Or logic design.

Supnik: Or logic design.

Hendrie: Yes.

Supnik: Building the last PDP-11 for-- or nearly the last PDP-11 for DEC - while the bulk of the resources were off on trying to put together the LSI VAX project.

Hendrie: Yeah.

Supnik: Well, we didn't have enough resources internally to even think of doing two projects in parallel. So, it was clear that the 11, as not being the strategic focus, would have to find partners. So, the PDP-8 group had been working with a company called Harris Semiconductor and had just produced a chip called the 6120, which was quite a nice PDP-8 and became the basis of all the small word processing style 8s they built. DEC made one, two, and three [the DECmates I, II, and III]. And they recommended Harris. So, we went to talk to Harris. And said we want to build this PDP-11, but we have a concern that CMOS, which was at that time not the processor choice, everybody used NMOS for speed, CMOS is not up to it. They said, "Ah, we think we can do it because we have this brand new process. It's going to be at four microns." Wow, today things are at one one-hundredth of a micron. Yeah, four microns and two levels of poly silicon for interconnect. Today, they are at like nine levels of real metal.

Hendrie: Yes. Okay.

Supnik: And so--

Hendrie: This is Harris, right?

Supnik: This is Harris.

Hendrie: Yes.

Supnik: So, we entered into a partnership whereby we would do the architecture, microcode, and logic design. They would do the circuit design, and the layout, and the fabrication. And we soon discovered that they had been more confident than they should have been, particularly in their circuit design capabilities. I cannot recall the guy's name. They had this big bluff designer, a true Florida-- dare I say it, a true Florida redneck.

Hendrie: Yes.

Supnik: Very large, thick neck, nice guy, but-- very confident. And he would come back with these designs. And we would look at them and say, "They don't work." "What do you mean they don't work?" I said, "They're going to go metastable." Now, metastability was not a new concept at that point, but it had only really become prominent in computer design around 1970/1972 when the 11/45 project almost died because of metastability issues.

Hendrie: Is that right? See, I never heard that story.

Supnik: Yeah.

Hendrie: I mean, obviously, I'd known about metastable--

Supnik: The 45--

Hendrie: Flops forever.

Supnik: The 45 was DEC's fastest machine ever. It was a hundred and fifty nanosecond microcycle.

Hendrie: Okay.

Supnik: And it was using, I guess, the H series of TTL chips. And it simply was unreliable. And eventually, a tech named Bob Stewart, who would go on to become probably DEC's premier processor designer for larger systems, uncovered and proved the problem and also, therefore, could show them how to fix it by stacking flops.

Hendrie: Yes.

Supnik: Dual ranking flops. But no, it was very close around there. It hadn't occurred on earlier DEC machines because the micro cycle was slow enough that things just settled before the next cycle was clocked or the next clock tick hit. So the circuit designer at Harris simply did not believe us that at the 200 nanosecond micro cycle we were aiming for, this was going to be a serious problem.

Hendrie: Mm-hm. He'd never experienced it.

Supnik: He'd never experienced it. So we invited him up, and Bob Stewart sat down with him for a whole day. And now, Bob Stewart was a guy with a very fiery temper. He did not suffer fools, and he had a bright red bushy beard, and a very sarcastic manner, but he was as polite as could be with the fellow from Harris. When the Harris guy would put up a circuit, Bob would say, "Okay. Here's the metastability problem. It's going to surface right here." And the guy would try another fix to it, and Bob was, like, "No. You've moved it over here." And they chased this circuit around the board, I think, for an entire afternoon, at the end of which, the people from Harris said, "Okay. We get it. Everything's going to have to be dual ranked." But it was in every sense, for both Harris and for DEC, a bridge too far, in terms of ambition.

Because we were still doing things entirely by hand, hand schematics, hand SPICE simulations, and so on.

Hendrie: I need to interrupt at some point. Can you go back and tell me, quickly go over, the history of the LSI-11 microprocessors?

Supnik: Sure.

Hendrie: So you said, "This is going to be the last one." But...

Supnik: Well, you missed the first two.

Hendrie: Yeah, I missed the first two, yes.

Supnik: Yeah, my fault.

Hendrie: Well, no. You weren't involved. So I understand.

Supnik: Actually, I was, peripherally.

Hendrie: Okay.

Supnik: So the LSI-11, which would completely revolutionize DEC's approach to processor design, was, actually, the brainchild of Western Digital and of Bill Roberts. And they approached DEC saying, "We think this 8-bit microprocessor with its programmable extensible Microm], the microcode ROM, could, with a little customization, be an 11. And DEC looked at it and did some trial coding and said, "Yep. We think you're right. Let's do a project together." So for this first project, for the LSI-11, which began in '72 and was delivered in '75, DEC only did the microcode and a bit of the architecture work, and then all the board level design and so on. And the microcode was done by a gentleman named Dwayne Dickhut, and as he proudly liked to say, he was also an MIT graduate - from the Milwaukee Institute of Technology. A wonderful guy, great sense of humor. And it was indeed a success. It was a pretty tricky chip to work with. The board design, which was done by Rich Olsen, had to go through a lot of hoops to get the clocking right and so forth. But there had never been a PDP-11 that small or that cheap.

Hendrie: Now, this was a three-chip? Wasn't it a three-chip set?

Supnik: Yes, it was a three-chip set, one of which could be replicated.

Hendrie: Okay.

Supnik: So it had a data path. It had a control store chip, which basically was the micro sequencer, and then it had up to four microcode ROMs.

Hendrie: Got it.

Supnik: So it could be minimum of three, maximum of six chips. I think the PDP-11 required...I think, four was the minimum. They required two microcode ROMs.

Hendrie: In order to do the full PDP-11 instructions...

Supnik: PDP-11 instruction set.

Hendrie: Now, do you remember what the process was?

Supnik: It was a 6-micron-- no, I'm sorry. It was a 7-micron NMOS process...

Hendrie: Okay. What is NMOS?

Supnik: From Western Digital.

Hendrie: Okay. Now, was it metal gate or was it...

Supnik: No, it was polysilicon based.

Hendrie: It was silicon gate, okay.

Supnik: Silicon gate, single metal, single poly.

Hendrie: Yeah.

Supnik: And...

Hendrie: Yeah, that was right at the edge in '72.

Supnik: It was right at the edge. And what happened is that the LSI-11 became a success beyond DEC's, sort of, wildest dreams, particularly after the original board, which was called a quad board (it was about 10 inches wide) got reduced to a half width or dual board, 5 inches wide by about 8 inches high. So now, there was a microprocessor, a PDP-11...

Hendrie: On a card?

Supnik: And one more board for the memory, and maybe one more board to have some interfaces on it. That could go almost anywhere. And people bought them by the thousands. They dropped them down wellheads. They put them into explosions. Because they were cheap enough that the data you got before it went up in smoke was invaluable. At the same time, Western Digital got into financial trouble. And it looked like they might not be able to make the chips. So DEC set up, in a great hurry, a second sourcing capability, out of the old Cumberland Farms milk plant, oh, I want to say in Southborough. It's on the East side of 495, I think, just as you got past Route 9.

Hendrie: Yes, I know exactly where that is. Yeah, it's off _____.

Supnik: So then, this was such a success that it was logical to do something as a successor, and this time instead of doing a PDP-11 that had no memory management and a small floating point subset, we wanted to do a full PDP-11 with memory management and the full floating point. So they teamed up this time with AMI, another semiconductor house, for a 6-micron NMOS single metal, single poly process. And again, sort of, stayed on the shelf at the time. But now, DEC would have to do more. We would have to do all the way through the layout. So the fab side of DEC that had been doing the second sourcing got together a small layout crew, and the small systems group that had done the LSI-11 added some logic designers and some circuit designers, and they also started working with a consultant by the name of Dan Dobberpuhl.

Hendrie: Oh. Mm-hm.

Supnik: Yes. And that name will recur repeatedly.

Hendrie: Yes.

Supnik: Dan was running...

Hendrie: I know him.

Supnik: Dan was running a consulting shop out of King of Prussia, Pennsylvania that did excellent circuit and layout work. So he was contracted for one of the chips in the second project, which was called the F-11, code name, "The Fonz." <laughs> It's 1975, after all.

Hendrie: Okay.

Supnik: It had three chips. It had a data path chip, a combined micro sequencer/ROM chip that could be replicated, and a memory management unit. And so Dan's crew got the whole memory management unit from the logic-- basically, from circuit design down, and the F-11 team, now somewhat augmented, still very small-- maybe, it's got to be fewer than 20 people, did the other two chips. And I came in at the very tail end of that. It was, sort of, done but there was a microprogramming effort required to implement something called a commercial instruction set that had been tacked on, in the hopes of making the PDP-11 into a COBOL powerhouse. Do not ask me why.

Hendrie: <laughs>

Supnik: So while I was doing the systems analysis project, I worked with the F-11 team, and wrote half of the commercial instruction set microcode, doing-- what was it? It took one whole ROM chip to do [decimal] multiply packed, one whole ROM chip to do [decimal] divide packed, and then, the third one was for [decimal] add, subtract, and shift pact.

Hendrie: Now, how many lines of microcode would fit into one of these ROM chips? Do you remember that?

Supnik: Oh, less than a thousand.

Hendrie: Less than a thousand, yeah, okay. I just wanted to get _____.

Supnik: But it was a fascinating design, one that had been created for the LSI-11, where part of the microcode was programmable logic array and part of it was ROM.

Hendrie: Oh.

Supnik: So you could read out a specific term from the PLA, but you could also put in terms that had don't cares. You could put in an address that was a don't care, and then, it would look at the instruction plus the address you put in, and select one or more lines of PLA to read out a micro word, and they wire-ANDed. So you could have a base microinstruction, and then you could zero out fields and then selectively, based on what instruction you were trying to execute, and then just compacted the microcode enormously. It's just a brilliant, brilliant design for a PDP-11. So I worked on that while I was doing this systems analysis report, and that shipped in late '78 or early '79. And it, too, was wildly successful. It was as fast as DEC's [best-selling] minicomputer, the 11/34, and it was, of course, only this big [a dual height module]. So once again, it was quite a breakthrough in price/performance. But by this point, the VAX had shipped, and it was clear that the [PDP-]11 was no longer going to rule the roost at DEC. And equally, it was possible to see that even something as big and complicated as the VAX, you know, the first VAX was 21 oversized circuit boards. They were hex width, so they were 15 inches wide. They were extended height. They were 12 inches high. It took 21 of them to do the 11/780's CPU and microcode, and then the memory was extra.

Hendrie: Yep. Wow.

Supnik: It took a double-width cabinet, lowboy cabinet. It's very hard to believe now. But you can still see the cabinets. People have hollowed them out and turned them into bars and refrigerators.

Hendrie: <laughs>

Supnik: They really look nice. So I was in charge of the PDP-11 project, from '79 to '81. It was quite a struggle, and we realized that we couldn't do it the way we had done it before. We could not do this work by hand. So as part of my training, I had worked with yet another microprocessor project called the T-11. And the T-11 was an attempt to produce a really small and inexpensive PDP-11 on a chip.

Hendrie: Single chip?

Supnik: Single chip. First single chip PDP-11. And again, it was one of Dan Dobberpuhl's designs, and it was just brilliant. I mean, it was fabricated in a 5-micron NMOS process. It had only 17,000 transistor positions and only 12,000 place transistors, and it was a full PDP-11.

Hendrie: That's brilliant.

Supnik: And it ran as fast as the F-11, just didn't have memory management. It was so small and so cheap that even DEC, with its limited abilities, was able to produce it in volume for 10 bucks or less. And we just sold it to any comer who wanted to use it as a controller, and in fact, it was the controller in Atari's

Paper Boy machine. But we, sort of, had realized that small and efficient was not necessarily going to win us the marketplace. Because in 1981, DEC made its first visit to the International Solid State Circuits Conference, to show off T-11. And in the microprocessor session, we were sandwiched between each HP's Focus chip, a 32-bit half a million transistor microprocessor, and Intel's i432, the so-called micro mainframe, which was a three-chip set with this incredibly elaborate bit oriented instruction set that-- invented by Justin Rattner that it was also half a million transistors per chip, and here we are showing off our 12,000 transistor miniature in between them, and the team came back feeling utterly humiliated.

Hendrie: Yeah, but they were doing the right thing, and those other guys were doing the wrong thing.

Supnik: That turned out to be true. Because Focus and the i432 sank without a trace.

Hendrie: <laughs> Exactly.

Supnik: And we sold hundreds of thousands of T-11's.

Hendrie: Yeah. But I get the psychology of your designers, though, yes.

Supnik: So in 1981, there were three separate pieces of the LSI design going on. There was the group that designed chips, primarily focused on the J-11 at that point. There was the backend team around the fab, and then there was the group off looking at how to do a VLSI VAX, and they were all separate. And the guys who ran it were good friends, and they said, "This makes no bloody sense. Let's get together. Let's form one group. We're going to be building a new fabrication facility anyway, because the milk plant wasn't really a suitable environment for going to better processes. Let's start a new plant. Let's have one group. Let's integrate manufacturing and engineering on the same site." First time it had ever been tried in DEC. "And we'll divvy up the roles." So that's what they did. Roy Moffa, who had been running the mainstream design group-- no, I'm sorry. Yeah, Steve Teicher, who had been running the VLSI VAX project, and Joe Zeh, who had been running the manufacturing, got together. They formed what was called Semiconductor Group, and they said, "Roy, you're going to run the business. Steve, you're going to run the engineering, and Joe, you're going to run the CAD." And then, they said, "The structure of the group's going to be this, this, and this. We're going to have a microprocessor design group. We're going to have a peripheral chip design group. We're going to have an advanced development group that will include process engineering, and we're going to have IT." So my friend Duane [Dickhut] became the head of microprocessor design because he had run the F-11 project. He had the most experience. And another experienced designer named Bill Blake became the head of the peripherals group, and they were looking around, saying, "Well, who's going to run the advanced development group?" "Well, who has advanced development in his resume?" "Oh, Bob, you do, because you did that in storage. You're it." So I was put in charge of the advanced development group, which now consisted of advanced CAD, advanced chip design, and process engineering.

Hendrie: All the advanced areas.

Supnik: All the advanced areas, and the J-11 was moved back to the chip design team, and I gave it to the head of the diagnostic group, Dan Casaletto, for which he has never forgiven me, by the way. Because it was a mess.

Hendrie: <laughs>

Supnik: <laughs> Such a difficult project. And so now, I'm in charge of process engineering.

Hendrie: Among other things.

Supnik: Among other things, right. So I go to introduce myself to the head of the process group...

Hendrie: Could I ask you to do a pause...

Supnik: Yeah.

Hendrie: ...for a sec? At some point, we need to go back and finish the story with the J-11 group that we suspended to have you go back and get the history of the 11.

Supnik: Okay. Right. Right.

Hendrie: So I don't know where you want to do it.

Supnik: Well, let me bring...

Hendrie: But I just want to...

Supnik: Let me bring...

Hendrie: Don't want to get you too far.

Supnik: Let me bring the 11 up to 1981.

Hendrie: Okay. And then, we'll do the process group.

Supnik: Right. So, as I said, the project was very difficult for Harris. They really didn't have the simulation capabilities to do the circuits. They didn't have the capability to understand their layout, and they had no verification. So one of the things that we concluded that we needed to have automated layout verification, what was called interconnect verification, or IV back then. The first programs for it were just coming out from the CAD companies, but there was a problem. We had hand-drawn schematics. It's really, really hard to verify your layout database, which was, at least, designed on a real machine, a so-called Calma Graphic system. But again, it's hand schematics. You can't-- you know, there's a step missing. So in, sort of, the last summer that I was there, we had an army of summer students and volunteers crawl over our paper schematics, and we hand transcribed them into a netlist. And then, we used the netlist for layout verification, and it worked. It worked. Because in comparison, T-11 had done its layout verification the old fashioned way, which is to print out these monster plots. And then, we all crawled all over it with colored pencils and rulers, measuring things and seeing that they went from point "A" to point "B" correctly. It took us...

Hendrie: Yeah, and then going to the schematic and say, "Oh, yeah, this is connected to this, yes."

Supnik: Right. It took us 3 months with 12 people to verify 12,000 transistors. It was clear that the J-11, which was going to hit over 100,000 transistors, at least in the ROM chip, it wasn't going to work. But we hadn't reached the fabrication stage, when the LSI group was formed in April of '81, and I turned the project over to Dan Casaletto. So we'll come back to its eventful history.

Hendrie: Okay. Good. Now it's integrated, now you're not in charge.

Supnik: I'm not in charge.

Hendrie: Okay. Good. Good. All right.

Supnik: But it will come back to haunt me.

Hendrie: <laughs>

Supnik: It really will. <laughs> So the advanced CAD group is small. The advanced chip design group is non-existent, so it's basically just me and whatever friends I could persuade to give me some time. But the process group was a fairly large group by then. I think it had, like, 40 engineers in it, with some significant capability in simulating processes and helping to understand circuits. So I went to introduce

myself to the person who was managing it. It was a woman named Ruth Rawa. Ruth was one of those very, very rare women who had entered material science in the '50s.

Hendrie: Wow.

Supnik: Unheard of.

Hendrie: Mm-hm.

Supnik: And she had followed it into semiconductors, and she had built the group at DEC from scratch. So she-- this is 1981. She had 25 years of experience, and I have none, and I'm her boss. So I would go in to introduce myself, as politely as I can, particularly because she's older than I am. I was only 35 at the time. So she looks me up and down, and she says, "Sonny, when I started in this field, you were in kindergarten." Okay. <laughs>

Hendrie: Yeah. "Yes, I think you're right."

Supnik: I think you're right. But I was able to work with her utterly, extremely well, first of all, because I never attempted to assert expertise because I hadn't-- I didn't have any. But also, because by this time, now, I was at group manager status, I had come to believe that management was fundamentally a service function. That the main purpose of management was to clear roadblocks for the people who do the actual work. If I had technical expertise, if there was a problem where I could dig my teeth into, that was great, always happy to do that and indeed, I would do microcode for the next 10 years on various projects. But in terms of being a manager, it was basically about keeping the workers enabled. And when the VLSI group was formed, because DEC was making now a major investment in a new plant and things, they went and hired an industry veteran by the name of Jeff Kalb to run the whole VLSI group. Jeff's a good guy, very smart...

Hendrie: I know Jeff very well.

Supnik: But he had his ideas. He had a lot more experience than all of us combined, probably. And as part of the VLSI VAX program, the team had recognized that the complexity of the VAX required greater flexibility in interconnect as we moved up to higher and higher transistor counts. So at a time when no one else was thinking about it, DEC had set off to develop a double metal process. Which went by the name of ZMOS.

Hendrie: Whatever.

Supnik: Whatever. It was a 3-micron NMOS double metal process. It was the brainchild of a process engineer named Andy Wu, who was basically the go-to guy in the process group, from the moment it was formed until about 1985. But Andy had no national reputation yet, or any recognition, and Jeff had never heard of a double metal process, and he thought the idea was patently absurd, from a manufacturing point of view. So the first thing he did, he comes in and says to the process engineering team, "Not going to do this. We're going to do double poly instead." Well, the engineering side knew this was going to wreck the program. So it became my job, as the head of advanced development, representing process engineering, to reason with Jeff and convince him that not only was double metal feasible, it was necessary. And I worked on that for four or five months.

Hendrie: Wow.

Supnik: Pretty incessantly, marshaling the arguments from what the VLSI VAX project was learning about the difficulties of complex VLSI designs, and eventually, I convinced him. And at that point, I was okay with Ruth, because I had saved her project. And after that, we never had a-- we never disagreed about anything. Even when she proposed moving us over to CMOS, after ZMOS. It was like, "If you believe it, I believe it." And I'll come to that. That's a little later in the story.

Hendrie: Okay. Did the double metal actually work?

Supnik: Worked like a charm.

Hendrie: Yeah, so it wasn't like it can't be done?

Supnik: No. It was just...

Hendrie: It was...

Supnik: ...that it hadn't been done.

Hendrie: ...just something that hadn't been done. I understand the need because of the issue of metal versus poly, for carrying signals around. I mean...

Supnik: Right. And we had...

Hendrie: Yeah, you know, the delays are just going to be ridiculous...

Supnik: And the struggles that J-11 would go through, once it was fabricated, to achieve anything like its design speed were just proof that you couldn't do this with two levels of poly.

Hendrie: Yeah, anyway, so keep going.

Supnik: All right. So the VLSI VAX group had put together a very ambitious program to revolutionize how we did CAD, to revolutionize how we did process development, to revolutionize how we did chip design, to revolutionize the architecture, because we were moving to 32 bits. They, sort of, did-- you know, if you remember "Hitchhiker's Guide to the Galaxy," it's like, you know, they wanted to do six impossible things before breakfast. And as a result, the project was going to take a long time. Not the least of the impossible things they had to do was to build a group from virtually nothing. Their design for the full VLSI VAX had one, two, three...four separate chips. We had the design resources for about one and a half. And we found out trying to recruit in '80 and '81 that DEC was no draw to experienced designers from Silicon Valley. You know, it was like, "Who are you? What are you doing? This? Bahahaha."

Hendrie: <laughs> Yes.

Supnik: So we got one or two experienced people and we found them and we said, "We're going to have to recruit kids." So we just went on a recruiting spree for master's level candidates with some VLSI background in school, who were really bright. And we brought them in and started training them up, except we were learning at the same time. So they were inventing and we were inventing. And that crop of, you know, kids-- they were about 10 years younger than me. I was 35. They were in their 22 to 24 range. They were just brilliant. They didn't know what couldn't be done, and they just dug their teeth into these problems as though there was no tomorrow, and many of them have gone on to be pretty well known in the computer industry now. Rich Heye, Hugh Durden, Victor Peng, are all out on the West Coast now, very successful. So after sorting process engineering, it was time to do something about advanced chip design. And once again, just as with the LSI-11, the idea came from somebody else. In early 1982, Zilog, who you may remember had the very successful Z80 8-bit machine, and the somewhat less successful 16-bit machine called the Z8000, approached DEC about building a single chip VAX.

Hendrie: Oh.

Supnik: And their reasoning was impeccable from a business point of view, which was the IBM PC was already out. People were staking out their claims. The Motorola 68000 was out. They didn't feel that at their size they could establish an independent standard at 32-bits. And they had done enough analysis to convince themselves that if you subset the VAX architecture, particularly the instruction set, you could make it all fit, at about the 3-micron level. So we talked with them and I became convinced they were correct. And so a program office was formed under Roy Moffa to investigate the business side of this.

Hendrie: Now, who was the-- was Bergard Pytell [ph?] at Zilog at that point?

Supnik: I never dealt with the Zilog business people, only with the technical people. But that name does ring a bell...

Hendrie: I mean, he was a technical-- yeah, he was a technical person.

Supnik: The name does ring a bell.

Hendrie: Okay. All right. Just curious.

Supnik: But Roy quickly, for some reason, decided that this-- he really glommed onto the idea of establishing an industry standard at 32-bits around a chip. But he didn't think Zilog was the partner for it. Don't know why. So he went on a hunting spree, and he talked to every major manufacturer of microprocessors except Intel. And Motorola was interested, but they didn't want to give up on the 68000, so they said no. And then, he talked to National. Now, National was run by Charlie Sporck. Let's just say Roy was not a match for the business wiles of Charlie Sporck. And instead of convincing them to build a VAX with us, they convinced him that we should port VMS over to their still hypothetical 32-bit chip, the 32032, and that was the conclusion from the business office. So I thought this was complete nonsense. I just said, "Not a chance." You know, "You're going to throw away the entire installed base of software." And I'd been working with the VMS people to define a subset that could both be implemented and still fit on a single chip, and I worked with, particularly, Dick Hustvedt, who was one of the prime VMS architects, along with Dave Cutler. Dave, by this point, had moved himself and a small group of people out to Seattle. They were, actually, more interested in building hardware. Dave was very active, as well, because he was thinking of trying to build a precursor system out of less-- you know, less ambitious chips, to what would be called MicroVAX. So we had it worked out. Then, comes along this confusion, you know, "Let's not. Let's use somebody else's chip." So...

Hendrie: Now, was there a way to subset the VAX architecture in terms of the running on the chip, in terms of what was implemented in hardware on the chip?

Supnik: Yes.

Hendrie: And yet, still execute through subroutines or some other way, the non-implemented instructions, so that any piece of VAX software would run?

Supnik: Yes. That's what Dick Hustvedt and Dave Cutler and I had worked out was...

Hendrie: Oh, okay.

Supnik: I did the microcode analysis that said, you know, 80 percent of the bits get used up by just 10 percent of the instructions. So if we can junk this 10 percent of the instructions, we can probably make the microcode fit. And it turned out that the instructions were the decimal and the packed and all the COBOL stuff, and for a low-end machine, it was, like, "Okay. We can live without that." And the other thing that tended to use up a lot of code on a VAX was the console for controlling the machine. I said, "Yeah, I think we can just write this as a-- in macrocode, have a macrocode console in and a very simple transition between microcode running instructions and microcode running a console well." So we hadn't worked out the details of what the trap mechanism would be for emulation. But we knew the VAX would...

Hendrie: And you knew you could do it.

Supnik: And we knew we could do it. So this is April of 1982, when Roy Moffa announces that we're going to use the National 32032. By that point, I had recruited Dan Dobberpuhl, basically full time as a consultant for my advanced chip group. And so I sat down with Dan and said, "Look, there's got to be a way we can do this. Let's floor plan a single chip VAX that steals as much as we can from the VLSI VAX four-chip project over there, and does the minimum amount of reinvention and still works." So we lifted the data path, the instruction decode mechanism, and we promoted what had been a first-level tiny, miniature TB to be the only Translation Buffer - TLB, Transmission Lookaside Buffer. But when we were done, the only thing we needed to invent really was the memory interface and the control store and microsequencer. We even used the same microword format from the big chip.

Hendrie: Oh, wow. Okay.

Supnik: Big project. And I hand drew a floor plan, using what we knew from the VLSI VAX project and then guesses about the rest, and it turned out to be about 330-- 300 mils on the side. I ran it by Dan. I said, "Dan, is this plausible?" He said, "Yeah, it's plausible."

Hendrie: Yeah, I mean, it isn't impossible. Yeah, absolutely.

Supnik: And then, I submitted a counter proposal through Jeff Kalb to Ken Olsen. And said, "We can do this ourselves, and furthermore," really, really sticking my neck out, I said, "We'll design it and get it to tape-out in 18 months." Now, you have to understand the VLSI VAX project had been running since 1979. It's 1982. They still have another two years before they're going to tape-out. So we're saying we're going to go-- we're going to take everything and get it all out in 18 months.

Hendrie: So they're going to be five years to tape-out? Yes. Okay.

Supnik: Well, they...

Hendrie: That's a long time.

Supnik: They were inventing the world.

Hendrie: Okay.

Supnik: And it's a good thing they did. I'll tell you an aside story, about Gordon Bell.

Hendrie: Okay.

Supnik: Okay. So it's pretty well known that DEC discontinued the PDP-10 in the early '80s. It's probably less well known that the reason DEC did that is that the attempts to design a successor to the KL-10 all crashed and burned because hand design of really complex, really high speed ECL machines was just not feasible, and there had been not one, but two separate projects that had crashed and burned. At the end of which, the 10 was effectively gone, because it hadn't had a new high-speed processor in 10 years. So Gordon was absolutely convinced that the key to successful machine design in the future was a CAD based strategy with heavy emphasis on simulation, at every level of the hierarchy. And so in 1982, while we're still ruminating on MicroVAX, he summons the brain trust of semiconductor engineering, Duane Dickhut, who's the head of the microprocessor group; Bill Johnson, who is running the VLSI VAX project; and myself, representing CAD and advanced CAD, to his house in Lincoln, on a Saturday morning, for a session of, literally, "You Bet Your Life," or "You Bet Your Project." And Gordon goes through an inquisition on "What's our CAD methodology? What's our simulation methodology? How do we know? It's expensive to do chip fabrication and it takes a long time. What are you doing?" All morning. And it later turned out he had done the exact same thing with a high-end VAX project called Venus, and a mid-range VAX project that was to be built in TTL gate arrays. And at the end of it, he concluded we had the right answer and the other groups did not know what they were doing. So he redirected the mid-range TTL processor group and said, "You're going to use ECL and you're going to use CAD based methodology, as insurance in case this high-end ECL project called Venus doesn't make it." And that was the result-- the origin of the project that later became known as the Nautilus or 8800 project, which is the most successful high-end VAX DEC ever built. But it's also the origin of the conflict that would just tear the company apart in the late '80s. And of course, Gordon left shortly thereafter...

Hendrie: <inaudible>

Supnik: He had a heart attack and decided DEC was too stressful, and so he never saw firsthand what happened. But I just remember that when the three of us got out from his house, it was late, we hadn't had lunch, and we were just completely drained. We could've lost our entire project that day. All right. So...

Hendrie: In hindsight...

Supnik: He was dead right.

Hendrie: He was right.

Supnik: He was dead right.

Hendrie: That was the amazing thing about Gordon. That he could figure out what was the right thing to do without some incredible two-year study program or something like that.

Supnik: The chart he wrote up in the early '80s showing the trend line for microprocessor growth and when it would cut through minis, mainframes, and super computers, was spot on. It was scary how accurate it proved to be. And had anyone paid enough attention, DEC could've avoided a world of hurt later.

Hendrie: If Ken had paid some attention.

Supnik: Yes. Well, yes, exactly.

Hendrie: All right.

Supnik: So with these two competing proposals of build-the-VAX or port VMS on the table, once it made it to Ken there was no decision. Ken said, "We're building the VAX. Just go do it." And DEC, in those days, was, you know, "All right. Ken said do it. It's the middle of the year. We don't have any budget and that doesn't matter. We're going to go hire a bunch of people..."

Hendrie: And do it.

Supnik: And don't bother me. <laughs> So I became the project leader. So now, I'm the head of advanced development, managing myself as the project leader, managing myself as the microcoder on what will be called MicroVAX II. I've never...

Hendrie: And that's what this project became known as...

Supnik: Became MicroVAX II.

Hendrie: Okay.

Supnik: So Dan Dobberpuhl joined full time, and we pulled in-- you know, the ace crew was all tied up, so we pulled in a bunch of really good people who were first timers, and so forth. The simulation and detailed microarchitecture was done by a guy named Rich Witek, who went on to be an immensely talented chip architect that-- and not only at DEC, but at Palo Alto Semiconductor, worked at Microsoft, worked at AMD, still active in the industry, a brilliant guy. But this was his first chip. He had been a CAD designer.

Hendrie: Wow.

Supnik: One of our principle circuit designers had never done anything except ECL. The layout team was-- you know, they were good, but I wouldn't say they were the experienced people. The experienced people were all of on V-11, on the VLSI VAX.

Hendrie: Was there an engineer by the name of David Grondowski that worked there?

Supnik: Yes, David and Bob Grondowski, the Grondowski brothers. Bob Grondowski worked on the project. Dave was a consultant.

Hendrie: He was-- yeah, he never worked for DEC.

Supnik: Right. But I believe he was a consultant on the project, too.

Hendrie: Yeah, he was good.

Supnik: And his brother Bob was really good. The...

Hendrie: Well, I have to tell you-- I have to stop you, just for a second.

Supnik: Sure.

Hendrie: David Grondowski went on to become a lawyer, and is now the museum's lawyer...

Supnik: No kidding?

Hendrie: ...out in San Francisco.

Supnik: That's fabulous.

Hendrie: Isn't that amazing? He completely changed career. He got to be a lawyer through being a patent-- between patents, patent examiner, then lawyer working on patents, then lawyer working on startups, then-- so...

Supnik: Wonderful.

Hendrie: ...that's where he ended up. I thought you'd-- and he's a very good friend of mine, because he worked for me at Data General, and I gave him the next generation Nova to design when he was six months out of WPI.

Supnik: Yes.

Hendrie: Because I said, "This guy just has it. It's going to work."

Supnik: Right...

Hendrie: And it did.

Supnik: Right. It was...

Hendrie: Anyway...

Supnik: That was the bits-- he did bits licensing, didn't he?

Hendrie: Yeah, yeah.

Supnik: Right. That was a fabulous machine.

Hendrie: I'm sorry. I interrupted you, but I...

Supnik: No, that's okay.

Hendrie: I couldn't resist.

Supnik: There's so many people involved in this...

Hendrie: So it was part of your-- yes, your plan. "We're just going to hire really smart young people, and they'll figure it out in real time." As opposed to, "They have to five or ten years' experience."

Supnik: And, pretty much, we did. You know, it obviously became all-consuming for me to try and get this done, and we ultimately could not make the 18 months schedule. I damned near broke the team trying. I learned a very important lesson about, you know, recognize when the schedule is broken, reset and do it right. But we did get it done in 21 months. And so it taped out in February of 1984. But there was a slight problem. We had already committed to go to the Solid-State Circuits Conference in February 1984, and at that time, there was this requirement at Solid-State that you had working silicon. We didn't have any silicon. Now, subsequently, when conferences like Hot Chips began to take over, because showing chips in advance became part of the marketing plan of tech companies, ISSCC relaxed this and now you can show of things that are in progress. But at the time, it was pretty strict. So our friends in the fab took our design in the state it was and fabricated metal-only mockups, of both MicroVAX and the full VLSI VAX project so that we could take photographs and submit them to the conference. And...

Hendrie: <laughs> Okay.

Supnik: <laughs>

Hendrie: Pretty funny.

Supnik: It is. And so we went, and we had the prime slots for the microprocessor section in '84, the VLSI VAX project and us. I actually gave the speech, and the only time I've ever spoken at Solid-State, and I was introduced by Bill Johnson, and Bill Johnson was so nervous, having just given the VLSI VAX presentation. He introduced me as Bob Supernik, and that became my informal nickname back at the group, for a long time. It was a great occasion. We ended up having a wonderful dinner out at the Mark Hopkins Hotel to celebrate, and all kinds of good things. So we got the real chips back very quickly. You know, having the fab onsite, having actually implemented our own mask making facility, we had an EB [electron beam] Lithography machine. We could turn prototypes in two weeks.

Hendrie: Wow. Wow.

Supnik: And we get these back, we fire them up, and they don't work at all. They don't work at all.

Hendrie: And you had this strategy with the simulation...

Supnik: Yes. Yes.

Hendrie: ...and the whole nine yards, and all the computer checking and the layout and the...

Supnik: Yep, and...

Hendrie: Got the schematics and...

Supnik: Yep, and we had applied our two and a half VAX's to it. So we had applied about as much computing power to it as [is in] your wristwatch. It didn't work. And, eventually, Rich Witek, who was leading the debug team, found out that by overdriving a certain voltage into the control store, he could, basically, single step the control store. So he could read out individual micro instructions, under control of his test apparatus, and put the data path and everything through its paces, so that we could debug the rest of the machine. But there was clearly something significantly wrong in the control store, and the problem was that the control store, as a circuit, was simply bigger than SPICE could handle or that we could model or could get to run, and whatever assumptions had been made were wrong. So Dan goes back to the drawing boards, redesigns. And we fabricate the second pass, great turnaround. It doesn't work, either. And by this point, we're six months into debug and we have yet to see anything that can actually run a VAX instruction.

Hendrie: Oh. Oh, well.

Supnik: But again, Rich and Dan go to work. You know, in those days, chips were-- the features were big enough you could actually put probes on them. They look at the circuit, they see that what's going on is there's a parasitic path that is sucking down the reference voltage, as a result of which everything reads out as zero. So they can overdrive the reference voltage on the tester and get the control store to work, and with that, it starts running real VAX instructions. So now, Dan has to go back and redesign the control store for the third time, and the simulation tools are still not good enough, still can't do a simulation big enough. So he has to make an educated guess based on his years of experience, and he guesses right, and the third pass works. And now, we can start doing serious debug at the system level, all through the fall and early winter. There's an adorable small system designed, by a fellow named Jessie Lipcon, that fits in a little mini tower, about the size-- a little bit taller than this personal computer, intended to sell at a really inexpensive price point. And we get that debugged, and sure enough, we're able to ship in May of '85, fifteen months after tape-out, and seven months before the VLSI VAX project will ship. Now, I want to say that this happened because, as I said, the VLSI VAX project tried six impossible things before breakfast, and they paid the price of being the pioneers. And furthermore, when they saw what we were trying to do, when their management team saw what we were trying to do, they did not treat us as competition. They treated us as brethren, and they made their resources and their designs available and they helped us with complex questions, and we reciprocated. We had designed a floating point unit that was a variant of the floating point unit for the J-11 that expanded to be a VAX, and it worked. So they took that back as their floating point unit, scrapped their design and took this one, instead, even though it was a little slower. I mean, it was a really, really, really selfless piece of work on their part. And then, there were kudos all around. Duane got promoted to be head of the semiconductor engineering group, as a whole. I got promoted to be head of the microprocessor design group, bringing MicroVAX with me. So now, I was in charge of MicroVAX, future VAX chips, but I was also in charge of J-11 and the VLSI VAX program, which hadn't shipped, yet. Or rather, the J-11 had shipped, but it had had tremendous problems. It had actually shipped in 1984, after probably seven or eight passes. The first pass had run at 30 percent of rated speed, and it had just been a nightmare, slogging through the circuits to get it up to, eventually, 75 percent of rated speed, which is where it shipped at, followed, eventually, by getting it to 90 percent. But it was a very troubled project, and when I became in charge of it again, it was now mine, and...

Hendrie: Yeah, oh, wow.

Supnik: ...Dan was-- so I paid my dues. I got to see to it that it was finally put out the door and put to bed. But it had been a wonderful experience _____.

Hendrie: Wow.

Supnik: And whereas, I was ready to do it again. I had started thinking about the next chip, six months after we started MicroVAX. I had been looking at the microcode. I was having difficulties with the

MicroVAX microcode, getting it to fit and getting it to run quickly. So I started a notebook in December of '82 on what would the microcode structure look like for a VLSI VAX chip that really ran well.

Hendrie: <laughs> As opposed to what you were working on?

Supnik: As opposed to what I was working on. Not that that was a bad design. The architecture, the microarchitecture of the VLSI VAX, which became the microarchitecture of MicroVAX, was designed by Dick Sites, a name you probably know.

Hendrie: Mm-hm.

Supnik: Brilliant computer scientist, who would go on to be one of the co-architects of Alpha, with Rich Witek. So it was fine, but it was complex to work with, and I was looking to see how I could make my life easier the next time around. Took me about six months of writing designs before I realized I was going after the wrong problem. The problem was not to make the microcode compact or easier for me. We were going to have a lot more space available in the next chip, because we would be going to a 2-micron process from 3-micron. I could add more bits. I could make the microword wider. I could have more words...

Hendrie: Even add more transistors _____.

Supnik: Could add more transistors...

Hendrie: To play with.

Supnik: The real issue was how to simplify the decoding of this microword from the over 500 PLA terms it took in MicroVAX, down to something more manageable, so that we could boost the speed. I'll come to that. So the high point is we've shipped both chips. It's time to celebrate. And in those days, DEC did not stint on letting engineers celebrate. So we rented a mansion in Newport, the Marble House. And we took the whole team and their spouses, or their significant others, down to Newport for the weekend, paid for their hotels and their meals, and we had this party at the Marble House. You know, first, with a string quartet to welcome us, and waiters serving hors d'oeuvres, and then a sit-down dinner, and then a rock 'n roll band to play until the wee hours. DEC knew how to celebrate, in those days. And then, the group began to move on, and people began to move in different directions. Rich Witek and Dan Dobberpuhl did not want to do another VAX. Done one; time to do something else. They got very interested in RISC technology, would begin-- they've become one of the threads in DEC's RISC technology story. I wanted to build the successor VAX, which would be our first internal CMOS process design, created by a brilliant process engineer named Rich Hollingsworth, who would be the guru of process engineering at DEC from

1985 until it disbanded. And what he demonstrated to the skeptical design community was you could make-- particularly, after the J-11 problems, you could make CMOS just as fast as NMOS, just use an n-well process, instead of a p-well process. Use a basic NMOS process with an NMOS process fast transistor and a pretty lousy slow PMOS transistor. Because you didn't need it, particularly, if you did what Dan had invented, which was domino logic. Domino logic is based on the idea that instead of the way most C...

Hendrie: Do you want to take a break? Get some water?

Supnik: Let me just come to the end of this, and then I'll go get some water.

Hendrie: Okay.

Supnik: The way most CMOS logic worked was that it was fully static. You applied levels that propagated through the true and complementary side, so at the end, you had a result. And the problem was that, either-- depending on whether you had an n-well process or a p-well process, either the N transistors were lousy and slow or the P transistors were lousy and slow. So what Dan invented was what he called domino logic, in which you effectively pre-charged the whole circuit to the high state, and then you selectively discharged just through the N transistors, the things that had to go to zero, using clocking. So it was not fully static, and the result is it behaved with the speed of an NMOS circuit. But it had the power dissipation of a CMOS circuit, because there was no leakage. When you pulled everything up to the pre-charge level, instead of being held there by a resistor that leaked, it was held there by a P transistor that did not, and it was just brilliant. Okay. We'll pick up our story where we left it off, in 1985. DEC is entering into the height of its success. This is when "The Ultimate Entrepreneur" will get written, when DEC's stock will get to 200, et cetera. And it's driven by the commercial success of the MicroVAX II chip and systems, and it's also driven by this wonderful high-end machine, the 8800, that was designed by Bob Stewart [and his team. But it's also, when I think the seeds of DEC's destruction were laid, and it has to do with business models. It has nothing to do with technology. So when MicroVAX was put together as a concept, it was supposed to be a chip that would be sold to all comers, and we were going to try and set an industry standard at 32-bits and head off what we viewed as a rising threat from the 16-bit PCs. At least, we viewed that as a threat in the semiconductor group. Now, who knows how practical it would've been. After all, VMS is hardly an operating system that most people would consider suitable for personal use. But at least, we were thinking in terms of changing the business model to a volume-based business model, where VAXs would be in as many places as we could put them, where other people would use them, where other software would run on them, including UNIX, and that's how the program had been initially approved. One of the reasons the VAX was-- MicroVAX II was constrained, in terms of its die size and its power consumption, was that we were trying to hit certain specific price and environmental targets. A chip would have to be less than \$50 and 2-1/2 watts or less. It had a very, very simple interface for working with industry standard chips and industry chips, and so on. Well, I later found out that this did not sit very well with Ken Olsen. Ken never believed in a model of sharing technology in

order to build volume and gain cost advantages. He believed in proprietary technology, sold at high margins, that people bought because it was superior. And Ken used to get very incensed when people would characterize DEC as a value play versus IBM as a premium play. There's a famous poster that the PDP-8 group once made, and it shows a PDP-8 banged into a shopping cart, thrust up against a particularly grotty exterior wall of the mill, and it says, "Welcome to Scenic Maynard, where thrifty people shop." Ken hated it. I mean, the engineers thought it was terrific, but Ken hated it. So after we had taped out and we were struggling to debug the chip, I get a call from him one day, and he says, "I'm glad you're making progress. We're not going to sell the chip." Now, you have to understand I'm not a vice-president. I'm not even a corporate consulting engineer. I'm just a group manager. And so when the president says, "We're going to do this," it's like, "Okay. What do I know?" But it was still okay to sell boards containing chips. So as long as we got some more added value out of it, Ken felt-- just like we had done with the LSI-11. We had sold boards. And people that built their own systems, that would be okay. So, now it's May of 1985, May 14th I believe, although that date should be checked. And it's announcement day for MicroVAX. I'm pretty sure we did this at the computer museum in Boston, which was then in Boston.

Hendrie: Yeah.

Supnik: And Ken pulls everybody together that morning and says, "I'm pulling the board. We're not going to sell boards. We're only going to sell systems." And he basically, by those two decisions, destroyed not only the business premise of the program, he destroyed DEC's OEM board business because there was nothing to sell now, and in my view doomed the company by committing it to a business model of high margin high-end products at a time when the industry at a whole was moving to volume products at lower margins. Now, not like today's PC non-existent margins, but Sun was already making noises. And it was clear that Sun was going to come after DEC by accepting a mere fifty percent gross margin instead of DEC's seventy percent gross margin. And we in the VLSI group wanted to fight back. I mean we had the technology, as they say, we can build it. But Ken was not prepared. So, that's the context of DEC's most successful years is that it is now, instead of growing the market, it is milking the market. And this attitude will create, first, Sun Microsystems, and then the 32-bit PC vendors as lethal threats to the company, my view.

Hendrie: Okay.

Supnik: But in the meantime, in the VLSI group, we're feeling our oats. We now have built two complete teams-- three complete teams. We built the VLSI VAX team, the MicroVAX team, and the J-11 team. So, we shuffle all the cards at the beginning of '86. And a lot of the people from the J-11, as reward for the suffering that they've been through, get to do CVAX, the next MicroVAX chip, the CMOS VAX, CVAX. Dan and Rich Witek from the MicroVAX team don't want to do another VAX. And they go off and do RISC technology. And the VLSI VAX team starts on the chip that will come after CVAX. So, we can actually now do two projects stagger-stepped. And I've laid out a five year, what I called a process/ product

map, because Rich Hollingsworth, brilliant guy, had given me his road map for what the processes could be from the three micron nMOS through two micron CMOS and then fundamentally stair-stepping down in a compatible way. So, one and a half, then one micron, then three-quarter micron, all compatible, all able to be optically shrunk so that we could reuse designs. And I lay out what we can do with this increased complexity at a micro architectural level and what the performance goals are going to be. So, the idea is that with CVAX, we're going to make a half-step forward. We're going to take performance up by a factor of two and a half. With the next one, we'll get another factor of two and a half. And then we'll get another factor of two and a half or more. And we'll do this by lifting the best ideas of the big VAX group. No point in reinventing the wheel. They've figured out how to make better designs or faster designs. We will figure out how to miniaturize those designs and implement them. Okay, so CVAX, which is our first internal CMOS project, two micron, double metal, n-well CMOS, it goes reasonably well. The-- we see how to meet the architectural goals. We make a few more architectural advances. We bring cache on chip for the first time. Cache is actually built out of dynamic RAM too, for density, which was kind of weird. But it worked. And the micro words are wider and simpler. The performance goes up from that. So, even though we're only going to-- and the other thing we finally understood is that we should take advantage of the fact that chips bin over a range of speeds. And therefore, the interface to the system has to be flexible enough to accommodate fixed system timing with variable chip timing. So, we fixed that. That was a problem in the MicroVAX. And so, the head start I had gotten on the MicroVAX project gives the project a head start. Dan Casaletto, who suffered through J-11, gets to do CVAX as the project leader. And it goes pretty well. We're late again. We're three months late. We still haven't quite figured out yet how to do scheduling on these big complex designs. But the team is really feeling a lot more confident because we've done it. They're veterans now. They have one under their belt. Everybody has one under their belt. And we're still hiring bright Master's students by the truckload. Except now we've got co-op programs. So, we bring them in from their junior year onward for three-month stints or for summer stints and train them or brainwash them in the DEC way, our way. And they hit the ground running. And they're just brilliant just like the original kids were. A lot of them come from the University of Cincinnati, Jeff Kalb's alma mater, where he set up really, really, strong links. But we also recruited from University of Illinois, Dan's-- Dan Dobberpuhl's alma mater, from CMU, from University of Washington St. Louis, Rich Hyde's alma mater, University of Washington, Washington State etc. Tended not to do MIT or Stanford very much not because we didn't want to, but we found students from these other universities were a bit-- I don't know, just a bit more interested in really getting their hands dirty quickly. So, that strategy continued to pay great dividends. And the group kept expanding. We started getting a modest number of women engineers, which became one of the group's focal points <cell phone vibrating> was could we find more women engineers. Do you need to take that?

Hendrie: No.

Supnik: Okay. So, CVAX is the last machine for which I will do the micro code personally. I'm still enjoying it too much.

Hendrie: You won't delegate that because it's too much fun.

Supnik: It's too much fun. But I sort of recognize that it's not appropriate, given that I have all these other things that I'm supposed to be doing. And the group has now grown to over a hundred. And we're-- what's happening is that the-- as our chips are getting more powerful, we're beginning to eat away the space that other processor design teams used to occupy. And so, the number of processor teams at DEC is beginning to shrink. At one point, there were as many as six. And obviously, the VLSI group first took over the low end. Then we took over the low mid-range. And people were really not territorial about it. They were-- when they saw it, they said, "Yeah, I can't compete with this with TTL gate arrays," or whatever their technology was. And so, could I either get into the system design aspect, or can I join the team? So, we were getting some really fine recruits. A gentleman named Mike Uhler, for example, who was part of the last PDP-10 efforts, joined and became one of our principal architects. He went on later to AMD, and MIPS, and a few other companies. So, but CVAX is sort of still my baby. And because it's-- these small VAXs are so successful, the software group comes in with a request. They say, "You know, we think we could actually make COBOL work just as fast on these machines as on the big ones with all their fancy instructions if you just give us a few more instructions. And not even decimal instructions, just give us a few more of the character string instructions." So, I do some trial coding. And it was yeah, yeah. It's-- it looks good. So, in like the last six weeks before tape out of the second pass, I insert these instructions in the microcode, do the verification, send it out in the field. I broke the machine. I had overlooked a sneak path where when they're interrupted and something particularly peculiar happens in the interrupt pack up process and, as a result, the interrupt never happens, and you're supposed to resume, the instruction goes haywire. So, I have now personally set the project back three months. And it's like all right, time to retire from microcoding. Time to let somebody else do this. But it's still out in time. And it's not only a tremendous success as a low-end machine, the mid-range group builds a really affordable multi-processor out of it called 6000 series, which goes from nothing to a billion dollars in sales in nine months.

Hendrie: Wow.

Supnik: So, this is a tremendous technical and business success. And things are still looking really good. But without Gordon's guiding hand, or perhaps I should say Gordon's dictatorial authority, either way, same thing--

Hendrie: Well, a guiding dictatorial authority hand.

Supnik: Yes. Things are beginning to get very anarchic in DEC. And groups are starting to fight each other in really unproductive ways. In particular, as the chips are taking more and more of the low end/low-midrange real estate, more and more groups are competing to build the highest end machine. There's the group that built the 8600 is now trying to build a machine called Aquarius, or the 9000, which will be an ECL multi-chip substrate mainframe kind of like the big IBM mainframes. And it was called Aquarius because it was going to be water-cooled, which DEC had never done. The group that built the 8800 was building a slightly less ambitious high-end machine called Argonaut. And then Dave Cutler was out on the

West Coast. And he wanted to build a RISC processor. And these groups are all competing for resources and budget. And there's no one to adjudicate. There's no one to say, "Yes, no, maybe." And Ken had-- the old DEC had thrived on internal competition. So, Ken was-- saw this as part of the process. But he hadn't really understood the capital cost issues that were coming up with these technologies, like the multi-chip substrate that the 9000 wanted to use was based on the technology developed by Gene Amdahl for his failed startup called Trilogy.

Hendrie: Trilogy.

Supnik: Trilogy.

Hendrie: Yeah.

Supnik: And the capital cost of developing that would rival, eventually, or exceed a semiconductor plant. So, all these groups are competing. And in the meantime, the semiconductor group is sitting on a roadmap that says ECL will be toast by the beginning of the '90s. And we can-- we'll show you. But everyone is proposing these thirty MIP machines. Seem peanuts by today's but-- and we're only sitting at three with VAX chips. And so, to say that we're going to go an order of magnitude in just two generations or so, or three, people don't credit it. Ken doesn't credit it. And in the meantime-- so, there's this rivalry among the ECL teams for building a big VAX. And then there's a rivalry about whether we're going to build a RISC processor to compete with Sun Microsystems new SPARC line, which results in an internal program called Prism. And Prism is a thirty-two-bit architecture, expandable to sixty-four bits. But it won't run VMS. It's going to run a new operating system that Dave [Cutler] is writing. And so, you've got this three-way war among the high-end groups for resources, and political influence, and so forth, and no one to hold the reigns. Into this volatile mix steps a fourth team, a group of-- a renegade group. That's the only word I can use. I mean renegade in the good sense of not being fully authorized. [It] decides that if DEC is going to get into the RISC game it can do it a lot faster if it uses an off the shelf chip like the MIPS R3000. And while the effort to build Prism is struggling for resources because of all the internal competition, they just go ahead, and they build a R3000 based workstation. And they port DEC's version of UNIX to it, called Ultrix, in six months. And they throw it out on the table and say, "This is what we should do." So, things are very complicated. Meanwhile, back at the ranch, back at the semiconductor group with the team that was following behind CVAX ultimately, with a little prodding from me, decided that we would build yet another sort of single chip model machine using the architecture of the highly successful 8800 as our goal. And we'll be targeting the same performance levels as the 8800. And this is called Rigel. And it's going to be built on our CMOS2 process, which is one and a half microns. So, it's a second-generation CMOS project-- process. And some new hires come in at the-- on the management side, or get promoted on the management side. In particular, a gentleman named Amnon Fisher becomes the head of the project. Moshe Gavrielov, who is now the CEO of Xilinx, is going to do the floating-point unit. And Bill Herrick, who remained a senior chip leader at Intel until he retired, will do the CPU. And one of our female engineers, Rebecca Stamm, will lead the cache chip. And Amnon says, "I'm going to solve

this problem of why we're always late. I'm going to make the project predictable." And so, while everyone else is studying the technology and the architecture, and I've done a first sketch of the micro code, which I've turned over to Mike Uhler quite thankfully, Amnon and his project leaders work on the question of why we get it wrong. And they come to understand that, in fact, it's possible to quantify the progress that we make even in this heavily custom, hand-drawn, hand-simulated methodology by counting basically the number of schematics that get generated and retired in a given period through-- generated through logic design, put through circuit design, retired through layout. By monitoring that and establishing what the run rates are, you can actually predict the course of the project quite well. And they put this methodology into place. And they nail every date, just brilliant, and real kudos to them for that. So, Rigel comes out, first the cache chip. It works. And then it's time to do the CPU. And the CPU doesn't yield. It simply does not yield. It's the biggest chip we've ever built. Not huge by today's standards, but for DEC, it was quite a leap in size. The CMOS2 process was brand new. We could not yield it. We could not get the defect densities down. We couldn't get the floating-point chip to yield. At the end of seven months of fabrication and debug, we had two working processor chips. And it's sort of like--

Hendrie: Oh my--

Supnik: Do we go for the hara-kiri knives now, or do we wait?

Hendrie: Yeah, so are you or do we build another two hundred and fifty chips to get another one?

Supnik: Right.

Hendrie: That kind of thing.

Supnik: Well, as I said, Rich Hollingsworth, who was in charge of process, he knew his stuff. He had-- he went through it systematically. He started cataloging the defects, where the process steps were not adequate, what controls were not good enough. And it turned a corner kind of in a period of like two weeks. Before that, it wouldn't yield at all. After that, it yielded beautifully, after he put these fixes through. And not only did it yield, it was an incredibly fast process. We had designed Rigel, this third generation VAX chip set, around a forty nanosecond microcycle.

Hendrie: Okay.

Supnik: Everything ran at twenty-eight or faster.

Hendrie: Yeah, okay.

Supnik: It was a brilliant process. Part of it was that in putting the controls in to be as precise as he felt he needed, he had made it possible to skew the process to the fast side and keep it there without variation. And indeed, when Intel would take over DEC's fab in 1996, they were flabbergasted by just how tightly we ran the process compared to the way they ran theirs. Yes, yes. So, Rigel-- when Rigel came out, it became actually now the fastest VAX processor then running. This is early '88.

Hendrie: Oh, my goodness.

Supnik: So, the whole high-end as we knew it is gone. And in the meantime, the high-end groups are still fighting for their space and their territory. And their designs are stretching out because they're so complicated. And so, everything came to a head in late '88. The stock market crashed in October of '87.

Hendrie: Right.

Supnik: You remember. DEC's financial results began to stagnate in 1988. Basically, revenue growth stopped at thirteen billion dollars and would never grow again. And finally, some decisions had to be made because there was no money left to keep all these rival projects going. So, the 8800 successor, called Argonaut, was cancelled in favor of the high-end project known as the 9000, or Aquarius. And the team was furious and heartbroken. The co-architects-- Bob Stewart was the project leader. But the key architect was a brilliant guy named Doug Clark who later became a computer scientist down at Princeton University. He was on vacation at the time in Hawaii. And Bob Stewart sent him a one-word telegram, "Stay. "

Hendrie: Okay.

Supnik: And Doug was so disillusioned by this that he left the company and, as I said, went into academia where he had a tremendous career. Doug was the person who pioneered statistically driven debug where you monitor bug rates. And if they start to fall, as opposed to proclaiming success, you look at your tests to see if in fact they are no longer driving the system adequately. And he wrote this provocative paper, which is true to this day, called "Bugs Are Good". And it became the basis of the way we thought about simulating and testing the chips and design. So, with that resolved, there was still now the conflict between the MIPS team and their microprocessor and the Prism team. And so, on one hand you have Dave Cutler, my colleagues Dan Dobberpuhl and Rich Witek, who were building a Prism chip. And on the other-- and we're still at least a year away from tape out because of resource constraints within the VLSI group. And on the other side, you have this group in California led by Carol Peters and some other folks from the Western Research lab who have built a working system that runs Ultrix and is actually ready to go to market. So, Ken commissions a team to look at it, and he puts me in charge.

Hendrie: Uh oh.

Supnik: Uh oh.

Hendrie: Uh oh. No win. Can't--

Supnik: And by this point, I am a-- what's called a corporate consulting engineer, which is like the second highest engineering title. And I tend to take my fiduciary responsibilities pretty seriously.

Hendrie: So, what's the highest one?

Supnik: Corporate-- senior corporate consulting engineer.

Hendrie: Oh, okay, just add senior in front of it.

Supnik: That will-- that comes later. So, I look at it. I look at the fact that Prism is still a year out, that people are beginning to understand about first mover advantage and how order of entry into the marketplace often determines market position. Where-- sorry, SPARC is out. I think about this. I agonize about this. But I come to the conclusion that Prism is going to be too late, that it will be the fourth entrance into the market because we knew IBM was-- had Power in the works. So, we'd be behind Silicon Graphics, which was a power in those days, and behind IBM. Oh, and HP the HP PA-RISC was also on the slipways. And so, I recommend that we go with the MIPS alternative. Ken accepts it. Prism is cancelled. My colleagues will no longer speak to me because I had basically contributed to ending the work they'd been doing for three years. Dave Cutler quits the company, goes to Microsoft. And so, this is in June of '88. In September of '88, the head of the-- the now head of the VLSI group, Bob Palmer, comes to me and says, "You can't be a manager anymore. You have lost the confidence of your people. You're going to have to do something else." And that something else was Alpha, as it turned out. So, when Ken cancelled Prism, he turned very casually to me and said, "Why don't you figure out if there is some way we can make VAX VMS run a whole lot faster? We are getting hurt on performance by these RISC machines. You know, go see if there's some way to make VMS a lot faster." So, I formed a task force of friends, the ones who would still speak to me. We called ourselves the EVAX, or Extended VAX, task force. It was me. It was Nancy Kronenberg from the VMS group. It was Dick Sites from the VLSI group. It was Rich Grove, who ran the compiler work for DEC, and some other folks. And we stared out looking at the question of whether there was a way to subset the VAX more radically to allow you to do better pipelining of the instruction to code and execution so that you could get into multiple instruction units and pipeline memory accesses, and so forth. And after a while, when we looked at the statistics, we realized that there were just certain features of the VAX architecture that really, really crippled it: the long and lengthy complex instructions that could be up to fifty-three bytes, the incredibly heavy procedure call and return instructions, which basically required full memory ordering through it. And it simply would not be competitive with the RISC technology as we knew it. And already DEC had built a couple of experimental RISC machines. So, we had real practical experience. Plus, we had the MIPS product line now running

UNIX. So, now the focus became was it going to be possible to run VMS on something else. Could you run VMS on a RISC machine? And the starting point was no. It's clearly impossible because VAX/VMS-- VMS was written entirely in assembly language. And it's a 1975 operating system. So, it must be a disordered mess internally. Plus, the VAX is full of architectural ornamentation, which is absolutely critical to the VMS interface. So, it's clearly impossible. And even if you could do it, people have lots of programs that they're not going to recompile. And you can't get VAX code to run. [My wife is home.] And so, we had the starting position was you can't get there from here. And then people were-- began, more or less on their own and with some encouragement, to see what it meant, whether these assumptions, these axioms almost of our thinking, were correct. And one by one they got knocked down. So, Nancy Kronenberg, who was a veteran at-- she had been one of the principal implementers of clustering. And she was a longtime colleague of mine dating all the way back to Applied Data Research.

Hendrie: Oh, really? Okay.

Supnik: And her husband, Paul Beck, too. Both of them, we had worked together in the little Cambridge and Boston offices. Nancy started studying VMS. And she discovered that, rather as you might expect given for a well-designed, well-architected effort by Dave Cutler and Dick Hustvedt, it was actually quite well structured internally into a layer that faced the hardware, a layer that faced the applications. And a big middle section that knew nothing about the architecture, just it's scheduling and the IO and stuff. So, she came back and said, "If you could figure out how to get the source code across, there's not that much rewriting we actually have to do." Okay. Dick Sites, in the meantime, had decided to prove that translating existing VAX binary code into a RISC machine simply could not be done with any degree of efficiency. We had already had quite a bit of binary rewriting technology developed by that point. A product called ATOM/OM, which would take a complete program, and now with knowledge the compiler didn't have, because it's got the whole program, could actually rewrite it, to be much more efficient. We typically got a 20 percent improvement in transaction processing benchmarks with Oracle just by doing that. Easy to prove that it couldn't be done. So, he goes off and he does probe coding and tries to look at how we would trace things. And he comes back after a couple of months and said, "I can do this, and what I need is a three-transistor hook in the architecture." Okay.

Gardner: What's that?

Supnik: It was a bit flag that you could set that would be cleared by any interrupt. That would permit you to simulate some of the non-interruptible instructions by basically testing whether you have gotten interrupted in starting over again. It's all he had wanted. So, Rich Witek joined the team now, and Dan joined the team as well, because we were seriously talking about building something. And our goal was to build something that would leapfrog ahead to 64 bits and just burn the industry down in terms of performance. And even though Dan and Rich were really not feeling very friendly to me, they joined the effort. Dick Sites and Rich became the co-architects of the Alpha architecture, and Dan became the project lead for what would become the Alpha chip. So, we proposed this, reported out, and got a go

ahead saying, "Yeah, you know, it's worth an advanced development project. Sure, go ahead and do it." But by this point this is now the end of 1988. I'm no longer holding a management position anywhere. I'm an individual contributor. I'm a corporate consulting engineer. And I realized that in this highly politicized environment with groups competed for resources, trying to form another independent team is pointless. So instead, I decide to construct a program office. And the program office will coordinate people who voluntarily sign on to help with this project out of resources that they can find in their groups. And I meet a key collaborator named Peter Conklin, who unfortunately is now deceased. Died way too young. And Peter understood this concept of what he called an enrollment management. And so, what we did was structure a program office. There were never more than half a dozen of us. We never had any money. And our job was to go around and influence people to get them to sign up with their own resources and their own thinking to join. Then we had some formal support. We were able to-- we had people out of the semiconductor group. We had Dan and Rich, so they can get some friends of theirs. We had some formal support out of the VMS group, because of Nancy. We had some formal support on the compiler group, because of Rich. But essentially over the first year or so, it was an effort of going to talk to people, making presentations, telling them how they can contribute and what they could do and convincing them to split off a piece of their budget to work on this. I think I delivered a hundred Alpha presentations in a year inside the company. Now, of course, it wasn't called Alpha to start with. It was called EVAX, Extended VAX. That was the name of the task force. But as we began to get some momentum together, people felt that was really giving away the store as a code name. So, we had to have a real code name. Jesse Lipcon, my colleague in the system group who had done all the small VAX systems, had named the very last one, the one that used the last VAX chip, called NVAX. He called it Omega. It was the end of the line. So I said, "Fine, we'll call this Alpha, because it's the beginning of the line." And that was it, and that stuck. And we had trinkets printed up and giveaways and doodads and so forth. Just that, "Hey, come join the team and get an Alpha Mood Indicator or an Alpha Memo Pad" or something like that. And Rich and Dan, Dan pretty quickly understood that he saw a way with the capabilities of the CMOS process. We were now down to one micron. CMOS3 and we're going to do the actual chip in CMOS4 in three quarter-micron. With the way Rich Hollingsworth could deliver precise, fast processes, he [Dan] saw a way to just shatter every known frequency record with this chip. So, what he decided to do was build a precursor chip, a chip that would be just the integer data path. We wouldn't have any on-chip cache. That [chip] would permit us to get software development started and would require a lot less effort. And then [we would take it to them] when we had more real estate with the next process, we would add a floating point and the caches and we'd be in business.

Gardner: Okay, really was real estate constraint <inaudible>--

Supnik: It was a real estate constraint, but also a time constraint. The floating point unit was going to be lifted from NVAX, the last VAX processor, and it wasn't ready. We had actually done quite a good job of designing a reasonable floating-point unit in the Rigel generation. We overdesigned the floating-point unit to be fully pipelined, knowing that we might need it later on, and it got used it four times. It got used in Rigel. It got used in the VAX vector unit. It got used by NVAX, which could use a fully pipelined floating-point unit, and it got used by the first Alpha chip. So, we got our money's worth. Now at this point, there's no real systems effort because there's no real chip specs or anything. So, Dan hooked up with the DEC

Research Group, run by Sam Fuller, and he persuaded a couple of the key engineers there to design a prototype system that could be used for software development. And he got three of the best people on the planet. He got Chuck Thacker, he got Dave Conroy, and he got Larry Stewart. So, Chuck and Dave did the basic hardware design, and Larry did part of the hardware design and a lot of the software tooling around it. [Although it was based on a CMOS chip, they ran it upside down. They interfaced it to ECL RAMs to make the external cache fast enough, and they kludged it all together with some I/O system. We hooked a couple of things up to it and suddenly we had a system. So this project, which was begun as a think program [thought experiment] in June of '88 and was not formally authorized until December of '88, taped out its first chip in January of 1990 and had a running system that same month. And then Ray Lanza of the ULTRIX group - the bulk of DEC's UNIX resources were working on something called OSF/1, which was this grand unification UNIX, which of course went nowhere. But the old ULTRIX, which was really just BSD4.2, was still around.

Gardner: The BSD4.2, yeah, emitting VAX code.

Supnik: Right. He ported it to Alpha with seven people in about two months and made it 64 bits at the same time. The first 64-bit UNIX port that we know of outside of Cray supercomputers. So we not only had silicon in early 1990, we had software, and we could start measuring. And even though it was running at half the speed that we had ultimately intended, it was a 100 megahertz rather than 200 megahertz, it just screamed. It really screamed. And the VMS team reached its goal a little later. This enrollment management had a couple of interesting effects. First of all, it got the people who still had initiative. And DEC was really hurting by this point. We had [had] our first layoffs. The behavior of the upper management was becoming increasingly irrational. One of the things I remember is the big ECL VAX called the 9000 was announced six months before it was ready. It was announced in late '89. It didn't ship until mid-1990. By that point, we had silicon on the last VAX chip we were going to do, NVAX. And we had beaten the schedule. It ran [on the] first pass and it had many fewer bugs than anything we've ever done. And we had measured it, and it was just as fast [as the 9000]

Gardner: As the 9000.

Supnik: As the 9000. Ridiculous.

Gardner: Yeah, I mean, just, in other words, they were taking the people doing discreet logic.

Supnik: Yeah.

Gardner: Were just--

Supnik: They were doomed.

Gardner: They were doomed. Yeah, they knew what they weren't doomed projectively, they were doomed. <laughs> They didn't ever know they were doomed.

Gardner: They never had a chance. So we have NVAX in hand, and I get a call from Ken [Olsen]. And Ken's very contemplative. And actually, this may have been a bit earlier. Well, this may have been when we were taping out. And Ken, he asks me, "Bob, is it true that this chip here you guys were working on is going to be just as fast as the 9000?" And I said, "Yes, Ken, it is. It's going to be just as fast and it's going to be 1/100 the cost," I said. He says, "Then this little board you'll design it into-- is going to beat all that cabinetry full of logic." I said, "Yes, Ken, it's going to," and he sighs and says, "I just can't get my head around it. I mean, I can intellectually grasp what you're saying, but I just can't get my head around it." So, as we enrolled people, they often began to come into conflict with their management. Their own management would have its own goals that you were supposed to be doing X but, no, you wanted to work on Alpha. And this became particularly acute in the VMS group, and ultimately the VMS-- by this point Alpha had been put under a VP named Bill Demmer, who also had the VMS group. And Bill said, "All right, we'll split the group. The tension's too much. We'll form an Alpha team, a separate Alpha team, and it can't be very big, but they'll do the VMS for Alpha, and the rest of the people will focus on the VAX where the money is." So the team was split, 400 people. Eighty of them joined the Alpha effort, and 320 remained with the mainstream group. The Alpha team was led by Jean Proulx, who had moved into general software after running the DEC Hudson CAD group. And they immediately started shaking up the way VMS had run for years. VMS used to build every two weeks. They built twice a day. They did a cross-compiler for VAX macro code to Alpha that was remarkably efficient. It used the same backend as all other compilers, and they had started with a very narrow range of goals. By the end of it they said, "We're going to deliver-- we'll be just one minor release behind the mainstream and we will cluster and inter-operate with VAX VMS." The one thing they wouldn't commit to was 64 bits at initial release. In the meantime, the remaining 320 people failed to deliver a new VMS release for two years, because the Alpha effort had taken only 80 people, 20 percent of the group, but it had taken the right 20 percent of the group.

Gardner: <laughs> They had taken the best 20 percent.

Supnik: The ones with initiative, and the rest of them could not figure out how to get organized. So, by 1990, I had moved out of the semiconductor group over to Bill Demmer's organization to better focus only on Alpha, and Bill--

Gardner: Now were you still a project office or--

Supnik: We were still just a program office--

Gardner: --had that disappeared?

Supnik: We were still a program office. We would be 'til the day it shipped. Bill calls one day. He says, "Bob, you're a vice president." "Why am I a vice president? I've got five or six people who work for me. One of whom is an administrator." "Well, there's this customer, and they want to see a DEC vice president. Here's your plane ticket." <laughter> So, both Jesse [Lipcon] and I who had at that point been holding primarily engineering titles (we were both corporate consulting engineers), suddenly became vice presidents, or as we like to say, "We were expendable vice presidents." You get on the plane. You go talk to the customer. But it did have one really neat side effect, which was when later when I would become interested in doing simulators, I could write licenses for obsolete DEC software and sign them as a DEC vice president, and they were binding. So, I personally released the PDP-8 catalog, the PDP-15 catalog, and the PDP-10 catalog, for non-commercial use, aided by some friends in the legal department who wrote up appropriate license documents. So, it worked out. But now it's 1990. We have working software, working hardware. VMS would come along by about June of '90 and get their first working system. We have demonstrated that we are going to destroy the VAX business one-way or another. And part of it is that, as a program office, we also took responsibility for writing the business case. And the business case was a pretty frightening document, because we acknowledged the change in industry economics and that, if DEC was to be competitive, it had to follow both the volume and gross margin trends for the rest of the industry. So, Alpha was going to be sold as a chip, and we never looked back. We never actually asked for permission, we just sort of did it-- licensed it to Mitsubishi, and eventually to Samsung. We were diverting people from the approved course that the vice presidents thought they were supposed to be doing. And Ken became increasingly alarmed that we had set up a rival power center as he was struggling to basically impose his will on the company. A will, that most of us thought, wasn't terribly rational. He was interested in doing proprietary versions of servers with Intel technology in them as opposed to commodity servers, and he was convinced that good packaging was going to be the differentiator. And he had reorganized, and these people would go out of favor and that person would go out of favor or in favor. And he started referring to us as the Alpha mafia. And really, that was the term. And we were, indeed, a merry band of--

Gardner: You were having fun--

Supnik: We were having fun, but we didn't think we were doing anything invalid. I had gone to the, then, vice president of engineering and manufacturing, Jack Smith, at some point. And he had said, "Okay, the company is in real trouble. We need to get this done sooner. Can you pull this in?" And I gave him a plan, called for 30 million dollars in incremental spending, and he had approved it and we had pulled in our ship date by six months, to the end of '92. But Ken was becoming increasingly irate about it. I never saw him, and basically Bill Demmer buffered me from everything. But there were some dicey, dicey issues-- dicey times, like we had a wonderful guy on the team named Rob Hannemann, whose specialty was micro packaging, which we weren't doing anymore, so he was kind of at loose ends. And so when he heard about this, we had had all kinds of little trinkets printed. He had a business card printed with a black hand on it, with a Greek Alpha in it, laminated, and he put it under the door of everybody on the Executive

Committee, including Ken. Fortunately, that day Jack Smith got in earlier than anyone else and collected them, but I heard about it, as you could imagine. And by this point, the semiconductor group was completely aligned behind Alpha. We now had made EV-4, the principle project. This was the real Alpha chip. So the Extended VAX name never really quite went away. The Alpha chips were always called EV-something, where the number was the process generation. So EV-4 was Alpha in CMOS-4. But we couldn't tell people it was EV, so we always said that the EV designation stood for Electric Vlasic. A Vlasic is a pickle, of course, and this was based on a famous experiment run at the Western Research Lab in which they ran electricity through a pickle and proved that it glows, it emits light. This was written up as the infamous technical note WR13, Western Research TN13, and so it was always EV, Electric Vlasic. But we were going for broke. NVAX would turn out to be a very fast chip and CMOS-4 ran it at 91 megahertz at its highest bin point. We were going for 200 [Mhz]. And this is at a time when Intel was at 66 megahertz and people were talking about maybe getting to a 100.

Gardner: And have you gotten to 64 bits?

Supnik: Oh yeah--

Gardner: Yeah, yeah. So this is 64 bits at 200 megahertz.

Supnik: At 200 megahertz. I mean, there's no doubt in my mind that Dan Dobberpuhl is the most brilliant circuit designer on earth. When you consider, between MicroVAX and Alpha, and later the StrongARM chip, and the Palo Alto Semiconductor chip, that unfortunately, never saw the light of day. So, as we roll into 1992, EV-4 is being fabricated and the systems are being prepared. And Ken is getting more and more agitated, and he's talking about closing the project down, and he's saying that we're unethical. He has the internal audit committee investigate us to see if we are spending money we're not entitled to. It's become a real Game of Thrones or Sopranos-type stuff. Very civilized of course, but very, very dangerous, and it all comes to a head in early 1992 when a couple of the corporate consulting engineers, we're now senior corporate consulting engineers, and we go and see one of the members of the Board

Gardner: Who'd you go see?

Supnik: Cannot remember his name. Bob something, or other.

Gardner: Not Bob Everett

Supnik: Bob Everett, exactly. And we tell Bob that Ken is going to sink the company. By now, he is into this line of very big Intel servers in tall, thin cabinets that actually are unstable, and that they're an

uncanny resemblance to coffins turned on end. We said the company is going to go under. So, Bob goes in and talks to bunch of people on the Board.

Gardner: Yeah, on the Board.

Supnik: Not only the board, but also in the company, because he was the one most actively involved in engineering.

Gardner: Yeah, a long-time engineer. He been head of MITRE

Supnik: Right, and he comes to the conclusion that Ken can't save the company and needs to be replaced. So, in the April board meeting he moves that Ken be replaced, April 1992, and he's voted down. Well, Ken, needless to say, wants to know who's behind this. Bob tells him. And within a matter of hours, the key ringleaders, if you like, or the key complainers, are informed that their careers are over and they should be looking for work outside DEC. Included me, Jesse [Lipcon], Don Gaubatz, and Bill Strecker. We're all informed we have no future at DEC, and it's unclear to me that Alpha's ever going to be allowed out the door now. So I take it seriously. I actually started interviewing. I was--

Gardner: Where'd you look?

Supnik: Intel. On the west coast. I figured Intel. I'm a microprocessor designer. I might as well go work for a company that does microprocessor design for a living, and I'm talking to them. I'm being interviewed by Patrick Gelsinger-- he was, then, at Intel. He had a key technical-management role. And I get back to -- this is now July '92. I get back to my hotel room and I get a mail message saying, "The board has ousted Ken." So, there was another vote three months later, and this time the Board did remove him. And I called my Intel people and say, "I'm sorry, I can't continue this interview. I'm going home." And I went right back. And Win Hindel took over as interim CEO while they looked for a full-time replacement, and Win had always been quite a supporter of Alpha and he told us, "Just get this thing done, you know, heads down, don't look." And we shipped it. And at the end of '92 we shipped, obviously, a new chip, a new architecture, a new chip, four new systems, 138 [130] layered products, two operating systems, diagnostic support, documentation.

Gardner: All the stuff together.

Supnik: All the stuff that goes with it. By the end of the effort, the program team was a little bigger. I think it was up to eight or nine people, but that's all it ever was. Peter was the heart and soul of the campaign, Al Avery, now, unfortunately deceased also, ran our performance marathon. Six months before shipping we were off our performance target by two-thirds and he ran just a regular-- I'd almost call it a sprint-like

methodology, where every week people were required to turn in new results and show what they had changed. And Dick Sites did all the analysis. He would take the output from the compilers. He would go through it by hand and say, "This is scheduled wrong, this is superfluous." And Rich Grove would take it back. Rich would go drive the people in the code generator to get better code out of them. And the performance goal was hit on the button.

Gardner: But it was the VMS performance? I mean, the chip was doing fine.

Supnik: Right, it was software performance.

Gardner: It was software performance, yes.

Supnik: It was actually application performance more than VMS. VMS'sported macrocode had done pretty well. And just -- so, Harriet Cohen was our sort of product-manager manager. She sort of coordinated all the product management. There was one person who did schedules, one person who did licensing. I'm sorry that I can't recall the names--

Gardner: That's all right.

Supnik: All my documentation is now in the possession of the [Computer] History Museum.

Gardner: Is that right?

Supnik: Yeah.

Gardner: Oh good.

Supnik: All my personal archives, my email archives. Everything is over there.

Gardner: And we know all the DEC stuff was moved to--

Supnik: Yeah.

Gardner: --the museum.

Supnik: Yeah, when [ph?]-

Gardner: That's excellent.

Supnik: And we disbanded the program office on January 1, 1993, because it had accomplished its goals. And it was kudos all around, but the capstone of which was again Solid State Circuits Conference. Solid State 1993, and Dan Dobberpuhl was going to present the paper on Alpha. Now it's shipped. It's a working chip. We have somehow managed to keep the technical details out of the press to meet the Solid State requirement of first real paper. And Dan goes up there and he shows how we have achieved 200 megahertz. He proves it with graphs and pictures, all in 20 minutes. And the audience is just stunned. And afterwards he is mobbed like a rock star.

Gardner: Well he is a rock star--

Supnik: Circuit designers all come pouring in, "How did you do this? What did need here?" and it was just the most wonderful experience. I was just on the sidelines watching, but it was truly, I think, the climax for my DEC career.

Gardner: Wow that's cool. Very good.

Supnik: So we're at 5:30 and maybe this is a good point to stop--

Gardner: Yes, why don't we do that?

Supnik: --because what we've got left to do is what I call the epilogue with DEC and then the 10 years I spent doing startups and my conclusion of my career at Unisys. I would add one thing to this transcript. I mentioned briefly, and we didn't follow up, that the Alpha business plan was a very daunting proposition, which might be why Ken felt he could not endorse it. What we said was that if Alpha was a roaring success, if we established the market we wanted to, if we got to 5 billion dollars in system sales within the first couple of years, DEC would only have to lay off an additional 25,000 people. If we failed, DEC would lay off 60,000, because that's what the numbers meant. That's what running on lower margins meant.

Gardner: And that's what this huge organization building outmoded products.

Supnik: And to me, the most amazing thing--

Gardner: <Inaudible>.

Supnik: --is that all the people who signed on to do Alpha, all of them knew this. We didn't keep this business plan a secret. So, they knew that their jobs, as well as lots of other jobs, were going to be in jeopardy, even if we succeeded. And that, to me, is a testament to the DEC spirit.

Gardner: Yeah.

END OF THE INTERVIEW