Oral History of Forest Baskett, Part 1

Interviewed by:
Len Shustek

Recorded October 28, 2016
Mountain View, CA

CHM Reference number: X7999.2017

© 2016 Computer History Museum

Shustek Okay, I'm Len Shustek, and it is October 28th, 2016. We're at the Computer History Museum with Forest Baskett, who has graciously agreed to sit for an oral history. Welcome, Forest.

Baskett Thank you.

Shustek In the interest of full disclosure, I should point out that I've known Forest for 40 plus years...

Baskett <laughs>

Shustek …dating back to the 1970s, when I was a graduate student at Stanford and Forest was a professor there, my thesis advisor. Why don't we start at the beginning? Talk about where you were born, and your parents, and your siblings, and early life.

Baskett I was one of two kids; I had a sister who was three years younger than me. It was a lower-middle class family. My dad had been a chemical engineer, and ended up mostly in sales-related jobs in oil industry-related businesses, given that that was Houston, and the oil industry was the big thing there. My dad was a huge influence on me. He would stop everything if I had a question, and try to answer that question, and in a way that would hopefully lead to other questions. He really wanted me to understand and explore the world. I think that was just a major influence.

Shustek Were you interested early in science and engineering? Did you have teachers who helped?

Baskett I think so. I was an ADHD kind of kid in school, as you might imagine. When I was in elementary school I got Ds in conduct, because I was disruptive. The teachers liked me and didn't like me, because I was frequently one of the first people to shout out the answer to any question that came up. It wasn't productive for the other students. When I was in high school, in fact, I had a math teacher who kind of took me under her wing, and thought I was a-- In those days people could start in half-year increments. When I started school I was what was called a "mid-termer". When I was getting to the point where I might graduate from high school, it's a problem, because then what do you do? So she said, "Well, what you should do is just go to summer school and graduate early, and then go on to college." She was very influential in recommending students to colleges. I ended up going to Rice, primarily from her recommendation. It was not a problem just getting in and going there.

Shustek Did your sister wind up going into technology as well?

Baskett No, she didn't. I was really raised by my dad, and my sister was really raised by my mom. My mom was more of a stay-at-home, homemaker kind of person. And my sister was kind of that way too. My sister ended up getting married and having kids, and being relatively not professional, but more of just a homemaker like her mom had been.

Shustek When you went to Rice, what did you major in and why?

Baskett Well, when I went to Rice, given my interest in things scientific, I started out as a physics major. I actually found that to be boring, and I didn't go to physics classes, I just read the book and went to the

 © 2016 Computer History Museum

exams. I wanted something more challenging, and so my sophomore year I decided to be a philosophy major. There were some really interesting philosophy professors at Rice, and in fact one of them, I think it's Alan Robinson, was a regionally well-known logician. I went through the whole proof of Godel's theorem as a sophomore; it was fascinating. It's a complicated theorem. But I quickly figured out that this was an odd place to be as a profession. So as a junior I became a math major, which I kind of knew how to do. I liked the analytical side of math, and at the time Rice had a really strong group of faculty in real analysis. That kind of fit in with the oil industry in the vicinity. But when I was a senior, the math department hired, I don't know, four or five algebraic topologists, and the whole character of the department changed just really dramatically. The real analysis people became second-class citizens. One of the famous ones went to the University of Chicago, or two of them, and started their careers there as opposed to being stuck with these weirdos at Rice. I thought about staying at Rice for graduate school, but because of that change I said, "Well this is not the right place."

The other thing that had happened at Rice, was when I was a freshman… I always worked part-time doing something, just to help support myself. The first year I worked in the library filing books. It was, ugh. <laughs>. But it paid. When I was a sophomore I decided that computers were really interesting. Rice at the time had a [IBM] 1401 and the first Rice Computer, which I think had started in 1957 is when that project started.

Shustek When were you there?

Baskett I was there from '60 to '65. I went to see a chemistry professor that I had had a class from and really liked, and asked him if he would hire me to help him with his project on the Rice computer. I didn't know anything at the time. This was the usual kind of ballsy thing. And he talked to me! He said what he was trying to do was simulations of the lattice gas. This is a model of gas where molecules can occupy lattice sites. Then you change the pressure and temperature, and see what happens. What he wanted to do was to make movies of phase transitions in the lattice gas. So he hired me to be the graphics guy.

Shustek Seems like a perfect fit, because it melds physics with mathematics…

Baskett It does, right.

Shustek …with computing.

Baskett It does. It was really fun. I had to learn how to program. I had to learn how to run that particular computer.

Shustek Were you programming in Fortran? Let me guess.

Baskett Mostly in assembly language actually.

Shustek <laughs> Okay.

Baskett There wasn't a Fortran compiler on that machine. There was a Genie compiler, and Genie was a lovely language. It was dramatically better than Fortran. It was an ALGOL-influenced kind of locally

 © 2016 Computer History Museum

developed language. The Rice computer had code words, so that it had descriptors and could describe logical segments of memory. The Genie language had matrices as primitives. You could multiply vectors and add vectors, and multiply matrices and add. It was just really quite..

Shustek The machine was a one-of-a-kind done at Rice? And the software..

Baskett Yes, it was one-of-a-kind.

Shustek …was also unique?

Baskett Absolutely, right. The movie-making equipment was interesting. It was a cathode ray tube with a 16-millimeter film camera attached to it. You could put a dot on the screen, and put other dots on the screen. Then you could advance the film by one click, and that was it.

<laughter>

Shustek The screen had persistence? You had to draw..

Baskett No, no it would just light up.

Shustek Okay.

Baskett So blink, blink, blink. But the film was exposing it, so you painted the image that you wanted, and then you advanced to the next image.

Shustek But dots only, not vectors?

Baskett Dots, yeah, not vectors. Right.

<laughter>

Shustek It's a challenge.

Baskett Yeah. It was fun. I had to learn the machine in great detail, and became really knowledgeable about computer architecture as a consequence, because the architecture of the machine was dramatically more advanced than most computers at that day in time. And it was for a long time, as a matter of fact. I eventually got to where I knew enough that I became part of the systems programming staff there. So I moved on from working for the professor making movies, to working for the computer staff.

Shustek What year was this? If you remember.

Baskett This was… '61 I think is when I started, and then went through '65.

Shustek So all during your undergraduate career.

 © 2016 Computer History Museum

Baskett All during my career at Rice. Naturally this really interesting new computer at the University of Texas really got my attention.

Shustek That was the 6600.

Baskett That was the CDC 6600. I think Texas had serial 10 or something like that.

Shustek Had you decided that you wanted computers to be your career, or were you still wanting to be a mathematician?

Baskett Yeah, I had, right. When I thought about graduate school, I thought a little bit about math at Rice, because I knew that. You know, do the same old thing. But I really wanted to be involved with computers, and I saw that computer science departments were starting to happen. Stanford had started one, UT started one, a couple of other places had. So they got the 6600…

Shustek This is UT Austin?

Baskett UT Austin. They got the 6600, and they started the computer science department. I said, "Ah, I'll do that." So I went to Austin, went to the computer center, asked for a job, which I needed, and got a job. I ended up, probably within two years, running the systems programming staff. I ended up having about 25 people working for me.

Shustek Which ordinarily would be a fulltime job, not leaving you enough time to be a graduate student.

Baskett Yeah, right.

<laughter>

Baskett Well, you do what you do. My professor -- the professor that I eventually decided on for my thesis work -- and a couple of us, three of us graduate students, had gotten together and decided that this primitive batch operating system that the 6600 was using could really be replaced with an actual timesharing system with online terminals. We went to Jim Brown, who was the early faculty member in Computer Science, and proposed that we build that system. He got the Computation Center to let us do that, and to think about that as their replacement.

Shustek So the impetus came not from professors, which is usually the case, but from the students in this case.

Baskett Exactly right. The three of us said, "Let's do this." We went to see Jim and said, "This is what we want to do, what do you think?" And he said, "Hmm, that could be fun. Let's try." <laughs>

Shustek Had you had any prior exposure to timesharing systems?

Baskett No.

Shustek <laughs>

 © 2016 Computer History Museum

Baskett No, not really. But I certainly knew about them. Project MAC was a big deal. Well, that's not true actually, now that I think about it. In '64 I had a summer job at SDC, System Development Corporation in Santa Monica, and SDC had a SAGE computer. I know you remember the SAGE computer.

Shustek Oh, yeah.

Baskett It was an interesting computer. In fact it had a timesharing system on it. That summer, that's where I worked and that's what I used. So actually I did know quite a bit about what made sense, and what you needed to do, and how they worked.

Shustek There's a SAGE computer directly beneath us…

Baskett Is that right? <laughs>

Shustek …in the exhibit. That's correct.

Baskett That's great. Really quite an interesting machine. I've actually forgotten all the details of what was special about it. One of the things I did that summer, because it was kind of an unstructured job, is I wrote a Lisp interpreter for that machine, because I got McCarthy's book, and I read it and said, "Hmm, this is really neat."

Shustek And it's very mathematical.

Baskett It's very mathematical. So I wrote a Lisp interpreter for that machine. That was fun. Looking at the 6600, that was kind of my model of what made sense. That's how people should do computing. One of the reasons that I think I was able to do that, was when I started work at the computer center, my first job was to write a billing system for the batch system that they were using. I wrote it in Fortran, which was what you did in those days. Not only did I write the billing system, but I was interested in exactly who the users were and what their characteristics were. I did a fair amount of data analysis on what the bills were like, and who was using what. As a consequence, I knew who the customers were, and I could talk to the center director about who was important. Issues would come up. The customers would always complain; you know, that's what customers do. I could look at the data and say, "Well here's how to minimize complaints."

<laughter>

Baskett I had the data, and that was powerful. That was really powerful.

Shustek That's an interesting optimization: "Minimize customer complaints", which is not the usual academic target.

Baskett <laughs> Right. Given that I knew so much about the customers, and knew how to really make operations run a lot more smoothly than they had, I ended up being in charge of all the systems people there. That's how that happened. In a similar vein, the timesharing system that we had in mind was aimed

 © 2016 Computer History Museum

at really making the system a lot easier to use. And a lot more enjoyable. And a lot more productive for the users. Something that they would complain about a lot less.

Shustek Who was the target audience? Were they students there, or researchers?

Baskett There were some students. But there were quite a lot of faculty who were trying to do computing, and supported the research. And graduate students who were working for the faculty, doing that same kind of programming.

Shustek How long did that project last? Was it successful?

Baskett I think we probably got the system running in a year and a half, something like that. Pretty quickly. The early versions were somewhat primitive, but it got better and better and better. That system, and revisions of that system, was actually in production for 10 years, which was astounding actually.

Shustek What were the terminals that people used?

Baskett We started off with Teletypes; that was the terminal of choice. They gradually got better, you know, it was that evolution. It was TTYs, typing kinds of terminals.

Shustek Could people compile and run programs, and do all of that? Did you have to write any of those components? Or did you use things that already existed?

Baskett We used the compilers that CDC supplied. We just made sure that they would run in that system. We had to have all the interfaces that they were used to, so that kind of defined what the requirements were for the operating system.

Shustek How did you deal with the lack of memory protection and other constraints?

Baskett The machine had memory protection, but it was on a per job basis. You got a block of memory that was allocated to you, and there was a base and bounds register, and that's what you got. But you couldn't maliciously or accidentally hurt anybody else. I remember Seymour Cray came to visit occasionally. I remember one visit, because we were deep in the bowels of building and changing and developing the operating system. I asked him, "Well can you change these systems so that when you execute a privileged instruction in user mode, it'll cause an exception, instead of being a no-op" And he said..

Shustek What was your motivation in asking that?

Baskett So that you could basically run operating systems as a user mode, and then take the exceptions and see if things were going right.

Shustek A virtual machine.

Baskett Standard virtual machine stuff. I knew people at IBM, and I knew about VM. So I said, "Well it's obvious, the right thing to do." And he looked at me and said, "No, I don't think so."

 © 2016 Computer History Museum

Shustek <laughs>

Baskett That was the end of that. That was fascinating. It would have been so much easier if we'd had had that one little change to the hardware.

Shustek Why do you think he was unwilling? Did he not understand why it was important?

Baskett He didn't understand why that was important, and why that would be useful. So he didn't want to bother.

Shustek Right. Did that timesharing system get used anywhere else other than UT?

Baskett No, but it was influential. CDC realized that they needed to do something like that. Livermore [Lawrence Livermore National Laboratory], which was a big user of the same machines, felt like they had to do something like that too. Several kinds of competing efforts sprang up. I think the Livermore system was the one that ended up becoming more widespread. But yeah, I feel like we had really shown the way; that you could do this and it was really, really helpful.

Shustek There was, as I mentioned before, a project at NYU Courant Institute to do something similar to that.

Baskett Mm-hmm.

Shustek That was at about the same timeframe, so I don't know who influenced whom.

Baskett I'm sure we were in communication, because we all talked to each other, went to the user groups, and those things. Although the details have faded.

Shustek So that was your fun project as an employee. What did you do next?

Baskett Once the system was running I continued my interest in collecting and analyzing the data from the machine. I continued my interest in making the system as productive for as many people as possible. One of the big issues in those days was job scheduling: how do you choose the next job to run? Because in a timesharing system you've got a queue of jobs, programs that are ready to execute. Which one do you choose and how long do you let it run? Somehow I figured out that. You know, timeslicing was something that timesharing systems almost always did. But that's only a very crude characterization of what you do. You've got a queue, but how do you sort it? The fact that you could sort it was an interesting observation. I spent many hours just over the big console, watching the system run. Trying to figure out what worked well, and when things got slow, and why. I observed that doing timeslicing in a round-robin fashion, with the slice being as small as possible consistent with overhead, was actually the most productive way to run the system. So that's what we did. That's how the system ended up running, and it ran really well.

Shustek Were all these tasks memory-resident, or did you have to do swapping to change context?

 © 2016 Computer History Museum

Baskett They were memory-resident. Just a huge piece of memory, jobs all laid out. You can move them around because of the base and bounds register architecture. I then had gone to a conference where I had presented some results about how the system had worked, and I met Leonard Kleinrock at the conference. I said, “Well I’m doing this kind of scheduling, and it seems like this is kind of the optimal way to run this kind of system. How do I analyze this stuff?” He recommended a small book on queueing theory to me. I went home, read the book, I said, “Oh, now, I see.”

Shustek Was that his book on queuing theory?

Baskett No, it wasn’t. It was a classic book that would have been, I don’t know, probably 10 or 20 years old. I spent many hours teaching myself queuing theory, and then developing models for how our system worked. Eventually I developed a proof that if overhead is zero, if you take the timeslice to zero -- the timeslice is epsilon -- you get minimal waiting times for all the jobs. Consulting with some other colleagues, one at UCLA in particular, we figured out that there were a couple of other versions of that theorem. We ended up with a pretty widely-cited paper in the Journal of the ACM. At first it was my thesis, then it became this really quite well-known, highly-cited paper.

Shustek Who was your thesis advisor?

Baskett Jim Brown was my thesis advisor.

Shustek Was he very helpful and involved?

Baskett He was helpful, but I wouldn’t say involved. Queuing theory was not his thing. He felt like I was doing the right thing, and he prodded me and poked me, and looked at drafts, and all that sort of stuff. But I was kind of really out there on my own, except for my colleague at UCLA that I worked with on it.

Shustek You mean Kleinrock?

Baskett No, it wasn’t Kleinrock. It was one of Kleinrock’s colleagues there, Dick Muntz. He was a younger person than Kleinrock. He was slightly older than me, but a young faculty member there.

Shustek When did you get your doctorate at UT? And what were you thinking about next?

Baskett End of ’70. Then in ’71, as I kind of finished up, I took a temporary job there as an assistant professor, from ’70 until ’71. I finished in December, and I think it got stamped by the university in January. Then the question is what to do next.

Shustek Exactly. Do you want an academic career? Do you want to do industry?

Baskett Right. I was already teaching at that point, in addition to running the computer center, which was challenging but fun. So that spring I interviewed for academic jobs. I can’t remember how many places I interviewed, but it was quite a number of places. There were lots of jobs; computer science departments were springing up everywhere. It was easy to get interviews.

 © 2016 Computer History Museum

Shustek Did you also consider industry jobs at the same time?

Baskett I didn't. I didn't know what industry could do for me at that point, or why that would be interesting.

Shustek You wouldn't want to work for CDC, for example?

Baskett No. I knew I didn't want to do that.

<laughter>

Baskett I had enough dealings with them to know that that was not a good idea. It turned out that one of my graduate students officemates had had summer jobs out in the Bay Area. He had become friends with a professor at Stanford that his wife had worked for. My graduate student friend called this guy at Stanford and said that I was interviewing for jobs, and if they had jobs they should talk to me because I was worth talking to. It was actually two of us. As I said, there were three of us that started the timesharing project, and we all kind of finished at the same time. So the three of us ended up often on the same interview trips together, when we were thinking about what to do next.

Shustek Who were the other two?

Baskett Herb Schwetman was one, and John Howard was the other. John and I ended up interviewing for a position in the computer science department at Stanford.

Shustek Who was the professor that you had been referred to?

Baskett Oh my god…

Shustek Do you remember who you interviewed with when you went to Stanford?

Baskett Well, I went through the whole rigmarole. I gave a talk to the department, and most of the faculty and most of the graduate students were there.

Shustek Do you remember how big the faculty was at the time?

Baskett Probably 10 or 12, or 13 people; 13 faculty members. It was an interview for a joint appointment, because Stanford had joint appointments between Electrical Engineering and Computer Science. There was a group called the Computer Systems Lab. Originally it was the Digital Systems Lab, but then it became the Computer Systems Lab. Ed McCluskey ran that lab; I don't know if you remember Ed.

Shustek Oh, yes.

Baskett They had interviewed a lot of people, and I think their first choice for the job. I turned it down and went to Xerox PARC [Palo Alto Research Center]. John and I both got offers. Because John was from Massachusetts and had been an undergraduate at MIT, and I was from Texas and was kind of a native Texas boy, all of our friends said, "Well, John will go to Stanford, and Forest will stay at UT."

 © 2016 Computer History Museum

Shustek <laughs>

Baskett Of course, exactly the opposite happened. John stayed at UT and taught there for quite a while, and I went to Stanford and became part of the Computer Systems Lab. Since there were two jobs, the other person that got hired, instead of John, was Vint Cerf. Vint and I were officemates in those early days.

Shustek Describe the atmosphere, and the machines, and the research projects that were going on at the time.

Baskett Well <laughs> this is a part that you'll remember well. When I first arrived, there was a still a Burroughs B6500 that was being used. I don't think I ever used that. There were minicomputers which had kind of showed up, and we had some of those in the lab. Of course there was a big IBM machine out at SLAC [Stanford Linear Accelerator Center]. And one more thing I forgot: there was the DEC PDP-10 at the AI Lab, which was out at Felt Lake. I ended up getting an office at the AI Lab and at SLAC, so that I could use those two machines.

Shustek At SLAC it was the [IBM] 360/91 computer.

Baskett It was the 91, right.

Shustek And it was the PDP-10 under McCarthy et al..

Baskett Under McCarthy and Les Earnest, and that whole crew [at SAIL]. Those were great days.

Shustek What did you do? What were you working on?

Baskett Well, I continued to do performance analysis, and thinking about what made systems work and why. And what kind of models you could construct that would indicate how they worked.

Shustek And teaching.

Baskett: And teaching, umm. One of the things that I had observed at UT, was -- and I've been struggling to remember the name of the company, it began with a "D" -- but it was one of the first CRT-based replacements for typewriters. It was a small screen.

Shustek: It wasn't Datapoint?

Baskett: Datapoint, yes. Thank you, it was Datapoint. I looked at that and I said, "Wow. This is where things are going." At Stanford, one of the things that happened at SLAC was that you and I connected. The Intel 8008 [microprocessor] came out, and Random Access Memory chips came out. All these things happened. Xerox PARC was building the Alto, and we decided that we could build something like that. I'd certainly had experience with CRT screens back when I was an undergraduate. <laughs> I knew how to put dots on a screen. <laughs> That just seemed like a fun thing to do, so we built a really interesting machine that had some architectural characteristics. I don't know if you remember how we structured the

 © 2016 Computer History Museum

memory banks so that the display would automatically do the DRAM memory refresh, so that you didn't have to have a separate refresh cycle.

Shustek: During the retrace intervals.

Baskett: During it, yeah.

Shustek: We called it the Video Graphics Terminal…

Baskett: Yeah.

Shustek: …and it had two modes. One mode that was text-only where it could store lots of text.

Baskett: Right.

Shustek: Then the other mode was a…

Baskett: Graphics.

Shustek: …graphics mode of 640 x 480 dots.

Baskett: It was 640 x 480, right. It was a P-39 Phosphor, because that had nice, long retention times, so you didn't have to worry about refresh rates very much. Yeah, those were fun days. It was kind of computer architecture on a small scale. What happened later, and I can't remember exactly how much later, a new graduate student in EE came to see me saying that he was interested in hardware projects. He was looking for a hardware project. He told me some of the things that he had done when he was an undergraduate at CMU. I said, "Well, you know, I have this thing in mind," which was kind of the... The [Intel] 8086 had come out as a processor, and it was dramatically more interesting than the old 8008 line. And of course memory had dramatically improved, and all these other things. I had become involved with Xerox PARC, and I knew how the Alto worked. I said, "Well, you know, I have a project in mind." So we sat and talked about it.

Shustek: And that graduate student was…?

Baskett: It was Andy Bechtolsheim. Right. Andy and I basically conceived of and designed most of the original SUN workstation as a project.

Shustek: Based on the 8086 or the [Motorola] 68000?

Baskett: Well, what happened was we started off thinking that it would be the 8086. But after we kind of started talking about it, the 68000 came out. I looked at the 68000, and I said, "Oh. Wait a minute." <laughs>

Shustek: “That's a real computer.”

 © 2016 Computer History Museum

Baskett: Reset. Yeah, this is a much better computer. So we switched to the 68000. I think we ended up with a 6810 as the original processor. I designed a memory-mapping system for that processor, which barely had the capabilities to do memory mapping. The original SUN workstation had a really fascinating memory mapping system. I was a little annoyed with Andy, because years later I discovered that he had patented that memory mapping system.

Shustek: <laughs>

Baskett: And not bothered to tell me about it or put me on the patent, because… And, you know, I had a paper that described it, so that was… But it didn't matter.

Shustek: That [paper] sounds like prior art.

Baskett: Yeah. It was--

<laughter>

Baskett: Right. But, you know, who knows that? It was no big deal.

Shustek: What was the attitude in the Computer Science Department and the Electrical Engineering Department to engineering-like projects like this? The Computer Science Department had been derived from…

Baskett: Oh, my God.

Shustek: mathematicians; George Forsythe and so forth.

Baskett: You know, I had a lot of graduate students. You were one of the best. Andy was amazing, although Andy left to form Sun before actually doing a thesis. He asked me to come with him as well, but that was not the right time for me. But the Computer Science Department had a history of hiring Assistant Professors in a systems kind of area, and then when it came time to promote them, they didn't make it. One of the things that I did was I took a leave of absence in my fifth year -- or sixth year, I can't remember --- and went to Los Alamos [Scientific Lab] which had just acquired that serial 1 of the Cray-1, to work on an operating system for that machine. And we did that. That was a really fun project. Not as successful as I would have liked, but it was a really interesting project and it came out with some pretty highly-regarded papers about about the structure and such.

Shustek: That system was called DEMOS.

Baskett: It was called DEMOS.

Shustek: What was unusual about it? What was special?

Baskett: The structure of the operating system was inspired by capabilities. But it wasn't hardware capabilities, it was software. It was a structure where things had tags, and the tags had meaning. You

 © 2016 Computer History Museum

could build general purpose utilities that reacted to the tags. It was a very clean, elegant way of constructing an operating system, file system, et cetera.

Shustek: Was the hardware sufficient? Had Seymour Cray learned any lessons from the 6600 days?

Baskett: Well, it was a lot better. Yeah, it was a lot better. But it was still a base-and-bounds kind of machine.

Shustek: Was it virtualizable?

Baskett: Yes, it was. I think it was; he had gotten to that point. We used that, as a matter of fact, in debugging the new versions of the system as we developed them. But anyway, then I was up for tenure. I was still coming back to Stanford and talking to my graduate students on a regular basis. I went to see Don Knuth. I said, "Don, you know my tenure case is coming up, and if you want me to stay at Stanford, you're probably going to have to get involved in this tenure case." I don't know if that was effective or not, but he did get involved, and I ended up being the first systems faculty member at Stanford to get tenure

Shustek: Wow.

Baskett: And then it got better.

Shustek: What was the department and the university's attitude in those days towards the connection between what their faculty did and commercial projects like Sun? We tried to commercialize our video graphics terminal.

Baskett: Yes. Right.

Shustek: We went up to Tektronix and we tried to sell it to them. But that's almost a commercial activity.

Baskett: That's a commercial activity. I think the department really didn't think about that much at all, other than it was a convenient way for faculty to consult, and to have summer jobs, and things like that. But not really in the way that we see it today, at all.

Shustek: Were they interested or aggressive about getting patents on work their faculty had done?

Baskett: No. Not at all. It never came up. I mean, you could do that, but it was just--

Shustek: You were on your own.

Baskett: You were on your own. There was the technology licensing office, but you had to kind of reach out to do that. The other interesting thing that happened was that Vint [Cerf] and I … The appointments were three year appointments, and so typically if you were going to go through the whole process you'd get two three-year appointments, and then you'd come up for tenure. At the end of our first three-year term Vint and I were both joint appointments in computer science and EE, and we said, "You know, this is kind of a double-jeopardy situation," which it was. Vint and I said, "Well, why don't we trade halves, and each of us be full-time in one department," so that we wouldn't have that. So we did. We went to the

 © 2016 Computer History Museum

various department heads. Vint said, "I want to be full-time EE," and I went to Computer Science and said, "I want to be full-time CS." They looked at us and said, "Okay."

<laughter>

Shustek: "Why not?"

Baskett: Yeah, they said to us, "Why not," because it didn't change anything. The departments cooperated really well.

Shustek: So now you have tenure. That's usually the introduction to a lifetime of academic service.

Baskett: It is. And there was a lot going on in terms of startups at the time. One of the things that happened to me -- it was a kind of a personal situation -- was that I had a daughter just before I went to Los Alamos. She got a little bit bigger in Los Alamos. Then we came back and she was still very young. My wife at the time developed bipolar syndrome, and I ended up being effectively a single parent, because she was so unable to participate in raising our daughter. One of the things that I tell people about being an academic is that it's a really hard job. My description is that it's really three full-time jobs. You have to be a full-time teacher, and teaching is very time consuming. You have to be a full-time researcher; that's what Stanford expects. And you have to be a full-time administrator, because there's a lot of administrative burden that goes with that job. So I took another leave, after a couple of years of being a tenured professor, and went to Xerox PARC. That's a fun place to be. A lot of smart people. Interesting projects. And extremely low stress, low demand. Your schedule is your own. You go to meetings, but there aren't that many of them, and they're fun. You work at your own pace. It was really wonderful to be able to have that kind of intellectual stimulation, but also to be able to get my daughter to a point where she was older and more able to make it.

Shustek: There were so many amazing things going on at Xerox PARC at the time.

Baskett: There were so many amazing things going on.

Shustek: What did you get involved in? What did you work on?

Baskett: Well, I met Bob Metcalfe at Xerox PARC. I met John Warnock and Chuck Geschke at Xerox PARC. Andy [Bechtolsheim] was taking the SUN project out, to do that.

Another thing that happened before I did that was: a guy had showed up at my office at Stanford; he said Bob Sproull had recommended that he come see me. I said, "Oh, Bob. Sure." I knew Bob from the Xerox PARC connection. So he sat down and said he was an assistant professor at UC Santa Cruz and he had this idea that he wanted to pursue, and he didn't think he could do it at Santa Cruz. He wanted to come to Stanford to work on it.

Shustek: And this person was?

Baskett: This was Jim Clark.

 © 2016 Computer History Museum

Shustek: <laughs> Yeah, all right.

Baskett: The project he wanted to work on was the Geometry engine, because one of the things that we had gotten involved in was Mead-Conway, the LSI revolution that happened about then. I think that revolution was really important, in that it enabled graduate students to build integrated circuits. Previous to that time, that was something that only specially trained wizards could do. That kind of democratization of that technology I think really changed the world in some really important ways. I'm not sure people realize how important that was, but I think it just caused integrated circuit development to explode as a consequence. And the whole ECAD world came into being as a consequence. One of the things that I did when I was on leave at Xerox PARC was to write ECAD programs for the integrated circuits that I was also working on. It was an amazing time. It really was.

Just to indicate how amazing it was: at Stanford, Jim had come into this research project that I was running. It was a big DARPA project, and he had started developing the Geometry engine as part of that project, which was very successful. He got to the point where he had a working system, a piece of silicon that did all of the 3D transformations that you needed for 3D graphics.

Shustek: Was he using the Mead-Conway system?

Baskett: Yes, absolutely. DARPA had a program where you could use Mead-Conway, and you could use MOSIS to get the chips made. You could build systems, silicon systems. At a university! <laughs> It was really quite remarkable. He got the original Geometry Engine and the original 3D graphics workstation running at Stanford with about six graduate students. I'd taken a leave at Xerox PARC and gone over there. Jim was over there quite a bit as well, for various reasons. And Jim said he was going to start a company to commercialize this technology, which turned out to be Silicon Graphics. He wanted me to join. I said, "Well, I'm busy being somewhat of a single parent. So I appreciate the idea, but I just can't do that."

Shustek: Did you think you would stay at Xerox PARC indefinitely, or were you going to go back to Stanford?

Baskett: I didn't know. I guess I thought I might go back to Stanford. But what happened was things were really changing, I guess, in my mind about what was exciting. Jim asked me to join his startup, Silicon Graphics. Andy wanted me to join him to do the SUN workstation and Sun Microsystems. Chuck Geschke and John Warnock called me into their office and said, "Well, we're leaving Xerox PARC and we're going to start a company and we'd like you to join that company."

Shustek: Adobe.

Baskett: That was Adobe. Bob Metcalfe called me over to his house and sat at the kitchen table with, I don't know if you remember, Ron Crane --

Shustek: He's a volunteer at the Museum!

 © 2016 Computer History Museum

Baskett: Yeah? <laughs> Right. Said that he was starting 3Com, and he wanted me to join that. And there's probably something I forgot. But it was…

Shustek: <laughs>

Baskett: …it was a crazy time. It was an unbelievably exciting, productive time.

Shustek: Weren't any of these appealing to you?

Baskett: They were all appealing to me. Except that I was really still seriously engaged in raising my daughter without a lot of help. I just didn't feel like I could leave that kind of safe environment for what I knew was a 24/7 kind of job, which startups are. But eventually my personal life got more predictable and schedulable.

One of the graduate students that I had had at Stanford was a guy named Sam Fuller. He was actually Ed McCluskey's graduate student officially, but we worked together and did a whole bunch of papers together. He then went off to CMU to be a professor. He was there for five years and had gotten to know Gordon Bell. At the end of five years he decided to join Gordon at DEC. So that was earlier. When I was at Xerox PARC, I was thinking, "Well, what do I do next?" My life is getting a little bit more under control. I still was in touch with Sam quite a bit. I said, "Sam, you know, [DEC] needs a research facility on the West Coast. You guys are missing out on everything that's going on out here." Which was tons. Sam said, "Yeah, that's a good idea." I said, "Well, why don't we do that?" So I started what we called the Western Research Lab for Digital Equipment Corporation. At that point I quit being an on-leave Stanford faculty member and visiting scientist at Xerox PARC, and became a full-time employee at Digital Equipment Corporation.

Shustek: What did your colleagues at Xerox PARC think? Are you starting competition for them?

Baskett: No, they didn't think that. They didn't think about competition, really.

Shustek: <laughs>

Baskett: They all wrung their hands about Xerox's ability to commercialize anything. They had a great job, but they didn't really think about… They'd gotten their fingers burned with the Xerox Star project, and so it was kind of not a big deal.

Shustek: Now, Xerox PARC had a sister organization that was supposed to commercialize some of the technology.

Baskett: Right.

Shustek: DEC did as well. Didn't they have the Systems Research Center? What was that? What was their relationship?

 © 2016 Computer History Museum

Baskett: Okay, that's an interesting situation. The commercialization unit for Xerox was a real live commercialization unit. When I was leaving to go to start the Western Research Lab, Bob Taylor, who ran the research lab at Xerox PARC, once he gave up trying to convince me to stay at Xerox, he said, "Well, Forest, can you hold a place open for me?"

<laughter>

Baskett: This is a true story, because things had gotten a little awkward at the corporate level at Xerox, and Bob was concerned about what their future might be. I went off and started the Western Research Lab. Then about two years later Bob wanted to talk to me, and said, "Things are just really falling apart. Is there anything we can do?" I talked to Sam, and together we decided that we should start another research lab for DEC in Palo Alto. That was the Systems Research Lab.

Shustek: What was the difference between the two?

Baskett: The Western Research Lab built hardware and interesting machines that were very relevant to what DEC was doing as a business. The Systems Research Lab was more abstract. They did a lot of interesting systems projects, but none of them were directly relevant to the mainline business of the company. They were more what might happen. It was really more researchy, whereas the Western Research Lab was more applied, I'd say. Which is kind of what had happened to me as I had developed various systems. I wanted to see them actually happen. I had gotten bitten by that bug.

Shustek: What was the track record? Did you wind up doing things that…

Baskett: Well, yes. Our big project…backing up a bit. Back at Stanford, when we started this big DARPA research program, one of the people was a new professor at Stanford, John Hennessy. He and I ended up being office mates a bit. When we started this project, I asked John to be part of the research team for this big DARPA project. When we got the award, John came to see me, and he said, "Well, what are we going to do?"

<laughter>

Shustek: What did DARPA think you were going to do?

Baskett: <laughs> Well, it was about VLSI. It was about kind of taking that whole Mead-Conway thing and moving it forward in a way that that was going to be helpful to the world and to DARPA in particular. I said to John, "Well, why don't we build a machine?" And John said, "Great. That's exactly what I wanted to do." <laughs> That was the MIPS project, as you know.

Shustek: Yet another company you could have joined.

Baskett: That's the one I forgot. Right. Because that happened too, I was invited to join MIPS .

 © 2016 Computer History Museum

Shustek: There were other activities going on -- as long as we've backtracked to Stanford -- that didn't involve doing VLSI but were still breakthrough architecture. I'm thinking of the S1 project of Curt Widdoes and Tom McWilliams...

Baskett: Right.

Shustek: …where they did a lot of design tools to create a high performance ECL machine.

Baskett: They did a lot of design tools, yes. That wasn't custom silicon, it was off the shelf silicon, but it was computer architecture, and it was building hardware. And they built the S1, quite an elaborate machine. Tom and Curt were graduate students of mine, and they did very well. When they finished their graduate work they started..,what was the name of their CAD company?...

Shustek: Was it Cadence or--

Baskett: No, no, no. It was after Cadence.

Shustek: Or Synopsis? We will think of it. <laughs>

Baskett: Yeah, we'll think of it... They built a bunch of CAD tools, and built a CAD company and it was quite successful. It was kind of what Cadence should have morphed into. It was Daisy? It was after Daisy. [It was Valid Logic Systems – Ed.] They used SUN workstations as their vehicle and they put software on that for ECAD. It was hugely successful.

Shustek: But that was a different world from the VLSI design CAD systems.

Baskett: Yes, it was. Well, no, it was ECAD. It included VLSI as well.

Shustek: Okay.

Baskett: But it was moved up a level, so that you could do complete systems with it. How did we get sidetracked on this?

Shustek: <laughs> We were going back to the Stanford and the work with MIPS.

Baskett: Ah, MIPS, right. When MIPS really got going, I had kind of drifted off into the Xerox PARC thing. I was really kind of an advisor to the team, although I did some of the early instruction set and design work on it. When that got going John wanted to know if I wanted to join that as well. I didn't do that either, but it was a great, great success.

When I started the Western Research Lab, the question was: what to do. When the RISC thing had happened -- the RISC project at Berkeley and the MIPS project at Stanford, obviously quite successful academic projects -- RISC had become something that Sun then started pursuing commercially. I said, "Well, DEC needs to really understand this, because the VAX is not a RISC machine, and eventually that'll hurt."

 © 2016 Computer History Museum

Shustek: Why not adopt one of the other architectures?

Baskett: Well, they were pretty primitive at the time. My thought was, in order to really influence a big momentum-heavy company that had the VAX and VMS being very successful revenue generators, that I would have to do something that made it clear that there was a better way to do things. DEC was working on an ECL version of the VAX, which they had never done before. My thought was, "Well, we'll do an ECL RISC machine, and it'll be faster and cheaper, and they'll notice."

Shustek: Most of the other RISC machines were not ECL, they were MOS.

Baskett: That's right, they were not. So we built what was called the Titan. I don't know if the Computer History Museum has any of those or not.

Shustek: I believe we do. They built dozens of them, right?

Baskett: Yeah. Many dozens, actually. It had about 60,000 gates and it ran at the same cycle time as the 8600 which was--

Shustek: A single chip processor?

Baskett: No, both of those machines were discrete. The 8600, the DEC commercial VAX, used ECL gate arrays.

Shustek: Right.

Baskett: And the Titan used just standard packages .

Shustek: So you weren't designing your own ICs?

Baskett: No, no ICs. But we had to make it work. You know, I ended up writing timing closure software <laughs> to figure out how to make it all work. The machine ended up being about twice as fast with about a half to a third of the components.

Shustek: Was the software incompatible with the VAX?

Baskett: But it was software incompatible with the VAX. It was incompatible in some easy ways. And then there were some stupid things that we did, that we shouldn't have done, that held us back. But nevertheless, it did get the attention of the DEC establishment. In fact, three independent competing projects sprung up in the various engineering organizations inside DEC.

Shustek: To build machines based on this design?

Baskett: No, to build RISC machines that were more compatible, But still, they were RISC machines. There was one which was an ECL machine, and there was another one which was kind of a more traditional TTL machine. And then there was one that came out of the people who were building integrated circuits for other parts of DEC, so they decided they wanted to build a complete

 © 2016 Computer History Museum

microprocessor. They eventually all killed each other off. But they were more clearly in line with what DEC was normally doing. Titan was simply something that caused that to happen, but didn't have a legacy of its own other than that.

Shustek: But did DEC decide on none of these RISC machines so far?

Baskett: Well, the design that came out of the VLSI group -- after they killed each other off -- got reborn and became the DEC Alpha. It was slightly different from what they'd done in the first version, but it was very good. A very good machine.

Shustek: What other projects were there at Western Research Lab in addition to Titan?

Baskett: What happened at that point was, we got to the end of that project and the question was what to do next. I was a little frustrated by what had happened with respect to these competing projects that then all killed each other off. It just seemed like such a waste, and so the corporate bureaucracy was beginning to bother me.

Shustek: Similar to what had happened at Xerox PARC.

Baskett: Yeah, right; similar to what had happened at Xerox PARC. A startup contacted me about trying to hire me to be part of their startup. It was a high-end graphics workstation thing. The guy who was running it was... I listened because I didn't quite know what to do next, and I was discouraged with DEC. So I talked to his founders, and what they were doing. It didn't seem like it was going to be the right thing, in terms of being a commercial success, to me. But the guy who was the lead guy, the CEO there, was at a party, and he thought that he was going hire me. He was basically bragging about it at this party. One of the people at the party was Jim Clark, and Jim said to himself, "He can't do that." So Jim calls me the next day and says, "Look, I know you turned me down the first time, but we're going great guns. You really should reconsider and come to SGI."

Shustek: Can you say who the competition was?

Baskett: Well, I'm trying to remember. It was the company out here that merged with the one in Massachusetts, you know, high end Apollo. Now the Massachusetts one was the one that came out of Apollo, and out here. Jesus. Can't remember the names.

Shustek: Okay.

Baskett: So Jim calls me the next day, comes over to the house and talks to me.

Shustek: If I got the dates right, this is about 1986. You had spent four years at the Western Research Lab.

Baskett: That's right. Silicon Graphics was four years old, as a matter of fact, as well, because all those startups happened around 1982, this amazing year.

 © 2016 Computer History Museum

Shustek: What were you hired to do? Was it an easy decision to go to Silicon Graphics?

Baskett: What happened was they were having a hiring open house the following weekend, which is something that companies sometimes do. They were shipping machines, and they had customers. I went to the open house, and talked to a lot of people, and looked at what the machines were like.

Shustek: They were [Motorola] 68000-based.

Baskett: These were 68000-based.

Shustek: Graphics workstations.

Baskett: Graphics workstations, that's right. I said, well, this is really, really cool stuff. This company can be successful, and they're doing interesting, interesting work. They wanted me to be a head of R&D kind of person, almost a CTO, although Jim was really the CTO. I thought, well, this is cool. These people all are on the same wavelength, you know. They're not playing politics. They're building systems. They're satisfying their customers. They've got a clear vision. Sun was still black and white [displays] at the time, for example, and this was not only color, but it was also 3-D. Sun was mostly black and white and 2-D.

Shustek: What stage in their corporate life were they? They had venture capital investment, but had not gone public?

Baskett: They had venture capital investments, yes. Mayfield and NEA were their lead investors, and they were getting to the point where being a public company was on the horizon. I joined, and probably six months later we started organizing for an IPO. We did that IPO; we became a public company. That was great.

Shustek: Good timing for you?

Baskett: Good timing for me, yes. Jim made that clear, of course. Then I had a very long career at Silicon Graphics: thirteen years. It was an interesting situation. It wasn't perfect from a corporate perspective, because Jim is kind of a volatile person. Really super creative, and insightful, and thoughtful, but also kind of emotional. So the board early on, almost from the beginning, had teamed him with a CEO. He was the chairman, and there was a CEO.

The first CEO was not terribly great at running such a dynamic business. The second CEO was Ed McCracken, who was hired away from HP. Ed was a really good CEO, actually. Really good. Knew how to build and run an organization. But Ed and Jim just didn't really get along that well. I mean, they worked together, and they collaborated, and they talked about the directions, and what the next product might be, and all those things. But they really didn't get along that well. Jim would always complain about Ed, and Ed, in his own quiet way would complain about Jim. I got along with both of them, so I was kind of an intermediary between these two explosive personalities. I did that for years and years.

Shustek: Now, you were in mainline product development. You were now not doing research.

 © 2016 Computer History Museum

Baskett: Well, that's not exactly true. I didn't have line responsibilities. But I was involved in all of the projects that were going on as kind of a "let's make sure we do this right". You know: who do we need to hire, and what are the technologies we need to incorporate, or not incorporate. I remember, after I'd been there about six months or a year, the guy who was the official VP of engineering came to see me one day and said, "Well, I've got this guy. I just don't know what to do with him. He's a really talented guy, but I just can't manage him. Any chance that we could have him work for you?" So I became the zookeeper, and I ended up having about twenty-five direct reports.

Shustek: Who was that person that he tried to give to you?

Baskett: That was Paul Haeberli. I don't know if you ever ran into Paul, but an unbelievably creative guy, technically talented but extremely artistic as well. He could make our machines do things that were just amazing. It was really wonderful, and delightful, and extremely productive to have around. But you had to let him do what he wanted to do. You couldn't ask him to do anything, and you couldn't tell him to do anything.

Shustek: I've known many engineers like that.

Baskett: Exactly, exactly.

Shustek: Did you take people from Stanford, from PARC, from DEC?

Baskett: Yeah, all those things. Absolutely, right. Silicon Graphics became one of the exciting places to work in Silicon Valley, and we ended up getting up to about, I don't know, thirteen, or fourteen, or fifteen thousand employees at the peak. It was really quite a phenomenon. It was a fun place to work.

Shustek: Let's talk about the products and projects that you were involved in. When you came there, they were building graphics workstations based on the 68000.

Baskett: Based on the 68000. We did a couple of iterations. When I came, the second generation product was kind of underway and getting ready to come out.

Shustek: That was Iris, or what came after Iris?

Baskett: Yeah, the first one was Iris 1400. The second one was the Iris 2000. I became heavily involved in the Iris 3000, which was the third generation product, and that was a big success, because we had smoothed off all the rough edges, and made things work really well. We had an excellent version of Unix running on it, and a really good file system. It was very high throughput, and had all the compilers that you might need.

Shustek: They were expensive machines, right?

Baskett: They were expensive machines, right. It was interesting how we designed the machines, and this was the design philosophy for the first four or five generations. The starting point was 1500 watts, because that's how much you can pull out of a high end wall plug. So the question, in terms of designing

 © 2016 Computer History Museum

a new machine, was what can we build inside that power envelope? VLSI was happening, and chips were getting more and more powerful, and more and more power efficient at the same time, so you could build bigger and faster and more powerful machines inside that power envelope in successive generations. That was where you started, and then you partitioned up the power into how much memory, and how much I/O, and what kind of CPU, and what kind of graphics.

Shustek: These were all based on Motorola 68000?

Baskett: They were, but MIPS was underway at the time. One of the things that I did was to lead the charge on switching from Motorola to MIPS, because the MIPS processors were about twice as fast and the same price.

Shustek: Had you been making multiprocessors out of the Motorola parts?

Baskett: No, the Motorola systems were uniprocessors, and the first MIPS processors, the Iris 3000s, were also uniprocessors. The thing that followed the 3000 was the first multiprocessor. I was heavily involved in the architecture of that, making that work.

Shustek: Did you worry, now that Silicon Graphics was a big company, about depending on single source vendors like MIPS?

Baskett: Well, yes, quite a lot. In fact MIPS became threatened as a company, because there was competition. DEC Alpha was one. Eventually we said this is really critical for us, and so we arranged to buy MIPS. MIPS became part of SGI.

Shustek: That was an expensive acquisition, in 1992, for three-hundred-some-odd million dollars.

Baskett: Is that right? I've forgotten the details. Yeah, cool. But it worked. It worked. One of the things that we did, after we had that team, was we started a very high end MIPS project. An ECL MIPS machine, but custom silicon. That was the R8000, a very, very effective, powerful machine. Tremendous floating point capability.

Shustek: Was that a successful project?

Baskett: It was a successful project. Yes, it was. It was short lived, didn't last long, because…

Shustek: It was a captive market. It was only used in SGI machines.

Baskett: …it was only used in SGI machines.

Shustek: What was the calculus about buying MIPS, which sold chips to other people who might have been competitors with SGI?

Baskett: Well our mission, and we were pretty true to it until we started building big multiprocessors, was graphic systems. 3-D graphic systems. We were fairly confident that no one else could build 3-D graphic systems like us. So selling MIPS chips for other purposes didn't seem like it was a problem.

 © 2016 Computer History Museum

Shustek: So that other business continued on after MIPS was acquired.

Baskett: It did. But I think it wasn't as vital as it had been, because it kind of tailed off.

Shustek: What happened after the MIPS acquisition? What was happening to the business? What was happening to the marketplace?

Baskett: We began to build all of our systems using MIPS chips. We had a great VLSI team, and they built great processors. We could get them fabricated at TSMC or wherever, so it was just part of our supply chain, so to speak.

Shustek: Were you building custom graphics chips, as well?

Baskett: Absolutely, yeah. We'd been doing that for a long time.

Shustek: You must have been watching what was happening to PCs and their graphics chips.

Baskett: Yes, right.

Shustek: Which market did you try to be in? Were you interested in higher end and lower end markets, or was it the middle market that you were interested in?

Baskett: We had pretty much of a stranglehold on the high-end graphics market. The aerospace companies, the auto companies, and all those people, they just bought our machines. We developed what we considered low end machines. Not PCs, but dramatically smaller and less expensive machines, so we had a low end. We even ended up with a team doing midrange machines, which was a questionable strategy, but we did it for a while.

One of the things that happened during that process was PCs became very powerful, and that was noticeable to all of us. One of the engineers said that he wanted to build graphics for PCs. His notion was that we could take the graphics that we had built for our low-end machine and simply adapt that graphics to a PC, and make a graphics card. We let him do that, and he built a graphics card. The original version of it was in a special deal with IBM using the microbus [Micro Channel], because they had their PC that used microbus, and nobody else used that. And he said, okay, we can easily adapt this to an EISA bus and then sell it to Dell or whoever. And we did that.

Shustek: This was called the "visual workstation?" Is that right?

Baskett: Yeah, right. The microbus version went into the RS-- what was the IBM RISC machine called, RS something, you know, out of Boston. [RS/6000] It went into that, and that was a successful product line. Again, they were newcomers. They didn't have all the software. They didn't have all the applications. They didn't have the ISVs, and so we didn't feel threatened. It was a nice contract, and it got us into a different kind of business. When we ended up with the EISA bus card, one of the things that happened just seared into my mind: the CEO said, "Well, this is basically a low margin business."

 © 2016 Computer History Museum

Shustek: CEO McCracken.

Baskett: McCracken. “This is a low margin business, and we’re a high margin company, so we’re going to stop doing that.”

Shustek: Did you disagree with that?

Baskett: I did. But it was a business decision, and I wasn’t the business person. What then happened was that a lot of engineers, who could see where the world was going, started leaving the company. There were about thirty-five, at one count, startups in Silicon Valley to do low end 3-D graphics.

Shustek: NVIDIA being one of them.

Baskett: NVIDIA being one, yes. There were thirty-five of them, and they were mostly populated by at least one, and sometimes four or five, SGI people. It was really tragic.

Shustek: The company must have seen that happening. What was the reaction?

Baskett: Well, yeah. There wasn’t a strong reaction, partly because we didn’t quite know what to do. One of the things that we did do is we spent time talking to Intel, and we did develop an Intel-based version of our low end graphics workstations. But the whole software environment was so different, it was hard for that to succeed, because it had to compete head to head with the rest of the PC business.

Shustek: But you had been selling or licensing software into the PC industry anyway. I think you sold, or licensed, the Iris graphics library to Microsoft for NT. The software you developed was getting out.

Baskett: Yes, we did believe that the graphics library needed to be a standard that would be part of growing the whole ecosystem around 3-D graphics. One of the things we did, and I was very heavily involved in that myself -- spent a lot of time in Redmond – was figuring out how to license the graphics library to Microsoft, and we pulled that off. OpenGL was what the result was called, and OpenGL is still supported by Microsoft to this day. It was a big success. I think that was absolutely the right thing to do. You had to build a bigger software environment, and you couldn’t have a closed system.

Shustek: You were doing that in part by acquisitions: Alias Research, Wavefront Technologies, Paragraph.

Baskett: Yeah. Those were more at the higher end of that graphics. OpenGL was the basic 3-D graphics library, and Alias and Wavefront were well beyond that in terms of the tools and the systems that they provided. They really served our high-end market, which was where most of our profits were. That’s where the real margins were, and the captured customers, so to speak.

Shustek: In the meantime, some of these competitors at the low end were using technology that was quite similar. Didn’t you sue NVIDIA over some patent infringement?

Baskett: Oh, probably.

 © 2016 Computer History Museum

Shustek: You weren't involved in that?

Baskett: I wasn't involved in that, right.

Shustek: Okay. You were making these kind of moderate- to high-end graphics workstations, and seeing competition come from below on the low-end PC end. But you were also getting into the high end supercomputers.

Baskett: We were.

Shustek: How did that happen and why?

Baskett: Because that was kind of a natural evolution. We had so many customers who were doing serious simulations, and associated 3-D graphics renderings of those simulations. Because of things like MIPS, we could build processors that were really state of the art in terms of their computational abilities. A lot of our software vendors would do the simulations on the workstation, and have interactive displays of those simulations. It was how we think about computing today, that it's interactive. It's not a batch thing, where you do a simulation, and then you do rendering. So we then we started selling systems primarily for computation, because they were so good at computation.

Shustek: Like the Power Challenge.

Baskett: Yes, exactly.

Shustek: That RISC multiprocessor could be half a million dollars.

Baskett: Right, exactly. We said there's a big market for this kind of computational capability, and we started building successively bigger, more connected, and integrated multiprocessor systems. It became really a big business for us. We started competing in the supercomputer world. That had some disastrous consequences when we acquired Cray Computer.

Shustek: Who were you competing with in the supercomputer world, and why was Cray thought to be a good strategy for dealing with that?

Baskett: Well, we had tremendous market penetration in the government labs and in the agencies because of this high-computation high-graphics integrated system. It was easy for us to understand what their computational needs were, and how much I/O they needed, and what kind of I/O subsystems made sense. And to design and build those things, and then have a fairly ready market for them.

Shustek: Was Cray one of those competitors that you were dealing with by buying, or were they a supplier?

Baskett: Cray was not really a competitor, because it was kind of a major step above us. But we ended up in some cases building thousand-processor systems for labs that were an alternative to Cray's. So in that sense, there was some competition. The Crays were big, honking machines, and we had these

 © 2016 Computer History Museum

clusters of very powerful and highly integrated multiprocessors. My argument was that the microprocessor engine was on the CMOS power curve, and that the other technologies just couldn't compete. If you just did the arithmetic, the Cray kinds of technological approaches just didn't make sense.

Shustek: Who was the advocate for the Cray acquisition?

Baskett: Well, to a certain extent I was. Not in terms of pushing it. What happened was Cray really got to a point where they needed help financially, because they had spent so much money doing development, and their sales were limited. The thought was that we could integrate a complete product line, and this would make sense. In theory that was an okay idea. In practice, the Cray culture was so different that it kind of meant the end of SGI as we knew it.

Shustek: Do you think it was the Cray acquisition which was it?

Baskett: I do, yeah.

Shustek: Financially it was a disaster. It was a seven hundred and forty million dollar acquisition in 1996, and four years later it was sold for thirty-five million.

Baskett: Yeah.

Shustek: And in some sense it wasn't the business that SGI was in.

Baskett: That's right. Right. Well, that's not true. Our computation business was at least half of our business by then. You have to really look into the business to see that that's not an accurate statement. We were in that business.

Shustek: What you're saying is that what killed SGI was a bad move into the high-end supercomputer business, as opposed to the conventional wisdom is that it was killed by the low-end graphics PC progress.

Baskett: Yeah, right. Well, both of those things happened. The PC started eating away the low end kind of early on. We pretended to participate, because we had an Intel version of our workstations, but it was so brutal down there. The alternative path for the company would have been to really dive into the low margin, high volume PC graphics business.

Shustek: To become NVIDIA.

Baskett: Become NVIDIA, exactly. That's what we should have done. It's easy to see that in hindsight, but we didn't. We took the high road. That's what Ed wanted to do, and it's part of why he got fired. Got to the point where anyone could look at it and say this was a mistake. So the board said, <snap>.

Shustek: In retrospect, what could people at SGI have done differently?

 © 2016 Computer History Museum

Baskett: Well, we could have not done the Cray acquisition, maintained our coherent approach to high-end multiprocessing, and maintain that business. There is a business problem, though, in blending a high margin, high-end business with a low margin volume business. That's hard for companies to do, so I'm not sure that one company could have done that.

Shustek: Maybe the right answer would have been to split the company.

Baskett: Might have been, and that's not something that we thought about in those days. But today we would say, yes, that's the obvious thing to do. Rather than kill the graphics card business, spin it out and form a technically integrated but business-wise separate entity. That would have been absolutely the right thing to do.

Shustek: So sales were plummeting. This was in '96, '97, something like that, and McCracken leaves.

Baskett: Mm-hm.

Shustek: And another HP veteran comes in, Belluzzo.

Baskett: Yes.

Shustek: You stayed through that era, as well. What happened then?

Baskett: I stayed for two years with Belluzzo in place, because I was trying. I thought, well... It didn't make sense to me in the beginning, but I had a lot of friends there. People that worked for me, and people that I just… They're really tight, close colleagues, and I wanted to kind of keep the ship running, so to speak. But it got to the point where this is ridiculous. This is not going to work.

Shustek: Any thoughts of taking some of those good people, colleagues, and starting a company on your own?

Baskett: No. I didn't. I didn't have an idea that was burning in my head that I wanted to pursue. I knew a lot, and I had been exposed to all this technology. One of the things that I did at SGI, because of my position, was I was involved in almost every technology discussion that we had. There was lots of new technology that was bubbling up, so to speak. Little companies would come to us and say, "We can do this. You ought to adopt this as part of your system." I was in ninety-nine percent of those meetings, and I saw all of that stuff all the time.

Shustek: SGI had done a number of acquisitions to solve these problems.

Baskett: And SGI had done a number of acquisitions. These were mostly not acquisition targets. These were "buy our product, our subsystem, or our component" kinds of discussions. One of the things that happened was that after these meetings, one of my colleagues would say to me, "Wow, you're really good at that. You should be a VC." It happened over and over. So when I got to the point where it was clear that SGI wasn't going to work and I couldn't do anything about it, I did resign and decided that I'd go out and see if I could be a VC for real.

 © 2016 Computer History Museum

Shustek: We'll get into that in a second, but just at the final end of your time at SGI, there were interesting decisions being made, right? They were dumping MIPS and signing on to the Itanium bandwagon.

Baskett: Yeah.

Shustek: Why was that?

Baskett: Well, Intel pushed us to do that. We had built x86-based workstations, and Intel encouraged us to think about doing that. There were a small group of people at SGI thought they could do it. That was a short lived project, I think.

Shustek: Had Itanium been successful, would that have helped save SGI?

Baskett: It might have, but I never believed that it would be successful. Never. Too bizarre. You know what I mean.

Shustek: So you decided to become a venture capitalist, which is a very different kind of activity, because you wind up giving people advice and watching them make mistakes, and maybe be frustrated that you're not in there helping them fix it.

Baskett: Yes and no. You know, my joke, even today, is that being a venture capitalist is a lot like being a professor. The difference is that when you give people advice, they tend to follow it.

Shustek: Because you hold the purse strings.

Baskett: Exactly. That's not completely true, but it usually gets a laugh. But it is a case where you get to see lots, and lots, and lots of new ideas from very creative, mostly young people, and you get to work with them, some of them, about developing those ideas. It's really, really terrific.

Shustek: Obviously you enjoy it. You've been there seventeen years.

Baskett: Seventeen years, that's right, right.

END OF THE INTERVIEW

 © 2016 Computer History Museum