SPARC Microprocessor Oral History Panel
Session Two
Business Development

Bernard Lacroute, Stephen L. Diamond, and Wayne Rosing

Moderated by:
Rosemary Remacle

Recorded: February 10, 2012
Mountain View, California

CHM Reference number: X6383.2011

Rosemary Remacle: All right, this is Rosemary Remacle on February the 10th, 2012. We are taping with the SPARC second panel. The first one was taped on June 9th, 2011. It included Anant Agrawal, Bill Joy, Robert Garner, David Patterson. That focus was on product and technology. The focus of this discussion will be more the business development of the SPARC Microprocessor. So we have with us today Steve Diamond, Bernie Lacroute and Wayne Rosing. And I would like each one of you to introduce yourselves, give a little bit of your background at Sun in particular, but then what have you done since then? Steve, why don't we start with you?

Stephen L. "Steve" Diamond: Okay, well, thank you, Rosemary. And I'm very excited to be here. I joined Sun from National Semiconductor, where I was running the 32000 Microprocessor Business Unit, and I was hired by Carol Bartz to spin up the SPARC architecture marketing program.

Remacle: Your official title was?

Diamond: I think I was Director of Component Marketing, which was kind of a radical title for Sun, who didn't make any semiconductor components in those days.

Wayne Rosing: Or marketing. <*laughter*>

Diamond: That's true. Well, the marketing people told me when I joined the company that they didn't really do marketing. They let the Executive Committee make all the pricing decisions, which was an interesting approach. I ran the SPARC architecture marketing program, and after SPARC got started, the board level business at Sun reported to me. After that, I was responsible for the industry standards group Sun, where we followed up with the open systems strategy of the company, and actually engaged with standards development organizations. After Sun, I went back to National Semiconductor where I was responsible for mergers and acquisitions. Then I joined a VLIW signal processor startup called Equator Technologies as VP of Marketing. Equator raised a couple of hundred million dollars in building a video optimized digital signal processor chip. After that I was elected President of the IEEE Computer Society. Then I went to Cisco and built their first Amazon Web Services compatible cloud. And after Cisco, I joined EMC. I'm in the office of the CTO. I'm the Global Standards Officer, and I'm GM of the Industry Standards Office. I am responsible for EMC's IP in relation to standards development organizations and standards strategy, in particular, around cloud computing and big data.

Remacle: Bernie?

Bernard "Bernie" Lacroute: Good morning! My name is Bernie Lacroute. I joined Sun in 1983. Before that I had spent 14 years at Digital Equipment Corporation in Maynard, Massachusetts, and was

instrumental in the VAX program as well as making Ethernet a standard, working with Intel and Xerox. I joined Sun in '83 as VP of Engineering. Then within less than a year, I became Executive Vice President in charge of Engineering, Marketing and Manufacturing, and was promoted to the board of Sun Microsystems as well. I left Sun in 1989, took a few days off, then joined Kleiner Perkins where I became a general partner from 1990 until about 2000, just at the time of the .com collapse. I decided I was going to go back to my roots. I was born in Burgundy, so I decided to spend pretty much fulltime on my winery in Oregon making Pinot Noir, Pinot Blanc, Pinot Gris. And my activity ever since has been to be the Chief Wine Drinker at WillaKenzie. That's how I spend my time, traveling, drinking wine, and a few other things of the sort.

Remacle: Sounds like just a terrible life!

Lacroute: It is! It's very taxing!

Remacle: Wayne?

Rosing: Hi, I'm Wayne Rosing. I joined Sun in 1985 after a phone call from Bernie. Bernie and I had worked together at Digital Equipment in the late '70s, and I was involved with the Intel/Xerox/DEC Ethernet stuff. And also was in charge of mid-range VAX development. And it's a curious little story, but Dave Patterson was going to be at Digital for a sabbatical, and due to changes of funding, it turned out I had the money. And so Dave joined my group, and that's where I began to understand the power of RISC. And we were looking at the idea of building a very high-speed VAX as an internal VAX, not an external product, that would be designed to support CAD tool development, because even in '78/'79 we were realizing the scale of what we needed to build. We needed CAD tools that didn't exist and computers that were Cray Class speed to do the CAD we wanted. I was essentially in charge of most of the SPARC program from '85 to basically about '89/'90. And then I moved on in Sun to Sun Labs, and then started the JAVA project. Then in '94, I left Sun, took some time off, did some personal stuff, then in 2000-2005 I was at Google as their VP of Engineering. And now I live in Santa Barbara, and am running an astronomical research organization, establishing a series of telescopes around the southern part of the world.

Remacle: Okay.

Rosing: So, I look up.

Lacroute: I look up.

Remacle: Let's start with a context setting discussion, and Bernie, I'm going to start with you. Describe where Sun was in the hardware, workstation hierarchy in Silicon Valley and the United States and the world at kind of the beginning of '83, when the discussions about whether or not to have their own microprocessor began.

Lacroute: Well, that's a very important point to set. The decision to go with a separate architecture was born sometime in the Fall of 1984. And Vinod Khosla was actually the individual who came up with the idea, and for a long time, championed it. Now what was Sun like at the time? It had finished the 1984 fiscal year at the end of June at about ten million dollars in revenue. And not much profit. Cash position, relatively low. Competition, extremely tough. Apollo, in those days, really dominated the workstation business. Sun had just recovered from the difficult times of building the Sun-2, which had significant engineering problem, as well as the CRT, which was not very reliable in those days. We wanted a 19" screen, and nobody even made those things. So they were kind of flaky. Some of them had an MTBF of about two hours. And so we had to go and fix them, and there's a long story there. But anyway Sun was a very small company.

Remacle: How many employees?

Lacroute: Oh, I don't know, maybe a couple hundred, 250 or something like that. Very, very small, even compared to Apollo. And certainly very small compared to Digital Equipment and so forth. So the engineering staff, there were a few very bright people, but it was pretty thin. There was not a deep pool of talent. Nobody had ever built a microprocessor, or never mind built any chip at Sun from scratch. So it was a really, really very risky, extremely gutsy decision. In the early day, Vinod really championed that idea, and I joined his bandwagon. And eventually after that Bill and Andy and then Scott rallied around that. But I remember the first time we went to the board and proposed this idea of building our own microprocessor, the board, and particular Dave Marquardt said, "You guys are nuts! You can barely make a 68K machine work, and now you want to build your own microprocessor?! It's craziness! It's foolish! It will never work! You guys are nuts, you don't have money; you don't have the talent. You are trying to take on something which Intel has taken years to go and build. That's foolish." So that was the context. Now why SPARC, why a new architecture? Because fundamentally the issue is how can you beat your competition? And you cannot beat your competition if you have the same product that they do! It does not work. You can survive, but what's the differentiation with Apollo, if all use the same microprocessor?

Remacle: So what was SPARC going to do that would allow you to differentiate your workstations from Apollo and DEC and so forth?

Lacroute: Speed. Originally ten times faster than the other guy at the same price. How is that for an idea?

Remacle: Nice value proposition.

Lacroute: That's a value proposition. But there was also a very important factor, which Wayne and I were discussing the other day, the idea of being able to port the software. I mean, those platforms were based on UNIX, and the idea is that the applications, the customer applications, could be ported from one architecture to the other in a somewhat painless fashion. And that, too, is a big part of the value proposition, because if you can guarantee that the customer investment in software will survive, and you can provide a machine at the same price as the other guy, and X-times faster, at least five times faster, then you're going to win. And that was the bet.

Remacle: Because they're important later on in the SPARC story—well, in the beginning they were, too—but let's be clear, who was sitting at the top of the company who had to bless this decision? You mentioned Vinod Khosla, but his title at the time was?

Lacroute: Well, I'm not quite sure if Vinod was still on the board. He no longer had an operating role in the company, but he was still on the board of the company. So ultimately, I mean, the way those decisions were made those days that was Scott, Andy, Bill and myself, we basically made the first recommendation to the board, then the board had to bless it, because it was a “bet your company proposition.” But I really want to stress the role that Vinod played in pushing that, because it sometimes not well-understood, and he is really the guy who said, "You know, you’ve got to go and do something different." And he was right.

Remacle: Do either of the other two of you want to add to that discussion from where you were sitting in the company?

Diamond: Well, I think one issue around SPARC was that it was a RISC architecture. And as Sun had done in the early days of open systems, it made a virtue out of necessity. Clearly, the company could not have, at that point, designed a CISC architecture processor. So RISC was just coming to the fore. It was gaining a lot of market credibility. The CISC vendors, and I was—before I joined Sun, I was at a company with a CISC processor, National Semiconductor—were really not even able to adequately design and simulate those CISC chips. We had, at National, at the time a significant percentage of the total compute power in Israel, where the design center was located, and they still couldn't simulate the entire chip set. So it was really impossible to build a CISC chip at Sun. But RISC provided an opportunity, an entrée to change the game entirely, and compete with the CISC vendors, because of some technology trends that were happening at that time.

Rosing: The other driver was back in the—I think it was around '77/'78, Carnegie Mellon, I think, and maybe MIT, issued a notion, a request for a product called the 3M Machine, which was a megabyte of memory, a megapixel of display, and a megabit of Internet. Excuse me, a megabit of networking. It was

sort of implied to be TCP/IP, but it wasn't 100 percent clear then. And that machine was sort of a destabilizing request. I remember we actually prototyped one at a technology fair at Digital. And a certain person I am sitting right next to said something about, "That's not quite in our strategy." <*laughter*>

Lacroute: That was not! <*laughter*>

Rosing: But out of that request, ultimately Apollo really found and serviced that market. But Sun came in with this Sun-2, and just had a sweeter design center. And so, even though Sun was one of the last people in what you might call the UNIX workstation business, they ultimately were triumphant, because I think they got that design center right, and in terms of what the company was, what the operating system was. And then the RISC thing was very simple tactically. It was, "How do we execute our software fast with the simplest possible system?" While everyone else was basically designing chips, if you will, in the sense of traditional design, Sun attacked it with a small number of gate arrays. And just solved the immediate problem of building the precisely correct product to do exactly what it needed to do. And that got us ahead a pretty good factor. And we were able to hold our own against MIPS, which was, at the time, a whole company that was founded on building a RISC processor. So the machine was less important to MIPS than the processor. We knew what we were doing, which was to deliver the machine.

Remacle: It seems like you would have an advantage that you had a company full of systems people who were able to work with the chip people to make sure that everything was optimized properly? Is that a...

Rosing: I think so. I think that—we can talk later about a lot of the optimization—but I think the key thing was that we had some computer architects—and Robert Garner, who's in the audience here, was one of them. We had Steve Muchnick, who was a compiler person. And we had a whole bunch of really excellent operating system folks. And so the team designed this thing knowing what they were trying to build. And they didn't lose that focus and over-concentrate on one technology at the expense of the overall vision. And that's what, I think, got us there. And then...

Remacle: You all indicated that there was a fair amount of internal discussion regarding whether or not you should have your own microprocessor. What about the decision for RISC? Steve indicated you really didn't have a good alternative. But was there much internal discussion, or was that pretty much a straight line between "microprocessor and it'll be RISC."

Lacroute: That part was fairly straightforward. I mean, Steve was right. There was no way we could have built a CISC machine. We had built the equivalent of a CISC machine at DEC, and I had a pretty good idea of what it took, and they were not chips, they were big boards, and so forth. And translating it into LSI was not a piece of cake. No, it could have not been done. But I want to make just a couple of comments about the business proposition. I don't think there was any way that you could justify the

decision just based on return on investment. If you did a classic business model, it would never have worked. It was intuition, gut feeling. "This is what we've got to do. And never mind it may fail, but that's how we can really build a big company." And the power of intuition, and saying, "This is where we need to go, never mind what the numbers tell you today. That's what you have to do, if you want to be different." I mean, Sun in those days was playing a game to differentiate itself. I'm going to give you a little idea about what we were doing. We were sorting the processor chips from Motorola, so that we could get the fast parts. And they were marked with a little red dot. And Apollo could never figure out why our machines were faster, because it was one less wait state. Well, we sorted those chips! That's one way to go and differentiate yourself, but long-term, you cannot play that game. So, you know, just doing a rigorous business case analysis, the project would never have flown.

Diamond: I think it's useful, just to follow up on Bernie's comment; RISC really put the CISC microprocessor companies back on their heels. They started having to justify their processor strategy, because people would point to Sun and they'd say, "Well, why are you selling these expensive complicated processors, when Sun is able to develop systems around SPARC and they're less expensive and faster?" And so in that early period of time when SPARC first came out and the systems were out, the traditional microprocessor companies were knocked back on their heels, and they started talking about, "Oh, well, we have a RISC strategy." For example, Intel announced a RISC processor that was really a graphics chip that they had previously designed, but Sun was such a powerful force in the high-tech world that Intel had to respond.

Remacle: Let's go back and talk about Motorola a little bit. What was Motorola's response when you guys started talking about doing your own processor?

Lacroute: The same reaction as the board, "You're crazy! You cannot design it. You cannot build it. You cannot manufacture it. Get lost. Stay with us." That’s simple. Nobody believed it!

Rosing: That's right. I mean, in many ways that attitude, existed on the part of a number of the companies we talked to about licensing, long after we had been successful. It was just very much, "We're semiconductor companies, we know how to do computer architecture." And we were a systems company, and we knew how to do computer architecture, and those two groups couldn't always get together and see things the same way. And we, of course, having the software, the compilers, the applications, the market, really did have that sort of—I won't say it's superior insight—but we had a keener insight on what we needed to do. As I recall, the Motorola thing was very much, sort of in the beginning, chauvinism of a big semiconductor company about this little operation in Mountain View. But later on, it just became very, very, very simple, proprietary versus open. No one really liked our open strategy. They all wanted Sun's business, but they wanted a unique lock on Sun's business.

Remacle: In Motorola's case, how much of it was a business pushback, i.e. what part of their revenue did Sun represent?

Lacroute: Well, in terms of the 68K family, the state-of-the-art microprocessor at the time, a very significant amount. I don't know the exact number, but it was non-trivial, because at some point in the '85/'86 timeframe, we had overtaken Apollo in unit shipments. So they were also depending on Sun showcasing the power of the chip.

Rosing: Sun was the flagship vendor.

Lacroute: Absolutely.

Rosing: For the 68000 product line.

Remacle: So Motorola had a lot on the line.

Lacroute: No question.

Diamond: They did, and Sun was a key part of their marketing strategy for the 68000. When they would try to design the 68K into a new customer, they would point to a Sun workstation and say, "See what can be done with this processor?" Their engineers used Sun workstations, their architects solicited advice from Sun system architects, from compiler writers, from software developers. It was a tremendous loss to Motorola when Sun went over to SPARC.

Remacle: Was there ever any discussion with Motorola once it was clear you were going to go ahead with the strategy that they partner with you in some way, shape or form to get the chip developed?

Lacroute: We tried very hard to convince them to go with us as a partner. Really a lot of discussion, negotiation, and some of the engineering vice presidents at Motorola were actually, like Tom Gunther, were kind of supportive of doing something with Sun and SPARC. But the upper management would never let that happen. And again, I think they were—you know, I'm guessing, because obviously, you can never put yourself entirely in the other person's shoes—but I'm guessing that there were really two factors then. One, they still didn't believe that it could be done. I remember a conversation with Galvin, who was the CEO of Motorola at the time, and he was respectful, but said, "You know, <laughs> it's not going to happen." So they didn't believe that Sun would succeed, and at the same time, the fear that they would no longer control the architecture. And that was a showstopper for them. They just couldn't see themselves not controlling the architecture.

Rosing: Right. These were the days where companies used the term "proprietary" almost as if it was a strong, positive marketing claim. And so that was one of our judicial moves, if you will, was to say "open."

Remacle: And would you say that Sun gets a lot of credit for Sun/UNIX gets the lion's share of the credit for flipping that proprietary advantage, if you will, on its head?

Rosing: I'm not sure if we flipped it. We certainly threw a great big stink bomb in it. I don't think most companies think that's a good marketing word anymore.

Remacle: I would agree.

Rosing: I mean, there's this...

Lacroute: But I think at the time, Sun really was instrumental in making that happen. And it adds a lot of value from the customer viewpoint. I mean, the portability of the application at the end of the day is what it's all about. Forget about the hardware, the software, whatever. Can I move my application software? Can I go from one vendor to another with the same application software, if one fails, or whatever.

Remacle: Sun benefited from being at the right place on that curve in the marketplace, because that was so obvious in the early 1980/'81 timeframe. It really was around '83/'84 people started to wake up and say, "Hey, you cannot change your architecture every time you come up with a new chip, because we've got this huge software investment." So at some level you kind of benefit? I'm putting a question mark at the end of that statement.

Diamond: Well, it also started earlier with busses. In the early days of Sun, they didn't have the engineering resources to design all the cards to make a workstation. The disk controller card, the network card, the memory card and so forth. And this was the time when standard busses were just coming on the horizon that were capable of the kind of performance that Sun wanted, like multi-bus and VME. And so again, like with SPARC, I think the brilliance was to make necessity into an advantage. And Sun was known as the open company. They used open busses, and they used an operating system that was UNIX, that it was available, and had been promulgated in universities. This was not the model back then. This was radical.

Remacle: Go ahead, Wayne.

Rosing: Something else we did at this time—I'm going to say I think it was '86—a couple of reporters from, I believe it was *Electronic Magazine*, were in my office and we were talking, and I was musing about basically how do we deal with benchmarks, which are just "lies, damns lies, and even more lies." And some of us cooked up—and I know Robert was very involved in it, and other people were very involved in it—cooked up the idea, "Why don't we have honest benchmarking?" And out of that came the whole SPEC consortium. And SPEC benchmarking. And we actually sort of did that in anticipation of what would

be a fight with MIPS. Anyone can claim anything, and you can probably write one loop that will substantiate your claim. And I thought, this was sort of a side thing that came out of what we all did, which was much more honest benchmarking. And if I recall, when the first Sunrise chips were sort of going through that, we were looking at eight. We were not looking at ten.

Lacroute: Nope!

Rosing: But we went to Fujitsu and said, "Smaller, smaller! Shrink the die!" etcetera, etcetera. And we tweaked everything, and we tweaked everything. And we ended up pushing ourselves from, you know, eight to twelve, fourteen, I don't remember the exact number. But the megahertz went up from eight megahertz and ten megahertz chips, to sixteen megahertz chips. I'll remind the audience I'm talking about *mega*hertz, not gigahertz. And this was really fast for computers at that time. This was really something special.

Diamond: Well, it was ten times a VAX.

Lacroute: Very, very fast.

Rosing: The goal was to be ten times the VAX 780. And we got there.

Remacle: Loop back a little bit on the issue of software. What was the response of the ISVs that Sun had been pursuing assiduously for, what, at this point, three or four years, to this decision to change the microprocessor architecture, the heart of your workstations?

Diamond: Well, there were thousands of applications on the existing Sun-3 68K-based platform. And I think it's fair to say there was a degree of skepticism that Bernie and Wayne are talking about around the ability of Sun to develop and promulgate its own processor architecture. I think that was shared by some of the ISVs as well. Because of what we positioned as ease of recompiling for the new architecture, we had a story for the ISVs and we had a program to help the ISVs move over, but I think they were also skeptical that this transition could be made. They were making a lot of money on the existing platform. We were presenting a new platform, which had performance, cost performance advantages to them. But it did require, I think, a jump in their thinking to understand the ROI there. Now, when we launched SPARC, we had a tremendous ISV support. There was a lot of work between the time that the decision was made and the launch. There were lots of ISVs...

<overlapping conversation>

Remacle: Excuse me. How long was the elapsed time between the decision that, "Okay, we are going to go ahead with our own processor, and it is going to be RISC," and the day that the first SPARC station was announced?

Lacroute: Two-and-a-half-years.

Rosing: Yeah, that's right.

Lacroute: A year too long.

Rosing: Yeah, a year too long. <*laughter*> I remember that feeling! I remember I was late the day I showed up for work. <*laughter*> And it was 7:00 in the morning.

Remacle: So you had a two-and-a-half year period, roughly, to bring the ISVs along, make sure they were clear on exactly what they had to do, what you were doing to support them, etcetera?

Lacroute: Yes, there was a program to really educate them, let them know what was coming. But there was also another piece, which was—the fear is that we would lose those ISVs to another platform like Apollo or whatever. "But we told the ISV's, "don't worry. If SPARC fails, we have a 68030 in development as well." And we did!

Rosing: We did.

Lacroute: Our approach was to say don't go away from Sun. We believe that the RISC machine will give you a tremendous advantage, but if that were to fail—and we don't think it will—here is the back-up strategy, the 68030." So that really was very helpful in keeping those ISVs in the Sun-4, which was the most important thing to do.

Rosing: Yeah, that's right, the Sun-4 was introduced as the Sun-4 260 and Sun-4 280. It wasn't introduced as SPARC. I mean, they actually downplayed—I wouldn't say we up-played, nor did we downplay—but SPARC wasn't the center of that introduction.

Remacle: It wasn't the center-point, yeah.

Rosing: It was launched, but if you were a Sun customer, you saw, "Oh, there's a workstation and there's a server. They're just like the other things, but they're a lot faster and priced in an interesting way,"

which Bernie can talk about. But we didn't completely lose the direction. And we had a 68000 strategy, which we spent real live money on. And does anyone remember what the MIPS turned out to be? Wasn't it about four? Wasn't that right?

Lacroute: Four, four-and-a-half.

Rosing: You know, the hype was going to be it would compete and it would be wonderful and it would be cheaper and it would be—it was more expensive to build, and it was slower. And that's just the reality. But it was there, and I think we may have even done a second generation of 68000 after that. And then it ended.

Remacle: So the 68030 plan or strategy was not solely a safety net for the SPARC program, but it was part of the overall workstation strategy.

Rosing: Keep it going, yeah.

Diamond: We announced, I think it was courageous to say this, but we announced at the time SPARC was announced that we would keep building 68000 machines as long as people wanted to buy them from us. So the ISVs that didn't port over would have a home if people wanted to continue buying the machines.

Remacle: This is a probably a good point to ask what was the salesforce's response to this strategy?

Diamond: Well, as you've heard, there are several parts to the strategy. There's the new architecture, new operating system, new compilers, new price performance, SPARC-4, SPARC/Sun-4 strategy. And then there's the open system strategy. And I think the salesforce had a very negative reaction to the idea that we would be putting people into—potentially into competition with them on the SPARC architecture. So I think those two things, you need to separate those two concepts.

Lacroute: You should tell your story about Joe. Joe Roebuck was the VP of Sales.

Rosing: Well, I don't know which Joe Roebuck story to tell! *<laughter>*

Lacroute: Oh, c'mon!

Remacle: Well, you can tell them all!

Lacroute: No, we couldn't tell them all! <*laughter*>

Rosing: No, we couldn't tell them all, not without a release from Joe.

Diamond: I think we'd need a button.

Rosing: I think the essence of the story was—Steve has said it best, you know, "By doing this, you're taking away from my grandchildren's education." You know, the sales force just hated the idea that somebody else could be essentially making a Sun workstation, and selling it against them, and they would lose their commission. The sales force at Sun was commission-based. They're not there for fun, they were there for earning their commissions. And they were very aggressive, and so that was always a very complex tension. I don't think it was ever fully resolved. As SPARC moved off into embedded applications and other things there was no pain and suffering around that.

Remacle: As long as it was going into other workstations.

Rosing: Where the problem was it could have been easy to take SPARC chips, go find yourself a way to get Sun OS, build a Sun clone and just undercut us on simple margin. And so the risk was that Sun would allow its core business to be commoditized. Sun sales force, I think, oversimplified that.

Lacroute: Oh, yes.

Rosing: But the commoditization didn't really happen. There were a number of companies who built early "workstations," but they didn't have the Sun OS. And so that just—you couldn't get here quick enough. And that commoditization was not—it didn't make it better for the ISV community. They didn't really benefit from that either. So it was just sort of we were in a different market where commoditization in the sense of the PC, there wasn't a business case for it.

Diamond: Well, I don't know if you remember this, but one of the potential licensees of SPARC was very skeptical that their engineers were as good as Sun's engineers. And so they thought if they built a SPARC-based system, they would never be able to compete with Sun. They weren't sure that they could even get a design to work. And I don't know if you remember that after we had gone on from the 4/260 and we had newer products in the marketplace, we actually gave one of the licensees the 4/260 design.

Rosing: Mm hm.

Lacroute: Yeah, that's right.

Diamond: And said, "Here, use this as an application node for building a SPARC-based system." Because they weren't sure that they could do it themselves. Sun's engineers being as talented as they were.

Remacle: Was it just talent kind of at an individual engineering level? Or was it also this knowledge that Wayne referred to earlier, the OS, plus the compiler, plus the hardware, plus the chip, plus—you know, that you had all under one roof? Kind of the Apple strategy, if you will.

Diamond: I think that's part of it, but I think the engineering talent, the culture, the engineering management, the world view of the company, the energy of the company. All those things together made an unstoppable force. And I think many in those years acknowledged that Sun Engineering was just a cut above anybody else.

Lacroute: It was this kind of—call it kind of a fad in the mid- to late '80s, or even '90s, open system, open strategy. So you need to examine it from different points of view. One, it's a brilliant marketing idea. Really brilliant. The sales pitch was, "Why, Mr. Customer would you buy a system which is not openand locks you into a specific vendor? I mean, it's brilliant!"

Rosing: "You don't want to be locked in!"

Lacroute: "You got to be crazy not to do that!" So that's one piece of it. The reality is a little bit different. Good example, we gave the design of the Sun-4 260 to a potential competitor. Well, it was last year's technology. "Go and copy that stuff, and you are going to be chewed like you've never been before! Because the guys inside Sun are two steps ahead of you!" So the very, very tight integration of the hardware, of the software, the chips was what made the Sun systems so high performing. And believe me, those engineers were competitive. There is no doubt whether it be hardware or software, they wanted to win! And win is to build the fastest, best system that you can possibly do. Forget about the open system. Forget about licensing this stuff to somebody else. And that's what engineers should do; they should do the very best they can! So there was kind of, you know, an interesting marketing ploy, but inside, man, there's no way anybody was going to beat those guys with a faster system. That was just not going to happen.

Remacle: This is probably a good place to loop back just slightly. Can you talk—Wayne, you used the term "Sunrise." Can you talk about Sunrise, how it worked into the SPARC, not the naming necessarily, but just kind of how it evolved. Who was the team? Who were the key players, both in the engineering, marketing, manufac—you know, the top level folks on the SPARC team?

Rosing: Oh, wow, now you're doing a memory test.

Remacle: Well, I didn't mean to.

Rosing: You know, Robert Garner, Anant, Don Jackson, Masood Namjoo, K.G. Tan.

Diamond: Ed Kelley.

Rosing: Ed Kelley, right, there you go. That was sort of the core bunch when I joined. And Bill Joy. Dave Patterson, of course. Then you had Steve Muchnick, Dave Weaver and Doc Williams and Steve Kleiman, and Mitch Bradley.

Lacroute: Who was the other UNIX guy that came from Beck? I can't remember his name now. I can see his face. And he was very instrumental in tuning the system.

Rosing: But the interesting thing is this was a company, an industry-changing gang of people. And it's remarkably finite. It was maybe 12 or 15 if you add everything in.

Diamond: Roughly in order of magnitude, less than you'd expect, at anywhere else.

Rosing: Right. And then the corresponding 68000 team under Howard Lee was on the order of three or four times the size of that. And so this was very focused. And I'm absolutely convinced to this day that if you want to do something like this, you want to only do it with the very smallest of teams.

Lacroute: Small team, very smart.

Rosing: Small team.

Lacroute: Driven.

Rosing: And a clear goal that doesn't get changed. With one exception: more megahertz, more megahertz, make it go faster, make it go faster.

Lacroute: And sooner and sooner and sooner.

Rosing: Yeah, that was his job.

Remacle: What was the Sun culture? You all implied there was a high performance Sun culture. Was it expressed as such, or was it just kind in the ether?

Rosing: It was in the air. It was urgent, focused and urgent, were the two things I take away from that period of time about Sun's culture. We knew what we needed to do, and we needed it now.

Lacroute: Well, there's another factor that is really important, and I think it applies, not just to Sun, but successful technology companies: engineering driven. And did I get beaten/criticized a lot ! In a startup/small company (and I also believe large ones) the focus had to be engineering. Get engineers. Get the best engineers that you can, and fund them like there is no tomorrow so that they build the very best products. And so the culture at Sun was very much engineering driven. And I think it's good. I make no apology for that, none whatsoever.

Remacle: You said you were beaten up for that. Who beat you up?

Lacroute: Well, I mean, Sun marketing people, "I don't get enough money. It all goes to engineering." Some, well, I mean, Carol, who happened to work for me, I mean, she was complaining all the time. And sales.

Remacle: This was Carol Bartz you're talking about.

Lacroute: Carol Bartz, right. And even the board said, "You guys are spending too much money on engineering. Not enough on this, that or the other."

Diamond: Well, I think in the days where we're talking about, and I was there from '87 to '93, there were two factors that really were unique. One is Sun engineers designed the next generation of workstation that they would use to design the next generation of workstation. So not only did they have the talent to do it, they knew what needed to be done. They knew exactly what could be done, and what needed to be done. And so there was no need for marketing to figure out what product to develop. And if you really had any questions, you could always go talk to Bernie, who would tell you what to do. So...

Lacroute: Faster, sooner. *<laughter>*

Diamond: So there was a laser focus on the next product and the next one and the next one. And when it was necessary to make this major change in SPARC, they had the engineering ability to do it. Which many other companies may not have been able to do. The other thing, we talked about this morning that, I think was unique about Sun in those days is when I was hiring engineers before I cam to Sun, the first

question anybody would ask was, "Do I get a Sun workstation?" That came before the salary, and perks and anything else—Sun workstations were the currency that engineers used to decide where to go. If you couldn't provide an engineer with a Sun workstation, then you weren't nearly as interesting as an employer.

Remacle: How did Sun—and maybe Bernie, you're the right person to answer this—how did Sun ensure that the workstation and the SPARC team were yoked together in a way that they didn't get out of sync? 'Cause it seemed like that could happen.

Lacroute: I'm not sure I understand the question. Help me.

Remacle: Well, you've got a system design going on over here of a new workstation, and then you're working on, what turned out to be gate arrays, but a chip, or chips, microprocessor. And there's got to be a feedback loop between the two of them to make sure they don't get out of sync. How did that get managed?

Lacroute: For the first product, they were the same people, so they could not get out of sync. It was a single focus, tighly integrated team. The only thing that they shared from the rest of the company was the box. You know, the sheet metal.

Remacle: The chassis.

Rosing: Yeah, they were designed as chips, excuse me, systems, not as chips. The Viking chip was sort of a different process, that was the first million-plus device, what you would call, integrated circuit standard type design. That was a longer process. It was more a chip and then the system showed up. That, I think, was a more ragged process if my memory is correct. But the early days, this stuff was highly optimized, and I think it's important to understand the decisions some of the designers made about the cache chips, and what were things that were not associated with the processor, but were associated at that time with the system, were just as important. Every last aspect of things at this particular time was optimized. I mean, it was optimized with what CAD we had, and it was optimized by sheer courage the rest of the time, because every nanosecond counted in terms of performance.

Lacroute: And I don't think you can do that if you separate the three pieces, chips, system, and software. They have to be one, and everyone must be thinking of the end result. And the end result is how fast can I make the system for the customer's applications?

Diamond: If you make the team bigger, then you have to split it apart into pieces, then you need interfaces between the pieces, and then the system design gets broken apart from the processor design

gets broken apart from the OS and the compilers. And I think it makes it slow down or become impossible.

Remacle: At a tactical/process level, did you have weekly status meetings? Monthly? Daily? Or did just people come to work and do their job?

Rosing: I think we sort of had all of the above. I mean, my style is go kick tires all the time. Check in with people, make sure they're happy and that they've got—they have what they need to do their job. And I would say, "And faster."

Lacroute: And management by walking around is really important. Go into the lab, talk to the engineers. Forget about the hierarchy for a minute. And find out what's going on. I mean, you get the sense for what's going on. So there were regular meetings, but you know, in crisis there were the pizza meetings every evening. I mean, well, 7:30 pizza meeting. "All right, until we solve the problem, we have a pizza meeting every day." That was the good old days when we did those things.

Rosing: Yeah, no, those are simple to do. You know, if you need to much management process in a project, you're probably already in trouble. At least when you want to do this kind of innovative stuff. I mean, this held true, at least in my career everywhere I've been.

Remacle: In addition to RISC and CMOS, what other evolving technologies did Sun employ in this process?

Rosing: Well, let's see, at the early times, the Sunrise times, particularly. And I think it may have been true...

Remacle: Wayne, let me interrupt you. Can you explain what Sunrise is?

Rosing: Sunrise is the first product that came out that were designed from '84, and came out in '87. Those times were relatively, from a technology point-of-view, the other thing we were very sensitive to was cache chips. Just really, really good intermediate memories, between the DRAM memory and the processor itself. And then later on when we got to the SPARC Station-1 generation, and SPARC Station-2 generation of products, which was more like '89/'90. What we were able to add to the equation was a graphic processor. There was just sort of this out-of-the-blue couple guys I found, who we hired, who just sort of went off and did this amazing graphics processor. So we were able to add one more thing to the workstation above and beyond the basic processor, and that really improved our graphics quite a bit. Then I would say, the next set of technologies that really enabled things, it would be moving into hard silicon and other major silicon designs. And from our point of view, the technology that happened was the

sophisticated CAD tools that then became available from ISVs on, guess what? SPARC stations. And now we had CAD tools that we could use that were as capable as the proprietary tools that were across the street at Intel. And we started with Daisy tools, which were relatively straightforward tools.

Lacroute: On an Intel platform! What a shame!

Rosing: On an Intel platform, that's right.

Remacle: Well, Phil Kaufmann left Intel and ran Daisy for a while.

Rosing: Yeah, but relatively straightforward tools. But then over the time Verilog came along, and we were a very early customer of Verilog's. And that allowed us to do basic system and logical simulation at a scale that made sense. And we could run the simulations on computers we owned that we had paid for, and they would actually come to a conclusion when you got to work the next morning, and you would know what you needed to know. And so we were constantly bringing in CAD tools from our own base, and using them.

Lacroute: Sun was one of the very first customers of Synopsys.

Rosing: Right.

Lacroute: And I mean, you know, it was a startup company, and they were very good in helping Sun build the next generation of chips.

Diamond: You know, Bernie mentioned sorting CPUs. Sun was not afraid to go to marginal limits in those designs.

Lacroute: Ah, yes. Sometimes even beyond the marginal limit.

Diamond: Well, I remember, I think I was talking with Robert, or maybe it was Ed Kelley, about the 4/260 design, and they were walking me through things, and there was some setup and hold-time requirement that wasn't being met, and I think Robert said, "Well, we'll just use a gallium arsenide PAL in there." And I said, "Well, what if that doesn't work? You know, there's nothing after that." He said, "Don't worry about it, we'll sort the chips." And so Sun had the ability...

Remacle: When you say, "sort," you mean physically hand sort?

Lacroute: Test.

Diamond: Yeah, set the tester up to sort out the chips that would go beyond the spec. And Sun was able to do that successfully, because of its CAD ability and its test sophistication, and manufacturing ability. They could do that.

Remacle: What was the response to—you hinted that Intel, in particular was forced by Sun to announce a RISC chip—what was the response of some of the others at that time, microprocessor companies? And MIPS,….Wayne you had alluded to the other RISC micro—I don't know why I did this—but the other RISC microprocessor.

Diamond: Well, first it was denial. You know, the first approach was, "This will never work. These people don't know anything. RISC is just an academic exercise. We know better." After MIPS and SPARC, it became impossible to deny that RISC was the right moment in time, and it was going to be successful. So then we had sort of the moment when Intel basically admitted that RISC was real, and they announced what had been built internally as a graphics chip as their answer to SPARC. Well, it really wasn't an answer to SPARC. Really, not nearly as capable. But for the customer base, it was an endorsement of RISC as a true option. And that really changed the game. And then Motorola had their own internal RISC product design, the 88K, which Wayne can talk about in great detail.

Lacroute: Oh, I mean, I think that's absolutely true. But other system companies played that game. I mean, IBM resurrected their RISC project, and I think it became eventually the PowerStation or something like that.

Rosing: Power PC.

Diamond: Yeah, that's right.

Lacroute: Actually, we had gotten some people from IBM original RISC project. DEC designed the Alpha RISC architecture, which is probably one of the very best RISC architecture of all times. Business-wise, they screwed it up. Big deal. But technology it was well done.

Rosing: Very well done.

Remacle: All right, to finish the discussion on what was the response on the larger market to Sun's decision to have, not only its own processor, but a RISC processor. Both of the two of you were talking

about the Intel response. What about, Wayne, you mentioned you had a conversation with Bob Noise on the topic.

Rosing: Right, sort of a little later in the process, when we were actually engaging in the, what was known internally as the Viking design, I was talking with Bob, we were scuba diving in Egypt, I believe, and we were talking off the record, but nothing too improper about what was going on, and I told Bob what we were trying to do, basically, and he said something interesting. He looked very please, and he wasn't in any way horrified, and said that he thought that there was a window. That we were in a window where it was possible for the systems companies to kind of take back their current future from being dependent on, if you will, semiconductor companies for computer designs. But we knew we were going to be dependent on semiconductor companies for process technology. No systems company, with the exception, really of, quite frankly, who? IBM, who had the critical mass to develop the semiconductor side, and the computer side. But this was—Bob was personally favorable and very encouraging. Because I was wondering sometimes were we biting off more than we could chew. But it all sorted well.

Remacle: What about the rest of the Intel executive team?

Rosing: I had relatively little contact. Some of the other folks here had some more interactions with them.

Lacroute: There are really two time frames to look at. The first time frame is before the first machine was built. Could we find somebody to build that chip, whatever that chip was going to be in the early days? So we asked Motorola, we asked Intel, we asked TI, we asked, I don't know who else, and they all said, "Well, forget it. Not going to happen." The only company that reacted favorably was Fujitsu. And so the decision to go with Fujitsu was one—yes, they had the right technology—but they also they're the only ones that were willing to go and build the thing. Later, many of the semiconductor guys came back and said, "Well, if you want us to build that the chip, that's so many million dollars." Well, Sun had no money to fund a chip development. You had to find a partner that'd be willing to partner and provide the technology. So in the early days, the Fujitsu decision was made by default, because there was nobody else.

Remacle: What was Fujitsu getting out of it?

Lacroute: That's a good question. One, I think they were intrigued by the possibility of a new microprocessor architecture. The idea that they could build chip, they could sell chips, that they could sell not just to Sun, but to somebody else was kind of intriguing to them. They really wanted to try. And they also wanted to showcase their technology. I mean, they felt that if this would work, their gate array technology would be shown to the world. Here is a company that could build something that is really state-of-the-art. And so I think it was big motivation there.

Rosing: I think we'd talk to the traditional semiconductor companies that had traditionally designed microprocessors, they would map our wanting a processor on their process. Whereas, Fujitsu had a gate array product line, and they were building very large gate arrays. And this, of course was before FPGAs. They were being invented at this time, and they were much smaller. So Fujitsu had a standard product answer, with the exception that they then designed a quarter of their array, a custom design, essentially register structure, internal RAM structure, surrounded by lots and lots of gates that you could basically program with a couple levels of metal. So the design of these first chips was very much a printed circuit board thing of laying aluminum down between things, and very hard to do, and done at a superb level. But Fujitsu had the advantage that they actually had, fundamentally, a standard product we could build on. And that succeeded.

Remacle: Who on the SPARC team was responsible for finding manufacturing partners? People who would fab your chips for you?

Rosing: In the beginning, as Bernie said, there was a choice of one. <*laughter*> And then when we decided to go open, Steve and I then started "the crusade" to go out to find more vendors, to see if Fujitsu would go forward with it, and others. And we talked to quite a few companies as we went through this process.

Lacroute: All of them, I think, you did.

Rosing: If anybody had something, we'd sort of talk to them.

Remacle: What pushback did you get? Was it consistent, or did everybody have their own different reason for saying...

Rosing: It was variable. The larger companies were still—they had their own architectures. And at this time, I would say, most of the larger semiconductor companies didn't see themselves as "fabs for hire."

Lacroute: That's true.

Remacle: And we're talking what timeframe now again?

Rosing: '86/'87/'88/'89. So you would talk to National; National wasn't a contract fab house. So we would be talking to the microprocessor people. So if you went to a big semiconductor company, you had to get through people who had a vested interest in their own product line.

Diamond: I remember after Bernie and Wayne kind of opened the door to another set of Intel discussions after '87, we had a deep technical discussion with Intel about them possibly building SPARC as part of the Intel product family, and as Wayne said, but we were talking to the Intel microprocessor guys.

Remacle: Which people were you talking?

Diamond: I think Jan Prach [sp?] was leading up the team from Intel.

Rosing: Dave House.

Diamond: And the answer came back, "You know, we really like your architecture, if you could just make a few minor modifications, we'd really like it." And the minor modifications really were to turn it into an X86.

Remacle: Not so minor.

Diamond: So we got that kind of response from, you know, the traditional semiconductor companies, because it wasn't like the TSMC model. We talked to the microprocessor guys at these companies, who already had an architecture, and already knew what they wanted to do.

Rosing: The breakthrough, in terms of the big licensing arrangement, turned out reengaging with TI, and we were very persistent with TI. And oh, I don't know, I cannot imagine telling you how many meetings it took. But TI, ultimately, at that point, did not have a commodity architecture. They had their DSP chips. And so they had a really good technology base and product base in their DSP chips. They had a production line in Japan that was tuned to doing memory, RAM. But they had sort of tuned it so that it was also able to easily be adapted to doing really high performance other forms of chips. They were using that for the DSPs, and then we came in and we provide a coherent plan for, "We can be a customer for you guys, and we'll be open, and you can even sell the chip if you want. If you don't want, you don't have to." But at that point, we were a big enough customer that they were interested in us primarily because we'd help drive their process line. And so we were able to finally break the code on what we really wanted, which was access to a manufacturer that would sell us chicken—silicon, excuse me, not chicken—*<laughter>*.

Remacle: Those are two different things.

Rosing: Silicon on a cost base. They weren't going to price the thing based on their perception of the value of that chip in our food chain. We were coming in, and we were paying a fair price to buy a fabbed product. And we would accept the yield, and all that stuff, in our cost structure. That was a really significant difference.

Lacroute: TI was also interested, because in the back of their mind, they didn't have a microprocessor. And they were seeing Motorola and Intel being successful with their microprocessors and —in the back of their mind they were thinking, "Well, you know, what have we got to lose to go with those guys? Maybe if things go well, we can have our own microprocessor line, and compete more effectively with Intel and Motorola." The situation was very different with Intel and Motorola. They already had their microprocessor architecture, and it was very, very hard for them to change. Especially because they wanted to control the architecture, which Sun would never let go, and should not. It would have been a mistake for Sun to let Motorola or Intel or whomever, control the architecture. But it was much harder with those guys who already had their own microprocessors to change. On the other end, part of the Intel discussion, at one point it got bumped all the way to Andy Grove, and Andy was intrigued. I mean, as we all know, Andy can be a risk-taker, and can experiment and try things...

Remacle: No pun intended.

Lacroute: Yeah, with a "k," not a "c." He's a guy who would be willing to try something different. But the discussion really ended up in a very bad place very quickly, and then Andy was the enemy. You can tell your stories. You sold a license, and I'll let Steve tell the story—and then came Intel behind you. When you make an enemy of Andy, it's not fun.

Diamond: Well, we were going around the world when—after we announced SPARC in July '87, talking with Tier 1 and Tier 2 systems vendors, who were, at that point, 386/486 customers. And whenever we would make a call, Wayne would visit somebody, I would visit somebody, then we would hear that Intel would arrange to visit them after we had, and the story—and of course, I have no way to know exactly what was said to these people—but there were rumors in the press that if people would move to SPARC, they would suddenly be walled off from early information on X86 processors.

Lacroute: Oh, there was also a shortage of Intel parts. Gees!

Diamond: There was a shortage of parts, and there's also a shortage of I think it was the "Red Book," you know, the information about the next version of the processor. They couldn't make enough of those books. And so if you were a systems vendor, and you expected to sell systems based on X86 chips, it was very, very dangerous for you to lose access to that information. It probably meant that you were going to be delayed months, if not years- in —developing a product based on the next X86. So that turned a choice to use SPARC into a risk—not to be punny—for X86 vendors. And we would hear this.

You know, we would visit Company X, and a couple weeks later, they'd call us back and say, "You know, we can't do this. It's too dangerous for our PC business."

Remacle: Bernie, you've said that Fujitsu was basically a decision made by the fact that they were the only game in town that was willing. They were the only ones who wanted to play your game. What kind of criteria did you have in your head—I don't mean necessarily your personal, but the Sun head—about if you could have picked a manufacturing partner, what criteria would you have had at the top of the list, and how did that match with what you ultimately got from Fujitsu?

Lacroute: Well, I think the criteria were again, very much driven by what we wanted to get out of the system. First of all, can you get the performance out of the silicon? This was of critical importance. Two, is the vendor willing to work with you to tune the design, the process and so forth, and really build the system as fast as possible? You alluded before to the chip, that the chip was the system. All the componenets, chip, system, OS had to be tightly integrated. So a vendor that was willing to work with us, to optimize the entire system and make it fast was critical. The triple ported memory on the gate array is a good example of that, as it was not a standard feature of the gate array. But they were willing to do that. And they were willing to accommodate us. Because I think they really, fundamentally, wanted to showcase their technology. You know, going to Japan was not without risk. In fact, for a period of time, Sun was accused of selling out to Japan.

Remacle: This was just on the tail end of the—or in the midst of the Japan chip war.

Lacroute: Yeah. And you know, there was the book, "Japan That Cannot Say No," and a bunch of other factors. And some of the people who had turned Sun down, not wanting to build a chip, would say, "Well, you're kind of traitors. I mean, you are selling out American technology to Japan. "If you put that in context, the military was absolutely livid because the weapons systems depended upon DRAM from Japan, and displays from Japan. "And here you go and you give them microprocessor technology?" So that was kind of an interesting pushback on Sun. That didn't last very long, but it went on for maybe half-a-year or something.

Rosing: Yeah, it was a consideration for a year or so.

Diamond: By the way, there was a time when SPARC was classified as a munition, because of the processor power in SPARC. Someone called me from one of the three-letter agencies before we had announced SPARC, came to my office and said, "I know all about SPARC. I know what you're doing, and I want to buy one of these things." And I said, "Well, we haven't announced the product, I can't really talk with you about it. We haven't set pricing and so forth." And he said, "No, I know the pricing. I know when it's going to be announced. And I want to buy one." And I said, "Well, can you give me a purchase order?"

He said, "I can't even give you my name." <*laughter*> "But I want to have the first one that can be sold when it's available."

Lacroute: And it turned out to be a very large customer for Sun for the early SPARC station. Very, very large customer. Even had to go and get a security clearance.

Remacle: Let's stay on the manufacturing track, although I keep being tempted—I have a couple other questions I want to go back to. What about some of the other manufacturing partners, or cross-licensing folks that you talked to? In particular, Cypress. You had some conversations with MIPS. Can you talk about some of those other folks that you talked to over the years?

Rosing: Well, I don't recall talking with MIPS. Perhaps someone else did. But we talked with Cypress, and Cypress ended up doing a number of custom design SPARC chips. We were very involved in that. They took a lot of the lead. And then LSI Logic also did chips. Those are two that are noteworthy. Along, of course, with Fujitsu. And then later, TI. Those relationships were really interesting. It was sort of in the beginning, these were hot semiconductor companies that didn't have a processor product line that thought they might want a processor product line. But more than wanting a processor product line, they wanted Sun's business. Very, very tactical. Particularly T.J. Rodgers at Cypress. The name of the game for T.J. was "purchase orders from Sun." And they got them, within reason.

Lacroute: For processors and SRAMS.

Rosing: And SRAMS, lots of SRAMS. Didn't get as much as I think they hoped, but they ended up, as I remember, they ended up doing value pricing. In other words, they were going to charge us what they thought it was really worth to us, as opposed to charging a fair markup over manufacturing cost. So that tension was always there. But there was a lot of business between Sun and Fujitsu and LSI and Cypress for a while. And then as it slowly moved, the locus of the business relationship for volume chips moved in time to TI. But that wasn't really until the '90s, I think. I think the business in the late '80s was all Silicon Valley and Fujitsu. They were our prime suppliers.

Remacle: Can you talk a little bit about the relationship between Sun and Dr. Yasufuku, and Sun and T.J. Rodgers?

Lacroute: I think the relationship between Sun and Yasufuku was actually, for a very long time, very good. I had met Yasufuku at DEC, and he kind of remembered me when we went to Japan. He was a good example of a non-traditional decision-making process. Yasufuku made the decision in three weeks to work with Sun. And we say the Japanese take forever. No! It took very little time. So I thought for a very long time, our relationship was good. It got a lot more complicated when Fujitsu became part of the TRON effort. TRON was a microprocessor, which was launched by the government of Japan, basically,

enunciated, "How are we going to build our own microprocessor?" And the Japanese government basically coerced every single company in Japan to support TRON. And Fujitsu had to support TRON. So that made negotiation a little bit more complicated. The other thing that kind of slowed down the relationship for a while is that Fujitsu wanted to invest in Sun. And it was not a great thing to do. But they ended up investing a few percentage points, which was kind of a token investment to pacify them. But they had seen the AT&T investment in Sun, and said, "Why not us?" Now that too made it a little bit more difficult. But overall, I mean, my recollection when I was at Sun, of the relationship with Yasufuku was good.

Remacle: What about with T.J. Rodgers?

Lacroute: You guys want to talk about that?

Diamond: Well, the relationship with T.J. started to create a full custom version of SPARC, called the 7C601. At the time, the Fujitsu gate array was the ten K gate array was the only...

Lacroute: CK22.

Diamond: Was the only SPARC implementation. And Cypress had actually started on the design even before we had a contract signed. So they were investing people and energy and money in that design. Of course, Sun was also developing the target for that chip at the same time. And T.J. was very forward-looking—as Bernie mentioned, he didn't have a microprocessor, but he was selling SRAM to system vendors who would populate that around microprocessors, and he wanted to move out of the—what was more of a commodity business, the SRAMs, into some of these more sophisticated chips like processors. So the companies were working very closely together, just on a handshake. And T.J. was very involved personally in that development. He would come to meetings. He would be personally involved in that process. He had actually come from a microprocessor background. Before he started Cypress, he had developed VMOS technology which could be used for microprocessors.

Lacroute: Oh, that's right.

Diamond: Which hadn't really been successful for digital logic, but is still around for power control. And by working with Sun, it was the shortest path he had to doing that.

Remacle: Bernie, you referred to the AT&T investment, do you want to elucidate a little bit on that?

Lacroute: Sure, in the '80s, when you build a computer, you had to manufacture it yourself, except for the boards, and you owned the inventory. Contract manufacturing did not exist at the system level. Especially in the case of Sun, which was trying to get the fastest machine, the fastest part, the inventory risk was enormous. Because you'd pay a premium for the fastest part, but if you did not sell the product within 90 days, your money has devalued a significant amount, since the same part was now selling for much less than you had paid for it. So Sun was burning cash at an incredible rate after it went public. We had tremendous amount of money tied up in inventory.

Remacle: What year did you go public?

Lacroute: It was late '85, I think. So we can double-check the date, but I think it was late '85 or so. And so Sun needed a lot of money, a lot of cash to build all those systems. And AT&T was struggling with their computer business, and they also wanted to make their version of UNIX, the dominant UNIX in the world. So we had a conversation with AT&T. And you know, it's very hard to resist a quarter-of-a-billion dollars in those days. So that was the deal. There was a price to pay for that. The price to pay was that we had to use their version of UNIX. And that was probably one of the most painful things that Eric Schmidt and a few other people did in their life, dealing with the AT&T guys to go and tune the AT&T Unix. But for Sun, it was a good source of money, a good backup. Again, a quarter-of-a-billion dollars in those days….I don't know, today that's probably worth two or three billion dollars or something like that. Lots of money. And the AT&T idea was that eventually they were going to acquire Sun. That was always in the back of their mind. Fortunately, we had a remarkable attorney, Larry Sonsini, and that agreement was crafted in a way that they just could not force an acquisition. Larry did an outstanding job on this agreement, so they could never acquire Sun. But on the other hand, I think that's the only time that they made money in the computer business, on the Sun investment. Every other computer venture they became involved with, they lost money. So it was good deal for them, too.

Remacle: Do either of you want to add anything to it? As long as we're on the subject of negotiations and so forth, Apple came up a couple of times in our planning discussion. Is this a good place to talk about Sun's negotiations, discussions, relationship with Apple over the years? Wayne?

Rosing: We tried hard. We saw this was at the point where it was now obvious that the 68000 family wasn't going to be able to continue to be competitive. And so we knew Apple needed to make a change. And I think the competition at the time was sort of Power PC from IBM, and...

Lacroute: 88K!

Rosing: Well, 88K a little bit from Roger Ross, but I don't think it had any credibility. That was Motorola.

Remacle: And again, the date? Date in this timeframe? About '87?

Rosing: No, this is going to be—

Diamond: '88/'89.

Rosing: '88/'89.

Remacle: Okay.

Rosing: So we made a good stab at trying to go to Apple. We were unbelievably clear about what we were doing. We didn't hold anything back. We explained exactly what we were doing, and what we were trying to do. My counterpart, who I did the direct sort of technical conversing with was Hugh Martin, who was there at the time. And of course, I knew Jean-Louis and other people did, but I had worked with Jean-Louis all the way back to my days at Data General. So there was a lot of ability to talk about this stuff, but there was remarkably little basis to ever really have a strategic alignment. It could have been good for both companies. Because we were fundamentally in different markets. Although maybe the SPARC Station was heading down towards where they wanted to go. But it could have worked, because I thought that the software was complementary, and the market was complementary between Sun and Apple at that time could have been very powerful. The cultures, I think, between the two companies just didn't quite—couldn't bridge that.

Diamond: I remember a meeting with Jean-Louis Gassée and Wayne, Bill Joy, Scott and me over at Infinite Loop. And I think they were very impressed with the capability, the architecture. I think the point was made that, "We're not in overlapping markets, that there's an opportunity." But they could not get past the point that we wanted this to be an open architecture.

Rosing: Right.

Diamond: Even though they were using the 68K and, ultimately, the Power PC, they were entranced with the idea that they could control their own microprocessor architecture. As, actually, we've seen them in their newer—you know, their current instantiations, of how important they see control of the microprocessor is to them. So really Sun's strategy...

Remacle: But that, in a way, it mirrors Sun's early decision to control the microprocessor.

Diamond: Yes, it does. Although, we were willing to let other people buy those microprocessors and build systems and compete against us. And of course, that wasn't—although, Apple had a dalliance with open systems, which is kind of an interesting...

Remacle: I think dalliance is a really good word.

Rosing: Yeah, it's a good word, yeah.

Diamond: Interesting story—in fact, I actually went—Joe Graziano, who used to be the Sun CFO and went over to Apple—asked me to come talk to Apple about open systems. And similar to what happened in this meeting, they really were not—they were not organically able to contemplate that idea. And their offer to us was along the lines of, "Why don't you forget about the open systems idea about SPARC? You can use the SPARC chips. We can use the SPARC chips, and we'll just cut it off." And Scott was very, very clear at that meeting that that was just not on the table. And that was really the nail in the coffin of that relationship.

Lacroute: The earrings were the nail.

Diamond: Well, the earrings came on the drive back. I was expecting Scott would talk about the strategic discussions that we'd just had at Apple, but he was very captivated by the fact that Jean-Louis had an earring. And he was French, and wondered, speculated as to the motivation for the earring. Was it some kind of signifier of something? And that really—instead of talking about the meeting, that's all we talked about going back to Sun. I was hoping we could have more of a substantive discussion.

Rosing: I remember that drive back, and I remember thinking about it. And I had been at Apple from 1980...

Diamond: So you already knew about the earring.

Rosing: Yeah, I knew, but it didn't faze me anyhow. But I had been at Apple for five years, and I really got the cultural difference of the need for sort of vertical proprietary control that's sort of in the DNA of the Apple culture. I think our need to control SPARC was very different. We weren't going to hand it over to another party that we didn't trust, because you simply have to recognize that the other party has a business interest that's associated with their shareholders. And if they're going to make a different business decision that you don't understand, then all of a sudden this baseline that enabled Sun to be competitive could be compromised. And it should be said that Sun was, at the time we got to a billion in sales, it was the fastest ever startup to a billion. And during the time we're talking about, Sun was basically doubling...

Lacroute: Every year.

Rosing: Every year. And I would go out for practically a year at a time doing these business trips. And I would be in-and-out of the office, obviously. But you'd come back—I'd sort of come back after an extended set of business trips and all of a sudden, I'd be around for a month, and I'd get around to the Sun buildings. And Sun was occupying twice, or two-and-a-half times as much as Mountain View as when I left on the prior cycle. So while this was happening, the rocket ship had left the ground in an extraordinary period of time. And even though I was one of the guys who was certainly pushing the open thing, I knew what my paycheck said, and whose name was on it, and we were in a period of extraordinary business success. And nobody was going to play with that.

Remacle: This is a good place to ask the question...

Lacroute: Can I get back on the Apple issue for a minute?

Remacle: Sure.

Lacroute: Actually, there were two sets of conversations with Apple. There was the licensing issue that we talked about. The other one, which I think is alluded to in the book around Steve Jobs is we don't know that the companies should merge. And those conversations took place as well. And didn't go anywhere. But that was...

Remacle: Who generated them?

Lacroute: Oh, actually, I'm not sure where it started. But you know, it was a time—it was in 1988, or something like that—when Apple was in a pretty precarious situation. They did not have much cash left. I remember a conversation, a meeting at Sun—not at Apple—with Spindler who was CEO of the day. And it actually could have been an interesting proposition. But it never materialized or turned into something useful. But that's the other piece of the Apple discussion.

Remacle: To summarize, there's been—you mentioned Viking …. there were earlier conversations about UltraSPARC and Spitfire. Can you put on a timeline for me for the future, kind of what were the major iterations of the SPARC chip over the years?

Rosing: You're best for that.

Diamond: Well, the first one, Wayne called it Sunrise, that was the internal name. That became the Fujitsu gate array-based SPARC, the ten MIPS first Sun-4/260. Then the next chip that came out...

Lacroute: And that chip—in 1987.

Diamond: We announced that in July, '87. By the way, just to mention, for historical purposes, there was another—whole other development on a bipolar technology with BIT, Bipolar Integrated Technology. They had a process that was supposedly as fast as bipolar, but as dense as CMOS. That didn't turn into a product. But that was announced actually in July, '87. We had BIT on stage with us.

Rosing: It didn't turn into a net product, but the chips got into Sun, the system was built, and it actually worked, ran Sun OS. But it just wasn't able to respond fast enough to the cost performance curve.

Lacroute: Right, price performance was not good.

Rosing: And at that point, Sun was in this sort of cash crunch phase, if you will? Where we were a little bit tight.

Lacroute: Oh, it was tight!

Rosing: But there were limits to what Sun could do, and going forward with this technology could have pushed us over the edge. And I think we were talking about 40/50 MIP class machine. So it was a step forward, but it was too big a leap. It fell down the cost performance curve.

Lacroute: Limited market, very limited market.

Rosing: Yeah.

Remacle: Okay, Steve?

Diamond: So then we had the Cypress-based 7C601 chip, full custom version of SPARC. And then during that time period, we also started what became the Viking design at Texas Instruments. And so those were sort of the first...

Lacroute: What was the timeframe on the Cypress chip?

Diamond: Well, we were designing that in '87.

Lacroute: Right. But when did it come out? In '89/'90?

Rosing: I think it was—Cypress chips were part of the SPARC Station production. So that was '89? Right?

Garner: <off-camera> The chip came out in '87, but production...

Rosing: Production was around '89 when the SPARC Station came out.

Lacroute: Oh, yeah, okay.

Rosing: That's when the volume really took off. And then there was the Viking came along. And then while Viking was underway, the UltraSPARC started. And the other one, what was the name, Spitfire? Which was it, Steve?

Garner: <off-camera> I think it was UltraSPARC.

Rosing: UltraSPARC. Spitfire, yeah.

Remacle: So Spitfire and UltraSPARC were the same.

Rosing: Right. And that was also the transition to the 64-bit world.

Remacle: Okay.

Garner: <off-camera> We had a gallium arsenide project thrown in there for a week or two. <laughs>

Rosing: Yeah.

Lacroute: I remember that, yeah.

Remacle: Can we switch gears a little bit and talk about the flat-out marketing stuff, most important, logo and SPARC name. How did those come about?

Diamond: Of course, those are the most important. I'm glad you got to that. So SPARC really started as Sunrise as we mentioned. And we decided that since it was our vision that it would be open, other manufacturers would be building systems, the fact that there was a word Sun in the world, we thought that that might be an inhibition for someone to put a logo saying "Sun" on a competitive product. So there was an effort to come up with an alternative name. And probably Wayne is a better one to talk about how we got to SPARC. It happened very quickly. It was in the musty recollection, one day it was Sunrise, and we needed a new name, and the next day we had all decided on SPARC. And I'm not sure how we got to that.

Rosing: Well, I remember sending out an email, declaring a classic naming contest. And I think there were lots of names, all very clever, but the one that I remember—we had started to converge on was ORA, for Open RISC Architecture.

Diamond: Right, I remember that.

Rosing: And ORA just didn't quite have the aura we wanted. And I do not know who sent it? Do you know?

Garner: <off-camera> Bill Brown wrote a program that scrambled all the acronyms. It was Sun Processor Architecture for RISC Computers.

Rosing: That's right. So Robert over here from the Audience is helping us on that. So one of the folks inside did a letter scrambler, and out came SPARC, and I remember seeing it and said, "Ah, that has a nice ring to it." And Scott took to it immediately.

Lacroute: Loved it!

Rosing: It just has that punch. It was neutral. It had punch. And there was an "S" in it. Which, of course, I always said was "scalable" processor.

Garner: <off-camera> You changed the acronym.

Rosing: Yeah, to protect the world from the Sun aspect of it. But that was it! It was not a cosmic decisions. It just felt right.

Diamond: We had a problem figuring out how to spell "scalable." There was the "scaleable" [*sic*] camp with the "e" in it; and the "scalable" camp without the "e" in it inside Sun, because it doesn't really look

right when you put the "e" in it, but that's how you spell it. So after the name came, we got the name cleared up, then of course, we had to do a logo. And immediately, Scott came up with the logo that he thought we should use, and it was a spark plug. Obviously, the pun is spark plug. And I thought that was terrible, just really terrible. Because here we have a game-changing processor architecture, a business-changing paradigm, Sun is a leading technology company, and we can't put a spark plug, old technology on our—so Scott said, "All right, I'll give you a month to come up with a logo in a month, or it's going to be the spark plug." And so we ended up hiring a very powerful design firm, Walter Landor Associates in San Francisco, and they came up with a suite of options. But very quickly, we condensed—we converged on the SPARC logo, roughly as it is today. I do remember that to get the ball over the end zone, Andy Bechtolsheim talked to one of his friends who is a typography designer. You know, Andy is a Renaissance man, so it's easy to get—if you need a typography designer, just check with Andy. And I think he talked to Hermann Zapf, who analyzed the proposed logo, and wrote up a justification for it, and we presented it to Scott, along with a color, which was special, and he bought off on it, and therefore we avoided the spark plug.

Garner: <off-camera> And there was issue spelling it backwards.

Rosing: Oh, yes.

Diamond: There was that issue.

Rosing: That came out in a presentation at some conference, and if my memory is working correctly, it was a fellow from TI named Peter van Cuylenburg.

Diamond: Ah, yes.

Lacroute: Ohhhh.

Rosing: Who announced, "SPARC is CRAPS spelled backwards." <*laughter*> Just in this big stately—and I was like, "Oh, my gosh!"

Diamond: I remember that.

Rosing: But it, fortunately, that seemed to have gotten overlooked, except for a few of us.

Diamond: We survived that.

Remacle: You're referred to the 1987 announcement. Can you talk a little bit more about that? Who all did you have on the stage with you? Why? What was kind of like the main message?

Diamond: I think we probably all have stories about this. I'll tell you what I remember. Of course, the key thing that I wanted to accomplish was to have Dr. Yasufuku onstage with us. Because this was '87, so we just finished a $200 million year, I think, something on that order with Sun.

Lacroute: Yeah, 250, yeah.

Diamond: And here we are announcing a new architecture, a new processor architecture, new operating system, new compilers, new system, and we were really just a very, very small company. And so I thought it was absolutely essential to have...

Remacle: Selling to large companies.

Diamond: Absolutely essential to have Fujitsu onstage with us. And I remember this was at—was it Hayden Planetarium?

Rosing: Hayden Planetarium.

Diamond: In New York. It was really a beautiful facility. We had a circular room. And we had something like 50 SPARC Stations, called Sun-4s in those days, around the circumference of the room. And we had two events. I think there was an analyst, first analyst event, and then there was a press event. And all these were set up, running. And an hour before T0, Conn-Edison had been digging in the street outside the Planetarium, and the power went off. So suddenly the room, you know, of course, things are at a fever pitch, people are plugging and putting up signs, and putting on their microphones. The power goes off. And I think the power was off for something like an hour. It seemed like a week. And the question was, "Should we cancel the event? Should we have—you know, bring in flashlights? What are we going to do?" And then the other question was, "Well, if we turn these 50 workstations on, how many of them are actually going to be working?" I mean, it really wouldn't do to have 50 workstations around there with only 20 of them with the screens on. And there really wasn't a way to get the ones that weren't working out of the room if they weren't working. So chaos was really reigning. Well, I'll just say that when we turned those—the power came back up, we turned the 50 workstations on. They all came up working. All of them. All 50 of them.

Remacle: Was there a cheer that went up in the room?

Rosing: Yep.

Lacroute: That was after the previous day incident where a bug was found. And talk about the bug, Wayne.

Rosing: Oh, well, there were many bugs, but there was a bug the night before.

Lacroute: A fatal bug. <laughs>

Rosing: Yeah. You want to tell that one?

Lacroute: No, go ahead.

Rosing: Well, that was—I don't remember the exact nature of the bug, but Steve Kleiman and a couple other people literally pulled an all-nighter, and found the bug, recompiled, I believe it was the OS and the basic libraries. Had that running. And there were various recompiles of various vendor's applications. And I know in that morning, before the announcement, I don't recall the exact vendor, Bernie thinks it might have been Computer Vision, I think it might have been ANSYS, it really doesn't matter. Some vendor ported a half-a-million line Fortran program. They just put it in. Compiled it. It ran, and we had yet another ISV on our list that morning.

Lacroute: Who could stand up and say, "Guess what I did this morning?"

Rosing: And say, "Guess what I did this morning?" So I thought the interesting thing about that announcement, compared to most of the ones I've done in my life, was Sun really had a lot of stuff there. It really worked. And real live, right there in front of the world, the strategy of portability was working. And I think it turned out to be, "Hey, it's just a Sun workstation, and it's really fast, and really well-priced, and therefore, this is great stuff!"

Lacroute: And Fujitsu added credibility.

Rosing: Yeah.

Lacroute: There's no question, because people were kind of worried whether or not where are those chips going to be available? I mean, it's fine, you can build one, but can you build a hundred, a thousand, ten thousand?" So having the Fujitsu presence was really critical and reassuring for the customers.

Remacle: Steve, do you want to finish your Yasufuku story? You started...

Diamond: Well, he was onstage. And endorsed SPARC, endorsed the openness of SPARC. Talked about Fujitsu being a future SPARC vendor, and in fact, ultimately, they really proved out the architectural flexibility by building embedded SPARC chips, which among other things, were very frequently used in cameras. microSPARC was an embedded version of SPARC. And so in the early days of SPARC, you know, the traditional architectures, like the 68000 X86, you would never have then seen those in an embedded application like a camera. SPARC was really scalable. I think it proved out the "S" in scalability, by being able to be scaled down to an embedded processor. And Fujitsu, of course, they didn't talk about the embedded chips then, but they talked about their commitment to SPARC.

Remacle: Just out of curiosity, was there anybody from Motorola there? What was Motorola's response after a couple of years of questions?

Lacroute: If I recall, we actually had invited Tom Gunter, who was VP of 68K development, and he didn't come. That was understandable, but I mean, it was the polite thing to do to go and invite him, but he didn't come.

Remacle: What was the response internally, after a launch, you talked about the chaos. There's usually kind of a fever pitch of excitement on a major launch. What was the response of the rest of Sun to that launch?

Lacroute: Oh, I mean, very enthusiastic. They loved it! We've done it! Here, we showed the world that it can be done, so it was a lot of positive energy. Yeah, it's there, it's real. Especially speaking for the engineering communities. "You know, we've done it!"

Remacle: "We did it."

Lacroute: I mean, it has been done. And so very, very, very positive. I think it was enthusiastic.

Rosing: Sales were strong.

Lacroute: Yep.

Rosing: It's interesting, if you look at unit quantities of computers shipped by Sun in this time, we're probably talking, what? Thousands a month?

Lacroute: I would say about 10,000 a year or something. So that's about it.

Rosing: It's still an interesting point of scale. Maybe five/ten thousand computers a year. I mean, I was at Sun when the millionth—was it the hundred-thousandth, or the millionth? I may be wrong. It was a SPARC Station that passed one of the big thresholds. SPARC Station 1.

Lacroute: Must be the millionth.

Rosing: Yeah, millionth, yeah. So that—by the standards we now live in, these were very small quantities. But it was just a different game.

Remacle: Speaking of quantities, does anybody have any idea how many SPARC chips have shipped to date?

Lacroute: I have no clue.

Diamond: I don't, but since it's also an embedded processor...

Remacle: It's got to be bazillions.

Rosing: Yeah.

Diamond: It's got to be tens of millions.

Rosing: Yeah, easily, I would think.

Remacle: Just another reflective question is how was the SPARC effort and team seen besides the marketing people who didn't get as much money as they'd like, but by the rest of Sun. Was it seen as kind of a renegade group of people? Or an integral part or...

Rosing: Integral. I think there was a little bit of competition, which is healthy between the 68K gang and if you will the SPARC gang. But once SPARC had launched, that went away, and then about a year or two later there was the great year of all the wood behind one arrowhead, where we sort of declared it a company strategy for a year that Sun was simplified, and we knew what we wanted to do, and we were going there. And that just kept the momentum going. So the company came around, because Sun engineers want to be on the cutting edge. They want to do performance-oriented things. And I don't think

we could have gotten good engineers to work on a 68040 system. They would want to work on the SPARC stuff then.

Remacle: Another way of approaching the same topic is how important was SPARC, or is SPARC, to Sun's long-term success?

Lacroute: Oh, I think it was absolutely critical. That's one of the things that made the company. The open architecture marketing approach, the portability of the application. It is SPARC, it is UNIX and it is the marketing campaign around open system. But without this key differentiator, which really made a big difference in terms of price performance, The company would not have been as successful. If you look at the pricing strategy of the original SPARC it was very disruptive and bold. The idea was to price the machine basically at the same price as a 68K equivalent, with four/five/six times the performance of that 68K machine. It really built the company and allowed Sun to continue doubling in size every year. And I don't think it could have been done with just another commercially available chip. So I think it was absolutely fundamental to the success of the company.

Rosing: And that particular period of time, in many ways, the Internet, as we began to enjoy it in the '90s, was still taking shape. One of Sun's key markets at the time was the University market, government researchers. And as Steve alluded to, it was considered chic that you had a Sun workstation as your computer. So a lot of what we think of as the Internet was prototyped and developed on Suns by hundreds, if not thousands, of people all over the world. So we had a really good track into that sector. And so that, if you remember, what, mid-'90s, one of Sun's marketing lines was, "We put the 'dot' in the Internet," or the ".net," or whatever that tagline was. I don't remember it exactly.

Lacroute: "The 'dot' in .com."

Rosing: But it came from the fact that Sun was the vendor of choice for a software engineer who wanted a high performance system to do their work on. We were the ones people wanted. And so that just continued to expand ultimately the application software that was available. And that's the ultimate fuel that just kept this whole thing going.

Lacroute: So to me, SPARC is the example of what a good start-up is all about. You take a big risk, and you succeed, and the payback is enormous. And somehow, my worry long-term is that we are unable to maintain that kind of culture, which has really made companies like Sun. But also, what made the US. I came from Europe, I came from France. You don't have those things, because people are not willing to take those kinds of risks, with a "k." And that is really fundamental to what this place has been all about. And SPARC and Sun are a good examples of this. So I think, you know, I'm very proud of what has been done there.

Remacle: Discussion was about what it took to get the SPARC effort to the success point, and how that reflected Silicon Valley in general, or what it takes to be successful in high tech and the world, in general, not just Silicon Valley. How else would you categorize or describe SPARC's contribution to changing or impacting both the semiconductor industry and the computing industry? What would you say is SPARC's, in particular, major contributions would be in that respect?

Diamond: Well, I think if you look other systems companies, and ask arguably, "Who has developed a successful long-term microprocessor architecture, who wasn't already a semiconductor company, like IBM, with the longevity of SPARC?" It's 25 years down the road, there isn't anyone. I think this is a singular unique achievement that Sun, which wasn't a semiconductor company, had the ability and the courage to develop a microprocessor architecture that's still here. It's still at the heart of leading edge Oracle systems. That's never happened. And I think it's indicative of the courage and capabilities that Sun had that they could start it off with such momentum behind it that it's still running.

Lacroute: I would like to put this in the overall context of the history of computer systems. If you look back in history about the great architectures, you certainly have the IBM 360, with OS 360 tied together. That endured. I'm obviously partial to VAX and VMS. That's another one that I think set the stage, so the first one was the mainframe. The VAX VMS was minicomputers, timesharing. And I think in terms of workstation business, SPARC, Sun, UNIX and Ethernet integration is another classic category of those architectures that are going to stand the stamp of time, because they represented a significant evolution in the way we did computing. Today, we have our cell phone. That's probably another one. But I put SPARC, Sun, UNIX in that category.

Rosing: Yeah, and I think that the PC evolved over time. It sort of came via a more complex route, but the commodity PC, you still see the same thing. Graphics, Ethernet now. It's networking. More or less standardized software, more or less standardized hardware. But there's still a very distinct difference between what you can get in the Intel Microsoft world, and what you can enjoy in the broader UNIX world. It's a different world, much more productive, primarily for software engineers. And generally more productive for CAD people. And it's enduring. It hasn't really been replaced, and markets have kind of found their equilibrium with each other, and neither one of them can all of a sudden take over the other one. They're just serving a purpose, and I like to think of this as we found a sweet spot that's generally really good for a large class of technical computing and high-performance computing.

Remacle: You know, Scott McNealy name has come up. There were a lot of other high-powered, high-energy folks, including the three of you at Sun. Can you talk a little bit about the personalities of some of the Sun leaders, and what was it about Sun that attracted so many people that have moved on to become important executives—again, you guys are all reflective of that—in other companies in the Valley?

Lacroute: Well, I think Sun attracted people who were bright, who wanted to succeed, who were competitive, and were absolutely committed to do the very best that could be done. I think it's true in

engineering, but it certainly was true in marketing, sales force. Those guys were driven! There's no question about it. I remember some of the manufacturing guys. They really wanted to win. So that spirit, that team spirit said, "We are going to be the very best, and we are going to do what it takes to get there." And I think the fact that Sun was able to attract those kinds of people is one of the things that made the company successful.

Rosing: Yeah. I think that with reference to Scott, Scott was always a—he served one very valuable purpose. Particularly at the time when I was there. It was when I had—I can only speak to my own experience. He was a clearly identifiable leader for the small number of really important key themes in the company. And Scott would get on a theme and he would stay on it. And he would stay on it with great focus, and very engaging. He was also passionate about manufacturing, very, very passionate about the salesforce. He certainly understood engineering, but I wouldn't call Scott passionate about engineering, but he was passionate about those other things. But he totally understood that the technical excellence of the company was the necessary to have a product. And so Scott got a number of things very right in these days. And he kept the focus up from, I think, his contribution to SPARC beyond obviously the early part of it as the CEO saying, "Okay, guys, go do this crazy thing," was the consolidation after we started to get successful. He was really critical in getting the whole company focused. We'd started an engine that could be successful, and then we revved it up and kept it there for many, many years.

Diamond: I think in Scott's case, there were a lot of forces that—we mentioned Sales, that could have derailed the SPARC effort. And he had to tread, you know, a very careful line between supporting the SPARC program, and making sure the sales force generated revenue to keep the company running. And you know, he had the courage of his conviction going forward with SPARC, but he also was capable of leading the company before and after the announcement so we stayed on track. Also, I think Bernie mentioned this, in the early days, we mentioned him earlier—Joe Graziano was a superstar CFO.

Lacroute: Yep.

Diamond: He went over to Apple...

Lacroute: No, he came from Apple.

Rosing: Came from Apple.

Lacroute: Yeah, he came from Apple.

Diamond: There was a rumored million dollar sign-in bonus he got at Apple.

Lacroute: Well, that was before Sun.

Diamond: And how many CFOs do you know by name? You know, the kind of people that Sun attracted are people who were superstars in their area, their field. And I think that's why the company was able to do what it could do. It had people who were capable of generating both the ideas and the execution.

Remacle: One last question for each of you. What did you take away from Sun that you have used in your other steps in your career after Sun? Wayne?

Rosing: Well, I certainly gained a great deal more appreciation for what it takes to be a leader in a highly technical culture. And when I had come from DEC, which has a tremendous culture, to Apple, Apple already had sort of a culture. So learning to transplant and bring some stuff in. And then when I went back to Sun, Sun was a smaller less-formed culture, and more technical. Sort of a lot of Stanford. It's no joke, the "S" in Sun is Stanford. And so there was a sort of, I'll guess, call it a Berkeley/Stanford culture at Sun. And then I helped, you know, drive it in my own way in terms of what we did there. And then when I went to Google, which was really a Stanford computer science dorm, I was able to sort of say, "Okay, this is a great culture. This is what Sun was. How can I, as a manager, preserve this culture, and then enhance it and make it go?" So what I got out of all this, for the SPARC effort, and the Sun time, was sort of an understanding of how to preserve and drive and create really high performance cultures, all built on the principle of the very best people in small teams with very clear goals that you don't mess with. You let them run, and then you rigorously measure success. It works because it's a good product. And if it's not a good product, you cut your losses and you move on.

Lacroute: Just a few things. The first one I would say is that whatever you do, you should want to succeed, you have to be passionate about it. So you have to build an environment and a culture of passion. Not number-crunching; passion. You've got to be driven. And the second point that I would make is that it's people. People, people, people. If you get the best people, you are going to do impossible things! And so never forget that you've got to get absolutely the very best. Passion, people—and in technology, products. If you don't build products, it's not going to work. You can have all the marketing around, and marketing is very important, and I use "product" in a large context. It's not engineering, but it's marketing, it's selling, it's building it and so forth. But products are everything. So those are kind of the three things that I would take away from this experience.

Remacle: Steve?

Diamond: I think what Wayne and Bernie have said most of what I would say. I think in typical engineering fashion, it's possible to get into analysis/paralysis. And I think, as Bernie mentioned, if we had just undertaken the typical spreadsheet approach at Sun, this never would have happened. Sun would never have grown as fast as it did, and we wouldn't be here today answering these questions. And so I

think it's the people, the leadership and the passion. And of course, you have to be right about some technological and business bets, otherwise they don't pay off. I think there are a number of them that Sun pursued and that did pay off. When you're able to put all the wood behind that arrow, those forces, the people, the technology, the right business decisions, the passion and the leadership, you can accomplish something that seems impossible. And that's what I've learned from Sun.

Remacle: Is there anything else that you'd want to get captured for posterity? Have we covered everything thoroughly enough? Anything we left out?

Rosing: Thank you to somewhere around a thousand to two thousand people at Sun who aren't on this stage, but contributed every much to this effort.

Lacroute: That's right. It's the people.

Rosing: Everybody in the company, almost to a person, got it and was behind it, and pulled on the rope.

Remacle: All right, on that note, I say, "Thank you all for all of the time, not just in front of the camera, but the prep time it took, and thank you very much.

Diamond: Thank you.

Rosing: Thank you.

Lacroute: Thank you!

END OF PANEL SESSION

