Oral History of Wayne Rosing

Interviewed by:
Uday Kapoor

Recorded March 7, 2019
Mountain View, CA

CHM Reference number: X8969.2019

Kapoor: On behalf of the Computer History Museum I would like to welcome Wayne Rosing for his oral history. It's a privilege to host you. My name is Uday Kapoor. I'm a volunteer in the Computer History's Oral Histories Program and it's an added privilege because I worked with you in the past and I've known you for so many years, so welcome.

Rosing: Thank you. I was just thinking today, it's 30 years since SPARC was introduced.

Kapoor: Yes. Yes. Exactly. So Wayne Rosing's training has been in mathematics, physics and astronomy. He's a computer engineer by vocation and has been programming doing computer and electronics and optical design and telescope engineering since high school. His career has spanned engineering leadership roles at Digital Equipment Corporation, Apple Computer and Sun Microsystems. Most recently, Wayne was Senior VP of Engineering at Google from 2000 through April 2005. Wayne and his wife, Dorothy Largay, refired to Santa Barbara as against retired, in 2000. In 1992, Wayne founded what is known as Las Cumbres Observatory Global Telescope Network-- excuse me-- LCOGT, now located in Goleta, California. At LCOGT he and his staff of engineers and scientists have developed scores of 1 and .4 meter aperture telescopes that have been placed in various locations around the world for research and education. So with that, let's get started.

Rosing: Okay, that's great.

Kapoor: I hope I said everything correctly, if you have any correction.

Rosing: That sounds great.

Kapoor: So let's start with your childhood. Where were you born and when?

Rosing: I was born in 1946 in Cleveland, Ohio and my family left Cleveland in 1950 and we moved to Phoenix, so essentially I grew up in Phoenix. And then I went to Berkeley in 1964 and my family ended up moving to San Francisco within a month or so of that.

Kapoor: So how was your childhood like? What were--?

Rosing: Well, childhood was I think pretty traditional. I had a lot of interests in science, electricity and mechanical things and doing mechanical things. My grandfather and father were both contractors and my grandfather took a particular interest in my interest in science and was very much supportive of that so I, as I was saying at lunch, my earliest memories are of going out in the backyard looking at the moon with binoculars and then a telescope with my grandfather, so I had a sort of classic upbringing of an extended family and somebody who was very encouraging of my exploring things when I was younger.

Kapoor: So what did your dad do?

Rosing: He was a homebuilder, custom homebuilding in Phoenix.

Kapoor: And you had siblings? Do you have siblings?

Rosing: I'm the oldest of four, so.

Kapoor: Okay. So what was your schooling like and were there any mentors, people that you followed?

Rosing: Schooling was, I think, pretty traditional. Started out in a Catholic school that was very overcrowded and in the sixth grade, at the end of the fifth grade it became obvious to my parents that the school wasn't able to provide at least what I needed and I think my brothers and sisters, so we all switched at that time to public school, Madison Number 1, for those who know Phoenix. And that was a pretty normal first year, but the second year I got a science teacher named Mr. Ferguson who really took me aside and really encouraged me on both math as well as generals, [ph?] my interest in science. And so I started to change from sort of being a student to studying at school more on my own sort of in the back of the room because I generally got to be ahead of the class most of the time. So I would take tests and do okay, but I didn't have to pay attention in class as much, so I could read in the back and that became my lifelong habit.

Kapoor: Very good. So when did you graduate from school?

Rosing: 1964 was high school and then I never finished Berkeley. I was started there and I got very, very interested in computers relatively early at my time in Berkeley. And by the way, this was 1964, so Berkeley was somewhat of a chaotic place. I found my place sort of working in Project Genie and then at the Lawrence Berkeley RAD [ph?] Lab doing a lot of computer work and that ultimately culminated in I think it was 1968 I got a job offer to go to-- from Wesleyan University in Middletown, Connecticut to go join a research project they were doing down at Tucson at the University of Arizona called SCLERA, Santa Catalina Laboratory for Experimental Relativity and Astrophysics, I think that's the approximate name of it. SCLERA. And at that project, I designed and built my first computer and that sort of really set my career vector from then on.

Kapoor: So you didn't actually graduate--

Rosing: No, no. Partially way through but then got distracted and--

Kapoor: I see.

Rosing: There was never any reason to go back.

Kapoor: Okay. So after the project that you mentioned, I understand you went to Data General, is that correct?

Rosing: Right. A little bit of a circuitous route. I finished designing that computer and building it and that was successful and operated for I don't know, eight or nine years, ran the experiment that was, it was designed for. I went back to Wesleyan, worked there in their electronics area and computer area for a

while. Sort of started my own company on the side and was doing custom electronics work for people, computer interfacing mostly. And then I got a contact from Data General to join them and be a part of their custom systems group which was doing custom add-ons to Data General computers.

Kapoor: I see.

Rosing: I ran that group for a number of years.

Kapoor: So would you say that you were primarily in the hardware design?

Rosing: At this time I would say from the-- with the exception of when I built the computer in Tucson, everything I did was hardware-based and then I would write some software if I was designing a product for somebody. I'd write the application software but it was not an initiative where you'd say you start with the software and add some hardware, it was the other way around.

Kapoor: Right, right. So the technology at that time was discrete devices or--?

Rosing: Well, let's see. Certainly that was pretty much discrete integrated circuits if I recall, like, for instance, everything we did at Data General you could arbitrarily say it was the TI 7400 catalogue. Now there were many other vendors besides TI, but integrated circuits was the game and gate arrays weren't really around very much. They were very much exotic and in their infancy at that time.

Kapoor: So how long were you at Data General?

Rosing: I was there until I think '77 or so and then I went to Digital.

Kapoor: Okay. And how did that happen?

Rosing: How did it happen that I went to Digital? Data General really didn't have their heart in expanding custom systems and the Eclipse project was well underway and staffed, so there was just sort of no real obvious place where I could do the next thing that I wanted to do, which was get involved in heavy duty computer design again. And I had a friend and he said, "Hey, Digital's interested," and I went up to Digital and we had a talk and I was hired originally to take over the essentially the PDP-11 product line from an engineering director point of view. And I started there and quickly moved into the VAX because I was taking the position that actually building yet another PDP-11 in 19-- what would it have been-- 1978, seemed like a bad idea. Digital should put all its energy behind the VAX effort. And I'm not saying I triggered that by myself, but as a result of that, the 11-- I think it was the 1172 MP or some product like that was announced. The first units were on the loading dock. And they actually made the decision and pulled the units back in and cancelled the project and the next day I found myself manager of mid-range VAX development instead of PDP-11, which I thought was good. And then out of that came the VAX-732.

Kapoor: Right. So how large a team did you have at DEC?

Rosing: Oh, I started with maybe 20 and that morphed eventually into an advanced development group that were maybe another 20 to 25. And so that would be '78 to '79.

Kapoor: So again, this is primarily the hardware design?

Rosing: Yes. At that time it was hardware and the two things that I was-- There were three things that I was sort of focusing on. One was this thing that I used to call a workstation and the concept was that you would take a small VAX. You'd put it in a desk. You would put a 19-inch monitor on top of it. It would have a keyboard, this little thing called a mouse. And instead of a, if you will, a display that's just a bunch of characters, it would have these abstractions on the screen called windows that would have characters in them. So we I think all recognize what that is today. And at that time I was also in charge of a project to sort of organize DEC's networking approach and I had been talking with Bob Metcalf who was doing a sabbatical at DEC. So added to my vision was this idea that you had an Ethernet connecting everything together and that at the other end of the wire there could be some VAXs. But this was sort of anticipating Micro VAX, which had-- was in its very earliest infancy at this time.

Kapoor: So the operating system was the VMS.

Rosing: The operating system was VMS. I wasn't very involved in that. We were all in the same building so I knew the people but I wasn't directly involved in that. But I said, there were the three things I was looking at. One was the idea of workstation computing. The second thing was we were thinking about, okay, we need to build bigger VAXs. The 11/780 had at that point been out and the 750 was coming out and where was this going from here was a question. And I saw that one of the problems in designing these large systems was computer simulation, that the CAD tools and simulation tools were critical to doing these larger projects rapidly and well. And so I was asking the question this way, why can't we build a 50 MIP VAX, not necessarily a product, but a prototype that we could use internally, and that was one of the questions I kept asking. Well, a number of things came out of that. The VAX 8600 sort of came out of some of those speculations. I don't think it would ever achieve 50 MIPS, but it was a fast machine in its time. And then that work led me to start to do some research and I read some of John Cox's papers from IBM Research and I sort of got the idea. Now I didn't know the term RISC at that point, but whatever you want to call it, it's now so indelibly RISC that for the sake of this conversation we'll use RISC even though Patterson hadn't invented it yet. And then very curiously, which closes an interesting loop, Dave Patterson came to Tewkesbury in I think '78 for a sabbatical with his family and he ended up in my group he was located physically, although he was much more connected with Bill Strecker, Dilip Bhandarkar and some of the other people doing more formal architecture work. But I talked a lot with Dave about the RISC ideas and so that was an interesting little rendezvous back in 1978 that laid the foundation for what happened later.

Kapoor: Right. So was Bernie Lacroute there at that time?

Rosing: Yes. Bernie was head of product management for the VAX. And there was a great little story about this. I was pushing for this idea of the workstation, the computer on your desk, and there was some strategic meeting that happened and some decision was made and it ultimately resulted in the, "No,

Wayne, we're not going to do your workstation." But the great quote as I recall it, I'm not going to say this is a literal quote, but at least it's a memory that may have been manufactured to suit the time, but Bernie's quote or his imputed quote at the time was, "Time sharing is not in our-- Time sharing is our strategy." And of course, I was rebelling against that in the sense of my whole career had always sort of been build smaller, faster, cheaper computers that I could have in my house, that I could have my own computer, that it wasn't an exotic thing. And so the idea of bigger and better and faster and more expensive was sort of the opposite vector.

Kapoor: So was that the reason that you ultimately left DEC?

Rosing: Yes. And again, to close an interesting loop, John Couch, who was the head of the Lisa project at Google knew Dave Patterson. Asked Dave Patterson if he knew somebody--

Kapoor: You mean at Apple. You meant Lisa--

Rosing: At Apple, I-- Excuse me.

Kapoor: Yes, mm-hmm.

Rosing: Exactly right. John Couch was the head of the Lisa project at Apple. And I'm going to say this would have been 1980 now, so the year after Dave Patterson went back to Berkeley. John called Dave, asked if he knew anybody who'd be a good engineering manager and Dave recommended me. So I get a phone call from John Couch and this happened the day after this decision had come down that no, Digital wasn't going to do this workstation style computing. So I was, I wouldn't say I was cross, but let's just say I was, you know, well, whatever.

Kapoor: Disappointed. [ph?]

Rosing: And John Couch describes the project, you know, 68000 computer, a bit map display, a mouse and it was modeled after the Xerox PARC work. I'm not sure the term STAR was used at the time. Well, I just said to John, "I'll come out tomorrow," and literally I flew out the next day and interviewed and was offered the job and then gave my notice and two weeks later I started at Google and that was that, so.

Kapoor: So, tell us about that experience, once you joined Apple.

Rosing: Well, again, I tripped up. I said "Google" and I meant Apple. Let's see. Apple was an extraordinary place. You had this clearly successful company that was on a very good vector. They had mastered their manufacturing. They had incredibly linear manufacturing of the Apple 2. They were building them by the ton. Everything was working. Their vertical integration was working. And the company was just doing well. They had struggled mightily with the Apple III which was a more difficult project and I think it was ultimately the issues were not technical, they were managerial and Apple didn't quite know how to manage, I think, big projects that all of a sudden now everyone is on your case all the time. Well, the next big project at that time was Lisa and so I took that over and I was the engineering

director. And that was a project, there was a mix of hardware which was developing everything from scratch and software, developed everything from scratch. Now I never really thought about that at the time but now as I look back, I really wonder what that was about. But there was really an ethic at Apple at the time about they want to own it, it's proprietary, it's theirs. Let's put it this way, it was the exact opposite of open, except the open-close dialogue didn't use those words. So, it was proprietary therefore you didn't license it, you built it and that's how it was done. So I was very involved from day one and we're building our, writing our own compiler and building our own operating system, developing our own windows system, developing our own graphics stuff. The mouse was a little different because it was sort of okay, how do we take this technology that we've been allowed to see and learn about and use and how do we go about designing a mouse that's rugged and important. But we did everything and it was just a ton of fun.

Kapoor: So this was your team that did everything?

Rosing: Yeah.

Kapoor: Okay.

Rosing: The Lisa was sort of built from the ground up with unique technology. It inherited an integrated circuit that Wozniak had built called the IWM, Integrated Woz Machine. It was a chip for encoding and decoding floppy disk bit streams. But other than that it was just, we just did our own thing and it was--

Kapoor: Right. So Steve Jobs took a lot of interest in that.

Rosing: And Steve, Steve was interested in it but Steve was not in this program. And then while we were doing Lisa as Lisa sort of moved from completely nothing to we had our very first prototype, then Steve conceived of the Macintosh as depending on how you want to call it, baby Lisa. It was clearly meant to have a design center that wasn't business-focused. Lisa was business-focused. Steve very much wanted the computer for the rest of us. Again, I'm now using all the later on marketing terms, but he had a vision of a very small, inexpensive computer and that was the Macintosh. And then everything flowed and a lot of our technology that we developed found its way into Macintosh because it just seemed logical and it was appropriate, so it was a really fun, creative time.

Kapoor: So you stayed at Apple until this transition took place.

Rosing: Yeah, when-- I stayed at Apple until 2005 and then I think it was in April I joined Sun.

Kapoor: So you thought that you didn't have a role at Apple, that you wanted to move on? Or how did that--?

Rosing: Yeah. I didn't feel I had a role and I was really interested in sort of the technological space. I wanted to get more involved in computer architecture and what it takes to build the next computer. So I

remember saying in my goodbye talk that I was looking forward to going backward to technology that I understood and that wasn't likely to be in--

Kapoor: That seemed to be what you were interested in all along.

Rosing: Right.

Kapoor: And that was where your strength was.

Rosing: That's right. And of course, the great irony is Bernie Lacroute gave me a phone call and said, "Hey, could you come over? We're interested." And I said sure and we interviewed and then changed how Google was organized. Excuse me, I keep saying, "Google."

<laughter>

Kapoor: Sun.

Rosing: Sun was organized a little bit and then Eric Schmidt and I together started running engineering.

Kapoor: Okay. So both of you were reporting to Bernie.

Rosing: We both reported to Bernie and we basically made a deal right up front that let's sort of share it more. So I would do some software and he would do some hardware. We would each be learning more. I'd be learning more about software engineering and he'd be learning more about hardware engineering and that worked just great.

Kapoor: Right. So at this stage at Sun they already had a product? Or what was the state of Sun?

Rosing: Sun at this time was shipping the Sun-2. The Sun-1 was the sort of metal square box that was Andy's original system. The Sun-2 product line had a workstation and some servers. And that line had taken off and was doing very well.

Kapoor: This is still Motorola-based?

Rosing: Motorola-based, 68000. I think 60-- I'm going to say 68000 but somewhere in there there was a 6810 or 20 but I don't quite remember exactly. And when I joined Sun I find out basically at my interview that they're working on a RISC machine.

Kapoor: Oh, so that decision had already been made.

Rosing: That had been made. It was in its very early form. The spec existed. There were a small number of people on the team but there was a compiler team, there was an operating system port team of two and then there was a team of I'm going to for the sake of conversation five or six engineers primarily

working on the RISC chips and then a couple engineers working on a computer board. And at that time, RISC was conceived of as a server. It was more server-oriented than workstation-oriented, as I recall, when I first got there.

Kapoor: So Dave Patterson was already--?

Rosing: Dave was a consultant.

Kapoor: Right.

Rosing: He was in and out. I met Dave again and that was very fun to close the loop on that again. And but we were going gangbusters. I don't think I added much to the project except became its strongest cheerleader and started figuring out, well, it's going to be successful so what are we going to do next and how are we going to take this thing forward.

Kapoor: So Anant Agrawal was there at that time.

Rosing: Anant was there. Robert Garner, Ed Kelly, Anant, Masood, K.G. Tan. I'm sure I'm short circuiting a lot of names, but.

Kapoor: That's okay. So the gate array design was then already done or was it started?

Rosing: The gate array design was being done and it was finalized. And it was really exciting when we got our first chips back and we were able to start testing the performance and making sure we could hit our goal. Our sort of internal goal was we wanted to be 10X a VAX-780, so that machine was nominally referred to as a 1 MIPS [ph?] machine, so we wanted a 10 MIPS machine.

Kapoor: Yes. When I spoke to Bernie when he was interviewed, he said the same thing.

Rosing: Right.

Kapoor: So that's very consistent, yes.

Rosing: And then we figured out, well, if you could make it run faster, it could be a 12 MIPS machine and then it could be a 16 MIPS machine. And so it was a really interesting time because it was sort of this peculiar window where if you will, a company like Sun didn't need to yield its product line to a semiconductor vendor designing a big CPU chip and so it was do you use the Motorola next generation or do you go do your own? And of course the point is Sun had a leadership position in workstations at this time, but it was still heavily competitive with Apollo and Prime and how could you distinguish yourself if everybody was using the next Motorola chip? Everybody's going to get the same frequency. Everybody's going to get the same briefing. But SPARC was one way that we could get a demonstrable technical advantage.

Kapoor: Yeah. So I heard of course, you joined when the decision had already been made.

Rosing: Right.

Kapoor: That there were a lot of internal discussions and I wouldn't call it battles to decide on the RISC approach because of the risk involved in doing one's own chip and so on.

Rosing: Right. There were. A lot of those had been fought and the blood had been spilled, but I just sort of, I got behind the project and because I felt it was, its design center was right. I did have one small reservation about it but I felt it was right and again, I have a tendency to pick up later on terms and reapply them to the past, but I remember very clearly articulating in my own way, "The company has to get behind this," and so I just became the cheerleader for all the wood behind one arrowhead. Now there were people who were very much in the Motorola camp and Howard Lee amongst others, and he pushed forward on Motorola, which is fine, because there was no point in being suicidal, <laughs> okay. But I tried to get everybody behind it instead of second guessing, well, do we do this or we do that. Do we stop it and use the MIPS chip? There was, you know, and then Motorola came up with a half-baked idea for a chip and I don't remember its name anymore, but they were going to kind of do sort of a RISC, sort of a RISCy thing. And so there was a lot of stuff, but I just said, "No , this is our project. We're going to stay behind it and we're going to stay on it."

Kapoor: So I understand there was also an opportunity with Intel as well. There was Intel architecture was considered but not really.

Rosing: I don't recall that one as much. I think that may have--

Kapoor: Happened before.

Rosing: Come and gone before my time, but.

Kapoor: Yeah. And then as far as the gatearray implementation, that came out and worked and that really got the RISC or the SPARC going. Is that--?

Rosing: Right. That got-- that certainly got SPARC going. And I'm not 100 percent sure but I think Sun was the first large company to really push a RISC-like architecture. MIPS was behind us but they didn't really have a computer, they had a chip and then that was the vendors had to then follow suit with the systems. But, yes, that's--

Kapoor: So yours and Eric's joint effort in terms of supporting the compilers and all the software to support the first RISC implementation was on target.

Rosing: Well, that worked and the most important thing that I think contributed to the success of the SPARC effort was this remarkably simple concept that we ported the OS that was the Sun OS. We didn't invent a new OS. The compiler people designed a good optimizing compiler that exploited the RISC

architecture but delivered the Sun software. And so even the day we announced it at the Museum of Natural History in New York City there were vendors in the basement bringing in their source code and compiling it so they could then be able to have their name on the screen that they're there. And you know, stories are legend about a half million lines of code that just compiled and operated immediately. Well, guess what? That was Sun's secret weapon because now we had our whole software catalogue virtually instantaneously on a faster computer at a much better price point and that was it. That made the RISC thing happen for Sun.

Kapoor: Right. So then you also had the networking infrastructure.

Rosing: We had the-- Of course we had the networking which was the other side of the equation. And we had server and we had clients', if you will, workstations. So we sort of had the complete picture fairly well together. And that really took off and then the transition away from Motorola became very simple. People just stopped buying those machines unless they had some compatibility reason and so that market just withered away. And then we just went crazy at Sun starting the Viking project--

Kapoor: In fact the Cypress--

Rosing: And the ECL machine and the Cypress project and--

Kapoor: And that was where I remember.

Rosing: Right.

Kapoor: I was involved and it was really exciting from my--

Rosing: I had a slide for about a year, this was the time when you still had foils and you'd use those silly pens that smeared, but I had this slide called "MIPS to the Moon" and it sort of took the idea of Moore's Law and sort of applied it to cache chips, to RAM capacity, to logic capacity and then sort of a modified Moore's Law having to do with disk capacity. And you could see when you put all that together that from a system point of view it literally was MIPS to the moon. There was no immediate cause to concern for there's going to be a big technological discontinuity. And I remember predicting that somewhere around 2000 we would have computers with 4 gigabytes of RAM in them. That was a little off but not by very much. And at that time, this would be 1988 when I was sort of on the evangelism side of the world, I saw that as hey, if the RAM can populate the address space then we have to have a 64-bit version of SPARC sometime in the late 1990s so that we're ready for the next generation. And at the time everybody thought I was crazy, but it turns out that was just a perfect extrapolation of Moore's Law that just worked.

Kapoor: Right, right. So a lot of other things happened. For example, you had the strategy of multi-sourcing--

Rosing: Right.

Kapoor: The technology when--

Rosing: The open SPARC.

Kapoor: Open SPARC.

Rosing: That was a strategy that in retrospect I wonder about a bit, but at the time the concept was very much we need to get the best vendors with the best technology available to build these chips and Sun would present a reasonably healthy market to these vendors and but we wanted it to be open. The traditional dialogue at that time with many of the vendors was, "This is great, but only me." Because then they would be the only vendor of the SPARC chip. Nobody really saw themselves as wanting to be part of a family. So those, some of those discussions didn't ever quite work out the way I think anybody expected. But it was the case that in many cases Sun ended up buying a big amount of chips from these various vendors and ultimately that culminated in a very long relationship with TI that was just very successful.

Kapoor: Right. So SPARC International was started as part of that community.

Rosing: Yeah. Trying to create a more open standard.

Kapoor: And the TI relationship happened because of the floating point, is that correct, that for the Viking you got the floating point from--?

Rosing: Right. TI was, what brought TI to the front for us was actually they had a really good fab. They had good geometry and they were competitive as measured in terms of memory technology at the time but because of their DSP work they sort of had a somewhat unique place where they were stronger in both floating point but sort of a mixed logic memory kind of way of doing chips. That's what I recall. And when we sat down with TI and really looked at the technology base and what we needed, you know, we needed more memory than other and we needed floating point and DSP kind of floating point. And so it just was a marriage made in heaven that worked beautifully.

Kapoor: Right. So I understand that the initial decision for the Viking was to go with bi-CMOS. [ph?] And that didn't turn out to be the best decision, if I'm--

Rosing: That's right. And it ended up being just more classic design, yeah.

Kapoor: And it actually delayed the project.

Rosing: It delayed somewhat, yeah.

Kapoor: Yeah. And then in terms of the server-client relationship of the architecture, you then had a relationship with Xerox PARC. There's a team from Xerox that--

Rosing: Right.

Kapoor: Collaborated with Sun on the multiprocessing.

Rosing: Right. Ron Ryder and his team, yeah. Yeah, and that was, that was borne out of it-- It's interesting how that sort of developed, if you will. As I recall there was a group of physicists at Caltech who were doing quantum lattice calculations with essentially little clusters of computers that were Sun Systems sort of networked together in a hypercube configuration. And I was introduced to this as, you know, normal customer things I get to talk to people. And I got to thinking about this and, you know, had not really entertained what I guess we would now-- We would now call that something like cluster computing or any of a large number of current buzz phrases but at the time it was just computers with over-connected, not on one wire but all connected together in whatever configuration you wanted. I thought that was a really neat idea and it struck me that that was a very cool way to think about how to scale computing beyond sort of a server. So I started talking about that and the general consensus at the time, I think even for Bill Joy and others, was, "No, no, no. More of a multiprocessing approach," a more classic multiprocessing approach was in order. So we started some of that but then in my touring around, I met the Xerox PARC people and Ron Ryder and they had done a system that was sort of a little bit of a mix of the different ideas and that really resonated with me. And they came in and they had software and they had an approach and it was working and so we just worked together and that became Sun's sort of muscular server business. Now I think in retrospect we might ask the question, should Sun have really made the leap to cluster computing, because they kind of missed that.

Kapoor: Right. So internally, how did things transition? You were working with Eric in managing the company's engineering team.

Rosing: Right.

Kapoor: And then because of the various projects and various directions, I understand that Sun was then split into planets.

Rosing: Right.

Kapoor: And so how did your role change during that time?

Rosing: Well, my role changed-- Okay, planets, I'm going to see if I can get this approximately correct. I'm going to say 1990, could be off by a year, but on the order of 1990, Sun was getting very successful and we had launched the SPARC Station I think or it was imminent that the SPARC Station was happening and Sun was just getting to be a serious company, I mean, multi-billions of dollars. You know, thousands of staff. It was just a really big company. And what happens in every company when you go through phases of your growth, at some point you turn a page in the management book and you do functional and you do divisional, you do matrix, whatever. You go through these processes of evolving how you deal with scale. And we decided to become planets, which sort of meant, like, independent corporations, independent presidents. So it was a sort of a strategy that I think sort of implied we were

going to try to clone the Sun culture into N-clones, put them in these petri dishes and let each one of them grown into another big Sun and then we'd deal with it next. So that was kind of the way it worked. And so, there was I think an extraordinary amount of sort of autonomy and independence on the part of the planets and that was an interesting change. At that point I became Sun Labs head.

Kapoor: I was going to ask that. So that's how it happened.

Rosing: Yeah.

Kapoor: Okay.

Rosing: Because I was interested at that time in sort of okay, what's the next technological place to go, the sort of same recurring question every five years in my life.

Kapoor: So essentially I would say you started Sun Labs or you took over the structure of Sun Lab?

Rosing: Yeah, I was the first-- I formed it, yeah. It was kind of forming on its own. There were various advanced development groups around, and I was always in touch with them, 'cause that's my first love, but then when the planet thing happened, I put my hand up and said "I would love to run Sun Research," and so that's how that happened.

Kapoor: So what were the key things that happened in the beginning or during the duration that you were managing it at Sun Labs?

Rosing: Well, Sun Labs, one of the big projects that we started on with Jim Mitchell was the Spring operating system, essentially the idea of an object-oriented kernel as sort of a successor way to think about how to organize the operating system. That was one major project. We had a number of other people, some people—Dave Ditzel so forth were working on SPARC ideas for different SPARC ideas. And then we had a smattering of people doing all kinds of other things, and then somewhere along the line, this notion of the successor to the NeWS window system. Those who remember it, NeWS was kind of a post-script-based, high-quality rendering, intensive window system concept, and Sun lost that battle to a different industry standard. But in sort of figuring out the details on that and where we went from that, James Gosling and Patrick Naughton came to me with this general concept of some sort of a new kind of computer language that would be much more tuned to the network--

Kapoor: Distributed.

Rosing: Distributed, more network oriented, more-- excuse, I said it incorrectly-- more Internet oriented, 'cause now we're talking 1992, and Mosaic and Netscape were doing their thing. The Internet was happening, and the web was happening, and so there was this notion that there was an opportunity here to do something unique for the web.

Kapoor: So initially I believe the project was called Oak.

Rosing: Oak was the first name that James gave it one day when we-- we moved offsite to sort of get out from under the antibodies <laughs> and let the thing prosper on its own, and that named Oak, and Green was another name that I think we used at the time.

Kapoor: So what was Eric doing at that time?

Rosing: Eric was running SunSoft if my memory is correct. He was running SunSoft. Ed Zander was running the computer company.

Kapoor: So in terms of Java, so I guess you really were involved with Java from very early stages.

Rosing: Yeah, more as its sponsor and wingman, making sure that it had good people but was left to grow.

Kapoor: So James Gosling was just recognized for his work with a Fellowship here at the Computer History Museum.

Rosing: Good for you guys.

Kapoor: Yes.

Rosing: Well-deserved on his part.

Kapoor: So tell us what happened after that at Sun as far as your career?

Rosing: Well, from my career point of view, as the Java thing matured, I was completely happy with that, but I didn't have a clear idea where I wanted to go next. It's sort of interesting in that, okay, you get into the research area. There's things you can do, but I simply didn't have something that was appealing to me at that time. And so I sort of thought about retiring and doing something completely new, and then the fork in the road would be do something technical or do something nontechnical, and I chose to take some time off and do astronomy for a while and sort of explore my other interests in life more. So I made that choice, but that didn't last long. Then I decided to get back and to do consumer software, and I went to a company called Caere in Los Gatos, which is doing optical character recognition. And so I managed their engineering for about three years and really learned a lot about sort of-- I learned a lot about optical character recognition, but then that's a esoteric-- it was interesting intellectually, but I really learned a lot about consumer software and about how you run a consumer software company. What is the pace of innovation? Which is basically very simple: You have to have your product ready to ship every September to fill the Christmas pipeline, or you don't make your year. So it's very different discipline, very interesting management thing to learn how to do, and by the way, that software had to work basically bug-free. So very different game, which I really enjoyed. And then, long and the short of it, in 2000 I got the job offer from Google, and--

Kapoor: So let's come back to that. So one other spinoff from Sun: FirstPerson. Was that a part of the Java work that you were doing?

Rosing: Yes, FirstPerson was effectively the planet. <laughs> We were an asteroid. Some people might've called us something else, but we were a little asteroid. We weren't a full-fledged planet.

Kapoor: So now let's go back to-- so after your work at Caere, were you once again thinking of retiring or looking for another opportunity before Google? Or did that--

Rosing: Okay, Ceare got bought out.

Kapoor: Okay.

Rosing: And I wasn't going to move back east. Caere was bought out and moved. So at that point, I was actually seriously ready to-- "Hey, I'm going to retire," whatever that meant. But what I thought it meant at the time would be I would move down to Santa Barbara where we already had a home, and I'd look for a job in one of the startup companies to be a CTO or a CEO or whatever, figure out something in Santa Barbara and have a good time and enjoy Santa Barbara. But I basically got approached for the job at Google, and then that all sorted out.

Kapoor: So you joined as VP of engineering.

Rosing: VP at engineering, that's right.

Kapoor: And so tell us about the environment at that-- this is the very beginning of--

Rosing: Yeah, this was November of 2000. My memory's correct, Google had had an investment round of 25 million. Prior to that, they had 19 million in the bank. There were about 300 employees I think in the company. All of engineering worked for me, meaning engineering, operations, the people who did the finance computing, who ran the finance computers. So it's sort of so to speak if you touch computers, you work for me. That was how it worked, so it was a very broad job. And I fell in love with the place, because as I used to describe it and I still do, it was like a computer science dorm, with food and people there coming and going all the time. It was just like sort of a Stanford dorm full of computer scientists with some people from other places and a big contingent from UC Santa Barbara. It was a very unique place, very unique culture. And I think my proposition that I proposed that I think got me the job was sort of saying "I think what I can do more than anything else is, if this company grows, I can help figure out how to preserve this culture and not let the big company culture take over too quickly and destroy this, because what you have here is neat."

Kapoor: So I assume Eric and Larry and Sergey were there at that--

Rosing: Larry and Sergey were there. Eric was not. Eric was one of my references, and when Larry and Sergey checked my references, they met Eric, and that then led in I would say less than two months, maybe three months, to Eric joining as CEO.

Kapoor: Wonderful. So tell us about your engineering experience, and how much did Sergey and Larry get involved in the management?

Rosing: I would say they were all over the place all the time, and that's not a criticism. They were kind of at one level involved in everything and interested in everything and particularly weighing in on product decisions and general strategic-- like what products are we doing? What projects aren't we doing? So they were involved with a lot of it, and of course they also had a fair number of responsibilities as executives to kind of the rest of the company and work with the board and all that, 'cause even in this time it was clear that, okay, Google was destined to be a public company, and therefore there was a great deal of work to do. So I didn't worry about that task. I worried about getting the engineering to go.

Kapoor: So I read an interview with you where you explain the environment, the development environment and the tools and how it was so different from when you came from Sun, what the environment was like at Sun. Maybe you can describe a little bit.

Kapoor: Well, I think the main thing that was very different about Google was, in many ways, there were no thought police. You could use whatever tool you wanted in whatever way you wanted to get a result and get a result fast, because the emphasis at that time was on fast and on improving the company. And it was a very interesting time, and I'm going to say-- well, this extended through 9/11. So 9/11 happened essentially about three-quarters of a year into my time there, a little more than three-quarters, not quite a one-year thing. So it's putting it in emphasis. We were still growing. We were a small fraction of the market share and search, but we were talking about deals with Yahoo, deals with AOL. These were very big deals that drove a lot of our demand, and so being able to execute was important. Now, one of the problems that was presented when I got there, which would've been November of 2000, was a very simple thing. At that time approximately if my memory's correct, about every four months, three months or so, Google would crawl a web, get all that data, and reindex it and then release a new index, which would result in different search rankings, because the web had changed in the meantime, and these things were more authoritative, and these were less. So Google changed discreetly on an approximate time frame if you will of quarterly. Well, when I got there, the index was late. The database had grown so big it couldn't get processed, and some of the procedures had broken down. The index was late, and then when it does come out, the search engine optimizers, which are companies that were there to sell to people how do you get to be first on Google would all go in a tizzy, and it would all be very complicated. So I asked the question one day "Well, why can't we do a real-time index?" I'm not sure I said those five words in that order that succinctly, but it became the question. Why can't we crawl an index continuously? And the first-- answer was that's too big a problem to solve or whatever, but after a while, people like Jeff Dean and Sanjay and some of the other folks started to say "Well, that is doable, but we need to do X and Y and Z to do it." And so we encouraged that. It was initially ideas that you got to say yes to. The most destructive word in business is "no." The most powerful word in business is "yes." Now, you got to be careful <laughs> not to dissipate yourself with too many yeses, but we nurtured that

idea, and then it really took hold. And out of that came the Google file system and map reduce, and at that same time we had sort of figured out how to do these incredibly cheap computers by buying commodity motherboards and Velcro and put 10,000 watts' worth of computing in a single rack, which is two or three times what other people would do, and just completely, totally change the economics of search and its utility, because once we could crawl the web uniformly, everything calmed down, and then we could then spend a great deal more energy on the quality of the search result and less on the I'll say gathering of the data to build that search result. So that was a fundamental change, and then that empowered the new ad system, which was the second thing that came out of that progression that essentially made Google what it is today.

Kapoor: So I think there was a good solid foundation and base that you established, but now I wonder whether that carried through to the size of the web and the data that they are encountering now. I wonder how things are.

Rosing: I don't have an opinion actually, an informed opinion about what Google's doing, but as a user of the web, it's in my DNA that I go to Crawl more; I go to Gmail. As a user of the web, I do find it's getting progressively somewhat more commercial than I like. I'll go to Google Scholar sort of knowing, hey, I want the scholarly literature, that I find that a lot of the stuff I get, if I don't sort of carefully frame my search, is ultimately sort of commercial content or it's shallow. And I don't think that's a reflection of Google. I think it's a reflection of there's an enormous amount of stuff that's being poured into the web that's not new; it's regurgitation of something else. And of course the commercial aspect of-- the web is for all practical purposes everyone's storefront now. What we need is a special search filter that says "I don't want any commercial stuff," or "Down-rank it to page three."

Kapoor: No, I think in that interview that I read, you had described the engineering environment and how the teams would work on a problem together and then go away when that problem is solved. It's a very different environment then.

Rosing: Yeah, we found ourselves in early '81 a little bit stuck, not knowing how to get stuff done. And so we sort of reoriented things to be super flat and to just say, okay, here are-- I can remember it schematically. My numbers are incorrect, but we sort of sat down, Larry and Sergey and myself and Eric and some of the seniors engineers, and we'd chat through all these ideas, and we'd come up with a list of 30 things. These are what we want to have happen; this is what's important. It was some number like that, and the arguments got quite heated over the last five at the bottom of that list. But we might have 70 or 80 engineers. Well, 30 things and 70 engineers is not going to work, okay? Because you won't have enough people on any project to accomplish anything. So we'd sort of float this list, let people join, and you'd end up typically having five or six maybe. I had kind of a rule at the time-- it's not a real rule; it's just sort of a guideline-- that the eighth person is a negative, because your communication burden starts to exceed your creativity, so try to keep a project to five. And we kinda had the idea we're bringing in lots of new people, so new people would be put in a project, so we'd be always training. And one of the things that used to happen is sometimes we'd bring new people into a project, and on their first day they discover they're the manager, because none of the other people wanted to manage, so the poor incoming person from school got to be the manager. Well, turns out that's not necessarily a bad strategy, because

that's often a very good way to get an engineer to integrate and learn with the team and work with the team, and then they're able to go off and do that again, not to manage but to be workers on other teams. So what would happen is stuff would start getting done, and then there'd be some emphasis, and then some people would switch around, and so it was very much a broad attack against a list, and a lot of expertise would switch around. That's how I saw it happening, and that really helped us. And we were also bringing in lots of people, so in these years, each of these years, Google more than doubled its engineering staff. So while this is happening, 80 people, then 160 people, then 300 and-- it just was ridiculous towards the end of my time there that just the sheer scale of people that were in engineering, and that was very small by today's standards.

Kapoor: So when you were there in the beginning of course, their search engine was still the primary focus, but were they expanding to other projects? That's why they were hiring more people?

Rosing: The search engine was the key thing at the time. A second indirect key thing was all of the activities necessary to run the search engine, in other words, operations, the data centers. Google had the presence of mind to buy a lot of dark fiber at the end of the '90s when dark fiber was literally purchasable for fractions of pennies on the dollar, so it had an extraordinary capacity to build a global network rapidly and have the bandwidth. So the problem was getting the AC power, the megawatts to fuel it, but the bandwidth was not the problem. So it just grew really well, and what can you say? It's a unique trip.

Kapoor: And then of course since then Google has expanded to different areas.

Rosing: Yeah, well, I think that some of the AI work was in its infancy with just a small number of people. Back then it was sort of called machine learning. I don't think anyone quite put AI in the acronym bin, but it was machine learning. There was really some fascinating work done there in search. So there were nucleations of a lot of stuff. The ad system was very sort of, in my judgment, traditional when I got to Google. It was basically sort of banner ads and very kind of traditional thinking. And it was the new ad system where we started to really target results to people's interests coupled with the bidding system that created competition that resulted in quality ads being available to be presented to people and then having a conversion rate, because you're presented with a quality ad, so you actually use it in your life. That was to me an extraordinary change. And then the third thing that I think of as hallmarks of what I saw us develop at Google was Gmail, was sort of a very early Ajax-based prototype for what I guess we would now call HTML3 or whatever the buzzword is or these days, modern HTML applications. But Gmail was one of the first of those done very early on and then launched almost immediately to an enormous scale. Now, I think that's one of the things that I'm most impressed with about what Google was able to do in 2004, 2005 was come up with an application like that and have it get to scale with almost no pain. That's rather extraordinary.

Kapoor: Now of course Dave Patterson and John Hennessy are both at Google, and so it's very interesting how things happen. And so tell us about your transition from Google, out of Google and then--

Rosing: Well, before I left, I had in mind I call it a refiremant project, not a retirement project. Originally in the '80s and '90s, the idea was to build six 24-inch telescopes and put them at six places, three in the south and three in the north, so that you could sort of cover the whole sky and look continuously to do a variable star photometry, just measuring stars that change. And that was a sort of obvious thing to think about. I wouldn't claim to have invented it, but certainly I was early on the idea just because I was working on Macintosh and Lisa and I sort of understood the Internet and integrated circuits and CCD chips, which is basically charge-coupled devices, which modern world now knows as electronic cameras. Well, that made it possible to do remote astronomy. Prior to CCD chips, you often did it with photographic film, which is sort of messy and difficult and hard to do that kind of stuff remotely. So I wanted to do this project, so when I left Google, I founded Las Cumbres Observatory in a serious way, and we've essentially built a series of observatories around the world at seven sites. We are on Maui as well. We had an opportunity to get a facility on Maui, so there's really seven sites. And we do scientific research. We have about 28 to 30 people in the organization, of which approximately half are staff scientists, post docs and grad students doing full-scale astronomical research. And then we get some support from the National Science Foundation, and we provide time for hundreds of astronomers around the world and many astronomers in the United States to take advantage of time-domain astronomy, which is something that we've sort of somewhat claimed to have invented. We will somewhat claim to somewhat have invented it. That'll be the most definitive way to say that.

Kapoor: So it's amazing that your love for astronomy started very early in childhood, and it's concluding with a very satisfying career in the retirement--

Rosing: Well, I feel very grateful. I had a I guess better-than-average track record picking good companies to go to. But somebody offered me some advice very early on. He said "If you're going to pick something to do, make sure it's important, and then do your best to do an outstanding job," and it's sort of something that just stuck with me. And so there's been a sort of loose-- every five to eight years there's some technological discontinuity that sort of showed up in the computing business that I was willing to say there's some leverage here, not necessarily for me personally financially, but in terms of the challenge and the work on a new thing, and that's been really gratifying.

Kapoor: So you're going to continue your involvement with the Las Cumbres I assume.

Rosing: Absolutely. I can't imagine not being involved. My current project is to put another set of six telescopes or six sites, 12 to 18 telescopes, something in that range. It will be smaller, and the emphasis now will be on education uses. It's not that they can't be used scientifically. It can be, but this is to make that capability available for educational purposes and then have the people involved-- I am assuming kids, but I think kids of all ages is the target, somebody who wants to learn this stuff, is curious that they can cooperate with scientists and do valid science as part of their learning about this stuff. So that's my current project, and the interesting technological bent on this is, of all things, Python. Python has become so dominant in the scientific community and so rich in its content of libraries that the leverage that's possible now to do some projects is really extraordinary. And curiously enough you don't think of Python as using it to program. You're actually using it to script libraries together. Now, that's sort of what it was designed to do up front, so I'm playing with the words a little bit. But there's these phenomenal tools that

allow two or three people to do projects that can operate over the whole world and reach hundreds to thousands of people just by simply doing them. And that's just extraordinary, because if I can do it, then you can imagine that a real software engineer or a scientist who has a very deep understanding of some area can get their technology and their work out to where other people can participate.

Kapoor: We recently honored the Python--

Rosing: Guido, yes.

Kapoor: Yes.

Rosing: I've never met him. I need to do that.

Kapoor: Yes, he was here for his Fellowship-- he was made a Fellow. So another question we normally ask people that we talk to is what would be the advice you would give to somebody starting a career? Of course you have set a very good example in this domain of astronomy, but any other other recommendations you would have for somebody starting a career?

Rosing: Well, I guess the first recommendation is to kind of restate a little more accurately what I related a minute ago, because I do get asked this regularly by young people starting out. And my first answer to that question is figure out what's important and you're interested in, and gauge that carefully. You really want to be enthusiastic and passionate ideally. And then jump in and do it, but keep learning, 'cause at least if you're in a technology space, the five years of head down solving a problem could be five years of you're not learning, and you may solve the problem and your head comes up, and the world has really changed. So I've always found it to be critical, at least in the computing space, to constantly look at the ecosystem around you and where it's going. And probably fair to say more than one person would say about me that sometimes I'm just a little too scattered in my perspective, because I'm trying to work this, trying to get this done but thinking about this and whatever. But that's me, but I think that's worked for me, because I try to look at the opportunities. I don't invent them. I can't say I've ever invented anything really. But I see the opening, and I see how it can help, and if you're in business, which most people are, the business of business is business. You have to make good decisions that can result in economic success, and that applies whether you're a startup person doing a startup or you're an engineer on a project. All these things now are the integral of a whole bunch of decisions, and so I think the myth is that somehow the leader does this magic thing and it all just is good and it all flows from the leader. I don't think it does. I've seen it too often. And managers have to figure out how they create that environment where the creativity flows up into the organization. It's very different than "I want you to do this. You need to do that." It's a different model. So I say just commit to do outstanding work on an important problem.

Kapoor: So this is a nonprofit. Las Cumbres is a non-profit organization, and are there any other philanthropic activities that you're engaged in besides this?

Rosing: Well, from the philanthropy point of view, Las Cumbres and this education project I'm working on now are my principle interests. My wife has a foundation called Linked Foundation, which deals primarily-- started with women's health in Latin America, issues around that. It's morphed a little bit into women's health in Latin America and women's health in America, 'cause it turns out there's plenty of need here that's still not being addressed. We're also very active in Santa Barbara in the arts and lectures program, which is a university program that brings artists and lecturers from all over the world into Santa Barbara, and it's created an extraordinary experience. You still go to London and New York for plays, but the quality of culture that's available in Santa Barbara is extraordinary. It's probably about the same as the total amount you could get if you lived in the Bay Area and you wanted to do a lot of driving, but in Santa Barbara you don't have to drive very far. I've now done my Chamber of Commerce <laughs> ad for Santa Barbara. And then we also are very active with a group called Direct Relief, which is a Santa Barbara organization that now is shipping over a billion dollars a year in retail value in medicines and medical supplies into the third world, so just an operation that's grown from incredibly little to extraordinarily impactful. So we have plenty to do there.

Kapoor: Any other thoughts before we terminate the interview?

Rosing: Well, it's interesting. The last time I was here in the museum and enjoyed the exhibits was approximately 2003. And it was really nice to go through the exhibits today, because there were many of them-- some of the things I saw reinforced what was in my mind's eye when I was in high school, when I was learning about computers and first got hooked. And I found it interesting to retrace the development with sort of my 72-year-old perspective on technology. And it's not only just sort of looking at the artifacts, but it's really seeing the extraordinary change in the world that has been driven by the combination of sort of digital technology-- we'll call it integrated circuits; you call it anything you want, Moore's Law, however you want to integrate it-- coupled with software. And software's interesting, because software is ultimately very much typically created by a few people, the core concepts. I'm not sure I am qualified to do a taxonomy and quote them, but there's maybe a few hundred critical things that have made modern software what it is today, and there's various people who've done it, and these ideas have moved from one person to another, and it's just built this incredible edifice of technology that's completely changed the world. I don't know what to say except I'm glad I had a chance to help make it happen, not as a creator but sort of-- I guess I think of myself as I'm just a great sort of a massage therapist that sort of gets the kinks out so that things keep moving. It's just been fun.

Kapoor: You started off in the hardware domain, but you're equally comfortable in the software. In fact, you've used the software in all kinds of domains to make the combination successful.

Rosing: Yeah. But my contribution if anything is leveraging people. A ton of people did all this work, and it could've been incoherent, and it could've been inefficient, but in many cases it was coherent and efficient because good leadership, and that's what I think is the fun thing. You asked me about advice. In the modern world, I don't think it's possible to contribute alone anymore. We're all building on so much what other people have done that good engineers have to be social animals. They don't have to be hyper social, but they have to be social.

Kapoor: I think the whole social media stands on the shoulders of the platforms that were created.

Rosing: Yep. There's obviously some burdens with some of the modern problems of modern availability of information. It's now often non-curated. So you have to lead a data-driven life, and you have to be really careful to make sure your sources of data are cast broadly.

Kapoor: And then of course the issue of privacy and security has been very topical.

Rosing: Yeah, I marvel at some of the problems that exist today, and they are a direct consequence-- I almost said "fault," and I don't want to be judgmental-- but it's a direct consequence of the ultimate democratization that the Internet has provided. Anybody anywhere can be in touch with virtually anybody anywhere for any kind of transaction no matter what it could be, legal or illegal or in good taste or in bad taste or whatever. I think that's a net positive for humanity, but it's like playing with gasoline. You just don't quite know. It's going to take some time to figure out, and we've got to be careful about it. But it's a challenging problem, and unfortunately I don't think cryptography will solve it. Cryptography is not going to solve it. You can't go secret again.

Kapoor: Yeah, they're always finding ways to get around that.

Rosing: Yeah. But it's an absolutely fascinating change to see. And I think one of the great challenges of our immediate time the last five to six years is this extraordinary flood of information. And we're all being lobbied. We have these lobbyists in Washington who seem to be in control of things there more so than our representative are. Those are lobbyists. They're paid by corporations mostly, and I've heard rumors and speculations for every dollar spent there's 200 dollars in profit made, so if that's the investment ratio, you can hardly blame any corporation for lobbying. They're doing what they got to do for their shareholders. But we are being lobbied, and that is an extraordinary problem, and everybody has to, I think, rethink their whole approach.

Kapoor: So thank you very much, Wayne.

Rosing: You're welcome.

Kapoor: It was a pleasure, and with that we'll say thank you.

Rosing: Good, okay.

END OF THE INTERVIEW

