Oral History of Richard Lyon

Interviewed by:
Günter Steinbach

Recorded October 10, 2016
Mountain View, CA

CHM Reference number: X7975.2017

Steinbach: Today is the 10th of October and we're interviewing Richard Lyon for his contributions to computing. Especially, my interest is the optical mouse. Thank you very much for agreeing to do this.

Lyon: Thanks for inviting me, Gunter.

Steinbach: I'll start with your background. Where did you grow up? What's your family background?

Lyon: My wife says I grew up in a Norman Rockwell painting. I was one of nine kids in El Paso, Texas. In 1957, we moved to a new house and both my parents are still living there nearly 60 years later, and I've got brothers and sisters all over the place. And we've been very lucky; we haven't lost anybody, no siblings, no parents dead yet, and we're a close family. We got lots of young'uns and young'uns' young'uns and so on by now so it's a good bunch. We were middle class, went to the local public high school and learned a few things, and then we all scattered to college.

Steinbach: What did you do for fun as a kid? Did you already start engineering or—

Lyon: Oh, I did. I was an electronics tinkerer. I used to make for example strobe lights for the local garage bands; they'd play their guitars and I'd hook up the strobe light and people would dance and things like that, but that was just one of the things. I was big into photography, did a lot of bicycling. We would ride bikes around out in the desert. We lived kind of on the edge of the desert where we could just go out the front door and hike for miles, climb mountains, ride bikes in and out of gulleys and all that kind of stuff.

Steinbach: That's very different for me but my wife tells me the same thing; they had tumbleweeds come in the front door.

Lyon: Right, plenty of tumbleweeds.

Steinbach: I guess then it was kind of natural that you got into electrical engineering in college.

Lyon: So my dad was—he spent his entire career as an engineer for the El Paso Electric Company—so he was always bringing home things, and he brought home manuals for how to program in FORTRAN one time, and that was kind of interesting because I got to read about FORTRAN before I had a computer to do anything with, and—

Steinbach: That was very early then.

Lyon: Yeah, early ’60s I think.

Steinbach: Early ’60s. Wow.

Lyon: Our high school got a computer of some sort in the late ’60s; I think it was a Olivetti Programma 101 and they had some other kind of little computer that you could tinker with and learn a few things.

Steinbach: My high school had an HP with a one-line display. How did you choose Caltech?

Lyon: Well, that was pretty easy. I had one older brother; he was the tormentor of my life and he went to Princeton, so I just had to look at the map and see how far I could go—

Steinbach: The other end of the country.

Lyon: Exactly; and when I was in grade school a teacher had mentioned MIT one time as —sort of put the image of—in my mind of Massachusetts Institute of Technology as a great school. So when I saw there was a California Institute of Technology I thought that must be great too, and that was about all I knew about it to tell the truth, and it was great.

Steinbach: You liked it, huh?

Lyon: Yeah—

Steinbach: But you didn’t stay there for graduate school. You went—

Lyon: No. I left after four years as most Techers do. They generally encourage people to do grad school somewhere else, get a little diversity of outlook that way, so I went to Stanford with the intention of doing the full grad-school thing; but Stanford—I don’t know—I got antsy and I decided I wanted to get out so after a master’s degree. I got some job offers and went to work in Silicon Valley instead of staying in school.

Steinbach: What was your work right out of school?

Lyon: First job was at Stanford Telecommunications, Incorporated, which was a little startup at the end of San Antonio Road not far from where we are now. They had a couple guys that had spun out of Ford Aerospace and Electronics where they had worked on developing signal sets for navigation satellites and things like that, so they were in the GPS business before there was any GPS. So I got to work on that. I got to work on some space-shuttle communication systems and all kinds of interesting kind of—mostly one-off kind of communication and navigation projects.

Steinbach: That sounds interesting.

Lyon: It was fun.

Steinbach: And then you went to Xerox?

Lyon: Yeah. I had a fortuitous meeting in—I guess like late ’76. I visited Caltech; this was maybe two years or so after I had graduated—two and a half years after I graduated and I ran into Carver Mead, who was one of my professor friends there, and he had a couple visitors that day, which was Ivan Sutherland and Bert Sutherland. He was in the process of trying to get Ivan to join him in starting a computer-science department at Caltech and Ivan’s brother, Bert, was working at Xerox PARC and they were looking to get some collaborations going between Caltech and PARC.

Steinbach: And there you were?

Lyon: And Carver introduced me to Bert, and he said, “Well, why don’t you come over and interview for a job with us?” or something, and it started with that and fortuitous that I was—I happened to be there that day visiting.

Steinbach: Yeah. Stuff happens.

Lyon: Yeah.

Steinbach: Was the optical mouse your first thing you worked on there?

Lyon: No, it was more like the last thing I worked on there so—

Steinbach: When did you start that?

Lyon: I started in April '77 and left in August '81. The optical mouse was sort of late 1980 so it was within my last year there. When I came in they had a—they were starting this project under Lynn Conway to do LSI System design, what they later called VLSI System design because it was about time to add that 'V' on there. And they had some other projects; like there was a project under George White doing speech recognition and part of the concept is that I would come in and help build the custom silicon to run the speech-processing stuff that they needed. So I was kind of the bridge between the LSI systems and the speech processing and I really didn't know that much about either, but I knew a little bit about each and enough to pass the interviews and get hired, so I worked on that. I did some custom chips for digital filtering and so forth, but within my first year there George White, the guy in charge of speech, got hired away to go start a lab at ITT someplace, and that left me kind of holding the bag on the speech project. One other guy who had been working on it moved away as well so I was now the speech-recognition team. They had just brought me up to speed and then left me to head off in my own directions so that was kind of fun. And I took a Stanford course; they had the honors co-op program where you could continue to take graduate courses while employed. I took a course in biological information processing and did a project on hearing as part of that, and wrote my first kind of outline of how to approach speech recognition through a signal processing model of hearing in 1978. And that document I wrote for that course is still sort of the outline of my career in hearing research that got me started down that path.

Steinbach: That's interesting. I thought it was Carver Mead who got you onto that.

Lyon: Well, it's more like the other way around; I got Carver into the hearing stuff. He was doing a lot of the biomimetic vision work and I was doing the hearing side and he talked me into coming and spending time at Caltech with him and we did some projects in that area together.

Steinbach: What was PARC like at that time?

Lyon: Well, it was an amazing place; it's hard to describe. There's so many amazing and famous people from those days that had just sort of invented interactive personal computing out of whole cloth. I think probably a bunch of them are already fellows of the Computer History Museum or something on that order so I probably don't need to review all that, but I didn't know quite how momentous it all was; I had kind of a sheltered existence there. And what I realized many years later is that those of us who were at PARC using Altos as our personal computers really never took the microcomputer revolution very seriously. Those were toys that you could get at the Byte Shop and you could poke around with but you couldn't do real work with them at that time, and we kind of went down our own path that kind of missed what was happening in the larger world of computers because we were so privileged; we were so ahead of our time that at least I kind of missed what was going on. So people left and joined companies that were in the outside computer business like Zilog. When Exxon bought Zilog—they made the Z80 microprocessor—and people went there and Exxon thought they were going to be a big information systems company. That didn't quite go, but just the fact that people were attracted away by these companies trying to make things out of little eight-bit microprocessors seemed very strange to me at the

time. It turned out that was kind of the way things went. It was the evolution of the silicon that really drove the industry more than the—I think more than the evolution of the software that we had at PARC. It was groundbreaking; it was too-far-ahead-of-its-time software is what it was. It took people like Steve Jobs to come in and kind of notice what was good about it and put together a team that could re-create a different version of it from the ground up running on little microprocessors that really made the revolution take off I think. So I kind of was out of it in that sense. We were too privileged with our Altos and all those other computers they developed—

Steinbach: You really used them.

Lyon: Oh, yeah. It was our day-to-day workhorse machines. From the day I started, I had a personal computer. I got to use a visual editor, Bravo, what you see is what you get, editor. We had the laser printers; we had the networks; we had the file servers. That stuff was all in place before I started in '77. I even had e-mail. I found at a used book shop one time, a copy of the 1978 ARPANET directory. It listed everyone with an ARPANET e-mail address and I was in there so I bought a copy.

Steinbach: Let's get to the optical mouse. How did that appear on the scene?

Lyon: So as part of our VLSI Systems project we had collaborations with people in various universities. We had Carver from Caltech, we had Carlo Sequin from Berkeley among others, and I spent time with Carlo. He had done a book not long before that on charge transfer devices, what we call CCDs usually, and including CCD imagers in there, so he understood very deeply about the light sensitivity of silicon and I didn't, but in a little bit of discussion with Carlo I, he kind of clarified for me how it works, how is it that the stuff we draw to make chip layouts makes photodiodes, where does the—what does the light do when it comes in, in terms of charges and so on, and what could you make with that. And once he explained it to me I said, "Oh, I could make some image sensors or something"; I thought that sounds kind of fun. So that was one thread of stuff that was going on. I could make chips with lots of photodiodes; I knew that because everywhere you put a diffusion is a photodiode whether you want it to be or not so that's something I could work with. The other thing that was going on is we were using mice all the time and they would get sticky and they would fail and they would have to be refurbished, and several people at Xerox had already filed invention proposals on how to make an optical mouse. Dan Ingalls had done that; Chuck Thacker had done that. There might have been others; maybe there were some co-inventors on these things. And in each case they never built it as far as I could find out and they never filed for a patent on it, and the reason was you could look at the idea and say, "Well, yeah, nice try but that's not really going to work." But that inspired me to think about it. So thinking about it some more I came up with a scheme that would really work, and I kind of chatted with some people about it, and it was partly again inspired by that class I had taken on biological information processing where I had learned about things like lateral inhibition in visual systems. I wrote a retrospective on the optical mouse recently for a book on advances in computer vision; they invited me to do that. And one of the things I emphasize in there is how my view of what lateral inhibition was about really formed the basis for how the optical mouse design was

crafted and how it worked. And it has to do with the different light sensors inhibiting each other in a way that leads to a stable binarized pattern; so basically it would see white dots in a dark field and even if there were not any white dots in a dark field it would hallucinate white dots in a dark field. It would say, "I saw light here, here, and here," and it would inhibit the ones around those. And it was straightforward to build it in such a way that once it got to a pattern that was stable it could tell it got there, and it would take that as a snapshot and then do it again and see if the dots moved, so it could track the motion of white dots in a dark field. And so I built it, I laid it out, and it took—I don't know—a month or two with our IC editor. I don't know if you've seen it, but in the museum main exhibit here there is a—there's a video of Doug Fairbairn demonstrating the ICARUS IC editor. That's what we used—

Steinbach: I did not see that.

Lyon: That's what we used for laying out integrated circuits and there's a nice—yeah, a nice video demo of it in there, fairly primitive. It's sort of an artwork system that lets you draw multiple layers of stuff, but it ran on the Alto; and Jim Rowson, one of Ivan Sutherland's students at Caltech, wrote most of the code when he was an intern with Doug Fairbairn, who was also in Lynn Conway's group there at PARC. And so I drew the layout. I sort of drew the logic first, kind of sketched it out, decided what I wanted to do, and I built this thing and I got some interest in it. And Chuck Thacker in particular was interested enough in it that he volunteered to help me go over the layout carefully to make sure that when I got it built it wouldn't have mistakes; and he in fact found two logic errors where I had interchanged a couple things and without his help it wouldn't have worked the first time out but with his help it did. So we fabbed it through one of our what we called multi-project chip runs; it was probably on MPC 80, that's multi-project chip 1980. So you can read about the whole history of this multi-project chip exercise in Lynn Conway's reminiscences.

Steinbach: It was a PARC thing?

Lyon: Yeah, it was a PARC thing.

Steinbach: Not like MOSIS; I think MOSIS came later.

Lyon: So we started up the MOSIS effort around that time, so MOSIS evolved out of this PARC effort largely—

Steinbach: I didn't know that.

Lyon: —and the—some of these MPCs were multi-university efforts where the people teaching courses from the Mead and Conway book would bring projects from their students together to PARC and we would put them together and fabricate them, and other runs were internal PARC runs and I don't actually

remember exactly how we did—which run the optical mouse was on but something in late 1980 because around December we got chips back and—

Steinbach: This was a fab at PARC or—

Lyon: No, we didn't fab them; we sent them out various places. Some of the early runs were fabricated for us by Hewlett-Packard and there were other companies we worked with various times but we didn't do any fab. So I got the mouse chip back and that's when I realized I kind of forgot certain details like this chip has no inputs—how am I going to test it? It was a chip with outputs only. It provided the signaling that the Alto expected to see from a mouse including the debounced switch buttons—button switches.

Steinbach: The whole system on one chip.

Lyon: Yeah. So it just had these outputs, and the only input was light, so I set one up under a little lamp and I put a pattern on some paper and moved it over and I watched the Alto, and the mouse—the cursor—moved, so "Okay, it works." So that was testing.

Steinbach: Wow.

Lyon: So yeah, it was kind of crazy that way. We did make a version later that had test inputs so you could run it through a regular tester and simulate light hitting the diodes and things like that.

Steinbach: From what you say, it was basically just a one-person project, just you. Right?

Lyon: So the first one was just me and then the second one got bigger because people could see the potential at that point, so in particular Martin Haeberli and Bob Garner and I think maybe another person or two helped develop the version that was kind of re-layed-out to have better optical magnification properties, and testability, and a smaller chip, and things like that so—

Steinbach: That's still exactly the same kind of logic—

Lyon: The same logic except for the testing features that were added, yeah. The layout was quite different but the logic was the same.

Steinbach: What you tried to do better was just the mechanical mice coming up?

Lyon: Yeah, largely. I wasn't really thinking so much about the cost of mechanical mice but that—they were quite costly at the time. I think Xerox paid 500 bucks each for them back in the '70s when they were built by Jack Hawley of the Mouse House—

Steinbach: Okay. And—

Lyon: Yeah, just to make them more reliable. So the optical mouse used a pad that's basically—

Steinbach: That was a downside I would think—

Lyon: Yeah, so dark paper with light dots. It was a downside but anyone that actually used it compared to the mechanical mice it was okay; it was still a big one. And when Xerox shipped office computers, and Smalltalk computers, and Lisp computers and things like that, they used the optical mouse because it was cheaper and more reliable, and more usable than the mechanical ones.

Steinbach: You just handed out the mouse pad with it.

Lyon: Yeah. They came in pads that you—if they got dirty you could tear it off and there was another one—

Steinbach: That's a piece of paper.

Lyon: —or you could tape it on your desk or whatever and each page had the reorder number on it so you can order more if you need them. They were paper, yeah. You could reproduce it on a Xerox machine if you wanted to.

Steinbach: Of course. Do you remember how large the mouse chip was?

Lyon: No, but I could look.

Steinbach: Okay.

Lyon: I brought a box of goodies here and I don't know actually if I have the original mouse chip in here or not. I have some of the mice that were made with the redesigned chip and there's a standard-size 14-pin dual in-line package here except it's molded with clear plastic and the chip upside down so it can see through the board where there's a lens that images the pad. I had some original—

Steinbach: It's only three or four millimeters. Right?

Lyon: Yeah, it's small. The original one was a little bit bigger but not a lot bigger and I don't seem to have it here.

Steinbach: What geometry size was the process at that time?

Lyon: It was probably five-micron NMOS; maybe by the time they went to production, four-micron or three-and-a-half micron NMOS; no CMOS yet in my life.

Steinbach: Not in your life.

Lyon: You could buy CMOS parts at that time like the RCA COS/MOS series medium-scale integrated circuits, but they were known for being slow. The development of fast CMOS for processors and so on was kind of ongoing but the processes were complex and expensive and the design was not yet simplified to the point where we could do custom CMOS chips. That's something that Carver and I and MOSIS and others worked together on: simplifying the CMOS processes and design rules to where the university systems could cut over and start doing work with CMOS.

Steinbach: The Magic editor also came out of that kind of thing. Right?

Lyon: Yeah. Magic was the editor from Berkeley; and Electric by Steve Rubin was the editor that came out of Schlumberger, which is where I worked after PARC so—

Steinbach: That was your next—

Lyon: So I used Electric when most of the world was using Magic. Yeah, in '81 I got married and moved to Schlumberger at about the same time, in the summer of '81, where I headed up a speech recognition research group. So this was a transition sort of nominally away from VLSI and more into speech and hearing, although I actually ended up doing a lot of VLSI design at Schlumberger as well, including building programmable SIMD machines for the speech and hearing processes.

Steinbach: I thought Schlumberger was more an instrument—

Lyon: So Schlumberger is an oil-field-services company; but in '79, I think, they had bought Fairchild Semiconductor, and around that time, maybe '80, they started the artificial intelligence lab at Fairchild

Research with four guys from SRI; it was—Peter Hart, Dick Duda, Harry Barrow, and Marty Tenenbaum came over and started this lab and started hiring people in to lead groups in different areas. So they got me to head up speech recognition and Andy Witkin to head vision, Ron Brachman for knowledge representation and some other people that aren't coming to mind at this moment, <laughs> and we had a great lab there for a number of years, and Schlumberger kind of went on a buying binge. They bought a bunch of information technology companies, thinking this was a good direction to invest at the time, but they didn't make great choices, and they didn't really know how to manage it or integrate the different companies together very well. And so while we had a great research lab it was sort of surrounded by businesses in decline for a number of years, and by '88 the whole thing fell apart. We did a lot of great research publishing in the meantime, but yeah, pretty much everyone left around '88—ended up with 13 of us going to Apple—and Schlumberger tried to take what was left and move it to the computer science research lab in Austin, Texas, and when they finally did they ended up only taking six people. I heard everyone else had left. So it was too bad. It was a great lab, across the street from Xerox PARC, and it did a lot of great things in artificial intelligence and vision and speech and VLSI and—yeah—

Steinbach: But they—

Lyon: It didn't last.

Steinbach: —couldn't make it work.

Lyon: It didn't last. I had some of the same problems that Xerox Research had, which is too disconnected from the businesses; and when we went to Apple, that was a little better, but it was still sometimes problematic that the advanced technology group was sort of set off from the main business groups and—

Steinbach: Even at Apple.

Lyon: Even at Apple. Even though we were not in different cities or something it was sometimes difficult to get technology out of research and into products; not as bad as at Xerox or at Schlumberger.

Steinbach: You still worked on speech recognition at Apple?

Lyon: Sort of. I worked on hearing. I had a group we called Perception Systems. We worked on other perception problems like handwriting recognition and shortly after I went there we hired Kai-Fu Lee to run a speech recognition group so I didn't work in that team. At some point, he got promoted and became my boss, but I wasn't part of speech recognition per se. I worked kind of closely with some of the speech guys but yeah, it was more—the big thing we did there was the handwriting recognition for the Newton.

Steinbach: I thought I had read somewhere that they got that from Russia or something.

Lyon: So yeah, the original Newton had a handwriting recognition system from ParaGraph International, which was a Russian company; and then my team in ATG did the recognizer used in the version 2.0 Newton. We call it the one that worked as opposed to the ones that Doonesbury—

Steinbach: I remember it not working. I remember putting in the password with handwriting and then never could get it right again.

Lyon: Yeah. Yeah, there was a good story that John Markoff did in The New York Times about us when we upgraded the Newton with a handwriting system that really worked, and yeah, so that was a fun project, and it used neural networks and other things and we got some good—

Steinbach: Simulated in digital logic on a computer?

Lyon: It was all software, yeah, and to get that to run in a Newton running at a hundred and some megahertz with—what was it—I think we got to use two hundred kilobytes of RAM and 200 kilobytes of ROM and that was the limit, and so it was tight; things were tough in those days. And nowadays you can have a thousand times that much memory and a thousand times that much speed; it comes a lot easier.

Steinbach: Would you say that Newton was also ahead of its time?

Lyon: Oh, yeah, it was ahead of its time and it was Sculley's brainchild and then it survived through Spindler and Amelio, but when Steve Jobs came back he killed it off and there's various theories about why, but I think fundamentally to him it represented John Sculley so it had to go. <laughs>

Steinbach: Anything else you worked on at Apple?

Lyon: Oh, there were a lot of things, but a lot of the work in hearing research never really connected much with products; but we did a lot of good research publications and developed good relationships with the larger academic and industrial hearing research field, kind of built some reputation on that that I kind of came back with later at Google. I kind of dropped out of the field for a long time and some of my previous friends and contacts wondered why I had sort of dropped off the face of the earth for so long when I kind of reappeared doing hearing research again but—

Steinbach: What Apple currently has does not go back to that time.

Lyon: No.

Steinbach: They have Siri and so on—

Lyon: Right. Well, they bought Siri; I think it was a SRI spinoff or something like that. They bought Siri from some outside company; I'm not sure what its history is. They do have actually a couple of speech guys at Apple that were there back when I was there that I—I was there when they got hired in the '90s and they're still there but most of the rest of the team is gone, yeah, not much overlap between people I knew then and people that are doing their speech systems now. A few like Alex Acero went back—from Microsoft back to Apple; he had been there back in the '90s. I think he's running a lot of that now.

Steinbach: What about after Apple?

Lyon: So in—

Steinbach: You went to Google—

Lyon: In '97, Apple was kind of in the throes of suicide and we—actually we had been trying to find various ways to save the company and in '96 we looked at various potential acquisitions, for example Be Labs of Jean-Louis Gassée. We looked at that as they had some pretty nice software that we thought might be good for building a future on and—but it didn't quite click and then I guess Steve Jobs got wind of this and came over and said, "You should buy NeXT instead;" so they did. So they bought NeXT and Steve came back and—but the company was in desperate straits and everyone knew it and they could see layoffs ahead and so the morale was pretty bad. So I think it was in April kind of the hammer finally came down and they had massive layoffs; they laid off—

Steinbach: Ninety-seven you say?

Lyon: Ninety-seven they laid off about half or more of ATG, the research people, at that time including myself and my team and my manager and his manager and so on so about five levels of us gone; so we had a big party, because they had good severance packages at that time, so generally it was good for everybody, I think. Everyone that left Apple at that time ended up doing other cool things in Silicon Valley, so a whole bunch of activity there in '97 that maybe became part of the big dot-com boom and bust but people found things to do; it was not hard to find jobs at that time. And that's when Carver and I got together with some projects he had been running at Synaptics and National Semiconductor and got those people and put it together and started Foveon to do electronic color photography.

Steinbach: And you were back with photography—

Lyon: Back to photography, one of my old hobby habits. I had been head photographer for my high-school yearbook for a couple years back when I was in high school, and my photography habit had kind of come and gone over the years, but I knew a little bit about it, and little bits about optics, and very little about color, but when we started Foveon my thought was that since I had designed this optical mouse chip I knew something about imager chips. I at least knew how a photodiode worked and how to lay out chips. And so I thought that's what I'll be doing; but as we put the team together it quickly occurred to me that all these other people that Carver had gotten from Synaptics and National were actually—they had an ongoing image sensor design project already and all of these people were more into designing imager chips than I was, so I thought well, that's interesting, so what's my role here? But that was a great opportunity, because my role became everything else, so <laughs> we were trying to build cameras—

Steinbach: Foveon—

Lyon: —so—

Steinbach: —wanted to make a whole camera.

Lyon: Yeah, so we did. Our first product line was cameras, not sensors, and so I knew something about photography, and I took what little I knew about optics and became the optics expert, and what little I knew about color and learned more and became the color expert, and what little I knew about software, and I can't say I really became the software expert, but at least I became the guy kind of guiding the functionality of the software for the image processing, and so on. And I got into writing some software myself, and learned C++, and sort of got in there and hacked with the other guys. And also because we had no one else to do it, I became in charge of the intellectual property, and we had a lot of it as it turned out, because when we got started both Synaptics and National put in not just their people but also some money and all the IP around this joint image sensor project that Carver had organized between them. So we had piles of patents, patent applications, half-written patent applications, invention disclosures, and all this stuff, and over the first half a year or so of the company I sort of worked my way through this pile of stuff, trying to figure out what's worth doing anything about in there. So while developing this prism camera—I should mention that—here, I brought this toy. The front end of the camera we made has a Canon lens mount on it that we cobbled off some Canon adapters and on the back it has this color separation prism assembly, sort of like what you find in video cameras of the day, only different, and on each face of this there's an image sensor that we designed. These are 2K by 2K image sensors, so that's 4 megapixels in—oh, what was that—1998 I think, so that was pretty much ahead of its time. To put three 4 megapixel sensors in one camera with a high-quality color separation prism was really way ahead of our time. Digital cameras at that time had kind of low-resolution Bayer-pattern sensors; they were not really good enough to use in a portrait studio, because they would make terrible moire patterns on people's suits and things, so we decided go after the portrait studio business. We figured they have

enough money to pay for a camera that's a couple generations' better quality than the competition. That kind of worked technically. We made a camera that really kicked and the first one we sold to Sheldon of Los Altos down there in my hometown, to the photographers who had taken pictures of my baby kids a few years earlier. And I went back there 20 years later and they were still using it and they said, "We've made so much money with this camera. We really love it." And they eventually moved on. But the trouble with that business was it cost too much. We sold the system for fifty thousand dollars. And people couldn't see paying that for a camera. What they didn't really understand, and what was hard for us to sell, was that the whole system is what made the studios profitable. It's the ability to take the picture and then sit down with the client and look at the images right there while they're still all dressed up and excited about having their pictures taken. You could make the sale. You could show them the finished digital picture right then. And there was no other digital competition for that at the time. And this was why Sheldon of Los Altos was so happy and made so much money. It's not because it was the greatest camera in the world. It was because it was the only digital camera that they could do the sales with right on the spot. But that got passed by eventually by other digital cameras that were coming along in that timeframe.

But back to my intellectual property story. Going through this pile of stuff, I found lots of invention disclosures by Dick Merrill. And as our patent attorney explained to me, if you incentivize a behavior, you get that behavior. And at National, they paid a certain amount of money for invention disclosures. So, we got a lot of invention disclosures. Well, a lot of them were great, some of them not so great. I found this one that talked about sensing light at different depths in the silicon to get a wavelength dependence. And maybe you could do color imaging that way. And I thought that's a good idea. We could take this Foveon advantage of measuring all the photons and all the pixel locations with separate sensors, put those sensors in layers on a single chip, get rid of this expensive prism that I had designed, and make single-chip sensors that could have the Foveon advantage. No moire patterns. So, I took this invention disclosure. And I went to Carver. And I said, "Carver, have you seen this thing? We should think about this." And he says, "Oh yeah, interesting idea." He says, "But Tobi and I did something like that. I don't think it works." I said, "Oh, really? What'd you guys do?" So, I went and asked Tobi, Tobi Delbrück, who is one of the other guys that helped us start the company. He says, "Oh, yeah. Here's what we did." And he showed me his reports and his measurements and so on. And I analyzed it. And I went back to Carver, and I said, "No, no. Here's what you and Tobi did. You didn't have the different junctions in there at different depths. And you're not going to get much color from just different depths of junctions if you don't have multiple junctions in there." And he says, "Oh, yeah. I see what you mean." He says, "Well, let's go talk to Merrill." So, we went to the inventor, Dick Merrill. We said, "Dick, can we actually build this thing?" And he says, "Yeah, I kind of poked at something like that. It didn't work. I don't think it will work." And I said, "Well, why don't we build one, try to make it work?" He said, "Yeah, okay." So, he had an idea how to build it. We wrote up the patent, filed it, built it, everything. Get the chips back, and he hooks it up in camera and makes an image of a color chart. And he gives it to me and says, "See if you can make color out of that from the measurements off the chip." And I tried. And I tried. And I said, "No. It doesn't work." And I went back to Merrill and showed him what the problem was explained it to him. And he says, "Well, what if we could do this?" And he shows a better way of layering the junctions with alternating P and N layers, epitaxial silicon and everything. And I said, "You could do that?" And he says, "Yeah, I think we

can do that." And I said, "If you can do that, that will work." So, he did. He set about making a chip now that was 2K by 2K by 3. So, it had as many photo diodes as these three image sensors I showed you but all on one chip in layers. He got the guys over at National Semiconductor to do this special process with multiple epi layers, and plugs, and all this stuff. And they built a chip. He put it in a camera, set it up in a studio, and took a picture. Gave me the data, said see if you can make good color out of that. A little bit of tweaking, I came back. There it is, good color. I mean that thing kicked, the first chip. It was so good. It was so amazing. Everyone go so excited. And it's a good thing, too, because we had already sold it at that point. Carver had already talked Michihiro Yamaki of Sigma into building a camera around it. We didn't even have the first silicon yet. But I was so sure and talked Dick Merrill and Carver into being so sure that it would work that he went and sold it. And we talked Sigma into being so sure that they'd started a camera design around it. So, yeah that's where the first Sigma digital camera came from, the SD9 single lens reflex. Not from the first silicon, which was 2K by 2K because they wanted a more conventional three-to-two aspect ratio that everyone's used to in an SLR, but from that first one that worked, we designed the second one that went into their camera. And it was great fun time. I mean those years at Foveon working with Dick Merrill and Carver Mead and other great people there was some of the most fun of my life, although I've got to say working at Xerox, working at Google, and working at Schlumberger, working at Apple, they were all amazing awesome experiences. But yeah, the early years at Foveon and doing both the prism cameras that made studio photographers happy, and the novel kind of image sensor that got Sigma into a completely new line of cameras, that was fun times.

Steinbach: Those cameras were still also expensive, right?

Lyon: They were not terribly expensive. What did the SD9 cost? There was a lot of grousing about the price. It was a little higher than some of the Canon and Nikon competition at the time. But it made way better pictures. It really did. And it got us following every—it kind of reminded me of what the situation had been at Apple where the Mac was so much better than the PC, but cost more. And people groused about that. And only the people that really cared about how cool it was would pay for it. Well, it was the same way with the Sigma camera. The photographers that just fell in love with the image quality and could do awesome stuff with it just loved it. And it had a loyal but unfortunately too small following. And Nikon and Canon kept improving in various dimensions making their cameras more feature-full, more faster, more higher resolution. They were just moving a lot faster than what Sigma and Foveon could do. So, there was no catching up with that. So, it just stayed kind of a niche product for the people that appreciated the kind of look that the Foveon technology gave.

Steinbach: So, the Foveon sensor was not very much more expensive than—?

Lyon: No, I think the sensor silicon was not much more expensive. There was nothing inherently more expensive in the process. You get a certain number of sensors per wafer. And we had some specialization in the process, but all the image sensor processes did. We didn't have the volume, the scale of volume that Nikon and Canon had. So, it probably did cost a little more.

Steinbach: And they didn't want to buy it?

Lyon: Oh, we tried. Canon made their own sensors. So, they weren't interested. And Nikon, we had serious discussions with. But they tried making their own sensors. I think that wasn't a huge success. They mostly bought sensors from Sony, which were good sensors, but they had to pay too much for them. They didn't really enjoy the fact that Sony became a camera company competing with them while being their supplier for sensors. And so, they were looking for alternatives. But I think it was hard for us to meet their cost goals. And our silicon and the way the kind of processing it needed was also different enough. I think that was scary for them. And to some extent, it just didn't seem possible at the time for a small American company to be a supplier to a big Japanese company. It just—that was not a relationship that we figured out how to make work. It worked for a small Japanese company, for Sigma, but not for Nikon. We had a great relationship with Sigma, great people there, good engineering. In the end, long after I left, they eventually sold Foveon to Sigma. So, Foveon still exists. Sigma owns it. They still—their digital camera line is still completely based on Foveon sensors. And they came out with a new one—a new line of cameras recently called Quattro that refers to a sensor technology where they have the layered sensors in the silicon, but the top layer has four for every one buried. And that was an innovation we had made before I left in an attempt to make really small sensors for cellphone cameras to make less silicon, more resolution per silicon and so on. I never thought I would see that in a big single lens reflex camera even though I think we knew it had some advantages. But it really complicated their marketing story, which was all about measuring every color at every location. But they did come out with it. And it's good. And it's, again, it's got a—it's got a following of people that appreciate it for its properties and also has its detractors of people that either don't understand it or don't like the way it behaves or one thing or another. But that was a thing that Dick Merrill and I and Paul Hubel and others invented at Foveon back around 2004. And it took it ten years, but it came out.

Steinbach: Okay so, after Foveon, you went to Google, or—?

Lyon: So, around 2005 or so, things were not looking so good. I could see Foveon wasn't really managing to find any big markets. And we were trying to get into the mass market of sensors for cellphone cameras. And it didn't really look to me like that could succeed; again, this problem of a small American company trying to supply huge volume Asian businesses. And the technology was not really a great fit in the low end. And I really couldn't see how Foveon had a lot of future. So, I started looking around for alternatives. I had some friends at Yahoo that wanted me to get—wanted to get me over there to work with them. And I was considering doing that. And my wife said, "Well, if you're thinking about Yahoo, why don't you see what Google has in research?" And I said, "Google? They do research?" And Google was not known for research at the time. They had research. But they kept it kind of quiet. So, she said, "Why don't you call Eric Schmidt and see what's up over there." And I said, "Okay." So, I talked with Eric. And he said, "Yeah, come one over. We do some research." So, I started going over there and interviewing with people. And it did look interesting. They had plans to make their research more open and attract more people and so on. And I said, "Yeah, okay. I can be part of that." It was kind of crazy

times. Well, at Google, it's always crazy times. But this was one of their high growth periods, I think. And—

Steinbach: That's 20—

Lyon: 2006. So, by the time I started, it was April 2006. I think I first talked to them around the end of 2005.

Steinbach: So, that's ten years already.

Lyon: Yeah, I've been at Google ten and a half years, amazing. Yeah, so I got over there. And at the time, everyone in research reported to Peter Norvig. So, he was the head of research. And that was fine except, as everybody's manager, he didn't really have a lot of time to put a lot of attention on what anyone should do. So, everyone just did what anyone did. And that was kind of cool except here I am working for an Internet search company. What am I supposed to do? I don't know what to do. I wanted to work on analysis of media, sounds, images, videos, whatever. They didn't have YouTube, yet. So, they didn't have a lot of videos. But they had some images. And I thought they ought to do more with sound. So, I started looking into that and poking around and discovered they actually had quite a few people already that knew about images and video and were doing some work in that space. And I said, "Okay, I can focus on sound." But—and that was my background and interest anyway. So, that was good. But being the only one there working on sound, I still had a real hard time getting started. And I found, down the hall, I was going—walking down the hall one day, and I see in a meeting room one of the guys that had interviewed me was in there having a meeting. And I said, "I'll just drop in and see what that's about." So, I just sat in on this meeting. And it turned out, it was the original six guys on the Street View team. And they were showing some pictures they had made by taking the cameras off the book scanning project and putting them on top of a van and driving around. And they got some pictures. They were having a heck of a time getting good pictures and didn't understand why. So, I looked at what they were doing. And I said, "I'll come back next week and show what's wrong and how to fix it." So, I came back with a presentation the next time and explained to them all about how full frame CCDs work and what the aperture of a pixel cell is about and how that relates to the exit pupil distance of a lens and all this stuff that was, like a few years before, I didn't know anything about it either. But with my experience at Foveon, I could see the incompatibility between their optics and their sensors and their readout methods and so on. So, I kind of explained this stuff. And they said, "Oh, I see." And I said, "And here's a lens that will work better for you." And I got my friends at Sigma to give me a discount on a batch of lenses that would work better with their cameras. So, they show up in the mail. We stick them on. Okay, it works better. So, that was progress. So, I kind of worked my way into this team, just sort of kibitzing with what they were doing. And a couple years later, they invited me to come manage the camera development effort for Street View, for making their kind of second and third generation camera systems. So, I did that for a couple years, which, at the same time, I had a small effort going in sound, splitting my time back and forth

between different organizations in different buildings, and so on. But we got some good camera systems out and made Street View a lot better over—

Steinbach: I'm surprised that Google would design their own cameras, that they wouldn't take a camera off the shelf. But would that—do you really need that for this purpose?

Lyon: Yeah, people used to ask me about that. Like what's so special about these cameras? I said, "Well, you have to be able to run them into tree branches a couple times a day and not damage them for one thing," things like that. But no, there was a lot of things about these cameras that were pretty special, being able to take several images per second in all directions while moving fast, and get all the data downloaded, be able to stitch it together in a unified panorama, and all this stuff.

Steinbach: Okay.

Lyon: There was a lot to it. And we did buy cameras. And some of the generations we bought parts of cameras or cameras. Then we did custom lenses, and so on. So, we did a lot of different things there. There have been quite a few generations, and it's a very interesting story. I don't know if that story's really been told anywhere. But some aspects of it are in an article we did a few years ago on Street View that shows some pictures of some of the cameras. Anyway, that was fun. But in 2010, I left that team, turned it over to another guy. And I went back to focus entirely on sound and hearing research because we had a lot of opportunities heating up there. So, that's what I've been doing for the last six and a half years. And about that time, also I started writing a book on hearing. It's called "Human and Machine Hearing: Extracting Meaning from Sound." And I've recently finally turned over the finished manuscript to the publisher, to Cambridge University Press. So, that's sort of in production. Hopefully, it will be out by the end of the year or so. So, six and a half years writing the book. After the first half-year, I had half the book. And I taught a course on it at Stanford in the psychology department. Brian Wandell invited me to give this course on machine hearing. And that was great fun with kind of half a book, which was enough for a one-semester course. But it took another six years to finish it. You know how books can go sometimes.

Steinbach: I notice your t-shirt is also concerned with sound, right?

Lyon: This is the "What's This Song" project. If you have an Android phone, and you say, "What's this song?"—

Steinbach: I recently installed that just like last week on my Android phone.

Lyon: Doesn't it come built in?

Steinbach: No, I had to put it on.

Lyon: If you say, "What's this song," it will listen for music and tell you what it is. It will look it up, take you to Google Play and offer to sell you a copy, or whatever. But this was a project that I—

Steinbach: It got it wrong, by the way, the one song that I installed it for. It got it wrong.

Lyon: I've never seen it get it wrong. It will sometimes say—says no match. But I've never seen it actually get a wrong match.

Steinbach: It got me a different song from the same group.

Lyon: Ah, from the same group. Maybe they used some of the same riffs in it that matched. Interesting. Yeah, a couple of friends of mine in Zurich did this project. And I just collaborated a little bit on some of the features that they use for matching. I think there's a patent issued on that. So, people can find out about it if they want to know what we did.

Steinbach: So, I assume that that looks for pitch and—

Lyon: It looks for certain patterns of prominent landmarks in a spectro-temporal representation. It's an approach that a lot of people have done. But we did our own version of it that worked pretty well. We had done other projects before this on matching songs for copyrighted content detection. We had done one on—we call Melody ID for matching the melody. So, if you sing a song yourself or play it on your guitar and upload it to YouTube, it doesn't violate the copyright of a music label because they didn't publish your song. But it might be that the underlying melody is copyrighted. And those copyrights are managed by what are called the music publishers. So, by being able to match melodies, you could bring the—you can respect the rights of the music publishers and bring them into the advertising game because YouTube can't sell ads on your video unless they know who owns the rights. So, if we can identify a rights holder, then we can sell ads and share revenue with the rights holder. So, that brought more players into the whole ecosystem. And that was kind of cool. So, I've done those kinds of projects, a lot of work with YouTube, other projects, other product areas at Google.

Steinbach: But mostly concerned with sound recognition.

Lyon: Mostly sound related, yeah. We did another thing recently that's in what's called the nearby—is that a feature or a product, whatever. In Android, there's a capability for phones to find other phones nearby using sound between them to—

Steinbach: Oh, really?

Lyon: Yeah, and you can also use it—it was first announced, I think in the Chromecast guest mode that you can—your phone can communicate with a Chromecast device on a television to establish that they're in the same room together. And this was a way of—so, you have a Chromecast on your television. And a friend comes over and wants to cast something to your TV. He'll say, "Well, what's your Wi-Fi credentials," and so on. You don't want to give him that. You want to use this guest mode that allows him to use his cellular connection to cast to your phone. And as long as they're in the same room together and we can verify that by the sound connection between them, then we can let that happen. So, Chromecast guest mode uses this nearby feature. There was a bunch of other things—

Steinbach: So, it whistles a tune, and that gets then verified, or recognized?

Lyon: So, what the nearby feature does is it just verifies that you're in the same room together by sending data over sound. It sends it like a token, a little code. And then Google verifies yeah, you heard the one that he sent. So, you're in that room.

Steinbach: I hadn't heard of that.

Lyon: So, and they have multiplayer games. If you want to find people near you that want to play the game with you, they'll show up in the nearby listing, a bunch of other things. So, that's been opened up to third parties. So, in general, being able to find devices near you is what that's about. And it uses a sound signal that's at high enough frequencies that you, and I, and most people over twenty years old can't hear. But some younger kids can hear it because it's at the top end of the band that the phones can send and hear, which is about eighteen and a half to twenty kilohertz.

Steinbach: But all phones should be able to deal with it?

Lyon: All modern phones—

Steinbach: I guess they're so small—the speaker sounds so small that the—

Lyon: The high frequencies are not that much of a challenge. The reason they can hear up to twenty kilohertz is because that's kind of the range of human hearing. So, to get a good spec on a phone, you say it goes to twenty kilohertz. So, we use the upper little bit of that band. Nobody can actually hear it unless they're a teenager maybe. And it's not very loud. And it's not very often, just when it needs it. So, we've done that. And there was also another audible version of that called Google Tone that's used

sometimes in classrooms. For example, if a teacher wants all the students to go to a certain page, she can send the URL with sound tones, kind of like the DTMF signaling on a telephone going boop boop boop boop boop. And everyone gets it. And they go to the right page. So, there's cool projects like that that have to do with sound but nothing to do with hearing, and lots of things that do have to do with hearing as well. A lot of stuff we're working on now that's quite exciting, but I can't really talk about it yet.

Steinbach: So, you're in Google's research lab.

Lyon: Yeah, we don't call it a lab. But there's an organization called Research that, unlike the research labs of a lot of other companies, it's not really very distant from the product groups. And we tend to have very good relationships. And researchers naturally tend to find problems to work on from the product groups in the company. So, they're usually working on things that are very relevant but far enough out to be more challenging and more than what a product group might take on themselves. But also at Google, a lot of product groups do take on research themselves. And so, we have a lot of PhD researcher types come in and work directly in the product groups rather than in the research organization. The research organization has kept its scope kind of narrow, mostly around machine intelligence, machine perception, natural language, and things like that. Whereas, a lot of the more computer system researchy kind of things will be done in the product groups. It's a good organization.

Steinbach: Yeah, well the key I guess is that you do have a good connection to the product groups.

Lyon: Yeah.

Steinbach: I saw one thing that you haven't mentioned that I found in some list of your papers or something was you—I think at Apple, you worked in shading in computer graphics things. That's another totally different—

Lyon: Yeah, I did this Apple tech report called "Phong Shading Simplification for Hardware Renderers" or something like that. So, for those who don't know—I'm sure most people do. But Phong shading is a computer graphics rendering shading technique named for a fellow named Bui Tuong Phong, who was a Vietnamese guy that worked as a grad student at the University of Utah when Ivan Sutherland had the graphics group there. And I had been, I think, really fortunate in my career to get to know almost everybody that worked in that graphics group under Ivan. Like at Xerox PARC, we had—let's see. Who all came from that group? Certainly, John Warnock and—was Martin Newell part of that group? I forget. Anyway, Jim Clark. And at Apple, we had others, Frank Crow. Alan Kay worked with Ivan. I don't know if that was at Utah or before that. But I had known all Ivan's ex-students for some reason, for various sundry reasons. So, I had a lot of exposure to graphics. I never really worked in that area. But I never knew Phong. It turned out, sadly, he had died shortly after this work. I think he actually got out of grad school and moved to Stanford, and then died of cancer, brain cancer or something like that. But he had

come up with this interesting scheme for rendering specular reflections off sort of semi-shiny materials. And it involved a mathematical expression that was the nth power of the cosine of an angle. I thought—and that thing's a bear to compute. And people that actually were building rendering systems would actually—they'd have software cosine routines and exponential routines and things that were like just taking all the compute time. Or if you'd made hardware to do it, it was horrible. And a number of people had looked at it and simplified it in various ways. So, like Jim Blinn, who had been at Caltech, and I guess he was another ex-Ivan student, had done some simplified versions of it. But they behaved rather differently and not entirely satisfactory to my way of looking at it. So, we had some people doing graphics chips at Apple. And I looked at what they were doing. And I thought, "This is horrible. Why can't we do something better?" So, I came up with a sort of a different way of looking at it that simplified the calculation to a handful of multiplies and adds including the vector normalization that was needed to do the approximate cosine operations. It was very sensitive. If things weren't perfectly normalized, then using a dot product to do a cosine doesn't give the right answer. And you raise that to a high power, it gives a terrible answer. So, I found out a way to make that much more robust using vector differences instead of dot products. And so, I wrote this report on it. And the cover of the report has like a block diagram. It's got adders and multipliers and a few things. And it's like the whole thing right there on the cover of the report. There's no cosines or logs or exponentials anywhere in there. So, it was very nice. That was all I ever did with shading. I mean all I did was write that report. I didn't actually ever build it. I wrote some MATLAB code to show how it behaved. I had pictures in there that I generated. But I never built a hardware rendering. I just sort of worked near the people who did.

Steinbach: Do you think Apple is using that?

Lyon: No, I don't think so. It's funny. There's a guy at Apple who—he had been at Apple. He credited me with helping him get a job there or something. And he said they actually used it where he came from in Germany. And they thought I was a graphics guy because I had done this report on this great thing. And they used it. But no, I don't think Apple—what Apple did is they put the report in the library. And when Steve Jobs came in, they gave the library away to Stanford. So, I don't think—I don't think it had any influence on Apple.

Steinbach: Okay so, we should also mention the awards that you have received. Right, you got—you're a fellow of the IEEE, of the Association for Computing Machinery—

Lyon: Yeah, I got the two fellow deals, the IEEE and the ACM.

Steinbach: And so, what did you get the IEEE for?

Lyon: Well, we'd have to go back and read the citation. I'm not sure what it said. But what I remember is it had a list of several things like signal processing models of hearing or something and sensory systems.

Maybe—I think it might have mentioned handwriting recognition. What I do remember is the ACM one specifically called out the optical mouse. For the computing machinery people, I think that machinery was a good thing. But I think if you have a career in Silicon Valley, and you do some good stuff, you should eventually be able to get to be a fellow of these professional organizations. So, that's not as special to me as the award we got from the Royal Photographic Society, which annually presents a progress medal to like the biggest thing that year. And they gave that to me and Carver Mead and Dick Merrill for the development of the Foveon X3 technology. And I was—

Steinbach: So, this is the stacked sensor?

Lyon: Yeah, the stacked photo diodes. And yeah, I got to be the one that went to London and go to the awards ceremony and get the medals and so on. That was a big honor. That was huge.

Steinbach: Okay, and finally, I saw on the net that you're a historian of slide rules, very simple computing machines.

Lyon: Yeah, how did that happen? So, it's funny. When I was in high school and even college, to some extent, I used a slide rule. That was what we did at the time. I got my first electronic calculator from Carver, actually, when I was a teaching assistant for his electronic projects lab class. He gave all his TAs HP-35 calculators as a door prize, I guess. So, I kind of went digital at that point and never looked back. So, the slide rules just—those are junk now. I've got digital. But at some point, somehow I discovered eBay. What was it called? In 1996, it might have still been called Auction Web, I think it was. And it was at that time when Apple was falling apart and people were demoralized. And I ended up spending time surfing the web because that was a new thing. And instead of paying attention to work, I kind of said, "Look at that. People are selling old slide rules. That's weird." And I ended up buying one. Then I ended up trying to figure out something about it like, "Where did this thing come from? Who made this one, where?" And I discovered there was this local organization called the Oughtred Society named after William Oughtred, the inventor of the slide rule in 1620. The guys that organized it are local. It's like the guy that edits the journal lives in Palo Alto. And I thought, "Huh." Oh, look. They have a meeting coming up. I'll go to that. So, I went to their meeting and met the old guard collectors and historians of slide rules. And I was the young guy there. I was only, what, forty-something at the time. And most of these guys were retired. So, I kind of got interested in it. And I ended up writing an article for their journal on the history and dating of the Otis King's patent calculator, a slide rule with helical scales. There had been three previous articles on the subject in their journal. And I looked at what they had. And I started gathering up new information from people, which I was able to do more easily now that there was the Internet and I could Google things. And I could email people. And I could start assembling the data. And I looked at the articles and the data, and no, they got it wrong. So, I put together a new history of this calculator and wrote a new article. And the article specifically called out in the title using the Internet because this was like a whole new thing to the slide rule guys. It's like, "What? You like email people. And you get information?" So, that was kind of cool. We brought a sort of modern information technology to

slide rule history. That's the only thing I've published in that area. That was in '97, I think. So, it's not like I've done a lot with in the last twenty years. But I am still a collector. I do still go to their meetings. And I've made a lot of good friends through that.

END OF THE INTERVIEW

