Oral History of Samuel H. “Sam” Fuller*

Interviewed by:
Douglas Fairbairn

Recorded: May 13, 2013
Mountain View, California

CHM Reference number: X6805.2013

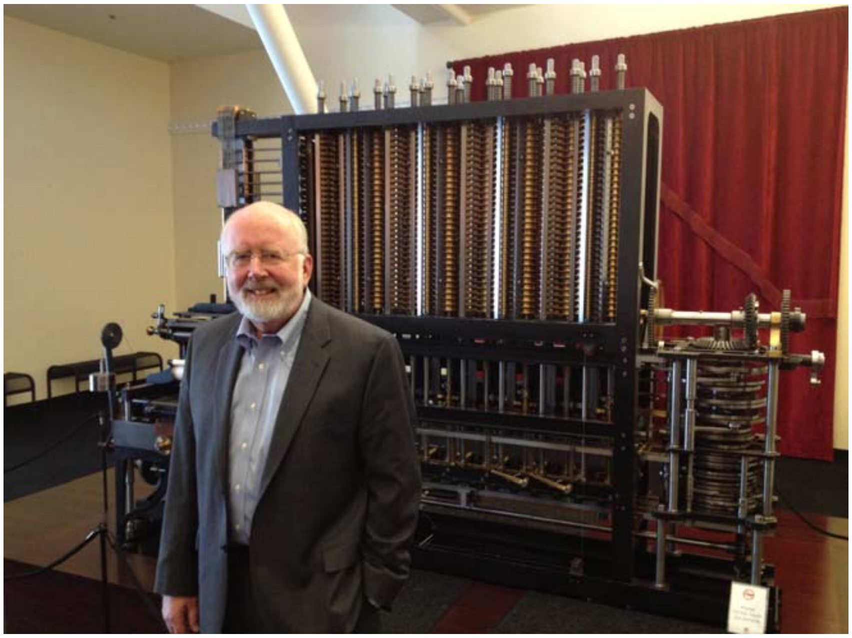

Dr. Samuel H. Fuller, May 13, 2013

*** Note from Sam Fuller:** I am Samuel H. Fuller 3rd, but have not used the 3rd in my professional life. My grandfather, Samuel H. Fuller, Sr. had only an 8th grade education, started as a coal miner in West Elizabeth, PA, but went on to become a successful salesman, and then founder and president of Fuller Instrument Co. My Dad, Samuel H. Fuller, Jr., had a BS and MS in Mechanical Engineering and worked on advanced technology for the US Army Tank Automotive Command.

Doug Fairbairn: Okay. Today is May 13, 2013. We're here at the Computer History Museum and I'm here to conduct an Oral History of Sam Fuller. My name is Doug Fairbairn. So welcome, Sam. Glad to have you here.

Samuel H. "Sam" Fuller: Glad to be here.

Fairbairn: So, Sam, as we've spoken about before we want to cover your—everything from your early life through your career to present day activities and what some of your views in terms of what's happening in the world of computing, and software, and so forth today. So to start off, let's just go back to the early days. Tell us where you were born. What the community was like that you grew up in, your family life, siblings, early influencers that may have or not steered you in the direction of technology.

Fuller: Okay. Well, that goes back a long ways. But I was born back in 1946, born in Detroit. And I grew up in the Detroit area. My earliest memories in elementary school, going to public school in Detroit. I just started out with a lot of interest in tinkering with mechanical things. I was taking apart and rebuilding my bicycle early on.

Fairbairn: Then Detroit was the capital of automobile manufacturing.

Fuller: Yeah, that’s exactly why we ended up there. My dad was a mechanical engineer who served in World War II. And once the war was over he was transferred from the White Sands Proving Grounds to the Tank Automotive Command up in Detroit, Michigan to design future military vehicles. So that’s how we ended up there. While he’s working at the Tank Automotive Command I was growing up in Detroit, tinkering with bicycles but then it turned out my grandfather who was a salesman had set up Fuller Instruments Company, which rebuilt and sold mechanical counters for gas station pumps, gas meters, as well as general purpose industrial counters.

Fairbairn: So there was an entrepreneur in the family?

Fuller: So there was in my grandfather. He was a great and generous gentleman. He got me at the age of eight or nine I think to start taking apart all the counters, cleaning them up, and then he and my dad would refurbish and resell them. So I got involved in interesting mechanical devices at that point.

And then as things moved on, along with my older sister our family at a small, two-bedroom house built after the war for returning veterans. By 1956 our family grew with two more sisters. So we outgrew that house and moved- out to Warren, Michigan about a mile away from what turned out to be the General Motors Technical Center. I was in junior high school and high school as GM was building the Tech Center and it was great fun to ride a bike over there and see what they’re building and ride past Fisher Design Studios, and GM, and Chevrolet R&D buildings. So, again, it was an interest ing seeing what engineers could do and what the fun there might be in engineering and design. So growing through junior high school followed several science fairs, and winning a couple of prizes along the way at the science fairs.

Fairbairn: Any teachers in that period that were particularly influential in terms of your--

Fuller: When I think back there was one math teacher, Nester Solamon, who was a great teacher and as I was developing an interest in mathematics he certainly encouraged that. And so I certainly have fond memories of him. But in my high school years I was— a lot of my interests was outside of the school itself in science fairs and making various projects. I made a small sailboat with my dad to go out on the lake by our vacation cottage and then by, I think, junior year my grandfather’s Fuller Instrument business had really evolved toward electro-mechanical counters and so of the work became building logic out of electrical relays. My dad realized he needed to build up an electronics lab in the basement to be able to do some of this work. So he began making instruments from Heathkits: oscilloscopes, signal generators.

Through that I got involved and learned about soldering and a little bit about electronics. I also helped him build the family stereo system.

Fairbairn: Yeah, Heathkits were an amazing thing in those days.

Fuller: They were a lot of fun back then when you had individual transistors, and resisters, and capacitors.

Fairbairn: And you could build stuff that you then used too.

Fuller: You could. Right. It is one unfortunate thing about integrated circuits is a lot of that has gone away. But maybe we'll come back to it. I'm working right now at Analog Devices to reinvent that experience for kids with some new kits. But working with Heathkits was a way to begin to understand electronics and I think that convinced me at that point that, "Well, rather than mechanical engineering, electrical engineering might be a better, more promising field of work." And so by my junior, senior year in high school I knew I wanted to go into engineering. I began looking at good engineering schools and Michigan had and has a very good engineering school at the University of Michigan. So I applied and went to U of M to do my undergraduate work.

Fairbairn: Now, did any of your siblings, your sisters, wind up in any technical area?

Fuller: No. None of my sisters. Our family had three sisters They went into elementary education, medical technology, and mechanical technician. But none of them in engineering design. One of my three children has gone into computer science and maybe some of my grandchildren will go into science and engineering. But my sisters, no, I was the only one in the family that really got focused on a long career in technology.

So that was the early years. And then the University of Michigan was a great challenge. It was, and is, a great engineering school. It also had a good math department. I think I went through all the math courses and most of the graduate courses in math while I was in Michigan. And I would say maybe the most seminal course in Michigan was in first semester of my sophomore year, I took my first programming course. It was building programs with MAD, which was the Michigan Algorithmic Decoder Language that Bruce Arden and Bernie Galler had developed. It was one of the precursors to ALGOL at that point.

Fairbairn: So what year did you enter the university?

Fuller: I went to University of Michigan in the fall of 1964.

Fairbairn: ’64.

Fuller: It was in the fall of ’65 that I took this first programming course.

Fairbairn: And what were you— what machine were you programming?

Fuller: It was an IBM 7090 at the time. We used keypunch to input the data. And I was fascinated that each programming assignment was a wonderful puzzle and after I got the program running, I went looking to ways to optimize it in various ways. So as I finished that course I realized that rather than a general interest in electrical engineering, maybe designing audio systems or TVs or whatever, that there was a lot of excitement to be had in computing and designing and working with computers. So pretty quickly I got a part time job in one of Michigan’s computer labs led by Professor Irani building simulation programs for a number of military applications. A lot of them were Markov chain models of various system designs.

Fairbairn: So is this unusual to have an undergraduate working in that kind of environment or are these mainly--

Fuller: Uh... maybe a little bit. I mean, as I think back with Professor Irani’s lab it was all graduate students except for myself. I was just a lowly programmer, debugger at that point running card decks back and forth to the computer center, getting the simulations to run. I stayed on campus after my sophomore year to work in that lab. I got to know people in the computer center and by the fall of my junior year I’d managed to get myself a job in the computer center working on the PDP-8. They were about to get an IBM 360/65 as an early timesharing system. And Michigan was developing MTS, the Michigan Timesharing System and that was going to be one of the early software systems on the IBM 360/67. But they needed a communications concentrator so there various Teletype terminals around campus could interact with the timesharing system. And they decided to make a PDP-8 that communications concentrator. So I got a chance to work on developing the interfaces and getting the PDP-8 to begin taking inputs from remote terminals.

Fairbairn: Both hardware and software interfaces?

Fuller: So in that computer project it was mainly hardware. It was a chance to after working in the software lab to work on hardware. So I did logic design and working with computer modules that are the logic modules that DEC had to plug into PDP-8s. And so and I spent my junior year’s spare time working in that hardware lab. And as my junior year progressed and I was really, really interested in working on computers— I saw this ad in the placement center, which said, “We have more computing under one roof than any other place in the world.” I said, “Okay. That sounds great.” And that turned out to be the

National Security Agency. So I got myself a job there in the summer after my junior year. That was great. I worked in the research labs at NSA.

Fairbairn: You have to get clearance at that point?

Fuller: Yes, they probably spend more money getting my clearance than they paid me for salary while I was there during that summer. But, yes, you had to go through a pretty extensive clearance process. But it was— it was great a great summer experience.

Fairbairn: And what kind of computers were you working on at those labs?

Fuller: So in the summer of ’67, that was the height of the Vietnam War and they had some special purpose military computers that they were using to collect what they called “SIGINTs,” or Signal Intelligence. So that summer I was working on the electronic front-ends that were transferring radar signals and getting them prepared to then be loaded into machines for further analysis. I went back to NSA again the next summer and began working on the diagnostic equipment for some of the special purpose computers that were being built there. One of my most memorable days at the NSA was when I learned that in one of their highly classified sections they had a “full” IBM 360. And what that meant was they actually had16megabytes of main memory. So the full 24-bit address space of a 360 was filled out in this computer. I said, “Gee, I’d really like to see that complete 360.” So it took a better part of the summer but I finally— my boss finally got me clearance to go and take a look at it. So it was the 360/65. It had the full 16-megabytes of memory. It was all magnetic core memory. It was these four rows of cabinets and the cabinet row went on for maybe 30, 40 feet, each of the four, 4 Megabyte sections of memory. But that allowed them with core memory to get the full 16-megabytes. You know, years, years later when I was at Digital I got a PC that had 16-megabytes in it on my desk. I said, “Wow, I finally got myself a real computer on my desk.” That was great to see. So that was my undergraduate years at Michigan. You know, I finished up electrical engineering at Michigan. I did quite well there and knew that I wanted to go on to graduate work.

Fairbairn: Yeah. You said you graduated Summa Cum Laude?

Fuller: That’s right. So I graduate at the top of the class at Michigan and so I began looking at graduate schools. They initially wanted me to stay at the University of Michigan but Professor Irani who had— the professor I originally worked for had said— I’d taken most of the math courses and much of the first year of graduate courses and computer science. And he said, “Well, you know, you’d probably learn more in another environment.” So I applied to Stanford and MIT. And I got into both places. I was arguing about— thinking about which way to go and one of the things I learned from that was while I was thinking about it, Ed McCluskey, a professor at Stanford, called and he was at that time the Director of the Computer Systems Lab. Over a couple of phone calls he talked to me about the kind of research he was doing at

CSL and what I could get involved in if I went out there. And I never got a call from MIT. And by then I was married. My wife was more eager to go to California than she was to go to Cambridge. So given both factors, I went to Stanford for graduate work.

Fairbairn: Was she also a native Michigan?

Fuller: No. You know, I should have said earlier that one of the things happened on my first summer down at the NSA was— I met this girl that was a friend of one of my fraternity brothers. And she became my future wife. I dated her that summer when I was down at the NSA.

Fairbairn: So that paid out in multiple aspects.

Fuller: Right. So she was actually a graduate of the University of Maryland. As I said, we decided to go out to Stanford. And it turned out to be great to work in Stanford’s Computer Systems Lab. There was McCluskey, which was one of the early designers of logic systems and hardware. But then Forest Baskett had recently joined Stanford as an assistant professor. He turned out to be a great person on my PhD committee and a great colleague to work with all my years at Stanford. Harold Stone was my other reader on my dissertation. So I was out there for four years. I originally went out on an NSF Fellowship but my wife and I had two children and Ed McCluskey said, “Well, why don’t you apply for a Hertz Fellowship because that could provide good support for your time here.” So I did apply and the final interview round was I think eight or nine candidates, and we were interviewed by Edward Teller. All of us traveled up to Teller’s house in the foothills of Berkeley to be interviewed by Edward Teller. And, as you may know, he’s a very intense individual. So he talked to each of us for a little while and then sort of, I think, halfway through the evening he said, “Well, look--

Fairbairn: Each of you being the number of students from Stanford that were applying for this?

Fuller: Right. Either all of us were from Stanford or some of them may have been from Berkeley as well. I remember we were sort of in the living room talking with Teller. And he said, “Well, look. I want each of you to think about a problem. I want you to think about the infinite series that consist of the sum of the reciprocal of the prime numbers. I want you to determine whether the series converges or diverges.” So we all went off to different rooms. I went off to his dining room. Now, years later I looked into it figured out that the series diverged. But none of us at that time knew that. So I probably did half a dozen pages of calculations. I did get some lower bounds and upper bounds. It seemed— and I knew sort of how dense the primes were as you went forward in the number series. So I couldn’t actually prove it one way or the other but I said, you know, to me the preponderance of evidence was that it diverged. And so I went back and talked to Edward Teller about this, as he did with each one of us. And I think he wasn’t so much interested as to whether I got the right answer. He was more interested in …

Fairbairn: How you thought about the problem.

Fuller: How we thought about it and how we came to our conclusions. So the net out of that is that I did get a Hertz Fellowship and that was terrific because I originally was paying my own way, getting some help from the NSF Fellowship. But I was going to pursue a PhD in three years and move on into industry. But with the Hertz Fellowship I could slow down a little bit. I took four years to get the PhD and went to a lot more seminars and took a lot better advantage of what was going on at Stanford and looking at various problems in some detail. I did a lot of my research up at SLAC, the Stanford Linear Accelerator Center where they had an IBM 360/91, which was at that time, a great computing resource. And I would take over the computer at night or on weekends when the physicists weren’t using it for some of their analyses or simulations.

Fairbairn: So what were you doing your research in? What was the topic of your thesis?

Fuller: I was looking at optimizing the performance of memory hierarchies. And actually one of the summers as a grad student I spent the summer at HP labs designing cache memories for the next generation of HP Computers.

Fairbairn: 2116, 14, 15, 16.

Fuller: Well, the HP 2116 was the workhorse I had to use to do the simulations. But it was for the next generation after the 2116’s. So a large part of my dissertation was then on higher levels of the memory hierarchy and the one that was most problematic at the time back in ’68 to 72 were the rotating storage of drums and disks. So I had looked at how to optimize the organization, the data, and the searching for data on disks and drums. And proved some things mathematically, built some simulations as well. I would say that one of the things I’m most proud of was doing some of the mathematical computations we with very large polynomial factors. I found out at that time that Joel Moses at MIT had developed a program called MACSYMA for doing symbolic computations. And so in, I guess it would be, in 1970 and ’71 using the early ARPAnet, I did a remote log-in to MIT and did a lot of my calculations on MACSYMA and that’s a whole chapter of my dissertation. And let’s see the other things of consequence with that— the graduate work. And met a lot of interesting people. As I said, I spent one summer at Hewlett Packard looking at their labs— working at their labs. And then I also went back to IBM Research for a summer and worked. They were starting their work on FS, for Future Systems, which was going to be the replacement for the IBM 360 and which ended up not going forward ultimately. But it was again a great place to work on computer architecture before I went back to Stanford. Upon graduation, I was probably headed either to IBM Research or Xerox PARC. But then one of the Stanford seminars was given by Gordon Bell who was at Carnegie-Mellon University at the time. It was about the multiprocessor research that was going on there. So I talked to Gordon after the seminar and Gordon invited me back to CMU to give a seminar. And, at the end of it I said, “What could be better than to go work on these advance multiprocessor systems at CMU.” And so I redirected myself from going into industry to going to a research university.

Fairbairn: Had you done any multiprocessor work? And is that a factor in the research you’ve done before?

Fuller: I hadn’t done any multiprocessor work at Stanford. But looking at computer architectures I realized that IBM and Burroughs and few others were beginning to look at, multiprocessor systems. And so I could see that as a real interesting direction for the future. And there were a lot of challenges as to how you would get multiprocessors to effectively work with large memory hierarchies and how the software would be constructed. So it seemed like too great an opportunity to pass up. So I had to convince my wife to move from Palo Alto, California to Pittsburg, Pennsylvania.

Fairbairn: To Pittsburg. Yeah.

Fuller: But we did. And, you know, Pittsburg is really a delightful place to live and to raise kids and so forth. So I went to CMU. It was a great environment. They had a large DARPA grant so as a young faculty member I didn’t spend all my time writing proposals. I could really focus on the research. And the other sort of plus, or minus, event was I was a month away from getting my degree at Stanford and going to CMU and Gordon calls. Gordon Bell called to say that, “Well, DEC realized that PDP-11 is beginning to run out steam. They need to begin looking at the next generation of computers.” And he needed to go back to DEC. So when I got there he was no longer going to be there. But he had, I think he had, something like 15 or 16 graduate students. And so between myself and another recent graduate from Stanford, Dan Siewiorek, who also went to CMU we divided up his 16 graduate students. And they’re people such as Andy Bechtolsheim, John Ousterhout, Doug Clark, Dileep Bhandarkar, Richard Swan, and Conrad Lee. A great set of students that I began working with the day I showed up, fully funded, on the DARPA contract. At any other university would have taken three to five years to even begin to get to that point. So it was with a great set of students we worked on C.mmp, which was fairly far along in construction when I arrived and then developed the architecture and implementation details for the next generation multiprocessor: Cm*.

Fairbairn: Can you explain what C.mmp is?

Fuller: So C.mmp was a multiprocessor consisting of 16 PDP-11 processors going into a 16 by 16 crossbar switch going into 16 banks of main memory. And a large part of the DARPA work was the software system built by Bill Wulf and his graduate students. C.mmp was fairly far along in construction when I got there. But we were ambitious and we said, “Well, what does the next generation look like?” And so I got in on the ground floor designing the structure of what end up being calling CM*, for scalable number of computer modules where each computer module was a full LSI-11 computer system LSI-11 with it’s own memory. And then the challenge was you couldn’t build at that point an effective 64 by 64 crossbar switch so what kind of interconnect network would you make? And so we came up with a scalable, hierarchal interconnect for these computers as what now would be call NUMA, a Non-Uniform Memory Access computer architecture, or memory architecture. So I was a principal in designing the architecture for that memory structure and extending the architecture for the LSI-11’s memory addressing so that they could access multiple megabytes of memory even though PDP-11 natively just could just address 64 kilobytes of memory.

CM* went on to get built and as it was getting built the other project that I got involved with at CMU was the Department of Defense, again through DARPA. Ed said, “We ought to begin taking advantage— the Defense ought to begin taking advantage of the rapid advances in commercial computer systems, rather than designing unique computer systems for military systems.” So then what they called their Military Computer Family project, which was evaluating all the computer architectures that were available at the time. Saying which ones would be most applicable to the future of military systems. So that was a great effort to run and I got involved in that in evaluating and developing metrics, benchmarks, in ways to look at the relative architectural advantages of the Interdata 8/32s, SEL 32, the DG Nova and the PDP-11 as well several of the militarized computers then in use. .

Fairbairn: So do you have simulators for each of these machines? Or how did you run benchmarks?

Fuller: We first came up with a set of metrics on how effective the architectures were in terms of instructions per clock tick, or in reality back, then how many clock ticks per instruction?. How efficient it was use of memory. And we then developed a dozen kernel benchmarks representative of military applications.

Fairbairn: And were those typically very heavily data oriented? What were their unique characteristics of these?

Fuller: Some of them had real-time constraints. So you would have data flowing in from some sensors and then doing the computation. Could they keep up with the particular data rates? Some of them were large matrix calculations doing analysis on large data sets that might come in from various areas. And so we’re going through that turned out to be quite a productive series of studies and we published a number of those reports in IEEE Spectrum, and other IEEE journals.

Fairbairn: Is there anything that particularly stood out from those things in terms of the results? Anything surprise you?

Fuller: Well, I would say the surprising thing was this was all pre-RISC. And the architectures like the PDP-11, which had a lot of auto indexing, so it had sophisticated ways to walk through memory, either single or multi-threaded … and variable length instructions, where you could really pack a lot of power into a single instruction. Certainly, on the benchmarks we had, came out very effective. And they actually— the results of that were to recommend the PDP-11 going forward out of the Military Computer Family. But a consequence of that was Gordon Bell who was at this time back DEC they were developing— or they were working toward developing the VAX architecture. And one of Digital's real concerns at that point was they foresaw that there would be a whole series of VAXs developed over the next 20 years. And they knew that unless they found some way to really have those VAXs be compatible going forward there was going to be a software nightmare trying to handle all the different models. They'd live through that with all the various different PDP-11's and the various variants that the PDP-11's had had. So Gordon and some others at DEC knew about the work on the Military Computer Family and so they ended up calling me and saying will I be interested in leaving CMU and going to Digital and setting up a VAX architecture group to try and control the evolution of the VAX architecture, and develop methods to ensure that successive VAX's really were software compatible with those going forward.

Fairbairn: So had you maintained regular contact with Gordon Bell? You took over his position essentially or part of his position going in. Was there regular communication back and forth?

Fuller: Yes and Gordon was great about staying in contact with us at CMU. II think he felt bad about leaving before I showed up to go back to DEC. And he would come back once a month. My wife would say, "Well, this is a Gordon weekend." So he would come back for a weekend and all the graduate students and myself and Gordon we just worked through the weekend either wrapping up the C.mmp or making presentations to him on our thoughts around the architecture for CM*. And it was terrific that he would dedicate one weekend a month for several years. And so we did stay in touch. But I would say through the— you know, the military computer family architecture work with DOD I ended up consulting with almost all the computer companies at the time, Burroughs, Univac, Honeywell and of course DEC.

Fairbairn: Each one was eager to make sure that their machine was shown properly right?

Fuller: They wanted their machine and they were often following work with that—they were trying to figure out what would come next. And I had gotten to know many people at those various companies. And so— but it was clear to me in working with those various companies that the place to be and where the excitement was, was at DEC at that time. And so DEC said, "There's a chance to set up the VAX Architecture Group at DEC." That was to me very attractive. And during the seven years I was at CMU I'd written a lot of technical articles. I think I published over 30 or 40 articles at that point with graduate students and so forth. And I was eager to see my work end up in machines that would be manufactured and produced in industry. This looked like a chance to do that. And so in '78 I left CMU. So the same problem that Gordon had with graduate students who had not completed their dissertations. I was going back on weekends to help some of the graduate students finish their dissertations.

Fairbairn: Those Sam Fuller weekends. <laughs>

Fuller: Yes. And but, you know, I got more and more engrossed in the VAX work at Digital and so I focused on that. I must say I hired some of my best graduate students at the time, Dileep Bhandarkar and

Doug Clark. As they graduated from CMU they joined the VAX Architecture Group and worked with me there.

Fairbairn: So when did Digital begin work on the VAX in earnest? Is that when Gordon went there or they already--

Fuller: Well, he went there knowing that there was serious competition at that point showing up from NOVA Computers, Interdata and SEL Computers. So he knew that the PDP-11 was in trouble. I think there were things like the PDP-11/74, which did extend the life of the PDP-11 considerably. There were some hacks that let it get access to a couple of megabytes of memory rather than just the 64 kilobytes. But I think after a year or two they realized that had limited legs. You needed to look at a machine that had a 32-bit address rather than a 16-bit address.

Fairbairn: Was it slow to make that realization or transition?

Fuller: It was— it was slow. There were 32-bit minis that were out there. And they were beginning to eat into the profits and the revenue of the PDP-11. So I haven’t got an exact date. But it was probably in ’75, ’76. I remember going down and consulting with Gordon and Bill Strecker and a few others. I think it was in ’76, on how you might extend the PDP-11 and how you might think about a 32-bit PDP-11. And ultimately, that evolved into the VAX— the VAX architecture first showing up in the VAX 11/780.

Fairbairn: This may be a diversion but do you know why that was? Why was Digital— DEC so slow to move to that next generation?

Fuller: I think it was like any company when you’ve got a very successful product line, as DEC did in the PDP-11. And it’s generating enormous revenues and profits— it’s very difficult to think about the product that will kill that cash cow. You know, we might come back to that in the final years of DEC because that is what happened at that point.

Fairbairn: Yeah, I was wondering whether there were— I was wondering actually what was in the back of my mind— was there a cultural thing that was actually visible then that played into what happened eventually?

Fuller: I would think— I think in the late ‘70s, early ‘80s DEC was still a very vibrant place where there were a lot of people with a lot of ideas. So I don’t think the problems had really set in at that point yet.

Fairbairn: So they were different. These are different issues.

Fuller: We’ll get to that, I think, shortly, right? So when I think about the several years I was running the VAX, building the VAX Architecture Group, I think the principle things that happened there were several fold. One, was just from scratch developing a set of methods for, “Well, how would you actually do the verification of a new VAX and prove that it was compatible with the prior VAX’s?” So fairly quickly after the 780 was developed, VAX 11/750 and 730 were in hardware development. So my first task was developing a set of verification tools that we could run on prototype 750s and 730s to determine how compatible they were with the 780.

Fairbairn: Excuse me for stepping back. In— was there a lot of controversy in terms of what the 780 ought to look like? Were there sort of major camps within the organization or was there a strong leader was Gordon or whatever saying, “This is the direction we’re going.”

Fuller: Well, Gordon did provide strong leadership. There’s no doubt about that. But there certainly was, -inside at Digital at the time, there was also a very profitable business in the PDP-10’s, 36-bit machines. And then there was a lot of belief there that, “Well, let’s just make an evolution of the PDP-10, a 36-bit machine. Let that be the flagship.” Gordon was the one that had realized that, “No, it really needs to be an 8 bit byte addressable machine. It needs to be a 32-bit machine, not a 36-bit machine.” Gordon did a lot of work on the PDP-6 and PDP-10. Its address was fundamentally an 18-bit address. And so again just being able to get to a quarter megabyte wasn’t what you needed. You needed to be able to address multiple gigabytes of memory directly.

Fairbairn: That’s right. So I had forgotten for the moment about the whole 10— I mean the 10 and the 20 and so forth. So they were developing an advance machine in parallel with the 11.

Fuller: They absolutely were. The KA-10, the KL-10 those PDP-10s were very effective machines.

Fairbairn: And extremely popular.

Fuller: Right. My whole time at CMU, all the real computing work was done on the KL-10’s and the KA-10. The C.mmp and CM Star were the experimental machines. But the company clearly was going to move forward with the VAX’s. And so a large number of the effort at Digital was focused on what became three systems, the 730, 750, and 780.

Fairbairn: So the 780 was the high end?

Fuller: Was the high end.

Fairbairn: And then that came out first. And then was--

Fuller: 750 and 730. And then ultimately they actually got the PDP-10 group to work on a higher VAX, the VAX 9000 but that came much later. So the problem then was how do we actually ensure there was compatibility? And I knew that just having a small set of benchmarks like I did on the Military Computer Family would ring out obvious changes. But these machines were going to go out into the thousands of copies and you needed to really try and get to the bottom. So what we did was develop programs that would run on the prototype machines. I took large applications but then more importantly we developed a set of synthetic applications, which walked through all the instructions, all the addressing modes, actualizing all the registers and different modes using various random number generators. Making sure that they were relatively prime so you weren't repeating things and we had run them for weeks on the prototype machines. And using various set of techniques, we can find out where there was a discrepancy and then we could unravel back to what the particular instruction was that was causing the difference. And what I remember about the 750 was we got to the point where when we found the discrepancy between the 780 and the 750 and it was as likely that the error was in the 780 as opposed to the 750, we said, "Okay. Well, it's probably okay for the 750 to go to the market." And the other thing I learned from that was you get any machine as complex as a 780 or a VAX or certainly any of the microprocessors today, they will never be 100 percent bug free. You can only hope that you provide strong enough verification that the bugs that remain are sufficiently idiosyncratic and rare that hopefully in the life of the computers, applications will never trip across them. But so that was the first was on the validation. Then I say the second thing that went on was another real competitive threat at that time back in 1980 was Tandem computer with their non-stop computer. And that was beginning to have a big impact within the financial markets for DEC.

Fairbairn: Now, Digital it came out with a VAX with a brand new operating system and brand new everything. How did they establish their base in that? I mean, what was the original target market for it?

Fuller: Well, the original target market was engineering, scientific programs. FORTRAN was probably the most important higher level language that was based on it. And the customer base was really the PDP-11 customer base. And many of those customers had PDP-11 model 74s, model 70s and those were the ones that were eager to get the VAX. And they quickly adopted them. They went to a lot of customers and then it spread.

Fairbairn: But it was a new operating system and new hardware, right?

Fuller: Well, it was but the VMS operating system on the VAX developed by Dave Cutler and his crew were the same team that had developed the RSX-11M on the PDP-11. It had a lot of the same feel.

Fairbairn: So it was very familiar and easy to port programs from one to the next?

Fuller: It was easier to move forward. And the VAX 11/780 had PDP-11 compatibility mode. So part of the selling feature was it could run PDP-11 programs but it could also run VAX.

Fairbairn: Right.

Fuller: But, then going back to Tandem which was a different type of threat, a number of us looked at what they were doing. And we came back and said, "Well, how would we get VAX to have this nonstop characteristic?" And if we copied what tandem had done it would have been a fairly major re-architecture of some of the fundamental structure of the VAX. And so, my insight at the time was well, we could actually build a computer interconnect which eventually came out as VAX clusters that could tie several 780s together, developed an architecture which had some roots back in CM* that allowed them to at a very high speed transfer large blocks of data between them so that they could stay in sync and operate as a redundant or be constructed to be able to tolerate failures. And the VAX clusters, I think, eventually got renamed VMS clusters. Turned out to be a very popular feature and many, I think maybe even most VAX or VMS installations as time went forward end up being multicomputer installations that had the VAX cluster interconnect.

Fairbairn: Now, the Tandem was aimed as sort of financial and business kind of applications. The VAX was aimed at scientific but there was still...

Fuller: But, you know the PDP 11 had gone into financial and business areas. We weren't eager to see those areas get eroded by Tandem, so...

Fairbairn: So, it was a big enough overlap that you really would be concerned.

Fuller: But then for engineering and scientific applications then they also found that well, gee, we need to use multicomputers to get the power that’s required. In fact, one of the things we discovered afterwards was that the Cray 1 had come out about that time. And as we visited the national labs, Los Alamos and Livermore, they said what they'd like to use 780s as peripheral processors for the Crays. And again, the VAX interconnects provided a very high performance way to be able to begin transferring data from the 780 front ends into the Cray 1’s going forward. So, I think, this had a real impact on Digital. Actually probably a larger impact at about that same time was I was working on the computer interconnects, Gordon said, " we need to understand what our whole interconnect structure looks like going forward from computer interconnects to local area networks to metropolitan area networks, and wide area networks." And DEC at that time in the industrial product line had the DECdataway as a credible local area network. So, I looked at that, but then I did some analysis against Aloha net from the University of Hawaii and Ethernet that Xerox PARC was working on. IBM had token rings at that point that they were looking at. And it was pretty obvious to me that contention arbitration that Ethernet or Aloha net used was a better way to go than the DECdataway which was a slotted time reservation local area network.

Fairbairn: What year was this?

Fuller: So, this must have been 1980 or '81 I think. We will be able to nail down the date because the fortuitous thing that happened, I had interviewed at PARC, I knew some of the folks at PARC and I called some of them. And about that time Bob Metcalfe decided to leave PARC and do something different. So, what clicked in my mind was “well, Bob, if you're an independent agent why don’t you come to DEC as a consultant and teach us about contention networks and Ethernet and maybe something really good can happen out of that?” And Bob probably already had that in mind to do something with Ethernet since he'd worked on it at PARC. He was eager to come do that. And so, for six or nine months, I forget exactly how long, Bob kept coming back as a consultant, talking with us about that and he was my agent in being able to convince folks inside Digital that Ethernet would be a much more effective to go forward then the DECdataway or the IBM token ring or other things that were going at the time. And Bob convinced us that if we're going to do this it needed to be an industry standard which we readily agreed to and we wanted another partner which eventually became Intel to be able to build integrated circuits to make it cost effective. Now, I would say Gordon Bell, he could see that this was a much better direction so he helped us sell that. And then, we began doing internal development on interfaces for VAXs and PDP 11 for Ethernet and Bob formed 3Com and he also began developing interfaces for VAXs and PDP 11s for their Unibuses and Qbuses. I remember one of the more memorable meetings, we were inside of DEC having some trouble developing the analog front ends for these Ethernet interfaces that would go out to the Ethernet cable. So, Bob came by with Ron Crane and a few others from 3Com and we had an afternoon working with the designers at DEC. I remember one of the designers at DEC stopped and asked Bob, "Bob, why are you teaching us all this? Once we know how to do this, Digital, we're this big company, we'll come out with these products, we'll just roll over 3Com. You'll be out of business." And Bob said, "Look, 3Com, we're just a little tiny company." He says, "If we can't move a whole lot faster than you guys we don't deserve to be in business." See, this had to be 1981 because he was, I think, already working on the interface for IBM PCs. And he would realize hey, there's a lot more business in PCs than there are in PDP 11’s and Unibuses and Qbuses. But, Bob and 3Com did help us get set up. Bob told me that in the first year of 3Com Digital was by far their largest customer because we were buying all the 3Com boards to help Ethernet networks inside Digital. I remember I was at that point I had moved back to Maynard headquarters. And I always found when people wanted to know what Ethernet was, I'd bring them to my office, I'd show them running some applications on a PDP 10 and then a VAX and it was just ASCII terminal so it wasn't that impressive on a demo. So, I had the end of an Ethernet cable in my office. I would unscrew the terminating resistor and everything would stop because the cable stopped being able to transmit the packets. Then, I'd screw the terminal.

Fairbairn: So, this really is happening.

Fuller: ...back on and things would pick up again. So, I said, "Okay, well, yeah, must be going over this cable." And this went on for two months and then I was at some meeting and somebody says, "Well, Ethernet's pretty good, but it's got this unreliability. Every week or so it kind of stops for 15 minutes." And so, I said, "Okay, that's probably me. I'll stop doing these demos." And we went forward from there. And other people figured out that you want Ethernet to be easily diagnosable and so this idea of having a multi-drop cable, which we did in the early day wasn't that great of an idea. You really wanted to have home-run wires from every desktop device to an Ethernet switch and then the switch would go out to the larger network. And so, if there's a failure or somebody else unplugged something they didn't disrupt the rest of the network. But, in those early days it was this multi-drop cable that went around.

Fairbairn: Yeah, that was amazing.

Fuller: Yes, so I think I probably had my largest impact on Digital was getting them to realize that they ought to tie up with Ethernet rather than try and push the DECdataway. But, shortly after that DEC was doing very well with VAXs and their other systems. And so, I remember it was in the fall of 1983 Ken Olsten and Gordon Bell sat down with me and said, "Gee, Sam, we really ought to begin developing some research labs at Digital to take a longer term look at the future." And they knew about my interest in working with universities and research programs so they said, "Would you like to set up or grow our research organization at Digital?" And I remember telling my friends, "I just got the best job in the computer industry." What could be better than this? So, that's what started in '83, that's what I did.

Fairbairn: Back east or out here?

Fuller: Well, it was meant to be back East so the first as I began working on recruiting and I went back to my earlier colleagues, Forest Baskett was one my readers at Stanford, Neal Wilhelm who was one of my office mates as a grad student and they were working on computer design and RISC processors were beginning to be talked about at that point. And so, Forest and Neal had said, "If you wanted to hire us to work at Digital, we'd be glad to set up a research lab to begin looking at RISC architectures." So, that sounded great. I made the proposal. I worked for a month or two to convince them to come back East. They made the argument true or false that well, it'd be a lot easier to hire a full research team to build these systems if we could do this in Palo Alto rather than back in Maynard. So, it took a bit of convincing of Gordon, but eventually DEC said, "Okay. Well, we'll do this in Palo Alto." So, we set up the Western Research \Lab for Digital. We did build the team. We did develop a RISC.

Fairbairn: Was Forest running that at the first? Was he the first head?

Fuller: Yes, Forest was the director. He was the director until he got recruited away to SGI. But, he was there until I think the initial RISC machines, which were called Titans, were built and came up. And then I started down this path of trying to convince Digital that this ought to be the future of the computers. So, I guess...

Fairbairn: So, what was the path at that time?

Fuller: ...at that point it was all VAX 11s.

Fairbairn: You had VAXs.

Fuller: The VAX architecture, VAX and VMS going forward. Separate time we could talk about PDP-10s, but they decided to stop development of the PDP-10s. They got the PDP 10 design team to begin working on a higher end version of VAX, which ultimately became the VAX 9000 to compete with IBM. So, DEC really was focused on VAX. I had to convince them that the Titan architecture would be a better way to go forward than just continuing with the VAX architecture. So, additionally I gave some talks and brought Neal and Forest back to give talks on Titan.

Fairbairn: So, they actually had a machine running here?

Fuller: Well, initially they had a design. They were working to develop their own design tools and then we ran some simulations. But, then ultimately we got some Titans running. Ultimately I think we got 12 of these Titans running out here in California. And we ran some small benchmarks. And it was quite an impressive difference, but the VAX advocates back in Maynard continued to say "Well, those are just benchmarks. They all fit in the cache, they all fit in first level memory. We demand real applications." So, eventually I got some folks in the VAX 8700 group. Doug Clark, one of the earlier graduate students I had hired and he was working on the VAX 8000 series. I got him and the Titan group to run a set of large applications on both the VAX 8700 and Titans. And the great thing about those they were built in the same ECL technology, they were both about the same size as a small refrigerator. Now, they both had about the same manufacturing cost. And I remember at the end of the day the Titans ended up being 3.7 times faster than the VAXs. And that report and that result convinced Digital okay, well maybe 10 or 20 percent doesn’t make a difference. But, if it was more than a factor of two, it’s really more like a factor of three, we have got to begin to work on that. And so, Digital got more serious about RISC. Initially, it was the PRISM work that went out to Seattle under Dave Cutler, the guy that developed VMS. Ultimately he left to go to Microsoft and it came back to the semiconductor division at Digital in Hudson, Massachusetts. And they ultimately developed the Alpha RISC architecture machine. And the tragedy around all of that was we started the RISC work in '83. I think we had a working machines by '85, '86. The company realized they needed a RISC machine by sometime in the late 1980s. But, until the board replaced Ken Olsen and brought in Bob Palmer as CEO, the company would not let the company release the Alpha architecture machine. One of the first things that Bob Palmer did once he became CEO was announce the Alpha architecture and announce it was available as a system as well as a chip. But, that was late in the day. If alpha had come out, I don't know, '87, '88 it might have been a different story. But, such was life. So, that was the Western Research Lab. And then the other thing that happened, a lot of things happened in '83, '84. But, in late '83 PARC had a change in management. Bill Spencer came in as the new director at PARC. There were disagreements between Spencer and Bob Taylor who was running the computing systems lab there and some combination of Bob Taylor resigning or Bob Taylor being fired resulted in Bob and then pretty quickly Chuck Thacker and others leaving. I said, "Well, these are certainly great computer scientists. Those guys could work in our research labs at DEC." And so, I began talking with Bob Taylor. Ultimately I convinced Digital that we ought to have a second research lab in Palo Alto, which became the Systems Research Center. And pretty quickly I think 12 or 14 individuals showed up. Most of them actually came from PARC, the Alto and Ethernet group. And so, one of my other memorable days was in the spring of '84, I think it was, I'm working in my office in Maynard and Ken Olsen calls "Sam, come down to the office." So, I said, "Okay." I go down to the office and in the office is David Kerns, the CEO of Xerox, with Ken. And David Kerns says, "I've just been talking with Ken." He says, "How could a great company like Digital pirate all these people away from PARC. It's not right." And so, I explained on the spot with Ken and Kerns. I said, "Well, look, we didn't hire Bob Taylor and those others away from PARC. They quit. And they're going to work somewhere. They're going to go work here or they're going to go work for one of our competitors." I said, "I think it'd be better to have them work at Digital than have them working for one of our competitors." And that to Ken made sense and Kerns didn't really have a response to that, because we had not triggered their departure in any way. It was just that we were fortunate enough to know about it and act quickly enough to hire them.

Fairbairn: Two questions. During this period when the Western Research Lab was doing what it was doing what was going on in the eastern time zone, whatever you established there, and secondly, was Forest Baskett still running the other lab in Palo Alto and how did that go down?

Fuller: Well, initially I thought "Well, maybe we'll just enlarge the Western Research Lab and have Bob and Forest there."

Fairbairn: That wasn't going to work.

Fuller: That wasn't going to work <laughs> and both of them.

Fairbairn: I know that it wouldn't work on the Bob Taylor side. Forest's side I was curious as to what?

Fuller: Forest was doing great work at WRL. He spent some time at PARC back and forth. And I think he valued independence of not being part of Bob Taylor's lab. So, it was pretty clear to me that we'd get more results out of these individuals if they just reported separately to me and not try to merge them into a single organization. So, yes, it was a bit of a sell back in Maynard that we had to have a couple of labs in Palo Alto. And to your point, at about the same time I did think it was important to set up a research lab back East. And there was a small group in Maynard working on it. But, one of my beliefs was if you could co-locate or get your research lab to collaborate more effectively with universities, it'll be a richer environment. And so, the decision was to move the small group that was in Maynard down into Kendall Square so we set up a lab there. We recruited Victor Vyssotsky out of Bell Labs to be the director of the Cambridge Research Lab. And concurrent with the development of the Systems Research Center and the Western Research Lab we began building up the Cambridge research lab under Victor. So, then let's see.

Fairbairn: What was their charter versus the western group?

Fuller: Well, their charter was actually by being fairly closely co-located to the software and other product development groups in the Boston area they were looking at how they might collaborate on some of the computer systems design or some of the software design. It was more directly connected to near term DEC products versus the further out for our products. So, then system research lab went forward. About a week after the meeting with Kerns and Olsen, Wyn Hindle who was second in command at that time to Ken Olsen came by to meet me and he said, "Well, Sam," he says, "we understand the logic for doing this. This is the right thing to do," he said, "but, I do want you within a year to have the majority of people at SRC to have come from places other than PARC," which was great because I was just given a hunting license to just go hire more people which was great.

Fairbairn: Go hire more people.

Fuller: So, that was a great development at the labs. And over time actually through connections at SRC and WRL, I met Patrick Baudelaire and we set up a research lab in Paris working on advanced computing structures based on FPGAs which became more popular over time. We did some original work there. We set up a lab in Karlsruhe, Germany working on Internet protocols, IPv6 was the particular focus at that point. But, maybe the one other thing that I think is worth mentioning in this is as time went forward and we move up at this point to 1995, I believe, the SRC began working on very large databases. So, there was a really bright researcher at SRC, Mike Burrows, who had come from Cambridge University in the UK and some of his colleagues there were working on the next generation of the Oxford English Dictionary and they had this enormous database. And they said, "Well, Mike, you're interested in databases, wouldn't it be interesting for you to figure out how to make this more easily searchable?" So, Mike thought that would be a good idea. It didn't, at that point, have a lot of relevance to Digital. But, it was a quite challenging problem in database retrieval and so forth. So, he began working on that and as Mike worked on that some others at SRC said, "this is a great way to search really large amounts of data." And they began applying it to searching their email systems. And then some individuals over at the Western Research Lab down the road, Paul Flaherty and Louis Moniersaid, "Well, the Internet was only an order of magnitude bigger than these databases that they were searching at SRC so what if we build a “spider” to crawl the Internet and we made that searchable database." And so, that was the genesis of Alta Vista. And so, let me check my notes here just to make sure because I actually looked up the date. So, I'm going to get this right because it seemed to be fairly important. I will get back to you, but I'm pretty sure it was the fall of '95 and not '94 that all the research labs were beginning to use Alta Vista. And the comment was well, they could pretty well find out anything they wanted to know. And even way back in '95. And so.

Fairbairn: And this is concurrent when Yahoo was starting up?

Fuller: Yahoo was starting up. And actually, through part of our university program, we donated some computers to Stanford for them to work on Yahoo. But, at that time I was thinking of Yahoo as sort of developing the table of contents for the Internet, but Alta Vista was the glossary. If you wanted to find out where something was you checked the glossary and it'd tell you what page it was on and show you the page. And it was all being done through this inverted text file system that Mike Burrows had developed for the Oxford English Dictionary. So, late in '95 we thought this is really great; we need to make this available to the community. So, I talked to the folks at Digital and they said, "Well, how are you ever going to make money out of this?" And we didn't know, so they said, "Well, Sam, it doesn’t make a lot of sense." But, then they said, "Well, look, we're not manufacturing anything." It really is just a service. So, all you had to do was just open up a port to make this available on a wider Internet. And the rest of the world can see this. So, they said, "Well, okay. If the research labs are prepared to support this. We're not going to support this inside Digital. You can go ahead and do that." So, we did in December of '95 if I got the year right. And it took off like wildfire. And so, my next marker was in March of April of the following year, '96. A PR person from Digital came by my office and he had this book with him. And this book was all the press clippings on Alta Vista. And the PR person said, "Sam," he says, "there's more column inches of press written on Alta Vista than any product DEC has ever announced."

Fairbairn: <laughs> And you hadn't even announced hardly, right?

Fuller: And it’s just been word of mouth. I got in with a lot of interviews. So, at that point, okay, we got to do something about this. And then at that point they did set up a business unit initially under Rose Ann Giordano and ultimately under Ilene Lang to try and figure out how to take advantage of it. But, actually to finish the Alta Vista story. Before they really figured out how to make money out of it, Digital was acquired by Compaq in '98. And the end of my story on Alta Vista is Digital was acquired by Compaq for $9,800,000,000.00. And within 12 months Compaq sold Alta Vista to CMGI for $2,300,000,000.00. So, about a quarter of the market value.

Fairbairn: The value of all of this.

Fuller: ...was this thing that started deep in the research labs, I believe, were only spending as much on the research labs at Digital as we were spending on aviation services at Digital.

Fairbairn: All one guy.

Fuller: So, for something that cost about $40 million a year for all the research labs, ended up being about a quarter of the market value of the company once it actually faded into history. The other great part about it is recently I was at Google and Mike Burrows and the group out of WRL and SRC that had originally developed Alta Vista and supported it, that team is largely inside Google now supporting the infrastructure for Google.

Fairbairn: So, they got to keep doing it.

Fuller: They've all ended up in good places.

Fairbairn: So, okay, so it was sold to CMGI and then it sort of disappeared after that? I mean, I forget exactly.

Fuller: Yes. CMGI didn't figure out how to make Alta Vista succeed and then the Internet bubble burst and they sold it. Ultimately, I think Yahoo ended up owning Alta Vista and folded it into something. But, it went through several sales. So, but the only numbers I keep track of is what Compaq.

Fairbairn: That it was a quarter of the value.

Fuller: Right, when Compaq sold it to CMGI. I would say the other important thing going on during those years in the research labs was continuing to develop connections to the research community in universities. And probably the largest effort there was what turned out to be Project Athena with IBM at MIT. And what happened there was I think it was in 1982 IBM and CMU had set up a joint project looking at distributed systems called Project Andrew. And when I realized that it got set up and I realized the quality at the work at CMU having come from there and the industry muscle of IBM, I said, "Well, this could potentially be a real threat to Digital because Digital fundamental was living on distributed systems based on the VAXs and the 11s and Ethernet and the network. And so, I got a hunting license from Ken who said, "Well, go find some comparable things you could do at universities." And so, I did end up talking to Stanford and CMU and MIT and the place where there really was a connection was with Michael Dertouzosat MIT. And Mike was a great leader of LCS, the lab of computer sciences at MIT. And Mike saw this as a chance to really do something great. And so, we talked through the better part of 1984 as to what became Project Athena would look like. Actually, in the end Mike convinced me that to really be effective it needed to be a more open system and actually IBM was another partner to join in Project Athena. But, ultimately as we got it set up both Digital and IBM, I think, contributed on the order of $20,000,000.00 worth of equipment. It was desktop workstations as well as back room computers. And one of the things that I certainly got involved with and we did right at that point with Project Athena was we had a number of Digital software engineers go down and become resident full-time down at MIT for the duration of the early years of Project Athena, as did IBM. The key things that came out of that collaboration, Project Athena, was the X Windowing System that was essential for DEC and other people as they looked at the early years of workstations, engineering workstations. And then also the Kerberos authentication protocols, which are still widely used, were also inventions that came out of the work at Project Athena. So, that was a great collaboration and that certainly helped move MIT to really be making more effective use of distributed systems on that campus. And it certainly enabled DEC to have a leg up as it moved into workstations and the distributed systems. But, we may be at a point now where I'd be talking about some of the final years at Digital.

Fairbairn: Yeah, I wanted to take sort of go through the state of things and where things were going right or wrong or whatever. And maybe we can take a break now.

<break in recording>

Fairbairn: I wanted to go back and talk about sort of the times when the VAX had been out for some time. It was very successful. They made an effort to create a LSI version of the VAX as well.

Fuller: The micro VAX.

Fairbairn: And somewhere in there and the Titan RISC processor, RISC architecture came up. Eventually it evolved into alpha. Tell me about given the benefit of hindsight and be inside at the time what were the critical missteps and so forth that led DEC to having to surrender itself to purchase by others and so forth?

Fuller: I'll give you my perspective. Other people may have a different perspective.

Fairbairn: That's why we talk, here at the museum we don't try to come to a conclusion, we try to present the various perspectives and let people draw their own conclusions. So, we're interested in gathering as many insights as possible and so all we're asking for is your insight, your opinion, your view as to what was going on.

Fuller: Okay. So, I think as you look through the 1980s as you say the VAX and ultimately the micro VAXs were quite successful as was VMS and then the networking software, DECnet were all doing quite well. The IBM PC came out in 1981. And PCs were in the process of developing. And maybe just one short story and then sort of my larger perspective on that. One of my friends at Digital was John Rose and ultimately John had the job of running the PC division, the PC product line for DEC. And over time, after John was running the division for a while, he ended up leaving DEC and going to Compaq. And at an industry meeting about a little over a year after he'd left DEC.

Fairbairn: Which was when approximately?

Fuller: This had to be the late 80s, '87, '88, I think that's about the right time. I asked,, "John, really sorry to have you leave," because I thought he was doing a good job developing the PC line. And he said, "Sam," he says, "I didn't want to leave DEC. I'd been there a long time, I was very loyal to the company," he said, "but, the third time after I got thrown out of Ken Olsen's staff meeting with my PC design saying that this isn't something that business would really want to rely on," he says, "just out of professional pride I had to leave," And he said, "So, I went to Compaq." They'd been trying to recruit him. He said, "I took the design that DEC wouldn't take onto the market place. I took that to Compaq." Basically they brought it out as the Proliant line of PCs. He said, "In less than 12 months the Proliant line was generating more than $1,000,000,000.00 annual run rate." "That was DEC's design." He said, "I just couldn't get it out the door." And so, I think what was going on in the 1980s was with all the success of the VAX and VMS and there were product lines in large organizations that had road maps carrying that forward and their careers were tied to VAX/VMS. They made the argument to Ken that they ought to continue that. And Ken listened to that. He was a strong leader. He would say about PCs that you wouldn’t really run a business on a PC and who would want a PC in a home. You really want to run a good timesharing system on a VAX. And UNIX began to really come forward as a dominant operating system. And he said, "UNIX, well that's a Russian truck. It's designed by a committee, you really are interested in industrial strength operating system like VMS." And from the research labs I could see the explosive growth of not only UNIX, but the TCP/IP protocols on the Internet. And said, "We ought to begin thinking about transitioning from the DECnet software stack to Internet protocols." And that didn't get much currency within DEC. And as I tried to step back and look at this I said, "The genius of Digital early on when Ken started it was he was bringing the power of computing out of the computer room, out of the mainframes to the individual scientists and engineers." The PDP 1, PDP 8, PDP 11, there the individual really got to use them. So, he was bringing the power of the computing to the individual.

Fairbairn: Those were single user systems, right?

Fuller: Right. They were. And the early work he did on whirlwind and the TX-2 and so forth. So, he was on the right track and in the early years, I think, with a lot of innovation that went on. Unibus was an open system, Ethernet, which we talked about, was an open standard. But, as you got into the 1980s one of the models of the company had become one architecture, meaning VAX, one software system, meaning VMS, one network, meaning DECnet, One Company. And so, it transformed from let's bring the power of computing to the individual to let's continue our product line of VAX, VMS, DECnet. And as the PC and x86 took over, and various variants of UNIX became popular and obviously as the Internet took over, DEC got pushed to the side. And the powers, this is I think, true of a lot of established companies, the powers within the company which were focused on the principal products of the company were more powerful than the disruptive technologies that were being talked about. And you always could argue well, but this year there's more revenue and more profit to be had with those than moving on to the newer systems.

Fairbairn: What about engineering workstations? Since they were such a major force in the engineering and scientific community did they see the MicroVAX’s being kind of the entry to that? Or was the engineering workstation a powerful idea?

Fuller: The hope was the MicroVAX, I mean, we talked about Project Athena and the workstations there were based on MicroVAX. But, if you look at the workstation, the engineering workstations, the dominant software system was UNIX. The tragedy is trying to look at engineering workstations going forward as part of the university program we gave what we called six packs of 11/750s to universities to develop single user software on them. And we gave one of those six packs to Bill Joy at the University of California/Berkeley; Bill Joy developed UCB Unix on it and ultimately went on to found Sun. So, maybe that wasn't the best of Digital's investments. But, we had the insight that we ought to do that. But, the company stayed focused on VMS and in reality we had Ultrix, which was our version of UNIX, but it didn't get the resources and the focus that VMS had and not surprisingly so. I think it's a common misstep that companies make as they stay too long on their existing product lines and don't realize what disruption is coming from below and with a couple of decades where Ken and the management at Digital had been quite successful following their own direction with the PDP 11, the PDPs of various forms, then when these other signals came along, there was these other disruptions going forward, they fell back on their earlier instincts. Earlier on we were right not to try and chase IBM and to try and go after these more open systems. And they lost track of that and ultimately got acquired. Ultimately when Bob Palmer took over from Ken Olsen in 1992 and came out with Alpha and, I think, had realized that there'd been some missteps it was late in the day. And DEC at that point, I think, didn't have the sales organization that could really drive DEC forward, turn this around, nor the customer base it had had a decade earlier. So, it’s a cautionary tale It’s trying to stay focused on well what’s the fundamental mission of the company and don’t get too tied up in the particular roadmap of a particular set of chips, processors and others things that are being sold today.

Fairbairn: Yeah, ironically as we’ve said, they were focused on single user kinds of systems, individual bringing the computing out.

Fuller: Absolutely were.

Fairbairn: Got sucked back into the more centralized computing resources.

Fuller: Right, and so it’s been a number of books written on Digital and the missteps in management, the missteps in the product line. In my case, in working in the research labs particularly all these labs were based on UNIX, we were the ones that were principle supporters of staying connected to the Internet and trying to see where that went forward. I talked about Alta Vista, but our voice wasn’t as strong as the divisions and the product lines.

Fairbairn: Yeah, the ones making money, right.

Fuller: They were making the money, right. Yeah, there was a great t-shirt that Ultrix guys, UNIX guys had which said "I work on UNIX, VMS pays my salary", you know trying to be respectful <laughter> of the main product line. But, yeah unfortunately DEC didn't make the transition.

Fairbairn: So, tell me about your transition during that period. What were you doing and what led you eventually to leave Digital?

Fuller: Well yes, it sort of is the ending years in and Bob Palmer came in and I think that was a new hope that things would be different, but it was very difficult at that point to turn things around. So, things continued to slide and so where I was in `95 or `96, I guess, independently of going to Digital I'd gotten recruited to be on the board of Analog Devices by Ray Stata, then the CEO and cofounder of ADI. He had realized that although they were focused on analog and data converters, AD and DA, converters. He knew the world had gone digital and he wanted some senior executives involved with the company that lived in the digital world to be able to provide some perspective. So, with support from Bob Palmer I joined the ADI Board, so that was going on in parallel, which I'll come back to in a moment. But, then by `96 and then in `97, but certainly in `96 various product lines within Digital began to get sold off. So, Palmer had sold the database group to Oracle where Mike Burrows and others had done some groundbreaking work, but that was gone. We did a lot of work in networking, particularly at SRC, the networking systems lab that was part of the Western Research Lab and some of their initial designs went into our networking division, but then that division got sold, I think it was Cabletron at the time. By early 1997 I realized that Digital was being taken apart and so I began looking at where can I go and work on interesting technology. I began looking at a number of firms out here on the West Coast and I probably was on track to coming back out to California when in the Fall of `97 Ray said he knew things were not going well at Digital. He asked where I was headed and I said I was probably going to be somewhere else soon and he had said well why don't you become our CTO or VP of R&D and so ultimately that's what I did.

Fairbairn: Now, you had no analog background?

Fuller: Right, no, my last analog was probably, I don't know, a sophomore at the University of Michigan <laughs>, but to Ray's credit he knew that he needed help shift the balance and so this would be a way to do that. But, when I agreed with Ray and by then CEO Jerry Fishman in September or October that I would like to do that, but I said I'm only a couple of weeks away from having my 20th anniversary at Digital, January 20th. So, I said I would like to stay till the 20th and then we'll make the change. He thought that was fine. So, I actually had an appointment with Bob Palmer on the morning of January, the same morning that he called this conference all with all the Vice Presidents to say that over the weekend he had concluded a deal to sell Digital to Compaq. So, I never had the meeting with Bob that morning because of all this turmoil going on and then, Jerry Fishman, at Analog Devices said well, "Sam you may want to stay at Digital because things will be different”. So, I thought about it for a day and I went back and said "Jerry when I joined Digital the company was investing 12% in R&D. By the time it got acquired by Compaq we were down 6or 5%in R&D. Compaq invests on the order of, I think, 2 to 2-1/2% on R&D, so I don’t think it’s going to be better.” I said Analog Devices, their targets are somewhere between 15-20% in R&D so I said I think there’s more chance to have an impact with technology by coming to ADI than staying and seeing what happened at Compaq. So, then by early February I had left Digital.

Fairbairn: So, you came to Analog Devices as VP of Engineering or?

Fuller: As VP of Research and Development initially and then a couple of years later it turned into CTO. And, I would say that the principle job there really is trying to understand what our product roadmaps look like going forward. And, the reality is as you move from individual IC components of AD, DA converters, or other analog components, ADI already had a pretty significant business in 1998 in DSP, or digital signal processors. There was really mainly ADI and TI that were selling DSP’s. But, going forward as integration continued, SoCs would become more and more relevant to these various signal processing applications. And those SoCs, they’ll have converters, DSP processors, but a lot of other digital content, a lot of memory, a lot of digital interfaces and so helping ADI understand how do we move into the age of SoCs and ASSPs is largely, I think, my contribution, , not let us get locked in to staying focused on components. And, like at Digital there’s enormous business in these traditional components. They will have a long life.

Fairbairn: Yeah, those products do tend to have an exceedingly long life.

Fuller: They will be great high margin products and we ought to have businesses that run these, but that shouldn’t blind us to the fact that to have growth going forward we’ll need to be involved in more integrated solutions, which are going to have a considerable amount of digital as well as analog and mixed signal content and an increasing amount of software will have to go along with those systems on chips. So, I sort of got into ADI, but.

Fairbairn: So, you joined them in `97.

Fuller: , `98 .

Fairbairn: Ninty-eight so you’ve been with them 15 years.

Fuller: Fifteen years now, right.

Fairbairn: Almost as long as you were with DEC.

Fuller: Well almost, another five years to go, but yes.

Fairbairn: That’s a whole new career.

Fuller: Unless it’s a new career, but at ADI as well as Digital I’ve been responsible for the university interactions at the places. And, remember I talked about Project Athena at DEC and there is a considerable number of pretty serious interactions now at ADI with university connections, both the major ones are at MIT and Stanford, but there’s a number of other schools that have got some fairly advanced research and advanced converter designs that we’ve also got connections with. So, let’s see I think it’s worth mentioning it took some considerable efforts both at Digital and ADI where the work is outside the company per se, but with the various industry consortiums. Back in 1983 I was in collaboration with Gordon Bell back at DEC. The worry was with what was happening with the Japanese fifth generation systems and how we were going to respond to that. And, ultimately out of that came Sematech and MCC (Microelectronics and Computer Consortium). I was pretty deeply involved in MCC. I was the Board Representative from DEC to MCC, less directly involved with Sematech then, more involved now, but less involved then. So, I sort of watched those two consortiums move forward. When I compare MCC in particular to IMEC, which is the European consortium, you know you see the differences between the organizations and what’s worked and what didn’t. MCC got off to a good start and I give part of its credit to Bobby Inman who said he really wanted first-rate computer scientists, computer engineers to help build the technology at MCC, yet as he did that, he was fairly difficult and fairly strong in rejecting potential assignees from the various companies that supported MCC to send their engineers to MCC.

Fairbairn: Is that because he didn’t think they were good enough or just didn’t think they represented the right direction?

Fuller: I think he thought that they were the second string, which is maybe some truth to the fact the companies didn’t want to let their principle designers go `cause that’s how they were doing their products, and he wanted the first string and so he hired many from the outside, which made it somewhat more difficult to get the right kind of collaboration going forward. So, that was one of the difficulties, but I would say there were some interesting technologies developed at MCC. But then, the problem MCC had was transferring those ideas back into the companies and by not having as much connection as they needed it became difficult to transfer it back. Ultimately, Bobby Inman left, after one or two other CEOs, Craig Fields headed up MCC, very bright guy. He saw that there wasn’t a lot of transfer occurring. He saw the support from established companies like Digital beginning to be reduced and his reaction was to take some of the technologies and move them into startups and the reaction of established companies was “so we’ve been funding MCC for all these years and the net result is that the technology that's worth pursuing commercially is wrapped up in and put into a startup, but then those startups become competitors to these established companies. So, that caused, I think, some of the support from the companies to erode even further and ultimately I think MCC, at this point, has become somehow associated with the University of Texas, at Austin. I’m not sure what’s happened to it, but that is what I recall.

Fairbairn: So, just briefly summarize what was MCC, what was its goals, I forget. I know what Sematech was, but I think for the purpose of this session just to sort of have it on the table.

Fuller: So, MCC was founded really in response to the Japanese fifth generation.

Fairbairn: And, it stands for?

Fuller: Microelectronics and Computer Consortium, and it was focused on those research avenues that back in 1983 people thought would be the future of computing. So, there was some focus on microelectronics for the continued advancement of Moore’s Law, a focus on computer architectures for what the subsequent generations of computers would be, a focus on artificial intelligence because the belief was that that was going to be important. MCC had programs in all those areas.

Fairbairn: And, it was funded by these companies as well as the government.

Fuller: Exactly, yes.

Fairbairn: And, Bobby Inman was part of DARPA [Defense Advanced Research Projects Agency]?

Fuller: Bobby Inman was, a former Director of the National Security Agency. I think he was still somewhere in the intelligence agencies at that time that he was recruited, a very bright gentleman, very effective recruiter, although he set the bar pretty high. I told you. He rejected a lot coming out of the companies. Very aggressive individual, and the Japanese fifth generation kind of faded as did MCC, and more recently at ADI I’ve made a couple of trips to IMEC, the research consortium in Europe in Belgium. They’re actually being successful. They initially had support out of the European governments, but now have support from many companies in Asia, Europe, and the US.

Fairbairn: Yeah, I was wondering. I remember IMEC being around, but there are no significant computer companies in Europe.

Fuller: No there aren’t. There are some great researchers in Europe and so I think the genesis of it is just like the fifth generation in Japan, just like Sematech and MCC in the US, Europe wanted to respond with some research initiatives to really trigger innovation in Europe and so they had a program, what do they call it, Esprit?

Fairbairn: There was Esprit.

Fuller: It’s gone through several generations of funding and I believe IMEC was initially funded out of one of those European initiatives and I don’t know all the history of IMEC, I think it would be interesting to sort of see. But, when I went there recently and realized that they had over a thousand researchers working on very advanced technologies, they reinvented themselves as an international research organization. They have a lot of support from TSMC and other companies in Asia, as well as, North America -- IBM and Intel -- and then ST Micro and others in Europe. So, they had found a way to actually get an effective research organization that didn’t gel at MCC. Actually, the one other short thing I might mention and we can sort of see where you want to take this. The other times in Washington was spent in the National Academies in CSTB, Computer Science and Technology Board, looking at various issues related toward the advancement of computer technology. One of the more significant ones was involved in cryptography and this was at the juncture of the Internet taking off, and you wanted to be able to have commerce on the Internet, yet the government was constraining anything stronger than 56 bit code, encryption. None of that could be exploited. So, the question was, well, what should the country do about that. So, we had a study and came to the fairly obvious conclusion that time was long past to end export controls around 56-bit DES encryption. You shouldn’t open it up to at least 128-bit key encryption. We wrote a great report, some great folks were involved with it. I remember giving a briefing to the White House and John Deutch, who was the Director of the CIA at the time was there. He absolutely rejected it <laughs>. Could not do that, and fortunately I think there was the change in the Clinton Administration and Deutch went back to MIT and the next Director of the CIA came in. The staff members had all supported moving it forward and they briefed the new director on that and within weeks he then did accept the fact that it went forward. I mean DES had to go forward (and latter even stronger encryption) for commerce to really be able to flourish on the Internet. So, that was a great involvement and in the last National Academies report, which we held here at the Computer Museum was on the future of computing performance. The fundamental view is that if you’re looking at Moore’s Law going forward it still has some legs left in it. But, back in 2004, the ability to continue to scale it to higher performance per individual processor stopped because of the underlying physics as defined by classic Dennard scaling reaching its limit. And in order to double the performance you basically had to double the amount of power consumed and so at that point you had to begin moving to multicore processors and the fundamental challenge of are we going to be able to develop software systems that take advantage of multiprocessors, not just dual or quad processors, but 64, 100 or 1000 processors. Now, we have bright folks like some of the folks from Google developed the Map-Reduce algorithm for use. There are clearly classes of problems represented by things like Google Maps and others that can really make effective use of tens of thousands of processors. But, at least so far, that’s a fairly restricted set of applications and so that fundamental challenge is this most recent study tried to address with the National Academies and so that’s, I think, been a pretty significant and important part of what went on in the sidelines as I tried to work primarily within the companies.

Fairbairn: So, is there some more you can tell about— you know you’ve been at ADI now for 15 years and you talk about the evolution of that company and where your major impact or focus has been in that regard. Maybe you could dig a little deeper into that part of the story.

Fuller: Sure, as I showed up at ADI I would often tell folks looking at the ADI DSP’s that they were selling at the time that geez here we are at ADI, but we’re selling these Sharc processors for $5 and they basically have the computing power and the memory capacity on a chip for what we were selling for $250,000 VAX11/780’s back in 1978. So, for the early years it was just getting a better understanding of how do we configure the subsequent generations of these DSPs and take them forward. But, fairly quickly we began to see that as Moore’s Law continued, customers were looking for integrations of frontend electronics that connected sensors to these digital processor into larger ASSPs (Application Specific Standard Products). And, so, now was the beginning of how do you begin doing integrated circuit designs that effectively include much larger amounts of a subsystem than they had before. And, for several years it really was looking at how things were being done with the larger digital SoC chips because they’d addressed large scale integration already. They developed the CAD tools and the methodologies to be able to develop these chips with hundreds of millions of transistors. So, one was bringing in the methodologies that were being used in digital processing into these mixed signal chips and getting the right on-chip infrastructure to be able to do that. Now, as we’re more recently looking forward we have to begin developing the collateral software that goes along with these chips so that it actually solves the problem the customer wants solved. You know, Digital said “we manufacture the iron and the software comes from heaven” <laughs> and I sometimes hear a similar argument at ADI which says “we’re really a silicon company and we give way to software to sell the silicon”. At the end of the day, I said you need to match what you charge the customer with the value delivered. And, if more and more of the value is in the software and you give that away for free you’re getting an imbalance which is not going to end well. So, we began looking at that. I say one of our most aggressive products right now is the front end of wireless base stations. It takes it from the RF signal through the analog and data conversion through the front end digital processing that happens at very high performance rates before it feeds into FPGAs or DSP processors. So, it’s a brilliant piece of mixed analog/digital design, but it’s got over a thousand registers that allow you to configure this for various different base station configurations in Asia, or Europe, or US. And, the challenge we’ve got is developing the software to be able to let our customers readily and effectively configure that and optimize it for their particular applications. And so, now we’re in the process of building up some software capability and system capability so as we go forward these ASSPs get evermore complex we’re able to benefit from that. Much like the comment I made about Digital, which was if the genius of Digital was bringing the power of the computer to the individual and we sort of got lost into VAX and VMS, but the power of ADI is connecting the physical world to the digital world of digital communications and computing. And so, if it’s building a signal processing chain that goes from the sensors into the digital world we need to look at whatever needs to be done within that signal chain. And, as things get more integrated and more of the software migrates into the signal chain itself, we’ve got to be prepared to do that. And, I give a talk each year to the engineering community at ADI of what the future challenges are. And, what I tell them is we’re really about to enter what I think of is the fourth generation of the information communications and computing technologies. If the earliest one was mainframe computing, over isolated, and the second generation was around time-shared systems like many computers such as VAX/VMS. The third generation was around PCs and cell phones, personal computing. The next generation is going to be around what people either call machine-to-machine communications or the Internet of Things, the vast majority of things that are going to be on the Internet are going to be sensors that feed back into various databases where you’ll have this large— what people call these big data challenges of how to take advantage of this massive amount of data that’s being collected. With those sensors and signal processing that goes onto those, so it’s appropriate to feed that data into the Internet and into these databases that are being developed is exactly the capability that analog devices has. And so, if you look at that larger picture and figure out how do we effectively do that. That does mean then in some cases we need to develop the sensors that go along with those frontend interfaces that ultimately feed the data in. And so, sensors we’re developing now are accelerometers and gyroscopes and pressure sensors and microphones and they’ll be additional ones going forward. So, I think that’s our larger challenge and we’ve got to be ever open I think to disruptive opportunities that might enable us to do that more effectively than we have in the past. This past two years, I’d say one of the more meaningful things that have happened is I was involved and ultimately helped describe to the company why we should acquire Lyric Semiconductors, a small 30-person semiconductor company down in Kendall Square coming out of MIT, looking at a very different way of doing signal processing. That lab in Kendall Square is developing a set of what will hopefully become potentially future very interesting products for ADI that are able to operate at much lower power. A lot of these distributed sensors will be at remote locations, not easily powered. We’re investing in some other research programs with universities and internal ADI looking at various energy harvesting methods so that if you’ve got these distributive sensors how can you collect energy from the environment that lets you continue to be able to run those sensors so that they can continue to feed that data back into the Internet and into the big databases that other various companies or others will use going forward. So, there’s a great challenge out there and I think part of my job is to try and in inspire people to look at the bigger picture rather than the small picture and try and look at those disruptive things that might be a real engine for growth

Fairbairn: So, in the R&D organization what percentage of the people would identify as hardware designers versus software people today roughly?

Fuller: Roughly, yes. Well the way I describe it to folks is if you look at our HR database at ADI, we’ve got about 3200 engineers. However, if you look at the engineers which have tools on their desk for either doing hardware or software design, there’s, I think about 1500 desks or engineers that are doing hardware design and there’s about 300 or 400 engineers that are doing software design and I think the rest are probably managers that are working with those designers.

Fairbairn: So, it’s still like two-thirds people that identify as hardware?

Fuller: So, it’s still two-thirds or three-quarters that were the hardware design. That’s right. And, if you look forward, there’s clearly a shift to more and more of that value being software — while I say software in reality if you look at the signal processing task a lot of it really is in the algorithm design. And so, some of that algorithm will get represented in software, some will get represented in hardware, but if you’re really in the business of doing signal processing what you need, I think, are people that really understand how to design algorithms that take however noisy, however difficult that signal is that’s coming in from the environment, from the real world, how do you process that in a way that becomes meaningful so that your computers can actually pull information and knowledge out of that signal. So, there is a growing number of of applied mathematicians that we’re employing to look at advanced algorithms. And, then as a second step saying what is most effectively done in software or what is most effectively done in hardware.

Fairbairn: So, on the hardware side just for my own curiosity, what is the current range of semiconductor technologies that you are designing for, you know, line widths, whatever, and what are you using for Fab resources?

Fuller: Right, that’s a great question. Analog Devices is quite different from let’s say Intel. You might think of. Intel’s as multiple Fabs (semiconductor fabrication factories) running one process in their copy-exact model producing enormous numbers of microprocessors. ADI has two Fabs, one in Massachusetts and one in Ireland and the one in Massachusetts has 200 different process variants for the various different analog or mixed signal devices from Operational Amplifiers through instrumentation amplifiers to voltage references to AD and DA converters, whether it be bipolar, complimentary bipolar, or CMOS and a similar large number of processes over at Limerick, Ireland for higher voltage power management processes and other things. There are these two Fabs internal to ADI, but we also have TSMC in Taiwan, which is a very important partner for ADI and a large number --about half -- of our products or about half of our wafers really come out of the TSMC advanced Fabs, so we go all the way from recently taping out some 28 nanometer CMOS parts with TSMC right up through 0.35 micron, 0.5 micron, and even 1 micron parts in some of our MEMS (MicroElectoMechanical Systems) devices and some of our high voltage devices. So, it covers this very wide range.

Fairbairn: So, are these 200 process variants running in your Massachusetts lab, for example, is this all computer controlled you run wafers sort of intermixed almost?

Fuller: Right, it’s a classic job shop and I love it when I go talk to manufacturing engineers of how do you do this. Oh well, we run on the Promise system that runs on VAX. I said I want to go see these VAX’s because this is 2012, 2013 and these VAX are not longer made or supported. So I go and, well there are no VAX on line. I mean they’re actually HP machines and what they actually mean is these HP machines are running VMS and Promise is an old job shop scheduling program, which runs on VMS. They needed a hardware platform to run VMS.

Fairbairn: So, it is the emulator or whatever that runs on the HP machines that runs VMS?

Fuller: Well, the VMS group ported VMS from VAX to the Alpha so they were running Alpha and then ultimately alphabetic ware by Compaq and then Compaq got acquired by HP so then HP ported, and I don’t know how good the port was, it was capable, but when they decided to end the life Alpha they ported VMS to an HP machine. And, I don’t know actually to be honest whether it was Itanium or whether it’s some older HP architecture. But, the guys on the manufacturing floor call it VAX, but <laughter> what it is is this one application, Promise, which runs on VMS and they run it on some hardware platform.

Fairbairn: Yeah, I remember in the 1980s I was at VLSI Technology and we were running Promise.

Fuller: Oh you were.

Fairbairn: Yeah. So, I haven’t heard of that name since then, but it’s like it’s still...

Fuller: I never heard the name until I went to ADI, but that’s where I learned about it. So, yes it’s a classic job-shop scheduling problem, but that’s used in Ireland and Massachusetts. But, as we move towards SoCs, or application specific large chips, these are all in CMOS and these are all fabricated at TSMC, or other foundries, in the finer geometries.

Fairbairn: So, how old is the Massachusetts Fab and are you going to be able to keep it going technology wise and so forth?

Fuller: Well, it’s gone through many revisions and they haven’t retired some of the older devices, or modules, in their equipment because it’s supporting old products, which continued to have very long lives. I mean just basic op amps. But, they have newer modules for some of the MEMS devices, some of these high performance Milli G gyroscopes that are available, or Milli-G accelerometers. And so, they continually bring in the new equipment and they continue to add, it’s more like dealing with a New England farmhouse, you know, they keep adding additional outbuildings to support the different modules as required. How long this continues, well it’s hard to say, but the products continue to be in real demand and have high margins so they’ll continue to be very well supported by the fabs. I’d say again the growth will probably be in the more advanced CMOS whether it’s with TSMC or working with other large foundries is the way we’ll be getting a lot of the more advanced, more complex products out.

Fairbairn: All right, well to wrap up is there anything more you can say about sort of the trends. I mean you’ve talked in general about the trends and maybe there isn’t a lot more to say about what you see happening in the AD/DA world, you know, in sort of the global computing world that we live, which includes real world inputs and analog and signal processing and digital computation all interconnected. Is there any other?

Fuller: Well, in many ways there’s a great future out there when we think about the applications that might get developed. I think the challenges that we face, one of them I mentioned briefly before, on this National Academy study when we realized that Moore’s Law can continue, but the Dennard Scaling and the classic Dennard Scaling which let you double the performance every couple of years at the same power level. That ended back more than five years ago and so now as the industry has come to expect maybe two orders of magnitude increase in computing power over a decade and that’s driven a lot of these wonderful applications. Well to get a factor of 100 increase in computing capability over a decade it is not going to be possible to increase the power consumed by a factor of 100. I think five year’s ago, it was estimated some several percent of the power in the US is running the various computer farms of Amazon, Google, and all the rest. Well, if you really want a factor of 100 increase you can’t use 100% of the power in the US to do that. Now, people are surely making some advances. There’re lots of inefficiencies in the way we build the computer farms. There’s even inefficiencies in the way we develop smart phones and others, so there is going to be factors of 2 or 4 or 5 that will help us increase computing power, but it’s not the Moore’s Law: 40 to 50 years of exponential growth, every 18 to 24 months a factor of 2. So, how are we going to address that? And then, concurrent with that even as we look for ways to move this forward we’re doing this with multiple processors, multiple cores per chip, and can we really rethink the way we do programming so that we can make effective use of 10s and 100s of cores rather than single cores. There’s individual examples, like MapReduce, that Google has developed. How widely can those be expanded to coverage all sorts of applications? So, those are two, I think, pretty serious challenges in front of us. And, I don’t think we’ll hit a wall in the next five or eight years, but if you think out 10 or 20 years we need to have something different or otherwise we’re going to begin to slow down.

Fairbairn: Yeah, we’ve been looking for that solution for multiprocessors for a long time.

Fuller: Yes, and people thought about that a lot. People said well if you’ve been working on multiprocessing and you haven’t cracked that over the past three, isn’t that going to be impossible. And, I think the positive thing you can say is well the problem with popularizing C.mmp and CM* and all the other multiprocessors that are coming forward is that by the time you built that multiprocessor artifact, Moore’s Law had taken the single processor ahead by a factor of 2 or 4 or 8 and so single processors were actually providing as much or more power than these multiprocessors and so that took away the incentive to figure out how to take advantage of multiprocessors. Well, the fact that single processors will now not continue to outstrip the power of multiprocessors. That’s gone away and so now there is a real incentive and so I think there is hope, with the right innovations and the right insights, we can make effective use of multiprocessors. These are my two sort of cautionary points, but on the other side of it, this ability to build these distributed sensor networks, I mean there’s nothing stopping that. And, you’ll be able to build these millions and probably over time billions of distributed sensors and they’ll probably be self-powered by various energy harvesting techniques so you will be able to have knowledge of an enormous amount of your environment. There’s lots of interest in either greener or more efficient use of our power. One of our projects going on at MIT right now is building a portable wearable vital signs monitor that people can wear behind their ear much like you’d wear a hearing aid today. And, you would be able to get six of your vital signs monitored on a continual basis wirelessly sent back to a database so that people that need continual care can see how their heart is behaving on a continual basis and those of us that are healthy, it’s a way to monitor your health and maybe do things to try and improve your health with exercise and so forth. So, there’s a lot of ways that technology can make an enormous positive difference. Then obviously there’s a dark side as well that we need to watch out for.

Fairbairn: Right. All right, well Sam thank you very much for spending the last few hours talking about your career and dreams of the future and it’s been extremely valuable.

Fuller: Well, Doug, thank you for your patience and willing to listen to all of this. I enjoyed it immensely.

Fairbairn: Thank you very much.

END OF INTERVIEW

