Oral History of Ken Poulton

Interviewed by:
Günter Steinbach

Recorded via Zoom May 13, 2021
Palo Alto, CA

CHM Reference number: X9479.2021

Steinbach: This is the oral history of Ken Poulton recorded on May 13th, 2021, for the Computer History Museum, but not at the museum because of the pandemic. I'm Günter Steinbach. So hello, Ken, and thank you very much for donating your history to the museum. I wanted to do this interview because of your work on massively parallel ADCs, but this history covers all of your work and career. Let's start with the background, like where did you grow up? What's your family background? Your childhood hobbies, maybe?

Poulton: <laughs> Okay, so my parents grew up in Piedmont. It's a small town enclosed in Oakland, and they were actually high school friends, not high school sweethearts. My dad became an orthodontist and college professor, taught orthodontics. My mom was a social worker and full-time mom, and then teacher. My wife Kate and I both grew up in Oakland, and we were born in '56 and '57, so we were growing up in the '60s, the early '70s. It was a little schizophrenic living in Oakland at the time because we were both right at a lot of the turmoil of the '60s. The Black Panthers were organized in downtown Oakland. The demonstrations and riots in [UC] Berkeley were going on just a few miles to the north. Kate actually had tear gas drift into her elementary school from some of the Berkeley demonstrations one time. But at the same time, we lived in the Oakland Hills, which was very white, very de facto segregated. We had school busing going on to do desegregation, but our schools, aside from the busing, tended to be relatively white. So it was schizophrenic in that way.

Poulton: So in high school, I was on the ski team for a while. I did a lot of the school plays. But the peak of my drama career was playing Duke Orsino in "Twelfth Night", and Sir Andrew was played by Tom Hanks. He always got the zany comic parts in our plays. I met Kate backstage during one of those plays.

Poulton: I was interested in architecture and chemistry, in particular, in high school, and I started programming actually in junior high school on a time-share computer. We had a Teletype terminal at our middle school, and it was connected to a computer at Lawrence Berkeley Lab. My mom remembers me writing out basic programs on paper during a family vacation or two, and we'll note that at exactly the same time Bill Gates was learning to program on a Teletype connected to a distant time-sharing computer. So we did exactly the same thing and our careers have very nearly paralleled exactly since then.

Poulton: Give or take a few billion dollars.

Steinbach: <laughs> You even had the same ruthlessness, huh?

Poulton: <laughter> I did not have the killer instinct that he developed.

Steinbach: So you did get into computer engineering at that early age. What about electrical engineering?

Poulton: I didn't do anything hobby-like in high school. One of my summer jobs in college, the first thing I did in electronics was working for a tiny company that made weather stations. This was doing PC boards, was the way they built the electronics for this. I got to do some PC board layout using I forget what the material was called. But the sticky traces that you lay down on Mylar in order to define the traces on these PC boards. So that was my first connection with layout, and that was kind of fun. It's not what really got me going into EE, though.

When I started at Stanford, I took classes in chemistry and physics, the classes that you would do for either a chemistry or physics major. Almost instantly I got turned off by chemistry because they started with organic chemistry, which was basically memorization. And the physics class, the introductory class for physics majors, emphasized that in each field you have a few basic principles and a relatively small number of equations from which you can more or less, in principle, derive the universe. Yeah, that was a pretty striking thing and pretty exciting. So I was doing my physics undergrad and as I went through that, it began to become apparent that I wasn't cut out for theoretical physics. I did lots of math, I was really good at math, but I wasn't good enough for theoretical physics. My physics advisor advised me to try out a class in electronics because for a working experimental physicist, having electronic stuff is part of what you need all the time. I was doing an introductory EE lab, and one of the physics major labs at similar times, and I was finding that the physics labs, they barely worked. I resonated very much with a quote from physicist Ernest Rutherford which was "If you need a statistician to interpret the results of your experiment, then you need a better experiment." That's really the way I felt about my physics labs. At the same time I was doing these EE labs, and you put together these components with 10% tolerances, and things worked, and you got something useful. I found that really appealing, and that was sort of the point at which I switched over to EE as my long-term role. I mean, the other thing that was appealing about EE was there was still lots of math. It wasn't at the level of theoretical physics math, and so it was a good fit for me. I finished my physics BS and there was a program where you could do a master's program called coterminal, basically overlapped with your bachelor's program. So I took an extra year and did both the BS in physics and the master's in EE, and that turned out to be a nice fit. I thought about doing a PhD and my advisor, of course, wanted me to go that way. But I looked at, again, what you can do with a master's in EE and the kinds of things you can really build and have a practical impact, and decided that the master's was really where I wanted to stop my education.

Steinbach: But did you get the master's?

Poulton: Yeah, so I got the master's. But I decided not to proceed with the PhD.

Steinbach: So what year was that?

Poulton: So that was 1980, got both degrees in 1980.

Steinbach: Okay, so then you must have joined HP right out of college, right? Because

Poulton: Right.

Steinbach: -last year I attended your 40-year anniversary. <laughs>

Poulton: That's right, that's right. <laughs> Right before COVID struck, and I remember being so pleased to see you. I rushed over and shook your hand and we were all we both realized right after we did that "Oh, maybe we shouldn't be doing that." <laughter>

Steinbach: Well, we didn't catch it from each other.

Poulton: Apparently not. <laughter> My 40th anniversary celebration turned out not to be a super-spreader event.

Steinbach: <laughs> Did you look at other companies also?

Poulton: So I interviewed with a bunch of companies, and the one I remember was Tektronix. They had really creative interview questions, and I very much enjoyed that. But what they were interested in hiring me for was to work on SAW devices, which was a summer job which had turned into a year-long part-time job during my last year at Stanford. I'd gotten this job at Ampex research. You have to be kind of old to remember Ampex. They used to be the leader in video recording, but the company is now gone. But they had a research facility in Redwood City, and I got a summer job there basically doing their first SAW devices. "This might be interesting, let's get a summer internship to try this out." I got to learn using Unix and C-programming to design these devices, because it's a relatively simple geometric layout, but there was some filter design involved. I also got to learn about what it was like working in a research lab that depended on government grants, and that made me run the other way from that kind of thing. Even though I wasn't directly involved with it as an intern, I could see these guys that I was working with spending a lot of their time churning out these grant applications that were two inches thick. Some small fraction of them would actually get funded, and that wasn't particularly appealing to me. So I knew I was looking for something that was more had more practical impacts. Professor Bob Dutton made a connection with me to HP, and HP Labs, just a mile from campus, actually in the Stanford industrial park. They offered me a job doing analog IC design, and that seemed pretty good, and forty-one years later, yeah, it seemed like a pretty good fit. <laughs>

Poulton: I started in 1980 and the first job there was working on track and hold circuits for building real-time ADCs to go into -- with the aim to put these into oscilloscopes. What was done at the time was essentially all analog scopes, and there was one digital scope product that was very low sample rate. Not at all the kind of real-time scopes that we use today. So as a department, one of our major thrusts was to try to bring scopes into the world of digital. The department at the time was run by Tom Hornak and he'd escaped from communist Czechoslovakia in 1968. They had a brief relaxation of travel restrictions, and he and his wife took a suitcase of vacation clothes and left Czechoslovakia and eventually made their way to the US. When I joined, I guess that was 12 years later, he was running our department. He was a great guy to work for. He had interesting stories and this really good grasp of everything technical, and a real vision of pushing what we were doing into HP products. That was I think the thing that was particularly great about HP Labs at the time, was that it was a place where you were doing something that was between research and development. The aim was always to make something that would eventually contribute to HP products. We were virtually never doing pure research, but we had the freedom to do things that were really long-term. I started in 1980 and the work that I did didn’t come out into a product until 1987. That’s the kind of time horizon that we were able to do there.

Steinbach: Great. What else have you worked on besides data converters?

Poulton: Relatively little, actually. <laughs>

Poulton: You actually gave me the question ahead of time and I forgot to research it. We’ve done data converters related to various instrument product lines; oscilloscopes, spectrum analyzers, RF sources, and we dabbled in things like image sensors. I’ve forgotten, there was a few other things that came along the way. But in terms of the major projects that I’ve done, and for that matter, most of our department has done, it’s really been focused on A-to-D converters, and for a very long time, serial data transmitters, transceivers. My work wasn’t directly very much on serial data, but for a while, about half of our department was working on that.

Steinbach: Okay, and I do remember you also doing system administration to some percentage of your time. How did you drift into that?

Poulton: So I got my feet wet in Unix back at Ampex in 1979, and that was something that I enjoyed working in. When I <laughs> got to HP in 1980, they were using these business computers, which were modeled on IBM business computers. Oh my god, it was just the most wretched computing environment you could imagine for technical stuff. But HP made them, so we were -- that’s what we had available. But by 1983, we started getting a few technical computers, and I think the original one was a DEC VAX within labs. I immediately got an account on that, but it wasn’t where I could do my real work yet.

Poulton: Actually, as I think about it, I ported a Unix-like programming environment into the HP 3000 [business] operating system in the first couple of years of being at HP. It basically provided a programming platform that was based on Fortran, because Fortran was what was available on virtually every computer, and a preprocessor called Ratfor, for Rational Fortran, that allowed you essentially some C-like syntax before Fortran had any structured programming capabilities. So I ported that to the HP 3000, and I used that for some of the -- quite a bit of the programming that I did. Basically, things like postprocessing SPICE runs. I’d have simulation data, and I’d want to be able to do some analysis of that, and this was a better environment for doing that kind of prototype programming. That turned out to be useful enough that I offered it to other divisions within HP, and not surprisingly, nobody was interested. So I got permission to actually sell that as a business of my own. So for a while, I had a business called Terminal Software, which I guess I ran until 1987. Basically, people who were stuck on the HP 3000 and wanted a more Unix-like environment, this was a stopgap that provided some of that support. So I sold that package. I think I sold it for like 150 bucks, initially. Maybe the price went up to $300, eventually. But that was mainly good because it was-- other people were actually finding what I had done there useful. Didn’t exactly make a lot of money.

Steinbach: <laughs> And another-- a little bit of a parallel to Bill Gates, huh? <laughs>

Poulton: Right, right, yeah. <laughs>

Poulton: So I guess at the time that I was doing that in the early ‘80s, he started this little company called Microsoft. I mean, god, I was doing something Unix-like, and you can -- I don’t know why he bothered with that DOS stuff.

Steinbach: Yeah. So I researched your patent and I saw two patents about mass spectrometers. How did that happen?

Poulton: So one of the things that was particularly neat about HP Labs, and to a lesser extent, Agilent Labs, was sort of the cross fertilization between different fields. At HP in particular, we had a company that had -- as HP Labs, we were serving a company that had a very, very broad portfolio of products. Basically, any of those areas were things that were fair game for us to consider working on. Time-of-flight mass spectrometry was something that we had some products in already, and being able to apply the high-speed analog-to-digital converters that we had was basically a spin-off of the work that we were doing for scopes. We didn’t actually do any A-to-Ds directly for that product line, but we figured out how to use those A-to-Ds for that product line. I should say that August Hidalgo is the one who figured out how to do that. We were basically providing support for his development within the mass spec.

Steinbach: Okay. Yeah, you mentioned Agilent, they had also -- you started at HP Labs, then HP spun off Agilent, and more recently Agilent spun off Keysight, and so each time you <laughs> were in the spun off part, as was I. What was your reaction to those spin-offs? Are you do you agree that they were a good thing? Or were they mistakes, in your opinion?

Poulton: So I really felt that the 1999 spin-off of Agilent from HP was a terrible thing. That we had a company that had provided some real synergy, and one of the things that came out of that was the optical mouse, that our department developed the navigation chip that enabled the optical mouse. Even though the original object for the navigation chip was a different application, it was a handheld scanner, handheld paper scanner[1]. That’s an example of the kind of synergy that we were able to provide, and the time-of-flight mass spec stuff was another. So I was pretty unhappy about their splitting, and basically, it narrowed the field of application for labs when they did that.

Poulton: It didn’t take very long to become apparent that well, at the same time as they did the split, they hired Carly Fiorina as CEO for HP. It did not take very long at all to become apparent that we were actually really happy not to be working for Carly. That she was, as near as we could tell, opposed to the HP way, really did her best to sabotage that, and was, in my view, one of a -- the beginning of a line of really disastrous CEOs for HP. The fact that we were not part of that was a really good thing because she probably would’ve just cut the whole instrument business if it hadn’t already been sold, or split off.

Poulton: While we were part of Agilent, we still had the chemical and biological things, and there was not a huge amount of things that came out of that, cross disciplinary actions, but there were some. When we

---

[1] [Interviewee’s note] And before that, linear tracking of paper movement in a printer or plotter.

split off as Keysight in 2014, I wasn’t too pleased about that either. But what that unlocked was the ability of the electronic management business to really grow and thrive as its own business. What Agilent had been doing was using electronic measurements as more or less the cash cow to fund biological expansion. It was a more quickly growing market, so presumably it’s rational from an MBA point of view, and that’s not a good thing from <laughs> my point of view. The result of the Keysight split has been that we are focused on our totally focused on the electronic markets, and focused on how we can grow, and it’s been quite successful. Really exceeded my expectations for the trajectory of the company. I’m finding that remarkable. It’s still going on at the moment. So that’s been pretty cool.

Steinbach: <laughs> Very good. Okay, I think we can get to the massively parallel ADCs now. As the first my first point is, so, as I said, I researched your patents, and there is none about massively parallel ADCs. Was that considered not patentable?

Poulton: We looked at it, of course, and we couldn’t find a dividing line between the massively parallel and the merely 2- and 4x parallel that would allow us to really patent the idea. It’s a pity, but <laughs> that’s the way it worked.

Steinbach: Okay. Okay, so, yeah, as you alluded, time interleaving of ADCs has been had been done before, including with the scopes, that you could parallel two channels to get one faster one. So how did you what made you propose, you and your co-workers, propose interleaving dozens of ADCs?

Poulton: So going back to my first project, which introduced as the one-gigasample digital scope in 1987, in that project, we interleaved four discrete -- four separate bipolar ADC chips. I made the GaAs track and hold front end that went in front of those chips. That was pretty much of a pattern that we followed in oscilloscopes up until 2000, that we would build the fastest bipolar A-to-D converter we could make in the processes that were available. Of course, the processes were advancing, and so that was a big part in what allowed us to advance the speed that we operated at. But basically, we would make the fastest unit A-to-D converter we could, and then we would interleave two-way or a four-way, and I think there was one scope product that ended up with sixteen-way. But basically, each chip was one, or later on, two-way interleaved on chip. That was sort of the paradigm that we, and for that matter everybody doing high-speed converters, was following.

Poulton: There was a development at Tektronix in probably the early ‘90s of an analog storage capture chip, where basically the idea was that you had a number of capacitors on a single chip, and you would store samples into those capacitors at a high rate. And I forget what rate they had, but call it a gigasample. Then you would read the charge out of those capacitors at a very low rate, and do the analog-to-digital conversion that way. We became aware of that because they were able to undersell our scopes pretty dramatically. In terms of how much high-speed circuitry you have, it’s a very simple system. We did some analysis and you could see just from their spec sheets, but you could also see from working with one of those scopes that you had a very limited storage, something like a thousand or fifteen hundred samples, and you had a bunch of analog artifacts that showed up due to the analog storage, and the analog readout process. It was clear that this was never going to scale up to the megasample or beyond rate that we really wanted to get to.

Poulton: James Kang and Jon Tani worked on this analysis, they were in our department at the time. They came up with the idea of “Well, could we use [CMOS]?”. We don’t want to do analog storage, but we’d like to use CMOS. What is CMOS good at? Well, it’s not very good at analog. The transistors are less accurate, and it’s not very good at going fast. The transistors are -- well, A-to-Ds were at the time 80 times slower in CMOS processes than we had in bipolar processes. Eight-zero. But what CMOS could do was give you lots of transistors, and so they said “What could we do with that?” The fundamental idea that developed out of that was that you optimized the data converter not for the highest possible sample rate, but you optimized the converter for the highest energy efficiency. You also need to make it relatively compact, and those two characteristics would allow you to put lots of these on a single chip. And the integration capabilities of CMOS were already at the level that we could put lots and lots of transistors on the chip, more than we had a need for, even in these massively parallel systems. So it was really this change of focus from the fastest unit converter you can make, and then we’ll somehow put -- interleave some of them together, and it was sort of a clumsy way that we did the interleaving, there, to really focusing on power efficiency as the main goal, and just put lots of them in.

Poulton: We talked with the scope division about that, and they said “Yeah, yeah, that sounds nice, it’s a good Labs project.” So unlike previous developments, they were not counting on us doing this, and it was risky enough that we weren’t counting on it working[2]. But we did a test chip with just one slice to prove out the current-based A-to-D architecture, and then we built a chip that had thirty-two slices on the chip, and would operate at four gigasamples per second. Four gigasamples per second was already the fastest rate of a bipolar A-to-D converter, and we went through that. Basically, we’re looking at doing this with memory on the chip. It’s sort of an obvious integration when you’re working in CMOS. But for the first chip, we decided that was more complex than we wanted to bite off. So we did a chip which was just parallel output of the data with no storage, and it worked. Our first silicon worked really very much, very close to the way we’d simulated it. We got that going on the bench and showed it to the guys in the scope division, and they said “Oh, really? Huh. I guess we better make a <laughs> scope around that.” So that fit into their mid-range scope line, but since it wasn’t something that they already had started, it took until 2001 for that scope to be introduced. We’d actually got first silicon on this chip in 1998, and it took until 2001 for that scope to be introduced, and so we didn’t publish this chip until 2002. While we were -- After we had finished that chip, and once they got started on the scope using that four gigasample chip, they looked at their whole product line, and there had actually been some consideration of whether we should still be competing with Tek[tronix] in the high-end scopes. We were, I think, not doing very well at competing, and they decided to really try to invigorate that high-end scope line, and said “What we need from you is a 20 gigasample per second A-to-D using this technology. We’d been thinking 16, and we really couldn’t goose it up another 25%. We ended up going to, instead of 64 slices, 80 slices. There’s a little factor of five that complicated a lot of things. <laughter> But we just needed that extra 25% in order to hit 20 gigasample per second. As it turned out, that was a really good call from marketing, that was going to in the long run, that was a much better seller than 16 gigasamples per second would’ve been. So we went ahead on that second chip. The first chip was basically a four-person project, which was already about twice as many people as we had on previous data converter chips. So Jon and James started that work but left the company, and Robert Neff, and I, and Andy Burstein, and Mehrdad Heshami built that

---

[2] We expected this architecture to take some iteration to become usable.

first 32-way data converter. The next one that was going 80-ways at twenty gigasamples was like a nine-person chip project, which is not that big a deal in CPUs. Teams were much larger already for CPUs. But for a data converter, that was the biggest project team. In fact, as far as I knew from publications, it was the biggest project team that had been done, which is <laughs> a dubious distinction, in some ways. But it points out sort of one of the things about the massively parallel architecture, is that it is just more complex. You have sort of the same level of complexity as a non-interleaved converter at the slice ADC level, and then you have the interleaving system, which has to have sampling, and clocking, and usually has some DSP, which is looking at all of your samples together. That is adding quite a bit of overhead and a lot of complexity to the chip. One of the things that comes along with interleaving are the artifacts that you have from interleaving, and when your ADC slices are not perfectly identical, and, of course, they never can be perfect, then you get interleaving spurs due to these mismatches. Basically, you can see that if you have an offset between the 32 different offsets in the 32 slices, that you're going to have a repeating pattern of offset voltages in the data that you recover. Similarly, if you have different gains or different sample timings from the many slices in your data converter, those cause different kinds of spurs. So figuring out how to do the calibration for those and whether to insert those corrections in analog offset correction, for instance, or in digital corrections became a major part of the development of these chips. That's a big part of why the teams had to get bigger in order to do them.

Steinbach: Okay, so the timeline was you got you taped out the first chip in the late '90s, 1998, you said, I think?

Poulton: Yeah, yeah.

Steinbach: And it became a product in 2001. And so that first one, did it already have DSP for correction? Or could you make them good enough within half an LSB, or something?

Poulton: So one of the things that we had to do in going to CMOS was we had to introduce calibration just for the unit ADC slice[3]. The gain in CMOS transistors is much lower than in bipolar transistors, and the offsets in CMOS transistors are maybe ten times worse than bipolar transistors. So you really had to rethink how you were going to deal with analog accuracy, and one of the changes we made was, in the bipolar ADCs, we'd really focused on making an ADC that was accurate. If it was an eight-bit A-to-D converter, we tried to make it accurate to eight bits without any trimming. In CMOS, that wasn't going to be possible, and so one of the things that Jon and James came up with was the idea of a current mode pipeline A-to-D converter, where rather than using voltage as the variable that you were operating on in each stage of the pipeline, that you were actually using current as the unit. Because the CMOS open-loop circuits could be more linear in current than in voltage. Another thing that we had to do differently than existing practice was, in CMOS circuits, and actually in a whole lot of bipolar data converters, closed-loop amplifiers were the norm, because that was how you could deal with the accuracy and poor linearity of CMOS transistors. In order to go fast enough, we didn't want to have closed-loop circuits, and so that was another reason why the current mode approach was valuable at the time. Then we went to reduced-radix architecture, where within your pipeline, you're no longer changing by an exact factor of two at each stage

---

[3] [Interviewee's note] This was not new in the industry, but it was new to HP ADCs.

of the conversion. You can think of a normal pipeline converter as long division, but in a base two. Basically, you're going to say "Is my voltage above or below the middle threshold?" and then take the residue from that conversion and do another conversion of "is it above or below the threshold". But each stage you're going to use a factor of two between stages, and what we did instead was a reduced-radix architecture where we actually had an intentional gain that was about 1.7, I think. Basically, we didn't rely on the gain of stages to be accurate at all. As long as it was somewhere between 1.6 and less than 2, it would be good enough, and we would do the correction in a digital circuit, which we called a radix converter, after the analog conversion. So that led to just within the ADC slice a digital support circuit which was correcting for all the offset and gain errors of the individual transistors.

Poulton: Have I answered the question? <laughs> I know I've gone a long way around here.

Steinbach: Yeah, I think so. <laughs> But the DSP, actually, that looks at everything, is that on the chip, or off the chip?

Poulton: Ah, okay. So the DSP for the radix conversion has been on the chip from the very first ones that we did. DSP that does correction for interleaving has been something that in scopes we've had outside the chip. Basically, in the environment of an oscilloscope, you know that you have a CPU in the box, and what we needed to provide on the chip was the insertion of the correction, whether it's digital or analog insertion. But the determination of the calibration coefficients would be done by software in the box. And that was a freedom that we had as designers, because we knew we were going into an instrument application, and not trying to sell to any kind of application where a CPU might not be available.

Steinbach: And you rely on calibration signals that go

Poulton: Yeah.

Steinbach: -to the input of the that are known signals that go to the input of the ADC?

Poulton: That's right.

Poulton: There are some kinds of calibrations that nowadays are done internally within a scope, and I have to confess that I don't know how all the calibrations are done in our most recent scopes. Last time I looked, there was a calibration output from the front panel of the scope that you would cable into the input channels one by one, and the scope would use that known signal. It at various times has been just a single signal, or sometimes it's been a series of different kinds of signals to do different kinds of calibrations.

Steinbach: So the user has to stop measuring for a while and do calibration?

Poulton: That's right, that's right. Which is what is known as foreground calibration, as opposed to the background calibration, where the chip, or the system, is able to do calibration while the user is getting their real data conversion done by the chip. There has been a lot of -- One of the things that CMOS converters and massively interleaved converters have stimulated is quite a lot of work in various kinds of calibrations.

Poulton: Background calibration has got a lot of appeal and has a lot of cases where it works really well, and then there's cases where it doesn't work so well. Because you either need to make assumptions about the user signal that's going through the A-to-D converter, or you need to be able to insert a signal along with the user signal. Background calibration always has some kinds of limitations, and so that has in addition to adding a lot of complication to your data converter. So we have stayed with foreground calibrations for quite a lot of our work for those reasons, particularly because we're trying to do things which are instruments which are general purpose and you don't know what signal the user is going to present.

Steinbach: How much work have you done on calibration issues?

Poulton: So I did some of the work on the original calibration that we did for the 32-way massively interleaved ADC.

Poulton: Mostly once we got that prototype running, then we were able to hand off that work to the scope division. Robert Saponas in Colorado Springs was somebody that we worked with quite a bit on that. I'm guessing that some of the code that he developed in 2001 probably is still living in some of our scopes.

Steinbach: <laughs> How many parallel ADCs do you have nowadays? If you can tell us.

Poulton: Let's see. I can't tell you because I'm afraid of giving the wrong answer. My involvement with the scope ADCs went through about 2005. Since then, my work has shifted to focusing on direct-to-RF DACs and ADCs. So I don't have in my head the details on what we're doing currently in scope ADCs.

Steinbach: Okay. And as order-of-magnitude or roughly do you know how many different Agilent and Keysight products have been or are using these massively parallel ADCs?

Poulton: At one point, somebody did a count and there were eighty-one scope products, this was several years ago; there were eighty-one scope products that used the original four gigasample per second ADC that we built. That was, it got paired with a custom data capture chip, and it started out in a mid-range scope and migrated into lower performance [scopes]. What were the low-end scopes as the performances came up to that level for what was defined as low-end, less expensive scopes. And the team in Colorado Springs was able to keep doing cost reduction on that, and it was really a remarkable thing that had a very long lifetime. Just sort of hit the right thing at the right time. I don't think any of the rest of our other of our A-to-D converters have had such a long life. But, basically, virtually all of our scopes since, I forget, about some time in the late 2000s have been based on the massively parallel architecture, various chips.

Steinbach: How would you compare time-interleaved ADCs with multi-core CPUs?

Poulton: They really reflect the same change of focus, that we made that change from “What is the fastest unit converter?” Or, in their case, CPU core you can make, to “What is the most power-efficient one you can make?” and then put a bunch of them on the same chip. Clearly we were driven by the same realization that CMOS was not going to provide massively higher clock rates, but it did have this whole other dimension that we could take advantage of. You gave me this question ahead of time, and so I went looking, and we built our first massively parallel interleaved ADC in ’97 and turned it on in ’98. It took us until 2002 to publish it, but that was two years before CPUs suddenly plateaued on their clock rate. So I like to think that we figured this out first. I can’t prove that.

Steinbach: <laughs> Very good. How do you see the future of ADCs? Do you think even slower ADCs will go to massively parallel? Or is this going to be really still more for the very highest performance?

Poulton: So at one point I was asked “When will every ADC be interleaved?” The easy answer to that is never. It’s basically another architectural tool that we as data [converter] designers have to work with nowadays. The problem of interleaving spurs is something that we were able to tolerate in oscilloscopes, and basically time-domain applications, also like serial data generation and capture. They’re not that sensitive to low level interleave spurs, where low level is sort of at the -60 dBc level. But when you go to the RF world, they’re a lot more sensitive to that kind of thing. So there’s a more careful trade-off that has to be made for, how do you use interleaving in RF applications? That will always drive quite a good fraction of data converters to not use interleaving because they just avoid that whole issue. That said, if you actually look at the specs of applications, there’s an awful lot of applications that can tolerate interleaving spurs at the levels that we can achieve. What we get in scopes is sort of at the -60 dBc level, give or take. We made an ADC for spectrum analyzers that got to -85 dBc. But, I mean, your spectrum analyzer can look down to 100 dBc pretty easily, and there’s a very strong cosmetic issue of “What the heck is that spur? I don’t want to see that in my spectrum analyzer.” Even if your application really isn’t sensitive to spurs at that level. We look at communications waveforms where we’re talking about 1% EVM [Error Vector Magnitude], which really means that if you’ve got something at -50 dBc, that’s already deep, well in your noise. But still, people will see that spur at -80 and say “Oh, wow, what the heck is that?” So there’s a cosmetic issue that I think is real in the sense of being able to use it in RF. Real in terms of marketing things for RF.

Steinbach: Okay. Okay, I think that’s all my questions about your work. What do you do outside of work?

Poulton: Well, before we leave this, I wanted to say something. You sort of got me to look back over my whole career, and I’ve been lucky to have my career cover several revolutions. One has been the change from analog scopes to digital scopes. We certainly enabled that, and the scope division, once that was possible, jumped on that big time. The movement from bipolar ADCs to the massively parallel CMOS ADCs has been a huge change for high-speed ADCs. And the movement from upconversion chains for RF to direct-to-RF, both for upconversion and downconversion is something that’s in progress right now. As I mentioned, because of spurious issues, it’s taking longer to happen. Again, that certainly won’t take place everywhere, but it’s something that we’re participating in, and it’s providing a lot of opportunities.

Steinbach: Okay. I guess and RF is going digital, so to speak, right? That’s

Poulton: Exactly, yeah, yeah.

Steinbach: Just like serial datacomm is going digital.

Poulton: Right, yeah. Just stick an ADC in front, at the front end. No problem.

Steinbach: Okay, so what do you do outside of work?

Poulton: <laughs> So, long ago, I used to do stage tech and stage management with TheatreWorks. After our first son was born, I found that I could do my job, and I could run shows with TheatreWorks, and I could be a parent, and oops; I could only do two out of three of those. So <laughs> that had to fall by the wayside. When my sons were a little older, I had been in cub scouts and boy scouts, and I found myself thrust into being a cub scout and boy scout leader for just 11 years. So <laughs> we did a lot of outings for that.

Poulton: Then around 1999, I worked with a local group that was working on fiber to the home for Palo Alto. The aim was to have the Palo Alto Utilities Department add another utility which would be a fiber to the home data network as an open network, and separate that from the ISPs. We actually got them interested enough to build a trial, and one neighborhood had a trial system I think in the year 2000. We had a plan that would cost about 50 million dollars and probably would've broke even in 5 years, and after that been part of the utility revenue stream that the city relies on. But "probably" wasn't good enough for the city council, they wanted something that was 100% sure thing. A few years later, they found their sure thing, they spent 50 million dollars on remodeling the or rebuilding the 5 libraries in the city, and that has a absolutely guaranteed return of 0. So <laughter> different priorities. Oh, and six months ago, I finally got fiber to the home from Sonic. Twenty years later.

Steinbach: That actually is not in my neighborhood, which is only a mile from yours.

Poulton: Yes. Well, call Sonic and ask. They did not, in fact this is running over AT&T's optical network. But it might be available in your neighborhood, I don't know. Fingers crossed. <laughter> Then as a recreational thing, I've been windsurfing since 1990, and kite surfing since 2001, when that took off, and then kite foiling the last 5 years.

Steinbach: What is that?

Poulton: Oh, you put a hydrofoil under your kiteboard. So you ride the hydrofoil instead of the board on the water.

Steinbach: Okay. And you can go even faster?

Poulton: You can go faster, or you can turn more smoothly. Different ways to enjoy that. Then a couple years ago, I got a Onewheel. So I've been using that Onewheel electric skateboard quite a bit during the winters. Last month, I got a wing foil.

Steinbach: A what?

Poulton: A wing foil. So it sort of looks like windsurfing, except you have an inflatable wing instead of a hard boom and hard mast sail. You do it on a hydrofoil also.

Steinbach: Okay. So you have a shaped kind of solid wing?

Poulton: Yeah.

Steinbach: Except it's not solid, it's inflated.

Poulton: Yeah, yeah, it uses inflatable leading edge, yeah.

Steinbach: Wow. And if I remember right, you have the scars and metal pieces inside you, to <laughter> to prove that you're doing it in a hard way. <laughs>

Poulton: I broke my right foot in one place, and a different time I broke my left ankle in four places. <laughter> Over the years, I've gotten three rides and rescues from the coast guard. <laughter> Your tax dollars at work.

Steinbach: Okay, okay, thank you very much for the interview, and good luck with windsurfing. <laughs> And

Poulton: I have to say one thing before we go. I alluded that HP Labs and our transformation into Agilent Labs, and Keysight Labs has been a good fit, but it's more than a good fit. As you know, since you worked in our department, it's this phenomenal group of people who are both really smart, and actually really good people to work with. That's been the thing that has been really the luckiest thing in my career, is to have this great group of people. You know, all this stuff that we've talked about has been the work of a lot of people. I've just been the one who's been spanned the most years to remain talking about it.

Steinbach: Yeah, I do agree that that was the best place I've ever worked, too, that department. Okay, thank you, and bye.

Poulton: Well, thanks. This was fun.

END OF THE INTERVIEW