AMD 2900 Microprocessor Family Oral History Panel

Participants:
J. Philip Downing
R.C. “Clive” Ghest
Anthony B. “Tony” Holbrook
Robert “Bob” McConnell
John R. Mick
Sven Simonsen
John Springer

Moderated by:
David Laws

Recorded: February 11, 2009
Mountain View, California

CHM Reference number: X5109.2009

David Laws: My name is David Laws. Today is Wednesday, February the 11th, 2009. We are at the Computer History Museum in Mountain View, California and we are going to record our memories of the development of an important product family from Advanced Micro Devices [AMD]. The project began back in the early 1970s, and had a major impact on the development of the high-speed computer business over the years. I'm going to ask each of you to start off by giving us your name, background, and very brief career overview, what your position in the company was and what role you played in the development of the Am2900 [bipolar microprocessor] family. Bob, let's start with you.

Robert McConnell: I'm Bob McConnell. In 1974, I was the design manager of the bipolar group at AMD. And I, in fact, stayed at AMD, with a brief excursion to MMI, for another 25 years. After that I did a number of other jobs in the development of microprocessors, both bipolar and eventually participated in the development of AMD's 386 and 486 devices. I ran the group that did that.

Laws: Thank you Bob. Tony.

Anthony Holbrook: I'm Tony Holbrook. I was with AMD for 21 years, from 1973 to 1995. I had the privilege of being the first managing director, or the business unit manager, for the 2900 family beginning in 1975, when we developed the first product in the family. And then continued on at AMD with a whole host of jobs in a number of different positions in the company. I look back very fondly upon that time.

Laws: Phil.

Philip Downing: I'm Phil Downing. I was the second managing director for the 2900 family. I think a lot of us worked at Fairchild before we came to AMD. I was the process engineering manager in 1974 - '75 and up until 1978, when Tony took another position with Advanced Micro Computers, was it? Can't remember exactly. But when Tony left, I became acting managing director. I was with AMD until 1989. Tony was responsible for bringing me over from a little company that we'd both worked at before called Computer Microtechnology.

Laws: Sven.

Sven Simonsen: I'm Sven-Erik Simonsen. I was one of the co-founders of AMD in 1969. In the early days, I was in charge of digital design engineering, later shifting over to being in charge of product definition and applications. I was with AMD from 1969 to 1982.

Laws: Clive

R.C. Ghest: I'm Clive Ghest and I joined AMD in 1970. I did a considerable number of jobs. I didn't quite sweep the floor, but did everything from designing circuits, writing test programs, data sheets, seeing customers, taking customer orders. That's how I started. Ultimately, I went into the applications side working for Sven. I did design of various circuits, logic design and the architecture, and ultimately the architecture of the 2901. And then later, I joined my friend, Bob McConnell, at MMI for a sojourn. I rejoined AMD first as a field applications engineer in Europe working for Sven, and then came back to Sunnyvale in about 1978, I suppose, and worked in product planning. And ultimately, became the VP of product planning, and then went into business development. I left AMD in 1979.

Laws: Thank you Clive. This is an interesting year, being 40 years after the founding of AMD [in 1969]. And 200 years from now, when someone is looking at this tape or reading the transcript, they'll probably wonder what was this AMD company? Sven as one of the founders, could you give us a brief thumbnail history of what AMD was, how it was founded, and how it developed?

Simonsen: The eight co-founders all had worked at Fairchild [Semiconductor, a division of Fairchild Camera and Instrument Corporation]. Fairchild was a pioneer in the development of integrated circuits and a major influence on the development of the industry in Silicon Valley] A new management team had come in, Les Hogan and his VPs from Motorola. And the result of that was [Jerry] Sanders at some point was let go. John Carey was let go. And the other six of us still worked there. We gradually got together and decided we wanted to start a company on our own, and it became AMD. The thing that is amazing to me when I look back now is that the total capital raised in 1969 was $1.55 million. When you think of these days what it takes to start a semiconductor company, with fab, etc., etc., that this thing could start at $1.55 million is a little unbelievable. But that was indeed the fact. We initially did second sourcing of Fairchild's 9300 series of MSI circuits and then later started branching out into more proprietary circuits.

Laws: The time of the conception of the 2900 product would have been '72, '73, maybe? Clive, I think you were involved at that point. Could you tell us what was going on in the microprocessor market at this point in time? And what is a bit-slice processor, and why should anybody care?

Ghest: Well, I think it was a very interesting time in semiconductors and computing, because I think a lot of people realized that ultimately if we could get the packing density that was required, we could put a computer on a chip. And so various companies were trying to build simple microprocessors. And the first company to do it, of course, was Intel. They had some custom programs going, 4004, the 8008. And they were surprised, even though there was some thought in the marketplace that it would be a good idea to have microprocessors, even they were surprised by the interest. And ultimately, they moved into building a better product than the 8008, the 8080. And around about that time when they were developing the 8080, I wrote a product planning document to senior management which was printed on pink paper because the secretary ran out of white paper. In that pink paper I advocated that we, AMD, get into the microprocessor business. Because we didn't have a lot of resources, as Sven said we started with a very small amount of capital, we didn't think that we could produce our own MOS microprocessor. So I advocated that we second source the 8080. But because we had a lot of expertise in bipolar, with the 9300 series and a variety of proprietary products, I thought that in the high-speed bipolar area, we could go our own way and build a bipolar microprocessor chipset. The difference between MOS and bipolar is the fact that bipolar is much faster. But unfortunately, it has a lower packing density, so you can't put so much on a chip. It also consumes more power. So it was inevitable if you were going to build a high-speed bipolar microprocessor in those days that it had to be several chips. And one of the elements which was, I think, the most interesting from the customers' viewpoint, was the arithmetic processing unit. Normally in computers in those days, the word length for the computer was anything from like, 16 to 32 bits. And you certainly couldn't put that much on one chip. So what you could do is to slice that unit into several pieces. And therefore one of the most logical ways of doing it was to build a four-bit slice. Then you could use several of them to build up an 8-bit or 12-bit, 16, or 32-bit microprocessor. So we decided on a 4-bit high-speed bipolar processing unit. And that became the 2901.

Laws: How much was this internally driven in terms of thinking “this would be a good idea,” versus reaction to competitive products in the marketplace or driven by customer inputs that they’d like something like this?

Ghest: Well, we knew of one company, MMI [Monolithic Memories Inc.] that was starting to develop a similar product. We managed to get some preliminary information, so at least we knew what that piece of competition was. And we had talked to a lot of computer manufacturers because we were visiting customers with our 9300 chipsets and proprietary AMD products. So we had a good idea how their computer structures were from Digital Equipment Corporation, Data General, Honeywell, and so on. And so they all had a lot of commonality. So what we tried to do was tried to build a chip which was very flexible, could be easily controlled by a microprogrammed system, and we could sell to multiple customers, both commercial and military.

Laws: So Sven, you were overseeing these activities at that time. What was the reaction to Clive’s proposal?

Simonsen: I would like to add if I may to Clive’s, that I think in product definition, one thing we always did when we talked with the engineering people in computer companies, telecom companies, was to ask them, “What would be the next level chip you would like from us?” And I think that was one of the instigations to how we started looking at if we could make a microprocessor in bipolar and how did we partition it. So the instigation came, I think, from customers. And of course the fact that microprocessors were in vogue and we matched that with the fact that we were good at the time in bipolar technology.

Laws: Was it influenced at all by Lee Boysel’s work at Four Phase [Systems]? He did an 8-bit slice in NMOS about 1969.

Simonsen: To my recollection, it was not.

Ghest: Not directly. I mean a multitude of people

were writing, and most people were going into MOS. And there you had the advantage of higher packing density, but much slower, much lower performance.

Laws: Did you find there was support for this idea in the company, or was it a maverick idea?

Ghest: I think for the people who were aware of what the customer [was thinking] - the marketing people and the technologists and the product planning people - I think there was great support. I don’t really know that there were any people who were adamantly against the concept.

Downing: I think there was a little apprehension on whether we could build it.

Ghest: Yes, I think that was the [key issue] -- yes. I think Phil is right. While not being the processing [experts], people on my side were saying, “Well, this is what we want.” And of course, the technologists who were developing the process were saying, “Well, that’s a lot of stuff. Can we do it? And can we do it and make money out of it?”

Laws: So how did you go through this process? Design little [test] pieces to see how much you could get on a chip?

Ghest: Well, I think the first thing we did was design the architecture. And we had an appreciation for the architecture of the traditional type of computers out there that were built with transistors and small scale integrated circuits. So we really knew if we were going to build a high performance product that it had to have a certain set of capabilities. And so I drew up the architecture and passed it by various other people in the company, and closely worked with Tom Wong and Phil. And we said, “Well, it’ll be a push, but we think we can do it. We think we can get it on a chip at reasonable size.”

Laws: Bob, was Tom working for you at that time?

McConnell: Yeah, Tom was working for me.

Laws: You were design manager at that time?

McConnell: Right.

Laws: So were you involved in working on this with Tom?

McConnell: Sure. There was a constant discussion. It was more like Clive lived in our cubes. And there was a constant back and forth on could we a design a chip and could we make it go fast enough. Other people had some positions on what the speed had to be. There was some competition that was sort of there with not so adequate products. And so we knew kind of where we had to be in terms of speed. So we cut the power back, and you still have to keep the speed. And that was the design trade off that Tom did a super job on putting together. And as I recall, we built it on an existing process while we were planning the next generation process so that it would finally be a reasonable size.

Downing: No, actually what happened was when we took at look at where we were going to build it, which was Fab 2.

Laws: Which was [producing] three-inch wafers.

Downing: Three-inch wafers. We were building low-power Schottky in there using a contact mask process. And it was messy. I mean MOS and bipolar were sharing the same fab. It was somewhat chaotic. And so we had to come up with a scheme where we thought we could get some die per wafer out of this. I think the estimates were like three or four die per wafer using the existing process. So we sat down and we came up with what was called a composite mask process where you'd put two masks in. And if you had a defect in one then you'd kill it with another mask. So there was a lot of double masking. But if you could group them in certain ways going through the fab, you wouldn't have to add double masks for every step. You could do a mask, which is a composite of everything, and then open up windows with subsequent masks. Then people were concerned about "You're going to have so many masks that it's going to be too expensive a process," and so forth. We had a lot of arguments about that. And [in answer to your question] "Did we work our way up to it?" my recollection is we just jumped in and did it. We didn't do any test circuits. We did some test chips with patterns on them. But we didn't do any intermediate circuits working our way from MSI to LSI. We went straight to LSI.

Laws: What [dimension] lithography were you working at that point in time? Two microns?

McConnell: Three I think it was.

Downing: Three micron.

McConnell: [That was] before microns were common in common use. It was still mils.

Ghest: It was better than inches!

Laws: And this was a Schottky process?

Downing: Low-power Schottky.

Laws: And roughly how many transistors could you put on a chip with that technology in those days and get any yield? You projected three die per wafer.

Downing: Bob, do you remember how many transistors?

McConnell: I don't know. But it was thousands and not tens of thousands.

Downing: Yeah, I thought it was like four or five thousand transistors or something like that.

Laws: In bipolar [process technology]?

Downing: In bipolar.

Ghest: So that was the challenge in the design. It was the fact that we want this architecture. But how can we design that architecture with so few transistors? So we worked very closely with Tom Wong. And so Tom came up with some innovative methods. I would say, "Look, I want this kind of function." And we would work together, and he would say "Oh, if I do it this way, I can save a transistor. And boy, that was important, you know. One or two transistors were important, that's how we felt. But he did a super job. He came up with the idea of this two-port memory, which is one of the good functions of the 2901, whereby you could get two pieces of data out of a memory at the same time, perform an operation on those two pieces, and then write it back into one of the upper end areas. I think that was unique at that time.

Laws: Tony, were you involved at all in this process? Or were you looking over their shoulders and wondering what was going on?

Holbrook: No. I was running a different group within the company at that point, the Interface Group, and having a good time. The group was successful and so on. I do want to go back to an earlier question you asked about whether there was resistance in the company or support in the company. I think it's important to recognize here a few things. AMD started with very few dollars, as Sven pointed out. And in order to get into the marketplace, was a provider of second source products predominately. And the desire in the company was to move into proprietary products, very strong desire to phase ourselves from being a second source provider to a proprietary source company. We were, I think from the executive management level, seeking ideas in order to do this. So I think there was an executive level support for this project from the very beginning, and I think Sven could attest to that.

Simonsen: Absolutely.

Holbrook: I mean we really were kind of desperate to find a way to put our own mark out there. And that was evidenced by the fact that we formed a separate business unit for this product line, before the first product ever came to market. But you know, all this work was going on to develop the product, but the recognition was there that it had good potential. So we formed a business unit around that before we even sold the first chip. So there was support internally.

Laws: So now you had a product definition and an architecture. You were working on the chip. Was there much contact with customers at this point? Did you try to build up any design wins before the product ever appeared, or was it kept under wraps?

Simonsen: We would typically, when we had an architecture or logic design, pass it by some of the key engineers of some of the key customers. Taking their feedback before we took the next step, or at least concurrent with working on it. And so I think we had some contact with customers. Would this be of interest if we could build it? And as I recall, we had very positive feedback that indeed, if we could do this at the kind of speed and prices we were talking about, they would be very interested in it.

Laws: Clive mentioned that there was an MMI [6701] competing datasheet, at some point in time. Was there any pressure to just go build that part?

Simonsen: Well, I think the G2 we had at the time was that they really couldn't build it. And I don't know if they ever shipped very many of them. But I think we felt on one side somewhat of a pressure, here was a part. But at the time, it looked like they really couldn't ship it. So maybe the pressure wasn't that heavy. It was more the [point] that Tony was talking about, be proprietary and to make a mark.

McConnell: I remember somebody, maybe it was even me, proposed to Sanders in a meeting that we should second source the MMI part as opposed to building a new architecture. And that ended very quickly. Jerry was very adamant that we were having a proprietary part. Like Tony said, the drive was for proprietary products. I don't think Sanders had invented asparagus yet. But this was the first seeds of those ideas.

Downing: But it was close.

McConnell: Yeah. Sure.

Ghest: But we convinced him that our part was better.

Laws: Now, of course, the family wasn't just the one part. There was the bit-slice, there was a whole family eventually. And a very important part was the sequencer, the 2910, I believe it was.

Downing: 2909.

Laws: At what point was that conceived? At the same time as the bit slice, or was that later?

Ghest: No, that was a later time. And John Mick is the expert there.

Laws: So at the point of definition, you weren't thinking of a road map with everything that would go along with it. You just wanted to get this arithmetic unit out, essentially.

Ghest: Yeah. We knew that was the core. I mean, originally we knew that customers would want additional circuits. But we knew this part would stand alone for a while and they could use their conventional techniques of providing the control for the arithmetic structure. So if we could do that chip, we could get a foot in the door and go from there.

Simonsen: I think it's fair to say that we didn't think 2901 by itself could be successful. But there was already the thought about a whole family that together would create a more total function than the customer wanted. So I think at least in the elementary sense, we had already idea of some of the next circuits in the family.

Laws: Getting a little bit more now into the product development phase. So Bob and Phil, you were working for Gene Connor at the time?

McConnell: Right.

Downing: Right.

Laws: Gene had all AMD bipolar [products] at that time.

McConnell: All the digital bipolar, right.

Laws: How was the product development managed? Was Gene intimately involved? Or was he hands off?

McConnell: I don't know. I think we were all hands on. I think Tom Wong was probably very unhappy at how many people were managing him, yeah.

Laws: What were the key design issues to be solved; I think we've discussed a number of those. How do you get the process available, how do you get yield? We've talked about composite masking. We've talked about a number of other trade offs. Were any really clever tricks in circuit design in order to make this thing work?

McConnell: Well, as Clive pointed out, the two-port RAM, Tom invented some very clever ways to minimize the number of transistors in the two-port. I think with multi-emitter transistors, if I remember right.

Downing: Yes, that's right.

McConnell: And Tom was just a super circuit designer. He used some fewer transistors per gate I think than most anybody could have expected at the beginning.

Ghest: So it was a very elegant, and I think also a very clean design. Later on, we talked about RISC processors. If you look at the structure of the 2901, it is eminently built for microprogram control. It's very flexible and it's a very clean architecture. So there wasn't a lot of odd transistors needed here and there. It was a very clean sort of architecture. And so once he came up with a technique of performing a certain function that was repeated across several pieces of logic. And therefore he could remove not one transistor but maybe eight transistors.

Laws: Were there any kind of CAD tools available in those days to help get the job done? Simulators of some kind?

McConnell: I think there was a sliding tool called a--

Downing: It's a slide rule.

McConnell: Did we have HP 35s at that point?

Ghest: No. I don't think we did.

McConnell: I don't think so.

Ghest: No, maybe just coming along.

McConnell: We were just beginning, as I recall, to use time sharing systems to run SPICE.

Downing: That's right.

McConnell: Probably in that year, 1974, when we started developing this product, we also started doing some timeshare simulations with SPICE.

Ghest: But we didn't do any logic simulations. I think they were just about coming along. Maybe a couple of years later we would have.

Laws: Did you build a breadboard?

Ghest: I don't recall. No. I don't think so.

McConnell: I think Clive had very high confidence in his logic design.

Ghest: Well, the reason why was that most of the pieces we had developed before. You know with the arithmetic unit, the carry units, and all the shifting and so on, etc. And therefore, it was rather like putting building blocks together and then minimizing it and making it clean, and making certain that we had the minimum number of transistors for that function.

Laws: Who else was involved in the product development? There one or two mask designers maybe?

McConnell: Edie and Helen.

Downing: Edie Pricolo and Helen Magnus

Holbrook: You guys are good.

Ghest: Were we still using rubyliths?

Simonsen: I think we were.

Laws: About 200 times [the final size] you'd draw it?

McConnell: 500 usually.

Downing: I think it was 500.

McConnell: Yeah**.** It fades into the background. When we started cutting ruby with machinery, then not very much longer after that, we started using pattern generators. But I think this was still a hand-cut rubylith.

Laws: The masks were designed for contact aligners? Or were they projection at that point?

Downing: Generally, it was all contact printing. We were just getting into proximity printing, which is where you don't quite bring the mask down and smash it against the wafer. But we were doing contact printing and were using some kind of polyvinyl alcohol to keep the mask from sticking to the photo resist. And that was a big problem. So basically, I think the biggest hurdle that I saw in this was that we thought maybe if we were at three or four or five die per wafer it would take us a lot of wafers to put any units on the shelf. And so having an inventory to ship, or being able to quote customers any kind of delivery was one of the big concerns. Is that your recollection?

Holbrook: Mine?

Downing: Yours.

Holbrook: I don't recall. I really don't.

Downing: Anyway, right up to the time when the very first wafers came out, there was, I think, a little bit of a reluctance to invest a lot more into the product line until we saw whether or not we could build the first part. And when it came out, the first lot got 30 die per wafer. And there was a big sigh of relief when [yield] hit that 30 die per wafer.

Laws: When did you see those first wafers, do you recall?

Downing: I think it was June of ’75, somewhere around there. Does that sound right?

McConnell: It would make sense.

Simonsen: Yeah, it sounds about right.

Ghest: I thought we had some earlier than that which weren’t quite working.

Downing: Did we have to do a rev[ision] on the mask set?

Ghest: I remember being in a meeting when we first got silicon. And it wasn’t working completely, but there was an awful lot of promise.

Laws: Any other memories from those struggles to get the first wafers out, Phil? You were pleasantly surprised by the yield, and that was good news.

Downing: Other than doing test masks, we really didn’t know what this process would do with real circuits. So we had a lot of transistors and metal patterns and things like that. It was single-layer metal, and that was one of the limitations. AMD hadn’t quite gotten to the point where they had mastered the dual-layer metal process. And so we had a lot of diffusion bridges and things like that that added resistance into the circuit and had to be comprehended. Tom had to comprehend all these sort of funny little extra pieces of circuitry in there that just came as a consequence of the process. So it was a pretty big die. I think the diagrams that John Springer was showing us shows what a huge impact it had when we could go to two-layer metal. But designing a high-speed circuit with diffusion bridges underneath the metal is a real challenge to say the least. And so there was some apprehension of course.

Laws: How did you test it?

Downing: How did we test it?

Simonsen: Very carefully.

<laughter>

Downing: And efficiently.

Laws: Were you able to use a standard tester, or did you have to come up with some special test boxes?

McConnell: I think it was standard. I think it would have been a Teradyne.

Ghest: Because the architecture was so RISC-like, the test program maybe wasn't as complicated as you might have thought.. You'd test the arithmetic unit, the adder itself. Well we had test programs to check adders. So we could check the adder. And then we had test programs where we could check say the shifting and so on. So we'd done all that before. So we could put pieces of this test program together. I mean, it was a big job, but it wasn't an immense task.

Laws: Now I'd like to discuss the product introduction and some of the marketing programs associated with the 2900 family. We have brought two new participants [John Springer and John Mick] onto the panel so I'd appreciate it if you could each introduce yourselves. John, tell us a little bit about your background and how you got involved with the program.

John Springer: I'm John Springer. I joined AMD in 1971, in the quarter that we had our first $1 million quarter in sales. I worked in product applications writing data sheets, customer support, that sort of thing. Then into marketing and then in bipolar memory for a couple of years, I guess. And then I became the marketing manager in the newly-formed bipolar microprocessor division in, I'm going to say very early 1975.

Laws: John.

John Mick: My name is John Mick. I joined AMD in 1973. For the 10 years before that, I had done a lot of logic design work in the aerospace business, primarily doing special purpose computers and digital signal processors for radar and video and various things like that. I decided that the semiconductor business was the place to be. So I joined AMD, and was fortunate enough to be the first person they ever moved from someplace to California. I got involved as the product definition and applications engineer for bipolar logic, which became the 2900 family. I had almost every data book from every semiconductor company in the business from my ten years of aerospace design. And I remember going through every data book to see what we were going to call this new family. And there was nothing that was called 2900. And I wanted to make sure that we didn't actually use 2900 for a part. So 2901 became the first part. The original family was going to be the 2902, for the carry look head. 2905, which was an open collector bus interface unit. 2906 was one of those-- an open collector bus interface unit but it had parity on it. 2909 was the microprogram sequencer. The 2915 was a TRI-STATE bus transceiver. The 2916 had parity and the 2918 was a really high-speed register with two sets of outputs. That was the original family. Then we got the 20-pin package, and we added the 2907 and 17 and the 2911 as a microprogram sequencer. And it developed on from there. But when we started the introduction we started out actually [presenting it] as a family.

Laws: At what point were these parts defined, John? You took over from Clive, when he moved on to MMI?

Mick: Clive was there working on what became the 2901. And basically I did the logic design and everything on all the other parts. And Tony I think said he was running the interface groups. Suren Kodical did all the bus transceivers with Stan-- what was Stan's last name?

Holbrook: Stan Wilson.

Mick: Yes, Stan Wilson. And Gust Perlegos did the microprogram sequencers. And so that was how the original thing started. So it was an exciting time.

Laws: You've got a wonderful recall for the product numbers, John.

Mick: When you worked with them as long as I did. We kept doing more and more parts in the family. And I was working with the circuit designers and trying to keep track of when they absolutely had to have the next circuit. So I always made sure I had the logic done so that they could start thinking about it.

Laws: John, at what point did you start putting together a product introduction plan?

Springer: Well, I remember the date we announced the product. And that was July 21st, 1975. And that date seems to be burned into my memory. And I don't really remember when we started putting all the programs together. I mean, the big deal was the announcement of the 2901 itself, and the rest of the parts were kind of the gravy that went with it. Partly because the 2909 at least wasn't out at the same time. I want to say that sometimes people say the Microsoft strategy was embrace and extend. And AMD did that with products. There was a series of proprietary AMD products. And they started with the 2500, which would take like an existing bipolar MSI chip and add something to it that made it unique and gave us a competitive edge. And so that became the 2500 series. There was a 2600 series, a 2700 series for memory. I assume there was 28, I don't remember what that was. And then 29 was the next one in line. So it was part of AMD's proprietary product numbering system from the beginning. And a lot of these parts, the 2915, for example, were parts we'd already introduced as a 26 something. It's a product that had already existed and then got put into the family to build a bigger family.

Mick: A couple of them did.

Springer: Yeah.

Mick: Some of them didn't.

Springer: Yeah. So there were a couple of those interface parts that we did that with. And it worked. We would get orders for that. But anyway, back to the original introduction. All the focus was on making a big splash when it came out. And so there was a data sheet to get out, but the basic marketing program was advertising and seminars and out talking to customers and product literature and all that kind of stuff, which I believe we put together pretty quickly in the early part of 1975.

Laws: So the date July 21st, is when an ad was run in Electronic News?

Springer: It was a three-page ad. I don’t think we’d ever run a three-page ad before. And the headline on it was “Bye-Bye MSI.” We were claiming a revolution in bipolar technology.

Laws: How was it received? Did the phone start ringing off the hook?

Springer: Yeah, they absolutely did. And we had product. As I recall, we had product in stock at distributors …

Mick: Had 5,000 units.

Springer: … the day that that ad broke so you could go get it. And probably one of the stupidest things that I’ve ever done, we actually got little boxes. Because this part, this was like a really special part. And we weren’t planning on shipping hundreds of them at this point, people were going to buy ones so we put each one in a little box by itself. And it was a cute box with the AMD logo on it with some padding inside and one part. So you got one, and it was obviously really special.

Laws: A jewel.

Springer: Yep. And $29.95 was the introductory price, 100-piece price.

Mick: That was the official introduction. But I actually have some data books from 1974, where we had datasheets and things like that in them. So customers could see what was coming.

Laws: How well characterized was the part at the point of introduction? Did you really know what you had?

Mick: When he’s talking about, 1975. It was well characterized by then. They were talking earlier, I think the first silicon actually came out in late 1974. That’s my memory, it was actually late 1974. Because I joined in 1973. And one of the reasons I was really excited about AMD at the time was they really wanted to do proprietary products. And that’s what I wanted to do. If you didn’t want to do proprietary products, you didn’t need me. Anybody can do second source stuff. Just copy what the other guy did.

Laws: Tony, this must be about the time you were appointed managing director?

Holbrook: Yeah, I think it was spring of 1975 that we formed the business unit. And I was asked to head it up. I want to go back to one point John Springer made here on advertising. You know, we put an advertisement out, “Bye-Bye MSI.” I thought our advertising approach was really effective, and one of the important ingredients to our ultimate success. There were a goodly number of components to our success. But the advertising was quite nifty. We did a block diagram of each part as we introduced it, which was stylistically done and I thought captured the attention of design guys in the marketplace. I mean, they saw this block diagram, they could relate to it.

Laws: These were the hand-drawn diagrams.

Mick: Yes.

Holbrook: Hand-drawn block diagram …

Springer: Like in pencil.

Holbrook: … of the parts. And the ads were never busy. They never were overwhelming. They were simplistic, but they captured the attention of the right audience. And they delivered a message that was really salient. And the advertising program actually was consistent throughout, right? I mean, we always did the same kind of thing. And it obviously, I mean in retrospect, it captured the attention of design guys throughout the world.

Mick: As I recall, that was Elliot Sopkin, wasn’t it?

Holbrook: It was. It was an outside advertising agency under the direction of Elliot.

Springer: We had an absolutely brilliant ad agency, Keye, Donna, Perlstein. And Paul Keye was the copywriter. They were a consumer ad company. They’d never done technology before. They went on after us to do Continental Airlines or something like that. And they were down in Hollywood. They injected a consumer ingredient in advertising that other semiconductor companies didn’t have. So it gave us catchy lines and provocative graphics and all that kind of thing.

Holbrook: Well, not just provocative, but graphics that the audience, our chosen audience, could identify with.

Springer: Yeah, yeah.

Holbrook: It had meaning to them.

Simonsen: And there was another thing that went through all of them. They always used a special type of individual letters.

Laws: Type font.

Simonsen: Font is the right word, yes. That went through all the ads. So people instantly recognized it.

Springer: Marketing the 2901 was an educational process. The key to getting people to use the 2901 was for them to understand how it worked and what they could do with it. So from the very beginning, the task in front of us was not price and delivery, it was teaching people how to use something that was new. Everything went on like that.

Holbrook: Let me expand on what I felt like were the component parts to our success. This family turned out to be very successful. I mean we really penetrated the marketplace. Why did it do that? It wasn't just simply a new widget, and the widget was nice. But it was a lot of other things as well. First, we got an advertisement out there that captured attention. The next problem is the customer base wants to use the widget, well how do I do that? So John Mick co-authored a book on how do you build these things? How do you build a system around this family of products? And that book became a bestseller. I mean, it was widely distributed. And in the years since, I've had any number of people come up and say, "Oh I, you know, I had this book, I read this book. It led me into this and this and this and this." We then fleshed out the family with supporting products around the bit-slice architecture so that it was more and more useful to people. It's a microprogrammable architecture, so there's software that has to be created by the designer. So we provided System 29, a development system, which let the customer base have the necessary tools to design code for the thing. Is there more there?

Mick: Well, we did the microprogram assembler. We did all that System 29 stuff.

Holbrook: And we provided software.

Laws: AMDASM.

Mick: AMDASM, right.

Holbrook: So we essentially educated and enabled the audience.

Laws: Now that was a major departure from AMD's traditional selling task. How did you go about addressing this?

Simonsen: Absolutely**.** You earlier asked, did the orders start coming in. Of course, the business we were in at the time was that it doesn't work that way. You have to go out and educate the customers and interest them in using this chipset for their next generation of whatever they are building. And so they may buy a few to start playing with it. But the real orders come when that next generation of their machine starts getting into production, which may be a year later or something. But we certainly went out, applications people, myself, field application engineers went out and made several large presentations at various customers. I did a lot of those overseas and people working back here did at many companies in the U.S. and gradually, we got more and more people to convince themselves they wanted to use this series for their next generation equipment.

Mick: And we had a really good sales force. I tried to figure it out once, and I figured I had given well over 1,000 seminars at different customers' and so forth. Sometimes you go out for a week and you do 15 seminars in a week, right? Hitting all these customers or different groups within the same customer, like at DEC or Data General, and so forth, you go see different groups.

Laws: And it was a different task for the sales force too.

Mick: It was a different task for the sales force.

Laws: They had to completely change their sales strategy.

Mick: And we had a really good sales force too, which was good. And then I also taught microprogramming courses through the University of California at Berkley. And I had typically as many as 80 students a semester who thought, "I've got to learn how to microprogram." And so I figured I had taught 400 or 500 engineers in the Bay Area-- I used to say I taught half the engineers in the Bay Area how to microprogram, which probably wasn't true. I had a lot of, lot of students go through my classes.

Laws: And at this point, you must have started adding field applications engineers?

Simonsen: Yes, yes we did. So instead of having to send people out from the factory where they were busy defining their next products. We started getting field application engineers that were out there and could more regularly visit the customers.

Mick: I was the first FAE manager. But it was basically under Sven. We started with six. Clive actually joined us in Europe as the FAE in Europe. We had Steve Campbell on the East Coast. I think it was Barry Fitzgerald on the West Coast. We had two people in Japan, one of their names was Watanabe, and I don't remember the other one.

Springer: Was it Hiro?

Mick: Hiro, yes. And we had one in Central [US], who I don't remember who. We had the first FAE conference, when six people came in. And for a whole week, we trained the FAE's on not only the bipolar products, but some of the MOS products too. And we used to have an FAE conference every quarter where we would bring them back for a week and really, really train them. And then they'd go out and train the sales force and things like that.

Simonsen: The sales force were really good. I mean, many of them had a bachelor's in engineering.

Mick: Right.

Simonsen: So they could relatively quickly pick up on some of the salient points of these products.

Springer: We had some early beginning application notes on the 2901. So we began a series on how to design a microprogrammable computer.

Laws: This was [eventually published by AMD as a series of book chapters under the title] "How to build a computer today"?

Springer: "How to build a computer." And we ran ads and people could sign up for it. And we would do like, I want to say one a month, but I'm sure it wasn't one a month.

Mick: I think it was one a month.

Springer: Was it? Anyway, there were probably a dozen of these things that came out over a period of time, a year plus. And they were all three-ring punched in a little binder, and you could put those in and you would have a book. And then it eventually became your book.

Mick: Became a McGraw Hill book, right. I think I had heard once, I don't know if it was true or not, that it had been translated into 18 different languages and it was used in 170 universities throughout the world, which created a lot of interest.

Laws: Were you orchestrating much of this marketing activity, Tony, or was it going on around you?

Holbrook: I was riding the wave.

Springer: Oh, I haven't heard that for a while.

<laughter>

Simonsen: Isn't it refreshing to have a real honest guy?

Downing: We also did some early second source deals as I recall, didn't we have-- what was it? Siemens and--

Holbrook: Raytheon.

Springer: Was it Raytheon?

Downing: Yeah. Raytheon was one.

Holbrook: Oh, I forgot that.

Springer: Yep.

Downing: And then we had SGS Thompson?

Springer: Yes, I think you're right.

Mick: I think you're right.

Springer: These were strategic relationships.

Downing: And then Siemens came in.

Holbrook: I think that the 2900 family's success and the component parts of its success was a great model for business even today. And you'll note that there's been a lot of semiconductor companies over the years that have dabbled in microprocessors, mainly CMOS, but with little success. There's only several architectures that are out there, one of them being the MIPS architecture, which isn't even supported. It isn't originating from a big semiconductor company. And why is it a standard? And it's a standard because John Hennessy, of Stanford, invented it and then wrote a book on it. He was a college professor, president now, of Stanford. The book got widely disseminated. Students got trained on it in universities all over the world in all kinds of languages. And it's a standard supported by a host of companies with software and design tools and things of that nature. So when large companies want to penetrate the market with a proprietary product family, there's a lot to learn by studying the 2900 family success, and the component parts of it and what's necessary to do that. It's a lot more than just inventing a widget.

Mick: One of the things that's kind of interesting in terms of the support was the System 29 development system. AMD actually had the first OEM contract for CPM with Gary Kildall at Digital Research. We were the very first actual OEM. And the day that IBM went to see him, for DOS (when he went and flew his airplane) he had his VAX at home in the basement. And they [IBM] went up to see Microsoft, and you know what happened. It was really interesting, all that history of how things happened and so forth.

Laws: So eventually, that [development systems activity] was broken out as a separate business unit Tony, right? The development systems unit was called Advanced Micro Computers [AMC], I think?

Holbrook: Well, yeah. Advanced Micro Computers was a separate business unit, but not just based on the 2900 development system. It was based on the several development systems in the company, as well as OEM board-level products. We were trying extend ourselves into the systems business at the low end, with modest success.

Springer: The Z8000 was the big push on that one.

Mick: Z8000 was the big push, yeah.

Holbrook: Yeah, it wasn't based [only] on 2900.

Mick: We did the original System 29 in my group, and then we turned it over to this new group, which was going to try to go off and do Z8000 stuff and so forth.

Laws: What other tools did you develop to sell the product? I remember you started some educational classes under Donna May White.

Mick: We had that. And we had some development boards that we did. So people could play with toggle switches and lights it, just to get a feel for what it would do. We did a couple of those.

Laws: And having put all this effort into marketing promotion and establishing a proprietary family, where did the product really take off? John, you must have been right at the heart of that.

Mick: You mean what company?

Simonsen: You mean what companies?

Laws: Which applications and companies?

Springer: I'm trying to remember where the first order came from.

Mick: DEC [Digital Equipment Corporation] used a ton of them. I think it was DEC.

Springer: DEC was big. But I think the first order may have come from Hughes Aircraft.

Holbrook: Yeah. It had a lot of success in the military market, aerospace and defense. You know, there was a lot of high-rel [reliability] product required. I think the success of the family wasn't tied to Digital putting a big minicomputer out or anybody else. It was a very broad base of design activity.

Simonsen: Cisco built the next generation router around the 2900.

Mick: But fast forward to 1991, I would bet almost all of the military equipment [used in Iraq] was bit-slice based. An awful lot of military equipment still is.

Laws: The advantage there was the ability to microprogram and emulate previous applications.

Springer: I know we put parts in the Space Shuttle. I remember that program.

Mick: Yep.

Springer: It took a year to ship two parts or something.

Holbrook: You know the microprogrammability feature was very attractive to design people because it gave them complete freedom on the instruction set. Today, if you buy a microprocessor, a Pentium or any other kind of processor, the instruction set is fixed. And now your hands are kind of tied on what you can do. There, you establish your own instruction set. So you can very efficiently tailor the computer or the computing engine, or controlling engine, to do specifically what you want. And as Clive pointed out on simplicity and elegance, I mean, it's the ultimate, because there's no fixed instruction set. You're not tied down by the decisions somebody else has made as to what this thing can do.

Mick: Yeah, the microprogrammability, I mean, you think of it as a computer, but in a lot of cases, like a disk drive, you could actually microprogram everything you wanted to control, all the sectors. There wasn't any real instruction set, and all the microcode was just to control the disk drive. So it wasn't like a PDP-11 instruction set or something.

Laws: So the disk drive was a major application?

Mick: That was one of the applications. And lots of those kinds of controllers, whether it was a disk drive or a tape drive or whatever, found lots of the microprogrammed controlled stuff. And then they could actually upgrade the microcode and add more capability and so forth.

Laws: Did you have to do things like [offer] reference designs that people do today John? Where you hand the customer almost a complete kit.

Mick: That was what the book was. We didn't actually ever hand the customer a complete reference design. But we handed them a book that showed some reference designs and showed the guy how to do it. And particularly in the Bay Area, as I say, a lot of people came and took the microprogramming classes and things like that. We gave enough seminars. And I think we can say we taught people how to use it. But we never actually, "Here's a working reference design," which you kind of see today in a lot of things."

Laws: Are there other applications that stand out in your mind as being really enabled by this family?

Mick: Well, lots of digital signal processing **(**DSP) systems, early on. We eventually did what was called the 29500 family, which was a set of circuits specifically for digital signal processing. So early on, we captured a lot of things there. Film processors. Had a couple funny different applications. The one that I always enjoyed was the guy who used one slice in a carwash. Used one bit for wax and one bit for soap. Had 16 registers, so he could queue up to 16 cars.

Simonsen: I think the truth is that a lot of applications came up that we hadn't even thought of when we developed the chip. And I remember a certain telecom company made switches that were built with this. Remember, we are talking about a time where computers is sort of a thing sitting here you use to generate some data and print stuff out, etc., etc. But it became more and more common that in almost any kind of digital equipment was embedded some kind of processor. So it found a lot of applications.

Springer: Certainly one of the major users was DEC.

Mick: DEC was a big user. Data General. Honeywell had their System 6, I think that they called it, computer, that was all bit-slice things.

Laws: What were the competitive reactions to the product? There was some second source activity I think John mentioned? Intel had some kind of a chipset?

Simonsen: Intel was really scared for the first time.

Mick: The funniest one I remember was Signetics, who was going to second source the part. They got left and right, shift left and shift right backwards. And they gave up.

Laws: They also had a product, 8X300 or some such number.

Mick: Yes. That never went anywhere. Intel had their bit-slice thing, which as I recall, really never went anywhere.

Simonsen: 3000.

Springer: 3001 and 3002.

Mick: And the sequencer was just awful. It was just terrible.

Springer: Our product was much stronger. The competitor who eventually actually amounted to something was National.

Simonsen: Microcontrollers?

Springer: Well, no, they did the 2901.

Downing: Second sourced it. That's right.

Springer: Yeah. And that was scary. Right, because National was a big company.

Downing: Well, it was faster than ours when they first came out with it.

Springer: But ultimately, it turned out not to matter. I mean, one of the things that started to happen around this time and became so much more obvious as we got into the Z8000 was that this thing about picking what product you were going to build a system out of became not just a technical decision, but a decision about what company you were going to stake your future with. And I think that's why National never really could penetrate on the 2900. We were clearly the leader. We were ahead of everybody else; we were doing the new stuff. And we're the ones people wanted to do business with.

Mick: Probably by then, [we were] talking already about the 29116, which was the 16-bit controller.

Springer: So National was something [for customers] to like, tweak us with.

Mick: Right.

Springer: But they never really took any significant business from us.

Laws: At what point did the 2901 A, B, and C [versions] start to appear? And was this using a different circuit design approach? Or a different technology?

Downing: Two-layer metal. What motivated the 2901A was that performance was being rather severely impacted by all the cross-unders and so forth in the circuit. So the designers were really anxious to get two-layer metal going. And it was right about 1976. We built the Fab 3 and we were able to separate out the MOS and bipolar [manufacturing]. Then we could really embark on the two-layer metal and we got something fairly decent going. That gave, what was it, like a 30 percent improvement in performance?

Springer: Mm-hm.

Downing: So that was a big deal going to the 2901A.

Laws: The lithography was the same? The only difference being the two-layer metal? Or did you tighten things up?

Downing: I think we went from three micron to two micron.

Mick: They were talking earlier about Tom Wong and his circuit design. One of the things I always remember, I was amazed at how he had this technique where he could figure out how to take one collector, put multiple bases in it, then put multiple emitters in each base and make this one thing that would do all kinds of neat stuff. He was really phenomenal at figuring out circuits and transistors that would do a bunch of stuff.

Springer: It was a different way of thinking. He was not thinking about transistors, he's thinking about P layers and N layers and how they work.

Downing: Right.

Laws: Phil, you had taken over from Tony by then?

Downing: No, Tony had the group until 1978.

Holbrook: I think 1978.

Downing: It was either late 1977 or early 1978 that Advanced Micro Computers was [formed] that's where you went from there. So I took over as acting managing director for the first half of 1978. And then became officially managing director in the middle of the year. And John had brought in another product marketing engineer.

Springer: That was Geoff Tate.

Downing: I was out of the process [engineering role] and so when Tony left I said, "What do I do?" I had it for about a year and a half. And as I recall, sales were $3 million or $5 million a year when Tony left?

Somewhere in there. So I was very fortunate to come in when the sales really took off in ’78. By the end of ’78, we did $15 million.

Laws: So it took about three years from getting that first product out to seeing a real kick in revenue.

Downing: Yeah. A [for] lot of customers. But the gestation period for the design had been about two years typically, before the customers’ products started to take off and we really got production going.

Laws: Especially if there’s a military component to it, it takes a long time for that to happen.

Mick: A lot of military we did. And a lot of commercial too.

Laws: Which revision of the 2901 were you at in late 1978, Phil? You’d [already] done the dual-layer metal version.

Downing: We’d done the dual-layer. I think that by that time, we were just doing the 2901B. Which was taking it down to like 1.2 micron. At that point, it was really all about shrinking it to take advantage of the lithography and keep doing the performance. And it was about ’79 I think when National came out with their part. And we had to respond by converting to ECL. The bipolar memory group, who was run by Morris Chidlow at the time, as I recall and he was saying “We want to get compatible processing, so please go into the ECL business, so we can build bipolar memory and bipolar microprocessors [in the same process].

Laws: That was using ECL internal circuitry with Schottky I/O levels?

Downing: That’s right. So the fastest, smallest [chip] that was there before I left, as I recall, was the 2901D.

Laws: And about this point in time, you started thinking about a next generation device, the 29116. John, when did that start?

Mick: Well, I would say after the 2901, we did the 2903, because we’d had a lot of people wanting BCD. So it had BCD arithmetic in it. And then, I would say it was about ’77 when we were working on the 29116, which was all ECL internal, as Phil said. And it was a 16-bit controller. I think, Tony, you were still running the group at the time when the first silicon came out. And that was again, a struggle to get all that on one chip. Fact is, we had too much stuff. And Tony said, “What can you take out?” And we took one of the accumulators out. If you look at the datasheet, you can still see the hole where it was in the provisional data sheet. I remember Tony coming and saying, “What else can you take out?” And we looked and talked, and finally said, “If you can’t build it, cancel it. We can’t take anything else out.”

Springer: It was really difficult though, to try to figure out what to do next.

Mick: You mean after the 29116?

Springer: No, even after the 2901, what the next generation should be. I mean, you remember we had all those meetings with DEC.

Mick: Oh. <laughs> Yeah. We used to have a quarterly meeting with DEC, with their evangelist. One quarter they'd come to AMD, and the next quarter we go back to Boston and meet with DEC. My memory is that we always thought the next thing was 16-bits. How we were going to architect it? What we were going to do with it and so forth?

Laws: One of the unique aspects of the 2900 family was its longevity relative to many other products in the industry. Bob, I think you retired the products just a few years ago, is that correct?

McConnell: Yeah. Right around [the year] 2000 the original 2901 was finally at end of life.

Laws: You took over as managing director of the product line from Phil, I believe, in the late 1970s?

McConnell: I think it was a new title called product manager. Working for John East in a new organization in the bipolar part of the company.

Laws: And about what size was the business at that point?

McConnell: I think the business at that point was growing into $10 million a quarter, $35 to 40 million a year kind of numbers.

Laws: And the Am29116 existed at that point?

McConnell: When did we introduce that? I think 1980. Clive?

Mick: I would have said '79.

McConnell: And the 29116 was very slow to take off. It didn't contribute a lot of revenue for two or three more years. It was a long design-in cycle. And perhaps incremental to the basic concept.

Laws: And what kind of technology were the products on by that point in time? They were ECL internal, probably one micron.

McConnell: ECL internal, about one micron. I think we were almost getting to microns at that point, away from mils. And I think we were three-layer metal by then.

Laws: On four-inch wafers?

Mick: Six.

McConnell: Think six? Would have been six, yeah.

Laws: An interesting transition here. Started on two, 10 years later you're on six.

McConnell: Started on three. Because Fab 2 was an unusual place. And fortunate for AMD that it was three inch. Because that gave it much more longevity. Yeah, so the 29116 was three layers of metal, one micron. Pushing 50,000 transistors, I think. And quite a lot more complexity. And brought with it all the additional challenges of designing that kind of stuff. And AMD was always at the leading edge of how big can you make the die? Oh, the reticle pattern generator won't go bigger. That kind of stuff was a constant battle with the various choices of mask vendors and simulation hardware and software, simulation tools. The vendors would say, "It'll solve your problem." No it won't, because we're doing more stuff, bigger stuff than most anybody in the industry.

Laws: And the days of rubyliths and stripping was long gone by then. What kind of tools were you using for layout?

McConnell: Layout was probably still by hand. I'd have to go search to find when the transition came.

Laws: Wasn't there the renting of a warehouse to layout the drawings for the 116?

McConnell: Maybe for the checking. The check plots and so forth. But at this point, I think we were probably moving from hand layout and then digitizing the layout to digitizing directly, just building it on the Calma. The basic tool was Calma and the output then went to a pattern generator. Those were big transitions. Lots of difficulties, lots of interesting problems to resolve on the technical level.

Laws: Did any of those tools have a 2901 in them?

Ghest: Good question.

McConnell: It probably did. Actually, the Calma probably had a Data General computer in it.

Mick: Probably. I was going to say I thought it was Data General, yeah.

McConnell: It probably did have a 2901 in it.

Mick: Very possible yeah.

McConnell: Never thought about that.

Laws: How big was the group? How many designers, product engineers? Twenty, thirty people maybe?

McConnell: Twenty-five to thirty people, in that range.

Mick: I was thinking Gust Perlegos did a lot of the 29116 [circuit design]. Bill Harmon worked for me, and he and I did the architecture. We hold the patent on it. And he did a lot of the actual detailed logic design.

Laws: And how did the definition come about? Customer input?

Mick: We talked to a lot of customers, yeah. I mean, it was always my belief the next thing had to be 16 bits. And so it was a matter of what was it and talking to a lot of customers and then working with the circuit guys on what could we actually build, you know, what made sense. Was this something they thought they could design? One of the things that was unique about AMD, I always thought, was that in the semiconductor business, you get so many different groups. You've got process guys, design guys. And in the process you've got metallization and all these other technologies. And the marketing and all, so on and so forth. But each group seemed to be very good in their area. So it was really like an excellent team that worked all together. So when we were doing the definition, this is what we thought we could do. And talking with Gus who was doing a lot of the circuit design, "Can you do this?" And so forth, and back and forth. Talking with customers then to get input on "Does this architecture sound good?" And so forth. Without necessarily giving away stuff that would immediately go to your competition. One of the things that I learned, I think at AMD was how to pick a customer's brain without always necessarily handing him a whole bunch of stuff that he can pass on. So it's a challenge, a skill I guess.

Laws: Clive, you were back at AMD again by now. And I think you were, I think the first FAE in Europe, is that correct?

Ghest: Yes. My job there was to build up an FAE team in Europe. And that was one of our big markets. Obviously the US was first and then I would say Europe was second. And later on, Japan and Korea came into it. But yeah, there was a big market in Europe. There were a lot of computer companies such as Olivetti, and they used the 2900 series. And an awful lot of military customers. Olivetti in Italy, and of course there was Siemens. And companies in Norway and Finland and so on.

Mick: Nixdorf was another big one.

Ghest: Nixdorf was another big one, yes. And then there were companies in the U.K. like General Electric U.K., and a lot of the military people there were using 2900 series. Ferranti was using 2900 series. So like John, I spent a lot of time giving seminars, especially on microprogramming. I can remember being in France and Paris. And always we had full crowds. People wanted to understand this

2900 series and how to do microprogramming. And it was enlightening when you've found some engineers and they said, "Well why is microprogramming better than what we do now?" And just with some cute examples you show them that if you do it in microprogram, it might take five seconds. Whereas if you do it with hard-wired logic, it can take you a couple of days to minimize things. And so it was an interesting time. All the Israeli companies, especially the military companies, were using it. I gave seminars from nine o'clock in the morning to nine o'clock at night. I used to go essentially hoarse. You know, you just ran out of talk. And even after the seminar, you'd have all these people coming up to you asking you specific problems and what were you doing next and so on. So it was a very exciting time, very interesting time.

Laws: Was anything different about the demands of customers in Europe?

Ghest: I would say that in Europe, and John might correct me. But in Europe, I think people used it in a much more flexible way. There were a lot of military applications. , One company would be using it for a navigation system in an aircraft, another company would be using it in a tank for guidance and ranging and so on. And then other people would be using it for disk control, like they did here. So it was a huge variety of applications. And of course the Russians wanted it for signal processing. I found fascinating, that you'd go to customers that you really wouldn't think would need it. You know, maybe somebody would be controlling a sawmill, and they'd want a small bipolar high-speed microcontroller to control his sawmill. And things like that.

Laws: So they really needed the speed there, it wasn't just the flexibility?

Ghest: They needed the speed. So it was very interesting.

Mick: We haven't mentioned it, but there was a 29110, which is also a microprogram sequencer, that went with the 29116, a companion part. Some of the same issues of teaching customers how the parts worked, what the new controller features were, and so forth. I always found the customers in Europe had really good questions.. But it was still the same teaching whether it was in the U.S. or Japan or Korea. You were teaching them how to use the parts and how they worked and all those sort of things.

Ghest: I think another big application was telecom. It was used in a lot of the electronic switches. They needed the high speed and they needed the reliability because there was dual processing. And so a lot of the French used it, Siemens used it, some of the Japanese used it. So that was another important application.

McConnell: Yeah, AT&T used it. In fact, they bought a high-rel version that was without a doubt the most lucrative business altogether that we had.

Mick: Cisco ultimately became a big user in the U.S.

Ghest: Cisco?

Mick: Cisco was a big user of the 116, yeah.

Laws: ATT used it in a switching system of some kind Bob?

McConnell: Yeah, it was in one of the basic AT&T switches. ESS [Electronic Switching System] something. And they paid very high prices. They weren't concerned about cost. They wanted reliability. And they would always be the ones asking question like, "Can you guarantee this part will continue to operate for 20 years?"

Ghest: They had criteria of one-hour downtime in 10 years. That's what they set their goals at. So you had to have extremely reliable parts and multiprocessors to do that.

McConnell: At this point, I think bipolar still had the reputation for being more reliable than MOS, and CMOS was still not a very big thing. And so we had the advantage in all the military and high reliability applications like telecom switches where people believed in general that the part was inherently more reliable.

Laws: Did you eventually do a CMOS version?

McConnell: We did a CMOS version. I'd be hard pressed to tell you that we ever sold any or not.

Mick: I was at IDT, and we did a CMOS version and sold a lot of them. <laughs> On the Am29116, as I recall, one of the really hard parts was the packaging. We finally ended up having to do a cavity down package with a heat sink on top of it that we spent a lot of time working on to dissipate the heat.

Laws: Now at some point the group was moved out to San Antonio, Bob. You took that out there?

McConnell: Well, part of the group I think was moved to San Antonio. In particular the 29300 family. And some of the multipliers and digital signal processing stuff was moved to San Antonio. I'm not sure I could detail the timeline of the transition to Austin. But eventually it ended up in Austin in the 1988, 1989 timeframe. One of the things that's interesting in the 1980 timeframe is there was still a problem at AMD, which was that the MOS group and the bipolar group didn't talk to each other, and didn't work together. And therefore, we were really pretty restricted in that we had to stay in bipolar technology. This was the definition of the product. And at some point in the early '80s, we should have been transitioning this family into CMOS and it wasn't happening. So there was a little AMD organizational issue that was in our way.

Laws: Was that resolved

McConnell: Yeah. But it was in the '88-'89 timeframe that one of the problems left. <laughs> And then it became easier. AMD was also late to CMOS. I think that's a well reported phenomenon. And the right technology going forward for this type of product was definitely CMOS to provide speed and lower power. And that technology wasn't really available at AMD until the 1987, 1988 timeframe.

Laws: Anything else about the development of the family at this point in time that's worth recording here?

McConnell: What I will say is, I think the marketing in this family was far more important than the development. Sanders had a line about punching above your weight that he used once in a while. And I think that's what we did with the 2901. We made more out of a family than we really had, and really promoted the company in a big way well beyond the accomplishments of the designers. And it was really kind of everybody working together and the result was bigger than the parts.

Mick: I always thought Jerry was probably the strongest proprietary product guy. He absolutely wanted proprietary products. That's how he knew he'd be successful. Letting other people second source you. So it was always a lot of fun to try to define products and talk with customers, knowing they were going to get built, that it wasn't pie in the sky. He was really strong in that area.

Laws: Okay, we're getting to the point where I'd like to start doing a wrap up. So Bob, obviously the 2900 family played a very important role in your career. And what did you think you really learned most from this experience that you were able to apply later?

McConnell: Well, there's a lot of things. You know, I think the close cooperation that we were able to achieve between process, circuit, and logic design was the best example of the right way to do things. I mean with John and Clive talking to Phil and I and the designers every day, right? About what could we do? How could we do it? How could we achieve what we needed to achieve? There was a level of cooperation between the elements of the design that I've rarely seen before or after. It's the goal that you set to try to be like that, because you know it's the right way to do things. And down the road at AMD, I think we were able to achieve some of that cooperation in the processor group as AMD started to build the 486s and K6s and move on. And I think we learned that that's how it had to be. The designers and the process people had to work together to achieve really good results in a good time frame? And when you do it in the compartmentalized, throw it over the wall methodology you just don't achieve the same results either in time or in quality. And so that's the thing I took away the most. Now, and then again, I'll also say what I also took away first and foremost is marketing wins. You take a product that's perhaps not especially unique, that other companies could have done, and you promote it in a way that makes the company and the product special. So those two things I would say are the things I took away.

Laws: Clive, how about you? What do you remember most about working on this product at AMD?

Ghest: Well, I'd like to endorse Bob's concept. I think it went well because we respected one another, we liked one another, and we knew enough about the other person's discipline so that we could have meaningful comments. So for example, when I worked with Tom Wong, he knew enough about the logic design aspect so that when I talked to him, he could understand what I was saying and could then go think about how he could do the detailed circuit design. And likewise, I would be able to say "Hey, if we did this, would that work?" Even though it was on circuit aspects and I wasn't a circuit expert, per se. And he would listen and said, "Yeah, maybe that's good. Yeah. I'll think about it." So the way we worked together throughout not only the architecture, the logic design, the circuit design, the processing, and ultimately, the marketing guys, everybody worked together as a team. And I think that is very significant. I didn't write a specification and then hand it over the transom. I was there talking to the circuit engineer every day. And the circuit guy was talking to the processing guy every day. And the marketing guy was talking to everybody. It wasn't on a formal basis, it was just, hey, this is the product, we went to get it out, we want the best job, and we all worked together. And I think that was a really good basis for how we conducted other programs at AMD at a later time. And yeah, to go along with the fact that the 2900 series is the best product line. Well of course, we worked on it, John and I, we think it was the best product line.

Mick: It was.

Ghest: But you know, even if it was the best, I don't think we would have done that well unless we had really had this great marketing and support of the product line. That was key. I mean, products are 10 a penny, but it's getting that team of people together who can get the right marketing strategies and get the customers oriented to use your product. And I think we did a superb job of that. And not only the marketing guys, but all the FAEs in the field. I mean, it was an enthusiastic time. We had sales guys and FAEs who wanted to educate the customer about the product because they were excited about it. And that rubbed off on the customer. And then the customers got on board and gave us ideas for new products. And we formed relationships with companies all over the world, didn't we, John?

Mick: Absolutely, yeah.

Ghest: And they felt comfortable to just call us. A guy like Paul Hart at Hughes Aircraft. He would be on the phone to me even years after I left the 2900 series program and say, "What are you doing? What's going on," you know, and so on. It wasn't so much that we were calling the customer, they were calling us. So I think that was a very enjoyable and very interesting aspect of this whole program.

Laws: When you walked in that day with a pink sheet of paper, did you have any concept that this was going to last decades and have the success that it did?

Ghest: Well**,** I thought that microprocessors were going to be successful. I mean in the early days, there were a few people who said that these are going to sell in hundreds of thousands. There was nobody who said they're going to sell in millions. And there were one or two people who thought that maybe that they were going to sell in tens, right John?

Mick: Yeah.

Ghest: But I think we all felt that there was a good market out there. People were going to go to computer systems, not only the general purpose computers that are used for data processing, but also for controllers and control all the various aspects of all these digital systems. I think we believed that. And we came up with this chipset to do that. And I think it was a good chipset. I'm sure in retrospect we could have gone back and improved it. But for the time, it beat the competition. And with the marketing push we had and the people we had there, it became very successful.

Laws: John, what are your most significant memories of working on the project?

Mick: Well, if I can go back a little bit. I was extremely lucky from the standpoint that my dad actually got me interested in electronics when I was in the third grade. So all through grade school and high school, I made my spending money fixing radios and TVs. Then after I got out of college, I spent ten years in aerospace designing all kinds of stuff, the most notable probably was the all the digital logic on the GSQ-113, which was part of McNamara's line during Vietnam for the drop sensors. And it would say where the enemy was. And then I also did all of the digital design work on the APS-94D side-looking airborne radar, and built all the test equipment for that. And I always remember when I finally decided to go to the semiconductor business, I came and interviewed Clive, and I had all these pictures in this book of all the stuff I had designed.

Ghest: I remember that.

Mick: After I joined the company, he said "Wow, I couldn't believe all these neat pictures you had," of the stuff I had designed. But the thing that was so good was that I joined the company knowing they wanted to do proprietary products. And to me, it was just unbelievably exciting to try to be as creative as I could dreaming up new products and knowing that if they were good, people would support it. Then to go out and talk with customers about your ideas and get the feedback from them and talk back and forth. And like Clive said, they got to know you, they'd call you, ask you questions and so forth. And then give you a lot of feedback, that it was just a really, really rewarding experience. And we were such a good team. And the thing I would go back to is, I thought that probably overall it was the best team that I had been associated with. Everybody worked together, trusted each other, asked each other questions, we never argued or fought or tried to do anything as far as I'm concerned that was not constructive. And so it was just for me an unbelievable experience that I really enjoyed. And it's nice to know it turned out to be successful. You know, we always thought it was unbelievably successful in terms of all the different places. Like, it wasn't hundreds of millions of units. But most of these, particularly the military systems and so forth, all the different companies that were using it, somehow doing it. And I enjoyed the teaching part of it too. Giving the seminars and the classes and so forth. So for me it was great fun.

Laws: John, you had some interesting career moves after you left AMD. I'm not sure we have time to go through all those right now. But tell us a little bit about what the AMD experience meant for you?

Springer: It was absolutely the most fun thing I've ever done. I mean, it was everything I wanted to do, I guess. It was high tech, it was state of the art, and it was telling people about it. And that was just a magical experience. I've been thinking back, kind of the history of the industry was we always replace two parts with one, all right? So the thing you always were doing with the engineers was trying to say, "Well, what's the benefit to you of using this part?" So it was "mil standard 883 for free," right? Or it had better output current, or instead of being a 4-bit part, it was an 8-bit part. But these were kind of like incremental things. And they didn't have huge impact. But with the 2901, you were actually giving people a different way to design stuff. So it was a much more fundamental change than traditionally what we'd done. And when people got it, they would just get excited about it. So it was fun to go out and give the seminars. And you'd just watch people and you'd just see the lights coming on in their eyes as you would do this. So I mean, what I remember afterwards was realizing that to me, marketing, at least the fun kind of marketing that isn't about price and delivery, is just about if you've got a really good product, all you have to do is teach people how to use it, and the product sells itself. And that's what the 2901 did.

Laws: Tony, you had a very different career path from John going forward, becoming a distinguished statesman in the industry. What did the 2900 family do for your career do you think?

Holbrook: Well, it diverted me from starting a surf shop or buying a motel on the coast of Oregon, yeah. I had a marvelous and a rewarding career at AMD, 21 years. Did a lot of things and enjoyed with the company a lot of success in different areas. The 2900 family was really the starting point of all that and was, with the rest of these guys, just a marvelously good experience, a very enjoyable experience. You asked earlier, what did you come out of this with? And to me, as I reflect back on it, we got everything right with the 2900 family. We created a design concept that brought benefit to the customer. We supported it with products well designed, well crafted. We were able to make them, and we've gone through all the difficulties that that presented. We captured the attention of the customer base worldwide with promotional material, marketing material. And then we enabled the customer to implement the concept to their benefit with support tools, with software, with application notes, with a whole book of applications, of applying it, knowledge, with seminars around the world. I mean, we got it all right. And that's a real credit to the people who were involved in that, people in this room and a host of others. And I think from a self-fulfillment standpoint, just having the interaction with the people was the best part of it I think. You know, the experiences that we shared as we strove to achieve. And I think, as you saw last night, we all look back on it with great memories. So there we are.

Laws: Phil, any thoughts?

Downing: Well, we talked about all this close cooperation. I remember that when I joined AMD, I had my office, Bob had his office. It was the same office. I had to go in one door of the office. To go to see Bob, I could either talk across the desk or I could go around to the other door to the office, which was his door to the office. So our desks occupied the center of the office. So that led to some close cooperation as I recall, very close cooperation. And I don't want to make too light of that, because that was the function of almost all the groups at AMD at the time, were boxed in, close quarters and so forth because we had limited resources. And so we couldn't afford to have people spread out, each person to an office and so on. It kind of tells the story that you can have too much money. And too little money. We seemed to have just the right amount of money. Just the right amount of investment to do the job. I do remember one story that I wanted to relate because it came up, I think last night. And that is, when we did the second source, and I think it was with SGS Thompson, they had reverse engineered the part. And they had found some vias , basically connections within the circuit, that weren't open, they weren't connected. And they said "You didn't happen to put those in as a blind, did you?" <laughs> And of course, we had. So it showed how much interest there was in this part. The guys had completely reverse engineered the part before they signed the second source deal.

Laws: Thank you Phil. Sven, you get the last word.

Simonsen: It's a little difficult to find something totally new. But I will say I second what has been said here about the team spirit, etc. But I'll add two comments to it. I found later, by going out in venture capital how difficult it is to create that same spirit in a company. And that made me understand how unique AMD was. When we were in the middle of it we all enjoyed it, but I don't think we realized how unique it really was when we were right in the middle of it. But I certainly, later, being involved with many different small startups, etc., trying to get this atmosphere going, I realized that that's not usual. I wish it was much more usual. But do we have a minute more?

Laws: We have about a minute.

Simonsen: I think Clive or somebody earlier on talked about how the Russians have copies of these things. Each year in the first week of April, there used to be an exhibition in Paris called Le Composants. And one year I was there and a couple of guys come up and asked questions. They were very clearly Russians. I mean, the way they're dressed, the way they talk, etc. And so they asked lots of questions. And, you know, I answered because they were things you could read in any datasheet or whatever. Afterwards, one of the salespeople, they said "You'll hear from CIA about this." And I said, "Yeah, yeah. Right." I had only been back in the U.S. I think a couple of weeks when two gentlemen came and visited my office. They would never say where they came from. I suspect they were NSA people or something. But they had clearly shadowed these people. They wanted to know what these two gentlemen had talked to me about -- what did they want to know, etc.

Laws: So the 2901 had international appeal!

Simonsen: For all of us, I'm sure there are many of these sorts of experiences. But all in all, it has been a great run, the time in AMD. I really enjoyed it.

END OF INTERVIEW

