Oral History Panel on the Development and Promotion of Intel Flash Memory

Participants:
Niles Kynett
Stefan Lai
Bruce McCormick
Richard Pashley

Moderator: Jeff Katz

Recorded: November 20, 2007
Folsom, California

CHM Reference number: X4364.2008

Jeff Katz: Today is November 20th, 2007. We’re in an external conference room at the Intel Folsom campus. We're going to do an oral history panel discussion about the development of the first Flash chips at Intel Corporation. We have with us, starting at the left, Niles Kynett, Dick Pashley, Stefan Lai, and Bruce McCormick, who were all members of the team at the beginning. They all did some things other than that and we'll have them tell you about that in a moment as well. The interviewer is Jeff Katz for the Computer History Museum. I'd like to start by asking each of the panelists to please give us a brief introduction to yourself with a little bit of biographical background, three to five minutes perhaps, describing where you grew up, where you got educated, what you studied and how and when you go to Intel, what else you worked on before you got to Flash. We'll start with Niles.

Niles Kynett: All right. Well, thank you. I grew up in northeastern Ohio. I grew up in a small town, Salem, Ohio, halfway between Cleveland and Pittsburgh in what's now known as the rust belt because it was all steel (industry). I went to college at the University of Cincinnati. I did that 1972 to 1977. Studied electrical engineering. The reason I studied electrical engineering is I was very happy with mathematics and that seemed to be the most math intensive field I could study and still be able to build something when I got done with my education. When I became a senior, we had a really interesting semiconductor lab. They were building bipolar but I realized that leading companies were building MOS So, for my lab project, we built an MOS (inverter) and it actually worked. That was kind of fun. I came to Intel as my first job, right out of school, worked in the D-RAM [Dynamic Random Access Memory] group as a product engineer. I worked for Bruce McCormick all those years ago. After that, I met a gentleman by the name of Dick Pashley, and I went into Technology Development designing SRAMs.[Static Random Access Memories] Then I went over to the microprocessor area and started with a co-processor called 8089, moved to the 80186 [microprocessor] team, then was on the 80386 team. At that point, I moved to Folsom and found Dick in the hallway again. I asked him what he was doing. He said he was doing this Flash thing. And I said, "Well, you're going to need a design guy, right?" And he said, "Yeah," so we kind of made an agreement. I moved over and worked with Dick on Flash. That was how I got over to the Flash group.

Katz: So did you come to Folsom with the intent of working on Flash, or something else?

Kynett: No, actually, I came to Folsom from the 80386 team, which I thought was one of the best jobs in the world. But I was starting a family and felt Folsom was a nice environment to come to. They were starting a logic group up in Folsom. I was trying to stay away from memories when I came to Folsom but it only lasted a year until I found Dick.

Katz: Thank you very much. Dick, tell us a little bit about yourself.

Dick Pashley: I'm Dick Pashley. I grew up on the east coast during my high school years, but my dad was in the military so we traveled everywhere. I've lived in Japan, I've lived in Germany and lived in over 15 states in the United States. I ended up going to the University of Colorado for my bachelor's and got a bachelor's in physics. For graduate school, I went to Cal Tech where I chose to major in electrical engineering. My thesis was on ion implantation. My last year at Cal Tech, I was asking why there was no class in the [design of] integrated circuits. Carver Meade was there at the time. I went to Carver, who said he would teach a class if we got enough students to take it. So we got a petition going. We got 25 students to sign it and Carver agreed to teach it. So the first class, all of us show up and we're ready to be taught how to design integrated circuits-- remember, this is back in the early '70s and there were no circuit design courses in MOS design. And Carver says, "Surprise, guys, I'm not going to teach you how they're designed, you are going to design one." Promptly three-quarters of the class left and there were eight students. We each designed a shift register and, eight weeks after we started the design, we shipped the tape out to Intel and the chip was made. If the part worked, you got an A and, if the part didn't work, you failed the class. I still have fond memories of that class. Carver was the one that convinced me to interview at Intel. I interviewed at several other companies but Intel was by far the most dynamic. Oddly enough, they had the lowest salary offer but that's a whole separate story by itself. I joined Intel expecting to work in ion implantation in Fab 1. But, when I arrived, that job was filled. I was told that my new job was to be in technology development, designing an SRAM and developing the fabrication process for it at the same time. But I didn't have a clue how to do this. Even though I designed a chip at Cal Tech, I was completely mystified. For about three months, I felt like Intel was paying me for basically fumbling around in the dark. Ultimately I designed the part called the 2102 SRAM. It was the A version, which was depletion load technology. That part was a real learning experience for me, but it worked on the first try. I was very pleased. After that, my group started expanding and Niles joined me. We worked on developing the Intel HMOS technology and a family of SRAM products. About 1980, I was given the opportunity to leave static memory and go over to non-volatile memory. I transferred-- still worked for the same boss, Gerry Parker-- to head the non-volatile technology development group. And, very shortly after, Stefan and I talked. He ended up leaving, I think it was Bell Labs, if I remember...

Stefan Lai: IBM.

Pashley: He left IBM and joined us at Intel. He started a group that was investigating new non-volatile technologies. I'm sure he'll tell us more about that in a second. That technology is what led to Flash. And so, with that, I'll turn it over to Stefan, if you don't have any additional questions.

Katz: Thank you very much. Stefan, tell us how he managed to talk you out of IBM.

Lai: That was easy. <laughter> But, anyway, I was born in Hong Kong and I finished high school in Hong Kong. At that time Asia was not the most stable place in the world. I came to the U.S. for my college education, and I went to Cal Tech. Curiously, Dick was one of my teaching assistants or lab assistants at the time in one of my lab classes.

Pashley: That's true.

Lai: I majored in applied physics when I was at Cal Tech because I didn’t think I was the theoretical type. I'm sort of a tinkerer, and want to do things with my hands, so applied seemed to better represent my interests. After Cal Tech, I went to Yale for graduate school. I was saying to Jeff on the way up that the reason for this was because my wife-to-be had some chance of getting to Yale later we thought. That was a very good choice because, at that time, my thesis advisor, Dick Barker, was starting work on a Silicon tunneling device. This is actually very unusual in the mid-'70s. We were way ahead of the rest of the industry in tunneling work. As part of my thesis, I actually discovered an optical memory device based on a tunnel junction. After Yale, I had the opportunity to join either Bell Labs or IBM. I decided on IBM.I continued to work on tunneling as well as Silicon dioxide interface and damage mechanisms. I was in San Jose giving a paper in the Spring of 1982. Someone in the audience looked at what I was doing and, you know, at that time, my understanding was that there was a lot of change going on within Intel. There were a lot of people leaving to form companies and they're looking for someone with my background. They brought me in for an interview. It was actually a struggle for me because my wife, at that time, didn't want to leave New York. But we decided that it was a good opportunity. So we came to Intel in California. Within about a year or so, I and some of the people working with me, together we conceived of this idea for an electrically erasable EPROM [(UV) Erasable Programmable Read Only Memory] . This eventually became Flash. So, for the last 20 plus years now, we've been just working on this. Now we're shipping the 10th generation of Intel's ETOX Flash memory, and the 11th is in the works. It's been a real wonderful 20 plus years of what we did.

Katz: Very good. Thank you. We'll get more on the actual “Aha!” moment in a few minutes. Bruce, tell us more about yourself.

Bruce McCormick: Sure. I was born in Pennsylvania and I spent my later childhood years in Indiana, where I went to Purdue University. At Purdue, I studied electrical engineering. My goal was to understand computing from operating systems all the way down to silicon processing. Interestingly -- this was in the early '70s-- I was very intrigued by Intel's 4004, which I studied at that time. I was tailoring myself really to come to work at Intel to be a microprocessor designer. And, once I got my master's degree, I discovered it wasn't meant to be. I ended up being with Niles doing DRAM product engineering on memories. This has kind of flavored my thinking towards things over the years. I moved from that to applications engineering for peripherals and microprocessors. Then I moved to this Folsom site as a product marketing manager in the memory organization. I also got an MBA along the way. I was pretty much getting my career oriented around the marketing aspects of the technology.

Katz: Okay. Very good. I'd like to explore how the Flash idea started. As a context for that, since you've almost all mentioned that you worked on some other kinds of memories before working on Flash, tell us how that experience influenced you and possibly helped lead to Flash. Anybody have a connection?

McCormick: I could start if you want. My first experiences with DRAM. The whole idea --of storing a charge on a capacitor, then it leaking off, and then you rushing in to charge it back up again -- to me, just seemed crazy. Early DRAMs had a lot of difficulties with data volatility, which tainted the technology in my mind. So, personally, a non-volatile technology had particular intrigue to me because it eliminated the DRAM volatility question.

Katz: Well, as a computerist or a would-be computerist, that made a lot of sense.

McCormick: Yes.

Katz: The computer system designer didn't want to have to be fussing with the memory all the time to keep it refreshed. Anybody else have a comment on how other memories influenced you?

Pashley: I think what was interesting is that, when I made the transition from volatile memory, namely static RAM, over to non-volatile in the technology development group, I immediately recognized that the non-volatile EPROM device had a major impediment. You had to basically pull it out of the system, shine ultraviolet light on it for 15 minutes in order to erase and re-program it. If it happened to be located in a remote location, like the top of a mountain, you'd have to send a service man up there. It was fraught with just all kinds of problems. Though it was a wildly successful memory for Intel. I mean, Intel made a lot of money off EPROMs. But, as a practical memory, it was very difficult to use. We had a memory that solved that problem at the time. It was called an Electrically Erasable Programmable Read Only Memory [EEPROM]. But it had two transistors in the memory cell, and thus it took up more than double the area of an EPROM memory cell. The way silicon economics are, even though it looked like it was only a factor of two in area, it was almost a factor of 10X in cost-per-bit. So customers came back and said, "Hey, yeah, I really want the functionality of this E-squared memory device, but I don't want to pay ten times the price per bit of the EPROM. I'd love to have the functionality of E-squared or close to it with the cost structure of an EPROM." Immediately there was this "Aha!" and -- Stefan really ought to address this better than I can -- we began the search for the holy grail of non-volatile memory.-- meaning a device that's cheap like an EPROM but can be electrically erased in-system rather than having to pull it out of the system. Stefan and many other people in the industry at the time were working on this, Companies like Xicor, SEEQ, Toshiba, were all searching for this ultimate device that could be a true, cheap, electrically erasable, non-volatile memory. To me, that was probably the wakeup call that steered me in the direction of what would become Flash, but I didn't know how to do it. That's where Stefan comes in.

Katz: So there was an apparent need for this capability to be nonvolatile and in-system changeable. How were you approaching the pursuit of that goal? I guess we'll start with Stefan.

Lai: When I came to Intel it was in some ways with a blank sheet, because I had been in academic research and then, at IBM, also academic. I had never done any chip design, chip manufacturing at all. When I came, I spent the first year trying to understand the E-squared technology. At that time, there was also a group called Advanced E-squared group, that were working on a novel concept to get a small memory cell. It turned out that the approach was not workable. One the thing that we found out over the years was, you have, really, two mechanisms you have to deal with. One is a programming mechanism. The other is the erase mechanism. The question is how to integrate the two mechanisms into one single structure. An "aha!" moment came when we looked at what we're doing. Basically EPROM already has the programming mechanism and works extremely well. And it was scaleable. The question, at that time, was how do we make use of this particular device? The whole idea was to add an electrical erase function to the existing EPROM. It was sort of a revolutionary idea in some ways because people did not believe it would ever work. But we were able to demonstrate it in a relatively short time. Because the structure is so similar to the existing EPROM, it allowed us to go from EPROM manufacturing to Flash manufacturing very quickly. A lot of the development work could be done very easily. We didn't need a lot of new and different processes and different designs. It took on a life of its own once we started working on it.

Katz: It appears that there was a concerted effort to try to solve this. It was not something somebody just stumbled over?

Lai: Oh, definitely. I can tell you that. When I was hired, I was told that my job was to find out how to make the E-squared PROM scalable.

Katz: I see.

Katz: Now, Flash, as we call it today, had earlier supposedly been invented by another company, Toshiba. But they didn't make much out of it and, apparently, it didn't work well. Can you describe how that affected what you guys did?

Pashley: The Toshiba product was difficult to use. And so, while it was a breakthrough in that it truly was a one-transistor non-volatile cell, it was extremely difficult to use and...

Katz: In what way was it hard to use?

Pashley: It required high voltages and -- Stefan ought to comment about the other side of it -- it was not manufacturable, so their yields weren't very good. They had a product that looked good on paper, and the press talked about it a lot. And you had other companies like SEEQ that were actively working on it. But their product, also, was very difficult to use and cost a lot. Remember, the whole name of this game was to build it cheap. Here you had two companies, Toshiba and SEEQ that were attempting to do it but couldn't really make it work in the real world.

Kynett: I think one of the things that we also did very early was particularly useful. Possibly I had a little bit to do with that, because I had come from the microprocessor design background. When we were starting to deal with the definition of our first commercial product, which was the 256k, one of the things that we worked really hard on was this idea of how do I make it easy to use in the microprocessor system. Dick pointed out you really don't want a repair man to have to go up to the top of Everest and change a program or be able to record some data up there. So it was kind of a simple idea: how do I get this part to talk directly to a microprocessor in-system? We came up with a command port idea, in which we were able to use the pins in a clever fashion to be able to allow the microprocessor, with standard microprocessor timings, to give command functions or signals to the Flash memory. This would allow the Flash memory to go off and do its erasing program and the microprocessor system and software would not have to have any other overhead circuits associated with erasing or writing the non-volatile memory.-- And you could do the erasing and writing right on your microprocessor system board.

Katz: And that on-chip erasing and writing mechanism did not exist in the previous efforts by Toshiba?

Kynett: No.

Pashley: No. Before we launch into the products, though, not that I don't want to do that at this moment, I think we ought to spend a couple minutes talking about what occurred when this "Aha!" moment arrived.

Katz: Please do.

Pashley: Stefan's right, it was a breakthrough. His group was all excited. As soon as I looked at it, I was jumping up and down saying, "This stuff, not only is it going to be cheap because it's a very small device, in theory it should also be very manufacturable. We should be able to take a factory that's building EPROMs and, since there were only a couple steps that were different, we should be able to wave a magic wand and immediately be building Flash memory in the same factory." Stefan and I went from our Santa Clara lab up to the Non-volatile Memory division, which was located here in Folsom, and gave a presentation. We were fully expecting to have the red carpet rolled out in front of us and people to be cheering and madly embracing what was, you know, the holy grail of nonvolatile memory, which had just arrived. But we were not greeted with cheers. Matter of fact, we were greeted with boos. Basically we were told that "customers would not buy it, it's not exactly what they need, plus it's probably not going to work, and, Gee, have a nice day but go away." So we walked away surprised and depressed...

Katz: Why do you suppose they didn't like it?

Pashley: Oh, well, that's a whole other story. <laughter> I have my own theory and it's this: The NVM group had spent their whole career developing a technology and developing a business, which EPROMs were -- and they were wildly successful -- there were probably 300 people working on EPROMs here in Folsom. Then here you have two upstarts come in the door and say, "Hey, we have a better mousetrap and this mousetrap can replace all the stuff you guys have been working on." There is a tendency, by the people that built this business, probably to think, you know, gee, that could replace my business. And, worst case, it could replace me. If you think about it, the EPROM group didn't necessarily know what that the new proposed technology was. When two new people came in with this brand new technology, it was a threat not only to the business but it's a threat to the people themselves. That's my personal speculation.

The bottom line is they said they weren't interested, and we were confused. We went back to the Bay Area. Intel's then-president Gordon Moore had graduated from Cal Tech also. I knew him fairly well, so I called him up and said, "You know, Stefan and I think we have this great thing, but we didn't get the positive response we anticipated from the Intel non-volatile memory division. We'd like to have you hear what it is, and give us advice on what we should do with this technology." We had lunch with Gordon. He listened to the technology story from Stefan. At the end, he didn't get all excited, he just said, "I'll take care of it". Stefan and I kind of looked at each other, thinking what does, "I'll take care of it" mean? <laughter> Gordon said, "Just sit tight."

Two weeks later, the phone rang. It was George Sneer, the manager of the Nonvolatile Memory Division in Folsom, who offered me a position to form a startup. Remember, I was a technical guy. This was, like, completely in left field. I said, "Well, wait, does that mean you're interested in the technology?" He said, "This means you get to put your money where your mouth is. If this technology is really, really good, you get the opportunity to lead the business." I had to think seriously about it because, obviously, I didn't have any training in being a business manager. I ended up making the decision to startup the Intel flash business in Folsom. But Stefan deserves most of the credit for getting the technology off the ground. And so does Bruce from the marketing side and Niles from the design side. Once the decision was made, I moved up to Folsom and we formed a group. Then we'll get into the details of how the product evolved. But Gordon Moore deserves the credit for recognizing the potential of flash in spite of Intel's "get out the memory business" trend in the mid 80s. If it was not for Gordon and a band of very dedicated pioneers, Intel would have never entered the flash market.

Katz: So let's pause at that point in the sequence of events and discuss how the rest of you got involved. We know that Stefan and Dick were instrumental in trying to sell the basic concept to senior management. Were Bruce and Niles involved at that time and, if not, how did you get involved?

Pashley: Well, there's one other thing that I'd like to add that I omitted just because I didn't think about it a few minutes ago. It's been, after all, 20 years. Even though we were all excited about this stuff and we got the approval to work on Flash, upper management came down right at the last second and said, "Oh, yes, Flash is exciting. But we've just signed a deal with Xicor Corporation to develop a single transistor E-squared Memory." We said, "Huh? What? We've got this stuff that's guaranteed to work. It's really exciting." But we had to work on the Xicor technology first. Stefan's project was forced to be what I call a "skunkworks" project, a back room project. He had to give up some of his people to work on the Xicor technology He had a very small group to continue to work on Flash. But the most of the people in R&D had to work on this Xicor single-transistor triple-poly cell. None of us personally believed in it, even Gerry Parker, my boss, didn't believe in it. But we all had to salute the corporate flag and work on this technology for a year and a half. In the end, it didn't work. Stefan knows the details of why it didn't work but...

Lai: Well, actually, I think, technically, it did work. We got it to be a working memory. In fact, that particular project spun off the first product by Sandisk, which happened to be based on the triple poly process. The person who worked on it was Dan Guterman. Before going to Sandisk he was one of the people who worked inside Intel on the Xicor triple-poly project. Another person at Intel on the Xicor project was Bing Yeh. He later went on to start the company SST. Their products were also based on a triple poly process. So that project, technically, was successful. It actually spun off two companies, in some ways.

The problem was the cost. Even though it was a single transistor in reality, it was three transistors merged together. So it was really not any smaller. Whereas our ETOX Flash was a real, true single transistor cell. There's no extra transistor. By the way, just double back, the Toshiba technology was a three-poly technology. It had an extra layer of poly for erase. When you had extra layers and the poly-to-poly erase, which is the Xicor technology, it really did not work very well in terms of the erase control. That's because it relied on the randomly occurring bumps in the poly for the tunnelling to work. It's very difficult to control those bumps. Whereas in our technology, with two flat surfaces it was much easier to control the tunneling process.

Pashley: Did it also have a problem with wear-out because of that?

Lai: Yes. It's easier to get trapping around the asperities, as the bumps are called, So the wear-out would be faster. That's why the Toshiba technology did not work well, and also that's why the Xicor technology is expensive. Dick was right. Now, I remember I was doing a pre-deal evaluation of the Xicor program on a Friday, talking to Gerry Parker at 5:00 p.m. from the Xicor site. I was having a phone conference with Gerry. I told him, "I think this thing will be very difficult. We'll have basic problems." He said, "Okay, Stefan, go home and sleep on it." And then, Monday morning, I walked in to Intel office and found out the deal was signed already. And, by the way, I was assigned to work on it. But Dick told me, "Stefan, if you manage to also work on the ETOX technology, put it under the radar screen, I will support you. I will fund you." So we ended up, for that one to two years, doing two parallel programs. Then, it turned out, because the ETOX Flash technology we had was actually simpler and easier to make, I could have fewer people working on it than on the Xicor technology. After two years, both programs were technically successful. But when we sat down to do the cost estimate of the two technologies; one was definitely a lot cheaper than the other one.

Katz: So you were working personally on both projects?

Lai: I was the manager for both projects.

Pashley: Yes.

Katz: Interesting. Okay. So how did the other two guys get involved? Niles, were you working originally on part of the skunkworks project?

Kynett: No, not at all. I got involved for a very simple reason: relationship. I'd worked with Dick before. I knew how he worked. I had a lot of respect for him. When I worked in SRAMs for Dick, I knew that, if we needed something, Dick found a way to get it. When I had been up here in Folsom for a while, I was trying to decide, right about that time, whether I was going to go back to the Bay area and get back into the microprocessors, or whether I was going to stay in Folsom and find a different job here. That's when I found Dick in the hallway. He described what was going on. It was pretty exciting to me for a number of reasons. One, I was going to be able to get in on the ground floor of a new technology, and that was pretty exciting for an engineer. You love to be able to do something nobody else has done before. That's the reason I came to Intel in the first place. Then, quite frankly, it was my relationships and my respect for Dick that drew me into Flash. As I mentioned earlier, I've known Bruce since my original first day at Intel. We're still good friends. And Dick, I hang out with him personally at his house at least on a weekly basis. So it was the the relationships...

Katz: At what point in the project did you come in? How far along had it got?

Kynett: I was working on the non-skunkworks project, on a single chip microcontroller called 80A2 that was using the triple poly E-squared technology. One day, Dick came in. I think he was giving us a presentation that he was giving IBM that had helped to put our executive staff's mind at the right technology. I can remember him coming in, saying, "We're going to do this Flash thing. It's really cool. Don't worry about this other thing any more. That one's cancelled. We're doing this Flash thing." When it came aboveboard, we went full bore onto the Flash technology. It was just night and day. One day, we're working on one project, next day, we're working on the other.

Pashley: And I think Niles brings up a real important point. During this time period, Intel was going through the painful process of killing DRAMs. They also were struggling financially. So businesses that were kind of marginal were on the chopping block. I remember going into the executive staff to sell our program to them. Right before us was, I think, the Bubble Memory Group...

Kynett: Bubbles.

Pashley: And the poor Bubble Memory guys got hatcheted. I mean, they walked in, came out two hours later, and they had big gloomy faces. Well, we went in there and we survived. But we were told, "You've got to make this stuff happen. We're not going to give you three or four years to prove feasibility and profitability. You've got to do it quickly." So all of a sudden, Stefan and Niles and Bruce had to gear up to build the product. Not just from a feasibility standpoint but get it to market and show that it truly was a revolutionary technology that could be built in high volume and had a market. that created a brand new atmosphere in the group.

Katz: I'd like to explore the market development in a few minutes more. But let's stay at the beginning, where the project was when it started gathering its own momentum. Stefan, you had demonstrated that the technology works and...

Lai: That's right.

Katz: ...it's manufacturable, and better than the Xicor equivalent?

Lai: That's right.

Katz: At that point, was there a product, or just test chips, or what?

Lai: There was a 64 kilobit test chip. Niles did the real first product. That was the 256k. But it was based on the 64k test chip, I think we actually qualified that product from a reliability product view but we never really shipped it in any volume.

Kynett: I think it was almost like an EPROM replacement.

Lai: Yes.

Kynett: It had the EPROM pinout.

Pashley: And customers were allowed to kind of evaluate the technology. We could give them the test chip product, they could put it into a system and say, "Oh, yeah, this is what I want" without us having to spend a lot of money on designing an official product." This way we got immediate feedback of what they wanted...

Katz: Roughly what time, what year was this, when you had demonstrated that the technology works with the test chip?

Pashley: 1986.

Katz: 1986. So then you started working on the full-blown 256k, which took how long?

Kynett: Oh, about nine months.

Katz: Nine months.

Kynett: Yes, we had less than a year to actually get this thing designed, out the door, and into customers' hands. I can remember Andy Grove coming up, -- you'll hear more about that, I'm sure from Dick and Stefan -- and putting me on the spot very early in the project. He said, "Okay, what's the die size, what's the performance, and when are you going to tape it out?" I thought for the two or three seconds that Grove would allow you to think, and I came up with an answer. That answer became a rallying point for the design team. We actually hit the specs exactly, and taped out on the day I said we would tape out. It was pretty intense.

Katz: Very good. You didn't have to delay it in order to hit that date, did you?

Kynett: Oh, there was no problem about delaying parts at all, never. <laughter>

Lai: Well, there was this one that goes from tape-out to customer's hand in 15 days.

Pashley: Oh, two weeks. I remember that.

Lai: It was either one Meg or the 256k.

Pashley: I think it was the one Meg.

Lai: It just blew all kinds of records, from taping out to a customer's hands.

McCormick: Thank you for doing that.

Lai: It was incredible.

Kynett: Well, we had dedication that you wouldn't believe. There were people that moved up from Santa Clara, ended up living in hotel rooms or apartment complexes while we were doing this. Everybody was working nights and weekends. It was around the clock on the one Meg project.

Lai: And we greased a spot in the fab. We literally had a test tape-out in the fab, to make sure all layers would be ready to process, and just fly through the fab in about nine days. It was an all-time record.

Kynett: The processes, as we were going through the fab, were timed precisely. They got the wafers to the point where they needed a new plate exactly when we had a plate ready and the plates were coming out from the manufacturer. In another 24 hours, you get a new plate and our process was just in time to use it. Everything was perfect. I've never experienced anything like that since.

Pashley: Actually, there's a memo that I wrote congratulating the team and it's probably worth reading, if I could?

Katz: Please do.

Pashley: It's relatively short. It starts off, "Congratulations on your one Meg first-run functionality. You have done the impossible by hitting the ISSCC," that's the International Solid State Circuits Conference, "deadline, and allowing us to get samples to customers before year end. Just two months ago, most of us thought that the Flash one Meg first run would miss the November ISSCC deadline. Then something I will remember for the rest of my life occurred: you, the one Meg team," and this is Niles and his guys, "rallied to the challenge and said the job could be done, and laid out a plan to accomplish the impossible. Through a lot of hard work and self-sacrifice like working weekends, overnight, and the Santa Clara team moving to Folsom for six weeks, you made it happen. In fact, you made the Flash one Meg a hard act to follow, going from design start to tape out in a record seven months." And that's pretty darn amazing, to do it in seven months. That's editorial comment. "I was especially impressed with the high level of teamwork displayed for such a large project. Your dedication has inspired others to go the extra mile on the one Meg. The tape out crew fractured the device in two days. The mask shop generated plates in three days and the Fab-1 team supported and processed your first run-out in a record 10 days. Putting it all in perspective, the Flash one Meg went from tape out to first Silicon in 12 days."

Katz: That is remarkable and a tough act to follow, I'm sure. Let's go a little bit to back to the 256k. Bruce, tell us how you got involved in that. Were you already on board at the test chip time or had something else brought you in?

McCormick: I was the product marketing manager for EPROMs in the early 1980s. As Dick pointed out, it was a very lucrative business for Intel. Part of the background here, and a key point, was that Japanese companies were targeting Intel, trying to get Intel out of the memory business. Prices had gone down precipitously. EPROMs became very, very inexpensive. Around the 1986 timeframe, I was the product marketing manager for N. V. RAMS and E-squared PROMs. My "Aha!" moment came in that year. When Intel, as Dick was pointing out, was under a lot of financial stress the company made the decision that they needed to get out of the memory technologies. We had all these non-volatile technologies but we weren't able to get any of them to be a real winner in the marketplace. We had Bubbles, we had N. V. RAMS, we had e-squared PROMs. The Flash that Stefan had with the test chip was obviously a breakthrough. You could see there was functionality taken from the most expensive technologies, as Dick pointed out, all the way down to what had become the cheapest technology because of the competitive pressures the business had gone through over the last couple of years on EPROMs. They had become highly commoditized. We were struggling, as a marketing applications team, asking ourselves “what can we do with this? We've got all this capability but nothing to go with it.” Finally I said, "Well, how would you design a computer around this new technology?" That was my background. You could see that you could boot from it, you could read applications from it, you could store files in it, you could update files in it, and you could turn it off and back on and find everything is still there. And you said, “Whoa, that's almost everything in the computer except for cache memory or very cheap archival storage.” The change that I had was I started thinking of Flash as a block-alterable, non-volatile RAM, and of E-squared as a byte-alterable non-volatile RAM, and of true non-volatile RAM as infinitely rewritable. It was a hierarchy of cost. You could actually build a computer system with only those technologies: with cache RAM, and byte-alterable and block-alterable non-volatile memory. I got very excited myself in 1986. We went out to convince customers -- well, we worked with IBM to see if the concepts made sense from an operating system point of view, and a computer architecture point of view. They agreed. And they also influenced Gordon Moore, during that time when Intel technologies were being shot, to further go after this one, because it was really the future and it wasn't just something in a lab.

Katz: At that point, had the 256 been done yet?

McCormick: No, this was prior to that chip.

Katz: ...Had the 64k test chip been done yet?

McCormick: I think it was between the test chip where we saw that the technology concept worked and the first 64k.

Kynett: Yes. I recall, at that time, Tom Cruise was doing the Top Gun movie, which just come out. We called this thing Top Gun. Dick came in dressed in his air force uniform and made this presentation for us all.

Katz: Very good. We should have had you dressed up for this session. <laughter> All right. So here we are now having proved the concept internally and no customers knew about it yet except maybe a few to whom you'd previewed it, like IBM. What pressures were there, if any, to make it work other than your own self-satisfaction? Was the company depending on it, or were the customers expecting it, or did you just think it ought to be done?

Pashley: Remember our comments earlier about Grove and the executive staff breathing down our throat to prove viability. That was definitely a motivating factor. If I place myself back in that timeframe, it was more pride, and it was more that we, as a team, felt like we could make this happen. If Grove hadn’t put that constraint on us, the people in this room and the people that we were working with were all determined to make this technology succeed. This was partly because there were a lot of nay-sayers that were remarking, "Aw, you know, this stuff just can't work -- it won't see the light of day." So there was determination in Niles' team and Stefan's team and Bruce's team and the manufacturing guys, to take this technology and get it out there as quickly as we could to show people that it would indeed work.

Kynett: Yes, I think the thing that's really important here is that we were a very, very small team. Everybody from the most junior engineer to the most senior person got to make extremely important decisions. And they had the ability to actually design something, to create something. You don't always get that. Looking at things towards the end of my career, the last microprocessor I did, there were 400 design engineers alone. It was huge. We're sitting there with maybe 10 guys on the Flash development team, including the manufacturing guys and the marketing reps., It wasn't because Andy Grove threatened, because there's always another job at Intel that good guys can go to. What really made these guys give up home life and work evenings, weekends, come up from Santa Clara and live in Folsom, was pride. They wanted to do the job. You can't force people to do that. We knew that this was an opportunity that, you don't always get in a career, to create something, create an entire market, create new capability.

For the first three generations of Flash, it was amazing the different things that we were able to add to the technology to make it easier to use, and thus create a larger total available market. People were inventing things to use it for. You would go to the grocery store and it was obvious that point-of-sales terminals were using this stuff. Medical equipment was using this stuff. You could see that you had impact on society.

Pashley: Probably one of the biggest impacts that I remember is the 256k. When we developed that chip, it wasn't to just sell it to the customer. We felt it was important to make a technical statement. He and his team presented this paper at ISSCC, which I mentioned earlier. You should comment on the paper, Niles.

Kynett: Okay. These papers were really big deals for the engineers, for the technical people. If your paper was accepted to this conference, you were really at the pinnacle of your profession because there were only so many papers every year that were accepted at this conference. The 256 was an industry-first. We got to introduce the Flash technology, which Stefan's guys had worked on. Plus we got to introduce the idea of the in-system programmability of a Flash memory, which nobody had done before. So we were feeling very proud, not only because we're being able to use Stefan's novel technology, which nobody else could do, but we had created this new computer architecture as well. And, relatedly, the very next year, we followed it up with the one Meg paper. As an engineer, the pay is good at Intel. But you can sell yourself for more at a different company. But the opportunity to be able to present back-to-back ISSCC papers with your name is them? The opportunity to have your name in the same book as Edison and Bell? And, be able to be an inventor and have patents? That means a great deal. In fact, one of the most sought after rewards for the technical team is a simple plaque that is a brass etched or engraved die photograph of a, 28256 or 28F010 chip and your signature, silk screened onto it. Those were very, very important. The pride of being able to present this paper at ISSCC to your peers, be able to answer the questions that came from the audience with actual lab data, this shows that nobody else in the world could yet do these things, that was cool.

Pashley: What was the reaction from your peer in the audience? That's where you find out if you've got something that was important or not.

Kynett: Shocked disbelief.

McCormick: Yes, definitely.

Pashley: What were some of the questions that you got at the end of the presentation?

Kynett: Well, there were a couple of questions that were very important to this technology. One is, can you control the erase thresholds and the program thresholds to a high enough tolerance, so that you don't over-erase the part and destroy it? And then there's another phenomenon that we have, program and erase cycling. We were able to control our Voltage. Thresholds and control our part well enough that we could demonstrate with data that we had that we were able to program erase these parts for over a million cycles without causing the parts to fail., To this day, that's something that's very, very difficult for competitors to do. Difficult even for Intel as well, later on as the technologies moved on. There was another comment that I had heard later, years later, from a person who was in the industry. He said that the 256k paper was a watershed paper. He said it was as if, before this paper, there were EPROM and then, after this paper, there was Flash. He knew it was an inflection point in the industry. When you hear comments like that from your peers, people you don't know, who are respected in the industry, that's pretty powerful.

Pashley: At this same ISSCC conference, there was also a panel session in which Bruce was a panelist. I chaired the panel but the topic was, if I remember rightly, “Nonvolatile memory, replacing DRAM” or something along...

McCormick: Disc drives as well.

Pashley: Bruce, you should give a flavor of what occurred during that session. <laughter>

McCormick: This was the one that Eli Harari was on.

Pashley: He was in the audience.

McCormick: He was in the audience?

Pashley: There was a disc drive guy, I think, Connor was there, we had Masuoka-san from Toshiba, we had somebody from Xicor maybe.

McCormick: Okay.

Pashley: Maybe it was Bill Owen. I don't remember all the people but...

McCormick: Well, one of the applications that got us very excited was the potential of going after replacing disc drives. The position I had taken at that time on that was that, to the disc drive industry, we will never really replace you but you will be very, very scared of us. And I think that has been true ever since then, because it's as a looming threat hanging over them. It's a very large non-semiconductor TAM that can easily be captured. All you have to do is take the cost-per-bit down and you could gobble it up and replace it with solid state drives. So we had some very interesting arguments over this and...

Pashley: And nasty notes.

McCormick: ...and nasty notes and comments and... <laughter>

Pashley: From Andy Grove. <laughter>

Lai: A side note is that this is happening right now. It took 20 years but the replacement of disks with solid state is happening right now.

Pashley: It has arrived. Yes.

Lai: So just a side note from my perspective: it was interesting. I wasn't at the conference but I heard a comment about it that people were saying that this technology would never work. For us, it was just pride on our side, because we knew it worked. We knew that we had...

McCormick: We believed. All of us totally believed.

Lai: We knew that we had the secret sauce. In fact, we had the secret sauce for a few more years. Nobody else could make it until maybe 1990, when they eventually found out what the secret sauce was. It was very exciting because we were doing things that nobody else was able to do at that time. I remember the revenue growth, $1.6M, $6.4M, $25.6M, $100M. That was the first four years.

Katz: It's remarkable that you remember that, too.

Lai: It's so easy. <laughter>

Pashley: It's a factor of four every year. <laughter>

Pashley: We even had a plastic thing built showing the growth, with the bars per year. He's absolutely right. It was quadrupled every year for the first four years.

Katz: So, Stefan, while these guys were off bringing your great idea into a real product, what we were you working on?

Lai: Well, we wanted to keep Gordon honest, so we just followed Moore's Law. We just know that, after the one micron technology F010, it's point eight micron after that, then it's point six micron. As I mentioned earlier, we're now down to the 10th generation. We're now shipping 65 nanometer so the cell size over the last 20 years has gone down about 1,000 times from the first day to now.

Katz: As has the cost.

Lai: As has the cost. And the cost went down even more because of the increasing wafer size and the efficiency. I just kept busy cranking out a next generation cell every two years.

Katz: You give them smaller cells, and they thought about all these extra gadgets to put around the cells to make them easier to use?

Lai: That's right.

Kynett: Absolutely.

Katz: Let's tell a little bit about those ease-of-use features, if you can.

Pashley: Let's spend one moment here just to comment on translating that cell size reduction to cost-per–Megabyte. I do happen to have a number on that. The scaling that Stefan was talking about has driven the price of Flash memory from $1,000 per Megabyte in 1988, that's when we first got started, -- the 256k was selling for, if you built a Megabyte board, $1,000 a Megabyte-- to 10 cents per Megabyte in 2004. That's a 10,000X cost reduction in 16 years. Pretty darn amazing.

Katz: With only 250X size reduction in memory cell area.

Pashley: Correct.

Katz: So that says there had to be other efficiencies.

McCormick: Well, there was one that Stefan's being modest about: MLC capability.

That's multi-level cell. It's a concept of storing more than one bit on each transistor, not just having it as a one or a zero but four states, essentially, so you can carry two bits on each cell. The technology was robust enough that we could entertain that kind of concept. It came in at the 16 Megabit level, I believe. That gave us another step function in being able to lower the cost per bit, which helped the application to grow even faster. That MLC concept was spearheaded by Stefan and his team.

Katz: Before we get into all the extra gadgets that went around the basic cell to make the thing easier to use, I'd like to learn about getting the customers to start buying the first one. Apparently all of you must have been engaged, in one way or another, in convincing the world that this was a good thing to do.

Pashley: First a comment on something that was unique to our group. In a lot of groups, only the marketing guys would go out and visit customers. But Bruce, very early on, since he's from a technical background, recognized that there was a real value in taking the circuit designers, the technology development guys, the manufacturing guys, to visit a customer. You didn't sell a product based on specs, you sold it based on getting the company to bring their engineers into the room. Then you didn't have just a buyer in the room. You had the equivalent of the Intel technical team Bruce would bring to the room, which is all these customer engineers. You'd have the project people at, let's say, IBM or whatever company you were visiting, maybe GM. They would bring their engineers and jointly there would be a discussion of this technology and how they, the customer, could build a product to utilize it. Iit became a joint development of the end product. Bruce just did a really good job of crossing that barrier by dragging all of us out to visit customers.

Kynett: That was one of the real key things that I'd never had any other experience in, being on the design side. Myself and several other design engineers did go out and meet directly with the customers. We got to meet with their people that were making their machines. And not only would we talk about how this current generation could work, we'd talk about their problems, we'd talk about how to fix it on this current generation and how we might be able to help them out in the next generation. Looking at the overall cost of what they had to deal with, in trying to figure out how best to come up with their solution.

Katz: So that made it pretty easy to figure out what next gadget to put on?

Kynett: Well, yes.

Katz: The customers just told you what they wanted.

Kynett: Well, yes, but the hard part was they never really wanted to pay for it.

Pashley: Yes. <laughter>

Kynett: I can remember a couple discussions. Dick used to have Friday meetings for his staff. His staff's staff would come in and talk about the projects and things they were working on. I can remember, on the 28F010, Mick Fandrich and myself, who had similar backgrounds, probably as much from the microprocessor as from memories. -- sometimes we described ourselves as very frustrated logic designers.-- Bruce was aware that we were often trying to spend a lot of his transistor budgeting in ease-of-use gadgets. So, finally, Bruce threw his hands up in a meeting and, in just total exasperation, he said, "Is this decision in the critical path? Am I holding you up from tape out?" I looked at him and said, "Every day you can't make up your mind about this, we're slipping tape out." He said, "Oh, go ahead. Put your stuff on this machine. It's all right." <laughter> We had to be right, I think we did a pretty good job of that.

Katz: Okay. I think we're almost ready to change tapes so maybe we should take a break.

[video off then on]

Katz: Before the break, we were talking about introducing the product line and we got the engineers out to help convince the customers. What were the customers doing with it, Bruce? How did they receive those engineers?

McCormick: Very well. It was a time when, as Dick pointed out, it was an inflection point When you have this kind of inflection point, you need to get people involved from all different disciplines to really understand the value of the technology and how it could be used. We went to manufacturing organizations, we went to software groups, and we went to people that were using all the other different kinds of memory technologies. We came up with a brochure that positioned the technology versus all the other memory technologies.

Katz: I'd like to hear how that positioning went.

McCormick: Now or...

Katz: Not now but just hold that thought.

McCormick: Okay. Prior to launching the technology, our minds were exploding with potential uses the technology. At first, we personally felt, this is really clever, that we've come up with-- this idea and that idea. Eventually, I personally felt like I was being led through it. So I have to personally give credit to God on all this, it was a religious experience for me. Some of the things we envisioned before we ever announced: we realized, this could displace DRAM, and hard disc. That had many application uses. Anything that stored information. We thought about digital photography prior to going out to market. I believe we went to Kodak before we launched...

Pashley: Right, we did.

McCormick: ...which was a very fun meeting. Dick and I were both there. We presented the concept of digital photography, and how this technology was going to change that business significantly. We felt, roughly in the back half of the 1990s, it would occur. There were things like CCDs (Charge Coupled Devices for image capture) that we knew were coming. We had Bell Labs talk to us about a capability called perceptual audio coding.

Katz: Let me back you up. How did Kodak react?

McCormick: Very interesting. This was a group of people, who were all chemists. Everything was built around film processing. Their money model was built around making prints. I think they believed us but, at the same time, they saw every aspect of their business, their skill sets, being disrupted or would be disrupted. This was a company that was 150 years old. They had a couple people coming in, laying in front of them something that is going to happen about five to 10 years in the future. They could not refute it, and then they started comprehending the implications of it. It was a real struggle, I imagine, for them. They could see that, beginning in the late 1980s, or early 1990s, their company would have to try to experiment to manage the kind of change that was coming at them that Flash represented. It was one of the most enjoyable customer meetings I've ever been in, because you could see the shock...

Katz: Did you leave with a design win?

McCormick: They were willing to work with us on some advanced products. One of the concepts we had was, maybe we can bring the digital portion of photography back to the United States. Because Japan was already starting to get leadership. We can work closely with you, Kodak, and make this a win for the United States. There were some early products that they developed around some of our early Flash chips but...

Pashley: Didn't your group have a partnership with-- I've forgotten the guy's name but-- that actually developed Flash-based cameras?

Tim, Treadwell?

McCormick: Treadwell. Yes, that was it. Thank you very much. Yes, we did have a partnership with a key development manager inside Kodak. This was all prior to launch.

Katz: Can you tell us more about what other customers were doing with it then.

McCormick: Well, once we realized that we were being a little bit too aggressive in suggesting re-architecting the computer around Flash, we decided to move more towards an evolving way of getting new markets to be established based on Flash. Even though we had all these visions burning in our mind as to what could be done in existing markets. We really launched the product in a big way, having products in distribution to help establish the credibility to the marketplace. We came out going after a wide range of embedded applications. Some of the very early ones were things like Schlumberger doing data logging in one of those down-hole data loggers. They had equipment full of Flash chips to record information.

Katz: Wasn't that a very high temperature application?

McCormick: Yes, but we had lots of experience from an EPROM standpoint in being able to go after those applications.

Pashley: That's right. Automotive, too.

McCormick: Right. Automotive was another key one. Tektronix embraced it for their oscilloscopes. Medical instrumentation makers embraced it. There were some also early adopters in the hand-held environment.

Pashley: What did we do to get immediate volume? Originally, we didn't go after all these exotic applications, said that you mentioned. But there were some product features we put in so we could go after a certain existing market and try to eat that up, right?

McCormick: The very, very first product was an EPROM oriented replacement. One of the awkward things about this product was, you could look at it two ways-- you could look at it from its beautiful side, and that was as a block alterable, nonvolatile RAM; but then you could look at its ugly side, and that is as a very difficult to update ROM. By going after applications that were EPROM in nature, where there was already the mechanisms to be able to update the EPROM and a business reason the customer had for updating the EPROM, we could target some early ones aggressively as well.

Katz: Give us an example or two.

McCormick: Okay. Postage meters are a good example of that. In the original postage meters, people could weigh their package and determine what kind of postage to put on it to send out. Now, inside this weighing mechanism was an EPROM with a window in it that carried the rate tables, which periodically changed. The way that was being changed prior to Flash was someone would drive out in a truck, unscrew the plate, pull it aside, de-socket the EPROM, reprogram it or put one that had been programmed back at the factory, and place it in to install the new rates. There was already a phone line in this machine for billing purposes. The need was there and Flash was a natural fit. But it was a very difficult sell.

Katz: So you went to Pitney Bowes and told them about it. And they said, "Go away, we’ve got all these guys who want to go out and change chips?"

McCormick: Right. The servicing people were their own profit center. The manager of that group didn't want to have anything to do with having these people...

Pashley: The whole service group goes poof. <laughter>

McCormick: Interestingly, we were either changing a business model or changing an architecture of a system, which required a lot of engineer to engineer discussions...

Katz: So you needed an evangelist inside every one of these 19th century companies?

McCormick: Absolutely.

To see the vision and to help get through the issues and the risk and the barriers.

Katz: How did you recruit evangelists?

McCormick: We'd go in with credibility, with these three people here. We'd go in with excitement, and somebody would catch the vision. We'd lock onto that guy and establish a relationship, until we got it to a design win. Another aspect of making it work is, once somebody sees it, does something with it, you immediately go to their competition and rub their nose in it. <laughter> Marketing by intimidation was something we used quite extensively, as well.

Katz: How did you do in the computer industry? PCs were going strong by that time.

Pashley: BIOS [Basic I/O System] chip. That was the first...

McCormick: The BIOS chip. We had some applications in the disc but they it was a small niche. The BIOS was the mainstream entry into the computer. The one Megabit was where we started that. It was interesting because it was, for a lot of people, a radical concept to be changing the basic firmware that they really did not want to be accessed. It was the software that determined how the system was configured, and they wanted that to be frozen and stable when they shipped the computer. But people were having trouble with that, keeping it stable. They were having recalls,, in some cases, which were very expensive., them...

We convinced them updating a Flash chip was a better way of going about it. There were other companies that saw it from a different perspective. They felt they could configure their platform for specific markets right at the very end of their production line. The notebook markets were more oriented that way, where they wanted to populate various functions at the very end of their line and have different versions go to different customers.

Katz: Who were the first computer adapters?

McCormick: Different companies. IBM was one. Compaq was another but they did it on different points in their product line. I believe, at IBM, it was work stations, because that's where they had a recall problem. At Compaq, it was on notebooks. It was a different point at different companies. Then pretty soon you got them to see, it makes sense all the way across, up and down the line. And, within a couple generations, we had 100% attach rate

for PC bios.

Kynett: And that lasted a long time, too. I can remember [much later] doing a firmware hub for part of a chipset. There were some discussions within Intel about whether the firmware hub was really part of the chipset, or it was a standalone. It was pretty much a moot point, because when marketing asked their customers, ”Who do you buy your firmware hubs from?”, They would ask their purchaser and the guy would say, "Intel." <laughter> Okay. Moot point.

McCormick: One of the tactics we used is, we went to the OEMs’ customers of the computer equipment. We saw their IT departments, and convinced them that this needed to be on a check off list of their computers they bought. It had to have a Flash BIOS. Otherwise, their total cost of ownership would have been higher. In this way we had end-user pull, through our OEMs, for this as well.

Katz: Did any of the initial customer activity result in feedback about the basic technology that required changes from your point of view?

Kynett: Yes, The original parts, even though we had a command port, the microprocessor had to control the timings of things like Erase. Almost 100% of our customers ended up over-erasing parts, particularly when they were first developing their software. I think that led directly to the 28F001B product, in which we put the Write state machine on–board. We realized that, number one, customers were going to make mistakes. Number two, they didn't want to tie up their microprocessor for the length of time it took the program to erase a part. The benefit that we got out of that is, if we took all the Program and Erase algorithms internal that left the technology development guys free to use some exotic algorithms that could be then programmed into the Write state-machine. This would help our yields, and it would help our manufacturability and help us do things like MLC.

Katz: You could do more exotic things on-chip than you could

depend on the customer doing in his software?

Kynett: Yes.

Pashley: Now you could just issue a command that says Erase, and the chip would go off and do it and then come back and say, "I'm done."

Lai: I think the most important thing for our success was that we worked closely together. As a result, design was putting in the things that could help the technology and vice-versa. We got all this feedback from customers, about what they wanted. But every change has an impact on the technology or on the product feature set. We worked very closely together to capture and assess that feedback, and that was very important.

Pashley: And that was unique at Intel. Remember, this whole project started in Technology Development. The way Intel typically operated, prior to this, was Technology Development would develop the Silicon technology. Then the design of the products was all done in the profit and loss division. Those two groups didn't necessarily communicate on a daily basis. They'd issue the design rules out of Technology Development, they'd have the process files, but there was not a day to day communication. One thing that I think occurred very early on was the communication between Stefan's group and Niles' group, and Bruce as well, communication that didn't exist for other groups inside Intel. This was just because it was a brand new technology. That communication, which started when there were only 20 people in the group, stayed there. Even when the group grew to 500 plus people, there was still that cross-discipline communication.

Kynett: From the very earliest of times, we didn't throw things over the wall to the next group. We had a complete cross-function development team in from the beginning. We had Technology Development people in there, Design people; we had the Manufacturing people in there, we had Marketing people. Even when we were initially defining the product, those product teams were already there talking to each other. And I do have to admit that there's another gentleman who really helped us out a lot. He came from the group that we threatened, the old EPROM group. It was. Sherif Sweha. He was one of the lead design guys in the EPROM group. He was instrumental, in the early Flash days, in saying, "Here's the institutional knowledge that we have in the EPROM products that you need for Flash." And a large part of our success was due to that relationship that we developed with him.

.

Katz: He was a believer, even though his management may have felt threatened by Flash?

Kynett: He was. He was one of the early Intel adopters. He saw what was going on. He was a very intelligent person, very friendly person, who was really, really key to helping us out. I think his name should be mentioned as one of those.

McCormick: I have another example of customer needs, working through the design as well as the technology. That was the concept of Blocking. One of the earlier issues that we had was with that updateable EPROM concept. It was a algorithmic, so you had to have a software program to actually execute the Write. The EPROM was where that program was stored. But if you are going to erase the EPROM that stores the program that must be used to rewrite the EPROM, how do you get out of this dilemma? An obvious strategy was to move the Write program into RAM and that's where you execute it. But, if you do that and the system goes down (killing the RAM) right in the middle of your update, you've lost everything, and now you're dead. One of the first end system needs that permeated back into the design and technology teams was the concept of a Boot Block. There would be separate parts of the Flash chip erasable separately from each other. In one part we could store just the basic code for turning on the screen or turning on a floppy or something like that. This would enable you to get the update loaded in and reprogram the other portion of the Flash BIOS memory, and still make sure you have just enough to bring the system back up if the system fails during the update.

That was important to getting the technology embraced.

Katz: Where did that idea come from? Inside Intel or from some customer ?

McCormick: It came from a basic concern from customers, I think.

Kynett: I think we always knew we needed Blocking. When we saw the original use of the Flash for BIOS and what the customers were going through this was a real barrier. Customers didn't want to update their BIOS for instance, have somebody take the power down on the computer, and have a totally dead computer. That was unacceptable. It was the third part that I worked on, the 28F001B that had four Blocks, all with independent Erase. It had a main Block, it had the Boot Block and it had two very tiny Parameter Blocks,

McCormick: Right.

Kynett: And then we also developed Lock Bits. Remember the Lock Bits? They made it possible to absolutely lock the Block, too. Then even through accidental Erase commands, you couldn't go in and destroy that portion of the memory without removing those Erase blocks.

Katz: Was there a particular customer that architected that, or did they just present their problem and depend on you guys to architect it?

McCormick: I don't remember one particular one.

Kynett: There was a lot of discussion when we brought those data. In Dick's meetings on Fridays, such things would be talked about and there would be presentations made. Because we were a startup group, we had a lot of communication. Mick Fandrich was one of the key guys who worked for me and who had lots of ideas. Don Larson was in our group, and he also contributed a lot of ideas. Basically we had more ideas than we had silicon real estate for. So Bruce was the one who was always saying, "No, no, no, no," until we finally forced him into saying, "Okay, I give up. I surrender. Go ahead and put these two things in." That's another thing Dick brought to the table, too, saying we have to manage our risk. We need to get the product out, and we need to take just enough risk. But we can't put everything in the next generation. The idea was, because it's a new technology, we should get these products out once a year with new features. Each year, we'll find out how the customer uses it and we'll build upon it until we really understand the technology. Take a certain limited amount of risk and turn it very quickly.

Pashley: That was a problem that we had in Technology Development early in my career. There was a tendency to put all the bells and whistles on something. I call them “hood ornaments” my acronym for the bells and whistles. I If you put too much stuff on a product, it never gets developed. Then, by the time it does come out, there's this pent-up desire by marketing to put a whole other bushel load of hood ornaments on the next product. So if you minimize the amount of new features on a product, you can treadmill the competition. You can put a product out every year and then marketing has something they can put their new features on. They know that a train, if you want to look at it this way, is going to leave the station every year. So they just line up with their new features every year and put them on the next product. That's something that I think Bruce and Niles and the people embraced, which was key to this technology ramping so quickly. And key to being able to meet what I would call emerging customer needs. It's kind of like peeling the onion. You come out there with an EPROM replacement and somebody says, "Well, that's kind of nice but I'd like to have a Block" or "I'd like to have this or I'd like to have that". As time passed, more and more features were added but they were done in an incremental fashion. And they were done very, very quickly.

Katz: So all these features must have been causing havoc to your competitors. What were your thoughts about the competitors?

Kynett: We wanted to beat them. <laughter>

McCormick: We wanted them to die. <laughter>

Pashley: Keep them in the dust.

Kynett: We wanted to be king of the mountain forever.

Katz: Which, if any of them, did you take the most notice of?

McCormick: Well, as Stefan pointed out earlier, we were alone for two or three years Because we had so many ideas of what to do, and because of the experience we had gone through in other memory technologies, there was some fear that we would have Japanese competition on us very quickly. But it turned out that AMD was the real competition for us.

Pashley: Toshiba continued to struggle. They were more of a paper tiger than anything else. SEEQ never made it in Flash. I think they actually gave up.

Kynett: SEEQ we were worried about because of the Block feature, which was really wanted very badly by customers. But they disappeared and the main competition became AMD for the first five or six years.

Pashley: AMD was in mostly a copying mode. I mean, we would come out with a product and then, nine months or a year later, they would have a similar product. That bothered us because we didn't like it that they were able to copy a product that quickly. But they were pretty good at coming up our tailpipe on that kind of stuff.

Kynett: As a compliment to you, Dick, I remember one conversation that AMD had with customers. The customer was asking them, "So what are you going to do next?" and they said, specifically, "Go talk to Intel to find out what we're going to do next." <laughter>

Kynett: It was fun being a small group and being able to invent things. That was really the real attraction of this technology all the way from the technologists to Bruce. We were creating something new that other people hadn't done.

Pashley: Probably the biggest market or customer application after EPROM replacement was cell phone, if I remember right. What occurred there and what made the cell phone market a natural for Flash?

McCormick: The GSM phones were just getting started at that time. The software code for that was problematic. They had to turn it quite a few times, so they needed an inexpensive way of doing updates after the sale of the phone. Originally, on the cell phones, their idea was; "Okay, we'll do this temporarily, but we're going to cost reduce to a ROM as soon as we get stability." Well, stability never happened. We didn't see customers ever move away from using Flash and go back to a ROM. Then they started being more aggressive about using some of the Flash features that couldn't be found in ROMs, so they couldn't even contemplate moving back.. Some of the things that we had envisioned earlier on, for example, parameter storage in small blocks, as Niles pointed out. They started to be able to store phone numbers and things of that nature in the Flash. And they also started thinking about software applications that they could put in Flash, which could be updated later as well. Today, it's being used also for picture storage, of course. Cell phone usage of Flash started probably in the 1992-94 timeframe...

Pashley: Maybe even earlier than that. Nokia...

McCormick: It was Novatel, Nokia came in later.

Pashley: I remember, in 1992, we were capacity constrained badly. We had this phone call from either Nokia or Motorola, who demanded that Craig Barrett empty microprocessors out of a factory and put the Flash 8 Meg in that factory. <laughter> I thought that was really a riot.

McCormick: That particular year, we had a lot of constraints. Motorola had designed Flash into almost all their communication equipment from police radios to cell phones. They were the first major adopter that I remember were really trying to adopt Flash for cell phone. We had Novatel prior to them a little bit.

Pashley: And Nokia.

McCormick: And Nokia were the ones that went into real production with Flash in cell phones. There were also lottery terminals that converted, do you remember that?

Pashley: That's true.

McCormick: We had the police over there storming our distributors to get Flash chips because there was a shortage. We had voice recorders, black box recorders in aircraft.

Katz: You created all this demand -- so much that there were shortages. How did you deal with that?

Pashley: Well, it's interesting how the problem was created. Early on, we were able to cannibalize EPROM capacity. The EPROM was a product that was losing money at the time for Intel. So it was pretty easy for us to come in and say, "Hey, we need that capacity", because our product had profit. We were able to cannibalize the existing EPROM factories. But here we were, quadrupling every year. You can imagine it wasn't very long before all the existing EPROM factories were bulging at the seams with Flash. At that point, we said, "Hey, Mother Intel, we'd like more capacity" and the response was, "Uh... you see the microprocessor? You see its profit margin? You think yours is good, look at that. You have Intel stock. What do you think the answer should be?" <laughter> And so we understood. We had to go external to Intel to get capacity. While that looked simple on paper, -- because, basically this was a widely available EPROM technology -- going out and getting another company, or foundry, to build your product was not easy and...

Kynett: Especially to our standards.

Pashley: Yes.

Kynett: We had very, very high standards for all our products.

Pashley: Stefan dealt with this a lot closer than I did, so I'm going to pass the baton to Stefan.

Lai: Actually it was in two stages.   In the first stage, one of a senior Intel VP claimed that the Japanese companies were making DRAM so well they should be able to make Flash. Just throw a technology into the Japanese company and...

Pashley: A ball bearing manufacturer! <laughter>

Lai: In essence, it was a steel company or something. They tried to make our Flash. They couldn't make it. So we had a year of zero output, a whole year, I remember. We tried and tried to get it to work. In the meantime, the other thing that happened was the industry was converting into eight-inch wafers. Intel, at that time, had enough money to convert one new Fab to eight–inch. Flash was not in that one Fab. We knew that we had to go to eight-inch for low cost. We didn't want a second tier and third tier fab partner. We went to Japan and looked for a partner. And it's an interesting story. I found out that most people who are looking for a partner, look at technical capabilities. But, for Intel, that's not the criteria. The criteria are that the two CEOs have to be able to get along together. <laughter> We ended up partnering with Sharp, which is a very good company. But in the pecking order in Japan Sharp is not in the first tier. But Tsuji-san and Grove-san, at that time, had met each other at a conference, and they liked each other. So we went to Sharp, and struck a deal with them. We gave them the Flash technology. In exchange, they give us eight-inch capacity. But we had to end up leading Japan in eight inch conversion. We brought up the first eight-inch process at one of the...

Pashley: It was a brand new factory so, in addition to...

Katz: Sharp had not done eight-inch before trying to do it for Intel?

Lai: They had not done Flash before.

Pashley: Right.

Lai: And they had not done eight-inch before, either.

Pashley: They were installing eight-inch equipment that they hadn't used before in a brand new factory. And, on top of that, they were bringing up a technology that they didn't know about.

Katz: So did it work?

Lai: Oh, yes it did. We started the project in February and, by July or August, we had working Silicon. We just did it right, versus doing it cheaply. There's the difference.

Pashley: Yes. Meanwhile, while this ball bearing factory wasn't yet pumping out Flash, Bruce's troops were building the demand. The demand was growing at 4X per year and meanwhile we were topped out on capacity. Yes, Stefan' was doing his part by developing the next generation that makes the cell size smaller, so we could squeeze a little bit more out of this existing Intel factory. But not quadruple every year. It was very, very difficult. That's why we had Motorola calling...

Katz: Very easy for the competitors, though?

Pashley: Unfortunately...

Kynett: We were creating a great market for competitors.

McCormick: That's when the competition came in.

Pashley: Yes.

McCormick: Our customers begged us to not let the competition come in.

Pashley: We had probably an 85 to 90% market share. But when you go a year and a half not able to grow your capacity very much...

Kynett: It makes angry customers.

Pashley: ...it makes angry customers and it makes competitors like AMD very happy.

Lai: Yes, we opened the door for them.

Pashley: Yes.

Katz: Was all that wild success the thing that was behind Intel's missing the NAND-Flash market cusp? Or if not, what was?

Pashley: Well, I think part of the reason for Intel's disinterest in the NAND cusp was that Toshiba was having a very difficult time making it work. Stefan will talk about the technical reasons for that. From a market standpoint, they talked about the technology but customers couldn't buy any of the product, and it was unreliable. It sounded very good on paper, but they couldn't build it. At the time, the applications that could best utilize NAND weren't there. The only volume applications were the ones that Intel had created for the NOR technology, like cell phones and stuff like that. The NAND technology wasn't really in the market in a big way, and really didn't have customers. That changed, obviously, with the digital camera hitting the scene in the late 1990s. It is ironic that Intel's first flash product introduction demonstrated the use of NOR flash as "electronic film." In April of 1988, we introduced the flash 256k memory at the Eiffel Tower in Paris. A picture was taken of journalist in the audience, and then electronically, the journalist was placed on a beach in Tahiti by using a "flash memory card" made of 16 256k flash chips. I think this was the world's first demonstration of a digital camera using flash as "film."

Lai: One of my mistakes in my career was saying NAND would never work. The truth is, it didn't work to our criteria because it wasn't reliable. But we had the wrong mental picture in saying that the memory had to be perfect for all applications. The NAND players basically spent the '1990s trying to sell the customer the concept that it's okay for the memory to be not reliable. You can have a bad block as long as you use it behind a controller that corrects the errors.

Katz: That's the same argument that disc drive guys had for 20 years.

Pashley: Yes.

Lai: Exactly. But we were looking at one specific application, cell phones. Because of that, you know, we were not interested in NAND. We looked at the market again and again and again and we were convinced that, for this cell phone application which, at that time, was a big business for us, they could not use NAND. And they still cannot, for the code storage application. But I think what turned the corner for the NAND business, was customers' realization that it can be very slow, can be off chip, and they could use a controller to control the errors. So, for secondary storage, it became a possibility. We missed that secondary storage concept. But, on the other hand, we never — had sufficient capacity to meet our own NOR-Flash demand. So for us to go into something that really had little demand didn't make sense. NAND did not take off until the year 1999 and 2000.

Pashley: Right. And, even then, it was losing money for probably three or four years. The crossover between film camera and digital camera in sales took place in 2003. If you look at the NAND sales, you just see a huge spike when that occurred.

Lai: On the other hand, this is a classical problem when you have a legacy business. At that time, we were already established. We had $2 billion in Flash business. It was difficult for us to change that and go into a smaller business which is NAND. This is a classical problem for disruptive technology.

Katz: But was there no Pashley and Lai, who's championed NAND at Intel?

Lai: You're right. We actually would have liked to do that but we did not have enough resources to justify the return on our investment.

Kynett: That's right. Part of the problem was: how do you maintain the competitive treadmill that you had to do with the NOR technology, while investing in the NAND technology? We were already capacity constrained. Plus, now, if we were to develop NAND, there would be people resources to be put in place just to develop the technology and the chips.

Pashley: It almost requires an approach similar to what we did with Flash, a little startup on the side to allow it to germinate and grow without the burdensome old business consuming the resources.

Kynett: Yes. At that time, I don't think Intel was interested in developing any more memory technologies. **Pashley:** Yes.

Lai: I remember, in 1997, I think, we would come to Intel’s senior management recommending to go into NAND. It was vetoed by management. The market was not there then and we could not project it.

McCormick: There's a difference in the technology. The NOR technology was a Random Access technology, which had higher value to code execution type of applications. Whereas, with Sequential Access NAND, you would have to have a duplicate RAM memory you move it into, to get the random execution; otherwise, you become very slow. So the market of photography, which isn't code execution, it's data storage, took years before it became a really significant market to justify the investment. The NAND people had year after year of losing money and investing into this technology, believing eventually it was going to pay back. It took quite awhile.

Pashley: I don't think NAND makers were profitable until the early 2000s, probably 2002, or 2003.

McCormick: And they started in the...

Lai: Late 1992 or 1993.

Pashley: That’s 10, 12 years of investment.

Lai: The company that drove NAND was really Samsung. Samsung was a memory only company. It made sense for them to invest in a new technology where there is zero or very low market-segment share in NOR anyway. Whatever was coming in from the DRAM base could go into NAND. For Intel, the microprocessor business was number one. Everything. We had this relatively small memory business trying to maintain its critical mass within the bigger Intel.

Katz: Interestingly, Samsung also had an end market system business in the very consumer applications that needed NAND, cameras and media players.

Lai: They do. They do.

Pashley: Right. One of the jokes we used to have in our group was about Charlie Brown and the football. Intel very often would say, "No, you can't have this capacity." We already mentioned that we needed more capacity beyond what we could cannibalize from EPROM. Every now and then, the microprocessor business would hit a bump in the road and demand would be a little lower than what the capacity was. Intel's response was interesting. <laughter> Memories were viewed as factory fillers. Suddenly we would be told, "Oh, you wanted a factory. Well, here it is. We're going to give you half of FAB 18 or whatever it is. Please fill it as soon as possible with Flash." So Stefan and his guys and the manufacturing guys would roar over there, they'd work, they'd spend six months getting it all ready to run the first lot of Flash wafers. Then, all of a sudden, we would get word that, "Oh, you know that factory capacity that you've been working on for the last six months and you have your first runs in the line? The microprocessor demand is back up and so sorry," -- Charlie Brown with the football -- "it's now gone." And then you fall on your back. We went through that I don't know how many times.

McCormick: A number of times.

Pashley: It was comical.

McCormick: We got to be the guinea pigs on new things, too, like bringing up a new assembly plant. The company would decide, “We'll let the Flash guys do it first and then we'll use it for something.” <laughter>

Lai: I’ve always said that we have the industry’s best Flash technology, but our business model is not for Flash. It's a second thought in the Intel environment, in manufacturing and a lot of other things. It's difficult to get this oscillating capacity and the cost will go all over the map. Whenever they make change, it costs you money. That's the basic difficulty for the Flash business at Intel.

Pashley: I think we even were allowed to build a factory once. It may have been FAB 18.

Lai: Yes, FAB 18.

Pashley: Our P and L (Profit and Loss center) had to fund that factory. We had to spend about $100 million a year that was going into building this factory. When it was about three-quarters built we were told, "Gee, sorry, we need that for microprocessors," — We said, "Do we get our money back?" <laughter> It was a big laugh. “We're all Intel”, pat, pat, pat. That was water under the bridge. We not only lost the capacity but we lost our investment as well.

McCormick: As the only guy here that's still working at Intel <laughter> I feel I should defend the company. To their credit, they've recognized the dynamics of this business and, over the last year or two, they have made moves to ensure the success of this dynamic being understood.

Katz: Well, very good. We're running out of our time allotted. I want to give you all an opportunity to answer something I forgot to ask. Anything else that we ought to hear, about the beginnings and subsequent success of Flash at Intel?

Pashley: I'll just make one comment. I reflect back on the early years, when we were developing Flash. It was probably the most exciting point in my career. The first three or four years, going from Stefan's group inventing the technology to introducing the 256k product. That was probably the period that was the most exciting because it was brand new. Everyone who touched the technology and who was part of the group shared an enthusiasm. They were doing something that no one else understood, much less could do. It was like a mission from God. You were doing something that you firmly believed in. Many people outside of the group would pick on Bruce because he was so enthused about the technology. But they failed to recognize “Gee, why is he enthused about the technology?” There must have been something there. The entire group shared in that. At least I know I personally did. I viewed it as probably the most exciting portion of my career. The only other time that I would say came close to it was the original development H-MOS, which was the technology that Static RAMs and microprocessors used in the '70s.

Katz: We'll hold that for another session.

Pashley: Yes. Yes.

Kynett: I reflect on what Dick says. There were two highlights in my career. One of them is the period from the 28F 256 through the 28F001B. The 28F001B may be one of my favorites because Mick Fandrich, myself, with Bruce's marketing people and inputs from the customers, had a lot to do with the definition. A lot of our names are in the patent office now because of that Flash stuff, and that's a huge source of pride. The other highlight was when I was on the 80386 design team These two periods were probably the best times that I had at Intel. It's awfully hard to replicate that.

Pashley: And it's not just the technology. I think something we should impart here, is it's the people that made it unique Yes, it was kind of great that it was Flash technology and we were doing something that no one had ever done. But the team really gelled, and everyone enjoyed working with each other.

Kynett: Yes, it was one of the hallmarks of Dick's organization that we tended to work really hard together during work hours. A lot of us formed friendships that lasted almost half our lifetime. Even in the SRAM days, Friday afternoon, Dick had his staff meeting. After staff meeting, we went out to pizza and to West Valley College for volleyball.

Kynett: And we spent weekends together going to the auto races at Laguna Seca. And the same thing happened with the Flash team. There are friendships that exist today (20 years later) that came out of that Flash team. Yes, we did a lot of neat stuff with the technology, but we never lost our humanity. We never forgot who was really important: the people that we'd hired to do these jobs.

Katz: And the customers.

Kynett: And the customers.

McCormick: There was no politics within the group. People worked really well within our group. It was a very exciting time. One thing I observe is the industry is not done yet. There are many things to come.

Pashley: Agreed.

Katz: Are you going to tell us about any of them?

McCormick: Not on this tape. <laughter!!>

Lai: Flash, from a business perspective, has a checkered history at Intel. Sometimes doing well and then we were always going through changes. From my perspective, since I spent my whole career in the TD area, we have developed a world-class team. We worked closely and well together. We are recognized by the industry as a real first class team. Everybody I talk to says, you're really a very strong team. We have been able to develop technology in Japan, Santa Clara, Albuquerque, we wandered all over the world for 40 days, 40 nights and came back to and continue to out-perform and lead the industry in technology. That really says that the people we have in the group can’t be beat.

Pashley: I think one engineer really characterized that and his name, I think, was Owen Jungroth This guy would work all the way through the night. We caught him one morning sleeping with his head on his keyboard. He got up and his-- the keyboard was imbedded on his forehead. Underneath his desk was a sleeping bag. Believe it or not, this guy would not drive home. He would just sleep at his desk and he'd do this for a week on end. And it wasn't because his boss told him to do it.

Kynett: No, we'd never do that.

Pashley: He just wanted to do it.

Kynett: Yes. Well, he did live a ways away. He had a home in Sonora. So, once it got to be too late, there was no sense in him driving at that hour through those mountain roads in bad weather and everything else. I think eventually we told him, Owen, go get a hotel room.

Katz: Well, with that, I think I'd like to say thank you all very much. We've captured quite a good range of stuff today. I'm sure that our future researchers and museum visitors will be enthralled with what they hear.

END OF INTERVIEW

