Oral History of Ken Ouchi

Interviewed by:
Robert Garner

Recorded: June 10, 2015
Mountain View, California

CHM Reference number: X7519.2015

Robert Garner: OK, well this is June 6, the year 2015. I'm Robert Garner. I work at IBM Almaden research in San Jose. I've had a long career in Silicon Valley. I started in 1977 working at Xerox, and then at Sun Microsystems, and then at Brocade, and then at IBM in 2001. And I worked with some people whose names will come up later in the interview. And we're interviewing Ken Ouchi. I think I pronounced it right, correct? "Oo-chi?"

Ken Ouchi: "O-u-chi."

Garner: "O-u-chi." See, I've always pronounced it wrong.

Ouchi: That's OK.

Garner: And Ken, what I'd like to maybe start with is the very beginning. So we've got to do the basics. When were you born? Where were you born? And tell me about you upraising.

Ouchi: I was born just at the outbreak of World War II. And of course, many of us now are aware of the photos of the tar papered buildings surrounded by barbed wire, or photos of little children carrying their belongings along with their parents waiting for the buses that were going to take them to they don't know where, because the United States was now relocating 120,000 American citizens and legal resident aliens off the west coast of the United States to inland locations.

Now people-- you look at the photos, and you look at the faces, and they look like Japanese, and people wonder how could the United States do that? And what most people don't know-- in fact, many people want to forget-- was that the legal and social prejudice and racism against Asians, and Japanese in particular, was increased-- was even worse than what was done to the blacks. Japanese could not become American citizens. They could not own property in California.

Of course, the children born to them were American citizens, so some families acquired property through the name of their minor children, but because of courts, they were sued, and they had to give up their property. The Chinese had been in the US for over 100 years, but they were still isolated in Chinatowns. The Japanese had been in the US are about 50 years at that point, but they, too, were still isolated in Japantown. So when the attack on Pearl Harbor occurred, it was actually very easy for those in power to now take these Americans and move them out of their homes, destroying their lives.

Garner: So can I ask you a question?

Ouchi: Sure.

Garner: So your parents were in San Jose? Or what city were they in?

Ouchi: My parents lived in San Diego. My grandparents, along with my aunt and uncle, lived in Wilmington. My father and uncle had a large wholesale produce business. The point I'm going to relate to you is that much of how this environment was in the history of my family at that point and what happened was not known to us. And in fact, most children of my generation never knew that this really happened. We never learned of this until we were--

Garner: Adults.

Ouchi: Adults. And of course, this was all on purpose. In fact, this was the decision that was made by my parents and their generation-- was they were going to be the best Americans, and they were also going to ensure that their children were going to be the best Americans, because they had this mass amnesia where they refused to relate the bitterness that they may have felt about how their lives were destroyed, but they were going to have their children see them rebuild their lives and succeed, and implant in their children the drive to succeed as Americans.

Garner: Wow. That's a very powerful motivator and story.

Ouchi: And because of that, when I went off to Cal. Well, first of all--

Garner: Can I ask one more question about your--

Ouchi: Yes.

Garner: So indeed they were interned then?

Ouchi: Yes.

Garner: In which camp? In the one in the eastern Sierras?

Ouchi: We were-- in fact, I later learned my mother was carrying me at that point.

Garner: Wow.

Ouchi: And they were fearful that I might deliver in the trip.

Garner: On the way to the internment camp.

Ouchi: On the way. And so I was born in LA County General Hospital.

Garner: Wow.

Ouchi: With an Army guard.

Garner: Wow. What a way to come into the world.

Ouchi: Yes. My mother had been forced to stay behind. My father, and her parents, and my uncle and aunt went to Poston Arizona, and they didn't see each other for perhaps over three or four weeks.

Garner: Wow.

Ouchi: And then my father related that, on the train ride-- he was very much aware of what was going on in Germany. And so the train stopped in the desert. They set up a machine gun perimeter.

Garner: Oh my goodness.

Ouchi: And he was scared that they we're just going to take them out there and shoot them.

Garner: Wow.

Ouchi: I mean, this was the--

Garner: I hadn't [heard] this kind of story. Wow.

Ouchi: This is the environment in which they lived. Now, what most people don't realize is that many of the Japanese and American citizens did not have to stay in the camp through the entire war. And so after about a year, we moved to Colorado.

Garner: Did your mother end up going to Poston Arizona, as well?

Ouchi: Yes, so my mother was in Poston.

Garner: And you were there, too?

Ouchi: I was there, too. I was--

Garner: Implication.

Ouchi: And I know my mother related the fact that the Quakers were extremely kind to them. And in fact, she said that there was a little box with a blanket in it, and that was my carrier.

Garner: Wow.

Ouchi: And so she later met them in Poston, and they were there for about a year. Cameras weren't allowed, so there are no baby pictures of me. The only photo that they have of me at about a year old or so was with a smuggled in camera with me and a couple other boys playing around. What they had done in Poston the course was create large gardens, and created all the food that they needed, landscaped a lot of the areas, and actually made it a very livable place as well as you could during that time.

But I think when I was about a year old, my family, as well as my uncle, and aunt, and my grandparents, were allowed to leave. And so we lived in Colorado, came back to California at the end of the war. My father--

Garner: Were your parents able to reclaim their property, or had it been lost?

Ouchi: It was all lost. So my father realized that individual families could move into communities and be accepted, whereas if you tried to move a group in, it'd be very much more difficult. So I grew up in a town just north of Long Beach, which was mainly people who were engineers and their families who worked at Douglas Aircraft in North America.

Garner: A large aircraft industry they had there at the time.

Ouchi: Right.

Garner: Massive, actually.

Ouchi: As well as a dairy, a wholesale cattle processing plant. And we had a wholesale poultry business where-- in the very beginning, we raised all of the chickens from chicks until they were large enough, and then they would process them. Later on, my father was able to contract out the raising of the birds themselves, and so we had this cycle where once a week we'd bring in big trucks of chickens and process them, and they would go out. And we were working 5 and 1/2 days a week.

Garner: Wow, it's hard to imagine that today in the Long Beach area.

Ouchi: Right.

Garner: But it was very rural there, it sounded like, where you were.

Ouchi: Right, so this was this pocket that is now-- the dairy is gone. It's now a large trailer park, and there's a large middle school now where the cattle processing plant was.

Garner: Well, it certainly taught you about responsibility at a young age.

Ouchi: Certainly did.

Garner: He gave you lots of responsibility, it sounds like.

Ouchi: So I was working. I thought at the beginning that I was going to become a medical doctor, and then--

Garner: Because the parents may have-- that's somewhat of a stereotype, right? Become a doctor.

Ouchi: Yeah, well it was-- a very good family friend was our medical doctor, and I thought it would be really great to be able to help a lot of people. And then I began to realize his entire life was so devoted-- he never married. His only vacation was once or two weeks he would take off and go fishing.

Garner: Wow. Another high standard set for you.

Ouchi: Yes. And he was doing essentially retail. He could only help one patient at a time. So I decided-- somewhere in there, I became very interested in physics. So I went to Cal as a physics major.

Garner: Was there a book that was around that got you interested.

Ouchi: Well, I don't--

Garner: Some friend of the family or a newspaper article? The aircraft industry maybe?

Ouchi: Yeah, the aircraft industry was there, but my father-- we had things like Science Digest laying around the house all the time. He was the one who was telling us that there was this amazing kind of stove that you could now put things in, and close the door, and it'll cook it in just minutes instead of cooking on top of the stove. These were microwave ovens, but he knew about them. And so this was the environment I was-- so he was running this wholesale poultry company, but yet was very much aware of what scientific things were happening.

Garner: So about how old were you when you got interested in physics? About how old were you when you became interested in physics?

Ouchi: Probably when I was a junior or a senior. I did well in science and math, had very good English teachers. It was a very good high school, but they weren't able to give everything that I needed, and somehow my father was able to arrange that several of us actually went to the local community college. I went to chemistry and calculus. So essentially I entered Cal as a sophomore, graduating from high school.

And then at Cal, took all the physics classes, and I thought they were a lot of fun. And then I had probably one of the most interesting jobs that anyone can have, at least at that time. At Berkeley, there's a large proton accelerator called the Bevatron. At that time, it was the largest in the world. 6.3 BeV, if you understand what that means.

Garner: A billion electron volts.

Ouchi: Yes, which when you look at what's going on right now-- the Large Hadron Collider, it's like 2000 times smaller than that. Anyway, so they had these jobs for people who would-- they have this bubble chamber full of hydrogen and liquid hydrogen under pressure, and when particles would go through, they'd release the pressure, and bubbles would form along the charged particle path. This was the bubble chamber. Three cameras would flash, track all the bubbles.

Garner: They would curve. You'd have these curved lines.

Ouchi: Yeah, because they have a magnetic field, and you could see the momentum of particle by the curvature of its path. In those days, they had people-- I call them trolls, because I was one of them-- who actually looked at this film, found the events, and then later on we had digitizer that would now digitize the bubble paths, because there was this big 7094 that crunched all this information.

Garner: So now this is a summer job, because you're--

Ouchi: No, this was when I was in Cal.

Garner: But I thought that was Cal in-- oh, it was Berkeley. It was not--

Ouchi: In Berkeley. And so what I would do was-- I was kind of crazy. I would go to work at 2 o'clock in the morning on essentially Saturday morning and Sunday morning and work till 10:00, so I'd get 16 hours that way. And then I'd work four more hours on a Monday evening. And so I was working 20 hours.

Garner: How did you have time to do work for school?

Ouchi: Well, you just kind of-- plus I took a long time getting through school, because I was taking a lot of classes. They were really interesting. And to be very frank, I think it took me a long time to grow up and figure out what I really wanted to do.

Garner: So you were very much imbued with physics then. So you took years of physics.

Ouchi: Yes. And I was about ready to graduate--

Garner: Was it mainly particle physics?

Ouchi: Yeah, it was all particle physics.

Garner: Did it include quantum physics?

Ouchi: Yes.

Garner: So the whole gambit.

Ouchi: The whole gambit. The quark theory was just starting to evolve at that point. We had particles-- the names of just--

Garner: Up, down, beautiful.

Ouchi: Up, down-- no, no no. At that time--

Garner: Beauty hadn't been discovered.

Ouchi: Hadn't been discovered yet. What we had was lambdas, and sigmas, and xis And of course, what we later discovered were those were different combinations and resonances of up, downs, beauty, charm.

Garner: Of course.

Ouchi: Top and bottom. But anyway, so what I had noticed was I would have one of two kinds of careers. One was to be like these research guys, and I've seen careers just go down the drain because somebody at Brookhaven had discovered this stuff first, and all you're doing was just proving that they were right. Or I was going to get a job at Livermore designing nuclear weapons.

Garner: Many, many physicists learn that it's difficult to get an exciting job.

Ouchi: Yes. And even in particle physics, I would be just one more cog in this giant team of people because it was getting much more complex. What was going to become really big physics-- was really big physics. The guys at the top got their names, but there were 2000 people underneath them.

Garner: Did you ever hear a lecture by Richard Feynman or any of the--

Ouchi: I did not hear anything from Feynman. Feynman was down in Caltech.

Garner: Wheeler? Any of the other names during that time?

Ouchi: Alvarez.

Garner: Alvarez.

Ouchi: In fact, I was working for Alvarez at that time. He and his son were the ones who had posited the theory of the iridium layer and that the--

Garner: The dinosaurs.

Ouchi: Dinosaur.

Garner: The meteorite that released the iridium that--

Ouchi: Correct.

Garner: Ended the lives of--

Ouchi: So that ended--

Garner: Dinosaurs. We started our life as mammals.

Ouchi: Yes.

Garner: So had posited that theory based on--

Ouchi: Yes. And his son essentially was the one who proved that the layer, and where it was in the KT boundary, and that whole thing like that. So somewhere very near there, I was deciding that I had better grow up and that I had better figure out what kind of job I really wanted to have. And so I had been taking some programming classes, FORTRAN.

Garner: Of course.

Ouchi: It turns out--

Garner: The 7090-- Fortran was written for the 7090.

Ouchi: Yes, so it was a big 7094. You put in your deck. You run. Stuff would come out. And I said, this stuff is too easy. It was really a lot of fun, and it was really easy. So I learned ALGOL. That was easy. So I went over in the engineering school, and they had started computer science. And so they said, well, if you come over, maybe you can graduate pretty quick, because by that time I had amassed something like, I think, 190 semester units. In fact, when I ended up graduating, I think I had like 220 or something.

Garner: What year was this then that you decided to move over to computer science area?

Ouchi: It was in about '65. In about eight or nine months, I kind of decided what I wanted to do, got married, and was really fortunate that IBM was exploding and were hiring a lot of people. So I was able to get a job down in the San Jose plant.

Garner: Now, before you, in the early 1950s, the San Jose lab was basically initiated or started with students from Berkeley and Stanford. Did you know any of the people who had--

Ouchi: No, I did not.

Garner: Gone from Berkeley to the San Jose lab.

Ouchi: No, I did not. In fact, the history of why they put that plant there was because of Stanford and Berkeley, because you notice that everybody that they hired was all back east guys.

Garner: Well, they couldn't hire enough back east, so they wanted to setup a lab

Ouchi: Want to set a lab out here. And they also wanted some crazy people. And so of course, that's spawned off the disk drive.

Garner: So you didn't know Al Shugart?

Ouchi: I didn't know Shugart.

Garner: Because he had been a student at Berkeley.

Ouchi: I didn't know any of those guys, because when I came into IBM, I was offered a programming position, but then there was this other intriguing position, which was in manufacturing engineering to help bring up the manufacturing line for the 1800 and 1130, which are very different machines from the 360. That's why most of my early career I was clueless about exactly how 360s operated, but the 1130/1800 was an incredible machine to learn how computers really work.

It had 16-bit word, a double, 32-bit, word. It had three index registers. It had an ALU. It had all the functions that you would think a computer would have. Plus it had interrupt levels, and it had what they call cycle steal, which is directed memory access. So once you begin to learn that machine, you've learned all the fundamentals of how real machines work as opposed-- 360 had all kinds crazy, strange instruction sets. But this one was one where you could really understand relatively quickly how a real computer worked.

Garner: Yeah, the 1130 could actually add, unlike the 1620.

Ouchi: Yeah, you could do an addition. Multiply took forever. What was really interesting was the fact that the index registers could be used to reference data, as well as, of course, the long words.

Garner: So you came on board then to help with the manufacturing. So came on as--

Ouchi: I came in to manufacturing, because what was intriguing to me was in manufacturing engineering, you worked with one another engineer or sometimes by yourself on an entire project. So you got to do everything. You get to do engineering. You get to do programming. You get to build real hardware that hooked up to machines. And so the job that I inherited was the 2310 disk drive that attached to the 1800 and 1130. Somebody--

Garner: 2310 or 2311?

Ouchi: The single pack.

Garner: Single pack.

Ouchi: It's called the RAM kit.

Garner: I see.

Ouchi: Somebody had written some of the very basic programs to bring it up, because you have to make sure-- if you realize the kind of logic in those days were circuit cards with the little SLT the solid logic technology.

Garner: SLT.

Ouchi: SLT chips on it, but the backplanes were all wire wrapped. And if you've ever looked at the backplane of a wire wrap, the first layer was done by Gardner Denver machines, and it was pretty accurate, but then all the engineering changes go on top of that. So wires are pulled off, new wires are put on. And many times there would be errors in that, so in the manufacturing process, you had to be sure that whole machine worked first. So what I inherited was the job of making sure that the disk drives all worked, and that they would all work in conjunction with all the other hardware then.

Garner: So again, very hands on experience with storage devices in this case.

Ouchi: And so we had the machine, and all we had was a very primitive assembler. It essentially did the accounting for you. It would translate mnemonics into the actual bit patterns.

Garner: The assembler?

Ouchi: And would keep the address, but there was no macros. There was nothing.

Garner: There was certainly no operating system for the 1130.

Ouchi: There was no operating system. The only thing we had was a one or two card loader, which it did was it loaded the card image into storage, and then the loader would load you wherever you were assigned to put it. And so we were all given a certain amount of space to do our job because they wanted to be able to run other programs when they finally ran the large system to be sure that the whole thing hung together.

And so using the index registers and actually overlays off the drives, I was able to run all three drives with a very small memory. In fact, at that time, I didn't realize what I had created was reentrant code and dynamic overlays, which I'd never learned when I was in school. I kind of created that on the fly.

Garner: But in your role in manufacturing, was that just something you did because it was fun, or did they ask you-- were you trying to contribute to utility software for the 1130 that would be shipping with it?

Ouchi: No, this was to bring up the machines.

Garner: Just for testing and bringing it up.

Ouchi: To be sure that they were manufactured properly. Some of the lessons that I learned while watching that is bill of materials and what goes into the machine typically have lots of errors in it. I remember teaching my development team that the machine that we have on in the lab that we've debugged is not the machine. The machine is that set of documents that you could now reproduce that machine, because there was more than one time where, out on the manufacturing floor, there would be parts missing or there would be extra stuff. Things would ship missing pieces.

And so I remember one superintendent telling those guys, you take the bill of material, lay all the parts on the floor, and you build a machine with what you see there, because he was right. The bills were all wrong. And of course, the bill of material being wrong would not order the right parts. You're always running short. You have excess. All of those kinds of issues.

Garner: So clearly the manufacturing process in San Jose left a little bit to be desired, it sounds like.

Ouchi: You will find that that was the way that the world ran.

Garner: No just IBM San Jose.

Ouchi: Not just IBM San Jose.

Garner: Because of the time pressure to get systems out the door, or lack of software that tracked all the components properly?

Ouchi: In those days, there was all of those things plus the fact that manufacturing, the materials organization, and the development organization usually did not communicate very well. I'll relate that a little bit more when I talk about very much later in my career.

Garner: So on the 1130, did you have a chance to meet any of the designers, or is that too far--

Ouchi: It was pretty far down the road. Most of those people were-- remember, I was in manufacturing. We were off in the boonies.

Garner: They never came to visit.

Ouchi: No, they never came to see us. They were all off in-- it was one of those things where you grew up here hoping that someday you can get into development or the--

Garner: There are many great stories of that in IBM. People start there and then move up. Were the any funny incidences-- funny in quotes-- where the machine smoked, or in test, or the disk drive crashed, or--

Ouchi: Never had-- well, the only lesson that I had was I had written my code so that it would run all three or four drives. I can't remember what it was. And it would really run well. Overlap. Did everything. And then I get a call on a Saturday saying, hey, your code doesn't run. What's going on here? It turns out that the thing wouldn't initialize properly with one drive, so the simple case wouldn't work. The more complex ones-- all the complex ones did. So it was kind of like the lesson was you've got to even look at every one of the cases, even the ones you thought were simple.

Garner: Did it find a lot of bad drives? Was there a reject rate through your testing?

Ouchi: No, the drives themselves were pretty good.

Garner: So you would just find assembly problems.

Ouchi: Yeah, it was assembly problems with the backplanes, mainly in the backplanes. Every once in a while, we would find a bad card.

Garner: But you would think that the Gardner Denver machine would not mis-wire a backplane. So maybe ECO, a change order would be put in wrong or something.

Ouchi: It's all of the changes.

Garner: Change orders were wrong.

Ouchi: Changes were wrong. They used to start putting on red wires, purple wires, blue wires. The Gardner demo would lay down a nice layer of yellow wires, but then things would-- and then there was one idiot engineer.

Garner: He probably wouldn't appreciate that, but it's probably true, though.

Ouchi: It was true. What he did was he was trying to trace out where all the wires were going using a ballpoint pen. It had electrical conduct. And what he was doing was blowing out transistors as he was trying to debug the machine.

Garner: Oh, no. Get out of here.

Ouchi: Get out here. It was actually kind of get out of here. So at that time, IBM had this incredible work study program and what they called the resident study program. So the work study program was you'd end up going to University of Santa Clara, and if you took one class, you went like one day a week for two hours. I was crazy, and I was going three days a week for two hours in the morning. We were working 60 hour weeks.

Garner: Already?

Ouchi: Already. And so I was newly married, and we're just having lots of fun. My wife got this job at the local elementary school, so anyway--

Garner: So you were working start up hours before there were startups and at a big company. You must have had a lot of energy.

Ouchi: And it was so much fun. It was one of those-- you just had so much fun. You didn't want to go home because it was just a lot of fun. And so I was taking a lot of classes at Santa Clara. So I got my master's degree in about a year and a half.

Garner: In what?

Ouchi: In EE comp sci. And then IBM had what they called the resident study program where you could now go to school for two years, and IBM would pay your salary, and all this other stuff like that. And so I thought, well, that would be kind of neat, except I was having too much fun at school, but I was able to enroll at Stanford on the PhD program. At about that time, Amdahl was building his supercomputer over on Sand Hill Road.

Garner: For IBM?

Ouchi: For IBM.

Garner: They were building a prototype of it.

Ouchi: Yes.

Garner: This was the ACS project?

Ouchi: Yeah, the ACS project. And that was an incredible--

Garner: ACS stood for America--

Ouchi: No, it was advanced computing system.

Garner: Advanced computing system. I'm sorry. That was the other project.

Ouchi: What they were trying to do was wire with coax-- teeny, teeny hair width coax-- the backplanes, because embedded LANs weren't well defined at that point. And so what they were trying to do was they would put cards in, put this backplane in, and then use ultrasonic bonding to bond these tiny little coax cables.

Garner: They wanted to have the least amount of inductance--

Ouchi: Exactly. And the trouble is, is that you wire this, and this falls off. You look at the back of it, and it looks like-- do you know what a Brillo pad is? It looked like a Brillo pad.

Garner: A rat's nest.

Ouchi: Because it had all these little tiny wires in the back. So they had put them out of business. They said shut it down.

Garner: IBM management put it--

Ouchi: IBM management said, first of all, Amdahl said he can't build a supercomputer that was 360 or 370 compatible.

Garner: Compatible, right.

Ouchi: And this technology looked like it just was not going to work. So they had shut it down.

Garner: And people back at IBM said we can design a 360 just as fast, and I think that was the other thing.

Ouchi: Yes, and so at about that time, some of the people there were being spun off to some offices in Sand Hill Road.

Garner: So let me understand something. Did you actually join the project, or you just went to visit it?

Ouchi: I was looking for a position there so I could go to Stanford and not have to commute back and forth. I'd go to work--

Garner: Because they were in Menlo Park or something.

Ouchi: On Sand Hill Road.

Garner: Sand Hill Road, right. Where the VCs are now.

Ouchi: Yes, in fact, it's a big white building where the VCs live. But I was able to join this other group that was just down the road for them, because they were just all spilling off, and people didn't want to move, and whatever. And they were going to build an integrated channel control unit so that, instead of having a separate channel box and a control unit box, you'd have this channel control unit.

Garner: This is the DMA function in it.

Ouchi: Yes, the DMA. In fact--

Garner: I always thought it was one box. You're saying that before on the 7090s there were split functions?

Ouchi: The channel was separate from the processor. The control unit was separate from the channel.

Garner: Just because physical packaging.

Ouchi: Physical packaging. Plus typically it was 200-- I think they finally got it up to 600 feet that you could be from the channel to the control unit.

Garner: Wow.

Ouchi: And then the disk drives were behind the control unit. I'll relate that in just a couple of minutes. Anyway, so this team was there, and they were working on this integrated control unit.

Garner: For 360s?

Ouchi: 360s or 370s.

Garner: Some high end 360s.

Ouchi: Some high end machines. But this was just a holding thing because I think they were looking for things for them to do back at the main plant site. But for me, it was really fortunate, because I could come in. I'd have an office. I could go to school. So I could still be on work study, get paid full pay, and do this other thing.

Garner: So now you went to Stanford.

Ouchi: So I'm going to Stanford.

Garner: To get a degree in?

Ouchi: In EE comp sci.

Garner: Again.

Ouchi: Working on my PhD. And of course, as you well know, the most critical thing in the PhD is to pick a project that's interesting enough for people to say, yeah, that's worth a PhD, but one that could actually be solved.

Garner: And you were used to being a student too long already, so you probably motivated to choose something that was solvable in a finite amount of time.

Ouchi: But I was also very interested in storage hierarchies. I was also interested in codes. In fact, let me relate one more thing--

Garner: Error correction codes.

Ouchi: Error correction codes. When I was in manufacturing, we were also building this thing called a 2305 drum. It had an error correction code on it, and the guy who implemented this code was-- the story goes-- was telling management that, unless you give me this increase, I'm not going to let you know the design of that.

Garner: A salary increase. He was demanding a salary increase.

Ouchi: Yeah, he was demanding a promotion and some other stuff, or else he would destroy all the documentation on the error correction code.

Garner: That's an interesting position to be in.

Ouchi: Anyway, so they said to heck with you. We'll figure this out.

Garner: Not rocket science, per se. Claude Shannon and a lot of research had been done on error correction codes. So what was the motivation for an error correction code on the drum? Do you know why they thought they needed one?

Ouchi: Well, what was happening is that you could not build a perfect drum. You could not build a perfect disk. By then, they had discovered that you could survive tiny defects if you could figure out how to get over them. And so there were a couple of things that were floating around. One of them was to just write right over it and put at that-- later on, mean the motivation for me was I said, wow, you could put these little bytes at the end of the record, and errors that were somewhere up here-- you could actually correct them. And that sure seems like magic. And so for some reason, that triggered in my mind this is really interesting stuff.

Garner: So these were Reed-Solomon codes, do you know?

Ouchi: This one was actually a hacked code. It was not a very good code.

Garner: Just someone made it up.

Ouchi: Just made it up. Typically what we would put is what's called a Fire code, which was a burst correcting code, single burst correcting code. I learned this later, because I thought this is some kind of magic. So I decided that I was going to take some classes on error correction codes. And so I had a coding background at that point.

Garner: Do you remember who the instructor was at Stanford?

Ouchi: Oh, it was some guy at--

Garner: Random guy.

Ouchi: It didn't matter.

Garner: So this was the late '60s.

Ouchi: This was the late '60s, middle '60s.

Garner: Middle '60s.

Ouchi: And so at that time, I had finally discovered how Fire codes worked, how BCH codes work.

Garner: BCH, block coded.

Ouchi: Yeah, they were block codes. And the other one, which I really-- was this thing called-- it was a block code, a byte code that this guy named Arvin Patel developed for the eight track tape. So what he was able to do was to do two double track corrections.

Garner: Oh, this was IBM's eight track.

Ouchi: Eight.

Garner: They're called nine track.

Ouchi: Nine track. Eight data, one-- in those days, it was eight data, one parity.

Garner: Then a longitudinal.

Ouchi: A longitudinal.

Garner: And that's what he was working on.

Ouchi: Right, and he said, well, if you intersperse the error correction bytes, you crank up the density so you could now afford to put in these bytes-- that I could now correct two tracks on the fly, and even if the blocks were bad because the blocks were small enough, you could now write over Swiss cheese, and you could have tape that you could now-- so I thought that's really interesting. It's called a B adjacent code. It turns out later on that I would use that for a double drive correction mechanism because the implementation of that turned out to be relatively easy when you think about read only memories, because you can map in and out really quickly.

Anyway, so at that point is when I became interested in error correction, and error recovery, and those kinds of things. And so it was in the back of my head, and I was-- as I went through Stanford, looked at-- at that time, I was really interested in cache. In fact, I'll talk a little bit about why it was so difficult to cache IBM DASD as opposed to-- if IBM had a fixed block architecture like the 2310 did, we would have had RAID a lot earlier, but because of the fact that IBM uses what is called count key data architecture on the drive--

Garner: Maybe you could describe it, because not everyone has heard of it.

Ouchi: Yes, because what this allowed--

Garner: This was on the 360 drive.

Ouchi: On the 360 drive. What it did was you would record a record with a count field which was fixed length.

Garner: A count field?

Ouchi: A count.

Garner: C-O-U-N-T.

Ouchi: C-O-U-N-T. Because what it had in there were two values. One of them was the value of an optional key field, which was recorded separately, and the length of the data field.

Garner: So a length of the key field and a length of the data field?

Ouchi: Yes.

Garner: So there are two lengths. One would be the length of the key field, and there would be the length of the data field.

Ouchi: Yes.

Garner: So you could adjust those two if you wanted to.

Ouchi: Yes. You would adjust those two.

Garner: You could adjust them.

Ouchi: And in fact--

Garner: Unfortunately, you would have to read those two numbers first to know how big they were.

Ouchi: Exactly. But say that you had to kind of find the beginning of the count field, and then from there, you would find you could now determine how you're going to read or write the key field and the data field.

Garner: And what was typically in the key field?

Ouchi: Typically it was the search key that was used in index sequential access methods or binary trees. So B trees. Because what had happened was, because core memory in 360 was very small and they were trying to process lots of information, it didn't want to read in all this information. They didn't want to keep tables, because table space was really expensive.

And so you could actually write a channel program, because what channels allowed you to do was not only read, and write, and compare key fields. It also had a conditional branch. And so you could have a search key, and if it's equal, you could transfer in channel to another place, or if it wasn't equal, you could ask it to jump over something and now execute.

Garner: So you had this processing capability in the channels, almost like associative memory, in a way.

Ouchi: Yes, and because of that, what they essentially did was put the access method in the channel tied to the geometry of what was going on the track. And so the track geometry was tied very tightly to what was going on in the access method.

Garner: Very inflexible.

Ouchi: Exactly. It was just locked in. So whenever we wanted to put in a new disk drive because of higher capacity that had a longer track length, then the access method, channel programs, and everything would have to be changed. It was a huge endeavor. Now, what was really hard for us, and why it was difficult to cache was because of this associative searching, and then finally, at the very end, you would discover if you're supposed to read or write that record. You could never tell ahead of time what you were going to do.

Garner: So you would have to wait for an entire disk revolution in essence.

Ouchi: Right, or essentially you'd end up with search key, key equals, next thing you get is a right. It was like maximum surprise, and because of that, unless you had a large cache, you could not cache an IBM

CKD data format. And so this became a huge impediment for IBM trying to cache their storage.

Garner: To back up, though. Back to your PhD-- so you were learning all about erasure codes and coding. So did you get a PhD in that topic?

Ouchi: My PhD was in storage hierarchies.

Garner: Oh, OK, that's right. Storage hierarchy, because you got interested in caches. Now, when you were working on your PhD, did you learn this about the IBM count key data as part of the process?

Ouchi: No, I knew about the count key data.

Garner: Before you--

Ouchi: Way back before, because in transition, I had--

Garner: So part of your thesis was talking about these problems probably in the count key. Tell me what your thesis was about.

Ouchi: My thesis was about how to characterize serial and random access storage in a coherent mathematical form so that you could have storage that was made of serial access pieces and random access pieces, and then to be able to now characterize how programs would execute given those kinds of memories.

Garner: And did you assume caches?

Ouchi: Well, these were caches, levels, and--

Garner: So you had a disk drive or disk drives.

Ouchi: It could be disk drives. In fact, at that time, I was really interested in bubble memories.

Garner: So any kind of memory.

Ouchi: Any kind of memory, any kind of random access memory. You could put random access memories in between in order to-- and so this was a set of mathematical expressions of this in that--

Garner: It was an ontology for describing how you could arrange-- many possible ways you could arrange storage devices with cache devices.

Ouchi: Correct. And then you'd just set parameters in this thing, and then you could do some mathematical processing on this and to be able to now predict how--

Garner: Latencies and throughputs.

Ouchi: Yeah, and you'd be able to now predict how programs would--

Garner: Might perform.

Ouchi: Might perform. And then--

Garner: But you wouldn't know about the cache hit rates and stuff like that. You'd have to assume from that.

Ouchi: No, no, no. You'd have to assume those kinds of things, but this was like average, longest, those kinds of things like that.

Garner: Now, in this PhD, you didn't consider RAID ideas at that time, during that period of time.

Ouchi: No, at that time it was just all-- because I was focused--

Garner: There was no reason to.

Ouchi: No. It was how to characterize hierarchies. And then I was able to--

Garner: Did you give them names?

Ouchi: No, they were just--

Garner: They were just--

Ouchi: They were just values in a vector, because you could make-- the thing is, if you could mathematically describe it, then you could begin to process and to be able to now predict how these would behave.

Garner: Did you write a simulator to simulate this?

Ouchi: Well, then I wrote a simulator.

Garner: As part of the PhD?

Ouchi: As part of the PhD program in order to do validate that some of this was actually true, that the math was actually coming.

Garner: Because by itself you can't tell.

Ouchi: Right. And so I was able to get a bunch of data streams, and you run it against error, and you say, oh, look what this one predicted. You do all that kinds of stuff.

Garner: So then you adjusted the math and then-- so it made it look like you'd really done a good job so you could graduate with your PhD.

Ouchi: I keep saying that my PhD thesis, which ends up to be about that thick, is so full of snopake that you'd go like this, and the thing would fall apart. Because I keep claiming that, if I had a good word processor-- if I did it today, my thesis would have been done in--

Garner: Was it done on a manual typewriter?

Ouchi: Yes, it was done on a manual typewriter. By the time I was starting to write it, I was back at work at IBM, and they gave me this secretary, and I was driving her nuts, because what you had to--

Garner: Was it an IBM Selectric typewriter?

Ouchi: Yes, it was. It was an IBM Selectric.

Garner: That would drive anybody nuts. Did you have a ball with different fonts on it and stuff?

Ouchi: Yes.

Garner: So she had to change the ball out all the time.

Ouchi: Well, I didn't ever have to do that. The problem that she was having is she wasn't that great a secretary. Anyway, because when your thesis advisor says, well, maybe you ought to put a little more stuff in here about that, or whatever, or change the wording here, it was how do I fit in this thing so it doesn't propagate more than--

Garner: One page.

Ouchi: Yeah, I don't want to do more than that. So I ended up bringing it home, and my wife and I actually did it at home. Pasted over. So anyway, I got it through. So then I came back, and the group that I had been working with--

Garner: Well, you never left IBM really.

Ouchi: No.

Garner: When you say came back, you mean back to San Jose-- you mean.

Ouchi: Came back to San Jose. I was able to leave for a year where IBM paid me to go to Stanford full time, but most of the time was sitting at home working on my thesis. So then I finished that, came back, and I was just cleaning up my thesis, because I wanted to get back to work because it was so much fun. And so I joined this group, which was-- they had all these ACS guys that had all these crazy ideas and wanted to do stuff. And so at that time, IBM had just developed what was going to become the 3340 disk drive. When you look at disk drives, in the early disk packs, you had to worry about head skew because you had multiple heads, and they would go like this when you seeked-- and when you go from one drive to another, they'd be a little bit off.

Garner: Because the packs were removable.

Ouchi: Yeah, the packs were removable, and the heads stayed--

Garner: Like the machine thing where you could put the whole-- how many platters were there typically?

Ouchi: It was like--

Garner: 10 platters?

Ouchi: 10 platters. 30 megabytes.

Garner: 30 megabytes.

Ouchi: And it had 20 heads, and you had to keep them all in alignment. And so they said, well, we could reduce a lot of these mechanical tolerance problems if we kept the heads in the carriage with the pack. So the 3340 had the head in the carriage with the pack, the actuator and everything else--

Garner: So the motor, actuator was part of the thing you lifted out, which made it really heavy, I assume.

Ouchi: Yeah, it was pretty heavy. So they had that, and they had another version of that that they were going to bolt down, because you could really increase the capacity then, because if it was not removable, it took a lot of mechanical tolerances out of the thing.

Garner: Does that have a number or a name?

Ouchi: Yeah, it turned into the 3350, but we'll get into that relationship in just a minute.

Garner: So anyway, you came back. Now you're an engineer at IBM.

Ouchi: I was an engineer. I was in development.

Garner: You're out of manufacturing.

Ouchi: I was no longer in manufacturing.

Garner: You've had your hazing through manufacturing.

Ouchi: Right, and I learned all about stuff that I would later learn is really important, namely keeping the bill of materials straight and how to keep engineering changes straights.

Garner: But now you're going to do engineering. So now you could--

Ouchi: I could now mess up. So I joined this is group, and it was kind of off in the corner, and we were looking at this 3340 disk drive. And we said, well, what can you do with that thing?

Garner: Now, this is early '70s?

Ouchi: Yes. About '72, '73.

Garner: I graduated from high school that year.

Ouchi: Jeez. So anyway, we also went and visited this IBM fellow named Walt Buslik, and he said, you know what we could do? Is record helicoidally on to tape. And you could lay down tracks and get really high capacity, much higher than you could with linear track tape. Of course, this is what VCRs were becoming and all this other stuff like that.

So we had this little thing called the Buslik tape drive that had a cartridge on it that held thirty megabytes. And so what we decided we would do is marry one of these helicoidally scanned tape drives with a disk drive and allow you to mount and demount the pack. Only the pack would be fixed, And you'd take this little tape cartridge, stick it in its mouth, and it would spool the data on to the disk drive. And about the time it would normally take to spin up a drive, we'd have the data on there. And you'd dismount the stuff, and you'd have the data on there. And we figured we could probably build this thing for like $1,000, something really inexpensive, because we were trying to get into the low end market.

Garner: Well that's an interesting statement. Low end, what? The 360 market?

Ouchi: 360 market, but also into the-- at that time--

Garner: 370 actually is the market now. The early '70s would be the 370.

Ouchi: Right, but there was another market that was below that where--

Garner: The business guys--

Ouchi: The business guys.

Garner: The system 36 and 37.

Ouchi: Yes, 36 and 37, which were fixed block machines.

Garner: So you were targeting working with those guys.

Ouchi: Yeah, even with those guys. The reason that-- in fact, they were fixed block storage, a because of that, that's where you saw RAID first, because they didn't have to contend with the CKD data format.

Garner: So that kind of opened the door.

Ouchi: Right, and so by the time Patterson published-- by then, RAID 5 was actually implemented in some of the IBM products.

Garner: In those system 36, system--

Ouchi: System 36.

Garner: Do you remember which one?

Ouchi: I can't remember which one.

Garner: Yeah, OK, we can check that.

Ouchi: But the RAID 5 patent is an IBM patent.

Garner: Yours, you mean?

Ouchi: No.

Garner: There's already another patent on RAID 5.

Ouchi: Another patent that followed mine.

Garner: That followed yours.

Ouchi: Right.

Garner: Yours was in '78.

Ouchi: Mine was '77, '78. We filed in '77, got it in '78.

Garner: But anyways, the point is, this project you're working on, these cartridges-- now, had you already thought about doing RAID at that point?

Ouchi: No.

Garner: Still not yet.

Ouchi: Still not yet. RAID hadn't come into the picture yet. And so what I did was-- so we figured out how to map these stripes onto a disk drive and figured out how to do the CKD because you had no large--semiconductor memory was extremely expensive.

Garner: Yeah, in the mid '70s, cores--

Ouchi: It was all by core memory.

Garner: Core was still-- it is getting a lot cheaper, but it was still the predominant form of main memory.

Ouchi: Still expensive. So you could probably have a track buffer was about the best we could do. But we could build this machine, and come in at a really low price. At the same time in Boulder, the Boulder lab--the IBM boulder lab, which was responsible for tape drives, was trying to develop this very wide tape with 72 heads on it, and what it would do is have 72 heads and would essentially pull a piece of tape out and run it back and forth like this. It was insane.

And they were in trouble, and they didn't know what to do. So somebody-- our management thought it would be a really neat idea if, in fact, we could sell them on this cartridge idea-- that you could now take a cartridge, a helicoidally scanned cartridge, which had very good densities, and stage it to disk, and you could now have essentially what appeared to be a disk drive with mounted packs. Only it was actually coming out of a honeycombed library.

So we said that's a really neat idea, because they'll have a high end. We'll have a whole family. You could have small libraries. And so they went off, and they actually sold the idea to them. They thought this was really neat. I guess they did.

Garner: So they started working on it in Boulder.

Ouchi: This became the 3850 mass storage system.

Garner: So that was its origins. This is the honeycomb system with the giant robot arm and big giant--

Ouchi: Giant robot arm, and it would pick cartridges. A cartridge had the effective capacity of one half of a 3330 disk pack. It was on demand staging, but what they did was-- of course they had this 2.7 inch width tape helicoidally scan, but instead of moving it smoothly like you would a videotape, those guys still had vacuum columns in it, and they would step that thing in synchronous with the rotating head. So this made a tape drive that should've been $100 to manufacture. It was now like $1,200. It was more expensive than our entire subsystem was going to be.

Garner: But again, the system was large. It had a storage-- a robot arm that went, and reached, and grabbed the cartridge, and brought it back to the tape reading station. So that made it expensive, as well, I assume.

Ouchi: That was expensive, too. That was now going to replace a lot of human beings, and a lot of disk packs, or tapes on the wall, and everything.

Garner: So this was a successful product for IBM, I would think.

Ouchi: For IBM, they sold several of those machines. Quite a few.

Garner: Yeah, quite a few.

Ouchi: But what it did was essentially killed our little project. So we were kind of--

Garner: You gave them the idea, and then--

Ouchi: We gave them the idea. It became the 3850, and so they're off now building these tremendously expensive tape drives. And then now staging to the 3330-- this is the other area we gave them. We'd now stage to 3330 so that the world would see what appeared to be a whole set of 3330 disk drives, but in reality, most of the volume was carried back in tape.

Garner: Was that the interface to the 3850?

Ouchi: Yes, that was the 3850.

Garner: So they took that idea, as well.

Ouchi: They took that idea, and they built that. They were going to build that machine.

Garner: Well, they did build it.

Ouchi: Yes.

Garner: And it shipped.

Ouchi: At that point, they were going to build it.

Garner: But then it happened.

Ouchi: So what happened to me it was we were now-- we said our little thing-- they essentially crushed it, because--

Garner: They were the bigger player. That was the home for tape development at the time or something.

Ouchi: Right, the entire lab was going to build that thing.

Garner: They didn't pat you on the back and say thanks for the idea?

Ouchi: No, they didn't.

Garner: No good deed goes unpunished.

Ouchi: Exactly. So then there were all these little cats and dogs engineers running around. One guy by the name of Chuck Frehlich had decided that if you take a disk drive and you read and write four heads in parallel, you can get high bandwidth.

Garner: Well, people knew that already.

Ouchi: Yeah, but he wanted to build a high bandwidth count key data machine, but in doing so, he had some spare heads. Now this is where the first--

Garner: I don't understand that. Why world--

Ouchi: Well, because the thing had a fixed geometry of 20 surfaces. One was a servo, and he wanted to group them in fours, so he ended up with 19.

Garner: Why did he want to group them in units of four?

Ouchi: Because that was going to give him the data rate that he wanted.

Garner: Interleaving.

Ouchi: Interleaving, and so essentially read four heads and parallel and give them--

Garner: Each of the sets of four could be in different positions.

Ouchi: No, they were--

Garner: All the same?

Ouchi: In same thing. What you did was made this appear to be a single drive with a single head, but four times the data rate.

Garner: He could've done eight times.

Ouchi: Yeah, he could have done eight, but read channels cost a lot. Every one had a serializer, de-serializer.

Garner: Oh, I see. Those are--

Ouchi: You had to have all that other stuff, because this is all logic that-- each one of those was not just a corner of a chip, but it was--

Garner: A suitcase of logic.

Ouchi: Yeah, so he had these spare tracks laying around. At that time, there was also a problem of every once in awhile a head would go out, and it would destroy the whole surface.

Garner: Head crash.

Ouchi: Head crash. So one of these crazy ideas-- I said-- and I can actually do this now count key data-- that I could now take this, and begin buffering parts of it, and create an exclusive-or across the bytes that were now flowing through these four heads.

Garner: Bingo. So that was the idea for RAID.

Ouchi: Right, so in the fifth head, I would now create the exclusive or such that, if one of the four heads were destroyed, you could now reconstruct the data using the remaining three plus.

Garner: So what year was about?

Ouchi: This was in early see '73.

Garner: That early? Because this patent is not until '78 or '77.

Ouchi: Yes, I'm going to get to why--

Garner: So early '73 you had this for internal proposal.

Ouchi: Internal proposal to provide this kind of redundancy.

Garner: So what motivated-- you already knew about erasure codes. You thought they were cool. So there was not ECC memory at that point, probably.

Ouchi: Yeah, there was already ECC memory.

Garner: The core memory--

Ouchi: Single bit correction.

Garner: And you know about that.

Ouchi: And I knew about that. This was just simple exclusive or.

Garner: Well, that's what simple RAID drive is. Just simple exclusive or. So you threw that out.

Ouchi: So I put this in a disclosure. A bunch of people signed it, and they said, OK. And of course the project died.

Garner: Sure, why not?

Ouchi: It sat in the patent organization. You fill out your notebooks. You do this.

Garner: And you guys did a lot of that back then.

Ouchi: Yeah, so when you--

Garner: Signed your name in the notebook and all that stuff.

Ouchi: And if you had a bright idea, and you fill out this form, and you throw it into the patent bin.

Garner: Hopper. But that one didn't get issued.

Ouchi: No, it did not. That went nowhere. It was a product that never went anywhere.

Garner: The patent did not get issued either.

Ouchi: No, it wasn't even-- it never got anywhere.

Garner: Neither the product nor the patent.

Ouchi: Neither the product, patent, or anything. The patent guys said, you guys going to do anything with this? We said, no we're not going to do anything, but in that disclosure was the generalization of it didn't have to be heads. It could separate failure independent units.

So now we fast forward for five or six months, because everyone was looking for something to do. The 3350 Madrid drives were non-removable, and it was a concern that the marketplace was not ready for a non-removable disk pack. There was also this problem of most of the processing at that time was batch processing. When you ran and a disk drive failed, you could always bring the tapes back, restart the job, get going again. People who were doing online reservation systems and those kinds of things-- when a disk pack failed, it was a real, real problem. They had to back up. They had to do a lot of things to protect themselves against disk drive failures.

Garner: These would be 3340s, you mean?

Ouchi: Yeah, at that time it was still 3330s.

Garner: And how often did they fail, do you think?

Ouchi: Not very often.

Garner: Once a year?

Ouchi: But when you have an American Airlines something like that.

Garner: All your key reservation technology.

Ouchi: You have reservation stuff, and you couldn't afford to have anything fail on you.

Garner: So they would back up once a day or something like that.

Ouchi: Yeah, but even the backup-- it was still a pain, and you couldn't afford to lose people's reservations.

Garner: Did they backup on tape at that time?

Ouchi: Yeah, they backed up on tape. They also dual copied certain things. They did a lot of things to protect themselves, but it was because of physical failures of the disk drive. So about that same time, the 3850 guys were starting to now gain traction with the idea of having 3330s in the front store and caching essentially back to the 3850 cartridges.

So Dick Arnold, whose initials you find on one of those, along with-- not on the patent-- and I and several other people thought, well, instead of putting in onto cartridges, what if we took an array of 3350s-- at that time they were called Madrid drives-- and we would essentially write RAID 3 across those drives. Fetch and store off the back store.

And this program was called Zark Somebody had written this paper about caching calling Zark. And it was called Zarking, where you would now stage, destage, use the LRU, do all those kinds of things. The other thing that we did, which was really crazy, was we were going to do gap suppression and also compression on the data, because a lot of the data was redundant information.

Garner: What do you mean by gap suppression?

Ouchi: When you write a count in a key in a data field, you had to write them-- first of all, because when you record magnetically you had to have what is called a sync field to get everything up, and then you finally write the field, and you finally write the end. And you had to have these buffers on the end, because there's always mechanical vibrations and everything. The other thing that was happening was, remember, that the channel had to do this essentially command processing during that time, and so we had to give the channel 30 microseconds of time. And so you had to give that much time in the gap for channel turnaround.

And so typically, these gaps were pretty big. That's why people didn't like to put the small records-- because you had these gaps. And so we said, why don't we smash those gaps? Plus a lot of data was redundant, and so we had this adaptive compressor that would essentially take that and compact so that we could take tracks, smash them together, smash together a cylinder's worth of stuff, and then write that as stripes across a set of Madrid drives.

Garner: Got you. So this was a paper design, basically.

Ouchi: At that point, it was a paper design. It was getting people interested in it.

Garner: We're talking mid '76 or something like that, or '75.

Ouchi: No, this was like '74.

Garner: '74.

Ouchi: In fact, you'll find one presentation by Dick Arnold, who essentially got the lab to fund this project.

Garner: Now you mentioned-- did I here you say RAID 3 earlier?

Ouchi: Yes.

Garner: In RAID 3, not RAID 5? So describe that.

Ouchi: Well RAID 3, what it meant was we would write the entire record across the multiple drives as opposed to a RAID 5, where you would go and update one of the records in the error correction piece. So if we were a block oriented architecture, we would have thought of RAID 5, because we would do individual block updates.

Garner: So you striped it.

Ouchi: Rotating. So we would write--

Garner: You striped it, basically, across all of those.

Ouchi: We would write stripes. And because we were compressing, we would have variable length strips.

Garner: Variable length stripes.

Ouchi: And we also had the fact that the storage back store was huge compared to, let's say, a cache memory or something like that. So we're trying to LRU as to which ones we should take out and where what should go next.

Garner: You'd need a garbage collector.

Ouchi: Yeah, so we had garbage collection. We had this time stamping mechanism so we knew which ones were now fresh, which ones were old, which ones were updated. So we had all these algorithms also.

Garner: Very sophisticated design for the time on paper.

Ouchi: It was. It was a really, really neat.

Garner: But the motivation was cost and performance-- what was the motivation? Was it cost or performance?

Ouchi: It was first to provide redundancy against failed drives, and then also to provide a whole lot more capacity.

Garner: But there was no-- there was checksum, let's say, but was there checksum on these stripes? What was the protection against a failed drive? Because you've striped across all of them, so one them-- the drive is going to go down.

Ouchi: Oh, well the striping would have the checksum in it.

Garner: There was a checksum?

Ouchi: Yeah.

Garner: That's what I just asked. So the stripe had a checksum.

Ouchi: Every stripe had a checksum in it.

Garner: Every chunk on a disk had a checksum or every stripe had a checksum?

Ouchi: Every stripe had a checksum.

Garner: So then if you lost a piece, you could use a checksum. How did you calculate the checksum?

Ouchi: It was just because it worked.

Garner: Again, it was because it worked. I've noticed you call them checksums. Even though you're striping, wouldn't that technically be RAID-- oh, because it's not fixed width, that's why you're not calling RAID 5.

Ouchi: No, well RAID 5--

Garner: It's very rigid.

Ouchi: Yeah, RAID 5 is imagine you have a set of blocks.

Garner: And a fixed number.

Ouchi: A fixed number of blocks. And you'd have the exclusive or, and what RAID 5 did was rotate that exclusive or so that you didn't end up pounding one poor drive when you're doing block updates.

Garner: But in your RAID 3, since you had variable length stripes, you're not necessarily lining up anything either.

Ouchi: Correct. So it was kind of a funny RAID 3.

Garner: Yeah, it's not really the original RAID 3, I think.

Ouchi: No.

Garner: It's more of a more modern concept, I think.

Ouchi: Right.

Garner: Now, one thing you didn't do is you didn't take a stripe and randomly put all the pieces. You made them sequential, probably, across--

Ouchi: Across these drives, right.

Garner: It's called declustering today, but you didn't take that step.

Ouchi: Right.

Garner: And you had to have a data structure somewhere that told you where all the stripes were.

Ouchi: Yes, and so we had all kinds of-- you'd put information inside of the records just in case because you can reconstruct.

Garner: So you could always reconstruct it.

Ouchi: Right. And so this was the machine that we proposed and had begun to build. It was built out of a standard 3830 control unit.

Garner: Oh, so you actually did start to implement it?

Ouchi: Right. Standard 3830 control unit, along with 3330 disk drives in Madrid.

Garner: So all you had to do is change the software, basically.

Ouchi: It was a microcode project.

Garner: Microcode project.

Ouchi: Now, what was helping us in a sense-- helped us, but then hindered us-- was that the 3850 had created a buffered 3330 controller, because it was doing tracks off of the cartridges. And so, we now had a piece of hardware that we could program to do our job. 3330s in front. 3350s in back. So we were off, and we started this project. And then the 3850 program got into trouble, and we were all pulled from--

Garner: Production units were having trouble.

Ouchi: No, the development was in trouble.

Garner: Oh, the development was in trouble.

Ouchi: And so we all had vacations in--

Garner: In Colorado?

Ouchi: Boulder. And so I spent a lot of time in Boulder flying back and forth, going, bailing them out, doing certain things like that.

Garner: Was it problems with their microcode or firmware?

Ouchi: Yeah, it was mainly microcode problems.

Garner: Firmware issues.

Ouchi: Of course, they also had drive problems, but we couldn't help them with that. There were two major microcode projects. One of them was the modifications to the 3830 controller that was controlling both the drives and the tape unit. And they had this other box called the DSM, which was the data staging manager, which was the thing that essentially mapped everything. And they were having real problems with that one.

So what happened was San Jose ended up taking that project-- part of the project-- and brought it back to San Jose. So I was involved in part of that, but then an opening occurred, and I became the manager of the product test organization. Now, in IBM, they've got an organization that, before a product is announced, product is shipped, they make sure that the design, in fact, meets all the standards, as well as it actually works.

Garner: Has some integrity.

Ouchi: Yes, so we're the integrity check. So I ended up managing that organization, and we did the 3350 drive and whatever else was going on at that time. While I was in that position, another group-- they had finished up their work on the 3850, had picked up Zark and it was now called Phoenix, because they want to resurrect it from the ashes. So that machine was--

Garner: You didn't motivate them to do that? They did that on their own.

Ouchi: No, they did that on their. There was actually some other people that were motivated. Dick Arnold was still pushing things, but by then, I was off on this other career being a manager.

Garner: In products?

Ouchi: In product test. And so they actually built a machine, showed it to the world--

Garner: Did they start with your old firmware?

Ouchi: Yeah, it was the same old code. So

Garner: They brought it out. They dusted the dust off of it.

Ouchi: It really came up from the ashes of the code that we designed. It actually did all those things. It did gap compression, but it did not do data compression, but it did the full stripe writes. It did the exclusive or's. It did all that thing. And while it was running, they actually went and pulled the plug on a 3350.

Garner: One of the disk drives.

Ouchi: One of the disk drives. The world kept going, did reconstruction, did all other stuff.

Garner: So the first working RAID system.

Ouchi: Yes.

Garner: So this is 1976?

Ouchi: Yes. And so about that time they said that was great, but we're not going to build this thing, because we're selling more 3350s than we could ever manufacture. So we're not going to do this.

Garner: You're trying to tell me that they didn't want-- they were so successful they didn't want to be more successful? Was there a manufacturing limitation on the number of--

Ouchi: No, it wasn't that. What was happening in their mind--

Garner: What kind of logic is this?

Ouchi: Is the complexity of a virtual storage system.

Garner: They were afraid of it?

Ouchi: They were afraid of it.

Garner: That it might crash and do weird things. Lose customer data, right?

Ouchi: Right.

Garner: It might screw up.

Ouchi: It might actually screw up.

Garner: But wouldn't enough testing show that it might work?

Ouchi: Well, the testing was testing of feasibility to show that it actually worked. We never--

Garner: So who was the doubting Tom? Was it some manager or some VP?

Ouchi: Oh, yeah it was the whole disk drive guys.

Garner: What did they believe? That disk drives shouldn't fail and that you don't need this?

Ouchi: No, but also that disk drives-- that they were selling too many already.

Garner: But what's wrong with selling too many? Why not sell more? So they must not have thought that the failures were a significant enough problem.

Ouchi: Correct.

Garner: It was OK to do backups and make extra copies.

Ouchi: Correct.

Garner: In fact, that makes you have to buy more disks if you make extra copies.

Ouchi: Exactly. Exactly.

Garner: And what you were doing is, if people were making replicas for extra copies, you were diminishing that need because you wouldn't need to have replicas?

Ouchi: Yes.

Garner: RAID 5-- how wide typically--

Ouchi: Well, we were doing like eight drives.

Garner: Eight drives.

Ouchi: Seven data and one exclusive or, but then of course, that moved around.

Garner: Was it still variable length stripes?

Ouchi: It was still variable length stripes.

Garner: So eight to what, then? There must have been a range.

Ouchi: I can't remember what the compression ratio-- it was like two to one just on gaps.

Garner: But if the stripes were all different widths, they can't be seven plus 1, or eight plus one.

Ouchi: No, no, no. They could be like one plus one.

Garner: OK, so whatever they were. And the maximum size could--

Ouchi: Max was seven plus one, and whatever is left over got wrapped around. But everybody was protected.

Garner: That was the max. But yeah, so you would result in better efficiency as well as data protection. And this was the problem IBM had with punch cards. Why should we get rid of them, because we make so many?

Ouchi: Right. More than we ever could ever make.

Garner: Right. So in fact, magnetic tape put pressure on punch card storage. And then disk drives eventually put pressure on tape, and you were part of that initial idea of putting pressure.

Ouchi: Putting pressure on disk drives.

Garner: But the upper management-- you don't remember any names of upper management.

Ouchi: I don't remember the names of the managers in those days.

Garner: So even though it was a proof of concept-- and this not research either. This was in the development.

Ouchi: This was in development.

Garner: Yeah, this was a development project. It just died then on the vine.

Ouchi: Right. There was a great, great. That was it.

Garner: But you did a patent though.

Ouchi: Yeah, well they said, well, shouldn't we patent something? And they said, yeah, maybe we should.

Garner: So that became this patent?

Ouchi: So they said you get to pick one patent. What is the salient thing that's important? And to me, it was this, the fact that we could now have failure recovery.

Garner: So the patent was titled “system for recovering data stored in failed memory units,” where memory unit was a disk.

Ouchi: In this case, it was a disk. It could've been anything.

Garner: Norman Ken Ouchi. Filed May 31, 1977. That was right when I was graduating from Stanford. Patent number four billion, zero, nine, two, seven, and three, two (4,092,732). The first RAID patent that later became a massive industry.

Ouchi: Yes, and the reason that it had difficulty in IBM is because of the count key data format.

Garner: Well, that's another reason that you're saying, because you were targeting this for the System/36, the small business--

Ouchi: No, this was targeted for the high end.

Garner: It was targeted for the high end.

Ouchi: Because the whole front end was 3330s. So it was still coupled in the 3330s, and it still had the 3330 in there, because we had to be sure that we would-- if somebody wrote something on a 3330 track, we had to scurry and get that back into the RAID part of the system, because we couldn't let it sit out there very long for fear that the 3330 would fail. So a couple of things had to happen before the kind of RAID that we have today was used in a company like IBM. One of them was we need to have a large enough storage memory so we could do this count key data stuff and make it invisible to the system above us.

Garner: But the System/36 guys wouldn't have this problem.

Ouchi: No, by then System/36 and those guys-- they had already fixed block-- they were already doing fixed block.

Garner: So like you said, were they the first commercial RAID product from IBM then?

Ouchi: Yes, and this was in the late--

Garner: Later.

Ouchi: Late '70s, early '80s.

Garner: Now, they didn't take the idea from you, though?

Ouchi: Yes, they did.

Garner: So tell us that then.

Ouchi: Well, what happened was this patent, now, became known, at least within IBM.

Garner: So it wasn't an unknown thing.

Ouchi: Correct. Because if you see what later would be called the RAID 5 patents and patents that came out of Rochester, they all reference this. This thing is referenced by a whole bunch of art. And so all the stuff that came out of IBM-- in fact, all the fixed block stuff-- references this, because that was seminal to--

Garner: Did you ever talk to the people over there?

Ouchi: Not really.

Garner: Didn't need to.

Ouchi: They need to. It's one of those kind of, wow--

Garner: It's obvious after--

Ouchi: It's obvious. You say, oh yeah, that's easy. So anyway, my personal career-- I had managed the product test organization, and then was able to go back east and work in the engineering programming and technology staff, working for Bob Evans. And of course, Bob Evans was one of the major drivers the 360 for IBM.

Garner: Bob's history actually even started before that as head of the organizations that included the 1401. His career highlight was certainly driving 360.

Ouchi: Driving 360.

Garner: So what was he like?

Ouchi: Oh, was incredible.

Garner: Do you have any stories about Bob Evans?

Ouchi: He was a--

Garner: Powerhouse.

Ouchi: Immense gentleman. If you notice that most of the people that are really good, they are really good with people, but if they didn't like what you were doing-- and I've only seen it once or twice-- they could get really ah. But that was it, and then you're moving on.

Garner: Some people said he had a timer on his desk that he would-- like a sand timer.

Ouchi: I never ran into the sand time because I--

Garner: You always were casually there.

Ouchi: Yeah, because I'd pop in and talk to him.

Garner: Yeah, what would happen is he would turn it over, and you had three minutes to make your case.

Ouchi: Yeah, you had to make your case.

Garner: If you hadn't made it, you were out of the room.

Ouchi: Right. One of the stories-- and I wasn't there when it happened-- was I guess he was somewhere in the San Jose research lab or something, and they were building something. And they were showing it to him or whatever, and he says, oh. And the thing fell apart. He says, I think you've got a reliability problem.

Garner: So what type of work were you doing in this office then? Or what was the office?

Ouchi: What we were doing is looking at all the IBM plans, and I had storage. And people had programming and processors, and of course, I am all these moving parts.

Garner: It's just a planning group.

Ouchi: And his job was to give to the corporate community his view of how all of this was coming together and how it would fit in the overall strategy of IBM.

Garner: Was Chuck Branscomb part of that at that time?

Ouchi: He was part of the line management.

Garner: Further up?.

Ouchi: Further up, well--

Garner: Further down or further up?

Ouchi: Evans actually reported in to the CMC, the corporate management committee.

Garner: And that's where Chuck Branscomb was?

Ouchi: Branscomb was there. All these other guys that ran the stuff. So Evans was not line. He was staff. So he can go in, and yell, and scream, and whatever, whatever, and they could say go pound salt.

Garner: Did you ever interact with Shel Jacobs in San Jose, because he was in marketing.

Ouchi: Oh, yes. In fact, for a while, Shel and I had a lot of fun. And I'll get to that kind of near the end, because I'm going to have to roll along. So anyway, I had that experience, came back, had the advanced storage management assignment when I came back to San Jose. And we were looking at what is the next generation of primarily the architectural standpoint of what we should do with storage. At that time, the guys in Poughkeepsie were enamored with this thing called a dual insertion ring. So fiber optic rings was where they thought the world was going to go, and I was pushing fiber optic switches.

Garner: What year is this about?

Ouchi: This would be like--

Garner: Early '80s?

Ouchi: '81, '82.

Garner: Fiber optic switches weren't really a part of the scenery in the 80’s.

Ouchi: No, they were fiber optic coming in. You turned it into electronics, and they you'd fiber optic out.

Garner: Did you know of any such beast at that time that existed?

Ouchi: Well, we knew that we had that technology.

Garner: You had the technology. IBM had the technology.

Ouchi: Yeah, because we were going fiber optic. What they wanted to do was have two rings, and when the token went past, then you could throw your stuff on.

Garner: This is token ring based.

Ouchi: Token ring only at 20 megabytes per second or whatever it was. I can't remember exactly what it was. And I told them this was insane.

Garner: But there wasn't a multiport general crossbar optical switch at the time, was there?

Ouchi: No, this was not going to be on an optical switch. It was going to be optical coming in, electronic switch--

Garner: No, I know. I understand.

Ouchi: Optical out.

Garner: But no such item existed in the market, I don't think.

Ouchi: No.

Garner: So my point is, I think you were coming up with first designs.

Ouchi: Yes, we were going to build this.

Garner: You didn't see one, and you said let's just do that. You actually conceptualized this.

Ouchi: Yes.

Garner: In 1982, because it could have been happening in other companies, but certainly that was very early in the existence, because I don't think a switch like that didn't come out from IBM probably until much later.

Ouchi: The ESCON switch.

Garner: ESCON.

Ouchi: I'm trying to remember when exactly it was actually shipped.

Garner: So you're in the planning group now when you were proposing this?

Ouchi: Yes, I had storage architecture. I was the advanced storage manager, whatever that meant. So what I was really trying to do was not only to get the switch, but I thought the switch would be the greatest place in the world to put the first level cache for storage. So I was also trying to convince the San Jose management to fight to get that switch, because it was mainly for storage, and let's put a cache in it.

Garner: So some of the channels would be disks. Some of the optical ports to the switch would be storage controllers, I guess.

Ouchi: Yeah, it was just mainly all storage, because everything else tapes-- maybe tapes-- but everything else was unit record stuff. It was too slow.

Garner: Was your vision that there would be a disk drive plus a controller would then be a port into this switch, or would it be--

Ouchi: Well, at that time, we were starting to get electronics that were small enough that I could now have a string and disk drives without a controller on it, and I could just pump that directly into--

Garner: You mean like fiber channel arbitrated loop.

Ouchi: Yeah, fiber channel into this thing. And that this thing could now manage what was going on down there in the disk drives. So essentially what I was able to do was spawn off this task force that had Poughkeepsie management as well as ours to agree on whatever came out would be really neat. And so I was able to convince them that they should go to a switch. So the ESCON switch-- and the main reason was I knew that it was because then I could cascade and do a bunch of other stuff which is really difficult to do with rings-- was with reliability and serviceability. It's difficult to isolate a problem in a ring. With a switch, you could tell.

Garner: Some of those guys didn't get it.

Ouchi: Yeah, but it a switch you could tell right away pretty quick what was bad. So I was able to now change the topology from rings to switches. Now, later on--

Garner: Change whose? That's a massive statement. Change whose topology? IBM decided not to do token rings? Is that what you're--

Ouchi: No, what IBM decided not to do was the dual insertion ring for--

Garner: For ESCON.

Ouchi: For ESCON. So now ESCON was the switch. And of course, there was other guys who said, yeah, switch makes a whole lot more sense. In fact, the one that I keep saying is, what does the telephone company do? They went from party lines to whatever. They never went to a ring. So then I was given another chance, and I ran off to Tucson and managed part of the test lab there. In fact, I ran test lab for Tucson.

Garner: So mid '80s now?

Ouchi: Yeah.

Garner: That was quite a switch. You had to take your family there.

Ouchi: I took my family there, just like I took them to-- for a period of time, every two years, we were moving.

Garner: IBM's ”I’ve Been Moved.”

Ouchi: I'd been moved. So I took them to New York, brought them back to California. Two years later took them to Tucson. Two years later came back and got to manage an incredible organization.

Garner: Came back to San Jose, you mean?

Ouchi: In San Jose. Always kept property in San Jose, because you hear all these stories about them. You sell your house. You can't come back. So things like that lot that my house was on-- kept that thing.

Anyway, I got to manage this organization that developed the tools to develop products for IBM. Mechanical tools-- and we were doing 3D solids along with Dassault, Electronic design stuff-- and that's just when EDA was starting to become large commercially. Also had a group that was pioneering how we would use PCs in manufacturing and in development.

Strange things like things we take for granted today. PROFS was the email system for IBM, and most of the people never used it. So what we did was we convinced the site manager to begin sending out all these notices and everything on PROFS. And we knew that a lot of secretaries were printing these things. The manager would read it, mark it up, send it back.

And then we taught the lab directors how to actually--

Garner: Type.

Ouchi: Type.

Garner: It was a serious problem back then.

Ouchi: It was. And so we finally got them to now begin to use PROFS as the communication means.

Garner: Yeah, that's amazing. We can't imagine it today because everyone has to either thumb or type, but back then professionals didn't type. They assumed administrative assistants or secretaries typed.

Ouchi: You had an assistant.

Garner: That's one of the impediments of the early PCs, by the way. And what we did at Xerox with Xerox Star-- people said, why would I buy this? I would have to type. So you were helping get people to at least understand--

Ouchi: To understand.

Garner: To make this transition so they can be more effective and use electronic media directly for shorter latency and turnaround.

Ouchi: And we were actually getting some of the first inklings of the kinds of problems you could have when people started to send out these emails with large distribution lists saying we're going to have this party, whatever.

Garner: Criticisms. Hard to interpret.

Ouchi: No, no, no. These were just-- they were nice things, but they were informing people about something that's going to happen, or something's going on, or something like that.

Garner: Well that's OK. Just announcing things.

Ouchi: Right, but this was chewing up a lot of bandwidth and capacity in those days.

Garner: Oh, it was. I see.

Ouchi: So somebody got this bright idea saying, well, we should send a note to everybody that you shouldn't do this. And of course, this note to everybody--

Garner: Went out.

Ouchi: Went out and it killed the system.

Garner: That's funny.

Ouchi: So anyway I had this organization that we were developing all these really cool tools. What was really neat was the 3D solid modeling organization. We finally got parts of the disk drive company to begin to use real good analytics tools, because we were able to now take-- when you think of a disk drive, you think of this thing that does this kind of thing or whatever it's doing. In reality, that thing is like recording on Jello. Everything is shaking. These dimensions--

Garner: Recording on Jello-- yeah, that's not quite an analogy I've heard before.

Ouchi: But it's everything is shaking. The actuator moves. Everything shakes. And so a real good design would try to move the frequencies high enough so it never bothers you. And of course, in those days, it was all cut and try, whatever. And we showed them these tools that you could now do this, and they actually designed heads, and arms, and drives.

Garner: To look for resonances and stuff.

Ouchi: Yeah, to get the resonances and move them out. And they could, in an afternoon essentially, take something that took them months before to close this in. And so now the disk drive guys were really sold on 3D solid models.

Garner: So these were commercial products available outside of that?

Ouchi: They were both IBM and additions to commercial products like the CATIA system. In fact, we got to go up to Boeing many times and present to the Boeing management about why it was a really neat idea to use 3D solid modeling.

So anyway, one of the other things that was going on in my little PC organization was the thoughts of we need to build-- I wanted to build a RAID device, because I could see that the three and a half inch drives we're going to start to take to market. You could now start-- at that time, get a gigabyte on a three and a half inch, and I knew that were going to be able to get a couple-- or three gigs.

Garner: So this is late '80s, early '90s.

Ouchi: This was the late '80s, early '90s.

Garner: And you were head of the PC product?

Ouchi: Yeah, I had the tools department.

Garner: OK, so PC stood for product--

Ouchi: No, no. This was just regular PCs.

Garner: Oh, OK. But you were still head of the tools department.

Ouchi: Yeah, I still had tools. We also did things like develop the shop floor control system and a bunch of other stuff that manufacturing-- those guys are going crazy because theirs ran on a 370. Anyway--

Garner: Going back to the RAID thing before because you just mentioned RAID. So Patterson published his paper in '87. So when was the first time you heard about it or you first saw it?

Ouchi: Well, I had heard the terms, and by '87, '88, we were well into trying to develop a product that was going to be RAID, but use IBM technology.

Garner: But it was kind of de ja vu for you in some ways?

Ouchi: Yeah, it was because I knew--

Garner: Patterson was popularizing in his work, and then companies were even being founded based on the principles, like EMC. And you guys realized maybe you better catch up with yourself.

Ouchi: Right, because we were being driven like we were a disk drive company, as opposed to a storage company if you can understand that.

Garner: Yeah, you weren't being driven with the right mindset.

Ouchi: Correct. And so they had spawned off another group to build an array, but it had a bunch of Poughkeepsie guys in it, and they were going to put a storage ring inside of it, and--

Garner: Do you remember the name of the project?

Ouchi: It was called Sea Star.

Garner: Well, I certainly know about that. Some people were in San Jose Almaden worked on SeaStar.

Ouchi: Yeah, it was--

Garner: Claude Barrera was one of those people.

Ouchi: Yeah. All the guys who later on became famous or whatever, but it was a huge sink.

Garner: It was a big project. It was big project.

Ouchi: It was a huge project.

Garner: Yeah, I thought you had probably 100 engineers on it or something.

Ouchi: Yeah. Anyway, so--

Garner: It was going to be based on a ring?

Ouchi: Yeah, it had a ring in it.

Garner: Even though that idea died for ESCON?

Ouchi: Yes. So I said let's put a switch in it. And they said go away. Go away. I was running tools. For some reason, I ended up running the test lab in San Jose, but I still-- I wanted to bring with this with me.

Garner: You're willing to do the hard work. Do whatever it takes.

Ouchi: Yeah, so I said, well, let me take my little group of PC guys with me, and they said, sure, go ahead. And so what I was able to do was get built an adapter for the disk drives using an Inmos transputer, which had serial ports on it, had a serial switch-- I think it was like a 20 or a 30 megabyte serial port-- had serial switches that could light, practically. And so I designed this system that had like 30 transputers in it that would handle 256 drives, and because we had these serial interfaces and everything, by using flex cable as opposed to a backplane-- this was a backplane-less array. You could build by just inserting in the front and putting little what later on would become USB ports. And so I had that machine designed.

Garner: So you guys built a prototype of it?

Ouchi: Yeah, and so--

Garner: What was it called? What was that project called?

Ouchi: It was called Cortez. And we actually built pieces of it, demonstrated it. It actually worked.

Garner: They must have thought you were nuts.

Ouchi: Yeah, they were.

Garner: The fact it worked didn't help, right? Transputer was looked down upon by many people as low performance, kind of wacky.

Ouchi: Yeah, you get to write code that was parallel code.

Garner: It was hard to read. Yeah, they had their own language.

Ouchi: Yeah, it was called occam.

Garner: occam. Did you write in occam?

Ouchi: Yeah, we wrote in occam.

Garner: Wow.

Ouchi: And so I had essentially programmers that were co-op students. Had a great time.

Garner: Students.

Ouchi: And the most difficult thing that we had was the fact that we had surface mounted components, and the guys weren't very good at building this thing. And it has wires on it. Wires would come loose. I'd have four cards. No four were the same.

Garner: Now, did this have RAID? What kind of RAID?

Ouchi: Yeah, this was a RAID.

Garner: How did it implement RAID in this case?

Ouchi: It had eight drives across, and you could have like eight-- I can't remember how it was-- front back, but we did one little drawer of RAID drives.

Garner: Eight? Seven plus one?

Ouchi: Seven plus one. And we would take count key data, put it into a semiconductor buffer, run the tracks, stuff like that, cache in and out of it.

Garner: So re-implementing Zark, basically, with---

Ouchi: Well, Zark only-- we were now running, instead of 3330s--

Garner: Well, of course. But the concepts, I meant.

Ouchi: Yes, it was Zark but we have now a--

Garner: More modern disks and more memory.

Ouchi: We had semiconductor memory.

Garner: Yeah, you had more memory now.

Ouchi: So now we were reading and writing out of the memory.

Garner: How much memory did you have?

Ouchi: Oh, I can't remember.

Garner: Several megabytes?

Ouchi: Yeah, it was maybe--

Garner: 10 megabytes?

Ouchi: No, it was larger than that.

Garner: 32 megs.

Ouchi: 32 megs. Something like that. It would have been what would have been a huge PC in those days. Anyway, so we were reading and writing out of there, getting good hit ratios out of it, and then storing back into these drives. And so we had our little eight cards with our drives, and hooked it up to a channel, and the thing would actually run.

Garner: And since SeaStar was the project of record, you guys were ignored?

Ouchi: We were kind of off on the side, and then they begin to realize that we were probably real. In the meantime, they had--

Garner: They, as management? San Jose--

Ouchi: The management, San Jose management, division management. Because what was really happening was you could see that this one box would be the equivalent of three and a half strings, and a string was a four box unit of 3390s. And so we were half of the size of this room, but now sitting in this one box-- that you could build all the drives in this box on this little mechanical line--

Garner: Did SeaStar-- did it assume 3390 drives, SeaStar? No.

Ouchi: No. It was using small drives, too.

Garner: It was using [commodity] drives, too.

Ouchi: It was using small drives, too. But they were still stuck in the mud. They couldn't even get a machine to work yet. And mine was up and running.

Garner: Did they ever-- I can't recall whether SeaStar ever did [work].

Ouchi: They killed it.

Garner: Before it worked, right?

Ouchi: Yeah, it never worked.

Garner: So maybe what you're saying is the effect of-- what was the project called again?

Ouchi: Cortez.

Garner: Cortez-- could have been that maybe there's a better way to do something, but then what came next then after SeaStar?

Ouchi: Well, what happened was I had gone off to contract manufacturers to actually have somebody build this for us, and I was really lucky in that my manager at that time was out of Rochester.

Garner: So he wasn't around very much.

Ouchi: Not only was he not around, but he was very supportive, because he had done Silver Lake-- which System/38. And he did that using all contract manufacturers.

Garner: I see.

Ouchi: The name is Dwayne Duker. I don't know if that name would mean anything to you. Anyways, so he took me around, and I met people like Plexus and several other contract manufacturers, and had designed it, in fact, to contract to Plexus to build this box for us. So we brought back all the stuff. IBM was all set to go, and then the components division and everybody said, hey, we want to get a piece of this. Let us do this instead of contracting out. We'll do it for whatever you thought that you're going to get them to do it for.

Garner: These are the guys back in the Poughkeepsie?

Ouchi: Back in Burlington.

Garner: Burlington.

Ouchi: They were were the semiconductor guys. Plus I don't think they liked having all this Inmos stuff inside of there. Looking back at it, what I was doing was just running right into the giant buzz saw.

Garner: Radical. Yeah, it was very radical.

Ouchi: It had Inmos stuff. It was being built outside, and when you looked at what was going on with the thing, it was going to essentially take the entire San Jose plant off the map, because you could build these boxes in a large garage, burn them in, and the entire capacity that they expected out of the Building Five plant you could do in the corner of a big garage.

Garner: Not politically correct.

Ouchi: No, so anyway--

Garner: Just like the issue IBM faced with risk in the 801 project.

Ouchi: Yes. Yes. So anyway, I had this project going. I thought, maybe I'm going to be able to save IBM. And then all these political factors came in and essentially said, well, we decided we're not going to do it.

In the meantime, I had been-- I had friends at Solectron-- Koichi Nishimura and Winston. And they had been bugging me, saying why don't you come and become the director for us. So finally when they killed my project, and I knew this was it-- so I actually took a corporate VP position as the chief information officer for Solectron. And this was a whole different kind of thing, but I had all this experience of knowing all these bills of material problems and all those other issues like that. They were a very curious company from the standpoint of every site had its own MRP system, and it was a different kind. Not two sites had the same-- you couldn't even call them ERP.

Garner: How many sites did Solectron have?

Ouchi: At that time, there were three.

Garner: All in the Bay Area?

Ouchi: No. We had the one in the Bay Area. They had completed the one in Penang, Malaysia.

Garner: I see.

Ouchi: And then they were in the process-- actually four-- in the process of acquiring two IBM sites in Charlotte, North Carolina and one in Bordeaux, France. And so they were all running on different systems. And so my job was to pull together a single view for the corporation. And of course, they were getting all this advice from consultants saying you need to move to a single ERP and to do all this stuff. And I told them forget that, because you can't afford it.

So what I did was essentially used Excel extracts and was able to use-- created a network of, at that time, dial up, and was able to create all the financial materials information into these extracts so that, once a week, our CFO and materials people could see a view of the company from a financial standpoint and a material standpoint, because with the margins that they had, you couldn't run very long because you'd go out of business before you know it. One of the things that the CFO demanded, which was absolutely correct, was that she keep all the materials as actual cost. None of this craziness of adjusting. You had weighted average actual cost. Anyway, so that was one of the first things that I did for them-- was to pull that together.

The other thing which was, in a way, really incredible, and it's too bad-- it's a good thing they didn't have that when I was still at IBM, because what I did was-- was able to find some tools that would allow you to take the CAD from a company and actually create the printed circuit board placement programs. Also had the ability to take the CAD and create the actual six layer, eight layer, whatever layer boards, because we had board manufacturers who will do that for us if we gave them the aperture stuff. We had the capability-- in fact, used a lot by Sun and by--

Garner: That's where I was at the time. We were a big user.

Ouchi: We could turn around a prototype in a day.

Garner: Very quickly.

Ouchi: And of course, it would pay dearly for that, because the turnaround time was just incredibly-- so if I had that capability when I was trying to build my little Cortez machine, I would still be at IBM.

Garner: And that project was still politically incorrect.

Ouchi: But it would've been working.

Garner: Yeah. After they canceled your project, they went off and did basically the Shark, right? The DS8K, which became a big monster machine.

Ouchi: Shark was essentially Cortez on-- instead of being a distributed processor, it had a single--

Garner: Single power processor.

Ouchi: Yes, as opposed to distributed machines.

Garner: You were ahead of your time.

Ouchi: Somebody said I'm 10 years ahead. It's too bad I can't make a lot of money at being 10 years ahead.

Garner: Yeah, that's probably about right. So yeah, Solectron became a very powerful force in Silicon Valley.

Ouchi: Yes.

Garner: A major manufacturer for most of the companies, but now it seems like a lot of that's gone overseas. Does Solectron still have a presence?

Ouchi: We were part of--

Garner: A transition.

Ouchi: We were part of that transition. Part of it was there was a lot of Chinese in our company, and so there were Chinese and Persians-- if you could believe that that was the core of the company-- so it was a very international company. And so we founded a site in Suzhou and then several more that they acquired in China. A lot of the guys that I knew are now extremely rich, because they founded companies in China.

And so one of the other pieces of getting the prototypes out in a very short time was the fact that I could now take the bill of material that the company had and the bill of material that was actual for the CAD, and I could resolve all that craziness on the phone with the engineers before we actually built the thing. So they had clean bills. They had clean machine programs. You can actually go into production the next day, which was-- if you've seen before, everybody had these problems with bills of materials. We also add the issue of part numbers. You run into all these administrative problems that are crucial for good quality manufacturing.

Garner: Very impressive. Then you finally retired from--

Ouchi: Retired from Solectron in 2000.

Garner: 2000.

Ouchi: I had thought that using internet connected relational databases would be a really neat way to provide capability to small companies. So I had built a shop floor system. This was one of those, well, what am I going to do? Two things happened. That plus my kids finally decided to have kids, and so my wife--

Garner: That's magic.

Ouchi: Yeah, it was really magic. Because I was offered a job at IBM, and she said, you can go, but I'm staying. You missed growing up with your kids. I'm not going to miss growing up with our grandkids. So I was at home, and so built a really high functioning shop floor system and some purchase order management stuff like that all on the web. Had a number of companies that were just running browsers. So they were running shop floor just shooting the bar codes just connected to the browser. But unfortunately what happened was they were mainly prototype shops. One by one they evaporated. They're gone. All that business disappeared. So then I decided to come wander into the Computer History Museum.

Garner: Good move.

Ouchi: It is. This is really fun.

Garner: So you've been trained as a docent now?

Ouchi: Yes, I am.

Garner: We can continue to inspire kids by seeing how much the industry has changed over the last 50 years.

Ouchi: It is incredible what has changed. And one thing I keep reminding people is, even though it's called history and museum-- that you could actually see people who actually did this work.

Garner: That's certainly true, and we have our running machines, 1401, and the PDP1, and the RAMAC drive.

Ouchi: The RAMAC.

Garner: Which gives some life to some of the things. So just one more question back on RAID. You said you finally met Dave Patterson at the museum here.

Ouchi: No, I met him at an event at IBM Research, and he kind of-- oh, OK.

Garner: OK what?

Ouchi: It was a, so you're him? Like to me.

Garner: OK, so you just acknowledged each other, but didn't swap stories or anything.

Ouchi: No, didn't say anything about anything.

Garner: Because for me, Dave had done-- some of the ideas that he-- I don't know all the stories, but he was still consulting with us at Sun Microsystems. He was in my office with a colleague of mine at Kelly every day for still many years after we did the SPARC work. And we were sitting around thinking, “Well, what could we do with Sun-4?” We could build a storage system out of them, and they [Berkeley] built their first RAID prototype using the Sun-4 workstations that we had just finished designing at Sun.

And that was his prototype of his RAID system, and it's hard to know whether the ideas came out of those discussions or he already knew about them, but it made for a great-- yet again, another great student project. And then his [1988 The Case for RAID] paper, a motivator from the creation of all [storage] industries-- that shows you the power of getting things published and distributed. So you guys didn't trade notes or anything?

Ouchi: No, we never did. He never said, “What made you think of that?” or anything?

Garner: We'll have to get you guys together sometime.

Ouchi: That would be interesting. And who knows because of everybody running around or whatever, because it was some IBM event that this kind of hard for him to-- in fact, I'm not even sure whether it even connected in his mind.

Garner: Well, we'll work on that then. That's one of the purposes of the museum-- is to get people together who may not have had a chance to meet otherwise. It happens a lot here. All right, well do you have any last minute words? Certainly your story is kind of from the root and soul of western America, and it matches with much of the history of what it was like to grow up in western America in California.

Ouchi: It was a very interesting time. One other piece. This was when I was already an adult, and I know my parents had run into my eighth grade teacher at some event at my high school or someplace like that. And he gave them a paper that I had written for career day, or whatever it was, was like that. And what I recall was you had to write a paper saying, what do you want to do with your career? And my dad said, “Why don't you write to IBM and see what they say?” And so I wrote this little thing, and they sent me back some brochures.

And so I had written this paper about all these great things at IBM that they'll send you to school and all stuff like that. And I said, well, what I'm going to do then is I'll go to some university, get my bachelor's degree, and IBM will pay for my advanced degree, and I'll get a PhD, and I'll invent really great things at IBM. So that was my eighth grade paper.

Garner: Really? Impressive.

Ouchi: Got an A minus.

Garner: Wow. Wow.

Ouchi: Not quite like the FedEx guy who got a C.

Garner: Right, yeah. All right, well any other last thoughts? I don't think I have any.

Ouchi: No, this is interesting.

Garner: Well thank you so much for all those stories, and thank you for continuing to contribute here at the museum.

Ouchi: This is fun.

END OF INTERVIEW

