Oral History of Gary Davidian, part 2 of 2

Interviewed by:
Hansen Hsu

Recorded March 14, 2019
Mountain View, CA

CHM Reference number: X8931.2019

Hsu: Today is March 14th, 2019, I am Hansen Hsu back with Gary Davidian.

Davidian: Good to be back.

Hsu: So, we're going to take up where we left off with the beginning of your work on the RISC emulator. As we were talking before, you had mentioned that right around the time of the 1989 earthquake that it was Phil Goldman and Erich Ringewald who were on the Jaguar Project and asked you--

Davidian: Yeah, I think they were working on the Jaguar Project and the Jaguar was using the Motorola 88000 RISC processor.

Hsu: Okay, so, at that time, maybe give a little bit of background of what sorts of RISC projects were actually going on at Apple?

Davidian: At that time, as far as RISC goes, there was probably two that I can think of. One was Jaguar, which was a major-- that would be a desktop, but the other was Newton.

Hsu: Oh, right.

Davidian: With the ARM processor.

Hsu: Right.

Davidian: I don't actually remember if at that time they were using the ARM processor. They had actually gone through a few iterations, I believe. I know at one time they were using-- they were looking at an AT&T processor. I believe it was called CRISP, if I'm correct. But it--

Hsu: I think it had a code name "Hobbit" or something like that?

Davidian: Maybe, yeah. So. But, in the end, they used the ARM.

Hsu: The ARM. Yeah. I heard there was also a project called "Hurricane", which became "Tesseract" later on?

Davidian: Yeah, that was Jaguar. Jaguar sort of evolved. There was some evolution there. Yeah, Hurricane was one iteration of that. It had-- they had a team sweatshirt made up. It said, "Hurricane: Prepare to be blown away." And we took that as a bad sign for a project.

<laughter>

Davidian: Then there was Tesseract and then after Tesseract there was TNT, which was the new Tesseract or the next Tesseract or something like that, which we just jokingly said, “Prepare to be blown up.”

<laughter>

Davidian: I was not really-- you know, the whole Jaguar and all those follow-ons are-- they reminded me too much of my Data General experience.

Hsu: Right.

Davidian: And Pink also. And then throwing everything together, you know, Pink-- they were going to use Pink as the operating system.

Hsu: Oh, Jaguar was going to use Pink as the operating system.

Davidian: They were looking at that.

Hsu: Okay.

Davidian: Yeah, that was-- yeah, that was one thing they were considering, I believe. But they were not-- you know, compatibility with Macintosh and Macintosh applications was not one of their goals.

Hsu: Right. So, these were all sort of completely new platforms, completely new projects, and no regard to compatibility or bringing on the existing--

Davidian: That’s a pretty good description, yeah, at least the way I understood it. I wasn’t actually on those projects.

Hsu: Right.

Davidian: I was avoiding them.

Hsu: Right. What was the difference between these various projects? Like, what were the changes between the--

Davidian: I don’t really know for sure. The way things tended to work at Apple was, you know, a project would get cancelled, but parts-- some of the components of it, maybe some of the ASICs or things like that would live on. Some of the technology would live on into some newly formed project that’s-- you know, somehow meets the criteria that the other one didn’t. So, some of it does carry on. I don’t recall if anything from Jaguar actually made it out. So.

Hsu: Okay.

Davidian: So, yeah, that's sort of my recollection of Jaguar. It was just-- again, it was, you know, where I know how the movie ends. And I didn't want to, you know, watch it again or live through it again.

Hsu: So, then-- but then these guys from the Jaguar Project got you to start working on an emulator--

Davidian: Well, it was--

Hsu: --for the 88000.

Davidian: They didn't actually-- they wanted me to think about it.

Hsu: Okay.

Davidian: It was-- because Phil Goldman, in particular, he worked on Switcher-- not Switcher. What-- MultiFinder is what Apple called it. Yeah. And compatibility-- you know, his experience on that was-- he didn't want to break applications. So, he really liked to make things as compatible as possible. So, I think his thoughts were, you know-- he was thinking about Macintosh application compatibility and emulation as one way to get there and was just-- you know, wanted my opinion on whether that was even possible. And, so, that's-- you know, that's sort of what prompted that whole discussion and-- but it wasn't actually part of my job. It was-- I wasn't assigned to that or anything like that. I was-- this was-- I was still working on the Mac IIfx at that time.

Hsu: Oh, right.

Davidian: But it did get me thinking about it. That was sort of the first time and, you know, it got me to-- I started getting familiar with the 88000 architecture at that time.

Hsu: So, then it wasn't until after you finished work on the IIfx. Was that when you started on this?

Davidian: So, what-- sort of the history of the original Power Macs-- oh, and I should point out that today is the 25th anniversary of the release of the first Power Macs, the 6100, 7100, and 8100.

<laughter>

Hsu: Wow, what a coincidence.

<laughter>

Davidian: It was-- yeah, it was March 14th--

Hsu: 1994.

Davidian: 1994. Yeah. So, you know, we were working on the IIfx and it shipped in March of 1990. And, as is typical there, then the team that finishes, then they're looking for what to do next. And, at the time, there was-- Apple had the Advanced Technology Group and there was a performance group within that and they had done some studies of where some of the time gets spent. They did performance studies of where the time is spent. Is it in the applications? Is it in the Toolbox, the OS? So, and they had produced some good PowerPoint slides of their analysis. Also, around that time Apple released a NuBus graphics card, I think it was the 8•24GC card that had an AMD 29000 processor on it.

Hsu: Right.

Davidian: So, that was another RISC project at Apple that I didn't think of initially. So, one thing we were—the remains of the IIfx team was looking at was there some way to integrate that card along with-- you know, we would be using the 68040 processor or maybe some follow-on to the 040. We were even-- we had meetings with Motorola looking at their processors beyond the 040.

Hsu: Right. Like, there was an 060--

<overlapping conversation>

Davidian: Yeah, there was an 060. I don't know-- you know they sort of skipped the 050 or there might have been some internal 050 design, but the 060 was the next thing that they actually produced.

Hsu: And Apple then shipped the Quadra and Centris line.

Davidian: Yeah, those were under development. I was also helping out a little bit with the Quadra development and some of the bring-up of that. And some-- yeah, I wrote the new block-move routine that—because the 040 had these move-16 instructions. But, yeah, mostly I was-- during that time period, right after the fx shipped I was sort of decompressing, too. So, we shipped in March. I think it was probably around June was when we started having discussions about what to do next. And, so, we started looking at the AMD 29000 at that time. And because that was what the GC card used. And they were looking at sort of having an 040 and a 29000 on the same board and trying to share memory and--

Hsu: Both being the CPU or--

Davidian: Well, the thought was that the 29000 would be a graphics accelerator and also a Toolbox accelerator. We would do more than just QuickDraw. It would be-- you know, we would try to optimize other parts of the Toolbox that the performance studies showed were using time.

Hsu: Oh, okay.

Davidian: So, it was really a Toolbox accelerator was the thought. So, the hardware guys were looking at that and they were saying, "This is really a pain trying to do this. Is there some way we can get rid of the 68000?"

<laughter>

Davidian: So, which, you know-- although we were still having discussions with Motorola. But, you know, so we started thinking about it and it was-- could we emulate the 68000 on the 29000? So, that was the first one that I looked at. Also, at that time, they brought in Brian Case, who was one of the architects of the 29000 when he was at AMD.

Hsu: Wow.

Davidian: He had also been at Apple working in the ATG group for a while. So--

Hsu: Oh! I mean, there was that attempt to come up with a completely new-- what was that called?

Davidian: Yeah, “Aquarius”.

Hsu: Yes!

Davidian: I never really was involved with Aquarius or it sort of disappeared and I never really got to see much of what it was, but some of the people that worked on it did remain and, as we went on to other projects they got involved in some of the architecture discussions and things like that.

Hsu: Yeah. So, was Brian Case one of them?

Davidian: He was hired for ATG for, I believe, for that project.

Hsu: Okay.

Davidian: So, it was really a very small group at that time that was looking at what to do next. And I’m trying to remember-- it was probably Jack McHenry was running it. Jon Fitch was the hardware lead. Bob Hollyer was the hardware designer and I was the software guy at that time. And, so, we started looking at that, but we also looked at other RISC processors just to see and we looked at the MIPS processor, SPARC, and then we also looked at-- it was seeming that despite whatever we wanted to choose, there was sort of this corporate pressure that if it’s going to be a RISC processor, it’s going to be Motorola and it’s going to be the 88000, because that’s what Jaguar is doing.

Hsu: Right.

Davidian: And, really, the 88110 was the processor that was still in development and that was what Jaguar was really intending to use. So.

Hsu: Was that because Apple wanted to continue to maintain that relationship with Motorola?

Davidian: I think so. I think that's probably what it was. Because Motorola came up with the 88000 architecture. I don't believe it was for Apple. I think that was their own doing. I think Keith Diefendorff was probably the force behind that at Motorola.

Hsu: Because, I mean, you would think that when you're switching architectures, that's a great time to change the supplier.

Davidian: Yeah. But I think it was that relationship that--

Hsu: Yeah.

Davidian: --had them looking at it. So, we were looking at all these different processors. I was coding up partial emulators for each one of these processors for-- limited number of instructions just to see what performance might be like and it-- and if there were any major showstoppers. There was one from MIPS. MIPS we had to rule out just because they divided their address space between user and a kernel mode, even at the user level. So, that would have been problems with us trying to map in the NuBus space-- you know, trying to keep the memory map that was used in the current Macintosh designs. So, we started-- I think we were leaning towards the AMD 29000, but we got pressured into looking at the 88000. And, so--

Hsu: Was AMD a better architecture, in your opinion?

Davidian: It's hard to say. You know, given that we never developed a complete emulator for it and, so, I—and we never built hardware to run it. I don't really know the answer to that and I don't know what-- if you started also looking at where the companies were going for follow-on products--

Hsu: Right. The roadmap--

Davidian: I think the roadmap looked better for Motorola.

Hsu: Mm-hm.

Davidian: And, you know, political pressure.

<laughter>

Davidian: So, we actually decided to build a prototype, but actually while the prototype was being developed, from-- one thing from the Jaguar Project was they had-- sort of their development environment was they had NuBus cards with the 88100 and 200 chipsets on them, so that you could actually do development on this NuBus board, NuBus card. This was also because the 88110 was not available yet. So, this was still using the original 88100 and 200. So, it was a three-chip set. There was a CPU and then two cache MMU units, instruction and data. And it was on a NuBus card. They had-- it was called the "Cub Card" for little Jaguar. And they had a design where it actually had-- it was a dual processor. So, there were actually these six large chips on it and some memory also. So, that was what

we started using. There was also a card, I believe, made by Tektronix that-- an 88000 card, a NuBus card by Tektronix that we-- that was also used somewhat interchangeably for development. So, that was the original development environment emulator. The first attempts I took was emulating the 68000, not the 68020 or 30. And emulating the Mac SE. So, I had a Mac SE next to, a Mac IIx with a NuBus card, and I was developing the emulator in that and it would actually use the I/O from the Mac IIx as the-- you know, as the I/O for the SE. So, I was just developing the CPU emulation and creating a memory map that looked like a Mac SE memory map. And it took a while, but I eventually got that up and running. And once it just started booting to the point where MacsBug could load, then I could sort of debug side-by-side with the Mac SE to find problems. But once it was actually running the operating system that's when I then started evolving the emulation to a 68020 emulator. So, I had an up-and-running Mac OS on this emulator and I could start adding 68020 instructions to it, but the system wouldn't use them, because it thinks it's a 68000. And, so, I could then write applications that would use them and test them and that's how I got confidence before actually trying to run an operating system that actually depended on the 020.

Hsu: Okay.

Davidian: And then, after that, I believe I actually had turned it into a Mac II emulator or Mac IIx emulator, again, using the I/O of the real Mac II. And, in parallel, the development of the-- what we call the RLC was going on. It was a RISC version of the Macintosh LC and it had an 88100 CPU and two 88200 cache MMUs. And I was able to develop the emulator for that, that ran the Macintosh LC ROM unmodified. So, it could run any version of the OS that the Macintosh LC itself could run. So.

Hsu: Okay.

<break in recording>

Davidian: So, this is a Macintosh RLC. It looks exactly like a Macintosh LC. In fact, it probably says "Macintosh LC" on it. It's the real Macintosh LC plastics. So, it had to fit into the same form factor. I think that's the 88100 and the two 88200s or maybe it's-- there's three of them.

Hsu: They're sitting on a daughter card.

Davidian: Yeah, I think space was a bit of a problem. And, but it's-- the card below it is-- the logic board below it is also a custom design. It's not the Macintosh LC itself. There's no 68000 processor in here. So.

Hsu: Yeah. And, again, there's really one CPU and then two support chips.

Davidian: There are the caches and the MMUs. So, yeah, the 88100-- you know, the original 88000 was designed as a three-chip set.

Hsu: Oh, okay.

Davidian: So, it was too big for one.

Hsu: So, the RISC LC was that-- that was a separate project that you were also helping with or working on?

Davidian: I was working on that. That was actually my job.

Hsu: Oh, that was your main project.

Davidian: Because this is-- I sort of stayed with the team that worked on the fx.

Hsu: Okay. And, so, was the next project that that team worked on.

Davidian: This was the next project after the fx or what we thought.

<laughter>

Davidian: Little-- things changed.

Hsu: So, then on the side you were working on the emulation stuff, the emulators.

Davidian: Well, this was the emu--

Hsu: Oh, okay, well, yeah, this is for that. Of course.

Davidian: Yes, this was the emulator at the time.

Hsu: Right, right.

Davidian: But this was really just a proof of concept. There was sort of concurrently with this was the-- on the software side was Macintosh System 7 was going on.

Hsu: Right.

Davidian: And the management did not really want to devote any resources to any of these hardware projects, except for me. So, I was the only resource from Systems Software that was assigned to this follow-on project.

Hsu: Oh, right, because-- okay, so, working out in terms of the org chart of the divisional, like, this was a hardware project, but you were part of Systems Software.

Davidian: Yes.

Hsu: So, you were supporting them--

Davidian: Yes.

Hsu: --but your line of reporting was to software, not to hardware.

Davidian: There was a lot that went on around that time. I don't remember the exact times it happened, but when I first started at Apple within the System Software group there was what was referred to as the ROM group, which worked on the software that was in ROM for the Macintosh products. But the ROM group worked very closely with the hardware groups, because we were supporting all the new hardware products. Some re-org that happened probably around the time that the fx shipped or maybe-- that general timeframe, it was sort of after the earthquake, moved the ROM teams into the hardware division instead of in software. And I felt that was a bad move and I really didn't want to move to the hardware group. I'm working on software, the software is where it belongs, and I also didn't want to move my office. So, I stayed behind. I stayed in System Software, but I was then-- my job was to work on RLC at that time.

Hsu: Was that the origin of the CPU Software division part of the hardware group? Because I remember when I was at Apple in the early 2000s, there was a separate CPU Software division that was part of hardware.

Davidian: Yeah, that is probably the origin of that. And, yeah, that's probably when it happened. So, I was in Systems Software, but for a while I don't quite remember who my managers were, because I sort of didn't have-- I didn't really report to anyone for a while it seemed. I think technically I did, but, you know, realistically I didn't, because I was separate-- doing something completely different than the rest of System Software. Everyone was focused on System 7 and I wasn't.

Hsu: Yeah.

Davidian: Because everything that I was supposed-- everything that I had worked on that was supposed to ship in System 7 actually shipped in System 6.

Hsu: Right, you mentioned the Time Manager.

Davidian: Yeah, so, that's why I was also free to focus on this and also just my background made me a perfect match for this. So, I think as things evolved over time, yeah, once-- yeah, while we were still working on RLC, I think it was-- I think, sort of time-frame, it was probably around mid-90 was when things started. I know it was probably around-- it was the end of '90, early '91 where I believe I had emulation up and running on a Mac IIx, you know, with a NuBus card. And I think it was probably mid-'91 when RLC was actually up and running. And then I think it was-- well, one thing I did, one I had RLC up and running, I actually did some experimentation to look at actually using-- running RISC code from Macintosh application. So, you know, I put in some hooks to actually run-- let you actually run 88000 code and I had some support for actually making toolbox calls from that 88000 code. So, I took one of those DTS applications-- I think it was called "Traffic Light" at the time and implemented enough of this interfacing back to the Toolbox so that that application could run and it was just a way to show that you could-- you

know, compiling for 88000 and running it with an emulated operating system. So, that was sort of proof of concept, but some of the stuff that I did for that did influence some of the Mixed Mode later on, but we're jumping ahead a bit. So, but that was sort of the first native applications.

Hsu: Right, yeah. So, then at what point did the PowerPC come into the picture. You know, the whole agreement with IBM had to take place first.

Davidian: I think it was around October 1991.

Hsu: Okay.

Davidian: Where it was just-- we were all called into a conference room and I think someone wrote on the whiteboard, "DOTC". And we were trying to [figure out]-- "What does DOTC--?" It was "Deal of the Century".

<laughter>

Davidian: So, you know, we were basically told, this is happening and PowerPC was-- well, POWER. There was no PowerPC really. The POWER architecture was not one of the RISC architectures that we looked at, because it was an IBM internal architecture.

Hsu: Oh. They weren't selling it to other customers.

Davidian: No. It was their internal workstation architecture. So, the gist of things was, you know, we have a very limited amount of time to determine if there's any showstoppers for switching to this processor. So, we very quickly had to get very familiar with the POWER architecture and I started, you know, seeing what it would take to code up an emulator in it. And I didn't really come up with any showstoppers. That was actually lucky. There were a number of things that they did that were good. They were big-endian. And because the 68000 was also--

Hsu: Were all the other RISC chips that you had written emulators for also big-endian?

Davidian: I believe they were or they had a mode bit.

Hsu: So, they could switch.

Davidian: Yeah. So. But, yeah, we did not come up with any showstoppers, but there was also-- this was sort of the formation of the AIM alliance. So, there was no-- there was no hardware either. You know, there was just starting on the definition of the PowerPC architecture, which, you know, evolving POWER into PowerPC. So, you know, it was basically a reset in some ways.

<laughter>

Davidian: In other ways, it was-- I knew it could be done. I knew the emulator could be done. That was one thing that the whole 88000 and RLC proved--

Hsu: Right.

Davidian: --it is possible.

Hsu: Yeah.

Davidian: And I knew all the issues that I had to deal with, the basic design. So, it was in some ways a porting exercise.

Hsu: Right.

Davidian: But it also required a lot of thought about how to map it all onto the PowerPC processor. So, I think at that time, you know, once the agreement was finalized, it was going to be about a year before-- first silicon of the 601. And the 601 was actually going to be produced by IBM, not by Somerset, which was the joint development facility in Austin.

Hsu: Right. So, IBM itself was going to do it, not the alliance.

Davidian: IBM-- yeah, it was IBM Burlington was making the 601.

Hsu: Right. So, it was a lot more like POWER than the later--

Davidian: Well, the 601 was actually-- the 601, it was PowerPC, but also had all of the POWER instructions. So, it could actually run both. And there were some MMU differences that-- where it did not follow the Power-- it was sort of unique in some ways in regards to the MMU. I don't remember all the details of that. It's been 25 years! But I do remember there were differences. And since this was sort of a sudden change there were other issues where we didn't have development systems. You know, there was no tools. So, there were going to be these RS/6000 workstations, that was going to be the development environment. And I did all my work at home and I really did not want to have an RS/6000 at home.

<laughter>

Davidian: The Development Tools Group at Apple was going to work on an assembler and compilers, but things were taking a while. So, eventually, I actually wrote my own assembler that ran on the Mac and that was so that I could do development on a Mac instead of on an RS/6000. And it was an assembler that was actually source-compatible with IBM's RS/6000 assembler. It was so that I could do my development, but the final build was actually done on the RS/6000. So, it was just a quick turnaround for me, but as far as everything else, it went through the development chain that Apple was really using internally. And before the 601 came out, there was a development NuBus card produced. Ron

Hochsprung worked on it. And it was called-- well, it used the processor that IBM developed, called the RSC processor. It was a single chip, POWER architecture processor. And I believe they used it in some of their RS/6000 workstations. And it had a lot of limitations, but it was-- it performed a lot differently than the 601-- than we were expecting the 601 to perform. But that's what I initially got the emulator up and running on. So, it was just like the Cub Card, except now it was an RSC card and I think probably by that time we were actually using Quadra 700s as the Mac to plug the cards into. So, you know, I developed the emulator again for that, you know, got it up and running. The card was somewhat-- the RSC card was somewhat flaky, but it was enough to prove that the emulator was working. It wasn't very stable for real long-term use, but it got the job done. And then it was probably about a year after the agreement was signed that the first silicon of the 601 arrived. That was-- I think it was probably some time around September 1992 and in the meantime, the development, the hardware team was working on PDM, which was the Power Mac 6100.

Hsu: Okay.

Davidian: 6100/60 at the time.

Hsu: PDM stood for "Piltdown Man"?

Davidian: Some people say that. Yeah, that's--

Hsu: That was never confirmed or--

Davidian: Yeah, it was Piltdown Man. That's what PDM stood for. And there was also the 7100 and the 8100 were in development at the same time. And I don't remember if there was the server also at that time. I think there might have been a 9100.

Hsu: Oh, really.

Davidian: I don't-- I'm trying to remember. Maybe the 9100 was-- no, maybe I'm thinking of the 950, which was the 040-based server.

Hsu: Yeah, the Quadra 950 and then there was later on a PowerMac 95--.

Davidian: Yeah, there was the 9500, but I thought there might have been a server version of the initial ones, but maybe I'm wrong.

Hsu: Well, I mean, I don't know if there was one, but-- the-- none ever shipped called the 9100.

Davidian: Yeah.

Hsu: Yeah. There was that famous Carl Sagan controversy over the code name for-- or BHA?

Davidian: Yeah, there were-- BHA, you see that on the ingredients of all sorts of things. It's a food additive.

<laughter>

Hsu: Well, "Butt Hole Astronomer" <00:40:49> or-- I mean--

Davidian: Yeah, there were a number of names. Code names change. Even on, you know, code names on all sorts of products changed over time. The IIfx had a number of different code names. I always remembered it as F19. But it, you know, it was called Zone 5. I think to developers it was called "Weed Whacker". It just had a number of different names.

Hsu: And, normally, these code names don't make it out to the press, right? And, so, how did this get out and how did Carl Sagan find out about it and then get mad about it?

Davidian: You know, before Steve Jobs came back, a lot of things got out into the press.

<laughter>

Hsu: Right.

Davidian: So, you know, there-- MacWeek used to have a column called Mac the Knife.

Hsu: Yeah.

Davidian: Which had all sorts of, you know, leaks and later on I learned actually a lot of those leaks were intentional. That was actually free advertising. But, you know, things leaked, I don't know how the names got out. You know, I don't know if developers were told of the names or what. But, you know, it's--

Hsu: Maybe it was to drum up hype about the products or--?

Davidian: You know, I don't know. When I used to see things leaked, it never really-- it never seemed to have come from the development team. It always seemed to come from-- you know, because they would get things wrong. You know, so, based on the way things were described, you knew that whoever was saying it didn't really know all the facts. So, yep. But I don't know.

Hsu: Okay.

Davidian: But, anyway, it was probably-- it was September of '92 when first silicon of the 601 was coming back. And the logic board for this first chip was being designed concurrently. I think we were a little shocked when the 601 silicon actually showed up, because we were pretty used to Motorola's scheduling, where nothing was ever early or on time. It was always-- it always seemed to be later and the first silicon you never really expected it to work. You know? But the 601-- you know, IBM was really good and they

were on time. And the first chips were a little interesting because IBM actually had a bug in them, a bug that prevented their internal scan interface from working and so they couldn't test the chips. We didn't depend on that. So we actually, you know, once the chips arrived it was like all-nighters for a day or two; just everyone who had a little piece of it, we were all in the same room in a lab, I think it was in VG5, and working to get the prototype up and running. And it eventually was up and running and I remember us the next day having a conference call with IBM and they were asking, you know, “Did you try out the MMU?” And I go yeah, I'm using it. Is there something I should know? <laughs> And they just want to know, “Did it work? Because we couldn't test it.” <laughs> So, yeah, you know, it was really good. You know, the thing really worked the first time. There were a couple little issues, though, that we had to work around with external hardware, some things on the bus at times they shouldn't-- you know, some unexpected bus activity. But, in general, it actually ran the first time. And we were also pretty motivated to get that 6100 prototype up and running because we were-- the team was-- there was this carrot that was hung out in front of us, which was if we got it up and running in time, we could go and demonstrate it at the Apple sales conference, which was going to be held in early October, I believe it was in Hawaii. So there was this incentive to get it up and running, you know, fairly quickly. And I think we had two boards built. I think maybe we had, originally, gotten five processors and we got things up and running and got to go to Hawaii. In Hawaii, we got things set up the night before and I was going to give the demo the next morning. And sort of to hide things there was this cardboard box we put over the 6100. And the next morning got there, tried things and it was completely dead. <laughs> We brought spare parts for everything. So at that time Keith Cox was actually the person who did the logic design for the 6100 logic board. And so he was still back in the hotel room. I called him and said, you know, “it's completely dead. Bring everything.” <laughs> He rushed over with all the spares and we're, you know, we're backstage and [Apple CEO Michael] Spindler is already talking <laughs> and we're trying to get this thing up and running…

Hsu: Oh my God.

Davidian: …and replacing everything. And we eventually did. We got it up and, you know covered it up again. And then not too long after that, I had to go out and give the demo.

Hsu: Wow.

Davidian: The first thing I did was I wiggled the mouse. I saw the cursor move and <laughs> and a big sigh of relief.

Hsu: Talk about flying without a net there.

<group laughter>

Davidian: Yeah. So we were never sure exactly what happened because eventually the board, you know, the board that wasn't working, it was eventually working, again. We think it was probably just heat from having the thing covered. But it was a successful demo and we were all happy and relieved. And we hung out for another-- for the rest of the week and toured Hawaii. So that was-- that was the first showing

of the 6100 and so that was around October ’92. Things progressed just hardware-wise there were probably a couple of iterations of the PowerPC processor itself, the 601 processor, 601 silicon. And it was also, you know, IBM produced very good hardware, but they always produced it for internal customers. They did not-- they were not used to having external customers. And, you know, they viewed a lot of stuff that we needed to know as sort of this “family jewels”, nobody gets to see that document. <laughs> And that's where the partnership with Motorola, you know, helped out because Motorola was very familiar with having customers for their processors. And so Motorola was producing the document the data books for the 601. So it was, actually, getting out some of this information in what was-- in the PowerPC architecture they were there was known as Book One, Book Two, and Book Three which described the architecture. And then there was a Book Four which was the implementation-specific characteristics. So the Book Four-- trying to get the Book Four for the 601 was always an issue getting it out of IBM, but Motorola eventually produced the data book, which had all the information that was needed.

Hsu: While your team was working on the 6100, the 7100 and 8100 teams were going on in parallel?

Davidian: It's really the same team. It was-- you know, they were just three different logic boards but they were all based off the same basic chipsets. And, you know, the main differences were clock rate. Although, the 7100 and 8100 had NuBus slots, whereas the 6100 didn't. It had a Processor Direct Slot and there was a NuBus adapter for a single card, but that was-- you know, they were all the same basic design. It was the same ROM for all of them, as far as we were looking at things. It was it was all, you know, the same development. But then going on in parallel with some of this was the whole software effort that was-- you know, after System 7 shipped, it was then the whole, you know, actually accelerate-- you know, starting to rewrite some of the Toolbox, having mechanisms for native applications. And that was all going on in System Software and there were teams that were working on that.

Hsu: I mean by this point, your emulator was solid, right?

Davidian: It was pretty solid. Yeah.

Hsu: So the system software started to depend on it.

Davidian: Yeah. And, also, at the time actually besides the PDM prototypes another NuBus card was produced, the Smurf card which had the 601 processor on it. It was actually called-- there were actually two Smurf cards. There was the original RSC card was actually called the RSC Smurf and then there was the 601 Smurf and Ron Hochsprung produced those. So for development people were using-- there were a lot of Smurf cards were used for people doing development. I don't know if those were-- I think some of those were seeded to developers, also. And then, yeah, the PDM effort was still going on in parallel. One thing that happened, you know, during the PDM—originally, PDM was designed with 4 megabytes of RAM. But System 7 started growing. So the emulator is designed for very low RAM usage, the emulator and nanokernel. And the emulator used one 4K page of RAM. The nanokernel used one 4K page of RAM, plus, whatever the MMU tables took. PDM even with 4 meg of RAM, it also, had-- it had RAM-based video, memory mapped video. So I forget what monitor sizes were then, but that may have taken another half a megabyte out of that 4 megabytes. So things were tight.

Hsu: Yeah. Generally, what 14-inch monitors, 640 by 480.

Davidian: I think there might have been a 16-inch by then. Yeah. Yeah, actually, that was another thing that PDM was forced into was there was this-- 61-- those initial Power Macs had this funny video connector on them. And that was this, I don't know what the monitor was eventually called, but we internally the code name for it was the Telecaster monitor.

Hsu: Is it the one with a built-in-speakers?

Davidian: With speaker, yes.

Hsu: For-- that shipped with the AV Quadras.

Davidian: It might have shipped with the AV Quadras, I don't-- yeah. But, yeah, we were forced into, you know, having that connector, which nobody really liked because you always had to have this adapter to use anything else. And then that connector eventually died when they realized that it was too wide for a NuBus card. <laughs> But, yeah, that was one memory from that project was that that video connector and those adapters.

Hsu: So were you still in System Software?

Davidian: I was still in System Software.

Hsu: But still working closely with the hardware groups?

Davidian: No. Yes. Yes, working very closely with the hardware groups.

Hsu: Whereas the other people on System Software were working on the other parts of the PowerPC--?

Davidian: The other part, yes. The big effort, you know, after System 7 shipped was then the whole, you know, the transition to RISC at the software level and development of Mixed Mode, development of accelerating parts of the operating system and Toolbox, development tools, debuggers. There was a lot that had to go on to really, you know, do this transition. But it was a very stable 68000 Mac and a very unstable RISC Mac because of all this new software that was changing it.

Hsu: And your involvement with those efforts were…

Davidian: I was somewhat involved with that. Some of the Mixed Mode stuff I was involved with some of the actual coding of that because those were actually opcodes in the emulator. Another thing I was doing at that time was, as far as floating-point goes we didn't-- I didn't emulate the floating-point unit, the Motorola floating-point instruction set. Apple had the SANE package. It was the Standard Apple Numerics Environment. It dates back to the original MacIntosh and it did floating-point in software. And that's what everything used before the Mac II came out. So there were a couple of reasons for not emulating the

FPU. One was that it predominantly used this 80-bit floating-point format, whereas that was not a floating-point format that was supported by the PowerPC architecture. So it wasn't that like there was this fast RISC instruction that I could use that would actually do this. So it all had to be done in software to emulate it. And then once you say you're emulating it you have to emulate everything, whether people use it or not. <laughs> And so there was a lot to the FPU that, there were all these transcendental functions like the sine, cosine, things like that. And then, you know, if you're going to emulate it you actually want to get the same results and getting those same results would be very difficult. <laughs> What I did instead was I implemented some instructions to accelerate the SANE package. And I rewrote the SANE package to use these new instructions and to produce a very fast software floating-point package.

Hsu: That was native to the 601?

Davidian: Well, it was actually-- it was 68000 code and the emulator had these additional instructions. So what it was-- and I actually…

Hsu: Oh, so you're still doing floating-point in software.

Davidian: It was still in software but there was, you know, I had used-- the 68000 had 8 data registers and 6 of them were-- 3 of them were used for a floating-point number in my instruction set. And so you could have two-- using six registers you had two floating-point numbers so you could do operations like add, subtract and things like that. So the values would actually live in these 68000 registers. This instruction would say, do a floating-point add between these two floating-point values. And so there would be this, you know, this PowerPC code that would implement that instruction that would do it faster. I also had…

Hsu: But you're using integer PowerPC [instructions].

Davidian: Yeah, it's all integer. The PowerPC floating-point was pretty much useless for doing Motorola floating-point. <laughs> And, you know, before implementing these instructions, I actually implemented the functionality of them in 68000 code, so I could actually test it out and verify that I-- do all the verification that I'm getting correct answers and things like that.

Hsu: So you're doing software floating-point using 68000 integer instructions that's being accelerated by PowerPC integer instructions.

Davidian: Yes. By special 68000 instructions that are implemented in PowerPC instructions. So it would be sort of like if, you know, that these were extensions to the 68000 instruction set that did these very complex, you know, floating-point operations, fairly lengthy. So I had considered, also, then doing a software emulation of the floating-point-- the Motorola floating-point instructions using the same instructions that I implemented for the SANE package. And someone else actually wound up doing something like that. I think I forgot what they called their product.

Hsu: It was a third party.

Davidian: It was a third party. They wound up doing it. It was just so-- there were some applications that actually required the floating-point unit. They didn't test to see whether it existed. And this was some way—this was a way to get them to run. It didn't get them to run quick. It got them to run. So, you know, but I always wanted to, actually, you know, sort of complete the emulation either through a software package to implement the floating-point instructions or in the emulator itself. But it was something that development wise would have taken too long and, also, trying to exactly match the results of the 68881-68882 would've been very difficult. And in the end, the transition happened very quickly to PowerPC.

Hsu: Right. So the third-party developers just ended up rewriting their…

Davidian: You know, that was something where they could get a huge improvement. You know, if they were floating-point bound in 68K code and they didn't actually need the 80-bit precision and could live with the 64-bit doubles in PowerPC, they could get huge, you know, huge performance improvements.

Hsu: Right. So it was a natural for them to rewrite that part natively.

Davidian: Yeah. Yeah it was.

Hsu: I mean you mentioned the lack of tools initially internally for Apple. That was a problem with for third-party developers, as well.

Davidian: Yeah. Metrowerks.

<group laughter>

Hsu: I mean Apple had not really thought that story through for its third-party developers.

Davidian: I don't know if it's not thinking it through. It was, you know, this was a sudden-- this thing happened fairly quickly because…

Hsu: Well, it was still three years between the announcement and the machine shipping, right?

Davidian: No. It was less than two. No. It was less than three. It was probably about two-and-a-half years. I think it was like around October '91 was the announcement and March '94.

Hsu: Yeah, two-and-a-half.

Davidian: Two-and-a-half. And, you know, it was in-- I think was probably June of '93 at WWDC in '93, we had a room set up with all with 6100 prototypes running, you know, running the emulator for developers to test out on.

Hsu: Oh wow. So they have less than-- they have six to nine months to…

Davidian: Yeah. Well, they were seeded but this was, actually, for them to test their 68000 apps.

Hsu: Oh, okay, right. Compatibility.

Davidian: Yeah, to try to break it. I was pretty confident that the emulation was solid, but there was this other little thing that we were being forced into, at the time, which was, if you go back to when I was involved in the ci and the fx, then the universal ROM there, that version of the universal ROM kept getting modified for a number of 030 and 040 products beyond that. So the basic Toolbox in that was not changing. A lot of that code was unchanged. So all the bugs were worked out. Well, they were-- for the Quadra 840AV they were starting it over and with a newly built ROM from new sources and incorporating a lot of the stuff from System 7. So I was more worried about compatibility from that ROM-- compatibility issues for apps from that ROM than I was from the emulator itself.

Hsu: So they were going to be using that new ROM for the Power Macs than the older one.

Davidian: Yeah. So there were potential compatibility issues just from that ROM development.

Hsu: Okay. But they could have just used the existing shipping Quadra 840AV to test that part.

Davidian: Yeah, but during the development I don't remember what date the 840AV shipped.

Hsu: Oh, that hadn't shipped yet either.

Davidian: Yeah that was going to-- and while we were developing the 6100 it was never-- at Apple it's never clear that a product in development is going to ship. <laughs> In my head I kept saying, you know, I forget what the 840AV was called internally. Cyclone, I think it was. You know, it was "Cyclone must ship" was the thing that always kept going through my mind was-- because they had to take the hit for the ROM change. <laughs>

Hsu: So they were the first product to come out with the new ROM?

Davidian: Yeah. It was called the Super Mario ROM. I don't know why it was called that but that was what it was called. So they were the first Super Mario ROM product. And so yeah, I wanted-- they had to ship and then, you know, System 7 was out before us and it forced us to bump RAM to 8 megabytes from 4.

Hsu: That would be the minimum requirement for running PowerPC was 8 megabytes?

Davidian: At that time, yeah. It's, you know, we were originally up and running in 4 megabytes, but it just—you know to have room for apps it was-- <laughs>. I'm trying to think of other things that went on during that time period was the adding support for virtual memory. That was, again, it didn't emulate

Motorola's MMU just because if you do that, if you're going to emulate something you have to emulate all the features of it and not just what's used. So there were some additional instructions added that made calls to the nanokernel for MMU management, things that were needed just for Apple's implementation of VM in System 7. And Phil Koch, actually, wrote that code and I incorporated it into the nanokernel. He was actually my manager at the time.

Hsu: I want to ask about the nanokernel. So I know that's a term that Apple was using.

Davidian: Yeah.

Hsu: What did it really mean? Because it was not a kernel in the traditional UNIX sense, right?

Davidian: Yeah, it's-- Phil Koch actually gave it that name. I don't remember what I was calling it. I don't remember what I called it in the 88000 emulator when I wrote that but Phil Koch is the one who started calling it the nanokernel in the PowerPC version. Basically, the purpose of the nanokernel was to do anything that had to be done in supervisor mode because the emulator ran in PowerPC user mode, even though it emulated 68000 supervisor mode, it ran in PowerPC [user mode]. So it did all the exception handling. It did anything that needed privileged instructions like all the MMU mapping, interrupt handling, and it also did context switching between the emulator context and the native context. So that's—mixed mode actually would go through that. And part of the reason why the emulator ran in user mode was during development of this there were two strategies that we were thinking of. One we called v0. One we called v1. And we, also, referred to it as, you know, who's on top? Which is does the operating system run on top of the emulator? Or does the emulator run on top of an operating system? <laughs> And v0 is what actually shipped which was the emulator really was at the bottom and the operating system ran on top of the emulator. So the nanokernel was below the emulator. But one of the biggest flaws that I felt in the original Macintosh was that they ran applications in supervisor mode on the original 68K Mac. And that was a decision that was made that had ramifications for decades-- well, maybe not decades.

Hsu: So all these CPUs had a separation between user and supervisor mode?

Davidian: Yeah. The 68000 had a user mode. It would’ve been much easier-- solving problems later on would have been much better if the applications ran in user mode and there was a user stack and a supervisor stack. And when exceptions happened then you ran on the supervisor stack and switched into supervisor mode. And it would’ve been easier to do things like MultiFinder and application switching.

Hsu: But then because all of the OS and third-party applications ran in supervisor mode, they could basically…

Davidian: They could do anything. Yeah. So v1 was sort of, you know, what evolved into Copland. There was NuKernel and the emulator would run on top of NuKernel and so it was running on top of an operating system. So I didn't know what-- you know, v1 was this unknown and so I didn't want applications to be taking advantage of something that we wouldn't-- that would be difficult to support in the future and we knew this future was going to be coming. It was always v1 was like six months behind

v0. And during development, you know, we [were] trying to make a lot of decisions about should we do v0? Should the initial product be v0 or v1? And, of course, the six months separation turned into a lot longer than-- you know.

Hsu: Right. So, okay wait. So v0 and v1 are-- this is referring to the kernel, the nanokernel?

Davidian: This was whole operating system.

Hsu: The entire operating-- okay.

Davidian: So v1 was a native operating system.

Hsu: Oh, okay. So v1 was the code name for a full PowerPC native operating system.

Davidian: With-- a native operating system with emulation on top.

Hsu: With emulation on top of it.

Davidian: And whereas v0 was an emulated operating system with native acceleration and applications.

Hsu: Okay. And v0 is really what ended up…

Davidian: v0 is what shipped.

Hsu: Shipped. Okay. So v1 was supposed to be what, Copland would eventually become?

Davidian: v1 eventually-- v1 continued on to become Copland.

Hsu: Okay. Okay. Which-- it's like Pink all over again.

Davidian: Yeah. Yes.

<group laughter>

Davidian: So Steve Jasik used to bug me, periodically, because he wanted to get into supervisor mode for his debugger because he wanted to use the single step instruction trace exception. And I kept telling him no because I didn't want anyone in supervisor mode because we don't know what-- I didn't want to make it hard for anyone for development in the future, operating system development in the future. So finally, you know, after bugging me I started looking at just how long things would take to actually go through-- you know to take all the code involved and just the processor timing for taking these exceptions and it was fairly lengthy. So I sat down with Steve and had this conversation about look how long-- how many clocks this is going to take? What you could do is you could actually emulate the PowerPC instruction set and interpret PowerPC instead of taking single-- instead of having the processor take

single step exceptions for you. And in talking to him more he started thinking about other things he could do there, where he could then like set breakpoints on data fetches, and memory location modifications and things like that. And he started thinking, well, you know, that was a much better way and it, also, solved this not giving-- not letting him, you know, run-in supervisor mode. So, yeah, that was a good compromise and I was glad that something actually positive came out of having this restriction that he couldn't run supervisor mode. <laughs>

Hsu: So I think this book, that I brought here, mentioned that some of your colleagues-- so you mentioned Bob Hollyer, Keith Cox, John Fitch, Carl Hewitt.

Davidian: Yeah, Carl Hewitt.

Hsu: Is this the same Carl Hewitt that did Actors?

Davidian: I don't know what Actors is.

Hsu: Okay. It's sort of a software paradigm and kind of like object-oriented but a different model.

Davidian: I don't think so. We all went on to work together again at Power Computing. Well, except for Keith Cox; he stayed Apple. And then after Power Computing Carl was one of the founders of Kerbango, along with Jon Fitch, and who else was it? It's slipping my mind right now.[1]

Hsu: And Kerbango did?

Davidian: An internet radio. And they were eventually bought by 3Com and shortly before they were going to ship 3Com killed it. But that's skipping ahead a bit.

Hsu: Let's see, we talked about the Hawaii demo. We talked about the WWDC demo.

Davidian: So, yeah, there's a big thing that happened in between all the-- year around this time, which was around, I think, it was August of '93, another clean room came up.

Hsu: Oh, really?

Davidian: Yeah. <laughs>

Hsu: This is completely separate from Apple. This is, once again, the Intel?

Davidian: Yeah. This was, yeah, an offer I couldn't refuse. <laughs> So AMD, the 486, and they had been talking to me for a long time. And never-- it's like there would be talking and then silence for long periods of time. So I never knew if this was going to happen or not but then it became real. So I

[1] [Interviewee's note] The third person was Jim Gable.

remember-- I knew that it was going to be something where I would have to quit Apple, at least, in my mind it was because it was not-- yeah, I couldn't really be doing both. I needed, you know, it was going to be a full-time effort. So I was prepared to resign. I called everyone on the PDM team into a conference room and said, you know, we need to talk and they thought that I had discovered some really bad hardware bug. <laughs> And I said no I'm leaving, you know, but I'm not leaving the area. And if anything serious comes up, you know, you can always call me but there's an offer I can't refuse. I went and talked to Dave Nagel and…

Hsu: He was in charge of software at this point?

Davidian: Yeah.

Hsu: He had come from ATG?

Davidian: Yes. And sort of the first question out of his mouth was, “How much?” <laughs> And I said, as a shareholder I would object. <laughs> But I actually, you know, I gave him a number and, yeah, he thought that was too much, also. <laughs> And it, actually, was less than what I would’ve gotten, you know, what I was going to get paid.

Hsu: For the clean room.

Davidian: For the clean room, but it was-- yeah. It was something that probably would’ve made me change my mind. But he wanted me to stay on somehow. He wanted me to, you know, somehow be available on an official basis. And so we tried to make that that happen. Phil Koch was my manager at the time. So there was a meeting set up with human resources. And so the human resources person wanted to understand what was going on and I explain this and, you know, that I have this other job and I want to leave. But we want to talk about it, a part-time arrangement where I stay on. And she said, “Well, you know, I don't see why Apple should do this.” And I said great. I'm done. I’m leaving. And Phil said, “No, no, no. Dave Nagel wants this to happen.” So there was negotiations going back and forth with, you know, my lawyers and Apple lawyers. But in the meantime everything progressed as if I was going to leave. We had a going away party and it was actually after the going away party that all the paperwork got signed. And I stayed on as-- I became a part-time employee, so I didn't have benefits anymore and didn't really have an office anymore. But I was available and I was being paid. So then I was doing the clean room project for about it was probably another year, close to a year. I did go to the introduction of the initial Power Macs in New York in March, but I was still, you know, part-time. And then I came back and, you know, things were a lot different then. It was…

Hsu: Could you talk a little more about the clean room project?

Davidian: Yeah, a little bit. It never got used. It was another one of these insurance policies. In this case, you know, again it was AMD was a second source. They had rights to the 386 and 486, I believe. But just like with NEC, Intel said the microcode is not included. And they worked out some agreement on the 386.

And from what I've read is that for the 486 the team that was working on the new microcode for that was given the 386 microcode to start with and they were referencing it which was not clean. <laughs>

Hsu: But because they had signed an agreement for it, they had the code…

Davidian: To use in the 386.

Hsu: Oh, but not to use in the 486.

<group laughter>

Davidian: And so not only did they not have microcode for the 486, they had a contaminated development team. So they approached me. Initially, it was actually surprising that I could-- that it would be even possible for me to do this because I had worked on-- you know, after the NEC case, I was given this exhibit that had the comparisons of the microcodes and had Intel's microcode in it.

Hsu: Oh, right, so you were now also contaminated.

Davidian: Yes and no, because one of the rulings in the NEC case was that Intel did not have a valid copyright. <laughs> So, yes, I was exposed to something that was not copyrighted. <laughs> So it was actually possible for me to do this. And this was a much different and much more complex project than the NEC clean room, and it took many months, and they invested a lot in the development process. I had these huge logic simulators that could actually simulate the processor and my microcode. They actually produced a 486 with a writable control store. I could actually boot Windows <laughs> on it with my microcode. But I eventually finished it, and sometime after that Intel and NEC [AMD] came to an agreement. <laughs> But then Intel—I'm sorry, AMD after that, that's when they started really producing their own x86-compatible, starting the K5.

Hsu: Right. So then they wouldn't need to-- yeah.

Davidian: Yeah, so that's when they split. That's when they stopped being a second source and really doing their own development.

Hsu: It just seems so weird that Intel would be playing hardball with its own second source. It doesn't make sense to me why they would license the 386 microcode but not the 486 microcode. Of course AMD is going to use the same development team on both, right?

Davidian: Yeah.

Hsu: What does Intel expect them to do?

Davidian: I don't know. The other thing about the fact that the negotiations went on so long and there were these gaps, and I started thinking about this many years later when there was another clean room

that I was involved in, which I'm not going to say much about. That one was another one where there was a very long negotiation with lots of silence, and that one, they finally decided to go ahead with it. They delivered all the equipment and specs, and the day after that they told me to stop, because they reached an agreement, an out-of-court agreement. And my feeling now looking back on all this is what they were really doing was just the threat that they have the ability to do a clean room--

Hsu: Gets them back--

Davidian: It was a negotiation tactic with who they're having a lawsuit with, so it was how to not pay me but use me as a threat <laughs>.

Hsu: So you became the clean room guy essentially after.

Davidian: Yeah, well, with the NEC, I think it was the first clean room in the whole process that actually made it to court. <laughs> And it was successful. <laughs>

Hsu: So they just kept going back to you.

Davidian: Yeah.

Hsu: You were the guy.

Davidian: Well, and I think in some cases it was also going back to the lawyers.

Hsu: Oh, the same lawyers, right, who would call you.

Davidian: Yeah. It was a small world. <laughs> So after that was done, I did eventually go back to Apple, and it was very different then, 'cause everything was focused on Copland.

Hsu: Oh, okay, so immediately after you got back it was all hands on deck for Copland.

Davidian: Well, while I was gone.

Hsu: Oh, right, while you were gone, yeah.

Davidian: In March '94, <laughs> V0 shipped. <laughs>

Hsu: So now all of System Software, yeah, is focused on Copland at this point.

Davidian: Yeah, which was going to be System 8.

Hsu: Right. What was your whole view of that given you've lived through all these similar failed projects before? Where did Copland go wrong?

Davidian: I don't really know. I don't really have an opinion. Actually, when I got back, I wasn't really directly working on Copland. The things I was working on then was the next-- the next PowerPC products were coming out using the 603 and 604 processors, and for those, one thing that was needed there was-- if you remember back, the 601 implemented both the POWER and the PowerPC instruction sets, and we used IBM's tools for that, and the IBM XL C compiler produced POWER instructions. So anything that was compiled with those compilers would actually have POWER instructions in them that are not implemented in the 603 and 604 processors, so there was this POWER instruction emulation that needed <laughs> to be done. It was in the nanokernel, so that was another thing that I worked on <laughs>.

Hsu: Wait, so the Mac OS was actually emulating POWER on top of PowerPC--

Davidian: Yeah, there were some POWER instructions that the compiler had generated and that they'd taken illegal instruction exception on the 603 and 604 <laughs> if those app-- the developers were encouraged to recompile once Apple had tools, a real PowerPC C compiler <laughs> and Metrowerks also.

Hsu: That was Mr. C, right? The Apple one?

Davidian: Yeah, I think Mr. C. I think there was another C compiler also, but I don't recall. I don't think they were using GCC at that time.

Hsu: No, they weren't. Yeah.

Davidian: But I think a lot of the developers were using Metrowerks instead of Apple's development tools, and Metrowerks, they generated pure PowerPC code. No, it was just all because of that rush into development and the need to use the RS/6000 as the development environment.

Hsu: Yeah, so once again, while the rest of System Software is working on Copland, you're working on the support for the new hardware, which was a similar arrangement as before.

Davidian: Yeah.

Hsu: When the rest of the team's working on System 7 and you're working on--

Davidian: Yeah, and also during that time while I was gone and when I came back was when Eric Traut started working on-- he started looking at the emulator, and he started implementing the DR emulator, the dynamic recompilation, which was similar to-- Connectix did Speed Doubler, and this was sort of Apple's implementation of the same sort of performance optimization. And this is something that really couldn't have been done in the initial products, mainly because when we were thinking 4 megabyte of RAM, there was no extra RAM, <laughs> you couldn't just grab another half meg or a meg of RAM to use <laughs> as this translated instruction cache, and also just development time and compatibility.

Hsu: But now Macs were, what, shipping with 32 megabytes? How much RAM were they--

Davidian: I don't think they were actually shipping with 32 at that time.

Hsu: 16?

Davidian: Well, 8 was enough.

Hsu: 8 was enough.

Davidian: Yeah, and it was upgradeable, so they would ship with 8, and most people would upgrade from there. And the DR emulator was optional. You could turn it on and off. It depended on the initial emulator, too. It was an acceleration on top of this, so a lot of the hard stuff actually always fell back to the interpretive emulator that I wrote.

Hsu: Yeah, okay. I guess this sort of-- because Copland never ended up shipping and so V0 stuck on until it was replaced by OS X essentially. That's why there was always still Mixed Mode and still a little bit of 68K code in the operating system, even Mac OS 9.

Davidian: Yeah. Well, it was in the applications. <laughs> Applications, there could've been Mixed Mode in applications, and it was, yeah, in the system. And, yeah, it still looked like the OS-- there was still support for 68K apps. <laughs> So the emulator lived on probably until I think it was Mac OS [X] 10.4.

Hsu: Right, Tiger.

Davidian: Tiger.

Hsu: Right, because that was the last version that had the Blue Box.

Davidian: Yeah, Classic.

Hsu: The Classic, yeah, emulation in it, yeah.

Davidian: And it actually 10.4 PowerPC. There was actually 10.4 Intel <laughs>.

Hsu: Intel, right. But that I don't think included the [Classic] Mac [emulator], yeah.

Davidian: No, 'cause it didn't have the PowerPC processor. <laughs>

Hsu: Right, exactly. Yeah, that stuff lasted quite a while. A decade, yeah.

Davidian: Yeah, it lived on a long time.

Hsu: Let's see. So you were back working on-- let me see here. Where were we?

Davidian: The POWER instruction emulation.

Hsu: Yeah, right. So you were still at Apple until ’95?

Davidian: September ’95.

Hsu: September ’95, okay. And so you were continuing to work on the stuff for the 603 and 604 until that point?

Davidian: Yeah, and I’m trying to think what else I was working on then. <sighs> Yeah, the 603 and 604 processors. I would go down to Somerset periodically. I was still working closely with some of the hardware teams and discussions about future PowerPC processors.

Hsu: So what was on the road map at that point?

Davidian: Well, the 603 and 604, and then 603e. I guess there was a 604e I think also. I’m trying to remember.

Hsu: There was a 620 server chip.

Davidian: There was a 620, but Apple never used it. I wrote the nanokernels thinking that there might-- the 64-bit architecture was defined in the initial PowerPC architecture documents, so I wrote the nanokernel assuming that there might someday be a 64-bit processor being able to support it, but Apple never used one. Well, I guess the G5 <laughs>, but that was a long time.

Hsu: That was a lot later, yeah.

Davidian: Yeah <laughs>.

Hsu: And not the classic Mac OS.

Davidian: Yeah, yeah, that was, yeah, the G5. Well, the G5 could run 10.4 I think. Yeah, so it could actually run--

Hsu: Oh, right, it could run the emulator, yeah.

Davidian: But not the nanokernel. <laughs>

Hsu: Right, okay.

Davidian: Yeah, and if I think back, I think some of the initial looking into multiprocessor support was going on then. I think there might’ve been a-- I think one of the 7500, 8500 series, there might’ve been a dual processor.

Hsu: Oh, yeah, there was the SMP, the 9500 SMP I think, yeah.

Davidian: Yeah, 9500 MP I think it was called.

Hsu: Or MP, yeah, yeah, yeah, yeah, yeah. I think I remember that. Yeah. Right. Did you ever work with René Vega?

Davidian: So I know René. After I left Apple-- well, actually while I was still at Apple, we interviewed people to help me out and work in Phil Koch's group, and there's one guy, Bob Abeles. He was a microprogrammer. He had worked for Hitachi and was microcode for IBM 370-compatible processors, so he was very familiar with the IBM 360, 370 architecture and also with microprogramming. And through the years he had worked with some other people at I think-- I don't know where they all crossed paths, but they were all working on the same type of stuff, and they'd worked at Amdahl and Fujitsu and Hitachi, and René Vega was one of them, and Bill Angel was the other. So René and Bill joined Bob Abeles at Apple I believe after I had left, but I had gotten to know them just when I would go back to visit. Yeah, René and Bill, they were very interested—they actually disassembled the emulator and nanokernel <laughs> before they had gotten hired, yeah, and they wanted to know how it worked. <laughs> So they were really interested in getting involved in the group where that was being done.

Hsu: Okay, 'cause I worked with René very briefly. I joined Apple in '99, October, on the file system team for OS 9. And then that effort went away, and I was briefly assigned to work with René, and he was at the time-- in 2000 he was in charge of the nanokernel and of the MP APIs, so that was why--

Davidian: Yeah, with Jim Murphy.

Hsu: Yeah exactly, yeah, Jim, yeah. Yeah, I worked with him, too. Yeah, okay, cool. A little bit of overlap there. So I guess should we talk about you finally leaving Apple?

Davidian: Yeah. Apple had started licensing, and the first licensee was Power Computing. Although they were headquartered in Texas, they had R&D in Cupertino, and Carl Hewitt had already left. John Fitch had followed. Bob Hollyer had followed. And I had worked with John and Bob since the day I started at Apple <laughs> and Carl Hewitt from the day he started <laughs> at Apple, so <laughs> we knew each other, and so it was just time to move on.

Hsu: Right, so they convinced you to join them.

Davidian: Yeah. Well, the big convincing thing was that it was just down the street. <laughs> It wasn't in Texas, and it was working with people that I enjoyed working with and not having to work on Copland. <laughs>

Hsu: Which you, again, felt like it would go the way of Pink?

Davidian: Yeah, yes, and I wasn't really in it from the beginning, and it was just sort of this thing that felt out of place, and I sort of felt that I had accomplished what I had wanted there, and it was time to move on.

Hsu: And you said they were located across the street--

Davidian: Well, no, down the street.

Hsu: Down the street.

Davidian: They were on Bubb Road.

Hsu: Oh, okay. That's still very close.

Davidian: I think the building, it may have at one time--

Hsu: Been an Apple--

Davidian: Been an Apple building.

Hsu: Right. Okay. I hadn't had a chance to ask about your thoughts of the management, but Sculley and Spindler as CEOs, what were your takes on--

Davidian: I liked Sculley. He came into the lab once to see a demonstration of PDM when it was in development. And Spindler, we had a meeting with him once, and I don't think he really understood what we were doing. He sort of went off onto this tangent about how much money Hewlett Packard makes on selling toner cartridges. <laughs> I was kind of sorry to see Sculley go. I was very sorry to see Gassée go. There was also other management that was brought in during the PowerPC development. A lot came from DEC.

Hsu: Oh, really?

Davidian: Yeah, well, Roger Heinen was sort of the first.

Hsu: Oh, really?

Davidian: Yeah, and I liked Roger. He understood what it was I did, and he was the one that needed to be convinced to let me go. I was in his organization during System 7 time, but I was working on RISC. And he had worked on the transition to Alpha at DEC. But then they started bringing in more people from DEC, and we were always told, well, this was the person that was responsible for the transition from VAX to Alpha, and then we find out, well, all they do is they hire another person that, oh, that was the person-- <laughs> so it was like we were getting the entire DEC management chain but no actual people that actually did anything <laughs>.

Hsu: Oh, 'cause they were all managers, so they all took credit for that transition.

Davidian: Yeah, yeah, although there were other people that came from DEC that were really big assets. Rich Witek was one of the Alpha architects, and he was working with Somerset. I think there were other people that came in from the Alpha proj-- I think some of the people that worked on Alpha also worked on the low-power ARM processor that was used in Newton.

Hsu: Oh, right, the StrongARM.

Davidian: Yeah, StrongARM, yes.

Hsu: Yeah, yeah, right.

Davidian: Which years later wound up at Intel. <laughs>

Hsu: Right, ironically. Yeah, actually, Allen Baum was I think one of those people. I think.

Davidian: Was he at DEC?

Hsu: Oh, he wasn't at DEC, he--.

Davidian: He was at Apple.

Hsu: He was at Apple, but then he had gotten involved in the StrongARM project.

Davidian: Yeah, I think he might've actually gone from Apple to DEC to Compaq.

Hsu: To Compaq to Intel. Yeah. Right, just stayed with that team the whole time.

Davidian: Yeah. <laughs> All in the same office <laughs>.

Hsu: So you knew him at Apple?

Davidian: Yeah, I knew Allen at Apple.

Hsu: He's a volunteer with us actually. We see him quite often.

Davidian: Oh, okay. Yeah, I've bumped into him a few times in recent years, but yeah.

Hsu: At the time the PowerPC was shipping, the head of the Mac division was Ian Diery, who's an Australian?

Davidian: Boy, I always thought of Ian as a marketing guy, but, yeah, the name sounds famil-- yeah, there were a lot of changes <laughs>. I don't know if I ever met Ian, but I didn't really have any opinions. And I believe I had the emulator up and running while it was still under Sculley. I know there was one sort of big management transition that was happening, and I knew it was going to be coming, and I actually wanted to make sure that I had the OS up and booting before then, <laughs> just so I could point out that, no, that new person actually had nothing to do with <laughs>--

Hsu: Right, to do with this? They don't get to create [claim] credit for this thing.

Davidian: Yeah, 'cause that seems to happen at a lot of companies where the people come in at the end to take the credit. <laughs>

Hsu: Yeah. Well, Spindler, he is in charge when the machines finally ship, so he does take the credit for that I guess.

Davidian: Yeah. I don't remember what the date was that the Sculley-Spindler transition happened, but--

Hsu: I think it was sometime in '93.

Davidian: Yeah, 'cause June '93 we had developers at WWDC running their apps <laughs>, testing them <laughs> on the emulator. <laughs>

Hsu: Let's see. So you joined Power Computing in September '95.

Davidian: '95.

Hsu: Okay, and Spindler's still with the company at that point?

Davidian: Yeah.

Hsu: Okay.

Davidian: Yes, yes.

Hsu: But is it already looking bad at that point at Apple?

Davidian: I don't think it was looking bad, because actually I think things actually looked good once the Power Macs came out, because they had performance; the raw performance of them was better than Intel. And the transition happened very quickly. I think the result of this was actually things were looking better, and then they started licensing.

Hsu: What did you think of the decision to finally do the licensing? Apple had been against licensing for so long. Why now? Why was this the right time to finally go into licensing?

Davidian: I think the thoughts about the licensing was it was actually going-- it wasn't really the whole Mac OS that they-- they were thinking about CHRP, and these would be boxes that could run multiple operating systems, not just the Mac. And so they weren't necessarily competing with Apple. It was another way for the Mac OS to be out there, I think is what some of the thinking was at the time. That's not necessarily how things turned out, but <laughs> I think that was some of the thoughts, 'cause the future was looking towards this open architecture. There was going to be a version of NT that ran on it and, I don't know, OS/2. There might've been OS/2 as another operating system.

Hsu: Okay, so then Power Computing, what was that like?

Davidian: In some ways it was a lot of fun, because it was working with people that were like family, and we all knew each other, and it was a lot of just doing stuff over again that we had done before. The initial products were based on the first generation Power Macs. Then we were getting into the newer ones, which were based on the PCI Macs.

Hsu: Okay, right. And the 603 and 604 chips, yeah.

Davidian: Yeah. And then processor development was continuing, so we were actually putting processors that Apple hadn't shipped initially with these designs, we were putting those processors into those motherboards.

Hsu: Oh, okay, so you were using existing Apple spec motherboards that-- licensed?

<overlapping conversation>

Davidian: Yeah, in some sense you could think of it as the processor was on a-- well, at least in the Power Computing designs, the processor was on a separate plugin card, so we were essentially developing new plugin cards for the basic motherboard with different processors. And so there was the 603, then the 603e. Then there was the 603ev, and that's where things got interesting. The 604 and the 603 designs, as they matured and evolved into the newer processors, Apple was looking at the 604 line here for their desktops. And we had been looking at the 603ev, and then the follow-on to that was the 750, also known as Arthur was the code name. And when I looked at the Book Four for Arthur, I said "This is going to kill the 604." <laughs> But Apple didn't seem to recognize that.

Hsu: So they weren't planning to move to the--

Davidian: They were focusing on the follow-- I think Mach 5 was the follow-on to the 604e or 604ev.

Hsu: Mach 5, okay.

Davidian: I think that was the processor code name.

Hsu: I'm not familiar with that one, yeah, 'cause the 750 became the G3, right? It was marketed as the G3.

Davidian: Yeah, it was known as the G3, and in my mind that was the killer, because I think it was in the 603ev they had added a little bit of a second arithmetic unit that could do just a few instructions, and in the 750 they added a complete second arithmetic unit, a second integer unit, and that was the big advantage that the 604 had was it had two integer units. But the big thing that Arthur did was they put in this tag store on chip to support an external L2 cache. That was the killer, and that's where it became clear that Arthur was the-- it was low power, low cost and high performance, <laughs> potentially high performance if you add on an L2 cache. So that's where Power Computing was focusing, and Apple seemed to be ignoring it. Either ignoring it or thinking of it only for portables, because I think with Apple the 603 was portables and the low end, the Performa-type <laughs> Macs, and so I think they probably would've continued on, whereas the high end was focused on the 604 and its follow-ons.

Hsu: What was the culture like at Power Computing?

Davidian: Especially in Cupertino it was like a startup, because we were separate from the rest of the company, which was in Texas, that's where all the manufacturing was and everything.

Hsu: So all the R&D was in Cupertino?

Davidian: Yeah.

Hsu: Okay. And that was deliberate, because you could be near Apple.

Davidian: Yeah, and we all lived there <laughs>.

Hsu: Right.

Davidian: Lot easier to recruit <laughs>.

Hsu: And all the recruiting had been done from out of Apple as well? Most of the engineers?

Davidian: Well, Carl Hewitt I think was the first to go, and I think he might've even gone before Apple had actually given a license to Power Computing, so I think he had taken the most risk. By the time I went there, they had already produced I think it was called the Power 100. It was the 601-based Mac.

Hsu: And the licensing was so Power Computing licensed entire motherboard designs in addition to the operating system?

Davidian: That's the way the licensing worked. It was the hardware design including the ROMs and the operating system license, too.

Hsu: Right, 'cause there was no CHRP yet at that point.

Davidian: CHRP was being defined, but Apple hadn't produced anything with CHRP. There was development going on, and the Developer Technical Support had seeded some CHRP boxes, and we had one. And they had a version of the Mac OS that ran on it if I remember correctly, but we never actually produced one. We were looking at it at the time.

Hsu: Did any of the licensees ever make a CHRP machine?

Davidian: As far as I know, nobody ever made <laughs> a CHRP machine <laughs>.

Hsu: Yeah. They were all essentially Mac clones.

Davidian: No, all the actual licensees made-- everyone did the same thing. They were licensed, like the PCI Mac design or the NuBus Mac design. Power Computing might've been the only one that actually did the original NuBus ones. I think everyone else might've gotten in later and then done PCI Macs, but I'm not too clear on that.

Hsu: Let's see. So once again you were working on the low-level software for the clone?

Davidian: Yeah, just also a lot of assisting the hardware team in board bring-up, and sometimes I would write sort of custom applications to test something out for them. I worked with Bob Hollyer pretty closely. One day he wanted an app to make the CPU hot and another one to make it cold, <laughs> 'cause he was doing thermal testing. <laughs> And I remember one day in the lab that he had the thermal probes glued to the processor or the heat sink and the box over the machine, and we're all around it watching the temperature readout and running my hot application, and all of a sudden the temperature started dropping. And we wiggled the mouse, and the system was still up, <laughs> and we're going "What's going on?" So we took the box off and looked at it, and I had actually melted the glue that was holding the thermal probe <laughs> to the processor, and then the probe fell off <laughs>. But also that app I wrote, not only did I melt the glue, but I actually got the current draw of the processor to be higher than what the data book said it could be. <laughs> So, yeah, they were a little surprised when it exceeded the data book <laughs>.

Hsu: So at some point, basically Steve Jobs returns to Apple in '97, and then after that Apple decides it doesn't want to keep the licensees anymore, and it buys Power Computing out.

Davidian: Yes. It seemed a bit more long drawn out than that, but yeah. <laughs> Well, the purchase of NeXT was sort of when it started.

Hsu: Yeah, so the deal was made in December of 1996, but it was finalized in February '97. But Steve had already gone on stage at MacWorld in January of '97 along with Amelio to sort of announce the merger, so it was public knowledge from January.

Davidian: Yeah, no, and I remember at WWDC I think it was the fireside chat with Steve Jobs.

Hsu: Right, which was, what, June or something?

Davidian: Yeah, I don't remember when it was, but I do remember that's when he referred to the licensees as leeches. <laughs>

Hsu: Not a good signal.

Davidian: No. <laughs> And it seemed like as negotiations were-- they pulled the whole thing with the-- System 8 was supposed to be Copland, and the license agreement included System 8, but then they changed System 8 to be something that looked a whole lot like System 7. <laughs>

Hsu: So then the license agreement did not include the new version of--

Davidian: So, well, Copland wasn't called System 8 anymore, so what we assumed was Copland was no longer part of the license. But from what I've heard of the license negotiations, Steve Jobs wouldn't license the Mac OS at any price. I think there were even-- said "What if we buy an iMac, <laughs> an entire iMac <laughs> and throw it away and just use the license?" <laughs> That wasn't good enough. <laughs> No, I think what was really going on was that Power Computing and all the other licensees, since their volume was actually fairly low compared to Apple, they could use the newer processors sooner. And they would also price lower than Apple, so they'd have price and performance better than Apple, but they couldn't do the volume that Apple needed to do. Apple couldn't match it, because they couldn't get the chips <laughs> in that volume, and it was putting pressure on their price for their high-ends, which was where they made all their money. So I think that was the thing Steve Jobs was looking at where the only solution to that was to kill the licensees. The licensees weren't adding value to Apple. They were competing with Apple in a way that Apple couldn't compete. And besides that, we had focused on the G3, Arthur, and Apple hadn't. <laughs>

Hsu: Right, at least not until Steve came back and it turned around.

Davidian: Yeah, that was very late.

Hsu: Yeah, that was kind of a crash course I guess.

Davidian: So that was really the end, I think what ended it. And at least the way that Power Computing got killed, looking back, was not too bad for the people that worked there.

Hsu: No, 'cause it got bought out.

Davidian: The intellectual property got bought out and paid for with Apple stock, so in 1998 I was forced to buy a lot of Apple stock.

Hsu: Okay, by Steve.

Davidian: Steve Jobs forced me to buy a lot of Apple stock in 1998, which I immediately sold <laughs>.

Hsu: Yeah, at that point in time, nobody would've known, right?

Davidian: Yeah, no, I sold at least enough to get back what I had to pay for it, 'cause I had to buy Power Computer stock options, which, if a deal didn't go through, would be worthless with real money. And that got converted into Apple stock, which was still falling. <laughs> So, yeah, but I didn't sell it all.

Hsu: So at the end of the day, what were your feelings about that whole thing ending? Did you feel like it was-- I guess from the competing perspective as an employee but also from the perspective of Apple, was it the right decision?

Davidian: Well, from my personal feelings, I felt relieved <laughs> somewhat, because the whole CHRP thing, I didn't have a good feeling about that being successful, and that seemed like just trying to develop this product that was actually gonna run this third-party operating system that we actually didn't have the source code for and just trying to bring things up, and it all in theory should work, but it was going to be a challenging thing, and I also wasn't seeing what using CHRP really brought to the customer. <laughs> So I was kind of glad that the CHRP stuff didn't have to be finished <laughs>.

Hsu: Were you actually working on that?

Davidian: I was working on some of that, yes, on the Open Firmware stuff for that.

Hsu: Right.

Davidian: So I was glad that was over. We were also working on some multiprocessor products. But for Apple, I think Steve Jobs brought a lot of focus to Apple. He also would make decisions. It seemed working at Apple, anybody could say no, but nobody could say yes, <laughs> and Steve could say yes, and Steve could say no. <laughs>

Hsu: Right, yeah. The culture that you described is just night and day from after Steve came back. You're describing a culture where there's all these competing teams, and things get killed, but then they still sort of live on in zombie form in a new project, and Gassée is letting you do whatever you want.

Davidian: Yeah. No, it was a fun environment. <laughs>

Hsu: Right, right, really fun, yeah. But it would take four years to ship a product almost, right?

Davidian: Oh, yeah, it was just amazing how long products were in development. Just even if you think about the initial, the very first few, it was a year between the original Mac, the 512K Mac, the Mac Plus, and then, yeah, the SE and Mac II came a year after that. But then the floodgates opened when it was all

those Performas. That I always thought was Sculley's thinking back to the Pepsi days where the focus was how much shelf space did you get in <laughs> the supermarket? You just have all these different products, so <laughs> someone's gonna buy one of them <laughs>.

Hsu: So that was partly Sculley's decision?

Davidian: The Performa stuff happened I think under Sculley.

Hsu: Oh, started under Sculley.

Davidian: Yeah, that's definitely when the explosion of products happened, and in some ways I feel like I'm somewhat <laughs> responsible for that. The universal ROM made it a lot easier to make a slightly different computer quickly <laughs>. Yeah, that's why there were things like the [Mac II]vx. <laughs>

Hsu: Yeah, I remember there was a cover of MacWorld with a vx on it, and I remember looking at it and going "Oh, that's the new Dream Machine," and then six months later the Centris 650 comes out and obsoletes the whole thing. Everybody who bought it were so mad.

Davidian: But once Steve was back-- and we got put out of business, and ironically, right after our last day, Apple wanted to talk to us-- well, some of us, about six of us-- engineers about coming back. So Jon Rubinstein first talked with us, and after what just happened to us, we weren't all that excited about it.

Hsu: Yeah, I could see that.

Davidian: And he said "Would it help if Steve talked to you?" So we did have a meeting set up with Steve, and he still didn't quite convince us.

Hsu: What do you think they wanted you to work on?

Davidian: That was one of the things. In our meeting with Steve, one of the things that was brought up by us was we're a team; we've worked together as a team; we'd like to remain a team. And his feelings were that "We need you wherever you're needed, and we can't guarantee you'll stay as a team." So that was a turn-off for some of us. Looking at it from his side, it makes sense, but from our side <laughs>--

Hsu: So then the entire team joined FB&G Engineering?

Davidian: Not the entire team.

Hsu: Most of it?

Davidian: Some of us, we formed FB&G Engineering. <sighs> It was John Fitch, Gary Rensberger, me, Bob Hollyer, and-- trying to remember his name. He was the marketing guy at Power Computing. Mike Rosenfeld. And we had actually been approached by Newer Technology to help out with their G3

accelerator card, so we actually developed the G3 accelerator for the original Power Macs, the 6100, 7100 and 8100. <laughs>

Hsu: So you were sort of a subcontractor working for Newer. You were contracted by Newer to produce these cards.

Davidian: Yes, we were a corporation, but our client was Newer Technology for this first project we worked on. So, yeah, we developed the G3 accelerator that they sold. They didn't quite follow what we were telling them as far as how to produce it, but they really needed to get into mass production in order to get the price points that <laughs> they were planning on selling it at. But they delayed things enough that it had to be essentially the prototype quantity production of some of the boards, so the build cost was higher than we had designed things for. It was very popular. It was a very good product, a lot of people-- very fast. When we were at Power Computing, we produced a G3 Power Mac based on the 8500 design, and that was a very fast computer. That's the one that said-- that was the last straw that got us put out of business, 'cause Apple had nothing that could compete with that. And we never sold it, but it was fully designed. So, yeah, the G3 accelerator was the first thing we did. Then there were a few other projects we did that were related to the processor but weren't actually a product. It was just evaluating silicon for PowerPC product. And I don't think we did much after that. I think we kind of disbanded after that. I think it was at least a year that we were involved in that whole thing.

Hsu: So why did the group disband?

Davidian: Project was done.

Hsu: Oh, I see.

Davidian: Yeah, and then I think in the middle of it, I think Gary Rensberger left to go move to Seattle and work for Microsoft.

Hsu: So I guess it wasn't really like--

Davidian: It wasn't a big company-- yeah, it was run out of my townhouse <laughs> in Mountain View--

Hsu: So it was really just a group of you got together to work together and sort of do these projects.

Davidian: Yeah, yeah, it wasn't--

Hsu: It wasn't like you had planned to create this startup that would--

Davidian: No, it wasn't going to live on. It was mainly created as a way for us to all work on this project together. Yeah, and then I guess after that, John Fitch went on to Kerbango. That was sort of the next step for-- and Bob Hollyer also. They both went to Kerbango. It did kind of just slowly dissolve. <laughs> I can't really think of when it actually ended <laughs>.

Hsu: So why did you decide not to go off and join Kerbango with them?

Davidian: I was actually doing a lot of other consulting. I was getting approached by a lot of people during the time I was at Apple and at Power Computing.

Hsu: More clean rooms?

Davidian: No, no, not always clean rooms. I actually did some work for Metrowerks.

Hsu: Oh. That makes sense.

Davidian: Well, when they were getting into Java, <laughs> their Java Virtual Machine, I accelerated the performance of it. I was focusing on performance. And I did some stuff with-- after Peter Hoddie left Apple from the QuickTime team, I'd done some stuff with the startups that he was at, although that's actually several years later. He was still at Apple for a while. Trying to think of some of the other projects I worked on.

Hsu: But you've basically been independent since '99, right?

Davidian: Yes. Well, I formed my own corporation, but, yeah, I'm the only employee. So after that meeting with Steve Jobs where as a group we said no, Apple's need didn't go away.

Hsu: So sometimes you took jobs with Apple?

Davidian: They would approach me. So, again, most of what I was doing was focusing on performance. So it was probably '99; it was '98 or '99 when they approached me again, and it was the Pro Apps group. It was Final Cut Pro. This was also about the time that the AltiVec vector instructions were introduced and G4 products were coming out, so they were interested in trying to accelerate some of their performance, so there was some stuff I did there. But really in Final Cut Pro, the app itself is not the performance problem. It was QuickTime. All the time was actually spent in QuickTime. Every frame of video that's compressed or decompressed is going through QuickTime. So they actually wanted me to work on the QuickTime components that they depended on, so I was actually working with the QuickTime group but being paid by the Pro Apps group and working on the DV codec that all the video cameras of that time used, which eventually iMovie was using it also. And another problem with the DV codec was just the quality of the compression and decompression of it. It was developed in Apple in Austin, the group that was working on FireWire, and it was just sort of part of the things they worked on. And it was a small team there, so there were actually several people working on different parts of the codec, and really what it needed was one person working on it and being responsible for everything and really looking at the whole thing, and that was something I could do. And I also addressed some of these quality issues with the compression and decompression. Through multiple generations it was losing a lot of quality in the image, and the customers were feeling that a lot, and Apple was getting a lot of complaints. And it actually had gotten to the point where there was a third-party company, ProMax, that sold a lot of add-on hardware for the pro developers. They actually had someone develop a new DV codec for them. Pro

video users were actually using this third-party codec instead of Apple's, and so Apple really wanted to address that problem, and that's one of the things I did. And eventually Pro Max stopped selling that. It got to the point where <laughs> they encouraged everyone to use Apple's native codec.

Hsu: So you took a lot of jobs with Apple, but you never decided to come back full-time.

Davidian: No, I didn't want to be an Apple employee. Part of the problem when I was an employee there was I didn't really fit into one organization, especially when working on something like performance, because that's all over the place, and when you're in one place, you're focused on what everyone else around you does. So the emulator was a performance project <laughs> really, but there were a lot of others. So I was-- between the vector instruction set, the multiprocessor and just general optimization of algorithms. Those were the areas I focused on, and it didn't really matter what the software was that had this performance problem in the beginning <laughs>. It was just applying those techniques to that software. So I just would describe myself as making slow things fast, and Apple had buildings full of people making slow things <laughs>, so there seemed to be an infinite supply. <laughs> So, yeah, Pro Apps was where they kept--

Hsu: Did you ever contribute to the Accelerate framework?

Davidian: No.

Hsu: 'Cause I think René Vega worked on that as well later on.

Davidian: Yeah, I don't know if he did or not. I know the group that that was in, and I don't think he was ever in that group. I had a philosophical difference with that approach to things. If you actually really wanted to get good performance, you didn't want to use canned routines like that, because everything had to go back to memory. It was just, okay, do the adds on the entire vector; now do all the subtracts on it, instead of load it once, do add, subtract, whatever, and then store the results. So, yeah, I was just focused more on implementing the entire algorithm in vector instructions. The Accelerate seemed to be more for people who did not want to understand or program at the low level, and that was not me. <laughs>

Hsu: Obviously, yeah.

Davidian: I'm sure it had its place, but not me.

Hsu: So you were sort of like a come-in-and-fix-it kind of a person.

Davidian: Yeah, it was "What's next?" so, yeah, I had worked on the DV codec. Then about the time the Titanium Power Mac came out--

Hsu: Oh, right, the PowerBook G4.

Davidian: Yeah, Titanium PowerBook G4. They wanted Final Cut to run on that. It had the bigger screen, but it had limited storage. So they actually were using Motion JPEG at that time for storing things, so I worked on the JPEG codec for a while. And I actually helped out the iTunes team when they were transitioning from MP3 to AAC about the time the iTunes Music Store was about to come out. Helped out a little bit in performance optimization for the AAC codec. And then some of the Final Cut team moved on to Aperture, and I eventually started helping them, too. So that was sort of the last thing that I was heavily involved in, just all the RAW [photo] processing. They would produce a reference design that actually you had to find what all the image processing was, and my job was to get the exact same results as fast as possible using vector instruction sets and multiprocessors. That started on PowerPC but continued through Intel, so at times I was actually developing on PowerPC and Intel at the same time, trying to do performance optimizations on both. <laughs> Yeah. And actually at that time, the G5 was the test environment where I actually could have more processors than Intel, 'cause I only had two cores on the initial Intels, and I could have four on the G5 <laughs>.

Hsu: So let's see. So was that the last time you did work with Apple, or have you continued to do contract jobs?

Davidian: Some of the Aperture stuff related-- again, it was not directly on the-- it wasn't the app itself that I was working on. It was the system software that they depended on heavily. But that was probably around the end of-- Aperture eventually got cancelled, so that kind of ended <laughs> Aperture.

Hsu: I had a friend who worked on Aperture I think, Mike Engber. Did you work with him?

Davidian: Know the name.

Hsu: I have written down that-- were you affected by Steve Jobs's commencement speech at Stanford?

Davidian: Yeah. He talked about connect the dots. That was one of the things he talked about, and I think a lot of the things we talked about in our previous session here was you could look back at that and connect the dots of how all of that led up to the emulator, all the experience that went-- all of my prior experience that made me ideal for producing that 68K emulator. Just the microprogramming—just the fact that I was at Data General, <laughs> that was a learning experience for what projects to avoid <laughs> in addition to the technical aspects I learned there <laughs>.

Hsu: The political--

Davidian: Yeah.

Hsu: How to navigate the political--

Davidian: Well, no, just FHP was something where it was just trying to accomplish too much and ignored too little of the past. It was ignoring the things that mattered in the past. So that was what I always looked at as-- because I didn't want to invest a lot of my time in something that was going nowhere, and I

had the same feeling with startup companies, too. You even mentioned Kerbango; why didn't I go work there? With startups, you could be there for a decade and have it turn into nothing. <laughs>

Hsu: It's such a roll of the dice.

Davidian: Yeah, and there's other things you could be doing during that time. So I like being able to work independently on small projects and see each of them complete, and maybe that little piece doesn't get used, but I didn't waste a big chunk of my life. <laughs>

Hsu: So what sorts of projects have you been continuing to work on since then?

Davidian: I haven't actually done that much. Personal life has taken a lot of my time.

Hsu: So are you semi-retired now?

Davidian: Yes. I've always said that if people keep calling me, I'll keep answering the phone, and if there's interesting projects, I'd like to work on them. But I don't really feel the need to go look for something that I might not enjoy and spend a lot of time on that. One thing in my personal life, my stepfather passed away a little over three years ago, and since then I've been going back about once a month to help my mother out. She's still living alone at home now. And when you look at it, I could be spending time with my mother, or somebody could be paying me to work on something that I don't really enjoy. I can't buy that time with my mother in the future. <laughs> And is this something that's being offered to me? Is it something that is worth sacrificing time with my mother?

Hsu: Yeah, okay. Yeah, that makes sense. But you've done well enough that you can afford to not work.

Davidian: Yeah, I'm fortunate to be in a position where I don't need to go look for work. When I was consulting, people didn't select me based on price. <laughs> It was based on track record and what I could do.

Hsu: It's a good position to be in.

Davidian: And a lot of it actually goes back to the NEC case where, besides the initial payment and the bonus payments, they had set an hourly rate for just follow-on work, just maintenance work and things like that, and that was set by lawyers who know what a lawyer's hourly rates are <laughs> like. So my hourly rate started there, and it increased over time. People needed me. They weren't necessarily happy writing the check, but they were happy with the results and thought they were getting their money's worth. 'Cause a lot with performance is, they can't buy a faster processor. <laughs> You have to make the software faster if you want to improve performance.

Hsu: So we're already sort of talking about connecting the dots. How would you maybe summarize looking back on your career? How would you characterize it?

Davidian: I think part of it is I don't think it's easy to define what it is I do. I started off as a microprogrammer, but really when I look back at what I'm doing, it's really performance is what I do, and it's also at a level where I'm sort of working between software and hardware. It's very low-level. And I guess these days microarchitecture is kind of where that all comes together, but that's really at the processor design level, and I'm sort of applying that at the software level where I fully understand how the microprocessor and its microarchitecture are designed, and I can sort of think like how the processor executes, and I can write software that takes advantage of all these components of the microarchitecture that contribute to higher performance. So I don't know how to really summarize it. I was just sort of this odd fit that-- I don't know if you can go to college to become-- I don't think there's a degree in what my life experience here has been. A lot of it's been just learning as it's been going on, as the technology's changed. None of this stuff has-- stuff I'm doing at the end wasn't being done when I was in college. <laughs> In some ways I'm a dinosaur, because writing assembly code or even microcode is rarely being done these days. So I don't know if that answers your question, but <laughs> those are some of my thoughts.

Hsu: What do you think of the trends, the hot things that have been happening? The move to GPUs and machine learning and cloud computing, mobile, all this stuff?

Davidian: Yeah, there's lots of stuff going on.

Hsu: Cryptocurrency.

Davidian: Yeah. <laughs> GPUs, they started as a solution to a specific problem, but people are seeing just how much compute power there is there that they can harness it for other things. It may not be optimal for those things, but there's just so much performance there <laughs> that even suboptimally it's an advantage. There's one other thing, not really reply to your question, but more just about the environment here in the Valley, not these technologies but just the way the hiring is done. I sort of think of it as there's kind of this arms race going on between all the big companies here as far as just grabbing all of the new talent coming out of college. Between Facebook, Apple, Google, whatever, they're basically hiring anyone with a degree that comes out. And all of these people, they want to ship something, and they also haven't had all the experience of using all the previous generations of the products that these companies have made. So when I look at some of the stuff that Apple's done and some of the newer versions of OS X, I don't necessarily see them as an improvement. It's more like somebody did something, and they want to ship it. <laughs> And I know the feeling of wanting to ship something. I was there, too. <laughs> But the other thing that's happening is all the senior people, they made a lot of money, and they've either left and retired <laughs> or started their own company. And I always viewed the senior people as sort of this output filter to say "That's a really stupid idea," <laughs> or "We're not going to ship that." And so the output filter is disappearing <laughs>, and there's this urge to ship stuff and not be familiar with what the old versions really did and why they did it. So I just find it frustrating to actually use some of the stuff, 'cause a lot of stuff I don't view as an upgrade. It's like what did they take away now? And it's not even just Apple. Even websites you're seeing this where "We've redesigned our site, and now you can't do these things <laughs> that you used to be able to do easily with one click." <laughs> It's just I think just too many people in-- I understand all the companies wanting

to hire them so that the other company doesn't get them, <laughs> but I don't know if there's actually the work there that needs to be done.

Hsu: So what would you say maybe to somebody who's just being hired, a young college graduate who's just being hired by one of these companies today?

Davidian: Sort of the first thing I did when I got to Apple was learn as much as you can about the existing product. Just dig in and understand every little corner of it and how it works and why. Try to understand the why and how it evolved over time going back to previous versions, and understand the thinking of the person who originally wrote the code. I tried doing that when I was working on the Mac ROM code early on. You could sort of see the style. Each author had a different style, different way of thinking. Bill Atkinson's code looked more like compiler-generated code. He wrote a lot more Toolbox code, which was callable by the-- had Pascal calling conventions, so his code looked a lot more like that. I used to look at a lot of Larry Kenyon and Andy Hertzfeld's code. But I would also just try to keep the same style that they wrote in. If I'm modifying some of their code, modify it like how I think they would've done it unless I was totally rewriting something, and then I'd learn from their code, 'cause sometimes it got to the point where something was modified so many times that it's better to really just-- it started with some original requirements and capabilities and ended with a lot more, and just start by looking at what it ended with and try to rethink how you would implement it from the beginning to support all of those requirements. So that's how I would think about things, but it's really just really wrap your mind around something and try to become an expert in it before you start changing it.

END OF THE INTERVIEW

