



## **3dfx Oral History Panel**

**Gordon Campbell, Scott Sellers, Ross Q. Smith,  
and Gary M. Tarolli**

Interviewed by:  
Shayne Hodge

Recorded: July 29, 2013  
Mountain View, California

CHM Reference number: X6887.2013

© 2013 Computer History Museum

**Shayne Hodge:** OK. My name is Shayne Hodge. This is July 29, 2013 at the afternoon in the Computer History Museum. We have with us today the founders of 3dfx, a graphics company from the 1990s of considerable influence. From left to right on the camera-- I'll let you guys introduce yourselves.

**Gary Tarolli:** I'm Gary Tarolli.

**Scott Sellers:** I'm Scott Sellers.

**Ross Smith:** Ross Smith.

**Gordon Campbell:** And Gordon Campbell.

**Hodge:** And so why don't each of you take about a minute or two and describe your lives roughly up to the point where you need to say 3dfx to continue describing them.

**Tarolli:** All right. Where do you want us to start?

**Hodge:** Birth.

**Tarolli:** Birth. Oh, born in New York, grew up in rural New York. Had a pretty uneventful childhood, but excelled at math and science. So I went to school for math at RPI [Rensselaer Polytechnic Institute] in Troy, New York.

And there is where I met my first computer, a good old IBM mainframe that we were just talking about before [this taping], with punch cards. So I wrote my first computer program there and sort of fell in love with computer. So I became a computer scientist really.

So I took all their computer science courses, went on to Caltech for VLSI engineering, which is where I met some people that influenced my career life afterwards. So from there, I went to Digital Equipment (DEC) and worked on facts [fabs?], chips, and tools. Off to Silicon Graphics, I worked on workstations there, mostly software and chip design software as well, CAD tools and the like.

And then, when was it? Around the late '80s the PCs started coming into play. And so that's when I first met Ross and Scott. And we started thinking about PC graphics and such, which was sort of a heretical at

the time I would say from the big workstations. But we went off to Media Vision. Some of them went off to Media Vision and Pellucid. And so then it became a reality around 1993.

**Hodge:** Where you at MIPS or were you directly with SGI?

**Tarolli:** Directly with SGI.

**Hodge:** OK. And were you in Boston this whole time?

**Tarolli:** I moved out to California when SGI was just getting going in 1983. And so I lived out here for about three years with fond memories right across the street from here and Moffett Field on the other side of Moffett at Clyde Court, so one of the very first buildings.

**Hodge:** How did you recruited into SGI at an early stage?

**Tarolli:** That was from people I had met at graduate school, who after a couple years, called me up and go, hey, you want to come out to California? We got this interesting start-up called Silicon Graphics. So it was the people that I met I would say at graduate school that sort of influenced my career afterwards.

**Hodge:** And what convinced you to come out here and do that? I mean, DEC at the time was a very stable and well-known company.

**Tarolli:** Good question. I think it was somewhat of my trust in them and their recommendation, and then coming out and meeting people. And you know, wanting to do something interesting in graphics also. That was the first that I had really gotten involved in 3D graphics.

**Hodge:** OK. Scott?

**Sellers:** I was born in Dallas, Texas and moved when I was very young to Colorado. So I grew up in Colorado. First cut my teeth on the very first computer in fifth grade. It was a Radio Shack TRS-80, programming into BASIC language, and did a bunch of stuff for the school district.

And I guess we were pretty advanced at the time. We were doing a lot of the taking textbooks and putting them on to the computers and developing tools and teaching aids. So that was how I first learned to program was back then. I went to school back East at Princeton.

And there I was an electrical engineer focused on computer engineering. So it's a combination of both VLSI engineering as well as computer science. Out of Princeton, I went to work for Silicon Graphics. That was my first real job, which was a fantastic breeding ground I think for learning engineering, learning obviously a lot about 3D.

At the time, I was working on memory subsystem chips for the workstation division of SGI and met Ross and Gary through Pellucid, which was a small start-up that was founded by some SGI executives that, as Gary said, left to start to think about developing graphics cards for the emerging PC market. Pellucid was acquired by Media Vision. And then from Media Vision, we left to start 3dfx.

**Hodge:** OK. Ross?

**Smith:** I was also born in Texas, although in Nacogdoches, which most people have never heard of, and haven't been back there since. But went to school in Austin, Texas and got my degree in 1983. My first computer was actually a Cyber 7600. But it wasn't really mine. It was a timeshare system that the school had. And we programmed that in BASIC as well. And my first program I think was Tic Tac Toe, which I think is a very common thing for someone to program in BASIC. So I've always liked games and computers and the combination thereof. And after school, I came out to California. And at the time that I left UT, there really wasn't a whole lot going on in technology in Texas besides oil. And that wasn't very interesting.

So I came out to California and got in the defense business, worked for Ford Aerospace working on satellites and things like that, and also working on big workstation products and stuff. So ended up leaving there and became a consultant to help computer companies get into the defense market. My first client was Sun. Actually, my first client was AT&T, but that's a longer story. Then I worked for Sun. And then at Sun, I met these people who had just left Sun and went to MIPS. And so they said, well, come on over and help us. I then went to MIPS. They made me an offer, and I got in the computer business at MIPS, where I did business development and marketing. And MIPS was acquired by SGI in '91, '90, something like that. And I went there and met these crazy guys that thought that the PC was going to take over the world. And that seemed logical to me. And so we spun off a company called Pellucid from SGI. And the mission there was to bring 3D graphics to the PC. So at COMDEX of -- was it '92? Really early on, we had a PC running real 3D graphics. And of course, we only had the Spider application, we had the Flight Simulator. We had like three applications for it. And that was the real dawn of the notion [to us] that applications are really important. So Pellucid wasn't a stellar success. But it gave us a good grounding as to what the PC could do and set the stage for the next effort, which became 3dfx.

**Hodge:** What was Pellucid's actual product?

**Smith:** We had the ProGraphics 1024, which was the first true color PC graphics accelerator high resolution and high performance for the PC. Apple had a great true color environment. The PC didn't at the time.

**Hodge:** What was true color in the context of that?

**Smith:** 16 million colors, 24 bits per pixel.

**Hodge:** OK. And the resolution?

**Smith:** 1024 by 768.

**Sellers:** Probably 1280 by 1024, too. [INAUDIBLE] listed to the [INAUDIBLE].

**Smith:** And it had three-way memory interleaving, because of the memory guy here—[points to Scott]

**Sellers:** That's right.

**Smith:** --which has probably never been done since or before.

**Sellers:** Never before or since.

**Smith:** And the cost?

**Sellers:** What was it? \$1,000 or something.

**Smith:** I don't remember. Yeah, it was about \$1,000. It was a volume product.

**Sellers:** We saved a quarter of megabyte of memory or something like that [by doing 3-way interleaving].

**Smith:** Memory was really expensive back then.

**Sellers:** It was a big deal, yeah.

**Hodge:** And when did you move from engineering to business development and why?

**Smith:** I was at Ford Aerospace. And we had invented a high speed network. This is 1984. We had a 100 megabit per second fiber optic network. And I thought there would be some commercial applications for it. So I went to the VP of my division, [and] said we should go sell this thing. And he goes, "well, we don't know how to do that. Go figure it out." So I went and sold one. And that was kind of my first entry [into biz dev]. And then, of course, after I sold one, they wanted to get rid of me. Because they said, how do we deliver this? And we don't have the documentation. And how we do customer support and all these things? And I was like, well, that's my problem. I sold it. You got to figure that out. So that was the transition to sales and marketing.

**Hodge:** Gordon?

**Campbell:** I'm actually from the other side of the United States. I grew up in Minnesota and spent some time in engineering, spent some time in sales. And then in 1976, I decided I was going to try marketing. And there wasn't much going on in Minnesota. So I went out to Silicon Valley. And I interviewed with Intel, National, AMD, and one other guy, and ultimately got hired by Intel. And I started their corporate marketing group for all their major customers and all their major products worldwide from scratch. Built that up, ran that for three years. And that was when Bob [Noyce] and Gordon [Moore] were still running the company. It was kind of the Camelot phase of Intel, if you will. Then [Andy] Grove got promoted and asked me to take over the marketing responsibility of our largest division, which was non-volatile memory. It was kind of fix-it assignment. And so I did that for two years. And I left and I started my first company, Seek Technology. And Seek was a non-volatile memory company. But we actually had the opportunity to, in conjunction with Silicon Compilers and John Doerr and 3Com and Ungermann-Bass, build the world's first Ethernet chip, which was a bit strange at the time. Because it was a chip nobody knew what it was doing or what you would do with it. Took Seek public, ran it for four years and left, then I started Chips and Technologies. And Chips and Technologies is the company that set the standard for the fabless semiconductor industry. The reason it did that is that I decided, having done the traditional approach with Seek, where you not only build your wafer fab, but you design your own process and you design your own parts, that's just way too complex. So I decided that Seek [should be C&T] would be a company that would never ever have a fab. And that was a bit heretical at the time. Jerry Sanders is quoted many, many times for saying, in some conference we were at, Gordy, real men have fabs. And then we also invented the chipset, which is how all PCs and Macs are built today. We did the world's first IBM compatible graphics controller. Later, we added the LCD capability to that, which is what enabled laptops and everything else. All the chipset and the graphic stuff we did back in those days is fundamentally what enabled all the smartphones, the tablets, the laptops, all that stuff. Then after running Chips for about 10 years, I decided I was going to do something different. So I decided I would create what was kind of an unknown concept at the time, an incubator, and I'd work with a lot of the other new entrepreneurs. After about a year of that, I decided that that wasn't really working that well. And we needed really to have our own fund. So we actually raised our first of three funds. And somewhere in that time frame, Ross was consulting at one of the companies that I was involved in. I guess we were interviewing you.

**Smith:** Yes.

**Campbell:** So I was interviewing Ross to come in and potentially help us at one of the companies. And it was the worst interview I think I've ever had. And so finally, I just turned to him and I said, OK. Your heart's not in this interview. What do you really want to do? And he kind of looks at me surprised and says, well, there are these other two guys, and we want to start a 3D graphics company. And the next thing I know, we had set up a meeting at the Tied house. And we had, over a lot of beers, a discussion which led these guys to all come and work at my office. And that set up the start of 3dfx. And there are tons of stories about all the things that transpired there. But that's how this all came together.

**Hodge:** When you were at Intel, were you part of the Crush marketing program there?

**Campbell:** Everybody was part of Crush at that time. Crush, I think it was one of the most amazing programs I've ever seen in terms of motivating an entire company to all have a single focus. But what I think is interesting about Crush, and it's not what you'd read in any of the books on it, is that the program was to go out and reverse a trend. Motorola was getting all the microprocessor design wins. We were getting zero. And Dave House was the guy running the microprocessor group at that point in time. And House and I would go in-- I was a non-volatile memory guy. He was the microprocessor guy. And so we'd go in once a month to executive staff and give our reports. And you know, mine, we were hitting the ball out of the park all the time. Because we had just introduced the first five volt non-volatile memory. And I mean, the world couldn't buy enough of them fast enough. But the flip side of that was that we were getting zero design wins in the microprocessor world, until we got this one design win, which was the IBM PC. And of course, that changed everything. But I don't think that we actually got that because of Crush. And as spectacular and as organized and as impressive as Crush was, I don't think that was really what flipped that tide.

**Hodge:** Sorry for that brief aside. So going back forward to the early '90s where we left this off, Gary, were you still out here in California at that time? Or had you reverted to Boston?

**Tarolli:** I was working back on the East coast, remotely for Silicon Graphics out of the sales offices and such. And then eventually, I left Silicon Graphics and did some consulting locally also.

**Hodge:** So what was the state of the art of 3D graphics at that time? There were several movies that came out in the early '90s that sort of brought computer graphics to attention, *Terminator 2*, I think, and then also *Jurassic Park*, which I know was worked on by SGI. That [have] any bearing on what you guys were doing?

**Tarolli:** I would say probably not, because that was sort of at the high end. And we were starting at that time to think about the PC graphics market, which was going to have to be really much more primitive

than that and simple. Although, the interesting thing that we always say is what we do in the movies with offline animation we'll be doing in real-time probably 5 to 10 years later.

And then 5 to 10 years later, we'll be doing down on a PC level. And then nowadays, it's like 5 to 10 years later, you're doing it on your smartphone. So it's just a matter of time that it progresses down.

**Hodge:** So what was the state of PC graphics when you guys, say, started Pellucid.

**Tarolli:** I don't think there was a state at that point in time.

**Smith:** There was, IrisVision.

**Tarolli:** Oh, yes, IrisVision. At SGI, one division had built an IrisVision board, which was Microchannel I believe for IBM PCs.

**Smith:** ISA.

**Sellers:** ISA or Microchannel.

**Tarolli:** Yeah. And that was on to 486, I believe, it first came out. So one of the problems with the PC is that the 486, I think, has just come out. I think I worked on it for just a few months or a year. And the PC was really pretty primitive. So even just coming from a workstation, it was just slow. But within a year or two, thereafter the Pentium 90 started coming out. And then it just really start steamrolling in terms of the CPU side of it as well. Dedicated PC graphics that started there, I would say there was probably none.

**Sellers:** There was really none. And SGI was the first to create this IrisVision, which was this massive board that barely would fit into a full-size ISA slot. But SGI, of course, had no idea how to sell that product. Because for them, it was all about very high end workstations. And so part of the formation of Pellucid, the start-up company that we all worked for was acquiring the rights for IrisVision. And so one of Pellucid's early products was actually trying to resell the IrisVision product into the PC channel. And that had limited success.

**Smith:** Actually, we sold them all. We sold every one that was ever made.

**Sellers:** Limited success.



**Tarolli:** Might not have been a lot though, right? I think they were \$4,000 to \$5,000 a pop.

**Smith:** Yeah, it was 2,000 units.

**Sellers:** Yeah.

**Tarolli:** But they were very expensive, right?

**Smith:** So it was Gouraud shading with a 16-bit Z-buffer maybe?

**Sellers:** Probably. No texture mapping.

**Smith:** No texture mapping. And 10,000 polygons per frame-- or per second.

**Tarolli:** Per second.

**Smith:** By today's standards, a software renderer on your iPhone would do a much, much better job. And there were no apps for it. That was the real issue. For 3D graphics, there was nothing in that time period.

**Hodge:** So when you guys were in Pellucid, were you trying to engage with application developers? Or were you perhaps too focused on hardware at that time?

**Smith:** We had no idea what we were doing. In general, I think. We had hardware. And it was becoming apparent to us-- now, I wasn't the marketing guy. So the marketing guy is now a big success at Google and shall remain nameless. We were struggling to find a market for this technology we have. And as we all now know, that's really not a good way to do things. Typically, you would find a market, then build some technology for it. We had this great 3D graphics board. We had a hybrid MIPS x86 workstation. And we had ultimately, this true color graphics board. And for each of those different pieces of technology, we didn't have any applications. We had a handful of apps. And so that was the real epiphany for me. Whatever we do next, apps are really the key. And there was good 2D graphics available in the PC space. There were lots of guys. Chips and Technologies and other guys made 2D graphics subsystems and chips for that marketplace, but not much in the way of 3D at all.

**Sellers:** I think ultimately, that's what allowed for Pellucid to be sold to Media Vision, this interest in this Windows accelerator market, which was really blossoming at the time. So Weitek was another one. We had all these other products, this IrisVision 3D product, which wasn't really going along. We had this idea

at the time also to take a MIPS processor and combine it with Windows NT. At the time, Microsoft was saying that Windows NT, which would run on all sorts of different architectures outside of the x86, that you could create a very cost-effective workstation products based on the MIPS processors. And so we had a design effort to build this combination MIPS x86 as Ross mentioned. And that didn't really go very far either.

**Hodge:** So you guys were really thinking of workstation non-gaming application at this point?

**Sellers:** Well, I think a lot of it was at least some of the folks at Pellucid came out of the workstation part about SGI. So this idea that you could build a "cost-effective"-- and I use that in quotes-- we're talking something that would maybe be in the \$5,000ish range, something that SGI never could build. Because every time they would try to build something that low end, well then the audio guys would come in and throw like \$1,000 of audio stuff into the box. And suddenly the \$5,000 workstation is a \$15,000 workstation. And we saw this cycle over and over again. But it was definitely this mentality that you could do a lot more in a low end workstation. And that kind of the common theme at Pellucid, what could you do? Whether it was 3D or this kind of hybrid x86 MIPS, or maybe Windows Accelerator Graphics products, that was all sort of in this common theme.

**Smith:** Value.

**Sellers:** We definitely were the classic unfocused start-up I think. We did way too many projects for the people that we had. But ultimately, we found something that worked in the Windows Accelerator. And that's what caused the exit to Media Vision.

**Hodge:** So how did you guys go from, I believe you described it as, the worst interview you ever did to a new company particularly with-- sorry, pointing at Gary being in Boston. And imagine you weren't giving him-- did you get him up to the Tied House to pump beer into him, I guess is the question?

**Smith:** Well, Gary was working at Pellucid as well. And so he would be out here [in CA] quite a bit. I don't remember the exact conversation we had. But there was another guy that we were involved with at the time, who was actually a founder of SGI, Herb Kuta. And Herb and I, at Pellucid, we had this prospect come by one day looking for 3D graphics. Because while there was 3D graphics in the movies, there was [also] the emergence of 3D graphics in video games, and in particular in the arcade space at that point in time. Was it Sega Daytona or Virtua Fighter? There was an early pretty nice video game that had come out. And all the competitors to Sega were looking for an alternative. And this guy showed up and we had a chat about 3D graphics. And he said, hey, there's a trade show in Anaheim this weekend. You should come down and check it out. So Herb and I trundled off to Anaheim. And we saw what was going on down there. And it was kind of amateurish compared to what SGI had done. And Herb and I looked at it and thought, well, this is kind of interesting. We might be able to do something here. And then we kind of

explored the market. And there were apps. Because there were these games, and people were really enthused about it. And so I don't recall how we actually started the conversation about doing something along those lines, but it was to build 3D graphics chips for games. And Gary's one thing was like "no CAD". Whatever we do, we can't do any CAD. And so the first chip didn't even have a line engine in it [which is essential for CAD].

**Sellers:** Well, the precursor to all that, though, was at Media Vision, after we had delivered these Windows accelerators, we were looking at what to do next. And I don't remember if that-- the AMOA [Amusement and Music Operators of America] show, right?

**Smith:** Yeah.

**Sellers:** That's the one you were talking about, before or after the Media Vision. But I do remember we had proposed a project within Media Vision to build a 3D game engine, a 3D chipset specifically for games. And Media Vision was seemingly a great company to do this for, because Media Vision was a consumer company. They sold these multimedia upgrade kits and these high end sound boards. And they had the retail channel. And they actually had pretty good relationships with the developers, because of all the audio stuff they were doing with whatever their Sound Blaster equivalent was.

**Smith:** Multimedia kits.

**Sellers:** So it made a lot of sense to actually build this 3D product as part of Media Vision. Except there was just one minor problem, which was Media Vision was run by crooks. [Good article about it on [http://en.wikipedia.org/wiki/Media\\_Vision](http://en.wikipedia.org/wiki/Media_Vision)] But that's another story. But the fact that Media Vision almost overnight cratered, because of the financial system that was-- [basically] old school cooking the books kind of thing. Of course, we were completely separate and had no knowledge of this or any involvement.

**Smith:** Not that we recall.

**Sellers:** Wake up one morning and in the paper you read about your company that seems to be run by crooks, obviously everything changed right then. And we actually were getting resistance within Media Vision to building this product. It was just kind of beyond the scope of what they could really grapple with.

So all this kind of came together pretty quickly that Media Vision was crumbling. We wanted to get out of there. We wanted to build this 3D game engine product. And we left Media Vision. And right around that time is I think when we met Gordon.

**Hodge:** So this was at the time in the consumer market then when a lot of PCs were coming without sound cards or CD-ROMs. So you'd go to the store, you'd buy a box. It would have a sound card, a CD-ROM, speakers.

**Sellers:** Multimedia upgrade kits as they were called. Yeah. You'd get your CD-ROM, your audio board, your speakers. It was a big market.

**Smith:** Some CDs with some games and stuff.

**Tarolli:** What was it? The X-wing fighter game with a joystick that you could fly down the Death Star channel. But it was all pre-recorded video. There really wasn't any [real-time] 3D graphics on it. So at Media Vision, we spent a lot of time trying to convince people what the difference between the pre-recorded CD-ROM game that had various avenues, decision points and true 3D graphics was.

**Hodge:** And Myst was kind of the paradigm wasn't it? The game that was on rails. And everything was pre-rendered.

**Sellers:** [Pre-]Rendered game.

**Tarolli:** Yeah, this was a pre-rendered game. And so we wanted to sort of free them from that. And that didn't play well with some people. Plus, the other thing is 3D technology then was expensive. The IrisVision boards were into thousands of dollars. So that's where I think the arcade business sort of made sense. Because those machines were in the thousands of dollars. And they could afford 3D graphics that was into \$1,000. But the PC market, that was I think out of the price range. Plus, there was no apps in the PC market. There was no Quake or Unreal engines at that time. No one was writing 3D software for the PC. So it was a sort of an evolutionary thing. But Media Vision going bankrupt was probably a serendipitous event. Because if we had stayed there, we never might have done what we really wanted to do. And that sort of forced us out onto the streets. And we said, OK. It's now or never. Right? It was like, the perfect time.

**Smith:** And having been shown the door [was a pretty good incentive].

**Tarolli:** So this is like the perfect time to do it. Because [our current] company [has] cratered. And we can't stay here anymore.

**Sellers:** And it was a fascinating time for what was happening in 3D. As Ross, mentioned, the PlayStation 1 was either about to come out or had just come out. And that showed, at a very rudimentary

level, what could be done with 3D graphics with very crude texture mapping. And it was pretty abysmal quality. But the consumers were just eating it up.

**Smith:** Loved it.

**Sellers:** And the game developers loved it.

**Smith:** It sold out. And we had had some experience at SGI, either directly or indirectly, on a product called RealityEngine, which was a very high end marquee product from SGI that came out in the '92, '93 time frame. Which as you mentioned, that was the product that was used for a lot of these movies and kind of put SGI on the map in terms of the Hollywood aspect of it. And so when we really looked at it and looked at the architecture behind RealityEngine, we believed that we could take a lot of the capabilities that existed in RealityEngine and bring those down to consumer price points taking a bottoms up approach, taking a much better than PlayStation was able to deliver with this RealityEngine type of technology. Not that we used anything from SGI. But it showed us what could possibly be done using very deep silicon integration. And I think the other thing that was also kind of the a-ha moment, if you will, in terms of this is all possible is coming out of the workstation group at SGI, Gary was one of the first to realize that the microprocessor itself could be used as a very effective geometry engine. And in 3D graphics, there's kind of two major pieces. There's a geometry phase, which is very compute intensive. And then there's the rasterization phase, which is your painting the pixels on the screen, if you will. And SGI's heritage was largely to do everything in silicon. So for Gary to pioneer this idea that you could really effectively separate it [and use a PC to do it], was pretty revolutionary at the time. Because it just hadn't been done. And we did it on some of the lower end SGI workstations.

**Tarolli:** Yeah. My last project there was the Iris Indigo, which we literally implemented the entire IRIS GL graphics library, which is now OpenGL, in software, except for the very last piece of it, which was just RGB shading. So I got a lot of experience there on what exactly it took to do all these algorithms and software. And then the chip technology was getting sophisticated enough where Scott could figure out how to actually build onto one or two chips enough that we could do a lot really fast on one or two chips.

**Hodge:** Did that become the division between you two guys? You did more software, and Scott did more hardware?

**Tarolli:** Yes, I would say so.

**Sellers:** Well, Gary did kind of all the algorithm development, where he and I would work together. And Gary would explain what the algorithm needed to do. And I would try to figure out how to build that chip.

**Tarolli:** How to do it.

**Sellers:** Right. Then Gary would say, oh, by the way, you have to--

**Tarolli:** Do this.

**Sellers:** --divide-- you have to--

**Tarolli:** You have to do a division. And how can we do a--

**Sellers:** And divide every single pixel clock. And at the time, that was a major, major--

**Tarolli:** Stumbling block.

**Sellers:** --challenge.

**Tarolli:** How to do a fast divide. I mean, it sounds trivial today, because you get it for free in a PC. But in those days, doing a division with 20,000 gates, which was half you chip if you didn't do it cheaply. So one of the things I had spent the last year or two at SGI was how do you do things really cheaply and really quick and inexpensive? And how do you approximate [something but still have it look correct]? I still do that to some extent today. It's sort of one of the things that my brain gravitates to. So I was very well suited for how do you "cheat," quote "cheat," and do something really cheap where it looks really good.

**Hodge:** That's why you wanted to avoid CAD, correct? Because that's not really acceptable in that arena?

**Tarolli:** Yes. Yes, if you cheat, you can't do anti-aliased lines. And all the workstation vendors are going to come back and say, "but that doesn't cut it". And they need 24 bit true color, which was out of the price range for the PC market [since it would require lots of memory and a] deep Z [or depth] buffer. But the last project I did at SGI, we had actually 4 bit RGB. One bit of blue, one bit of red, two bits of green. You would not believe how many people just laughed when they heard that. And then I showed them the images, and they said, oh. Oh, I guess I can see what that is. But basically it was almost like the laughing stock of SGI, 4 bit. But you could double buffer it. And it was actually tolerable to look at.

**Sellers:** Because memory was crazy expensive then.

**Tarolli:** Yeah. I mean, all we had was 8 bits in a frame buffer for the low end product, which still cost \$9,000. So you know, getting down to \$100, \$200, was a large step. Some of it took time. I mean, there's no way we could have done this five years earlier. So some of it was just waiting for the right time with technology. And that's where Gordon came in was important also. Because you had to build a semi-custom chips. And I don't even remember all the details of the technology. But we had to use as many gates as we could possibly do. But we weren't going to have 20 people laying out the chip by hand. So a lot of it was what's feasible to do.

**Hodge:** Was it just literally the four of you to start?

**Campbell:** Actually, these guys all came in and camped out in my office. I think you were there for like six months.

**Sellers:** Yeah.

**Campbell:** And the reason it kind of took six months was that there was a mindset that this was going to be a motherboard.

**Smith:** See, he always tries to hang that on me.

**Campbell:** It was kind of a mindset of three against one. And I kept seeing these proposals to do this motherboard. And I kept saying, no motherboard [because motherboards were a commodity by then]. And so finally, after about four months or five months, I saw a proposal that said here's an add-in card [that would plug into a motherboard] . And that kicked everything over the edge. And all of a sudden, the plan came together. We went out and did a bake off to pick our VCs. And everything just started dropping in place.

**Hodge:** So I'm going to editorialize here a bit and say the current Silicon Valley VCs aren't particularly charitable toward companies that spend four months and not really having a business plan they like. What made you patient with them?

**Campbell:** I mean, there were no VCs involved at that point in time. But basically that was kind of a superficial perspective on it. There were a lot of issues and I think Gary was touching on some of them. You are right on the cutting edge of what level of integration and how many transistors and what you could put together. And there is no way you could build a full-fledged CAD quality 3D graphics chip. And so what we did is we cheated. We made something that was good enough, as Gary was talking, from the color standpoint, and enough polygons and enough everything else to be a game machine.

**Smith:** So during that time, these guys are working. So we developed a full simulator that actually became really important to raising money. Because we could actually demonstrate the quality--

**Campbell:** You going to talk about that big Reality Engine?

**Smith:** That came later. That came later. That came later. No, but they developed and they had in software, it was a C simulator.

**Tarolli:** Yeah.

**Smith:** I assume.

**Tarolli:** A complete working chip, in essence, in software.

**Smith:** And we rendered these movies that showed what the quality would be. And we had several examples of that. And we had a Pentium 90, right?

**Tarolli:** Yup.

**Smith:** And at the time, a Pentium 90 was a big deal-- with all the memory we could find.

**Sellers:** And expensive [INAUDIBLE].

**Smith:** And Scott and I lugged that thing across Japan on the subway. And Gordon's is saying, we have to be here, here, and here. You had to go to Japan in those days.

**Hodge:** What was the reason for that?

**Smith:** It's where the games were.

**Tarolli:** The arcade games, right?

**Smith:** Exactly. And consoles.



**Campbell:** It turns out that I had had Japanese investors in Chips and Technologies, which is a different story. But so when the opportunity came to do something here, we welcomed Japanese investors into this and got quite a few of them. So that helped us both from the business standpoint and from the standpoint that we became more of an international company right away.

**Smith:** But while they were working on the simulator, Gordy and I-- he was beating me up about the motherboard business plan. And we kept iterating.

**Hodge:** Was that a motherboard plan, because you were focused on arcade machines and not computers?

**Smith:** The PC had no presence in the arcade market at that time. They were all custom workstation [type single board computers] with MIPS or PowerPC processors. And we were building a [PC-based] product because we ultimately wanted to get to the PC in the home.

So [what we were developing] was a PC-based arcade or coin-up architecture. And so the motherboard was a way to do that. And finally, through the keen insight of our leader, we gave up on that. So we ended up doing an add-in card.

And ultimately, we made it compatible with MIPS. The Midway and Atari game systems were all based on the MIPS microprocessors. So it was actually a good revelation.

Because [our approach] used the same basic concept of a PC, [with a] standard kind of architecture. [However] the packaging really mattered. And how you presented it to the market really mattered. Because motherboard companies had a zero valuation. The Valley was full of them. And Taiwan was full of them. But this kind of special add-in card had a very attractive nature to it to the investment community.

**Hodge:** Who were your first investors?

**Smith:** Is it Mitsui?

**Campbell:** No, actually it was--

**Sellers:** USVP.

**Campbell:** Yeah, USVP was the guy we picked.

**Smith:** US Venture.

**Hodge:** US Venture Partners?

**Smith:** Yeah.

**Sellers:** They were the only round A guy, right?

**Campbell:** No, I think we had a little bit in there from Chase at the time, too.

**Sellers:** That's right.

**Campbell:** Because the introduction to USVP ultimately had come through Chase.

**Sellers:** We brought Tony Sun from Venrock [INAUDIBLE].

**Tarolli:** Venrock was the other one.

**Sellers:** Is that [INAUDIBLE]?

**Smith:** Mitsui, Comtech?

**Tarolli:** I don't remember that.

**Campbell:** Yeah.

**Smith:** Maybe that was our first pitch was to Mitsui.

**Campbell:** But it was kind of interesting. Because when we went out and we picked a small number of VCs, we went out and did the pitch and came back pretty favorable. It was a pretty short process actually.

**Hodge:** So what year was this now? Was this '94?

**Campbell:** About that. And what was interesting about it is that the business plan that we ultimately came up with was one that I think it's the closest business plan I've ever seen from a start-up that never really changed much. And there are very few start-ups you ever say that about.

**Tarolli:** After you fix[ed] the Excel spreadsheet error, right?

**Hodge:** What was the Excel spreadsheet error?

**Smith:** I think our first pitch was to Mitsui. We'd been pulling all-nighters for like days. We're at my little apartment in Palo Alto. And we did have a colleague of ours, a lady who's no longer with us, Alma Ribbs.

And we were working on the business model. Because you have to show what the revenue forecast looks like and P&L and all that. And things looked great. It was like, Jesus, we're just going to print money.

And then we noticed that it didn't matter how we changed the cost, we always had the profit. I was like, something's probably wrong with this. And we found we had a big spreadsheet error. So we fixed it, got everything done in time, and went and made the first pitch. But it was one of those hair-raising moments of oh, jeez.

**Sellers:** We went from a great business to it's a disaster in one formula change.

**Hodge:** What was that climate like for venture capital of the Valley at that time?

**Sellers:** I think that was really right at the beginning of when venture capital blossomed in Silicon Valley, culminating in the peak of the late '90s. It's funny. We joke today that it's still called Silicon Valley, but there's almost no silicon left in Silicon Valley, right? There's very few chip companies that get funded anymore. But at the time, pitching a chip company while it was expensive and they needed to raise a lot of money by those days' standards-- VCs loved chip companies. Because they felt that there was a lot of differentiation. There was IP protection, high-barrier to entry. So at the time, there was definitely money available for innovative companies that were focused on chip technologies. Probably at that time, you saw a lot more chip companies than software companies.

**Campbell:** Well, I think you're right on the-- I mean, Chips and Technologies, S3, there are a whole bunch of companies that have been very big successes. And so I think the VC community, which tends to herd a little a bit, hadn't gotten out of the herding phase yet.

**Hodge:** Have they ever?

**Campbell:** No, probably not. But the other side of it is that doing a chip company, hadn't gotten super expensive yet. And so I think today, there's hardly any appetite at all in Silicon Valley for chip start-ups. And probably rightfully so, because they're just so damn expensive. There aren't many markets that can support the cost of chip development today. So that was a different time and place. And I think everybody was still optimistic there were going to be things like the previous wave of success.

**Hodge:** And you guys were slightly, probably about three years, before the internet wave really hit. And everyone was doing a web start-up I imagine. Is that roughly correct?

**Sellers:** More networking start-ups probably. That was our big competition. This was after we'd gone public and were trying to continue to grow the company. But '98ish time frame, '99-- just keeping our engineers, recruiting engineers when every other building was yet another networking start-up company that was being sold for a gazillion dollars. So that was very challenging at the time.

**Hodge:** And what were your guys' titles in the company? Gordy was CEO, correct?

**Campbell:** Until we went public. Then we brought in another guy to take from that stage on.

**Smith:** I was VP of sales and marketing.

**Sellers:** I think both Gary and I were CTO and Chief Scientist.

**Tarolli:** Something like that.

**Sellers:** Something like that, yeah.

**Tarolli:** CTO Software, Hardware.

**Sellers:** Yeah.

**Tarolli:** Or some breakdown.

**Hodge:** And so where did the IP come from to design these chips? It seems sort of hard to just have two or three guys get together and design that graphic accelerator.

**Tarolli:** I would say that the basis for the algorithms had been around for 20 years or 30 years in the flight simulation market. And SGI built them into workstations. And so how to do 3D graphics was sort of old hat. The problem is how do you do it cheaply.

So in terms of the algorithms, there wasn't any new things invented in terms of the implementation of how you actually do it inexpensively. I would say that's where some of the innovation came from.

**Sellers:** I'd say that was the simulator we mentioned before. This thing was just all software that Gary created was the sort of the research part of how we would develop all of it. And Gary would map an algorithm the right way.

**Tarolli:** Right.

**Sellers:** I remember it had all these different--

**Tarolli:** Oh, yeah. Do it the right way.

**Sellers:** --flags, where you could do full floating point calculations and do everything kind of the SGI right way. And then Gary would use that as a kind of apples to apples comparison against OK, here's the cheap way. And since this was all about gaming and consumer use, there wasn't a perfect answer, right?

Because it ultimately comes down to does it look good enough? And that's very subjective. And so you could really go to the extreme of when you can start visually seeing artifacts and visually seeing something that's not quite right. And then you just kind of come back slightly from that. And that was a constant balancing act--

**Tarolli:** Yeah.

**Sellers:** --where we tried to make it look good enough. Everything had its limits. And so even in some of the early chipsets, developers couldn't do anything that they could do on an SGI workstation. For example, the depth of the scene, because of the Z-buffer limitations and other things.

It was constrained, but it was perfectly acceptable for games. It probably wouldn't be perfect for high end flight simulators or SGI-like things. But I would say from an IP prospective, a lot of it was Gary just sitting in the simulator--

**Tarolli:** Seeing what you could get away with.

**Sellers:** --creating these algorithms. Yeah.

**Tarolli:** And asking Scott, can you build this? How expensive? Because how do you do a divide? It's well-known how to do division.

But how do you cheat? Well, do a table look up. Scott, how big is a table? What does a gate count look like if I ask you for a 12 bit table?

And then we're going to do some linear interpolation between it, like you learned in high school trigonometry. You get our sine tables out, and you literally interpolate. That's all. So that's where some of the IP came from.

That's how you basically do a cheap divide. And so Scott would go back and forth. It's like, OK. For 20,000 gates, what can we do? And then divide, right?

And how does it look? What happens if I shave a bit off, right? So he would do the cost end of it.

And I would try to figure out the algorithms. And then we'd look at the results, and say, well, what do you think? Well, I can see the difference. But it really isn't bad.

**Smith:** I think the key that we had is we had a mission. And the mission was to build this game thing. And so you'd look at the pristine method or the most precise method, which was full floating point and 24 bits of Z or whatever.

And you start dialing it back saying, well, what can we get away with? And knowing that mission, saying the mission is for games, and knowing that cost was really important, was equally or of greater importance than having all the precision that you would normally need, that gave us insight as to how to proceed. And [the lack of that insight] failed a lot of other companies.

Because they would try and bring a purest 3D graphics mentality into building their chips. And you couldn't get there from here. Because we sacrificed all kinds of things for cost. And ultimately, we'll probably talk about this.

We sacrificed 2D, which was very heretical at the time. We left it out of the chip. And we were told by many, many experts in the field you can't sell a graphics chip without 2D. How do you run Windows?

And our answer was... "you don't". We were only there for games. And if you want to run Windows, buy that Chips and Technologies [or S3] chip. It's a great chip for that.

**Campbell:** The interesting thing about that from a chicken and egg standpoint is we were all about 3D games. But there weren't any.

**Smith:** That's a minor detail.

**Hodge:** So one thing you brought up, Gordy, was that VCs at the time still liked chip companies, because they perceived there was a high barrier to entry. But if a lot of these graphics techniques were already known and it was just implementation you guys were doing, what gave you the confidence that you would be able to beat out the new entrant into the market and stay successful?

**Campbell:** Well, I don't think anybody else was doing what we were doing at the time. The only other 3D graphics guys were taking the CAD approach. And so there really wasn't anybody doing this.

**Hodge:** What was Nvidia doing at the time or ATI?

**Sellers:** So Nvidia was founded in the early '90s, '92. They were around the Pellucid time frame, as I recall.

**Campbell:** A year before us I think.

**Sellers:** Yeah. Their founders were predominantly from Sun, Sun Microsystems, the graphics group of Sun. And then Jen-Hsun, the CEO, came from LSI Logic. The architects from Sun had this very different way of rendering 3D that was-- what do you call it?

**Smith:** Forward texture mapping.

**Tarolli:** And quad patches or something.

**Sellers:** Quad patches or something. And it was not a technique of development that really anyone knew anything about. So we took a more traditional approach, which as Gary said, there's textbooks on this

stuff dating back for decades about when you're trying to create a 3D object, making it look 3D, you break it down into individual triangles.

And the more triangles you create or tessellate an object into more triangles, the more realistic it looks. Well, there was another technique, which is what they use, which was this quadratic patch thing.

**Tarolli:** Yes. And some were forward texture mapping. It just was not terribly user-software friendly.

**Sellers:** Not very intuitive.

**Tarolli:** Yeah, and not very intuitive unless you just wanted to write a game for that or an application. But there were some existing 3D libraries, like SGI, IRIS GL, and other things that were sort of in a different paradigm of how to do 3D graphics. So we stuck with the more traditional one.

**Sellers:** So when Nvidia was trying to do their first chip, they were trying to be more of a multimedia card as opposed to 3D. So they had 2D Windows acceleration. They had audio acceleration or audio capabilities.

And then they had this 3D like capability. And that was what was in what they called the NV1, the original Nvidia chip. And ultimately, the market was not particularly favorable to that. It was sort of a tweener.

It wasn't the best graphics. It wasn't the best audio. It wasn't the best of anything. And they went down the direction that just ultimately was not the right one for the time.

**Tarolli:** I think they were a dead end I would sort of [INAUIBLE].

**Hodge:** So did you guys have-- what came first, you guy's first sale of your product or your first product being finished?

**Sellers:** We actually had a purchase order before we had a product. In fact, we had this thing called the monkey. The monkey would move around the building a little bit. And we did have a purchase order, I think it was from Orchid, before we had chips.

**Hodge:** This was a physical stuffed monkey?



**Smith:** Yes. Physical stuffed monkey. And I think it might have moved to either Scott's or Charlie's office by that point in time. But we had a purchase order. And then we had chips. And then we had-- you all probably remember the bring-up process better than I do. But I remember it was a Super Bowl Sunday.

**Sellers:** It was Super Bowl.

**Tarolli:** Super Bowl Sunday.

**Sellers:** Yeah. Gary and I in the lab--

**Tarolli:** One of those proverbial all-nighters I think.

**Sellers:** The architecture that we ultimately created was very different than the traditional consumer graphics chips in that we ultimately decided that a two chipset was the right way to go. And in terms of competition, you had the very, very high end guys that could not figure out how to build something cost-effectively.

And so they were, almost by cost, just necessarily out of the equation. Then you had the very low end guys that were coming from the 2D graphics world, this Windows Accelerator space. And this is again the S3s and the Weitek's. And they were trying to take a fundamentally 2D engine and kind of bolt on 3D.

And it was kind of pseudo-3D. And we found the sweet spot in the middle. But what we found was that by building a two chipset, you actually could have enough gates to pull it off what we're doing.

**Tarolli:** And memory pins I think.

**Sellers:** Yeah.

**Tarolli:** Memory speed, at that time, was very slow. And that was one of the keys that totally efficiently used memory.

**Sellers:** That's right, yeah.

**Hodge:** So were these custom fabricated or were they FPGAs [field-programmable gate array]?

**Sellers:** Fully custom.

**Hodge:** OK.

**Sellers:** Yeah.

**Hodge:** Did you have your own fab?

**Sellers:** What's that?

**Hodge:** Did you have your own fab?

**Sellers:** No. As Gary said, that's kind of what Gordy really had a tremendous amount of experience was that you don't want to get into the fab business. And Gordy had connections with all the major third party fabs.

And TSMC ultimately was our fab of choice, which continues to be the world's largest third party fab I believe. So Gordy really helped steer us down that path, so we didn't have to waste all of our time figuring out what our fab strategy was going to be and having to reinvent the wheel in that regard.

**Hodge:** How well did that work out?

**Sellers:** I think it was central to our success. I don't think we ever would have been able to build the company as it was without third party fab capability.

**Campbell:** Yeah. No, there was no reason to do it any other way.

**Hodge:** What would be the length of the design cycle, not the design cycle, but just the end stage and having something ready for the fabbing, then getting something back from it?

**Campbell:** Too long.

**Sellers:** You mean, in terms of when you release it to fab and when you get samples back?

**Hodge:** Hm-hmm [Yes].

**Sellers:** I think at that time, it was about four to six weeks.

**Hodge:** What is it now?

**Sellers:** It's actually about the same.

**Tarolli:** Still four weeks I think.

**Smith:** Costs a little more now. Ultimately, it depends on how much you want to pay to accelerate the fab. But that part of the process actually has not changed significantly in terms of the time. But getting back to the Super Bowl story, so we had this two chipset, and one chip was what we call the frame buffer chip. So that did the Z-buffering and the Gouraud shading. And then we had a second ship that was called T-Rex for texture mapping chip. And T-Rex ultimately was kind of the real secret to our success, the fact that we could do real-time, perspective correct, bi-linear, filtered, texture mapping.

**Tarolli:** 16 bit RGB textures. High quality, so the developers didn't have to jump through hoops.

**Sellers:** I mean, if you can't do texture mapping in a visual pleasing manner, it's almost a black and white thing. You either do it or you don't do it. Like the PlayStation did not do it well. And you just would see it. And while it sort of understands what you're trying to show, but it's very visually unappealing.

**Hodge:** Would one of you guys mind explaining roughly what the graphics process looked like at the time from start to finish, sort of a verbal flow chart of it?

**Tarolli:** Oh, OK. The life of the triangle. So like Scott mentioned before, I realized at SGI that you could do a reasonable amount of geometry processing on the CPU. So the Pentium 90 was out, which did fairly fast floating point for the first time in a PC.

So we used that to do all the typical front end 3D work, which was take your triangles and rotate them and do geometry operations through 4 by 4 matrices all in software, and then even set up the triangle. There's a bunch of calculations you do to find the slopes across a triangle.

**Hodge:** So this is very much like the wire frame portion of the--

**Tarolli:** Yes. It's just three vertices. And you transform them. You compute them in perspective. And you do some other operations to prepare for the actual rasterization, which is filling in all the pixels on a triangle.

At that point, we literally send that data over the PCI bus to the first chip in the system, which was a frame buffer controller, which would sort of do some more processing on it. And then it converts it into scan lines from top to bottom and across, and keeps the other chip busy doing a texture mapping.

It says, OK, I'm working on this scan line. And it's been such a long time, I'm having a hard remembering how they kept in sync. But they basically just went through it.

And the texture mapping guy would give the texture mapping pixels back. And then it put them together and did the frame buffer stuff, hidden surface removal, Z-buffer, and alpha blending. We did alpha blending on the first chip.

**Sellers:** Yup, first one. Yup.

**Tarolli:** We did a lot of things that other people who were trying to do \$20 chips coming up from the bottom would not do. Because it's too many gates, it's too expensive. But we had a million transistors on each chip. It was more expensive, but we could do enough of the bare bones minimum that it could do like real SGI style graphics for cheap.

**Hodge:** So you were using the PCI bus on the card between your chips as well as--

**Sellers:** No.

**Tarolli:** Not between the chips. Just between the Pentium CPU and our card. And then on the chip, it was all [a] custom [interface].

**Sellers:** So there was this frame buffer controller chip, the FBI chip as we call it, the frame buffer interface. That had the PCI connection, so that the CPU would calculate the triangle coordinates and send down commands to the FBI chip. There was a private bus between the FBI chip and the T-Rex chip that just used a proprietary protocol, because we didn't need an industry standard protocol there.

So we created this very high speed-- it's pretty novel, actually, to be able to get the amount of bandwidth out of the number of pins that we had available. And then the texture mapping chip had four memory channels, which was important. Because in order to do bi-linear texture mapping, which is basically when

you look up the texture, you're looking up its four surrounding neighbors and taking a weighted average of those four, that creates a much more visually appealing texture map result.

So again, having those four memory channels on the texture mapping chip was a critical piece of being able to this in a performed way. You would do this look up, do the blending of these four texture elements, and then pass that back to the frame buffer controller that would then render that into the frame buffer.

**Hodge:** So it sounds, sort of the running theme, is that memory was your main limitation throughout.

**Sellers:** Memory bandwidth. I mean, there was absolute memory cost issue. But memory bandwidth was really the major issue. And even in today's world, that's still the same issue.

**Tarolli:** And different interleaves I think. Because people coming up from the PC market had a chip that had memory out there. And that was it. And they just used it for doing all this stuff.

And what was innovative about our design was because this chip wasn't doing traditional 2D graphics, we could do anything with the memory we wanted. And so Scott did a lot of innovative things. You had two different interleaves, two or three for the frame buffer controller, one for the Z-buffer, and one for the color, and in four-way on a texture mapping.

Because we set the bar a little higher. We said we want to do all these nice features inexpensively, but still fast. We're not going to sacrifice too much performance. And that required six different memory interleaves across the two chips, where most people had one. It required additional cost as well, but it was still sub \$1,000.

**Smith:** I mean, our memory bandwidth that we had at the time was possibly an order of magnitude more than other people. And that's why we could achieve things that no one else could do.

**Hodge:** So when you had your first purchase order, how big had the company grown to at the time? And when was this?

**Tarolli:** Well, chips came back during Super Bowl '95 – so it must have been [early '95], right?

**Hodge:** Yes. I'm sorry, I keep interrupting a Super Bowl story.

**Tarolli:** Or no, '96.

**Sellers:** '96 I believe.

**Tarolli:** We certainly got together in '94. I was living on the East Coast, but I would fly out here for two weeks. And Scott would come back there for a week or two. So I was commuting. And they were working out here in various offices and living rooms and such. But then we got funded in January--

**Sellers:** I think it was March.

**Tarolli:** --March of '95.

**Sellers:** Yeah. Yeah, right

**Tarolli:** And then proceeded to kick it high gear, because now we could payoff all our credit card loans that we had used to float the company for six months.

**Sellers:** And hire people.

And we could hire employees that we could pay. And so then we taped out in the fall. October, November, December, it must have been. Because Super Bowl '96 was when the chips came back a couple days before. And then we called Gordy up at 3:00 AM in the morning.

**Smith:** We were celebrating.

**Tarolli:** We said, time to see them.

**Sellers:** Because we had to tape out two different chips, we got the FBI chip back first, the frame buffer chip.

**Tarolli:** Oh, that's right.

**Sellers:** And we keep talking about this simulator that Gary had written. So Gary had this ability to instead of doing all the rendering in pure software simulator, he could nest and strip all that out and send the real commands down to the actual graphics. So when we got the hardware back pretty quickly, we could actually do something interesting with it.

And one of the demonstrations that we used for a long, long time was a flight simulator that Gary had worked on actually at SGI and had evolved from then. We used the flight simulator a lot in these VC pitches to show what this thing can do. So Gary and I were in the lab on this Super Bowl day.

And something was not right. And we were debugging the C simulator. It wasn't sending the right thing to the chip, and whatever, whatever. Finally, we were doing it step by step.

And you could almost see individual pixels being drawn or individual triangles being drawn. And we said, well, you know what? Maybe we actually have this right, just let it run. And Gary removed all the debug statements or something. And the thing just--

**Tarolli:** It just worked.

**Sellers:** --took off. And you could see the flight [simulator] on the screen. And whole thing was like, oh my god.

**Tarolli:** It works.

**Sellers:** It's actually working. And that was a very exciting time for us.

**Hodge:** Do you remember how big you were by that point?

**Sellers:** Pretty small still.

**Smith:** We were still tiny.

**Sellers:** 12, 15 people maybe.

**Smith:** Yeah.

**Campbell:** It's interesting, we had been using simulations to show what this was going to be when we eventually got it. So it was hard trying to tell customers as well as investors and everybody else what 3D on a PC was. They were used to these big SGI workstations and, to a lesser degree, maybe the PlayStation.

But people just didn't know what it was. And so after we got the chip back, I said, OK. What we need to do is we need to program up a real-time 3D environment. And the thought process was that we would actually do a presentation in 3D.

And so I still remember the first attempt at this was we basically simulated a cubicle. And basically, we were running a PC on a projector. And I could fly around the cubicle in real-time. And we turned the cubicle into a financial presentation.

So I think it was the Hambrecht & Quist financial conference. And I still remember this, because there was this great comparison. McCracken, who was running SGI, was upstairs in the ballroom. And he had an \$80,000 SGI workstation giving a PowerPoint presentation, which incidentally crashed about three quarters of the way through.

I was down in the basement in a little tiny room with a PC and a projector and our first chip on a card. And I was flying around this cubicle in real-time and giving a presentation. People were kind of blown away.

They were all going, you've got to see this and dragging people in and all kinds of things. But the comparison between this little tiny start-up who was actually doing something really neat in a basement versus SGI, who was just doing another PowerPoint up in the ballroom. You think about that and you go, you know, things aren't always what they appear. It's pretty amazing.

And then we took that even a step further. After we had been fairly well along on the revenue ramp and we were taking the company public, they cubicle had changed. The road show that we went on the IPO tour with is they have folding screens that are 8 feet by 8 feet that you can pretty easily pop up for trade shows.

So we took three of those, so we had a screen that we set up in the front room that was 8 feet high and 24 feet wide. We had three projectors that were all synced together. And we created a simulation of our plant.

And so the way this worked is you drove up in the parking lot in a red Ferrari. And because it was 8 by 24 feet, it looked like it was actual size. And then you got out, and you walked in.

You walk through the doors. You're in the lobby. And the environment went right with you. And then every room in the plant was a different part of our pitch until you got out in the back and you saw all the 3dfx boxes being loaded up and shipped out.



Now, when we did this in Tokyo, because we did a tour through Japan, Europe, and the US, the way we closed it is on a 8 by 24 foot screen, we had one of these giant wheel space stations that rotates. And this thing was just rotating slowly up there. And space ships were coming and going out.

And everybody was going like this. Now, in a financial presentation, what happens usually is you have people come running in late. You give the presentation. And lots of times they leave before you're even done.

So they're all sitting here watching the space station. And I got up and I said, thank you very much for coming. The presentation is over, thank you very much and walked off.

And they're all just going-- and I said, the presentation is over. Thank you very much. And they were still sitting there like this.

And finally, I said, OK. Kill the projectors. And everybody was going, what happened? Where'd it go?

But it was probably the most unusual roadshow I've ever seen. And it was graphically dramatic. I mean, you could not help but understand what our product was when we started doing presentations like that.

**Hodge:** So when you guys started with the roadshow, were you still in the arcade mode? Were you in the arcade mode but going toward the PCs? What was the status?

**Smith:** Something happened along the way, which was the price of your memory fell through the floor. When we were first doing the business model, I think we anticipated the initial graphics boards were going to be like \$1,000. And by the time we had silicon, we were using EDO memory, right?

**Sellers:** Yup.

**Smith:** It had gone from \$100 a megabyte or something-- I can't recall what the numbers were-- to something like \$25 a megabyte. Hey guys, we could actually have a consumer product here. And by the time we got to market, what were the Orchid boards, maybe \$250 or \$300?

**Sellers:** Yeah.

**Smith:** It became very possible for us to be in the consumer market much earlier than we anticipated. And so by the time we went public, our revenues were dominated by consumer business not arcade business. The arcade was a very important component, because it legitimized what we were doing.

**Hodge:** What were your first customers in the arcade space?

**Smith:** I think it as Midway wasn't it?

**Sellers:** Midway. We had Capcom.

**Smith:** Taito.

**Sellers:** Taito.

**Smith:** Konami.

**Hodge:** Do you remember what the first game was out that had 3dfx?

**Smith:** Was it San Francisco Rush?

**Tarolli:** SF Rush was one.

**Smith:** San Francisco Rush.

**Tarolli:** Gretzky Hockey, was that an early one? I think so.

**Smith:** NFL Blitz.

**Tarolli:** NFL Blitz.

**Sellers:** Midway did a lot of games.

**Smith:** Yeah.

**Sellers:** I still see them. You know, Shoreline Theater over here, sometimes you'll see a Midway game that still has 3dfx.

**Smith:** Arctic Thunder was a good game.

**Sellers:** Yeah.

**Smith:** Hydro Thunder.

**Tarolli:** Hydro Thunder.

**Hodge:** Arcades were still fairly viable at this time, correct?

**Sellers:** It was a \$6 billion market when we started.

**Smith:** At the time.

**Sellers:** But it declined quickly.

**Smith:** It became a \$0 billion market by the time I had a business focused on that-- note to self.

**Sellers:** We used to say that the arcade market was bigger than the movie market.

**Smith:** It was at the time.

**Sellers:** But it was changing dramatically. The PlayStation and the consoles just killed it.

**Smith:** The fact that you could get good equivalent content at home-- and the arcade guys didn't keep up. They were still really focused on being hardware guys. And they should've transitioned into being software companies. Because [software] was their value proposition.

**Hodge:** So what did you do when you started moving into the PC market? What did you do to get applications on the platform?

**Smith:** That's possibly, at least in my opinion, the single most important thing that we did, which was connecting with the developers early, early on and evangelizing them. And that was very radical for a graphics company to be doing that. Because normally, you would just use whatever API Microsoft published or whatever.

And you'd build this hardware and hope for the best. We couldn't do that, because there were no APIs. Actually, we approached Intel.

They had a rendering library called 3DR. And it was a great library. But we told them, look, we need this for our graphics board. But we have to run DOS.

Because in those days, Windows had really become mainstream. But games were all still on DOS. And we said we have to have DOS. And Intel said, well, you can have it. But you can't run DOS.

And who was it? I think, did we talk to at Microsoft? [The Argonaut or BRender folks?] We probably talked to every API guy out there. We couldn't do OpenGL, because we had no CAD. And so we crafted our own API.

And it was really simple. We called it GLide. And we went out and pitched it to game developers. Before we even had chips, we would go on the road. And we'd take the rendered videos and we'd show it to people.

And [the game developers would] go, "yeah—if you could build that, we might be interested." And actually, we bought a couple of RealityEngines. If you all remember that.

**Hodge:** From SGI?

**Smith:** Yeah, we did.

**Sellers:** Gordy loved that. Quarter of a million dollar spend.

**Smith:** But we built an emulator basically. We throttled down the performance of the SGI machines to match what we could do in the real hardware.

And we gave that to a couple of key developers. And we had one internal one as well. And by the time we had silicon, we had content. We had software that we gave developers. We did all kind of developer stuff.

**Sellers:** We did a lot of things, I think, kind of the right way, which now seems sort of obvious. But at the time, the chip companies were not developing their own APIs. As Ross said, [at the time] it was largely just sort of build to what Microsoft says you should do.

And so we created our own API. We got sort of rave reviews about the level of interface. It wasn't super high level. But it wasn't so, low, low that you're programming registers and things like that.

So GLide was a very popular interface that people used. But we also hired graphics artists and demo guys, basically game developer types that worked for us that would just crank every day some new demo. And as Gordy mentioned, those were the guys that did this great financial presentation.

But what we ultimately delivered when we had the hardware was a developer kit that had, of course, all of the ways for the developer to interface their games to the hardware. But the set of demos that they could just see, what that showed them was the possibilities of what they could do with the hardware. And I think that was really, really important to show the possibilities.

Because you can't just read about this and say, oh yeah, and if you do texture mapping, then you can do this. Until you see it, you just can't appreciate it. So I think that was a great decision to spend--

**Smith:** We spent a bunch of our money on that and time. There's an interesting anecdote. We went to Japan on one of our first trips.

And [in Japan at the time] it was all about fighting games. And we didn't have a fighting game demo. And I think it was Scott and me. I think Scott and I were the road warriors at that time.

And we came back, and we said we have to have a fighting game demo. And how do we do that? We actually built one that became very, very popular with our developers.

And some of our game developer guys went on to become really a big deal in the game industry. The Anubis on Gary's shirt was this game demo that we did called Valley of Ra, which had a reflection mapped Anubis, so you could see the environment around him as he's fighting.

What was he fighting? Something. Some-- I can't recall now. But it was a really cool game demo. And it lit up a bunch of developers including the Japanese. I mean, [with our demo], everyone knew that you could do a high end fighting game as good or better than anything Sega had on the Voodoo Graphics chipset.

We had flight simulators. And we had underwater scenes and all kinds of things that showed the special effects you could do. And the thing that we like to say about the architectures is that there were no apologies. You could turn everything on within our chip, and it would still sustain high performance [because of the memory architecture]. [We designed the chip to operate at full speed with all the effects turned on, which was unusual back in the day.]

And everyone else's chip, because it didn't have the memory bandwidth, if you turned on this or turned on that, the performance would degrade. And in games, if you can't sustain the performance, 30 Hertz or 60 Hertz, depending on what the game is, it's not playable.

And so [with a competitor's chip,] the fact that you have all these features that you can't use [because the performance degrades], who cares? So I think that [having full performance with all the features turned on] was really essential. Showing the game developers these features and them knowing that they can use them without penalty became very important to us.

**Hodge:** How important was it to your performance that you decide to build a card that was 3D only?

**Sellers:** I think it would have been exceedingly difficult to try to do both. And that was the big value differentiation for us for a long time, not worrying about all the 2D stuff. And as we said, the 2D guys that were trying to come up from that kind of legacy architecture and then add 3D, we would clobber them.

They didn't have enough gates. They didn't have enough memory bandwidth. And we loved to do side by side comparisons, because it was just a joke. Most of our side by side comparisons were actually on this \$250,000 RealityEngine that we bought.

And this Valley of Ra demo, we had it running on both. And you would sit and stare at this thing, and you could not tell the difference. It was really remarkable, some PC graphics board that was that big compared to some monster RealityEngine. And that was a really cool thing to show.

**Hodge:** Did you have difficulty selling-- I understand the original version of the 3dfx card physically sort of "clanked" when it was flipping relays--

**Smith:** That was a feature. Because you knew it was time to have fun. To Scott's point about the 2D, 3D thing, part of that was because internal to the team, we didn't have a strong VGA capability. And we ended up licensing that I think, didn't we?

**Sellers:** Eventually, yeah, we did.

**Smith:** But the problem with that is that you don't get paid for it. If we put 2D in our chip and Word Perfect 2.6 didn't run flawlessly, people would return those boards to Fry's... because Word Perfect doesn't work. We wouldn't get paid an extra dime for that.

Whereas, if we focused on games, it's a pure value play. And so people would buy our board to play Quake, or Tomb Raider, or Virtua Soccer, or whatever it was because that game really made them happy. And so that clunk thing was the 2D to 3D pass through [circuit]. We had a relay on the board.

And you would put your 3D graphics board, your Voodoo board, into the [PCI] slot. And you'd cable the VGA cable over to it. Then you'd run a cable from the Voodoo board into the monitor. And when you played a 3D game, there was this click.

And the 3dfx logo would spin. And it was time to fight something or chase something or fly something. And people loved that.

**Hodge:** Choose this spot to take a quick break.

**Tarolli:** Sure.

**Smith:** The point I'll make here is about scalability, which is we took it to new levels.

**Hodge:** All right, we're reconvening after a break with some props this time. The panelists will describe them at the appropriate points: one of the original 3dfx boxes, [a] quad board system, another board in the back I can't see, and apparently an iPhone, which I didn't realize was 3dfx.

**Smith:** No. This has about the performance probably of Voodoo 2 maybe?

**Sellers:** I don't know actually what's in there.

**Tarolli:** Some PowerVR.

**Smith:** It's a PowerVR chip.

**Tarolli:** I've lost track, yeah. Or more. I mean, it's amazing--

**Smith:** It probably has better geometry performance.

**Tarolli:** --where things have come.

**Hodge:** So one of the things we sort of skipped over, PCs at the time were transitioning from, I believe, it was the ISA bus over to the PCI. And I think VESA--

**Smith:** VL-Bus.

**Sellers:** The VL-Bus. The VESA Local Bus.

**Smith:** Dreaded VL-Bus.

**Sellers:** Yeah.

**Hodge:** So how did these buses impact you guys?

**Smith:** Given that we had a software geometry pipeline, we had to take the geometry work that the processor had done, and move it over to the frame buffer chip, which meant we had bus traffic. And the PCI bus at the time was a great inflection point for us, because it was significantly faster and lower latency than the precursor buses.

And that was another thing that made it important for us at the time. We had a faster processor that could do floating point, which was important to us. We now had this better bus architecture, which was important to us as well.

And then I think the third thing was this Windows 95 transition, that happening about then, which meant a lot of consumers were buying PCs. And these PCs had ready made sockets for our stuff and a pent up demand for cool games. So I think the timing was very good on when we came to market.

**Hodge:** And Windows 95 was the first version of Windows that could actually conceivably run a game without choking, correct?

**Smith:** Well, people would still argue about that I think. But there were successful Windows 95 games. But DOS was still very important at the time when we were first coming out.



**Sellers:** I think Windows 95 just really drove a lot of new PC sales that had PCI buses on it. And that was a big deal-- I remember great debates when we were designing the first Voodoo about what buses we need to support. Do we need to support ISA, or VL-Bus, or PCI?

And those are always the hard things when you're developing new technologies. How far back do you need to support? Because you can really get bogged down in supporting legacy. And supporting PCI only was I think a really important decision.

Because it enabled lots of other things. It reduced our market initially, because not every computer had a PCI bus. But pretty quickly, that got taken care of, largely because of Windows 95 driving new sales.

**Smith:** It becomes one of those things that you remember now, which is in a start-up, often times what you don't do is every bit as important as what you do do, and leaving things off, like this legacy support for ISA or VL-Bus [was one of those things]. Because I'm sure I was probably arguing that we should support it, because we're going to lose this installed base. And the engineers would be saying, yeah, but the future is this thing. And now we know.

**Hodge:** Was that hard with your investors? Because it sounds like you made several decisions like that to leave legacy things behind. You left behind 2D and you left behind buses. Were your investors happy with your decisions? Were they involved enough to care?

**Campbell:** By the time we chose the initial investors and then added to that, I think they were pretty convinced that we knew what we were doing. I don't remember them ever being overly involved in what we decided to do product wise or market wise.

And by the time the products have launched, we were on a pretty healthy growth ramp. So investors always like that. So if you're delivering, they're happy people usually.

**Hodge:** And Gary, you had mentioned some point on software you wanted to--

**Tarolli:** Yeah. Just that when we started, there were no real 3D games on the PC. And so it took a while to get going. I was trying to remember some of the early game companies that we had been talking to, like Papyrus and id Software that had software renderers.

And some of the transition was made easy due to our software library, which was easy to use. But to get them to use it, they'd go, what's your installed base? Well, we don't have any cards out there really.

**Sellers:** Trust us.

**Tarolli:** And then the board vendors would go, how many games run on your stuff? Well, none yet. Take advantage of it. So it was hockey stick growth getting going. But as soon as you get the first couple, then it starts picking up.

And so there was a transition period there where it was very slow to start with. But we did something we with the Doom engine at the time. We sort of ported it over to our hardware early on ourselves and got it running. One of our engineers was just a fanatic about the software.

So he figured out how to get it running. And then we got Quake on our hardware very early. And that was a big selling point.

**Smith:** The other thing we had going for us, though, is that we did have arcade wins by then. And so the fact that Midway was going to use our technology for arcade games, gave us this air of legitimacy that meant this must be good if Midway's going to use it. And then we had some Japanese companies come on board as well. So that really helped with the PC game developers.

**Hodge:** So did you have a formal relationship with id Software, since they were one of the most popular developers in that time frame?

**Sellers:** It evolved. Early on, we were chasing John Carmack and id. It was clear that we needed a killer app to really drive the sales. And Doom came out while we were still developing the original Voodoo 1.

And it was just a runaway hit. It wasn't really 3D per se. It was kind of 2 and 1/2 D as we like to say. But it was so popular amongst the consumers and really created this kind of first person genre shoot-em-up game.

And we really wanted to forge a relationship with John Carmack, who was the founder of id. And there was a lot of hype behind this second game called Quake. And Quake, when it first came out, was pure software rendering. So it didn't take advantage of any of the hardware acceleration.

And as Gary mentioned, well, pretty much all of our engineers were gamers. They played Doom and Quake into all hours of the night. And that was a big part of our culture was sort of this gaming culture.

But John Carmack was always very good about releasing parts of the engine and some capabilities to allow, in essence, the Quake community to extend the rendering platform. And so one of our engineers

actually figured out a way to create a renderer for Quake, so at the very least, we could actually show John Carmack what Quake would look like if it indeed took advantage of the hardware.

And that was a huge step for us. Because once he saw it, then his mind just started cranking. And he very quickly went from being kind of a software rendering purest, to one saying, hey, I don't need to do all this stuff anymore. I can take advantage of the hardware. And once Carmack did the version of Quake, which supported Voodoo 2, that was--

**Tarolli:** That was it.

**Sellers:** Our sales just went crazy.

**Tarolli:** I think he converted it to OpenGL. You know, he had had a major financial success with Doom early on. So I think he took a financial risk, but very forward thinking, and said, I'm just going to write to OpenGL.

And this is my minimum level. There's at least one piece of hardware out there that does this. Everyone else is going to come along sooner or later. And for him, he didn't care it took a year or two.

I think he could afford it and whatnot. So he sort of, I think, set the bar for what hardware should do. And that helped us tremendously. Because it looked gorgeous on our own hardware. And it was sort of a perfect fit for our hardware.

**Hodge:** So that was one of the major applications on the platform. Was Tomb Raider the other major one?

**Smith:** It was a very popular game, too. And it was beautiful. It was more, in some ways, family friendly I guess.

**Hodge:** That's a very relative thing to say about Tomb Raider.

**Smith:** Yeah. But there were a number of games and people had their favorites. In the hardcore gaming community, Quake was the thing. When Quake came out, it drove everything. And Unreal came out after that, I believe.

**Tarolli:** Right. And once there's one or two of them, then the other guys are looking and go like-- the lead game developers don't want to come out with something that now looks worse than all their competitors, right? So they start pushing themselves, I think, to look better.

So it very quickly went from no hardware support to quite a lot. And nowadays, it's like everyone assumes that there's hardware out there, it's all going to run in hardware. This happened probably 5 or 10 years ago.

But it took a little while, a couple years for it to catch on. But then when it did, we saw the results. You see the results. So it was worth it.

**Hodge:** So is this box on the table, then, the first Voodoo card you released?

**Smith:** I think it might have been.

**Tarolli:** I think it was.

**Hodge:** Doesn't seem to have the Voodoo branding on it.

**Tarolli:** No, it doesn't. But this was Orchid, right? That was before Voodoo, right?

**Sellers:** Yeah, this was an original. This is a Voodoo 1. Orchid was the name of the board company. We have our little logo on the back. But this is before we really developed-- oh, I guess it's on the front, too-- the 3dfx brand.

So we were just some largely unknown chip company at that point. And as things evolved, especially with the Voodoo 2, we got more of the mainstream board manufacturers, for example, Creative Labs and Diamond Multimedia. They were truly the major players in the graphics board space at that point. That's when we our brand was much more pronounced and a bigger part of their offerings.

**Hodge:** So how long between the Voodoo 1 and 2?

**Tarolli:** About a year?

**Sellers:** Yeah, about a year. We did it pretty quickly thereafter. Changed silicon process, so we got a big speedup. I think the first Voodoo 1 ran at--

**Smith:** 50.

**Sellers:** --45 or 50 megahertz. And Voodoo 2 ran at 90 plus.

**Tarolli:** Same with the memories as well, right?

**Sellers:** Memories all speeded up.

**Tarolli:** Half of the stuff that was done on a CPU we did move into hardware. So that got another big speed up.

**Smith:** Big geometry speed up.

**Sellers:** That we demoed at COMDEX. When was that? Maybe late '97?

**Smith:** Is that the million triangles per second?

**Sellers:** That was at GDC.

**Tarolli:** That was the Voodoo 1, I think.

**Sellers:** We first showed Voodoo 1 at COMDEX. And I remember Creative Labs wanted the exclusive-- remember that-- to show it.

**Smith:** Oh, yeah.

**Sellers:** And Creative Labs at the time was just a monster huge, huge company and very, very widely followed. And they were kind of the marquee upgrade company at the time. And they had this massive-- COMDEX at the time, was a big deal. I mean, I don't know, how many people would go to a COMDEX?

**Smith:** Oh, 100,000 people or something. It was crazy.

**Sellers:** And almost the entire Creative Labs booth was about Voodoo 2. And we had this demo showing Quake running on a Voodoo 2. And we would have people stacked 20, 30 people deep looking at these videos, just goggling over how you could actually do this. It was very, very cool to see.

**Hodge:** Now, I understand some of your demos and trade shows actually were using, I think you mentioned this before, the SGI hardware to emulate yours. Were you away from this at that point and using your own?

**Tarolli:** Yeah.

**Smith:** Oh, yeah.

**Sellers:** Yeah. That was all pre-Voodoo 1 stuff.

**Smith:** That was before we had hardware.

**Tarolli:** That must've been '95.

**Sellers:** '95.

**Smith:** The chip was probably a little behind schedule. And we had a COMDEX event. Whose booth were we in? We were in Orchid's booth?

**Sellers:** Yeah. That's right.

**Smith:** I think it was Orchid's booth.

**Sellers:** Yeah.

**Smith:** And Orchid was next door to Diamond. And it turns out, several of our colleagues from Pellucid were at Diamond at the time. And we didn't have silicon. So what are we going to do?

So we said, well, let's put an RE [RealityEngine], and it'll be the Voodoo graphics simulator. And we were very clear that it was a simulator. And we showed an NV1-- I think it was an NV1 at the time-- running like a Sega virtual fighter game, which looked awful.

**Hodge:** And they had an NVIDIA video card at the time?

**Smith:** Yeah. And then we had Valley of Ra, which was our environment mapped-- and you can find videos of this on the internet-- game demo written on the RealityEngine. And the results were remarkable. People said, well, why would you buy that [NV1] when you can get this thing [Voodoo Graphics]?

And all the Nvidia guys were just screaming about, well, it's not the real thing. Yeah, but it's a simulator, it's fine. And it turns out when we brought out the real chip, we ran the thing side by side. And as Gary said, you could not tell the difference. In fact, in some ways, the real hardware actually outperformed the emulation mode that we had implemented on the RE.

So I mean, it was very accurate. And really what we did was probably a little close to the edge in terms of "guerilla marketing"-- Gordy approved it. He thought it was a great idea.

**Tarolli:** But the good thing is we delivered what we promised. And that was the most rewarding part of it, right?

**Hodge:** So your company culture sounds a tad bit irreverent. You've named the chips the T-Rex, the FBI. There's a video floating around with an Elvis painting in your headquarters as well as, I think, a Gorn from Star Trek. What was [the] company culture like?

**Smith:** It was fun. It was work hard, play hard taken to the logical extreme. Because we were focused on games. I mean, in our first building, we quickly ran out of space. And so we had like a dungeon that we put the software engineers into and the game developers. Was that like a basement?

**Tarolli:** It was like a shipping, receiving warehouse.

**Smith:** A warehouse or something.

[INTERPOSING VOICES]

**Campbell:** I never turned the lights on.

**Smith:** Yeah.

**Sellers:** It was perfect for them. Because they didn't like the asbestos falling down from the ceiling.

**Hodge:** Is this the building that had birds nesting on the inside?

**Tarolli:** Yeah.

**Smith:** Yeah, yeah. I-- We had a swallow's nest, a barn swallows. It would come every year.

**Tarolli:** Back at the Sunnyvale golf course.

**Sellers:** The eighth hole or something like that.

**Tarolli:** The eighth hole.

**Smith:** We'd sit out there and drink beer. And incoming golf balls you'd have to dodge a bit. But the culture was great.

**Sellers:** Very young. Most of the software engineers were early 20s gamers through and through, loved games. Would code during the day and play games at night. I mean, we would have these company-wide Doom-fests. And at the time, that was very new that you could hook up--

**Smith:** LAN party.

**Sellers:** --multiple PCs through the Ethernet and actually play against each other in these first person combat games. And so we'd have all-company Doom-fests. So it definitely created a culture of fun. And it was all about the game.

**Hodge:** Was that a selling point in competing against the networking companies for talent?

**Sellers:** Oh, I think it was a great selling point anytime you're doing something that is visual that you can actually see and look and see what you're going to be working on. In fact, I remember one of our first salespeople we actually had one of these mini mock-ups that we did.

We took an arcade box and inside it, put a PC running this video loop of Gary's simulator. And this was actually in our lobby of that first office we had in Sunnyvale. And one of our first salespeople, we interviewed her and ended up hiring her, she actually didn't realize until she was there that--



**Tarolli:** Oops.

**Sellers:** --we had no product at that point.

**Tarolli:** It wasn't real as she would say.

**Sellers:** Yes, it wasn't real. But I think it was a huge selling point, the fact the we were consumer focused, and gaming focused, and fun focused. And we had a very unique blend of doing really innovative technology and doing silicon engineering and a lot of software, and mapping that into really thinking we were a game company, thinking we were an entertainment company.

So we did a lot more marketing than probably any chip company had ever done. Ultimately, after we'd gone public and other things, we did television advertisements and we did all these things. And for a chip company in Silicon Valley, that was almost unheard of at the time.

**Campbell:** And everybody really got into the swing of it. I mean, you'd go down into the area that was never lit up and every time you go down there, they'd show you something new, a new demo, a new mapping technique. There was always something. And it was kind of fun. I mean, it was a very creative environment. People loved to try and figure out something new that we could do.

**Hodge:** What was it like to be in charge of a company with 3dfx's culture?

**Campbell:** Well, I think the work hard and play hard thing, that to me was kind of Silicon Valley. You went out and socialized with your crew and had beer fests and did all that kind of stuff. And a friendly environment where everybody knew everybody and everybody was not in a hierarchy so much as part of the group or the team was, I think, the way you did that.

I think the thing that was added here was it's the gaming industry. And that was a whole new twist on it. I mean, if you go to the trade shows, a good example, you'd have guys that would show up at our booth with Dracula capes and pointed teeth. I mean, it was just crazy.

**Hodge:** So when did you realize that you weren't quite in the computer hardware business per se anymore?

**Campbell:** Well, I think we actually were. Because if these guys hadn't absolutely been on the cutting edge, then none of this would have been successful. The trick to this was we went out and we were pretty creative on the customer side, putting a huge developer program in place and things like that.

I mean, if you didn't have a cutting edge technology product to work with, you wouldn't have been this successful. So I think we were a cutting edge computer company in this sense.

**Tarolli:** Just going back to trade shows for a minute, the one that I remember that I was most proud of was a SIGGRAPH in New Orleans. Remember that, right, where we showed the Granblue demo?

**Smith:** Oh, yeah.

**Tarolli:** Right? Because this is not a traditional PC show like COMDEX.

[INTERPOSING VOICES]

**Tarolli:** --workstation.

**Smith:** --thing.

**Tarolli:** Right? And so we brought our best and baddest hardware at the time.

**Smith:** Was that Voodoo 2?

**Tarolli:** I think it was Voodoo 2, where it's showing multiple texture maps at the same time. Because we had designed the chips to be expandable, where you could do two or three of the texture mapping chips and actually do two or three textures simultaneously.

And one of these boards here is sort of built like that. And so at that SIGGRAPH-- I don't know if it was '96 or '97, that summer-- we showed technology that was rivaling SGI's--

**Smith:** High end.

**Tarolli:** --high end. Not quite, but it was pretty amazing though for a little PC.

**Hodge:** So the board that's on the right of the table then, the four boards linked together, was that something you introduced in Voodoo 2, the ability to link boards together?

**Smith:** Well, actually-- and this is an interesting side note as well-- Voodoo 1 had scalability built into it. Gary talked earlier about scaling. We reduced everything to scan lines. And we had included the capability to do scale and interleaving.

So you could actually effectively double the performance. Now, the marketer at the time said we don't need that. Let's just focus on schedule. And get the thing out. And forget that feature.

And of course, the engineers, as they often do, didn't listen and put it in there. And so even Voodoo 1 had scalability. And it turned out to be probably one of our most important features ever. So it shows how much marketing really knows at times. But then we extended that to different levels.

So this beast over here was actually a board done in a spin-off of 3dfx called Quantum3D that built boards for visual simulation. Actually, the old SGI market-- we kind of went full circle-- we went after that.

And so here we took the scalability of the Voodoo architecture to its logical extreme. So there are four SLI boards all SLIed again together to do really high performance levels. And so their product would actually truly compete with a high end SGI machine.

In fact, in the Viz-Sim market, it chased them out of that. So it could do full scene anti-aliasing and all kinds of features. And it just took advantage of the scalability and the horsepower or the underlying graphics architecture.

**Campbell:** So part of the thought process that led to all this stuff was that as we started to speculate on what was really happening behind 3dfx and its market phenomenon is we speculated that the PC would really drive 3D graphics and that companies like SGI would not be able to sustain the pace of development dollars necessary to stay in front. And as it turns out, I think that was correct. We did a spin-out company called Quantum3D to actually go after that particular market more directly.

And I think while we were right about the speculation that the PC would drive 3D graphics, I think we underestimated the complexities of getting into the high end of the market. But it was pretty interesting. I think the PC volume was definitely what drove 3D graphics.

**Hodge:** Ross, perhaps you might want to talk about this. If you're just reading the transcript, there's a box on the left of the table. It says Righteous 3D Orchid. It's the original Voodoo box.

And there's a 3dfx logo somewhere on it, but from 10 feet away, I can't see it. The box on the right is a Voodoo 5 box, where it's 3dfx branding all over it. And I can't [tell] you who the OEM that make it is. How did you go from you guys being an afterthought on the box to the OEM being an afterthought on the box?

**Smith:** I think actually by the time this board came out, we were actually in the board business. So when we started the company, we sold chips and development kits to board makers who would make their own boards. That's Orchid, and Diamond, and Creative Labs and all the other board guys.

We had every board maker, Canopus and all kinds of board makers, making boards using our chips. And we did a good job of enforcing our branding. So it was very small there, but over time it grew substantially. Scott mentioned, at COMDEX we dominated the Creative Labs booth one year. Probably-- was it '98 that we bought STB, or when was that? There, somewhere around there. I had actually left the company by then. I was off focusing on this craziness at that point in time [points to Quantum3D board]. But we acquired a board company and went after the board business in full pursuit. And so this box here is just all 3dfx. It was a 3dfx branded product.

**Hodge:** OK. Was there a plan even before you bought the board company to get your branding out there above that of that OEMs?

**Smith:** I'm not sure above, but certainly equivalent. Because we were the underlying architecture. And what we had as our North Star was Intel Inside®. They had a very successful branding program.

Of course, they had spent hundreds of millions of dollars to do that. We couldn't afford that. But we knew we could drive branding substantially with the power of the consumer that loved 3dfx branding. So shortly after this product came out, the Righteous 3D product, we had a lot of other products come out that had much more prominent branding.

**Hodge:** And GLide, if I remember correctly, was a very big part of the branding you did as well, trying to get things out that said they were written for GLide?

**Smith:** I do remember that.

**Tarolli:** The early days.

**Sellers:** Yeah. I don't think we tried to brand GLide per se as much as we branded "Powered by 3dfx"'s logo on the games themselves. So if you flipped through the gaming magazines at the time, you'd see some hot new game coming out.

And down in the corner would be a logo saying "Powered by 3dfx." So it was very similar to the Intel Inside® kind of idea. But it fed upon itself. The game developers and publishers actually liked it. Because pretty quickly, it was obvious that the hardcore gamers all had 3dfx boards.

And that's the same audience that the publishers were selling to. So the publishers wanted that logo. They wanted to market the fact that the game supported 3dfx. It was a selling point for them. And it obviously allowed us to raise our brand as well.

And that did, as Ross mentioned, culminated in our moving the business model away from just chips to selling the entire board. The thinking there was we could really control the brand. So instead of sharing it with a Diamond or a Creative Labs, we could control everything. And that was where we eventually evolved to.

**Hodge:** When did you guys IPO?

**Sellers:** 1998.

**Hodge:** And what was your main product on the market at that time? Was it still the Voodoo 2?

**Sellers:** The Voodoo 2 was just ramping. And there was another product that we told the public markets about, which was ultimately branded the Voodoo 3. It was internally called Banshee.

And that was our first product that incorporated VGA and 2D graphics and 3D all onto a single chip. And so the public markets were very interested in that product, because that was the continued evolution of 3dfx, getting us into more mainstream, particularly being able to sell to computer manufacturers as an OEM chip. So that was obviously where a lot of volume was.

And we continued to really be focused on the upgrade markets and selling boards through the retail channels for people to upgrade their PCs. But there was a whole other space where Nvidia and S3 and others were playing, where you would sell chips directly to the OEMs, the Dells, the Gateways, and the Compaqs at the time. And Banshee was all about that market. So it was a brand new market opportunity for us.

**Hodge:** How big were you guys by the time [you] IPO'd?

**Sellers:** Well, I remember our revenue went from 4 to 40 to 200 [million]. And somewhere between the 40 and the 200 I think is when we went public.

**Hodge:** Did you have a particular reason for going public, just to get liquidity for your shares because everyone does that when the company reaches a certain size?

**Campbell:** No. I think what you do is you try and do the IPO so you have some cash resources. And that's what we did. I mean, we went public so that we had resources in the company. Because you can't always you count on everything being up to the right all the time. It's like if you look at Apple today. I mean, Apple has more cash resources than almost everybody.

**Smith:** Some countries.

**Campbell:** But they're also a very secure company because of that. They can weather an awful lot of problems. And every succeeding generation gets more expensive and takes longer to do. And if you slip in the development cycle on the next generation or, God forbid, if it comes out and doesn't work, then you'd have some safety in your cash reserves to be able to weather that storm.

**Hodge:** So by the time you guys IPO'd, around that time, you left, Ross?

**Smith:** Hm-hmm.

**Hodge:** [to Campbell] And then you did not continue as CEO after the IPO?

**Campbell:** Yeah, no. We hired a new guy to come in and be CEO.

**Hodge:** What was the reason behind that? Were you just tired of it at that point?

**Campbell:** I was primarily a venture capitalist at the time. And I actually had a fund that we were investing out of. So this was something that I came over to do not as a full-time job, but came over to do in addition to the job I was actually doing. In retrospect, I think it probably would've been better if I had stayed. Anyway, it was never really intended to be a long-term deal for me.

**Hodge:** I'm sorry, [INAUDIBLE]. What was the year again of the IPO?

**Smith:** '98.

**Hodge:** '98. So not particularly long after that, by about I believe near the end of 2000, 3dfx wound up filing for bankruptcy and had all of the assets substantially bought by Nvidia. I'm correct on that?

**Sellers:** Reverse order, but it doesn't really matter. We sold it first. And then we basically sold all the assets of the company. And what remained was what ultimately filed for bankruptcy. Nvidia did not have an interest in being in the retail board market, like we talked about selling through the retail channels.

What Nvidia really wanted was the technology, the chip capability, the engineers, et cetera. So we constructed a transaction. So you sort of sold all of the good assets that they wanted and left some of the pieces of the business that they didn't want. So it was the latter part that ultimately filed for bankruptcy.

**Hodge:** So what happened over those intervening two years that took you guys from the top of the market into basically being sold to your competitor?

**Sellers:** Well, I think probably the biggest decision that we made that ultimately became probably the anchor of the company was this decision to get into the board business. We talked about this at length in the management team and the board.

You hear about some of these acquisitions and decisions being make or break for the company. And we talked very specifically about it being a make or break decision. And ultimately, I think we assumed too much. We started competing with what used to be our customers.

So the Creative Labs and the Diamonds of the world, who used to be our customers, once we got into the board business, they were now our competitors. And so on the one hand, we were competing against the chip guys like Nvidia and S3 and others. And we were then, once we got in the board business, competing with all the board guys. And so it really was a we're-going-to-do-it-all kind of strategy.

And that's a big bet. And when-- just a little bit of slip as we did-- we were a little bit late coming out with some of the next generation products and didn't have the runway to come up with the next generation products, which I think would have been very compelling on the market. But ultimately, we ran out of time.

**Hodge:** Had Nvidia switched to more traditional solution in the graphics space by that time?

**Sellers:** They had. Yeah, they took a very different approach, where, after they realized the mistakes with the original chips, the third generation, NV3 as it was called, was very specifically designed for the OEM space. So instead of worrying about the high end upgrade market and the gaming market, et cetera, we were focusing on they designed a chip to really go directly after the S3s of the world, OEM-centric business model.

And what happened is those models ultimately converged. So they drove a lot of success and were able to grow their company very effectively with an OEM chip model. And we ultimately kind of converged directly competing against them. We came from the high end 3D space and ultimately got into the OEM. And that was where we collided.

**Hodge:** This was also sort of beginning [or] middle of the first internet boom at the time. Did that have any impact on what you guys were doing?

**Sellers:** Well, I think the environment certainly did. I mean, it was crazy in Silicon Valley. And so it was a difficult time to try to grow a company from a recruiting perspective, a lot of distraction. Because there was a lot of other things that were going on in the market outside of what we were doing, networking companies, the internet wave.

We tried to do things through our ISV partners, the game developers and other, to try to do 3D acceleration for web browsing and stuff like that. But that was way too early for its time. So I'd say, really more on the business side, we had to compete against everything that was going on in the internet wave. We didn't really take advantage of that very much.

**Hodge:** One of the things I've seen in a few people's history of 3dfx, claiming that was one of the company's downfalls, was excessive employee perks, which frankly sitting like 100 yards away from Google is almost a laughable idea-- at the time you guys were spending \$50,000-- the number I heard is \$50,000 a month on lunch. Do you think that was in any way an impact on you guys? Or is that sort of a red herring people are throwing down and not looking at the real structural issues?

**Sellers:** I don't remember our perks were particularly extreme.

**Campbell:** I don't think there was anything really in that. I think it was just a really bad decision to get in the board business. I was on the board at the time, and I argued very strongly against that.

But ultimately, the board chose to actually go ahead and do that. And you can't take a high-flying, high-multiple, high-technology, sexy 3D graphics chip company and marry it with a marginal, piece of crap board company and come out of it with anything that's good.

**Hodge:** Well, this would almost be like them buying their own foundry in a way, wouldn't it?

**Campbell:** No. The foundry would have been OK. The problem was that you were in the business of selling chips to all these guys that put them on boards. As soon as they figured out that you were in the



board business, they didn't want to buy from you. You either made the transition very, very quickly and became a dominant guy in the board market, but if you didn't do that, you died.

**Hodge:** I guess then, it sounds like the group at the table was pretty much against the decision that happened anyway. Did you guys continue to stick around for a while after that? Ross was already gone. You had gone back to the venture capital side.

**Campbell:** Yeah. But I was on the board throughout this whole saga.

**Sellers:** Gary and I were there to the bitter end.

**Hodge:** And so Gary, you're at Nvidia now, actually, currently?

**Tarolli:** Correct.

**Hodge:** Did you go to Nvidia as well, Scott?

**Sellers:** I did not. No. I just felt I wanted to do something outside of 3D graphics. It had been a fun ride. And I wanted to try something different.

**Tarolli:** Myself, when I think about it, almost 100 engineers were hired by Nvidia in the subsequent month. And there was another group of about 20 that weren't that went off to various start-ups. But the interesting thing is most of those ended up failing in the next year. So we got to say hi to some old friends again about a year later, some very good friends that came back after a year or two in the market.

**Hodge:** So to what extent then, with Nvidia absorbing so many 3dfx engineers, did following Nvidia products become sort of the 3dfx products that were never released?

**Tarolli:** We tried to think of some ways to marry the technology. But that really wasn't necessary. It wouldn't [have] been a good idea anyway. Because you can't try to bolt two things together. It was just A versus B, and both were good.

But the one thing that they did do and now market is they got the trademarks, all our trademarks for SLI. And so that's still one of their big marketing things. They were one of the first companies after 3dfx to come out with something called SLI.

**Smith:** It wasn't the same SLI. But they used the brand.

**Tarolli:** Right. It was done a little differently. Same concept, I think, technically done very differently. But people still remembered the trademark SLI from these days. It's right on that board there, on this box.

**Hodge:** It was originally like scan-line interleaving?

**Tarolli:** Interleaving, yes.

**Hodge:** And it became with scalable link interface or something like that?

**Tarolli:** Yeah, scan-line interleaving is what it was. And so they used that in their marketing and such. And yeah. So that term still lives on today.

And it's a big part of things. They sell PCI cards with two chips on them, I believe now. They've been doing that for a number of years.

**Hodge:** I actually owned a pair of Nvidia cards in SLI. They were both fast and loud.

**Tarolli:** Yes.

**Hodge:** Very, very loud.

**Tarolli:** Well, we're getting them quieter now.

**Smith:** Some things don't change.

**Tarolli:** Yes. We're now working on a quiet and low power.

**Hodge:** So we've sort of gone through the history from beginning to end and have skipped over probably quite a lot in doing so. So what haven't we talked about that we ought to, people, funny anecdotes, competitors, random name dropping of other intersections with computer companies through the time you were there?

**Smith:** Well, just as a side note, these guys were peripherally involved, but the whole Quantum3D experience was quite interesting. Because we took this PC graphic stuff into this professional market. And we clobbered some of these proprietary guys in the flight simulation market with 3dfx technology.

We changed that industry as well. Some would say that we ruined it, because we brought the PC in that marketplace. But I mean, SGI, which had a multi-hundred million dollar a year business selling image generators to the flight sim community, we put them out of that business with Voodoo graphics basically.

**Hodge:** How did you continue to develop it afterwards? Did you just continue down the--

**Smith:** Well, one thing is the final chip that we used was called the VSA 100. Was that Voodoo 5?

**Sellers:** It was part of this, yeah.

**Smith:** Yeah. And we had built this massive thing called Alchemy, which is another four chip monster like this thing. There's one over there actually. We won every major flight sim, driving sim, sim whatever.

And so the Boeing F15 [simulator] uses our stuff, and the F16, and F18. All these aircraft and ships and everything made the transition to PC-based image generation based on that technology. And it'll never go back. It's all PC-based now.

So that was very interesting. Because a feature that we had put in-- admittedly, that I was initially against, but now fully embrace-- was this SLI stuff. It enabled Quantum3D to actually be successful. And so that was a really interesting aspect. And we rode that wave for a long time.

**Hodge:** How long were you with Quantum3D?

**Smith:** Well, in dog years it would be 70. I was there 10 years or 11 years, long time.

**Campbell:** Yeah, people don't realize, but the PC graphics migrated [to] a lot of different places. And the cockpit was one of them. And it was interesting, because as the cockpits went from 2D to 3D, and then the GPS worldwide database came out, now all of a sudden you could have an entire worldwide database of all the geography, so that you didn't fly into mountains and you didn't do a lot of things. And that's all PC graphics again. It's been absolutely astounding how some of those things happen.

**Smith:** We sold millions of dollars of GPUs to all the avionics companies, Honeywell, Garmin, Bendix. And at the time, we became kind of the dominant GPU supplier into the cockpits of most of the military aircraft, and now commercial aircraft as well.

**Hodge:** So were you still using late '90s GPUs into the 2000s? Or did you switch over to Nvidia?

**Smith:** We switched to Nvidia after Nvidia acquired 3dfx. So in like 2001, we did our first embedded system based on the NV17 chip, which was a great chip. We sold a gazillion of those.

And then we also switched the image generator business over to Nvidia with a product we called Independence, which did the same kind of trick of synchronizing multiple Nvidia chips before NVIDIA even had that. Before NVIDIA had SLI, we were able to do that usually at FPGAs in our own software. So that became very popular as well.

**Hodge:** Did any of your products ever get used in early applications of what would now be scientific or general purpose GPU computing? Was that even thought of then? Or were you guys just gaming graphics?

**Sellers:** It was predominantly gaming. We didn't really have the ability to program the chips, per se, what you would need to have that kind of flexibility in terms of the GPU-like capability today. The product that we were working on at the very end before we sold to Nvidia, we did have a separate geometry chip that we were working on that perhaps could have done some of those types of things. That would have been a very high end geometry chip. That hadn't actually been released to fab I don't think by the time--

**Tarolli:** No. And it probably wasn't 32 bit. Was it 32 bit IEEE [floating-point]?

**Sellers:** Maybe not, yeah.

**Tarolli:** The technology is similar to 10 years before that. The technology was just getting to the point where you could do floating point in these chips at reasonable speed and accuracy to start then addressing these other markets, which had real requirements for like medical simulation, seismic simulation, and all those kind of applications. But now, a lot of the top supercomputers are built out of GPUs now.

So in the last 10 years, it's made that transition. But that was from 2000 to 2013. I guess, the first five years of that, the technology was just becoming feasible.

**Hodge:** I'm going back to movies again briefly. In the late '90s, Pixar became well-known. Lucasfilm, they did the first of the new Star Wars movies with like the most CGI anyone had ever seen in movies up to that time. Was there any overlap at all between that stuff and you guys? Or was it still pretty much entirely separate?

**Campbell:** I talked to Tom Porter.

**Tarolli:** Yeah. We used to talk to them. But it was still pretty much separate. Because they would do things offline for the final rendering.

**Campbell:** We had always dreamed that one of these days, we'd have 3D graphics good enough to have a real-time movie capability, where you could actually exist in this 3D environment. And it was kind of interesting, I remember one of the financial conferences I went to-- this is after Jobs got bounced out of Apple.

And he was focused on Pixar and I think a little bit of NeXT at the time. And I saw him at a table putting out the brochures and stuff like that. So I went over and chatted with him a little while. And I said, you know, you've got Pixar doing some fun stuff.

And we've got this great 3D graphics engine. You know, why don't we get together and do some neat stuff? And Steve always was the kind of guy that was on the cutting edge, doing this, doing that.

And it's the one time I can ever remember talking to him and his response astounded me. He says, well, that'd be fun. But right now, I'm just working on basically our results and trying to make sure that we have performance and good numbers.

And I couldn't believe that this was Steve Jobs saying this. Because it was about 180 degrees from what you'd expect. But no, we never really seemed to be able to get any traction on that side.

**Smith:** But in general, though, if you look at a frame of animation for a movie, those things still today render in hours. We had 17 milliseconds to render a frame. And so what you do in 17 milliseconds is you do the best you can to get the visual complexity that you're trying for.

But when that clock signal hits, you got to go. And so that's very different between real-time and the non-real-time market. You know, the movie market is still not real-time.

**Tarolli:** And the requirements, the technology that they're using to render these films are somewhat different and in different league than the stuff we do for the real-time. Although, now they may be converging somewhat, because of programmability. But they would use techniques and sampling and stuff that we just couldn't do in hardware easy enough.

And it didn't really matter to them. Because they had enough compute power. They would burn 30 minutes per frame rendering it. What did matter to them was the workstations and the CAD stations that they used to design it. Because they'd sit there and wire frame and do it.

And then they would send it off to get rendered. And a couple hours later, they could see if it was good. Now, the difference is they can view it in really high quality, not what they're going to show in a theater, but good enough to get a really good idea of what it's going to look like on the current PCs and graphics, and then send it off for final rendering.

So what they could preview is much, much closer to their final product. Because if you've seen the animations in shows now, it's like, is that from a movie? Is that real-time? Or is that whatever?

It sort of doesn't matter. They can view it good enough now that they can get an idea. So it's vastly helped them. But it's still in the design on their table, not the final product you see.

**Hodge:** At 3dfx, at any point, did you guys do a division, group, whatever that was just sort of focused on research?

**Tarolli:** We had advisory council with some game members on it, some game developers on it. But I wouldn't say we really had a research group. A few of us would go off and think about the next product. But we'd only been through about two product cycles. So we weren't quite there yet where we had enough resources and time to do a research group.

**Hodge:** And I understand you had a video game museum in, I guess, at least one of your buildings?

**Tarolli:** Do you guys remember that?

**Smith:** Yeah, we did.

**Sellers:** Yeah.

**Tarolli:** We had old video games.

**Smith:** We tried to collect as many as we could. I'm not sure what happened to that though. It's kind of a tragedy.

**Tarolli:** Yeah.

**Sellers:** We had a bunch of old consoles and such.

**Hodge:** What prompted you to do that?

**Smith:** We were in the game business. And so knowing the history of games was really important to us. And so I think Alma brought in like a Neo Geo or something. I can't remember. It started a trend. We got a Sega Genesis. And we had this whole display case full of these old video games.

**Hodge:** Why did you feel the history was important? I ask this because some companies in the Valley, most famously Apple, claim not to care about history and almost that it's an anchor on them as they go out to create the future.

**Smith:** Well, if we looked at our customers, the same people that bought that Neo Geo, maybe their children or whatever, those were our customers. And so understanding that what really drove our customers wasn't the technology, it wasn't the SLI, it's the game.

And so having that connection to our customers about the fact that people bought our stuff primarily to play games, I think was a really important aspect of our success. And that was reflected in everything that we did. From game developer recruitment to our packaging, to our naming, to everything was all about being really customer focused. And so those old game systems and arcade games represented an interesting aspect of what are customers were all about.

**Hodge:** Now, about your customers, people hear games and they tend to think kids. That wasn't your customer demographic, was it?

**Smith:** There were some kids. There were some interesting phone calls we got from 12-year-olds. When they would come to the moment of purchase, they would have to go get their mom to give us the credit card number and things like that. But there was a wide range of audience. But mostly, it was kind of the older teenagers and 20-somethings that had the disposable income to buy a beefy PC and upgrade it and had the time and the wherewithal to play Quake all night.

**Tarolli:** Because that was one of the revelations that we learned. A lot of these gamers were actually 20s and 30-year-olds, some of them 40-year-olds at the time that had disposable income and that were just avid gamers. They'd spend \$5,000 a year on building custom PCs with boards. And they still do, except now they're 60 and 70-years-old.

**Smith:** Yeah. And there are companies like Falcon Northwest and Alienware that sprung up building these bespoke PCs for the game market that were tuned in ultimately cryogenically cooled and crazy.

**Hodge:** That's a good point we haven't talked about. We've been talking about boards being sold retail. How much did you get sold into pre-built computers.

**Smith:** We had a pretty good amount of that. Early on, we had a couple of companies that built game PCs. NEC had one. HP had one. I'm not sure how successful they really were in terms of their overall volumes.

But as soon as they saw this trend for game PCs, they kind of jumped on board. And they were the specialists. They were like Falcon Northwest, who really had a passion for building these things. And they're still around. They still make good products today.

**Sellers:** Voodoo 3 was the first product that was really designed for the computer manufacturers. So that was the first time that we actually had real design wins with interesting volumes with the more mainstream computer guys.

**Smith:** Yeah. The other part was kind of a boutique-- it's be in one game SKU that HP had or NEC had.

**Tarolli:** So the really interesting thing is people are still building these boards or trying to. I still get emails from people. Do you have any spare chips left over? Do you have a printed circuit board design? Do you have any documentation on how to write software for them? Because they're still building boards.

**Hodge:** Physically building Voodoo boards?

**Tarolli:** Physically doing them and modifying. They unsoldered all the RAM chips, and replace them with 200 megahertz RAMs upping the speed. These are guys in France a lot of them. I get emails from them about two or three times a year asking for documentation. They wanted to license the design from Nvidia, which owns it at one point.

**Hodge:** Just out of curiosity, why are they doing that instead of emulating it?



**Tarolli:** I have no idea.

**Smith:** I think why is a silly question with this one.

**Tarolli:** It's interesting. Because they can. I mean, you can go on YouTube, you can see videos of old games running on Windows 7 on our boards. So they've done Windows 7 drivers and things, just amazing things with them.

I still haven't figured out why. But it's a hobby I guess. It's a good hobby of theirs.

**Campbell:** I've got a jacket that's got a 3dfx logo on it. Every time I wear it out almost, somebody will come up and say hey, I had a Voodoo board.

**Smith:** It had a real passion. And quite a following.

**Sellers:** It's like a legend in the consumer space.

**Hodge:** I just forgot my question.

**Tarolli:** It's OK, I forgot my answer.

**Hodge:** So skipping ahead to what you guys have done since, Gordy, you've gone more back into the traditional venture capital role?

**Campbell:** Actually, I spent about 15 years on the venture capital role. But now I'm back on the company side again.

**Hodge:** OK. And Ross, you've done a variety of start-ups since then.

**Smith:** Yes.

**Hodge:** Scott you've done one start-up?

**Sellers:** Same one. 70 dog years, as Ross likes to say and still going. Yeah.

**Hodge:** [to Tarolli] And you've stayed with Nvidia?

**Tarolli:** Yup. Easing into retirement and enjoying the benefits of 3D on my cell phones and whatever.

**Hodge:** So would you guys like to talk about the spectrum of how it is that you were happy to stay with a large company? You've done one start-up and you've been sort of a serial entrepreneur. Serial or parallel, I think, in your case.

**Smith:** Well, now to be fair, I had the 11 year stint at Quantum3D as well. So only after the board of Quantum3D elected to do a kind of a different direction as well-- maybe there's some of that one of my portfolio basically-- I left Quantum3D.

I've worked in a number of projects, many of which I've still been involved with Gordie. A kind of a web TV kind of company. It's now called Leonovis. I've also worked at a microprocessor company or two. I actually get back into something that I did a long time ago, which was 3D microscopes. Worked at Bruker and managed a division there.

And then now have a couple of start-ups, one doing interesting DSP technology called RADX. Kind of the same concept as Quantum3D, which is using the PC to distort what is otherwise a perfectly healthy proprietary market for RF communications and measurement instruments. And also I have an app company [called Jigsaw] building, apps for these things [smart phones].

**Hodge:** Why do you still want to do start-ups? It sounded like it was a very impressive amount of work. What makes you want to still do it?

**Smith:** It is a lot of work. But it's the most rewarding thing you can do. Building a company from ground up is a unique puzzle. If we look at the 3dfx experience, we had this huge technological challenge of building this advanced chip or chipset being delivered to the consuming public very inexpensively and also building a business model around that.

Then you overlay that with having to build a company at the same time. It's a great puzzle. And it doesn't always work out. You can't always solve it. But when you do, it's very rewarding.

**Hodge:** What's it like in Silicon Valley building a company now as opposed to building one in '93, '94?

**Smith:** I'm more tired I think than I was then. I look at pictures of myself. And I have darker hair, and a lot fewer pounds, and a lot more energy.

**Hodge:** In terms of the Silicon Valley ecosystem, working with venture capitalists, acquiring talent, that sort of thing.

**Smith:** The talent part-- both my companies are at points where we haven't engaged the venture community quite yet. We're getting pretty close on one. And there one of the things that we learned from Gordie is that don't be so focused on trying to raise money the old fashioned way.

There's a lot of other possibilities where you can raise money, including overseas, as well. And sometimes that's easier to do. But the talent thing is still hard. It's really hard with Apple and Google and Facebook all here. Getting decent engineering talent here is very difficult.

So in fact, in both my companies, my core development teams, they're not in Silicon Valley. One is in Los Angeles, and one is in San Diego. And one is in Serbia. So you kind of have to do what you have to do to make that happen. We have good marketing and finance here and a few engineers. But in general, it's really hard to go head to head against Google or Nvidia or any of these guys to recruit talent. Scott probably has a good perspective on that as well.

**Sellers:** Yeah. I think the difference now in terms of starting a company is that it's such a distributed world. It really did used to be that if you wanted to build a company, you needed to be in Silicon Valley. And sometimes maybe you could do it in Austin or Boston or some of these other places.

But all the great ideas seemed to consistently come out of Silicon Valley. And I don't think you can really say that's the case anymore. The world is so connected that you can start companies all over the place.

Certainly, the venture capitalists invest in companies far and wide, not just in Silicon Valley. I think the VCs also expect companies to be a lot further along than they used to. Because the cost of starting a company have plummeted. So what's expected for you in terms of proof point is a much more mature state than what used to be write a business plan and show some PowerPoint slides.

And that was kind of where most companies were when they got funded. While that may still happen from time to time, most of the time now you have to have some proof point. Your website already exists and you've got so many users. Or your software exists and you've got so many downloads, or what have you.

So I think that in and of itself has changed a lot. Because it puts a lot more of the burden on the entrepreneurs. They have to do a lot more of the heavy lifting before they raise dollar one. And as Ross said, there's an awful lot of alternative funding that exists now, whether it's people that have been successful and now are interested in angel funding and those types of things or alternative column VCs or private equity firms or what have you. But they're all over the place now. It's not just Silicon Valley.

**Hodge:** Do you see angel funding as in some ways sort of doing the same thing that VCs used to, being perhaps more hands on in a more early stage?

**Sellers:** Well, I mean, you see certainly a number of companies that raise several million dollars of angel funding. And that used to be a real pretty meaty ground A from the VC world. So yeah, in certain cases.

**Hodge:** And are you having the same problem that Ross described in terms of finding talent?

**Sellers:** Yeah. I think any company in Silicon Valley would say the same thing. When you start companies here, it's very hard to recruit. It's very hard to retain good quality people.

So you have to find that balance of deciding what you're going to try to do in Silicon Valley and deciding what you're going to go elsewhere. I think it's pretty rare these days that you try to do it all in Silicon Valley. It's just not feasible anymore.

**Smith:** I mean, I just assumed with my current start-ups that I couldn't. I just said, there's no way I can-- I didn't have enough money or attraction to get people out of these big companies. So we went in with the assumption that we're going to have outsourced development. And like I say, I found a good team of engineers in Serbia.

**Tarolli:** They say 9 out of 10 start-ups fail. So I had a very good ride with Silicon Graphics for about 10 years, and then 3dfx for about six or seven years, eight years from pre-founding. And those were probably two of the most rewarding work experiences I've had.

And so I feel very fortunate to be two for two on those two experiences and just don't have a great desire or need to want to do that again. Because I had two really good experiences. So it's like quitting while you're ahead almost.

Those were great. And I have a lot of great memories. And they were the definitely the best times working wise. And many people don't even get to experience it once. So I just feel very fortunate to have been there sort of twice.

**Smith:** You should've asked Gordy. He's back doing it again.

**Tarolli:** Now, view it from the outside. But it's a matter of personal preference. And so I'm happy I've done it twice.

**Hodge:** You have a new company you're doing as well?

**Campbell:** Actually, involved in a couple. The problem with venture capital is that you make an investment in a company. And you can suggest what they should do and shouldn't do.

But they don't need to listen to you, because all you are is an investor. And the only threat you really have is maybe you won't invest next time or whatever. And there are a lot of aspects about venture capital that aren't very, very rewarding in that sense.

It's a lot more fun actually to be in a company where if you identify what you want to do, you can go do it. It's a direct action response satisfaction sort of thing. So I think actually doing something in the company is in a lot of ways more rewarding than doing it from a venture capital standpoint. I think it's a lot more fun doing the type of thing that you're doing on the company side usually.

Because as Ross said, when we did 3dfx, we had to take basically a new concept of technology and not only build it, but actually go out and market it and develop the customer base for it. And those kind of challenges are a lot more fun I think than just trying to decide which new company would I invest in.

**Hodge:** What was the most surprising thing you guys learned during your time at 3dfx?

**Tarolli:** Well, one of them-- I don't know if it was the most-- that gamers were worldwide and older. We get emails from people in Malaysia and all over the world. I had no idea that it was that many gamers out there and that they were, like I said before, in their 20s and 30s. And I think today, that gaming market has grown even more substantial to now, where I think it literally dwarfs the movie industry, doesn't it?

**Smith:** Oh, yeah.

**Sellers:** It's huge.

**Tarolli:** The movie industry's suffered due to other things. But gaming market's probably subsumed the movie industry, the arcade market, and the gaming industry, and a few other things as well.

**Campbell:** And now the performances that are on the portable devices, you've got pretty quality gaming experiences in the mobile market as well.

**Hodge:** Things you would do differently, besides buying the board company?

**Smith:** Still wouldn't do CAD.

**Tarolli:** It comes for free eventually when you get there. But it definitely was not something to make a goal.

**Hodge:** Well, as a thought experiment, if you guys had never bought the board company, what would be your plan for continuing to compete?

**Sellers:** Well, I think even looking back it's still unclear what the right thing to do was. We had such a very different strategy than Nvidia did. And Nvidia ultimately was the competitor. And I don't think we really had the DNA or the culture of a company to build an OEM chip company, which ultimately is what Nvidia did and obviously has proven that's a very successful model.

I think in terms of looking at things from a-- could 3dfx have grown bigger and had a more successful future? It would have had to have been finding a way to either penetrate other markets, which I think at the time that was pretty small potatoes. We were looking at some IP licensing stuff where we would license some technology to some more consumer companies.

But it didn't really materialize until much later. But to really figure out a way to get into the OEM space much more aggressively than we did. It's a lot easier said than done, though.

Because you almost really do have change the whole culture of the company. And we are all the high performance gaming. And you just can't say, oh, now we're going to add this OEM chip project. It's just a very different mentality.

**Smith:** You mentioned the licensing. I forgot about that. But one of the things that we did almost successfully do is we had a license deal for Sega, for the Dreamcast.

**Sellers:** For a design deal.

**Smith:** Yeah. That might have also changed the company. Because that fell apart through some-- ultimately it ended up in litigation that we won against Sega. And was it NEC as well?

**Sellers:** Yeah. They were involved with it.

**Smith:** Basically, the next generation Dreamcast was going to be based on 3dfx technology. And they changed at the last moment over to NEC and to PowerVR--

**Sellers:** PowerVR.

**Smith:** --as a matter of fact, which now powers a lot of smartphones. But that would've given us an interesting position as well. Because if you look at both ATI and at Nvidia, they've had significant revenues and profits that have come from being part of the console market. So had that stuck, that might've been a very interesting aspect of the company.

**Sellers:** It might have changed Sega's fortunes, too.

**Smith:** Oh, yeah.

**Sellers:** That decision to go with the different graphics technology I think was--

**Smith:** Killed them.

**Sellers:** --sort of a poor decision on their part, too. That could've changed it.

**Tarolli:** Was that their last console?

**Sellers:** Yup.

**Smith:** Yeah.

**Tarolli:** To answer that question, which was what would you do differently, I guess I would transition to being a more schedule focused company earlier. Early it didn't suit us, because you just got to get done what you've got to get done. But eventually, as you make that transition to what Scott said, to like the OEM company, you have to meet their schedules not your schedules and be more Windows friendly and compatible.

Towards the end, we solved that problem I think. But it was schedule-wise. And at Nvidia, that was sort of their mindset all along. And it's interesting, because engineers never like to leave something out of a product.

But what I learned was you just have to change the way you think. You're not leaving it out. It's going into the next one. This guy's gone.

**Sellers:** A treadmill.

**Tarolli:** Yeah. It's not a problem. It's just like keep working on it. Oh, you got a month now. And you go and work on it for month, but then pick it up again, because it goes into the next product.

So for an engineer, for myself, it was very easy to rationalize the decision not to put it in this product just to meet schedule. Because at the end, you could just cut all these things out. And you could just rationalize it as an engineer.

It's just going into the next chip. We're not dropping it forever. This one's gone, next one. It's like you missed the bus, take the next bus, right?

**Smith:** I think we should've done motherboards personally. Anyway that's my--

**Hodge:** Would you like to each offer a piece of advice to your choice of future engineers, marketers, developers, entrepreneurs, company builders?

**Tarolli:** It's a great time. It's a fun time.

**Sellers:** It never stops. It's always a great time. The market's moved so quickly. I think there's always an opportunity for great entrepreneurs and technologists to set their mark. It's an awful lot of work, there's no doubt about it. Gordy's seen it more than any of us probably combined.

But sometimes you look back and you said if I ever knew how much work it was going to be, would I ever do it? Usually the answer is yes. But you can never overestimate the amount of work it takes to really start a company and take it somewhere.

**Hodge:** Since you guys are generally having the hardest time finding engineers, programmers, that sort of thing. What would you recommend younger people interested in technology to major in should they go to college, et cetera?



**Sellers:** Well, there's an awful lot of computer scientists out there. That's not the problem. The problem is where you are and who you're competing against. I think what we've done is the most interesting balance, which is be trained from a technical perspective and then bring a business side to it.

Because that's really what Silicon Valley is all about, mirroring the technology with creating markets and create business opportunities. So I think from an educational perspective, that's one of the things I certainly recommend young people is getting both sides. Don't just be an engineer.

Don't just be a business major. Really try to understand both. Because that's the way it is in this day and age. It's very hard to do one without the other.

**Campbell:** The flip side of it is that there are very few places in the world that you can get almost anything you want. And Silicon Valley is one of those places. It may be expensive.

And it may be hard to attract them. But they're here. And there are a lot of places it's hard to do things, because you just can't find the people.

**Sellers:** Yeah. That's not just the engineers.

**Tarolli:** It's everything.

**Sellers:** It's the lawyers and the VCs--

**Tarolli:** Manufacturing people.

**Sellers:** --and the connections. That's a very good point. Everything can happen here much faster than probably anywhere else.

**Campbell:** And I think with the change in communications, as Ross said, it is possible to do people located in a lot of different places more easily today than ever and maybe a hybrid strategy. You do here what you have to do, and you find places like he did in Serbia and other places like that where you can find a very cost effective capability. But I think there are lots of strategies that you can take advantage of. But Silicon Valley is very hard to be in terms of resources and people.

**Smith:** I think what I would offer entrepreneurs is if you're going to do this, don't do it alone. Try and find a team of people that you trust implicitly and who are smarter than you. And that's a really good angle.

Because when you get into the trenches, you need people that can come up with different ideas and aren't afraid to tell you, hey, you're crazy.

Or you can tell them they're crazy, because you live together. It's a marriage. And ultimately, you can produce lovely children from it.

**Campbell:** The other side of it is that a good commentary on 3dfx is you have pretty much the founding team here and everybody still likes everybody. I mean, there are a lot of start-ups that you go through that whole process and they never want to see each other again. I mean, it's just a horrible experience.

This was a really good experience. And I think we all really appreciated what we respectively contributed. And we had a lot of fun doing it.

Everybody worked hard. But it was a great experience. And I think everybody came out of it really well.

**Hodge:** Was it hard being a team of-- depending on how you want to count it-- three or four co-founders? Because the "traditional" Silicon Valley model is two co-founders, usually one tech guy and one business guy. Did you find that at all difficult being three? Or was there really no frame of reference to know it was--

**Sellers:** I think we all brought very different skills and perspective. It would have been very difficult to do the company without each of what we brought to the table. Because we were very ambitious. We needed everything from the algorithm side to the chip development side to the business side to the--

**Tarolli:** Software side.

**Smith:** The software side.

**Sellers:** Everything.

**Tarolli:** It was really a good fit.

**Smith:** I don't think it ever occurred to us that this wasn't the right model. We clicked.

**Sellers:** I think, too, a lot of the entrepreneurs when they meet they get so focused on maintaining the maximum equity that they're going to hold in the company that they lose a lot of time and a lot of expertise that they could have simply by reaching out a lot more.

**Smith:** And sharing.

**Sellers:** Yeah. I mean, bringing people on, of course everyone's in Silicon Valley is going to want equity and going to want a piece of the company. That's just the way that it goes. And too often I think people spend too much time worrying about how much ownership they have in the company as opposed to really what you want to do is build a great company.

And there's plenty of equity to go around. And there's plenty of upside to be had in that regard. But you waste time if you try to keep everything and think you've got all the answers.

**Hodge:** Any other concluding thoughts from anyone? Well, thank you so much for your time. Gary, I know you flew in to come out here.

It's very appreciated, and the rest of you guys making the time from your schedule. And hopefully you guys found this fun and enjoyable in much the same way that debugging a chip on Super Bowl Sunday is.

**Tarolli:** Yes.

**Sellers:** Thank you.

**Smith:** Thanks.

**Sellers:** Thanks.

END OF INTERVIEW