Oral History of Thomas Albrecht

Interviewed by:
Roger Wood

Recorded via Zoom May 26, 2021
Boston, MA

CHM Reference number: X9485.2021

© 2021 Computer History Museum

Wood: This afternoon, we're going to do an oral history for Tom Albrecht. This is on behalf of the Computer History Museum, which is in Mountain View in California. And it's been organized by the Data Storage History Group, which is part of the Museum. My name is Roger Wood. Tom and I have worked together at the same companies for much of our career. We've never worked on the same projects quite, but I'm very familiar with all the great things that Tom has done.

Today is June[1] 26th, 2021. I'll just show one slide here to summarize Tom's career, to put things in better perspective and then we'll start with the interview.

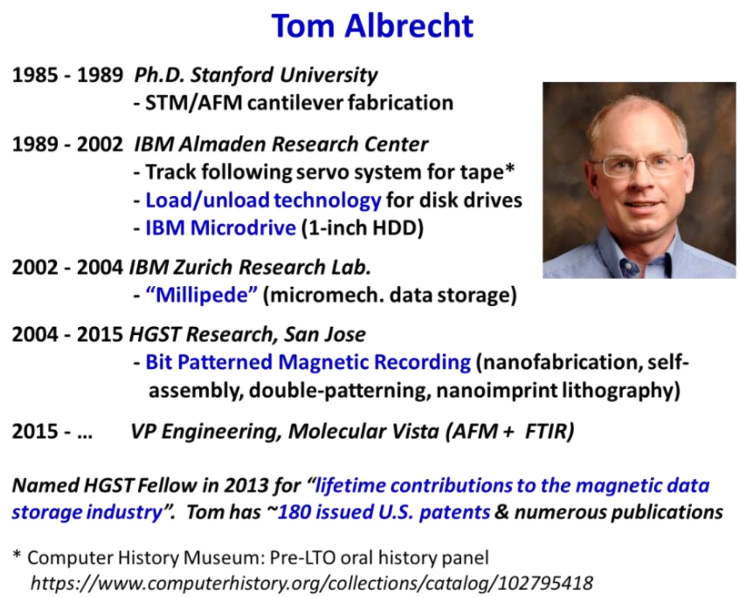

So, Tom, Dr. Tom Albrecht, has a PhD from Stanford University. His topic was Atomic Force Microscopy. He was then hired by the IBM Almaden Research Center. He made contributions in a number of areas. I've highlighted the ones that relate to data storage on this slide. He made major contributions to the LTO Tape Standard. And that has been discussed in quite some detail in a separate oral history, conducted by Chris Bajorek. I put a reference to that at the bottom of the page. So, we may not cover that in quite as much detail. The next topic is load-unload technology for disk drives. It turns out to be a bad idea to leave the read/write heads on the disk surface when the disk is not rotating. And you have to get them off there very carefully or deal with that somehow. And Tom came up with the methodology for doing that, which is now universally used in all disk drives. The next topic is the IBM Microdrive. This is 1" HDD. This was state of the art when it was introduced and was a very successful product. Had a relatively short life because of the advent of Flash memory. But we'll talk about that in some detail. Tom then took a break at IBM Zurich, a Research Lab there. He worked on a technology called Millipede, which is basically the idea of making a very small nano indentations in a recording medium. Again, it was a data storage technology and I'll let Tom describe what happened to that. Then the company changed to Hitachi Global Storage Technologies, HGST, still on the same site or a little bit later moving to a different site in

[1] [Editor's note] Please note that this oral history was recorded May 26th.

Evergreen. Tom during that period, about ten years, led up a project on Bit Patterned Magnetic Recording, also called Patterned Media, and that involved making tiny, tiny domains on a disk surface and using each domain to record one bit of data. And there's a lot of very advanced technologies involved in doing that. That's still on the HDD roadmap, the industry roadmap, but it's in kind of a hiatus at the moment. In 2015, Tom left to join a startup company working again in Atomic Force Microscopy. Tom was named an HGST fellow in 2013 for Lifetime Contributions to the Magnetic Data Storage industry. Tom has an amazing 180 issued U.S. patents, and a whole number of publications as well. There's a picture of a quite young-looking Tom up in the top right corner of this slide, and I'm now going to pass it over to the real Tom Albrecht. Perhaps you can start by telling us a little bit about your family background and how you got interested in technology.

Albrecht: Yeah, thank you, Roger. First, a big thank you to you and to the Computer History Museum for being willing to record these oral histories. Not that mine is so significant, but I think the total effort together of all the things you're recording should make for some very interesting viewing in years to come. So, with that, yeah, a few words about my family background and so forth. I grew up in a small town in Wisconsin and like many typical Midwesterners of German and Norwegian descent, but quite a few generations back, so I was just basically an All-American kid from a small town in Wisconsin. My father was a medical doctor and my mom was a stay-at-home mom. Our family was mom and dad and four boys, and the nice thing was is that some of us kids really enjoyed working with things and building things and experimenting and so forth. And my father had a modest tool shop that he let us play with more or less unsupervised. I don't know how smart that was, but we learned a great deal that way and built kinds--all kinds of interesting things. So, that's sort of the beginnings of my technical inclinations.

Wood: So, Tom, you studied Physics at university. And chose Carleton College. And how did that come about?

Albrecht: Yeah, so during my high school years, this trend towards working on technical hobby type things progressed a bit. I was into ham radio and so forth and enjoyed all those kinds of things. Did a bit of repair work on radios and TVs. It wasn't all just electronic either. Mechanical things were also fun to--we built motorized go-carts and things that typical kids in that area would do at that time. So, when it came time to go to college, I had been a good student in high school. Not just in technical things, but in general subjects. So, I wasn't necessarily seeking to jump directly into Engineering at the expense of everything else. So, I ended up choosing to go to Carleton College in Minnesota, which was a very well-regarded liberal arts college, a small four-year liberal arts college, for that area. And studying all kinds of things. I did know that at heart, I was really an Engineer. And Carleton, being a liberal arts college did not offer degrees in Engineering, but they did have what was called a 3/2 Engineering Plan, so you could attend Carleton for three years, and then go on to an Engineering school for your final two years to really get an Engineering degree. But while I was at college at Carleton, I became quite comfortable with the place, as people often do with their colleges, and being a Physics major there, had enough hands-on in technical, practical things that the engineering side of me was quite content with that, so I never did follow through with the 3/2 Engineering Plan and graduated from Carleton with a degree in Physics.

Wood: And then Stanford University, how did that come about?

Albrecht: Yeah, so when I finished up at Carleton, again, I still understood that I was really inclined towards Engineering. But had a Physics background at that point, so I applied to several universities for graduate school, all in Physics or sort of closely related things. And I had sort of two things I was looking for. One was, indeed, I wanted a program that was sort of engineering oriented that was more Applied Physics, or device-oriented things. And the other things, for better or for worse, I had the perception that some people don't have any fun in graduate school, that it's a dreary time or a difficult time in some people's lives. So, I sort of had my eye on that. And as I visited different schools, I came to different conclusions about those things. I did visit the very prominent graduate school in Physics on the East Coast, one that starts with a "P" and that one sort of fit my-- for whatever reason, the impression I got while I was there is that the grad students weren't having that much fun, and that was tremendous pressure on them. I remember I visited during Spring Break, and the students, instead of going to do anything fun were spending their whole Spring Break studying for a qualifying exam. And this wasn't a qualifying exam that was coming up in a week or two. This was a qualifying exam that was still a year away! <laughs> And I thought, "Boy, that doesn't really look like how I want to spend my graduate school years."

So, anyway, shopped around a bit. Stanford looked a little bit different. There, people really seemed to be enjoying themselves. Coming from the Midwest to this area in March the weather was just absolutely beautiful, so, it felt like going on vacation. And perhaps most importantly Stanford had an Applied Physics program that wasn't really like Applied Physics at some of the other schools. At some other schools, that simply meant Solid State Physics. But at Stanford, it really meant a mix of Electrical Engineering and Physics and lots of device-oriented projects. I had also done a study as an undergraduate, sort of as a senior thesis kind of a thing on the Free Electron Laser, which was a new research direction in those days. A type of tunable laser, pretty exciting, quite interesting, too. Very different than other kinds of lasers and that it used electron beams and wiggler magnets and so forth, too, to create laser radiation. And the team that was leading that effort was at Stanford. So, that was also something that was attracting me to come to Stanford. But once I got to Stanford, I sort of checked out that Free Electron Laser Group and, again, I wasn't quite sure that that looked like a fun group that I wanted to be a part of and spend a bunch of time with. They were going to assign me to go off and do a bunch of theoretical calculations as sort of my initiation to the group.

And looking around, then I found Calvin Quate's group, which was pretty interesting. They were working on Scanning Probe Microscopy, STM, Scanning Tunneling Microscopy was the most successful form of that up to that point. The Atomic Force Microscope, which is what I ended up spending my time on, had just been invented, basically, right around the time I arrived, that work had been done. And perhaps very importantly, the students in Cal's group really looked like they were having fun, and that was true. I had a great time being a member of Cal Quate's group at Stanford and working on Scanning Probe Microscopy. Now, while I was there, I did a bunch of things with AFM. Now AFM, as I mentioned had just been invented. It was invented by Cal Quate, together with two colleagues. And those colleagues were Gerd Binnig from Germany, and Christoph Gerber as well. They came from the IBM Research Lab in Zurich. And Gerd Binnig had been the inventor of the Scanning Tunneling Microscope, which was a very exciting new technology in the 1980s. It was the first technology to really have Atomic Resolution Imaging. You could see the atoms on a surface, and there are famous images from that time. And it was really very exciting to see the arrangements of atoms that all of us had only sort of dreamed about or seen in textbooks up to that point. So, that was very exciting and Cal Quate was a very early adopter. He saw the importance of that, and immediately wanted to steer his group in that direction. He was well-positioned to do that, because prior to that he'd been working on Acoustic Microscopy, using acoustic waves and liquids, including in superfluid helium to image things. So, they were good at Advanced Microscopy already. And he had already steered a couple of his graduate students into Scanning Tunneling Microscopy before I got there, and he had invited Gerd Bennig and Christoph Gerber to come over from the IBM Zurich Lab to spend several months working with the students. And I had that privilege as well to work with Gerd and that's really wonderful. Gerd, of course, went on to win the Nobel Prize because of his invention of the Scanning Tunneling Microscope and he won that very quickly after, you know, on the time scale of Nobel Prizes at least, which shows just how exciting that particular prior work had been.

In any case, the Atomic Force Microscope was invented. A big difference between the Tunneling Microscope and the Force Microscope is the following. The Tunneling Microscope uses a very sharp conducting tip very close to a surface, and electrons tunnel across the vacuum gap between this sharp tip and the surface and you get a little current flow and you can use that to regulate the spacing of the tip, between the tip and sample. And if you do that, you can trace out the contours of a surface, and as I mentioned you can even do that right down to the level of individual atoms, tracing out the contours of the atoms on a surface. But the Scanning Tunneling Microscope cannot look at insulating materials because it relies on that flow of electrons between the tip and sample. And the Atomic Force Microscope was a way around that. They used a sharp tip on a cantilever. You can think of that as a little diving board, or a little spring, and it can be deflected with very small forces if it's soft enough, and basically you can trace out the contours of the surface very precisely with a very sharp tip on a little force sensor like that, provided you can use something to measure that force at a very low level. And what Gerd and Cal did is they used a Tunneling Microscope to look at the backside of a cantilever while the front side of the cantilever had a little sharp tip on it to trace over the surfaces. And so, they had just demonstrated that when I got to Stanford and they immediately recognized that this thing would work better if they had a good, controlled method of making these little cantilever sensors with sharp tips on it. The first one they made was a little hunk of gold foil, and then they smashed a little diamond and picked up some of the diamond shards with a tweezer, and glued them on the end of this, and the whole thing was maybe a millimeter long, so kind of fine surgery to create that. And so, that's not really very reproducible. And Cal had been well aware of microfabrication technology, because Stanford was quite active in that, and they had a nice facility for sort of integrated circuit technology as well as micromechanical stuff. Some of the early work in micromechanical devices had been done at Stanford. And so, he introduced me to one of the guys[2] that was working in that area, with the instructions that my project was, at least at the beginning to make these microcantilevers that would be a big improvement for the Atomic Force Microscope. If you want to put up Slide 17, it'll show an example of some of these microcantilevers that we worked on constructing.

[2] [Interviewee's note] Phil Barth

 © 2021 Computer History Museum

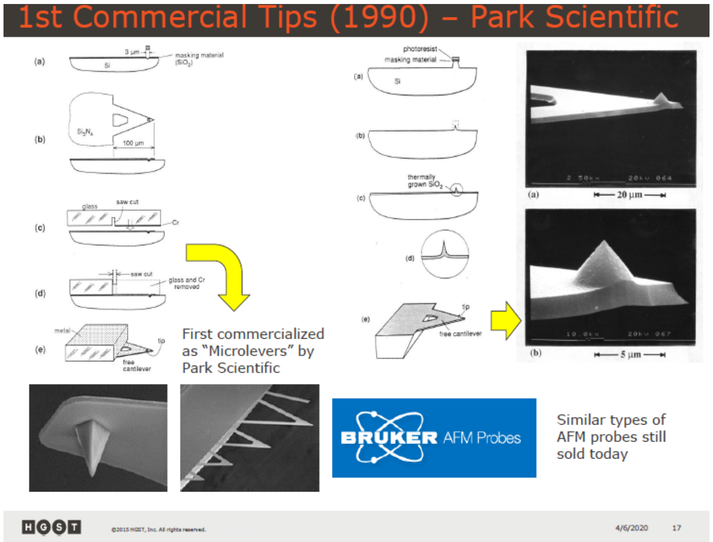

And sure enough, these little cantilevers did in fact turn out to be quite an improvement for the Atomic Force Microscope. And basically, universally after that time, microcantilevers have always been used in Atomic Force Microscopes. I did send some of these microcantilevers immediately off to Gerd Bennig in IBM Zurich. He had gone back by that time. And then he promptly was able to demonstrate atomic resolution with an AFM, using these little cantilevers. Now I also built up my own AFM and spent quite a bit of time. My research project at Stanford wasn't just making cantilevers but really all about everything AFM. I was the first graduate student working in the field of Atomic Force Microscopy. And then shortly after that time, with my own cantilevers I was able to demonstrate atomic resolution on an insulating surface. It happened to be hexagonal boron nitrite, a layered material. Those types of materials turned out to be especially easy for demonstrating the atomic resolution with these instruments. So, here you can see some pictures of the cantilevers down at the bottom, sort of these V-shaped things. Those are 100 to 200 microns long, and about one micron thick. And then there are three other pictures there showing sharp tips of various types that are built onto the end of these cantilevers. The diagram's there, which I won't go into a lot of detail right now. Those are diagrams of how the microfabrication was done to create these cantilevers. Those eventually became commercialized using the same process that I developed. Bruker[3] still makes those today. The company was Digital Instruments and Park Scientific Instruments that originally introduced those cantilevers at that time. And then if you want to go to Slide 16, that will show the atomic resolution imaging that I was able to do on the surface of an insulating material and that was the first example of atomic resolution on an insulator.

[3] [Interviewee's note] Other companies today making AFM microcantilevers include Nanosensors, Nanoworld, MicroMasch, Budget Sensors, and others.

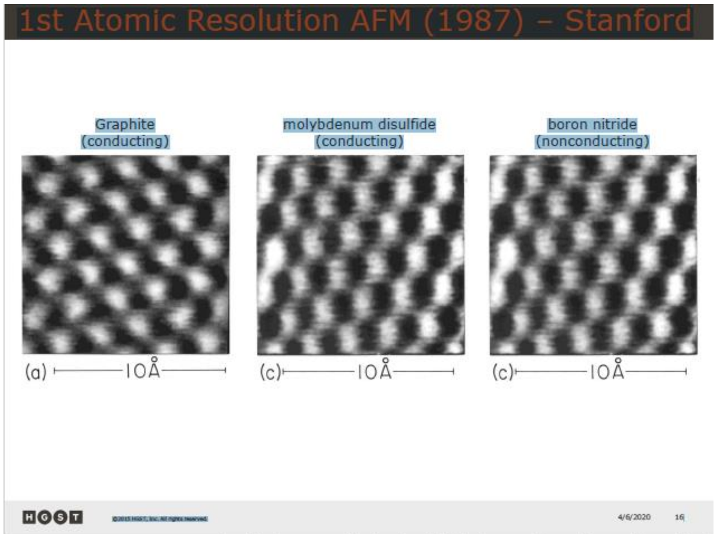

So, yeah, here's the atomic resolution images. Boron nitrite on the right there is a non-conducting material that we were able to show that. Later, people would come to see that perhaps achieving atomic resolution on layered materials didn't really count. It was a little bit too easy, so a number of years later did somebody else[4] finally demonstrate what they now call true Atomic Resolution. Using a later invention, we had which is called FM Detection. We'll come back to that a little bit later. The last thing I wanted to mention about that is after I finished grad school and was recruited by IBM, I'll talk a little bit more about that in a moment, I didn't want to go to the big corporate entity immediately. I somehow sensed that that was moving from the freedom of graduate student days to some kind of rigid corporate structure. And so, I delayed my entry into IBM by the summertime and spent that working with my friends at Park Scientific Instruments. Some of the other students[5] in Cal Quate's group had formed this new company. They had been making Scanning Tunneling Microscopes, one of the first companies doing that. And I worked for them for the summer and we did two things. One is I helped them take the process I had developed for microcantilevers and then commercialize that with a commercial company[6] that would do the fabrication, so then those became available to customers out there. And the second thing I did is I designed the AFM instrument that they first introduced, and that was the first Atomic Force Microscope sold commercially. And the design is still influencing AFM designs today that the way I designed that and how the sensing of the motion of the cantilever was done[7] has been copied for many years, and there's still products out there doing it in a very similar way today. So, that was a lot of fun. That was a summer where we got a lot done, but then in the fall, I did indeed finally move on from AFM and I joined IBM.

Wood: So, Tom, you ended up at IBM Research, obviously. How did that happen?

[4] [Interviewee's note] Franz Giessibl and others
[5] [Interviewee's note] Sang-il Park and Sung Park were the founders; Mike Kirk was another student from Quate's group that joined the company in the early days.
[6] [Interviewee's note] The name of the commercial company was IC Sensors.
[7] [Interviewee's note] To clarify, the basic sensing method (Optical Lever) was introduced by Prof. Paul Hansma's group at UC Santa Barbara; here I refer to the mechanical architecture of how the basic concept was made practical and easy to use.

Albrecht: Yeah, so IBM had been-- especially IBM Almaden Research Center-- had been a favorite choice of Stanford grad students. And in particular, prior members of Cal Quate's group, the group that I belonged to. There were two people from Cal Quate's group at IBM working on Scanning Probe Microscopy. One of them was Dan Rugar, who is still there today working on Atomic Force Microscopy. Really impressive guy, great guy to work with and John Foster, who had done a little bit of STM Data Storage, sort of tacking down little molecules on a surface in an attempt to try to see if you could store data at super high density with Scanning Tunneling Microscopy. Both of those guys had grown up in Quate's group doing Acoustic Microscopy but switched over to Scanning Probe Microscopy upon graduation. They, like Cal, sensed that this was really the future. So, in any case, IBM seemed like a natural place to go. I knew a bit about it. I had visited a few times with these prior students that had gone there. I did look at a few other IBM sites as well. I had some offers from Yorktown[8] and Burlington, Vermont and Rochester, Minnesota. But in the end, it was just awfully comfortable to go from Stanford and just move 30 miles south, and go to IBM Almaden, where my friends were already doing work on Scanning Probe Microscopy. It all looked like a very comfortable transition. So, that's how I ended up at IBM.

Wood: Okay <laughs>. And then I guess there were several projects ongoing that you were perhaps invited to participate in. And tell me about how that worked out, and which ones you chose to focus on in the end.

Albrecht: Yeah, it's a case of - for whatever reason they were very generous when they hired me and wanted to give me freedom to select a project that would be of interest to me as well as of course being something useful to IBM. And so, when I first joined them, they had me not report to a first line manager, but to report to the functional manager or what we call sort of the Director of Data Storage Technology Research at IBM Almaden. His name was Heiner Sussner. And so, I reported to him for several months and the idea was that I could try out, say, two or three different projects, work on each one for a month or two or three. And then after a period of time, say certainly within a year, I should pick one of those and make a career of it. Or at least a start of a career. So, I didn't follow the directions very well on that. I sort of failed to carry out what they had hoped I would do. I did in fact work on several different projects. I worked on three different projects, more or less simultaneously. You know, a few days on this, and then I'd go over and work a few days on something else, and then rotate. Around and around, week-by-week doing all kinds of things. The three projects were the following. One was, of course, working with Dan Rugar on Atomic Force Microscopy. And a second one was working with Jim Eaton on Tape Technology. And then the third one was working on sort of radical new ideas for hard drive technology with John Foster. So, those were sort of three interesting projects that I got to kick off the time at IBM Almaden with.

Wood: You made some very major contributions to the tape world. To some extent that's covered in a different oral history. But do you want to briefly tell us about that as well, Tom?

Albrecht: I can! Tape wasn't something that was-- that I naturally would have gravitated towards, but in-between the time where I interviewed at IBM Almaden and when I started, I didn't actually interview with

[8] [Interviewee's note] IBM Thomas J Watson Research Lab in Yorktown Heights, NY

Jim Eaton, but he had somehow become familiar with the fact that I was there, and he thought maybe there might be some interest in working with his new team. He was putting together a brand-new effort on Tape Technology for storage. Tape was, at that particular time, a relatively neglected, perhaps fading away magnetic data storage technology. The storage density was quite low. Maybe something like 100 times lower than what was going on in hard-disk drives at the time. And tape was slow. It had a long access time. You know, you get out a tape cassette and it would take you maybe 30 seconds to wind to the position of the tape that you wanted to access something. So, for lots of reasons, tape was a pretty low performance and low-density storage medium at that time. But Jim Eaton had a real vision for what tape could be. And he envisioned something with much higher aerial density, and also something with much faster access time. And he wanted to build a new generation of tape drives that used small cartridges with two reels in them, sort of like a music cassette, where you'd park the tape in the middle of the reels, so that the access time to any part of the tape would be at most half of the total time to spool through things. And not too big so that the reels wouldn't be too large. So, you could access very quickly. And anyway, he came up to Stanford and visited me. He sort of sought me out and brought his little satchel of parts with motors and tape cartridges and talked about this. And I was intrigued enough that I thought, "Yeah, I should learn more about this." It was all electrical and mechanical <clears throat> technology, so it was sort of a natural fit for the Engineer in me. And as I mentioned, I didn't stop doing that after sort of my couple months trial period with that. Stayed with that one actually for more than ten years! I worked on a number of mechanical technologies to do with tape guiding and cartridge design and so forth. But the really, the big impact came when we wanted to introduce track following servo to tape. Hard drives had gone through this about ten years prior to that. Where instead of just parking the head at a particular position and hoping the track on the disk would line up with it, you actually followed the track wherever it went, and you know, as you scale the track smaller and smaller, that becomes very important. Same thing with tape, at that time, the tracks were wide enough that if you simply put the head where the track was supposed to be on the tape, you were close enough to it that everything worked fine. But they knew that to go to a far higher areal densities on tape, they would need to go to very fine tracks and introduce track following servo for tape as well. Now the first inclination had been, of course, to simply copy what was being done for hard disk drives. That was obviously the natural place to start. But there were a bunch of reasons why that wasn't a very good choice for tape. Primarily because tape heads see quite a bit of actual wear and removal of material over their lifetime, which changes their spatial response function; and then also tape is a pretty dirty environment. So, little chunks of debris float through the head tape interface and lift up one side of the head versus the other side of the head and so forth. And all this meant that you would end up with some bad errors in position signals derived from servo heads in tape drives if you simply tried to copy what was being done for hard disk drives. So, Jim Eaton and I and Rob Barrett, who was another Cal Quate student from Stanford who joined IBM Almaden shortly after I did and joined the Tape Group, the three of us worked on what we called Timing Based Servo. And since, as you mentioned, that's been handled in a different oral history, I actually won't go into all the details of what that is and how it worked. You can look up that other oral history to find out all about it. But it ended up becoming the dominant and standard method used in all tape drives going forward, including in the LTO standard, which has become the dominant tape standard in use for tape drives. And indeed, this whole activity, although it turned out a little bit different than Jim Eaton's original vision, did breathe new life into tape, and tape is still today a very important storage technology. Interestingly when IBM sold its hard disk drive business to Hitachi to make HGST, Hitachi wasn't interested in the tape and basically told IBM to keep that old, outdated technology, and IBM has kept it and actually has made plenty of money on it over all these years, and they're still very active. And the Zurich Lab is busy perfecting new inventions in tape and down in Tucson is where the products are developed. And so, that's all a very active area still today. So, anyway, I had a great deal of fun working on tape over those years. And served on the LTO Standards Committee when we wanted to introduce the Timing Based Servo. Also, played a role in the design of the cartridge and the drive interfaces and so forth. So, ended up being able to leave a recognizable mark on that industry and something I can look back on and feel good about.

Wood: It's quite a remarkable story, isn't it? And the Timing Based Servo has really proven its extensibility to this day, several decades of increases in aerial density on tape. And yet it's still the basic kind of Timing Based Servo that you invented. So, it really is a remarkable story. So, Tom, we should switch to HDD stuff. There's a-- I wrote a little note down here saying, "Tahoe." And I don't know where-- I know roughly what it is, but I don't know where the name came from. But tell us all about Tahoe.

Albrecht: Yes! So, Tahoe, the name Tahoe has everything to do with skiing. And John Foster had gone off to a Lake Arrowhead Conference. I never attended one of those myself, but I wouldn't be surprised if you went to one of these or more over the years. And apparently, among the things that happened at these Lake Arrowhead conferences is people would show perhaps new ideas, maybe even some crazy ideas, for new ways to do hard disk drives, new technologies to put in. I'm sure they covered all kinds of standard stuff, as well. But John had gone to one of these, and he had seen a presentation by a fellow named Jim Lemke. The thought had been at that time that in order to really make a quantum leap in areal density for hard disk drives, that one of the most important things to do would be to reduce the magnetic spacing, and not incrementally, but go for all of it at once. As everyone knows well, hard disk drives generally float the head on a thin film of air over the rotating disk, and that way you can avoid friction and wear. Works very well, still done today. But anyway, at that time the thought was that getting rid of that air, and bringing the head directly into contact with the disk, would be some kind of a holy grail. It would really allow you to raise the areal density and achieve success in some fashion. A fair amount of that thinking ended up being a bit misguided, and we'll get into that in a moment. But where Jim Lemke and skiing come in is that the idea that Jim had presented was instead of using an air bearing, to use a very thin layer of liquid on the disk. Because you don't really want solid-solid contact; the solid head banging into the solid disk is just asking for trouble in terms of friction, and wear, and debris, and all kinds of problems. The thought was to put a thin liquid layer in there, and John, being a waterskiing enthusiast himself, it sort of clicked in his mind that when you are skiing on water, you're able to have a relatively small area of contact between your ski and the water, and actually achieve a rather low friction level. I'm a terrible water-skier myself, so I've never really achieved that low friction myself. But I think John was quite good at it, and he really had a vision that a hard disk drive head could ski on a liquid-coated disk in exactly the same way. If that liquid were super thin, say just a few nanometers thick, that somehow this would all work and give you the very low magnetic spacing, and then allow you to raise the areal density a great deal. We had a lot of fun with this ski type terminology in the project, and this figure just shows some of the little pictures we drew for ourselves, and we stuck on our office doors at that time.[9]

[9] [Interviewee's note] These cute "ski"-related pictures were made by Larry Best, a member of our team that helped with firmware development among other things.

I guess I'm referred to as a lift operator, that refers to load/unload. We'll get to that <laughs> in a moment. But any case, that's all I wanted to show there, is the ski concept was a lot of fun, and the group was a lot of fun. John was just a very enthusiastic manager, and really had a healthy ambition for wanting to achieve more than just an incremental gain in areal density.

Wood: Tom, fill in the names for us, please.

Albrecht: Yes, indeed, I should have mentioned that. So John you see in the lower-left. That's John Foster. As a ski instructor, of course, he's the leader of the group and the project. Moris, that's Moris Dovek, yet another Cal Quate student who joined IBM. He didn't join at Almaden, he was at the plant site, working in head technology. He was responsible, as part of the team, to really take care a lot of the head-related things, the air bearings, and the head overcoat issues, and so forth. So I guess he got called the ski designer in that little graphic. Jackie, that's Jackie Spong yet another Stanford-- not from Cal Quate's group, but from a different group, but part of the Stanford applied physics alumni. She also had worked a bit on load/unload. In fact, Jackie was the one who said "Why don't we just do load/unload?" when we got to a certain point of the project, or we had troubles that we were going to solve that way. I was much too practical and thinking "Let's not do something crazy, let's solve it by other means." But Jackie had the vision to say "No, we should just do load/unload." It sort of became my job, eventually, to make it all work. But that's who those four are. Moris today is still very active in the magnetic recording industry. He is maybe the-- I may not have my information quite straight, but I'm thinking chief technical officer at Headway, or TDK Heads, basically, at this point.

Wood: That's the latest I had, as well.

Albrecht: Okay, and Jackie, she did not stay in the hard disk drive industry, so I've lost track of where she has gone off to now.

Wood: It looks to me, Tom, like you were having far too much fun.

Albrecht: Well, we were, and maybe that's trouble, because this project, of course, didn't succeed in the manner that it was intended, in that skiing drives with liquid films became commonplace and successful, and used throughout the industry. That did not happen.

Wood: I knew Jim Lemke quite well, and among his myriad talents, he was also a wonderful salesman. I believe he had a little company called VISqUS that did this Tahoe [ph?] skiing technique. I think it was sold to Maxtor, if I remember rightly. But, as you say--

Albrecht: I think you--

Wood: -in the end it did not come to anything. But I always reckoned Jim could sell the hind leg off a donkey. <laughs>

Albrecht: Yes. <laughs>

Wood: <overlapping conversation> amazing fellow and technically brilliant, as well.

Albrecht: Yeah, and John Foster also had the powers of persuasion, and however you want to put it. Real optimistic, could get people to try things that they might not otherwise be willing to work with. The company, VISqUS, had already been formed when John saw that Lake Arrowhead presentation, I believe. Because he came back talking about both Jim and VISqUS when he was excited about that technology.

Wood: Very good. So I'm hoping that some of the younger engineers watch this video, because it's nice to hear sort of all the interactions, and the enjoyment you obviously-- you got out of these projects, as well. Okay, we should move along. I think the next topic on my list here is the load/unload, which you mentioned briefly. So tell us about the-- first of all, the perils of leaving the read/write head on the disk surface when the disk is not rotating?

Albrecht: Yeah, absolutely, and I'll do that in the context of this Tahoe program, because this was a direct outgrowth. Although the Tahoe program and the ski technology itself didn't come to fruition as it had been envisioned by John, it did end up leaving its mark on the industry, in the form of load/unload, and a few other things. Now, when you're trying to reduce magnetic spacing a great deal, with or without a liquid layer on the disk, one of the things you need to do is make that disk very, very smooth. Because any texture, or roughness on the disk, tends to add to your magnetic spacing that you need to fly over that stuff at a height higher than that. So we had a lot of incentive to make the disk as smooth as we possibly could, and then this, of course, would be no surprise to you, Roger, that we started to have stiction problems. This was a known problem for many years, that if a head landed on a smooth disk and came to a stop, it would tend to stick. Due to capillary reaction of just sort of any contaminate materials that are present on the surface, it'll really suck the head down with quite an attractive force, and hold it there quite solidly. In fact, in severe cases, if you apply enough torque to turn the disk, it'll rip the head off the suspension. So that's how serious that problem is. Now, if you go back to the very early days of the hard disk drive, the RAMAC, which was the initial IBM hard disk drive product in the 1950s, or, really, the first hard disk drive ever, they already recognized that problem in those days, and used a load/unload system on the RAMAC. It had a statically pressurized air bearing, so it actually pumped air through a little hose to the head, and then the head floated on the disk very much like the puck on an air hockey table. And then when it was time to shut the tool down, obviously they wanted to turn off that air pump, and they couldn't let the heads just land on the disk. So they had a little mechanism with pulleys, and levers, and cables that would lift each head up off the disk and hold it suspended in mid-air. So load/unload was not exactly a new idea, and it was well known. But those earlier what I will call load/unload contraptions weren't really well suited to small form factor, inexpensive disk drives that were becoming dominant in the late 1980s. Here, just a very quick aside, in the early 1990s, IBM really went through some turmoil, and one of the things that caused so many people to lose confidence in IBM as a business is that it was perceived that the day of the large-scale computing device, the refrigerator- sized server was going away. That didn't really turn out to be entirely true. IBM is still making those today, and still making some money off of those. But what did happen, and this was very definitive, the refrigerator-sized disk drive did come to an abrupt and rapid stop right around 1990, shortly after I joined IBM. Because people figured out how to use consumer-level, or, at least, small form factor, relatively inexpensive drives to replace what the large drives did. So eventually, arrays of little drives replaced the big drives, and they are gone forever by about 1992. So, with that aside, where are we here? Load/unload, although a number of these larger drives had used various load/unload contraptions over the years, the small drives couldn't afford to do that. That was too complicated of a mechanism. So this alternative idea came along, which was called ramp load/unload, and this also we didn't invent. There was a little company called PrairieTek that was introducing some of the first two-and-a-half-inch mobile hard disk drives. They chose to use ramp load/unload, and demonstrated that, and we were aware of that. So when we had our problems in the Tahoe project with the heads sticking to the disk because they were too smooth, that was a bit of an inspiration for us to see what PrairieTek had done. So we studied those drives and then we made a number of improvements and inventions on our own to create a load/unload system to use in Tahoe, and, indeed, we did that. I'm going to show here one of our Tahoe drives. This doesn't show it as well as I'd like, but there's a window in that drive, and you can see the head parked on this little white plastic thing, and that's the load/unload ramp. You can also see the funny sticker on there that says "Ski" inside. That's us continuing to have fun with the Tahoe skiing theme. So this actually was one of those drives that had liquid film on the disk, and tried to be achieving a low magnetic spacing in that particular way. It did sort of work, but in the end, didn't turn out to be the right way to do things. Keeping the air actually turned out to be a good idea. You can make that air arbitrarily thin, as we've learned over the years, and you can work on the magnetic spacing in terms of reducing the overcoats, and lubricant layers, and so forth. That's where that's all gone, and then people were able to keep a few nanometers of air over all those years, so skiing on liquid never turned out to be essential. But, the idea of going to a super smooth disk as part of the solution for achieving very low magnetic spacing did catch on from this project. We did those Tahoe prototypes with the team in IBM Fujisawa, Japan, which was responsible for making our mobile two-and-a-half-inch disk drives. They saw what we were doing, I think they properly recognized that skiing on liquid was a pretty crazy idea, but they were very intrigued with the load/unload system that we had working at that time. Stiction and start/stop wear had been these perennial problems for disk drives for a long time. They often ended up in tough, difficult things that would delay programs, or cause programs to fail. A lot of money and time was wasted, or I should say invested, in start/stop solutions. Usually, in those days, the way that was done to keep the cost manageable is they would intentionally roughen the disk, either the whole disk surface, using basically lapping particles. Or just what they would call the start/stop zone at the ID[10] of the disk using sort of a zone texturing idea, which could either be a zone particle texture, or laser-induced bumps that they would form. But still, that was found to be very challenging, to get just the right tradeoff, with enough texture to prevent it from sticking, but not so much that you had to design your head to fly very high. Also, trying to start/stop on a textured zone required a pretty robust overcoat, so that was hurting your magnetic spacing, as well. There was a lot of difficult testing. You'd want to run these things for tens to hundreds of thousands of stop/start cycles, depending on the particular market segment you were working with. Especially for mobile, where people were shutting down drives left and right to save battery power, and you could actually have hundreds of thousands of start/stop cycles over the life of a drive, this became really an intractable problem, the start/stop in contact solution. So when they saw this crazy group doing load/unload for their liquid ski drives, and the project eventually went belly up and disappeared, and John left IBM, and all kinds of things happened, they were interested in that. So I continued to work with Fujisawa. In fact, I had many, many trips to Fujisawa those years, I counted something like 40 trips over there, over a period of a number of years, to help them take what we had learned in the Tahoe program for load/unload, and introduce it into conventional two-and-a-half-inch disk drives. This allowed them to go to very smooth disks, get that texture down to a very low level, start to optimize things for signal-to-noise ratio, and high-density recording, rather than worrying about start/stop friction, and stiction. That led, actually, to sort of I would call a golden era of rapid areal density growth in disk drive technology, and the way was led by IBM Fujisawa two-and-a-half-inch drives. No one in the industry could match what IBM could do at that time in areal density on those drives. In fact, the mid 1990s we were able to achieve-- and, of course, this isn't only because of load/unload. There are a few other very important things coming in at the same time, like MR head is not to be underestimated by any means, and various other technologies[11]. But we achieved about 100% compound growth rate. So every year we could double the density of a two-and-a-half-inch hard drive, and then the larger drives would follow on shortly after that, as they adapted the technology to work in the larger form factor drives.

Wood: One of the interesting aspects, Tom, from the electrical engineer perspective is what happens if somebody accidentally unplugs the drive? How do you get the heads off the disk?

Albrecht: Yeah, right. So that was one of the most important problems to demonstrate a good solution for. Because, indeed, it all sounds fine, in theory, to talk about this ramp load/unload system, and when you want to shut down the drive, then move the heads up the ramp, and park them, and everything is safe. But especially in the mobile environment, where your batteries simply run out, or whatever, there are unplanned power outages, and the drive, of course, needs to be able to survive those very well. Now,

[10] [Interviewee's note] ID = inner diameter

[11] [Interviewee's note] A number of remarkable new technologies were introduced in HDDs in the late 1980s and 1990s, including thin film and magnetoresistive (RM) and later giant magnetoresistive (GMR) heads, sector-based track-following servo, partial response maximum likelihood (PRML) channel electronics, and thin film magnetic disks. These four major new technologies were probably more important and certainly more technically impressive than load/unload. Nonetheless, load/unload played an important role in bringing down the magnetic spacing, which enabled HDD developers to maximize the areal density gains from this wonderful suite of technologies, giving rise to what I consider the "golden age" of explosive areal density growth in HDDs.

even from earlier days, there had been what's called a power-off retract system, and that is in the old days, without load/unload, you'd always park the drive, the heads at the inner diameter of the disk, because the start-up torque would be the least for the motor if the heads were parked there. So there was always a system that would move the heads to the ID of the disk if you pulled the plug and removed the power. That was done by capturing some of the energy in the rotating disk pack, the spindle motor. When the power was turned off, you started to use that spindle motor as a generator, since it had permanent magnets in it, and then you'd rectify that and apply it to the actuator to move the head in the desired direction, and in the old days, that was towards the ID. Of course, we wanted to do the same thing for load/unload, use that to move the heads to the outer diameter, where the parking ramp is. But one thing we discovered is that the existing circuitry for capturing that spindle motor current was not efficient enough, and we couldn't get enough current to reliably put the heads up this load/unload ramp, and park them safely. So I went down to the lab, and I remember a number of late nights fiddling with circuits. <laughs> Could I do better than what was being done? You want to find a way to get this fairly small voltage and current from the spindle motor out. Especially in a mobile drive, the motor is running on five volts. So you know that the back EMF being produced by the motor when you turn it off is going to be less than five volts, then you want to rectify that, and apply it to an actuator. You've got to have some kind of rectifier, that's going to have its own voltage drop. The old way of doing that had simply been to use silicon diodes, each one drops about .7 volts. If you make a full wave rectifier, then it drops 1.4 volts. Sure enough, you don't have enough voltage left to get the job done. I did figure out an interesting circuit using bipolar transistors instead of diodes. It turns out that the collector-emitter saturation voltage of a turned-on transistor can be very low, only about 0.1 volts. As long as you've got enough total voltage to turn on the base of the transistor, and we did, then as long as the main current goes through the emitter collector path in the transistor, you can win. So, in fact, the first generation of load/unload drives produced by Fujisawa used this little circuit that I developed in the middle of the night in a lab in Almaden, using six bipolar transistors to rectify that current for the power-off unload. Eventually, people could do much better than that, and they started making active rectifier circuits that were very fancy MOSFET circuits, and very sophisticated, and then those did better. So that solution that I came up with only lasted one generation long, but it was a critical piece of getting this into products. Because to the extent that those active rectifier circuits didn't really exist yet, getting a custom chip made is too big of a barrier to introduce a new technology. That six transistor thing, which you could build yourself with a few parts from Digi-Key, was perfect to get something in quickly with a very low barrier to entry. So, yeah, the load/unload power-off retract rectifier ended up being a vital piece of the puzzle to get that technology in.

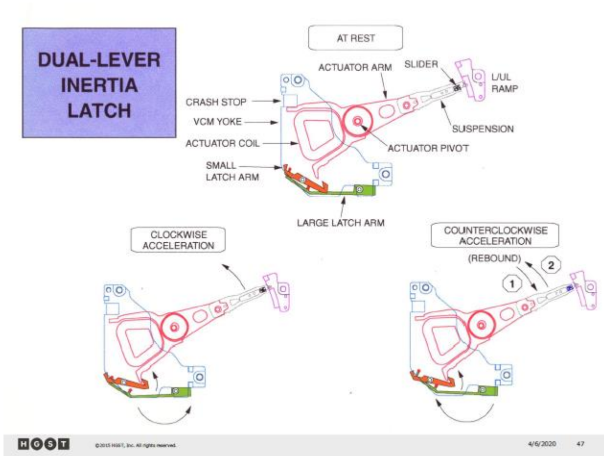

There's one more issue, of what goes wrong in terms of things that can destroy disk drives. One is, of course, if the head lands on the disk due to the power being removed, that was solved by the retract rectifier that we just talked about. Another thing that can go wrong, especially in mobile drives, is the customer drops his computer, and subjects that hard disk drive to a big shock. Now part of the beauty of load/unload was by parking the disks-- parking the heads off the disks over on this load/unload ramp, even if things jumped around and jiggled a bit when you'd dropped that disk drive on the floor, along with the rest of your computer, the heads wouldn't bang against the disk and cause a big problem. But one thing they could do is if they were not sufficiently well latched in that position, and there was a rotational component to the shock, which there almost always is, then the whole actuator could rotate, just like seeking out a track, and put those heads back on the disk. Then they would land, and they would stick, and that drive would be dead. So another thought was to use what's called an inertia latch, and there had been some inertia latches introduced prior to us doing this. PrairieTek and also the follow-on company, Connor-- no, Integral, <laughs> sorry. Integral was the company that also had load/unload in some of these small form factor drives. They had a type of inertia latch, because they recognized this problem, but they didn't do a good enough job. We were easily able to make their inertia latch fail. By having just the right combination of linear and rotational shock, you could make it fail every time, and that was viewed as not good enough. So along with some guys in Yorktown Heights, I think Sri-Jayantha[12] played a role in this particular type of dual lever inertia latch, as did Kohji Takahashi in Fujisawa. The idea of an inertia latch is if your drive undergoes a rotational acceleration in the direction that would cause the heads to want to rotate towards the disk, then as long as you had another body in there being a latch that would rotate the same way, and catch it before it could get on the disk, then that's a latch. But unfortunately, that thinking doesn't work well enough, the situation is quite chaotic during a shock event. It turned out that shocks of all different kinds could cause the heads to end up on the disk, and the inertia latch not to work. Part of it was the simple case of if you cause the actuator to bounce off the OD crash stop, then it would also end up back in the disk. So that's sort of the wrong direction of shock, one that was actually pushing the head off the disk. But that could cause a rebound from the bumper and put it back on the disk. So bottom line is we needed an inertia latch that would close in both directions of rotational acceleration. I won't go into the details here. If you're curious, you can kind of figure it out studying this picture. But this is a mechanical rectifier, that if you put a rotational acceleration on that green body and cause it to rotate

[12] [Interviewee's note] Full name Sri Muthuthambi Sri-Jayantha, an impressive mechanical engineer and inventor.

either clockwise or counterclockwise, then it exerts a little force on the orange body that always rotates it in the same direction. And causes that little hook you see on the right side of the orange object to hook over a little tail on the actuator, and hold it where it is so that it cannot move. So that's the idea. The dual lever inertia latch is a mechanical rectifier that takes whatever kind of shock you apply and locks the actuator firmly, and this worked very well. It was actually not at all easy to make this particular latch fail, and this became standard usage on disk drives that used load/unload. I have not checked today whether this is still what's used now on drives, but it might be. It certainly was used for many years, the dual lever inertia latch.

Wood: Very nice, Tom, very ingenious, is that. Some of these things seem not that important, but the idea of the drive failing from a shock, I mean, it's paramount that that has to be fixed. And this is really ingenious, very nice.

Albrecht: Yeah, all those details that are the difference between a good idea and something that you actually <laughs> want to commercialize.

Wood: Absolutely, yeah. I'd like to move on to-- this is probably about the main topic. The IBM Microdrive, you, and I think it's Tim Reiley, played a really major role in the development of this product. I'd like you to tell us all about the IBM Microdrive, please.

Albrecht: Ah, this is a fun slide (62) to look at. I mean, this is what the Microdrive is all about, making a very small disk drive. A bit about this particular project, indeed, Tim Reiley was the father of the Microdrive concept.

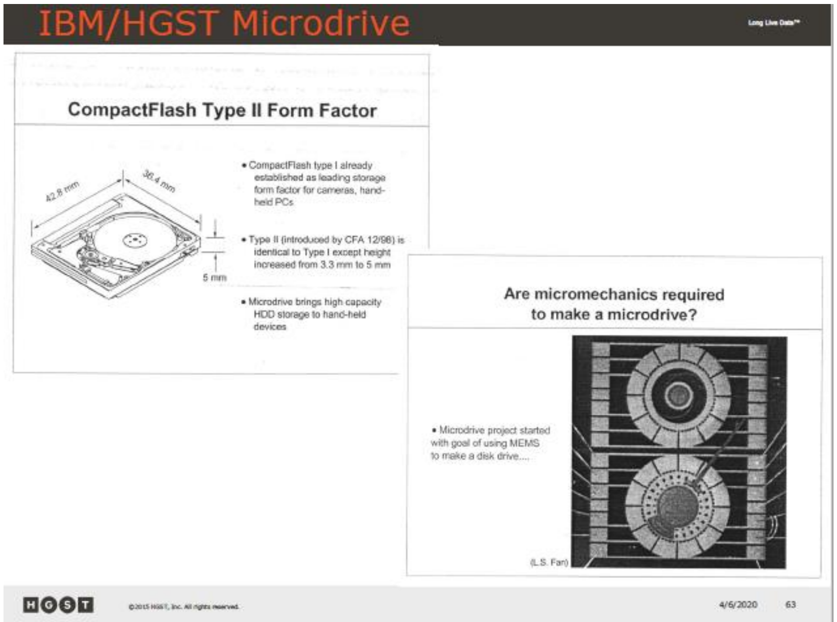

Tim was interested in micromechanics. He had seen all of the things that people were building using integrated circuit technology, and silicon, to make tiny mechanical devices. He had a vision of making a disk drive that way, that you'd have a tiny microfabricated motor, and a tiny microfabricated actuator with an arm on it and a head. They made some conceptual devices in the lower-right, those didn't really work in any kind of complete fashion. But down below right there, you see actually at the top some kind of microfabricated rotational motor, which was intended for the disk. At the bottom of that picture, an arm, it looks a little bit like a clock hand on another center of rotation, and that was meant to be a rotary actuator to move that arm, or the clock hand, to access a disk that would be spinning on that upper motor. They worked on this for a few years. Tim was working with Long-Sheng Fan, who was a microfabrication expert from UC Berkeley who joined IBM Almaden, and they produced some of these kinds of structures.

But after a few years, the manager, Francis Lee, was their manager, felt that this wasn't practical enough, and it wasn't likely to produce a drive that really worked; a neat concept, really cool to look at, but too far from anything real. For whatever reason, they asked me to take a look at this project, and I was immediately sort of attracted to it. I knew Tim quite well, I got to know him much better through this project. One of the things we recognized is that if you followed through on this concept as they had in mind, to make something truly tiny-- and this we'll come back to with the Millipede as well, because it ends up being a fatal flaw for that project, as well-- you end up with too small of a piece of magnetic media. In other words, the disk diameter's going to be too small, so that even if you have a state-of-the-art areal density, you won't have very many megabytes stored in this device. So, really, this concept was too small to be reasonable, because there's always a certain amount of fixed overhead cost, and you could not drive that to zero just by microfabricating things. To have too little media there would mean that the cost per bit would be too expensive, and even in those days, they could see that solid-state storage

 © 2021 Computer History Museum

would probably do a better job than a disk drive this tiny. So we set about to look at whether we could make a more conventional hard disk drive that was very, very small, and then still make it commercially viable. This was during what I call the golden era of areal density growth for hard disk drive, where we were achieving 100% growth per year. So it was reasonable to suspect that form factors smaller than the 2.5 inch and 1.8 inch drives that were the bottom end of the hard disk drive world at that time could be commercially viable. Also, solid-state flash memory, although it was existing in some of its infant forms, was not very high-density, and was quite expensive in those days. So, indeed, we were able to come up with a technology package for this to make a small drive -- you could see in that picture with the hamster, and so forth, that it had sort of a conventional technology.

If you look at that, you can basically see what appears to be everything that a normal two-and-a-half-inch drive would have, just a lot smaller. That's really what it is, it has a normal voice coil rotary actuator, it has a load/unload system for the head. It has a small disk, about 27.4 millimeters in diameter. That big white C-shaped thing is a breather filter to let filtered air in and out of the drive. So we were able to find ways to simply shrink all the components of a conventional drive. That was done together with IBM Fujisawa, again. We had made some acquaintances there, and they were very enamored with this concept of introducing a tiny disk drive. In particular, the top manager of the team in Fujisawa, and his name was Ino-san I forget his first name, which I apologize for (Hideya Ino). But Ino-san really was the person who decided that they were going to do this, they were going to help us turn this concept into a drive that really worked.[13] They were successful with that. There were a number of challenging technical things that

[13] [Interviewee's note] The contributions of the IBM Fujisawa technical team to the development of the Microdrive cannot be overstated; without their major investment of people and effort, the project would never have happened. I cannot do justice in giving credit to the many Fujisawa engineers that were vital to this project. A few that come to mind, however, are Keishi Takahashi, and Kenji Kuroki who were leaders in the project. David Albrecht, an assignee from San Jose working in Fujisawa at the time, also played an important role in the Microdrive's design. Tim Reiley from Almaden, the father of the Microdrive concept, remained active over the entire life of the project as well.

needed to be done to succeed in this miniaturization, and two of them were exactly the things we were just talking about a moment ago. It turns out that the load/unload system had some special challenges for this very tiny drive, in that the power-off retract wouldn't work. You simply could not get enough back EMF from this inertially rotating spindle motor power-off to move the actuator back onto the parking ramp. But we were able to come up with some very efficient and clever capacitive retract systems, where we just stored some energy in an electrolytic capacitor, and this thing is small enough that you don't need too large of a capacitor to make this all work. One of my colleagues, Erno Klaassen, who was only at IBM a short time. He was the son of Klaas Klaassen, who I'm sure you'll remember very well. He was an electronics wizard, and came up with some very clever ideas for a pulsed, a load/unload retract system that maximized how much you could extract from a capacitor and apply to this actuator and then get it parked again. So that was one of the clever things. The spindle motor had special challenges. There was a fluid type spindle bearing used in this, and that was a very early use of that, because it was easier to make that compact than the necessary ball bearing designs that were being looked at. And also, the inertia latch was a challenge because inertia scales with some high power of the radius. When you try to miniaturize an inertia latch, you find it has very little inertia and doesn't exert much torque under shock conditions as you make it too small. So we actually had to put a very much over-size inertia latch in this thing. There's a great big counterweight that goes up under the disc, and it looks rather disproportionately large. But indeed, we were able to make the little dual lever reversing inertia latch in this drive.

So technically we were able to make all of that work. But how do you get this concept into the marketplace in a successful way? Well, Hewlett-Packard had introduced a small drive a little bigger than this one called KittyHawk a few years prior to this. And they actually had failed in the marketplace. It had this nice little drive, it actually worked, but there weren't really any killer apps for it. There weren't enough of them. Digital cameras did not exist yet when Kitty Hawk came along, nor did music players. But they did in the timeframe of this little Microdrive. And this ended up being the marketplaces that were vital to making this all happen. But you can't just put out a drive that's small and tell everybody to use it. We had one very lucky break, in that the flash memory format that camera users were using was something called the Compact Flash form factor. That's this little thing here. This one's actually a Microdrive. But this size and shape and the particular connector that's on the end of it. That's a standard form factor that was designed for flash, and about the size that we wanted. And one really cool thing about this Compact Flash standard that the flash people were using, is that for whatever reason, they chose to use a PC-AT interface. In other words, precisely the same interface which became IDE, which is what hard disk drives were using. So here, there's this flash standard out there in this little form factor that already has the right interface on it for a hard disk drive. The only problem was that the version that the camera people were using, and this was used in all kinds of devices, was too thin for us. It was 3.3 millimeters thick. And we could not get our Microdrive to work in a 3.3 millimeter high form factor. But we could make it work in 5 millimeters high. So we actually approached the Compact Flash Association, which was the standards association for this Compact Flash form factor, and asked them whether they would be willing to make a variant that was thicker. And we did exactly that. We came out with what was called the Type II compact flash form factor, which was the same in every way, but 5 millimeters thick. And this was really done for us. And I was the chair of the Compact Flash Association working committee for that Type II standard. So we worked through all of the stuff and wrote all the specs and all the definitions. And in the end, the association adopted that and created it. And then camera-makers started to make their slots compatible with both the old Type I Compact Flash standard and the new Type II Compact Flash standard. And the first kinds of devices that really became a home for the Microdrive were digital cameras. I still have one, to this day, that uses it. And I still use it all the time. But there's a little Microdrive in here with the Compact Flash form factor. If you buy one today, they no longer use this form factor, and you can't use the Microdrive, nor should you. Because flash is a much better idea now. But around the year 2000, and this was the type of camera you could buy. And they have these slots. And like I say, just for old time's sake, I still use that camera all the time. And the Microdrive in it still does actually work, believe it or not.

So we created that new version of the compact flash standard to house the Microdrive. It went through a number of different generations. You can see here's a six-- no, this is a four-gigabyte version. And I think the slide that was shown earlier actually shows a six-gigabyte version. I think that was probably the last. The first one was the 340 megabyte Microdrive. Now, just one more word about the marketplace for Microdrive. Because this was quite fun to see. Cameras were an interesting marketplace for us, but it wasn't an overwhelming success. There weren't enough camera people wanting the extra capacity to buy enough microdrives to really make that into a profitable business for IBM or what later became Hitachi's product. But what was happening at that time was music players, digital music players. At that time, we called them MP3 players, were becoming very popular with young people. And they tended to use just a few tens of megabytes of flash. Because that's all they could afford. They were trying to make a little device that might cost 50 or 100 dollars, and at that time, you couldn't buy very much flash memory. So storing music in the MP3 format is one megabyte per minute. So if you made a 30-megabyte version, you'd get 30 minutes of music on your little music player. And that, in fact, was very popular in the late '90s. Little players with that kind of capacity. Now, we recognized right away, well, people would sure like a lot more music on their portable player than 30 minutes. So this looked like a great space for the Microdrive. So Tim Reiley and I, with the help of some of the marketing people from the mobile HDD business, and some people from Japan. We went around and hawked this Microdrive to consumer electronic companies. We visited Sony. We visited Matsushita or Panasonic. We visited Thompson Consumer Electronics in Europe. And we also visited Philips in Europe. These were the really heavy hitters in the consumer electronics world. They were all making digital music players at that time. And we thought they should love our Microdrive and buy a ton of them and build these drives. And not a one of them bit the bait. They all had their own plan. Their plan was wrong, but they all, together, did exactly the same thing. They knew that music players needed more capacity. And they looked at the marketplace and said, "What can we do really cheaply?" And they decided that recordable CD-ROM was the answer. And so all of those companies introduced recordable CD-ROM music players in around the year 2000. And in the long run, every one of those failed to succeed after sort of a soft start. So we were all quite dejected. We were really disappointed that they couldn't share our vision and make this cool little device with a hard drive in it. It might cost a few hundred dollars to do all of this, but you'd have a ton of music in there and people would love it. None of the consumer electronic companies went for this. And so we were depressed and nothing happened.[14] But then much to our surprise, Apple comes out with the iPod Mini.

[14] [Interviewee's note] There was one very small company called e.Digital that actually introduced a Microdrive-based music player to the market called the MXP 100. However, they were completely uknown and lacked the marketing muscle of the larger players; their MXP 100 sold poorly and did not survive in the marketplace.

And this really came out of the blue. At that time, Apple was not the wealthiest company in the world. It was a company in fact in quite significant distress. The 1990s were not a good time for Apple. And Steve Jobs was back at the company and they were doing well. And they were coming out with some very innovative and interesting products. There certainly was nothing like an iPhone on the market yet at that time. And even for them to introduce a music player was very surprising to us. But a good vision on Steve Jobs' part. And they came out with this little player, and they put the IBM Microdrive inside it, and eventually the Hitachi Microdrive inside it. And as Apple is wont to do, they made various colors on the product and packaged it beautifully. This was a real hit. This was the most popular music player on the market for a while. And our Microdrive was a big part of it. And those were successful years for the Microdrive and profitable years for the Microdrive. I do still have one of those here that I use once in a while. Not so often, because I don't listen to that much music in this particular way. But it's fun to have the little device and still use it.

. One moment that I remember very distinctly about all of this, I was in Fujisawa on one of my Japan trips, as I had done many at that time. And on a Saturday, I was looking for something to do. And that particular Saturday I was walking around the Ginza district in Tokyo, sort of the high-end shopping area. And lo and behold, of course, there's an Apple store there, because they were just starting to plant themselves in all the really cool shopping areas of the world and put their stores in. And this was the weekend that the iPod Mini was introduced in Japan. And as often happens for Apple products, there were people lined up outside the store and around the block. And I was just really struck by that moment, because this was a consumer electronics device, not made by a Japanese company, and in Japan, they were lining up around the block to buy this thing that Apple built and had our drive in it. At that time, this was still the time when as far as consumer electronics was concerned, the view was simply that Japan had won the whole show. There weren't many other companies in the world other than, say, Philips in Europe, that were making consumer electronics devices of any kind. Radios, TVs, tape players, whatever they were. All of this stuff was owned by Japan. And yet, if you look today, Japan has lost almost-- well, a great deal of the consumer electronics business. And companies like Apple took a great deal of it. The phone-makers basically stole the whole show. I guess they still make TVs. They do that well. But this was very interesting. This was the beginning of a trend where a consumer electronics device not dreamed up in Japan became a worldwide hit and a technical success. And so yeah, that was a moment to see that thing introduced in Japan and how popular it was.

Roger: Thank you, Tom. I really wasn't aware of a lot of that story and the detail. It really is a remarkable story. The thing I remember in particular is I believe IBM had set up some large-scale production of these drives, only to learn that Apple had decided to use flash memory in their next product. So tell us about your perspective of this. It all seemed to happen very suddenly.

ALBRECHT: Yes.

ROGER: Go ahead, please.

ALBRECHT: Yeah. So that's an important part of the story to tell here, because of course the Microdrive is no longer a success. In fact, its success was rather short. Just a few years that Microdrive was made and maybe only one or two years that it made any money. So yeah. What happened is that after the iPod Mini, Apple came out with the iPod Nano or something like that. I forget exactly what the next generation was, or maybe it was even a follow-on of the mini that just used a flash memory. I was a little bit surprised by this, because we still had higher capacity and a lower cost per megabyte than flash. But apparently Apple understood that the companies wanted a certain amount of storage, but it didn't have to be huge for all the customers. They still sold some larger ones that had capacity and had hard drives in them for a while, for a year or two. But the market shifted towards smaller music players that used flash. So that was interesting, that there was not unlimited growth potential, the more gigabytes you shipped, the more they will buy in that segment.

And then of course on top of that, right away, came the next phenomena. And that is, although during the mid-1990s, we had our golden era of areal density growth, which you can also interpret as the golden era of decreasing costs per megabyte or costs per gigabyte for hard disk drives, that slowed down starting around the year 2000. And flash memory was really getting into its stride. And it had its golden era of increasing density and decreasing costs right around that time. And just as our curve was slowing down, theirs was speeding up. And that indicated what would happen to the Microdrive that eventually, in small form factors, magnetic data storage could no longer compete with solid state. And the reason for that is one simple fact, that with a solid-state storage, if you want to scale it small, let's say you want to buy just one gigabyte of storage, you can make that very cheaply. And there's very little fixed controller and packaging overhead that goes with it. Maybe as little as one dollar that needs to be wrapped around the flash to make a device that you can sell. But a disk drive is different from that. A disk drive, at least at that time, had about 35 dollars of cost. And that was the same for the Microdrive as it was for larger drives. Maybe a tiny bit cheaper, but not much. A lot of that cost was in electronics. And although a lot of

 © 2021 Computer History Museum

mechanical components were scaled smaller, smaller doesn't always mean cheaper. Sometimes it means actually a little harder to make. So the fixed costs wrapped around the storage medium in a hard drive was around 35 dollars, and that meant that you couldn’t scale it down to a cheap price point. It only made sense to put a fairly large amount of storage in a hard disk drive. And eventually, what that means is that you have to amortize that fixed cost over a larger and larger piece of media, as flash memory consumes the low end of the marketplace, eats up more and more of these smaller devices. Because if its low overhead and its relatively lower cost per byte. So those two dynamics changed the future. Toshiba had actually tried to come out with a drive even smaller than our drive. They had a 22-millimeter drive. That failed immediately and never had any success. And shortly thereafter, the Microdrive failed in terms of no longer being continued as a product. And as you know, there are tons of music players and iPhones out there now. And there is not a disk drive in any of them.

ROGER: An interesting aside, Tom, I took my wife to the Computer History Museum when it was opened. It was only a couple of years ago. But she came away with a piece of jewelry, which was a Microdrive hung on a chain. And it made quite an attractive little piece of jewelry. I don't know if you've seen them.

ALBRECHT: I have not. Was it a real drive or just a facsimile?

ROGER: I'm not quite sure. I think we have-- I think we bought three of them. I don't know if she gave them to her friends or not. But I'll look at it more closely. I'm not sure, but it was an interesting discovery was that.

ALBRECHT: Very interesting. For those of us that like hard disk drive technology, that would indeed be a fun piece of adornment to have. I don't know how wide its appeal would be with the general public, but somebody who visits the Computer History Museum would be a good candidate to go for something like that.

ROGER: So now for something completely different Tom. This brings us up to 2002. And you took a little adventure in 2002 to IBM Zurich Research. Tell us all about that.

ALBRECHT: Yeah. So this was the Millipede Project.

 © 2021 Computer History Museum

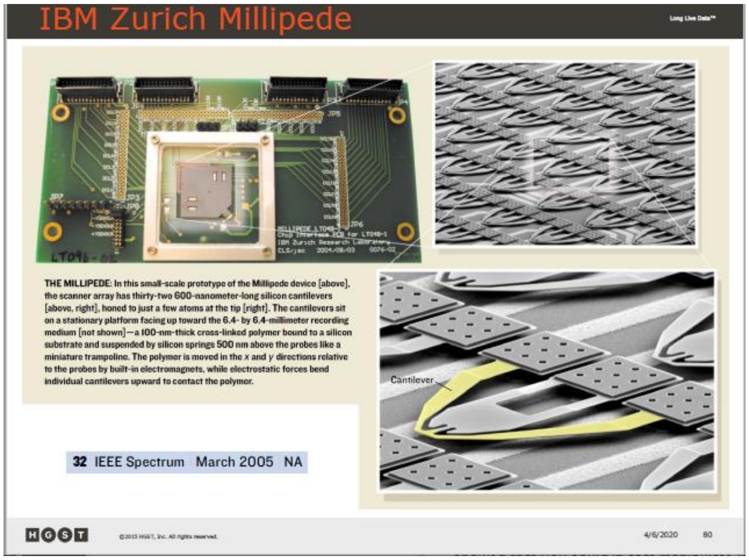

So this was a interesting technology that drew on a number of areas that I had some familiarity with. The team was the IBM Zurich team that was very familiar with scanning probe microscopes. It included Gerd Binnig, the inventor of the scanning tunneling microscope and atomic force microscope. And they used microfabricated arrays of cantilevers with tips on them, just like an AFM, except there's now a large number of these things, to store data on small pieces of polymer recording media. Now, it had been a dream of scanning probe people way back in the '80s, that eventually somebody would figure out how to use scanning probes to store data at a much higher density than ever before. And you could even envision storing data as single atoms. And for example, the work of Don Eigler at IBM Almaden later showed that you could in fact manipulate individual atoms with a scanning probe microscope tip, put them where you wanted to. And if you could put atoms where you want to, then you can fashion that into some kind of data storage. That, in itself never became practical because it was too slow and probably not reliable enough to move individual atoms.

But the guys in Zurich had a little different approach that they were excited about. And this was thermo-mechanical recording. And that is the tip on this sort of cantilever, you see sort of this yellow shaped V in the lower right there. You can't see the tip there, but on the bottom side, there's a sharp tip. That could be bumped against a piece of plastic and also heated at the same time. There's also a little heating resistor built into that cantilever. And if you do that and the polymer is properly engineered to have the right glass transition temperature or melting point or whatever, you can leave a little divot in the polymer film on demand. You can move the probe to where you want and make a little divot. And this divot is quite small, maybe a few tens of nanometers in diameter. And not only that, you can erase this little divot, make it flat again, by just hovering the tip near it while it's hot, and it sort of reflows the plastic or polymer material. So you have a writable and erasable storage medium that you can store data at a very high density. So they wanted to commercialize this and turn it into the next big storage technology. Now, even before I went there, I knew this had a lot of risks. That this was not by any means a slam-dunk success. A lot of technical challenges to make this kind of recording really reliable. And while I was there, I actually came to see very clearly that this would have difficulty competing on price. But we'll come back to that in a

 © 2021 Computer History Museum

moment. As a technical feat, and here I don't take so much credit. The team that was already in Zurich[15], just a wonderful, talented bunch of people in microfabrication and scanning probe technology, built these beautiful devices and these arrays of cantilevers with the sharp tips and even made controllers for them that would talk to all these things, and a nice little microfabricated or at least miniaturized actuator that would scan all of these things relative to their pieces of media. And in any case, pushed it quite far along, towards a relatively mature technology.

We eventually, while I was in Zurich, were able to do a recording demonstration modeled very much after what was going on in the hard disk industry at 600 gigabits per square inch, which at that time was significantly higher than what a hard disk drive could do. Later, they did one even over a terabyte per square inch. But, it wasn't enough. One thing that became clear, one thing I helped them do is sort of envision what would the storage device look like? And sort of modeled on our Microdrive success, we identified what was, at that time, a popular flash form factor, and it still is today, the SD card. And it looked like this device could fit into an SD card and you could store a certain amount of data in there. But, right away, you could see the problem that was eventually was going to kill this thing. And it's actually the same problem that killed the Microdrive. And that is there's a certain amount of overhead cost here. And it's actually very similar to a disk drive. There's about 35 dollars of cost in this device to make everything you need to make. All these microfabricated arrays and whatnot, and the electronics that drive it all, and the actuators and the packaging. And the area of media is very small. I forget what the total area available in this thing is in the square millimeters. But the problem is even though you have areal density that's significantly higher than hard disk drive, you have not made a device that stores bytes more cheaply than a hard disk drive. In fact, you have not even made a device that stores bytes more cheaply than flash memory at the time we were doing this. So we could see that we were going to need to hit some pretty astronomical areal density numbers. A lot higher than what hard disk drive people were used to hitting, to make this thing commercially available and to be able to compete with flash.

Alternatively, you could have tried to figure out a way to put more media in this. Make a larger chunk of material, but that really went against the whole architecture envisioned here. If you were going to make some kind of larger media or removable media, there was no easy engineering path to do that. And so unfortunately, this never went to market. The team was quite good at solving the technical problems. There were tip wear issues and there were error rate challenges. They did a pretty good job of addressing all that stuff, but at the end of the day, the economics did not work. And this was, to some extent, a bit of a shocking failure on the part of the team. These were smart people, but I do remember very distinctly Peter Vettiger and Gert Binnig talking to me about this project one day, very proud of the areal density they could hit. And the comment they made was, "How could anything compete with this? We have the highest areal density." But as you know, it's not areal density that the customer cares about. It's dollars per gigabyte or cost per bit. And these devices did not succeed well. They had simply observed that those of us in the hard disk drive industry used areal density as our metric of success, and it's still done today. If

[15] [Interviewee's note] The Millipede team in IBM Zurich was led by Peter Vettiger, with Gerd Binnig providing technical guidance. Several other prominent team members come to mind – Evangelos Eleftheriou (electronics and data channel expert, and later the manager of the team after I left), Michel Despont and Ute Drechsler (amazingly skilled nanofabricators), Mark Lantz (actuator designer), Urs Duerig and Bernd Gottsman (polymer and thermomechanical recording experts). There are more whose names escape me now.

you talk about having a terabit per square inch, if somebody's got two terabits per square inch, they can probably beat you commercially, provided the costs are all about the same. And as long as you're comparing devices that are very similar, a disk drive to a disk drive, then simply looking at the areal density probably gives you a very good insight into which one's going to win. The one with the higher areal density. But when you compare a disk drive to a milliped to a flash, you should never talk about areal density, because it is not the metric that the customer cares about. They care about dollars per gigabyte, and that's a very different calculation and does not track areal density one-for-one when the host device is completely different, particularly when the area of available media is not at all the same. And here we had a very tiny piece of media, so you simply could not store enough on this to really outshine what flash could do. And of course not only this failed, but eventually the Microdrive failed which was sort of the next larger increment over this that was sort of trying to do the same type of storage. So that's a little bit the story of what went wrong and why this couldn’t succeed as a data storage technology. That was quite clear to me at the end of a little less than two years there. And although they would have liked me to continue there, I also could see that I didn’t want to go through whatever pain was coming here to figure out how this technology would be used, eventually it was used as part of a new 3D printing technology that's actually commercialized today in Switzerland. But it was never commercialized as a data storage device.

ROGER: I hadn't realized that, Tom. So there was some kind of spinoff from it, sort of microfabrication.

ALBRECHT: That's right. The company is called SwissLitho, and I read within the last year that they were now acquired by some larger company. I forget what it is offhand (Heidelberg Instruments), but it was a success. They managed to produce that startup company based on this technology and turn around and sell it to a larger company.

ROGER: You never know quite where these things will lead. So you came back and spent about the next ten years on a very major project for IBM and HGST. A very challenging project and made a lot of advances in that work. Tell us all about pattern magnetic recording, or patterned media.

ALBRECHT: Sure. So this project, from a technology point of view and sort of a technical accomplishment point of view, was the most exciting project that I ever had the opportunity to work on. To really get out and do things that were far and away different from what anyone had done before. A little background on what patterned media is and why this project was started. Of course, Roger, as you're well familiar, around the year 2000, people are becoming very well aware of this phenomenon of thermal stability as a limitation on the areal density that can be achieved with magnetic recording. And we're starting to see the improvement per year in magnetic recording drop from that peak of 100 percent per year in the mid-90s to smaller and smaller numbers. I don't know what it is now, but it's maybe ten percent per year or five percent per year or something. And a key factor, of course, is that-- well, here I can step all the way back to something like the Tahoe Contact Recording Project. In the end, magnetic spacing was not the key issue limiting areal density of magnetic recording. It was important to get that spacing down. But eventually we hit the wall with thermal stability on granular media. And that is that standard sputtered media, when you produced it, made an array of grains of various random sizes and shapes. You tried your best to control it. You tried to make all those grains as uniform as possible, and very importantly they all needed to be large enough to be thermally stable. If they get too small and they have magnetic properties such that you're still able to write them with a normal read-write head, then once they fall below a certain number of nanometers, they erase themselves spontaneously over a time period that's too short for reliable data storage. So the industry recognized this, and although there were some very nice technologies that were successful in extending conventional granular media, mainly perpendicular recording and then things like anti-ferromagnetic coupling layers, very good inventions and really helped extend this game for another ten years. Eventually all those levers were pulled and there wasn't more to do.

So there were two ways to solve the problem. One is, well, let's go to a type of media that's magnetically stiffer. Something that can't be written by a normal read-write head, that would, even at very small grain size, be thermally stable at room temperature and successfully store those bytes for many, many years. And that, of course, became heat-assisted magnetic recording. And that is the one which, at least up until now has been the early winner and the technology that the industry is really trying to get commercialized successfully and launch in a big way. And there, they use special granular media. I guess I probably won't venture to say that I know what they're using anymore on that. But I guess cobalt platinum at one time, but I'm sure there's new formulations that I don't know anymore. But indeed, it can't be written with a normal head. The coercivity is too high, but that's exactly what you want for a very small grain that can support super-high recording density and stay thermally stable. They put a laser in the head to heat up the media while they're writing. And that's all a very challenging technology integration package, but the industry is working very hard on it and it is delivering.

The other approach, and this is the one that I worked on for ten years, is patterned media, which said, let's keep media that has a coercivity low enough that we can write it with the normal read-write head, but let's get away from the concept of sputtered media and the random grains that are produced by depositing a normal sputtered magnetic film. You get this distribution of sizes and shapes that creates what's called disk noise because the size and shape of those grains don't correspond very precisely to what you're trying to record and all of that works together against you. But if you can replace all those random grains sputtered in a blanket deposition process with individually fabricated magnetic islands that we control, we make them a specific size and shape and magnetic property and we put them exactly where we want them on the disk, that's another route to making an extendable technology that's thermally stable and could take you another decade or two in the magnetic storage business.[16]

---

[16] [Interviewee's note] The preceding explanation of the technical motivation behind patterned media is a bit difficult to follow (sorry!). Basically, on conventional sputtered (granular) media, there is a size distribution of magnetic grains on the disk. Each grain functions as an indpendent magnet. To record a desired data track on such media, the grains need to be substantially smaller than the individual bits so that the size and shape of the individual grains doesn't overly distort the shape and position of the recorded bits too much. If the grains are too large, the recorded bits will have ragged edges, which give rise to timing jitter and amplitude variation that can cause data errors when the bits are read back; this is referred to as "disk noise" that places a limit on how small the bits can be (the areal density) for a given magnetic film. Making the grains smaller reduces disk noise, but when the size of grains falls below a certain threshold, the grains are no longer magnetically stable – they can flip magnetization spontaneously, which degrades, or even completely erases the recorded data over time. To solve this problem, either patterned media or heat assisted magnetic recording need to be used. Patterned media solves the problem by replacing a collection of grains in a sputtered magnetic film with an individual, intentionally fabricated single magnetic "island" per bit. If the islands are arranged in circular tracks with uniform spacing of islands along the

So, we set about to do this. Now, already when we started, pattern media was viewed as perhaps the underdog technology. Many people in the industry favored heat-assisted magnetic recording. But I was willing to sign up. Maybe at the time I would have preferred to do heat-assisted magnetic recording, but that isn't the job I was offered. I was offered patterned media and we sank our teeth into it. We formed a wonderful team of about 20 people that really went after this. The key problem was always viewed to be a "Could you actually make this stuff with islands small enough and do that in a cost-effective fashion that would deliver a cost-effective solution at the very high density for the industry?" and my feeling was this is just a matter of building something. When you give engineers the problem set that's so well-defined, can you make this stuff? Can you make it cheaply? Can you make it small enough? It's very easy to judge failure and success and know whether you're doing what you want to do and everybody understand exactly what the problem is and over the years, traditionally, engineers have been awfully good at delivering on building stuff that seems impossible, given enough attempts and sustained effort.

So, we launched into that and tried very hard to do it and we did in fact come up with a particular method to make patterned media,

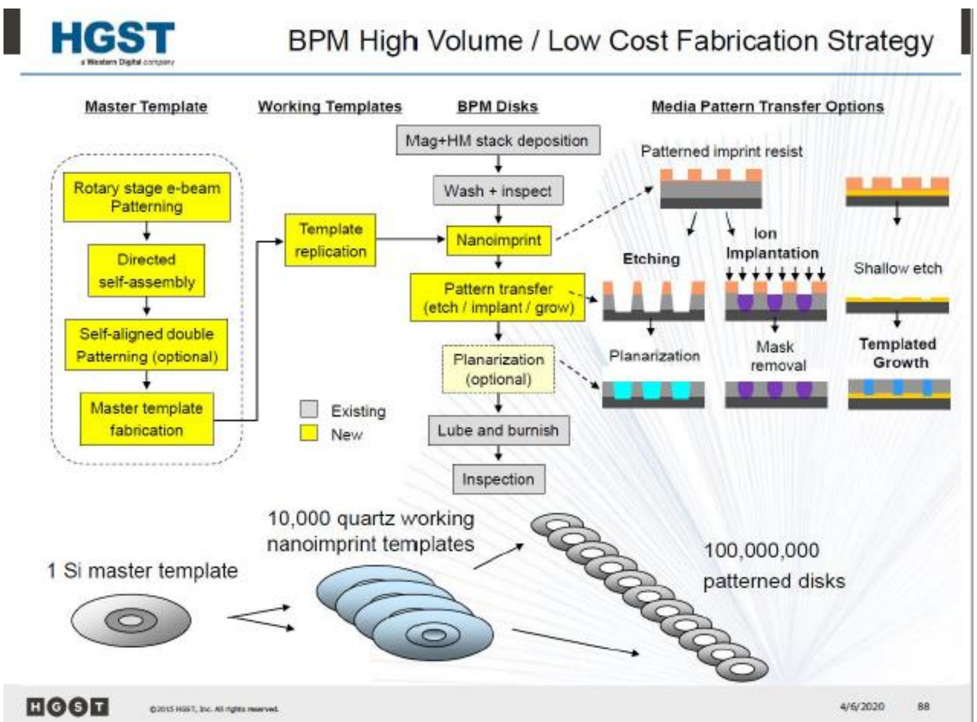

So, a lot of tricky technologies here coming together. But one reason why this was viewed as so daunting is that the islands that we needed to make on a disk, arranged in circular tracks needed to be a lot smaller than the smallest features the semiconductor industry was making on their silicon wafers at the

tracks, thermally stable data can be recorded with one bit per island. Since these bit-sized islands are much larger than the grains of conventional granular media, they are large enough to be magnetically stable at much higher density than can be achieved on conventional media. In theory this all works very well; the big challenge is cost-effective fabrication of patterned media.

time or even smaller than they could envision making by any means in the near future. So, normally, that should give you a little pause. The semiconductor industry has been investing billions for decades to make smaller and smaller ICs, but yet, we took on the task of making something a lot smaller than what they wanted to make and we did it by a bunch of technologies that are quite different than what semiconductor wanted to do and in fact, would be a viable choice to do it cheaply enough and it's simply the following.

You start with one master template that has the patterns you want on it. That can be very expensive. If it costs you millions of dollars to make one, you can still do everything you want to do, as long as you can copy it a lot and our intention was then to create tens of thousands of working copies of that and what are called nano-imprinting templates and then those templates can each print about 10,000 disks before they wear out and you can end up with 100 million patterned disks for a generation, which is enough. That's enough for a viable product generation. Nanoimprinting was already invented at that time. It was just coming on the scene and becoming of interest to the semiconductor industry and it was clear that it had the potential to print very, very small features. The challenge was to make that initial silicon master template and there, we had to go very unconventional as well. If you look at that column on the left of this slide on the dotted line[17], that's talking about the master template fabrication.

It started with rotary stage E-beam lithography. That in itself was quite a project because although E-beam lithography is out there, it can't make features small enough of what you need for pattern media. So, it alone isn't going to be the lithographic solution and not only that. Every decent E-beam lithography machine out there is in an XY rectangular format and we really need to make circular tracks on a disk and we cannot do it by patching together little rectangular areas because what's called the stitching errors. When you patch one rectangle with another, it creates little offsets in the patterns that would have in fact been fatal for us. We understood all that right up front, took the bull by the horns and persuaded people to invest in rotary stage E-beam lithography machines and we, in fact, worked with a Japanese company, Elionix, to create the necessary machines, which they did and we bought one and we had one in the basement at Yerba Buena Research Center.[18]

Then, recognizing that E-beam lithography could not make features small enough for us, we needed to multiply that density further. So, then we used what's called directed self-assembly. This was another technology in its infancy at that time, but we recognized that it had the potential to do what we needed. There are these chemicals called block co-polymers, basically as a polymer/chemical chain, where half the chain is made of one type of polymer and the second half of the chain is made of another type of polymer. If you choose those two polymers to be immiscible, they don't like each other in solution, they will phase separate, even though they're attached to one another and so, if you make a thin film of these things in liquid form on a disk and you anneal it properly, the ends of these molecules will self-align themselves so that it creates alternating stripes of material A and material B corresponding to those two polymers that you've combined and people had been getting good at doing that. They had various ones out there, polystyrene and polymethylmethacrylate are the two polymers that were very popular and quite

[17] [Interviewee's note] Within the dotted box on the left side of the slide.
[18] [Interviewee's note] Key members of our e-beam lithography team were Elizabeth Dobisz, Alexei Bogdanov, and Gabriel Zeltzer.

well-established and we made one important breakthrough for the whole self-assembly field. Self-assembly on its own produces relatively disordered patterns. They look like fingerprints. You do have these alternating stripes of the two different materials, but they are laid out like a spaghetti heap or they look like a thumb print when you get done and that's no good. We need these all to be straight lines or better yet, arced lines that are going to follow these circular tracks that we need to make. Now, they already had developed prior to us coming along methods to guide self-assembly and they called that directed self-assembly and they tended to use shallow topographical features, like a little step edge.

If you did this annealing of the block copolymer film on a substrate that had some straight step edges in it every so often, you could get these things to line up pretty well and the semiconductor industry was already looking into that as a method for making very small structures for their future. But they tended to make one to one guiding, one guiding structure for every stripe of the polymer that was going to come along and we speculated that that should be possible to have sparser guiding than that and not only that, to be able to do with chemical patterns rather than topographic patterns and this was successful. We worked with a team at the University of Wisconsin-Madison and Professor Paul Nealey to apply this idea to make a pre-pattern that had some stripes of different chemical species on a substrate and then when you annealed the block copolymer on top of it, it lined up with what was underneath it and did it in an orderly fashion to create structures much smaller than the guiding patterns. So, in other words, we could create a guiding pattern that only had a guide for every ten tracks, but then when we annealed the pattern, it would produce something that had a stripe for every track that we wanted. This was a very significant breakthrough and is now widely adopted by everyone doing directed self-assembly and is thought to be an important direction for self-assembly in the semiconductor industry going forward. So, that was really fun to be able to make a breakthrough in that area.[19]

Now, after doing all that, we still can't get the job done. We're not small enough. Block copolymers, as they existed a few years ago, couldn't go small enough to make what we wanted. E-beam lithography was way too large. Directed self-assembly brings us a step closer, still not small enough. We adopted one more technology that we stole from the semiconductor industry that is double patterning. They have a method where you can use single ridges and coat both sides of them and then do a bunch of etching so you can end up with the two side walls left from an initial ridge and that's a way of doing line doubling. That's in fact widely used in semiconductors today to make structures smaller than their lithography will support. So, we combined all three of these things-- rotary stage E-beam, directed self-assembly, and self-aligned double patterning to make the master template.

[19] Key members of our self-assembly team included Ricardo Ruiz, Lei Wan, and Julia Cushen. This team was absolutely outstanding; for a time we likely held the world record for making small patterns by self-assembly – a surprising feat for a small team from the hard drive industry in a field dominated by the semiconductor industry. This team, collaborating with Dr. Nealey at Univ. Wisconsin, also invented density multiplication via directed self-assembly using chemo-epitaxy, one of the most important breakthroughs in the field of self assembly in this time period.

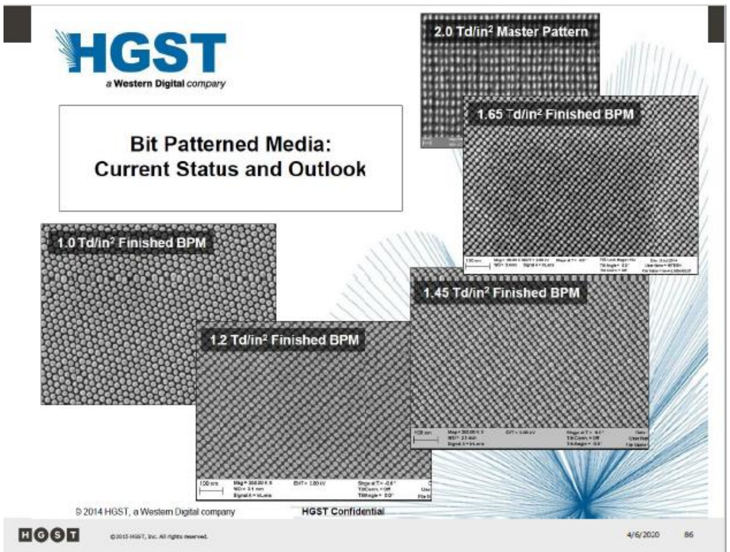

This one right here and these-- some of these are master patterns. For example, in the upper-right, that's a two-teradot per square inch master pattern and all the rest of those are actually finished disks. The only reason is that we were doing it a generation at a time and at the time, the project was cancelled, we had the two teradot per square inch master pattern produced but we had not yet successful transferred that via nanoimprinting to finished media. But this gives you an idea of how far this came along and how real this was. There's a bunch of things to observe here and that is-- especially in that two-Teradot master pattern-- self-assembly doesn't like to make rectangles. It likes to make circles. It likes to make hexagonally packed things or it likes to make straight stripes and we learned to do the self-assembly in a way that we could use the intersection of wide stripes and narrow stripes that were oriented perpendicular to one another to produce rectangular islands and that's how we got those. That turned out to be an important part of the strategy of being able to use this media in real disk drives because you want a bit aspect ratio that's a little bit higher than one and hexagonally, close packed media doesn't do that very well. So, here, we were working towards high bit aspect ratio media that would really work in a regular disk drive.

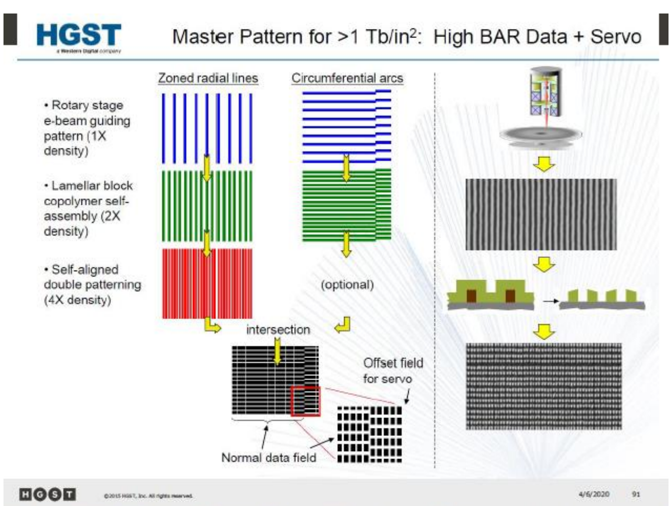

Here's a little bit about how do you really make a disk pattern with the circular tracks and even servo patterns. We made what we called zoned radial lines using all those techniques we just talked about and then we would make those lines finer using block copolymer directed self-assembly to double that density-- that's the green lines you see there-- and then finally, self-aligned double patterning to make them one more factor of two smaller and then at that point, you can make them about ten nanometers apart. So, that's the radial lines and then we came back and made circumferential lines that were actually circular, followed circular tracks and that was done with just the E-beam lithography and the block copolymer self-assembly. That got us small enough to hit the track densities we wanted to get and then with some processing wizardry[20], you could get those two patterns to form their intersections and create rectangles and then finally, down below, this will probably only be entertaining for people that know a little bit of the details of the innards of disk drives, but you could produce offset patterns to create servo patterns as well.

So, we really tried to solve all of the problems needed to make this technology viable for a disk drive. The high bit aspect ratio patterns, a plan for how the servo patterns would be taken care of, made the investments in the E-beam lithography tools. We bought the necessary nanoimprinting tools[21] and learned how to transfer those patterns into the disks, also very specialized processes for such small patterns. Furthermore, we were able to do that and transfer it to magnetic films that had a high enough uniformity so what is called the switching field distribution of the individual islands was tight enough. You want all these islands to be the same. So, as the head passes over it, you know exactly when the individual islands will switch and the neighbors don't switch when they're not supposed to and things like that.[22]

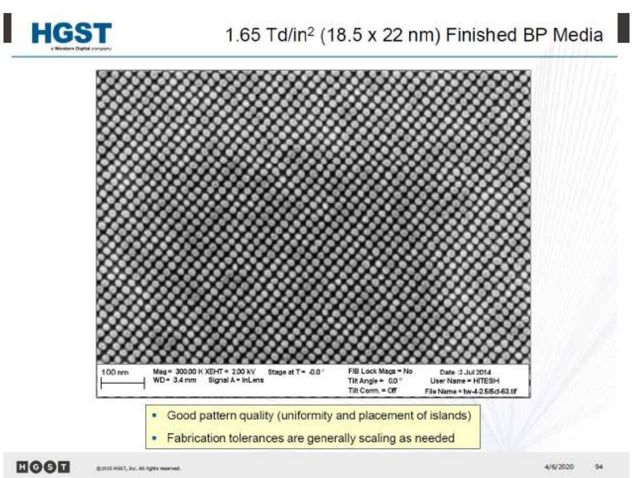

[20] [Interviewee's note] Ricardo Ruiz, Lei Wan, Harry Gao, Greg Doerk, and Tsai-Wei Wu were among those team members responsible for this "wizardry."

[21] [Interviewee's note] Molecular Imprints, Inc. partnered with us to develop nanoimprinting equipment specifically designed for patterned media. They delivered two generations of tooling to us, and we used their tools to perform thousands and thousands of imprints using templates that we fabricated ourselves. Tsai-Wei Wu was the leader of our nanoimprinting effort and he was the main working contact with Molecular Imprints.

[22] [Interviewee's note] Manfred Schabes and Terry Olson were magnetic recording theorists and magnetic modelers who helped us understand the requirements in terms of magnetic properties that had to be achieved. Olav Hellwig was our magnetic film deposition expert who created the magnetic films with the target properties to make it all possible.

 © 2021 Computer History Museum

Okay. Here's one. It's just a nice close-up of some of our 1.6 Teradot media. This was some of the very best media that we had made. This was also high bit aspect ratio. It's not so obvious, but it had an 18 by 22-nanometer bit cell and what you can see there is that the placement tolerance of the individual dots is actually quite good. Those islands are where they are supposed to be to within a nanometer or two. That's the kind of thing we really paid attention to. To really make a commercially viable technology, it's not just "Can you make something small?" but "Can you make it small and hit all the tolerances you need to hit?" Put all these islands, get these all the same size, close enough, get them in the positions they need to be, make them magnetically uniform, and have relatively few defects, only about one in a hundred islands can be defective and we, in fact, succeeded with that at 1.65 Teradots per square inch.

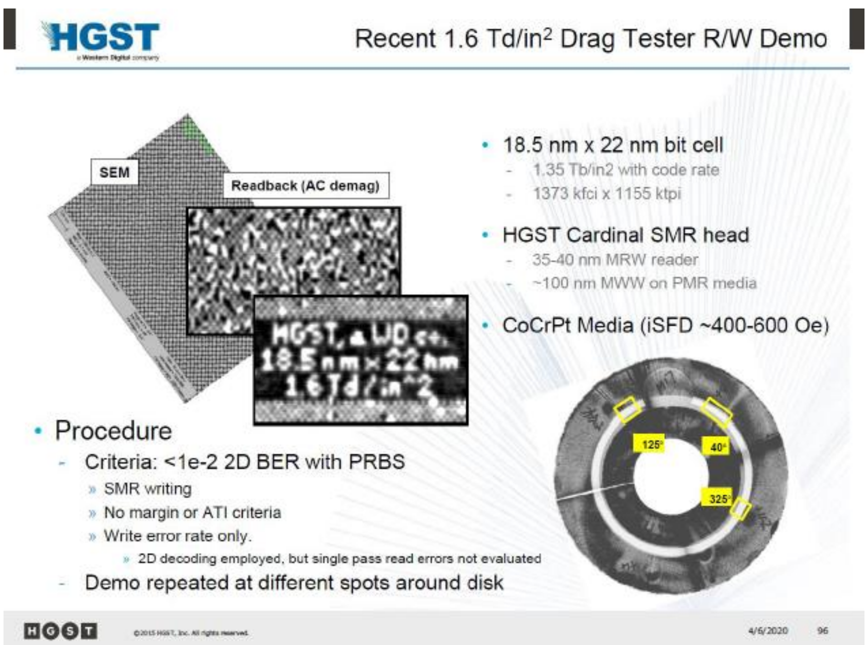

This shows some examples of magnetic recording on this media. This was the 1.6 Teradot per square inch media. At that time, that was much higher density than commercial disk drives could do and it was higher density than any laboratory demonstration with conventional media. You see sort of three pictures there arranged on the left side. One is sort of an SEM of a close up of the media, but the next two pictures there that look a little bit out of focus, those are actually magnetic read back using what's called a drag tester. You take a magnetic read/write head and you put it into my favorite device, a scanning probe microscope and you slide around on the surface of the magnetic medium and this way, without building a fully successful spin stand, you can actually do a read write testing and test things like error rate and adjacent track writing and so forth and we did achieve reasonable error rates, probably less than one times ten to the minus two, which at that time was felt to be sufficient to make a working storage device. The last picture there that shows writing, "HGST, a WD company, 18.5 nanometers by 22 nanometers, 1.6 teradots per square inch," that's just a pattern we wrote so that it would be visually confirmable that the error rate was really achieved. You don't have to do a computer analysis to see that all these bits were written the way they were supposed to be written. So, at this point, I was actually quite optimistic about pattern media because the drag tester was showing that you could demonstrate read write performance that we needed on this medium. Unfortunately, though, this was not at all successful in convincing the leadership of the company and this is one of the things that went wrong for our project is that we did not yet by this time have spin stand up and operational and functioning so that we could show this on the spin stand. My feeling was that if we had got far enough to do this on the spin stand, we actually would have succeeded or at least we would have persisted for quite a few more years because it's awfully hard to kill off a technology that's showing the highest aerial density of any technology and doing it on a spin stand. But nonetheless, we did not yet have spin stand and that's just sort of a sad story of project leadership, failures on my part and so forth. We tried to get the necessary resource for spin stand early on and we could never hold on to it. The company always took our spin stand person away and applied to conventional media and finally, we brought up our own guy, a young guy from scratch, straight out of graduate school, Daniel Bedau, very smart guy. He did all this stuff together with Mike Grobis, another very smart guy, and they got this drag tester working well enough and Daniel was very close to having the spin stand ready to repeat these kinds of measurements on real spinning disks and like I say, had we gotten that-- and I think we were within a year of showing that-- it would have been much harder for this project to go away. But where we were at that time and with Western Digital having bought the company and Western Digital not being a company at all with a history of deep investments in expensive new technologies to do the impossible, they really came in with a preconceived idea that patterned media had no future and eventually, they executed on that plan and we were put out of business. But the nice thing was they allowed us a month or two to write it all down and make a nice paper and we did publish that. If you look it up, I'm the first author and the title is "Bit Patterned Magnetic Recording: Theory, Media Fabrication, and Recording Performance," and it's in IEEE transactions on magnetics[23]. We really got to show everything we accomplished with no censorship at all because the company at that point had decided the technology wasn't going anywhere. It wasn't good for anything and they just let us publish all of it and so, the story is well documented there and has good readership. I can sort of track who all is downloading it and citing it and so forth and see that this still has an interest out there in the world and maybe it will come back at some point. So, let me stop there for a moment and see what other questions you may have about patterned media.

Wood: Thank you very much, Tom. That's a fascinating story. It really is. You overcome some real challenges and good use of new technologies there. My understanding is-- I'm out of date, of course, because I've been retired for several years, but my understanding is that this is still on the HDD roadmap, but in conjunction with heat-assisted magnetic recording. So, they'll be asking for even smaller islands and they'll be combining this with the laser-assisted recording...

Albrecht: Yes. Yeah. That is true.

Wood: If anybody's working on it actively, I do not know.

Albrecht: Yeah. I think the answer to the last question there is not very much, but a few comments on that situation are warranted. So, heat-assisted magnetic recording (HAMR) did sort of win this race and

[23] [Interviewee's note] IEEE Transactions on Magnetics (Volume: 51, Issue: 5, May 2015)

 © 2021 Computer History Museum

was selected as the next technology and WD's withdrawal from bit patterned media sort of finalized that choice and made it clear that there would be no competition for heat-assisted, although, I guess microwave assist is sort of a related technology that people are trying to do with some possibility of success. But the challenge for heat-assisted is its granular media. So, all the things that were going wrong and making it difficult to extend granular media, except for the thermal stability problem, are still with it and engineering all those random grains to do the things you want them to do has turned out to be very difficult for the type of media that's used for heat-assisted magnetic recording and as a result, the aerial density progress on HAMR media is actually quite poor. They're down in the 5% or 10% per year. They announce 20% but do not achieve it and so, that's too bad, but it does mean that the target we eventually need to hit with pattern media to come back in isn't that high. They're nowhere near getting to two terabits per square inch, at least for shipping product yet and that's where we were on this technology when we left off from it. So, I suspect that if this comes back in a combination with HAMR, it may be somewhere in the three or four terabit per square inch regime and based on our experience, I don't feel that that's an unreachable density point or dimensions that we could not deliver on. So, it may come back. However, I don't know if it will come back.

The whole hard drive industry is slowly losing the game to solid state. All of the form factors except server are now pretty much kaput. Now, that-- of course, don't underestimate that. The cloud has an awful lot of HDDs in it and it's been a good business, but still revenues are not increasing for the HDD industry and it's not a growing happy place at this point. So, I don't know whether the investment will materialize regardless of what happens here, even if it looks like it's a pathway that makes sense. HAMR does run into its own thermal stability problem somewhere in that three or four-terabit per square inch regime and invoking patterned media in combination with thermal-assisted writing would be a route to go well beyond that, probably to ten terabits per square inch or more. I've seen people throw out numbers like 50, which I don't take so seriously, but we'll see. It all depends on what happens with solid state storage and as I look at the whole field now as a spectator without a bias towards one or the other, it does seem likely to me that solid state actually has the bigger chance of achieving more density gain and cost reduction than hard disk drive at this point. I might be wrong, though. We'll see. So, I don't know whether patterned media will really come back. I do notice from what I've seen that some of the key patterned media people at Seagate, which was the only remaining company to be working on it, have left and are doing other things now. I just saw the leader of the Seagate patterned media team, David Kuo, pop up at Water Harvesting, which is a startup in Palo Alto. So, that also-- I haven't talked to David and I don't know what's going on there, but the patterned media future must not have looked particularly bright, even at Seagate.

Wood: People will look back on this interview in ten or twenty years or whatever, Tom, and they'll know the answer to what happened exactly.

Albrecht: Yeah.

Wood: You're right. I'm afraid that solid state is going to win the day eventually. The thing that brought it home for me was looking at flash memory and the fact that it's 3D storage, basically. They're up to, I think, 192 layers or something crazy like that and that's something that the magnetic recording business cannot do. So, it's going to be interesting in a few years to look back and see how things have evolved, but as you said at the moment, I think 80% or 90% of the data on the web is still on hard disk drives and

 © 2021 Computer History Museum

is changing only slowly. So, we will see. Okay. Just to finish off, you left HGST or WD and joined a startup company. Just for completeness, do you want to say just a few words about that?

Albrecht: Yes. I'll be happy to say just a couple of words on that. The first thing I would say is that the run that I had and also that you had, of course, Roger, in the hard disk drive industry, this was a wonderful run to see, to live through these years and experience all these new technologies and the phenomenal growth, that 100% per year we did in the 1990s never was matched by the semiconductor industry at any time and that was exciting to do all that. So, no regrets at all. I look back on that as just a great experience to have gone through, also the fact that it's a mechanical device makes it a very interesting thing in the whole information processing world. There aren't many mechanical devices left in that sphere of influence anymore.

So, as far as what I moved on to-- once the patterned media project was laid to rest, I did look at some new projects in HGST and WD and eventually, some of my friends from graduate school wanted to start yet another AFM startup company. They did in fact start it a few years before I left HGST and I did-- I sort of helped them a little bit on the side, giving them some advice and pointers on designs and so forth, but it occurred to me at that time that I'm living here in Silicon Valley and I've worked for a big corporation almost the whole time, except for the little bit I did with Park Scientific Instruments back in 1989, I should probably try out the startup scene just to do what people in Silicon Valley do and since this fit my skillset quite well. It is atomic force microscopy. I had the right background in it. But it's a very interesting new twist and I didn't talk so much about why I didn't continue AFM at Almaden back when I first started, but the reason was because Dan Rugar was going off into a brand new, very ambitious project, which was going to use NMR spectroscopy to identify materials, perhaps down to the single-molecule level and they invested a lot in that project. I knew it was a really big project. I chose not to get involved at that time, but the desire-- I certainly understood why they wanted to do that. That's always been sort of a Holy Grail of scanning probe microscopy is "Can we not just see the stuff, but can we actually identify the individual molecules and atoms that we're seeing?" and the particular NMR approach that they were going after at that time never really turned out to be viable. It's just too hard and it required too much sensitivity, not enough signal. NMR signals are really, really weak and they couldn't do it.

But what I'm doing now is looking at this in a different way and that is to combine infrared optical spectroscopy with atomic force microscopy and it turns out that if you shine infrared light on a sharp metal tip, you get what's called tip enhancement of the electric field that's in that light and it concentrates the light fields into a very tiny zone right at the end of the tip, nanometers, just what we want and people had already found that you could use this for spectroscopy to try to identify what's under the tip. You can check whether the particular molecules there absorb that radiation wavelength or do not. But up to this point, people have been doing optical detection. You'd shine the light in and the way that the light scattered off the tip is affected by whether or not the material on the substrate right below the tip or the sample material absorbs or not and they were able to kind of pick this up and it sort of works, but it's really hard. Then it turns out that yet another IBMer-- you'll probably recognize the name, Kumar Wickramasinghe-- he was an IBM Fellow from Yorktown and he was very active in the early days of AFM, just like I was, and he's a professor down at UC Irvine. He figured out that when light is being absorbed by a sample right below a sharp tip that's field enhancing your incoming light, there's actually a

 © 2021 Computer History Museum

mechanical force exerted on the tip. This is very surprising. Most people think this sounds impossible. You're going to do spectroscopy and sense light absorption with mechanical detection, but it actually turns out to be the best way to do this, the signal to noise ratio is far better with mechanical detection than it is by looking at the scattered light. So, this actually really works. We're down to about ten nanometers of spatial resolution with the ability to take IR spectra of what's under the tip and it looks like it can go farther than that and we'll keep working on it.

The reason to use IR is that molecular materials all have what's called a "fingerprint" spectrum in the IR range from about 5 to 13 microns' wavelength and we have this tunable laser like we do that operates in that range, then we can do spectroscopy and really identify these materials and you can imagine there are materials labs all over the world that would like to be able to do that well to figure out what contaminate is messing up your process or what in your device fabrication didn't work out quite the way you wanted it to. So, we're hoping it has a bright future in failure analysis and all kinds of cases where you want to do chemical analysis at the nanometer scale.

Wood: It sounds to me like you're still having a lot of fun, Tom.

Albrecht: Yeah. It's still fun. Still fun.

Wood: Any closing remarks to young people looking forward to a career? You seem to have had a very successful and enjoyable career?

Albrecht: Probably the little bit I can pass along, whether it's really useful or not, I have no idea, but growing up in an environment where I could simply fiddle around a lot with technical things, to work with tools and electronics and mechanical things, served me very well when it came to be time to do experimental physics and really make things work in the lab for products, learning to make things work very crudely on all the junk I worked on in high school and before was useful, and the other observation is that there are opportunities, special opportunities out there for people who are experts not just in one well established field. Like, if you become an expert in one particular branch of physics, you'll find there's lots of other experts in that same field and to really break out of that and have an influence, well, some people do, but your chances aren't that good. But if you happen to be able to cross over, you know a bunch of stuff in physics, but you have to be good at building electronics or you know some mechanical engineering or some chemistry, there are interesting crossover opportunities that nobody is properly exploiting because they don't fit neatly into one of the research siloes that are established by the known fields of research and sort of the copycat effect that occurs. The more people work on something, the more people want to join it and work on the same thing. But the real opportunities are in the things that they miss. So, my advice to young people is look for the spaces between the siloes and you might find something and good luck.

Wood: That's good advice, Tom. One last question-- do you remember your ham radio call sign?

Albrecht: Well, I've had two of them. My ham radio call sign from high school days was KA9BJO and then eventually, that lapsed and about ten years ago, I decided I should get the license again. So, now, I'm K6VL, Victor Lima, and I believe you're also or at least were a ham radio operator, right?

 © 2021 Computer History Museum

Wood: I'm G3WEW from the UK.

Albrecht: Yeah. Very nice.

Wood: I went back about ten years ago and got a US license, which was KK6RW.

Albrecht: Very nice.

Wood: I'm not active, I'm afraid. But it's nice to go back and I participated in some of the local ham radio meetings for a few weeks. But anyway, nice to meet you as a ham radio colleague.

Albrecht: Yeah.

Wood: Just to close, I must say, it's been delightful actually hearing you talk and hearing all the experiences you had and I'm sure that this will be one of the most valuable of the Computer History Museum oral histories all together. So, thank you very much, Tom.

Albrecht: Thank you very much. You've managed to tolerate this long discussion, I commend you.

END OF THE INTERVIEW

 © 2021 Computer History Museum