Oral History of Roger Wood

Interviewed by:
Thomas Gardner

Recorded via Zoom March 28, April 7 & May 14, 2020
Los Altos Hills and Gilroy, CA

CHM Reference number: X9301.2020

[hh:mm:ss] Timestamps approximately every 5 minutes ±15 seconds

Gardner: Good afternoon, I'm Tom Gardner, representing the Computer History Museum here with Dr. Roger Wood for his oral history, "Roger" hereafter. This is an unusual oral history. Ordinarily, we record them in a studio with professional videographer, but we happen to be in the middle of the COVID-19 virus lockdown, so we're not in a studio. I'm here in my home office in Los Altos Hills. Roger, where are you?

Wood: I'm in Gilroy, California.

Gardner: In your office, kitchen?

Wood: Office, yes, kind of office.

Gardner: Yes, my wife calls this my "pit." But it's where I do a lot of work for the Computer History Museum. This is one of many oral histories we've been conducting in the area of data storage. Roger has an extensive background in data storage. This slide gives a summary of Roger's background and experience, born in 1951, I guess that makes you almost 70 years old, with 40 to 50 years of experience in storage.

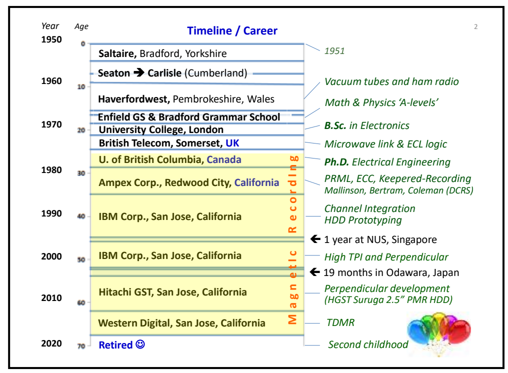

Wood: Yes, I highlighted in yellow an area there where I've been involved with magnetic recording. All the time from British Columbia down to when I retire. So it's more than 40 years or so.

Gardner: Started in England, educated in England and Canada. Then experience with Ampex, really not in storage, more in video recording, but definitely in latter part at Ampex a key storage technology. And then a long history of storage experience at IBM and then at IBM successors versus Hitachi GST, HGST, and then Western Digital. Roger, tell us a little bit about your background.

Wood: Okay <laughs>, as you can tell between the two of us we carefully prepared a whole series of charts, which is unusual. And part of the experiment is trying to figure out how exactly to cue each other on these charts. But this is sort of my early journey around the country.

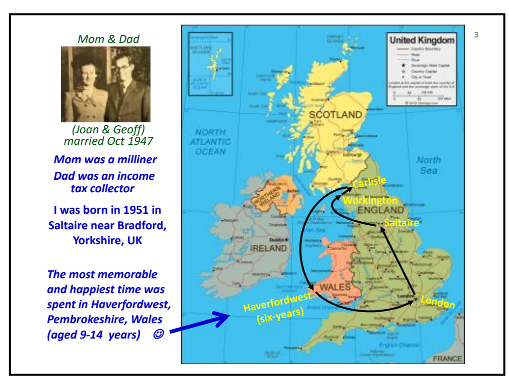

That's my mom and dad up there. Actually, I found a nice little picture of their wedding day. Mom was a milliner; that means she made hats, typically for the [posh] ladies around Bradford, Yorkshire, where we lived. Dad was an income tax collector, which didn't always make him very popular. As you said, I was born in '51. The little village or town was called Saltaire, and that's near Bradford, Yorkshire. And the name comes from Sir Titus Salt, he used to have a big woolen mill there on the River Aire, and so “Saltaire”. So this is where I started, and then went up to Cumberland, Workington, and then Carlisle for a few years. And then my dad got a transfer down to Haverfordwest in Wales. And I was very worried that I'd have to learn how to speak Welsh, but that wasn't necessary in Pembrokeshire, that's sort of half English and half Welsh. And as I point out here, I think those were certainly some of the most memorable years there, and probably about the happiest time. I think my family was particularly happy. My dad had got a promotion, so he was in charge of the tax collection office in Haverfordwest. And so he didn't have an immediate boss in the neighborhood, which must have been nice. And we kind of took advantage of the situation. Haverfordwest is sort of surrounded by sea and beaches, round about 270 degrees, because it's on a peninsula. So very often he'd try and come home early in the evenings. We'd all bundle into our little Ford Anglia and head off in some direction towards the beach. So those are some of my happiest memories there.

Gardner: Where were your grandparents from initially?

Wood: Grandparents from Bradford area, from Yorkshire area.

Gardner: So you're a Yorkshiresman way back.

Wood: I'm a Yorkshireman through and through, yes! Very proud of the Yorkshire Cricket Team.

Gardner: What was your grandmother's maiden name?

Wood: My mother's maiden name was Squires. And which grandmother do you want?

[00:05:00] **Gardner:** Both actually.

Wood: <laughs> Okay, stop and think for a minute. One was Sampson, and the other one was-- oh, I forget. That's awful, isn't it? I do know really, but it doesn't come to the mind at the moment.
[Violet Louise Wood (nee Sampson) and Alice Squires (nee Mathews)].

Gardner: We can always edit it into the transcript, if not the video**:** And I ask you not for the Museum so much as I ask you for your grandchildren to know who their great, great grandmothers’ names were.

Wood: One thing that you're already aware of is I did look up my ancestry, and it turns out that two of [my], is it, sixteen great, great grandparents were consecutive mayors of Bradford, Yorkshire. And it's kind of a nice story, because it was William Willis Wood’s son, and Thomas Speight’s daughter that must have met and liked each other enough, and hence they're in my family tree. Kind of a sweet story.

Gardner: And to set the record completely straight, you can look that up on Wikipedia.

Wood: <laughs> Right. I made a Wikipedia page for each of them.[1]

Gardner: Childhood interests, you must have played football, soccer in America.

Wood: Played soccer in England, and cricket, of course.

Gardner: Ah, cricket.

Wood: But I didn't play much soccer. It was rugby and cricket at my school. And I enjoyed the cricket, but I didn't enjoy the rugby, a bit too rough.

Gardner: You don't look like a rugby player.

Wood: No. <laughs>

Gardner: That is a brutal sport.

Wood: It is, yes.

Gardner: How did you get interest in electronics?

Wood: Right. <laughs> Yes, I started off in the days when portable radios had just come into vogue and they were based on vacuum tubes with a big 90-volt battery and a little one-and-a-half volt battery to drive the filaments. And I was sort of fascinated by that. And then somewhere in a library book I noticed something about electronics, and a picture of a triode. I went down to a local radio store, and asked them if they had a triode I could buy. And was told, "No, we have no triodes, but here's a pentode." And I happily took that away, and of course, it's quite easy to set up, if you can get a 90-volt battery and a one-and-a-half or a six-volt battery. And I started playing with this vacuum tube, and so that was the start of the interest. And then I became interested in ham radio. I don't know how I became aware of it, but I became interested in ham radio.

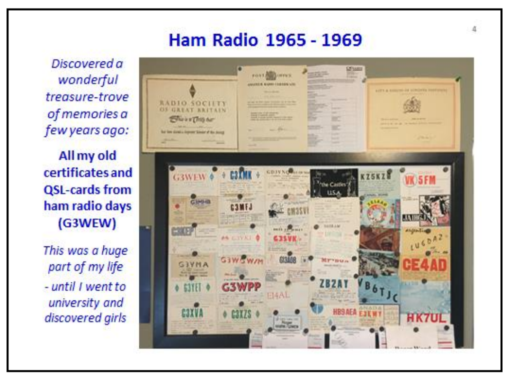

And that kind of consumed me for many years. And as I say on this chart here, we came across what I call a wonderful treasure trove. All my QSL cards, which are the cards that you exchange [to confirm] when you've had a conversation [over the radio]. So the cards in the picture here I'm rather proud of either because they're a long way away from England, so it's a long distance conversation, or they're my best friends - the G3s [call-sign prefix], particularly - these are people I knew in the ham radio club in Bradford. So that was a huge part of my life. I operated mainly on high frequencies, 20 meters, all in Morse code.

---

[1] William Willis Wood (mayor): https://en.wikipedia.org/wiki/William_Willis_Wood_(mayor) and Thomas Speight: https://en.wikipedia.org/wiki/Thomas_Speight

Gardner: I was going to ask: Morse code as opposed to voice?

Wood: Yes, so it was voice with my friends on 160 meters. And Morse code on 20 meters.

Gardner: How far did you get? What's your furthest QSL card you retrieved?

Wood: I think there's one there from Australia or New Zealand.

Gardner: It's about as far away as you can get!

Wood: It's just about the antipodes, yes, from the UK. I was especially proud of those. And as I say here, this is a huge part of my life! I spent hours and hours working ham radio, and then, of course, eventually went down to University and eventually forgot all about ham radio.

Gardner: Yes, before we move on, tubes rather than transistors. By '65, at least everything but the power should have been transistors.

Wood: Everything was tubes.

Gardner: Really?

Wood: I bought a transistor, just out of curiosity, and almost immediately blew it up! <laughter> I was absolutely furious! I like vacuum tubes a lot, and of course, there's a main power tube that does the transmitting.

[00:10:00] **Gardner:** True.

Wood: And if you carefully tune the tank circuit, as they call it, [that's] the tuned circuit on the anode of the vacuum tube. And you tune it such that instead of the anode being bright red, you tune it until it's a dull red. And that means that you've got about the right loading on the vacuum tube.

Gardner: Really?

Wood: Yes, exciting times, they're very forgiving, the vacuum tubes. So, as I say, I was not impressed with transistors when I first came upon them.

Gardner: So at some point around 1969, you went down to London to attend the University?

Wood: I did, indeed! And I went to University College, London. And the reason I went there was because I wanted to do Electronics. I didn't want to do Electrical Engineering, because that reminded me of sort of heavy-duty transformers and power stations, and electrical grids and things like that. I wanted to do Electronics. And the only place I could find to do Electronics was University College, London. I think in retrospect, it wouldn't have made any difference which college I went to. They all advertised Electrical Engineering. And of course, all included a lot of Electronics. And I've got a note there that I enjoyed the University time. And I met a very good friend there, I lost contact with him, unfortunately. But he was a London taxi driver! It was Barry O'Dwyer, and he was a mature student, I was just a youngster. So he taught me to drive in London, which is quite a challenge. And in return, I helped him with his homework in various classes we were doing together. So that was fun.

Gardner: Were you taught tubes ahead of transistors, or transistors ahead of tubes when you were at the University?

Wood: I think it was all transistors at the University.

Gardner: Not even tubes.

Wood: Certainly we played with transistors and built circuits for transistors. We did not do anything like that with tubes. We probably learned them. Just didn't actually practice any tubes.

Gardner: So I'm a few years ahead of you in the same educational experiences. At my school it was called Electrical Engineering, even though it covered Electronics. My class was distinguished as being the first class where amplification was taught with transistors before tubes. And then tubes were introduced as sort of a transistor with high impedance. But you were probably ten years later—it sounds like tubes as a curriculum item had disappeared. At least at the University College in London.

Wood: Yes, what was of interest, perhaps, is that shortly before going down to college in 1969, we got a letter saying, "Please come down two weeks early," maybe it was one week early, I think it was two weeks early. And we've organized a short course, intensive course in programming, programming FORTRAN. That's what we did. It was the first time I'd been exposed to anything like a programming language, and that was fun, too.

Gardner: Just FORTRAN, not Basic, too?

Wood: Just FORTRAN. I don't remember anything about Basic in those days, but it was FORTRAN.

Gardner: Okay, I think at this point, we would probably terminate the oral history for a time being, and then we'll resume it after looking at how well we've done, or not done.

Wood: Okay.

[00:14:00] **Gardner:** Hello again. This is Tom Gardner. I'm still in Los Altos Hills talking to Roger Wood, who I presume is still is Gilroy. How're things in Gilroy, Roger?

Wood: Things in Gilroy are just fine. I haven't been out of the house for three weeks but I keep myself happy with my garden and we have chickens so I get fresh eggs every day, which is nice.

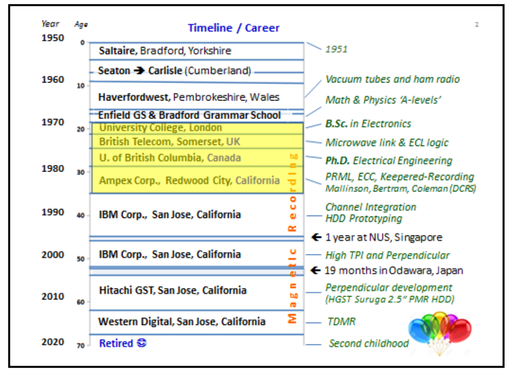

Gardner: We're both coping with the current COVID-19 pandemic, and this was an interesting test case for the museum to see if we can do oral histories using social media applications. We're going to resume the session with Roger's history from his career at the University College in London, continuing through his first employments, first at the British Telecom then going on to the University in British Columbia, and then his business career at Ampex. While you were at the university, you got married, didn't you, Roger?

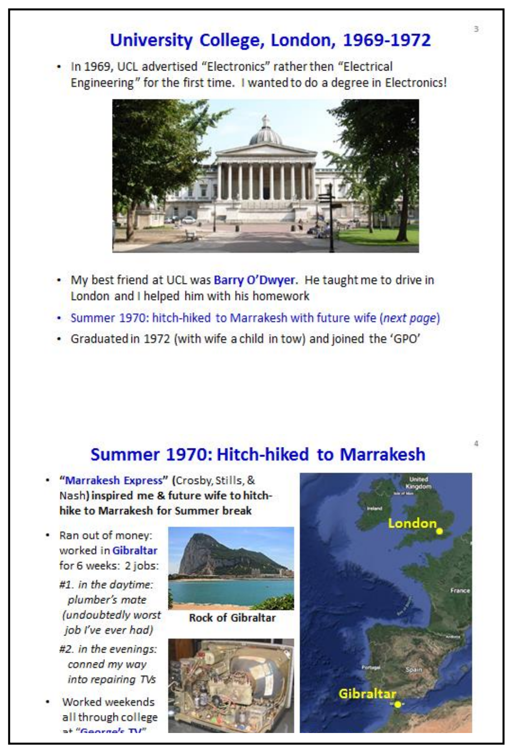

Wood: yes, I did get married at the end of 1970. In the summer of 1970, we had a wonderful adventure, and it was the summer that there was a song on the radio called "Marrakesh Express," which we enjoyed. So we were inspired to go hitchhiking. We hitchhiked all the way from London down through Gibraltar to Marrakesh. We'd run out of money in Gibraltar because we'd spent several days in Paris, which was probably a mistake, so I took two jobs and my future wife at that time was a bartender, a barmaid, I guess. The two jobs I took were a plumber's mate, and that was the worst job I have ever had. I won't go into detail, and the other one, which was valuable, in the evenings, I persuaded the local TV shop to let me repair TVs and somehow or other they believed that I could do it and I think I was relatively successful. I got 10 shillings, I remember, per television set, about a dollar per television set, and that stood me in good stead because I worked weekends after that, all through college at George's TV on Holloway Road. So that [big adventure] was the first year. At the end of that year, basically I got married - within a few months of meeting my intended. And just before I graduated in 1972, I had the first child.

Gardner: How many children do you have now?

Wood: We have three children and four grandchildren altogether, and there's a nice picture of them at the end of the series. The last slide has a nice picture of all my kiddiwinks.

Gardner: We'll get there. Now, you didn't quite wind up where you wanted to be when you graduated from the university.

Wood: No, I was hoping to get a job with the BBC. I'd always admired what they did. I was, of course, interested in television, and they demonstrated a longitudinal videotape recorder, which was quite a beast, when I visited for an interview, but they didn't offer me a job. I was quite disappointed. But British Telecom's [did], or the General Post Office, as it was then. I was a civil servant. I joined British Telecoms, and this picture here is a cross-section between two sister stations. These are tiny little outposts that were set up during the war to transmit back and forth. So I was posted in Backwell on top of this little hill and on the other side of the Bristol Channel over here was Castleton. Each station had about 20 people, relatively self-contained. It had a little workshop and there was a cook to cook us meals every day and it was all to do with testing the microwave propagation for a 20 gigahertz link between the two sites. And there's some interesting things here. The Bristol Channel has a huge tidal range, and if I remember rightly, there was discussion about whether the multipath interference from this specular reflection would change because of the tidal range. And I think the conclusion was, in the end, that it was shielded by these hills -- I won't call them mountains-- in between. But Bristol Channel has 47-foot tidal range.

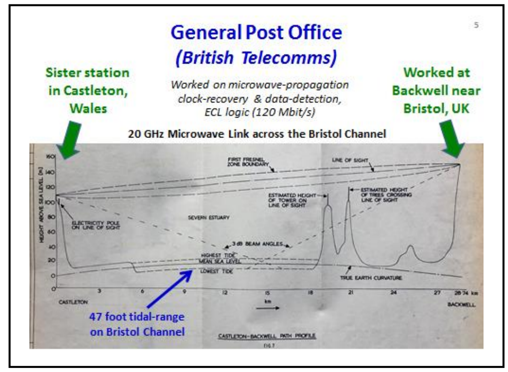

[00:19:00] It's the second to the Bay of Fundy, and it's not really a pretty sight. There's acres and acres, probably square miles of mud that are exposed at low tide, but that's where I worked for three years.

Gardner: Now, the GPO, isn't that the, sort of like the British FBI or NSA or something?

Wood: No, not at all. I think you are thinking of the GCHQ which was controlled probably in the military. I'm not sure how things were organized there, but no, it delivered letters and it ran the telephones. And they sent me on a training course for six weeks and I was climbing up telegraph poles and going down in manholes and using time-domain reflectometers to figure out where the break in the wire was many miles down the telephone line.

Gardner: So this was an information-transferring application, it says there, a microwave [link] with clock recovery and data, so this was digital transfer for information purposes: audio, video ...

Wood: Yes, It could have digitized video or audio, probably not video in those days but [it was at] 120 megabits per second, so that's where I learned all about ECL logic, emitter-coupled logic. That was my first introduction.

Gardner: After three years, you returned to the university and not in England.

Wood: I did. I wanted to go on an adventure. I'd always wanted to go back to graduate school and wanted a bit of adventure, as well. So I applied to several colleges in North America and ended up at British Columbia largely because they offered financial support, which was important but also family housing. So I've highlighted there on the picture, that's where I lived for four years, in a little, tiny, two-up, two-down, terraced house - in the end, with three children. And we had just a wonderful time there because it's beautifully located. There are beaches all around that peninsula, and I think in the background there, these are the mountains that we'd go skiing on in the winter. On this picture, that's probably Mount Baker in Washington, so it was a wonderful time we had there. The children were very young, so we'd go exploring every weekend in our little VW Westphalia camper van, which was fun.

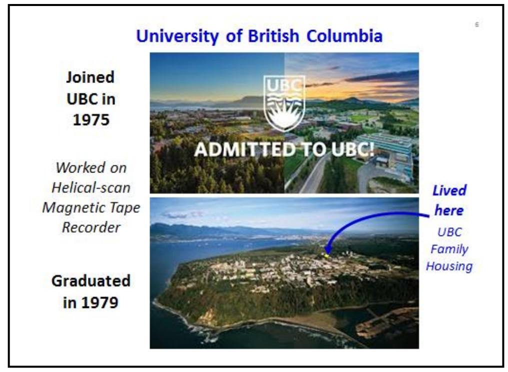

Gardner: Electronics engineering again?

Wood: Yes. I cast around for a good thesis topic, and there were several suggested between me and Donaldson, who was the advisor and they started with fairly exotic things to do with image processing and vision and at the bottom of the list was this boring topic called magnetic recording. And finally it bubbled its way up to the top of the list and partly because there was a local company called MacDonald-Dettwiler who were doing satellite downlinks, and they were very interested in using a tape recorder to record their high data rate, which was 20 megabits per second.

Gardner: This was a helical scan tape recorder that your thesis was on?

Wood: Yes, Type C [reel-reel format], I think they called it. And I forget who made the transport now [IVC 825A], but I spent the best part of four years working on that, looking at various ways of writing and detecting data on that. So it was a good introduction to magnetic recording.

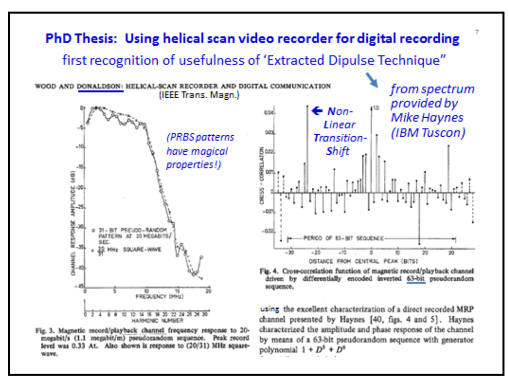

These are pictures out of my thesis. So I became very interested, actually, back at British Telecoms in "maximal-length pseudorandom sequences" because that's what we used to transmit back and forth between the two sites and these are just sequences that have very nice properties, very random-ish properties for testing out a system. And in fact, they have some quite magical properties which basically translate nonlinear effects into linear effects with a big time shift and I didn't appreciate at the time how important this would become. But this is an example here of an "extracted dipulse" which certainly was widely used and probably is still used to figure out how much distortion there is on channels. This [graph on the right] is the cross-correlation function. This peak [at position -24] here represents a nonlinear transition-shift involved in recording. That's the effect of a preceding transition on the placement of the following transition and the interesting part about this, perhaps, was that I was recording, I think it was, 31-bit patterns on the left here and I could not measure phase response which is obviously very difficult on a tape recorder. Mike Haynes at IBM Tucson had two heads on his machine and he wrote with the first head and a couple of inches downstream he read back. So he could get a phase response, as well, and you need both amplitude and phase response to do this cross-correlation properly and get this data. So I remember painstakingly looking at his printed copy of his amplitude and phase responses, carefully measuring them and translating them into numbers I could process and that's where this picture came from. Of course, years later, I met Mike and we became good friends but he asked me in great puzzlement: why didn't I just ask him for the data and it didn't occur to me at that time. I mean, these folks, when I was university, these famous names, they were kind of like gods. I certainly wouldn't have approached one of these people, but yes, that turned out to be quite useful, that extracted dipulse approach.

Gardner: So you were working with IBM Tucson even when you were at the university and before you ultimately wound up with IBM?

Wood: I suppose so, in a sense, but they were not aware that I was working with them. I was just looking at the data they produced.

Gardner: I'm still a little confused; elaborate a little more about what "extracted dipulse technique" means.

Wood: As I say, these PRBS patterns have these quite magical properties. If you take two of the patterns and multiply them together-- this is in the format where they're plus and minus one, each bit-- you get another pseudorandom sequence, exactly the same pseudorandom sequence but shifted to a different position. Multiplying two of these sequences together is a nonlinear operation, so what we're doing here, basically, is translating nonlinear effects into linear effects and this is time on the horizontal axis here and what happens is this nonlinear effect appears as an echo down here or as a pre-echo down here because it's translated the nonlinear distortion into a linear echo sometimes before the main pulse and you can also explain why these two things [responses] are large here. They correspond to different nonlinear distortions, but yes, I became fascinated with these sequences way back at the post office, and I still think they're absolutely fascinating.

Gardner: Now, you finished your education, got your thesis written, and you didn't go back to Britain. You went on to Ampex.

Wood: I did. Yes, I had a job held open for me at the post office. I felt quite guilty about not going back, but one day, I think it was a phone call I got from one John Mallinson[2] of Ampex saying he would be driving up there, and I'm never quite sure whether it was on a vacation or a business trip but he drove up there with his librarian who was his next-to-be wife. But he had seen some of the publications made during my thesis and basically he made me an offer there and then that I couldn't refuse and in Redwood City.

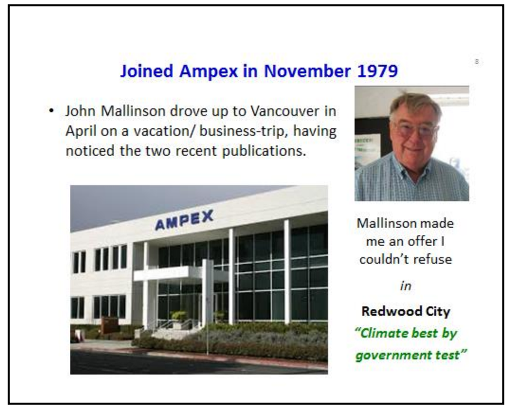

[00:29:00] And I remember looking at Redwood City and I somehow had this idea that it would be giant redwood trees everywhere but in particular I noticed that it claimed to be the "best climate by government test", and having endured England and Vancouver and all the rainfall and the snow, I thought it would be a good idea to move down to Redwood City. So indeed I did, I think, probably in November of '79, I joined Ampex.

Gardner: Mallinson was one of your several mentors and friends at Ampex?

Wood: Oh, absolutely. Yes, I kept in touch with John a long time. I won't talk to this chart for a moment, but yes, I kept in touch with John. We had dinner together every month or every couple of months throughout the time I knew him, and actually I traveled down to see him in Southern California just a few weeks before he passed away, sadly. But yes, he was a wonderful mentor and father figure, and he would spend hours explaining one-sided fluxes and all sort of magnetic curiosities that he knew about. He was very much a physicist and was a great teacher, as well. So yes, I really appreciate John Mallinson.

Alex Maxey, on this picture, I didn't know so well, but I spent time working on electronics basically, part of the electronics for this machine, the Ampex DCRS Digital Cassette Recording System, which turned out to be one of Ampex's most successful products. So it's a transverse recorder, and that's what this illustration shows. It's not a helical recorder. It's like the old-fashioned quad recorders. The heads are mounted on this drum that rotates and the tape is formed into this cup shape and the heads scan almost perpendicularly across the tape. DCRS had six heads on there, one in contact at a time. It ran 117 megabits per second in a single channel and had PRML on it and had a Reed-Solomon code on it. This is one-inch tape so [it's] different to the quad machines. Alex Maxey designed, I think, pretty much all the transports, all the quad transports for Ampex, and he was a well-known figure.

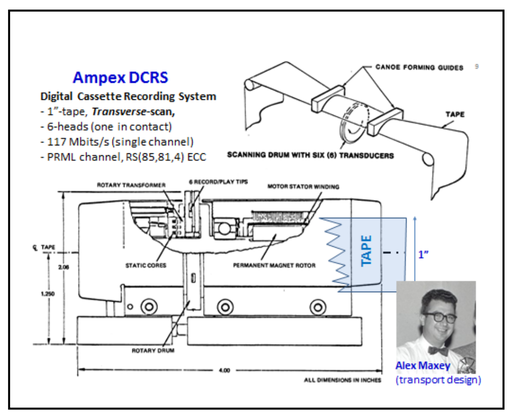

[2] https://en.wikipedia.org/wiki/John_C._Mallinson

Wood: DCRS was unique in a couple of ways. It was the first place that implemented PRML, which we'll talk about probably separately, but it was also one of the first places to implement an MR head, certainly the first on tape, but it was used in a very unusual mode. They wanted some sort of confidence in the recording. There was a unit which didn't necessarily have all the readback electronics in it, and that was for airborne applications. It just recorded, but they wanted some confidence that the recording was actually happening. The heads ran across the tape vertically and so the cross-track direction was down-track [along the tape] here and so there was an MR head here. There's a little inset picture here. There's an MR head here which picks up these tracks as they go underneath, and the longitudinal speed is quite slow. So an inductive head doesn't work very well but we had an MR head here and basically all it was looking for was that there was some kind of signal on the tape that had been recorded by the heads and it was successful in doing that. It wasn't trying to resolve anything that was being picked up. It just wanted to know was there a signal there or not but it was one of the first applications, probably the first application of an MR head on tape but it was a rather strange application.

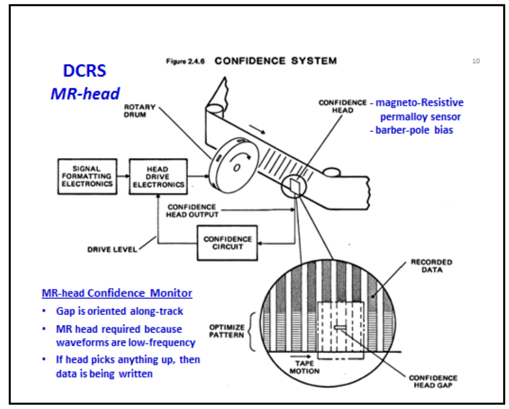

Gardner: Where'd you get the head from? Did Ampex make it?

Wood: The head was made at Ampex. We were able to deposit permalloy, able to do the etching to create the Barber-pole structure. Yes, it was all made at Ampex by the head department. And the AdTech department, which I was part of, did the sputtering. Yes, it was all Ampex, the whole thing.
[00:34:00] We also made inductive heads, but that didn't come to anything, thin-film-inductive heads.

Gardner: No.

Wood: And Ampex abandoned the disk drive business.

Gardner: Now, my understanding in tape technology is read after write is fairly common, but I guess it wasn't practical in the rotating head.

Wood: Yes, I think that's very true. You'd have to put 12 heads on there.

Gardner: So there was a conflict. It was the issue of 12 heads, not so much that reading after writing on a transverse head is not possible?

Wood: Yes. This was not what I would call a computer tape drive. It was a digital instrumentation tape drive, so it would be used for sort of long, continuous recordings and not so much for short records that would be rapidly accessed. It was a digital machine but not a sort of rapid-access, computer-style drive. It had a computer interface, but I'm not sure it could ever have been interfaced with an IBM computer. So the applications were mainly military and civilian aircraft doing test flights. That was one of the favorites.

They went into submarines, as well, so a lot of interesting applications. I think, at one point, most of the test flights were instrumented with this recorder.

Gardner: So it was long streams of data, and therefore there weren't the issues associated with short blocks.

Wood: Right.

Gardner: You have a block diagram of the recorder.

Wood: Yes, and this is somebody I'd like to mention, as well, Charlie Coleman[3] or Charles Coleman, Chuck Coleman, he usually went by, another one of my heroes and he basically was responsible for all the electronics on DCRS and you can see a rough picture here of all the bits and pieces. There's the confidence system which was the MR head up here and this is a segment used on the playback machines and on the record-only machines this section would be omitted altogether but let's see if we can pick out some of the features here: Record Electronics, so there was an error correction code which I'll mention again in a minute. The record electronics included encoding for PRML and the playback system-- where are we-- it's probably just [part of the] Reproduce Electronics, but the PRML channel is in here, as well, and then the error correction system with the large delays involved. And this is shown as analog in, analog out but it was much more general than that.

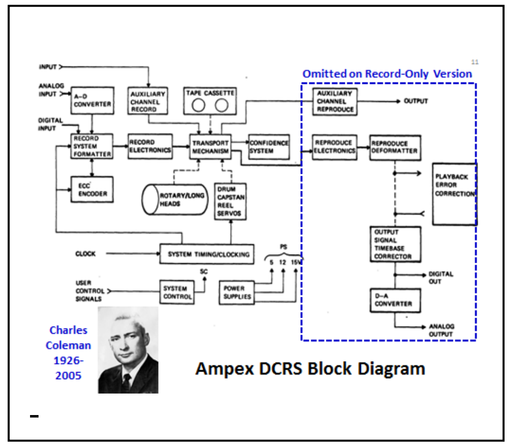

Gardner: So any key contributors that you'd like to mention before we move on to the next slide?

Wood: Well, me, of course.

Gardner: You, of course, and you did specifically?

Wood: Charlie Coleman really had the overall responsibility. We used to call him Mr. 0.1 Percent or Mr. 1/10th dB, something like that, because he was a perfectionist and his approach was interesting in that he liked to divide the system into blocks and each block independently had to perform perfectly and then he would put all the pieces together afterwards. It was probably overkill, but he did a wonderful job. As I say, he was one of my heroes. I even wrote the Wikipedia page on him, just for your interest.

Gardner: It sounds like structured circuit design, analogous to structured programming today.

Wood: Yes.

[3] https://en.wikipedia.org/wiki/Charles_Coleman_(engineer)

Gardner: Maybe while they find the interfaces, make sure they all work. Any other folks worked on parts of it that you'd like to talk before we move on to the ECC?

Wood: Okay, let me talk about the error correction, and this is interesting because when I arrived, they'd engaged Elwyn Berlekamp[4] at Berkeley and he was one of the very famous people in the field of error correction. They'd engaged him to do an error correction system for this DCRS machine, and he created one or created a design which we tried to implement and basically it was not really successful. And the reason was that he was trying to do six-error correction, and to do that you'd need to implement the Berlekamp–Massey algorithm, which means basically a little microcomputer, and it did not operate in real time

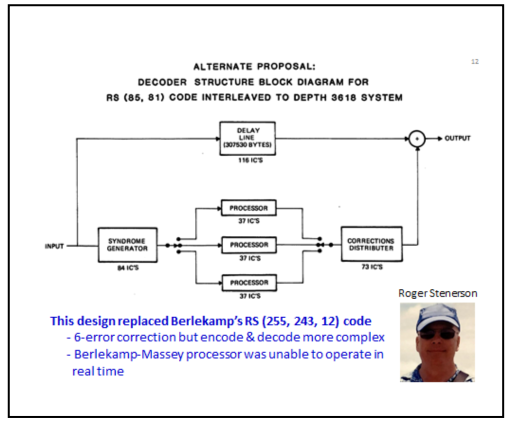

[00:39:00] so it easily got overwhelmed when there was simply too many errors to cope with. So it was not very successful, and so we replaced it with this structure here. Fortunately, 255 is divisible by 3, giving you 85 so you can make a nice length-85 double-error-correcting code instead of a length-255 6-error-correcting code and that's what we did. This thing could operate in real time. Roger Stenerson was the person who put it together and got it all working and one of the interesting features of this, which I'll show on the next page, and one of the reasons it was much more successful than the Berlekamp version and I don't know-- this probably wasn't my idea, I'm guessing it was probably Charlie Coleman's idea, but the idea was to put the bytes of the codeword physically onto the tape in a certain pattern and so the pattern of crosses you can see there, horizontal is down-track, vertical is cross-track, and this is the physical representation of those bytes as they appear very deeply interleaved on the tape and the interleave is very, very carefully chosen to combat against three different types of error events that can happen. One is that you can get a scratch, a longitudinal scratch, horizontally on this picture, down the length of the tape. The next one is that there may be an entire head pass missing. That would be a vertical stripe on this [picture], and those tended to happen. It [head clog] often cleared up immediately because of the sort of trauma of the head re-entering the tape, as it hit the tape. Sometimes, they clogged and then cleared immediately. And the other thing was there were lumps [asperities] on the tape, large lumps like a kind of tent pole. The tape would rise up on top of this tent pole and so there'd be a large, relatively circular area where there was no tape in contact and so you just had errors. This pattern was very carefully chosen to combat all those three different types of error events. And we'd feed in sort of random errors at the error rate of the tape into a simulation and predict a certain corrected error rate from that and lo and behold the actual tape recorder worked much, much better than this because the errors weren't random. They were bursty in these different types of events. So it worked extremely well. So we were absolutely delighted with that result. And it was simpler, as well.

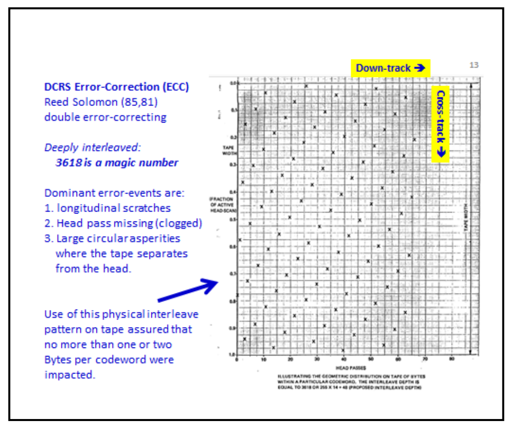

[4] https://en.wikipedia.org/wiki/Elwyn_Berlekamp

Gardner: Roger, this reminds me of what was done in the CD audio and CD-ROM technology more or less contemporaneously. Any thoughts on that?

Wood: Yes, that hadn't occurred to me, but they also had a very deep interweave. I don't know if they worried about two dimensions[5]. I don't know how long their codewords were but I should try and find that out so yes, there is an interesting comparison there. Presumably, they also tuned their system to the types of error events that they would get.

[00:42:00] **Gardner:** I recall seeing demos of drilling in a hole in a disc and then playing it back without any error. I have some papers on that I can send to you separately. They were very worried about the same sort of scratches, radial scratches or circumferential scratches on a rotary medium and blots caused by dust and debris essentially wiping out a section. Sounds, very much like the types of problems you folks were trying to resolve at Ampex. And to best of my recollection, it was almost about the same time. Now, you also worked on hard disk drives at Ampex, didn't you?

Wood: I did, indeed. Ampex was trying to get into the hard disk drive business, trying to mimic some of the IBM products. We were set up as a small "Recording Systems Group", we used to call ourselves. There was only half a dozen people but we set up a spin stand to rotate a disk and, in particular, we focused on the electronics, mainly, and we felt we were very successful. We called the project El Cholo because one of the things you tried to do, of course, is to fly as close to the disk as you possibly can, and somebody had the idea that this was like a low-rider car or automobile. Anyway, we felt we were very successful and this was published at INTERMAG in the end in 1984 and we managed to pack 100 megabytes on one surface which was quite a lot at that time. Enablers were-- Ampex was in the thin-film disk business. It was a plated disk, not a sputtered disk, which was interesting. Eventually, it didn't go into production, but it got pretty close, I think. Low-flying height, of course, you're trying always to get the head as close to the disk as you possibly can. PRML channel with precomp even at that time, long blocks, two-kilobyte block length, double error correction similar to the system on DCRS, just a 16 times interweave to keep the block length within bounds. It's interesting, though, the sort of ignorance we had in that we were very pleased that we wrote the entire disk with data and recovered all the data successfully. You can see the errors on there. All these were corrected, and this is a picture of the disk, basically, but all the errors were corrected and we were delighted with that but there was no attempt to do off-track testing which was the type of testing that IBM had pioneered. They would do these things called 747 curves[6], but we had no idea about this so we just felt it was the right thing to do was to try and write a whole disk and get all the data back. And the other comment there is that it didn't have a real servo system. It just had a single sector, but I think it was a voice coil motor but a very low bandwidth system. It was on a spin stand so no external

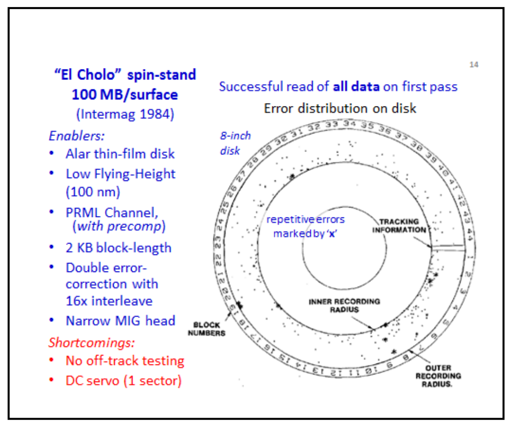

[5] The CD error-correction codes did not span across multiple tracks (i.e. not two-dimensional)
[6] https://www.computerhistory.org/collections/catalog/102790969

vibration, et cetera. But yes, we had great fun with it, a small team of people, and that might be the next slide, I think.

Gardner: Alar has a very interesting and checkered history in the hard disk drive business. Did you have any interesting experiences with Alar?

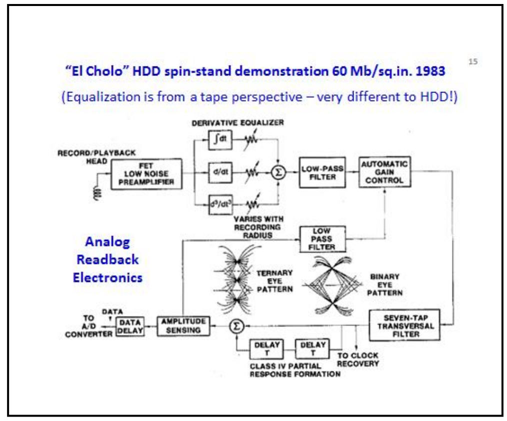

Wood: Well, it was interesting. Marv Garrison was the person behind the Alar disk and I shouldn't say this but I remember him being referred to as a snake oil salesman, which is probably totally unfair, but yes, we didn't end up with a very good reputation. I think we tried to sell disks to IBM promising that there was no chemical overcoating of any sort whereas, in reality, they were coated with, I think, sodium stearate, soap, basically. And I think IBM realized that there was some mysterious substance on the surface of these disks which was probably causing all their heads to crash and I think we were immediately quite unpopular with IBM. But they worked quite well, I think, magnetically. It was the mechanical problems with head disk interface and so on. And I think they were relatively inexpensive to make, as well, which was good. But yes, I think they sent a lot of samples out and you probably know better than I did, Tom, but I think somebody actually made drives with them.

Gardner: Oh yes, there were several companies that shipped significant product using it, and several of them had major recalls. In particular, I believe Maxtor had some serious problems with the Alar disk as did IMI[7]. I visited IMI once at their Oregon manufacturing facility,
[00:49:00]
and I saw a warehouse full of canisters of rejected Alar disks for, as they described it, white worm problems, which may have been that sodium stearate you were talking, the soap somehow collected and then caused head crashes.

Wood: Not one of Ampex's proudest moments.

Gardner: No, everybody said the recording performance, the magnetic performance was superb, much better than the oxides prevalent, but they just couldn't reliably produce a head disk interface. And in the longer run, plated might not have been a viable medium, given certain noise characteristics, but they never got there. The world went sputter.

Wood: You're right. Probably a lot less flexibility and capability in producing different structures, as well.

Gardner: I'm also interested in the single track servo, one servo per revolution, closed-loop system. Is that what you just said?

[7] International Memories, Inc., 5000 Series, announced February 1981

Wood: Yes, so there was no attempt to do a real servo system, in other words. It was basically a spin stand. It could probably have been a lead-screw but I think it was actually a voice coil motor but there was no external vibration so very little to compensate for and I'm guessing, even at that time, it was probably an air bearing spindle.

Gardner: All right, so the spin stand got you close, but the runout from changing disks was taken care of by the single servo. Once per revolution could remove runout and then interchange but still read data written at a different time having dealt with the runout problem.

Wood: Yes, and I think there was thermal drift. It would gradually drift off track due to temperature changes.

Gardner: Okay, and this was the recording channel.

Wood: Yes, so I just wanted to mention this because the tape world and the HDD world have had very different perspectives on how to do the electronics. So the tape world comes from probably audio recording where you want to try and equalize the signal, modify the signal with pre-emphasis and so on to try and get it as much like what you recorded, and that typically involves boosting high frequencies and also boosting low frequencies to try and get the whole response flat. The disk world operated off changes in magnetization, so they were trying to identify the sudden pulses that come out of the head when you get a transition in magnetic recording. So they would get pulses for every transition whereas the tape people would think about NRZ bits, just positive or negative bits in the bit stream so that's reflected in this diagram. There's what's called a derivative equalizer here, includes an integral to boost the low frequencies, a derivative to boost the high frequencies, and a third derivative to give a bit more boost at the band edge and so that was the starting point. So you end up then with a binary eye. It's sort of reproducing the plus or minus bits, plus or minus magnetization that's written on the disk and then that's taken and formed into a ternary eye pattern ready for the PRML channel after that. But it's always been one of the interesting points, talking about bits. Very often the disk folks would think of the bits as being the presence or absence of a pulse representing a transition in magnetization whereas the tape folks would always think of NRZ bits. It was either plus or minus magnetization, not the changes between the bits.

Gardner: As a disk guy, you just blew my mind, so I'm having to process, in an analog way, your statement because you're right. As a disk person, we just look-- we look for the transition. Right, and this worked.

Wood: It worked very well, yes. Yes, it did indeed.

Gardner: Okay. The next slide actually shows some of the folks who worked on it.

Wood: Yes. I don't know where we found this picture but we had a little reunion quite a few months ago now and several of these people made it to the reunion which was nice. So if I can run through them quickly.

[00:54:00] There's Linda is the secretary on the left-hand side. Vinnie Wolf did the ECC before Roger Stenerson came on board. Steve Ahlgrim did a lot of the electronics. Dave Peterson did the PRML channel. On the right-hand side at the back is Kurt Hallamasek, who did the Servo and standing next to him, the young lady there is the librarian and also Kurt's girlfriend and I forget the name of the guy in the front there. He was a summer student and I'm kind of guessing that it was Jim Belesiu but I'm not quite certain about that. But it was really nice. We hadn't seen each other for like 30 years, so it was really nice to have that reunion.

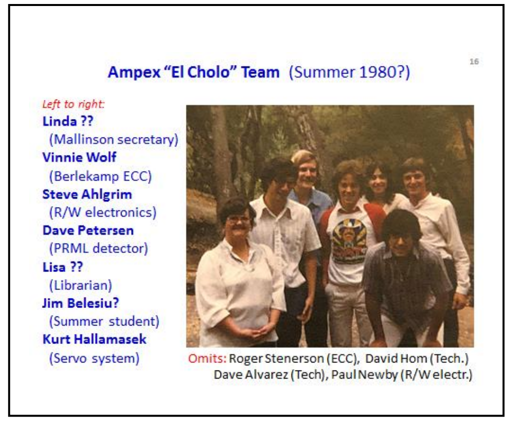

Gardner: This was the summer of 1980. We were all much younger. They look like a very young group.

Wood: Yes, I guess I'm not in that picture, but yes, I was much younger at that time, too.

Gardner: Now, this was separate from or part of the ordinary hard disk drive business of Ampex at the time?

Wood: Yes, I think this was either called the research department or the AdTech department, something like that, and we interacted. I think most of the development activities, formal development activity was in Milpitas or no, in Sunnyvale or Cupertino and we interacted with them. I remember, in particular, they wanted us to make a 1,7 encoder, and somehow, mysteriously, the Ampex library had got hold of an IBM confidential document ...

Gardner: Oh.

Wood: [document] describing a 1,7 encoder or at least a state diagram for it, so we were able to very rapidly translate that into a state machine and I think within a couple of days we'd programmed a PROM and sent it down to them and they were amazed. Of course, it's not suitable for a product. You have to actually design the circuits, the read-only memory or the logic circuitry in a real product, but I'd like to think they were impressed with the speedy turnaround.

Gardner: Ampex exited the hard drive business shortly thereafter or even right at the El Cholo project?

Wood: Yes, I think it was like '85 they exited it[8]. Before we leave Ampex, let me talk to this picture, as well. There's another of my heroes at Ampex, Bev Gooch, whom I still get Christmas cards from. He must be very elderly now, I think, but he was the head expert, and he had the bright idea that in a transverse or a rotary recorder, the head physically rubs against the tape and wears, which is not good, particularly. So he had the idea of why don't we have the tape resting on the metal film and back off the head a little way and, of course, there's so much spacing loss it doesn't work at all except that he happened to try it with permalloy and all of a sudden it started working. These are all inductive read-write heads. The theory was with a little bit of DC into the head, you formed an effective gap within the soft film that was wrapped around the drum and it was remarkable. He was able to record through, if I remember right, several mils, several thousandths of an inch of permalloy with wavelengths of similar numbers which is normally totally impossible. And I think it was Rex Niedermeyer, then, who came up with the idea of instead of putting this thin magnetic film, instead of associating it with a head, let's put it on the tape or, in fact, let's put it on the disk. So we made disks with a thin, soft film on top and, of course, putting a soft film on top of the disk means that nothing comes out of it because the soft film, that's a keeper. No flux comes out of it. There's no readback whatsoever until you put a small DC current into the readback head, and then this effective gap forms in the soft film where there's a saturated region.

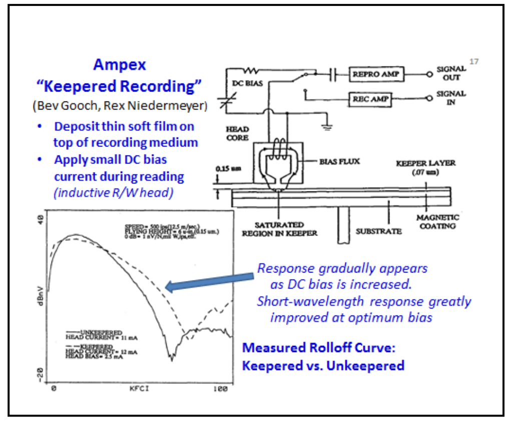

[00:59:00] And as you turn the DC card up, you see this picture in the bottom left here where the response gradually changes. As you turn the current up, this gap null, which corresponds to the effective length of the saturated region, gradually shifts lower and lower in frequency as the effective gap gets longer and longer. And the region we were interested in was this short-wavelength region where you get enhanced response. And it seemed to work quite well, actually, and we were getting good results and good advantages with it but it didn't work with MR heads and that was the death knell because MR heads were introduced at that point to disk drives. That was the death knell, unfortunately. Ampex pursued it for quite a while, foolishly, and even persuaded DSI in Singapore do to some work on it, foolishly, but MR heads killed it.

Gardner: So I think you're going to tell us about some of your Lake Arrowhead experiences.

[8] According to 1986 Disk/Trend Report Ampex exited the disk drive business during 1985

Wood: Absolutely. I was so lucky in those younger years in my career for having wonderful mentors all around me. There was John Mallinson, Charlie Coleman, Bev Gooch, Neal Bertram, of course, a host of people. And the other thing, I think it was John Mallinson probably who, he really took me under his wing in many ways. He got me invited to this Lake Arrowhead workshop organized by Harvey Mudd College and Jim Monson., I attended a lot of these over a period of about 20 years and I don't like to admit it but this hairy one in the middle here is me and I think a lot of people will recognize a lot of these folks here. I've got names underneath for those who are interested, but it was just such an eye-opener to me. I mean, there were all these people whom I considered sort of extremely famous, and I was there in this very informal conference sort of tossing arguments back and forth and just soaking it all up like a sponge, which you do at that age. So I learned so much from all these people, and a lot of them I've kept in touch with, Jim Monson in particular who organized these, is still a good friend. So yes, so it was a fun time. On the extreme right next to the bearded guy, which is Jerry Miller, on the extreme right is Mike Haynes, who we mentioned before. He's the person who got the nice amplitude and phase response to allow me to do the extracted dipulse. And this is probably where I met him for the first time and he asked me why on earth I didn't ask him for the data rather than trying to measure it from the printed page and there's a young-looking John Mallinson at the front, in front of Mike Haynes.

Gardner: Now, Lake Arrowhead had a series of workshops. Did Al Hoagland's workshops replace the Harvey Mudd ones, or are they two separate things?

Wood: No. They ran in parallel, to a large extent. I think this was called the Interactive Workshop on Magnetic Recording, and I don't know what Al Hoagland's was called. I think this one started first, and then Al started a series with, I think, a slightly difference emphasis, perhaps shifting more towards manufacturing and markets and business end of it. While this one was basically a bunch of physicists and engineers with probably not a lot of interest in whether or not the stuff actually got manufactured.

Gardner: Yes. I attended several of Hoagland's IIST Workshops [Institute for Information Storage Technology]. They were also magnetic recording workshops in some form, but you're right. There were more than just physicists, magneticians, and engineers at Al's workshops, so you typically had some senior executives and I believe Jim Porter[9] was always there with his market review. So they were same field, maybe a little less technical, but there was some pretty heavy technical presentations at Al's workshops.

Wood: In case anybody wonders who we're talking about its Al Hoagland[10]. Hopefully, you're seeing my cursor at your end.

---

[9] https://computerhistory.org/blog/jim-porter-and-the-history-of-the-global-storage-industry/

[10] https://ethw.org/Oral-History:Albert_Hoagland

[01:04:00] **Gardner:** The red sweater with the yellow shirt in the right center, last row.

Wood: And Neal Bertram standing next to him and one of my mentors and Dave Thompson, of course, Barry Middleton, a good friend, Jim Lemke[11] who passed away about a year ago.

Gardner: Eric Daniel at the one end, another good friend of ours who passed away.

Wood: And sorry, Eric Daniel, Jeff Bate, Jim Monson and that's a young-looking Gordon Hughes.

Gardner: I'm impressed with the names.

Wood: So yes, it is, but a lot of these people are still around. A lot are lost but a lot of them are still around and I count them as friends. And I've put together a whole series of pictures, actually probably about 15 pictures altogether, year after year after year and with just a few exceptions I've got names by each person and you can see each person, of course, getting older and older and sadder and sadder.

Gardner: Indeed we are. Now, your last job at Ampex was director of advanced recording?

Wood: Yes, John Mallinson left in 1984 to become director of the Center for Magnetic Recording Research at UCSD, and he left me in charge. So those were heady days because I was being asked to do invited talks about PRML at various locations and I got a promotion, basically, to director level, which is like a second-level manager and there were three groups altogether. There was the electronics group, a systems group at the bottom in the circle. These are all first names but Chuck Coleman, of course, was very central to that. On the left-hand side, heads and media, the facilities to sputter and etch and make thin-film heads and plated media and sputtered media in that small team. Each team was probably half a dozen people, maybe a few more, and the top right is Neal Bertram and some of the folks who worked with Neal Bertram and yes, those were heady days. I enjoyed that. It didn't last very long but that was a very interesting time and then my boss was Michael Felix, but he was replaced. He retired and was replaced. I'm not sure he was retired. I should say he probably was retired rather than he retired. He was probably pushed out but he was replaced by a gentleman from the NSA called Bill Mehuron and I will never forget that name. I think he was brought in specifically as a hatchet man, and his job was to reduce the expenditures. Ampex has been described as a company of engineers run by engineers for engineers. I can't quite remember how the saying goes, but Ampex had very little interest, or at least the people I knew in Ampex had very little interest in making a business or making a profit. So Bill Mehuron came in and he basically dismantled a great team and I had to lay off a bunch of people which was sort of the saddest time I came across.

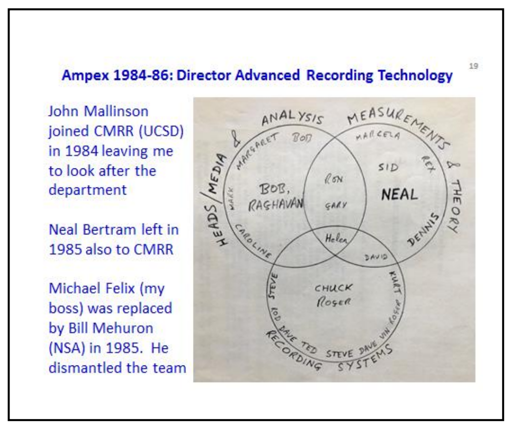

[11] https://en.wikipedia.org/wiki/James_U_Lemke

Gardner: And that's what led you to leave Ampex?

Wood: Yes, it was. I was totally disgusted at that point, and that must be the last slide.

Gardner: So this brings us, I think, to the end of this part of our oral history, Roger. Anything you'd like to add about Ampex? Ampex was in the hard drive business quite a while. It actually started in the early '70s when they acquired a Memorex spinout, and of all the people that wound up there, you and I have one shared acquaintance, Dave Jepson, who was a recording engineer. Anything else to discuss about Ampex?

Wood: Oh, you know, you mentioned it, but I wasn't really aware that Ampex was already in the drive business. I thought it was a relatively new venture; that they were trying to replicate some drive, probably a 3380 at that time. So I thought it was a relatively new venture. I know they got a bunch of new people on board to try and drive it and it lasted probably three or four years altogether, probably dismantled by Bill Mehuron at the same time in 1985. But yes, Ampex is a wonderful company in some ways but yes, not terribly successful. In some ways, it was because they were perfectionists and they knew what a good video recorder looked like, what the pictures looked like, and they turned their nose up at these silly things like Betamax and VHS. They didn't have the quality, so they weren't very interested and it drove them out of business.

Gardner: I think to this day, Ampex recorders and players are still used in professional video, and there's a business of refurbishing heads for the machines that haven't been produced for 20 or 30 years.

Wood: Yes, that doesn't surprise me.

Gardner: So with that, we will end this session and I look forward to next session in a few days. I'm going to stop recording right now.

[01:10:00] **Gardner:** Good afternoon again. This is Tom Gardner sheltering in Los Altos Hills and I'm here with Dr. Roger Wood who is sheltering in Gilroy. Hi, Roger, how are things?

Wood: Very good. Thank you, Tom. Ready to go again.

Gardner: How you doing after eight weeks— eight weeks of sheltering in place?

Wood: Oh, no problem. I'm quite happy here. As I said, we have our fresh eggs and our fresh vegetables and I haven't really been off the property for several weeks now.

Gardner: Yes. But we are a bit hairy as we haven't had a haircut in a long eight weeks or more.

Wood: <laughs> I did comb my hair especially for the video.

Gardner: As shown in the yellow highlights today's session will cover Roger's career of beginning with his move to IBM and continuing with the various companies that replaced IBM as his employer, that is Hitachi GST, HGST, and then Western Digital. And hopefully, we'll finish today with discussion of his retirement. Roger, what prompted you to move in 1986 from Ampex to IBM?

Wood: Well, as I was saying at the end of the previous segment, I was quite disgusted with Ampex in the end because they destroyed my nice little research department. So I picked up the telephone and called Denis Mee[12] and asked him if he had any jobs at IBM and very kindly he organized, I think it was, two days of interviews all together with various people. They offered me a job which I certainly appreciated. And I'd like to mention Jim Belleson because I spent many years working for Jim and he really was one of the best managers and such a nice person as well and he taught me a lot with over those years. I started at the Cottle Road site and I've got the little picture here. For those ex-IBMers, they probably will recognize Building 25, Building 28 and Building 50. Building 50 still exists; 25 and 28 have been torn down. And for some reason this cafeteria building, which I think was Building 11 or 13, still exists. It's sort of been cordoned off with a fence all around it and it's totally overgrown.

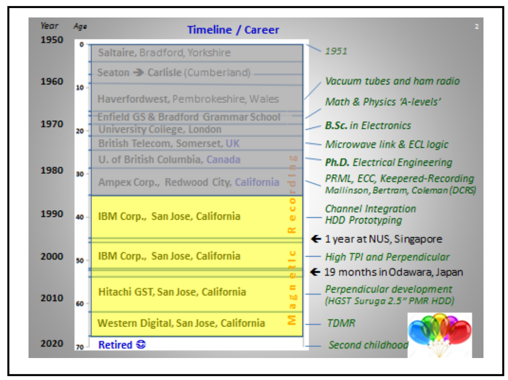

Gardner: <laughs> In interest of full disclosure, Denis Mee is a good friend of both of us and the author of a number of seminal books in magnetic recording — that is a picture of a very young Denis Mee.

Wood: <laughs> I think he's 91 or 92 now, isn't he?

Gardner: Yes. I also note that the fountain to the left in your photo, I believe it's a fountain that Jack Harker[13][14] walked on water on when the Winchester disk drive program achieved a cost reduction goal[15]. Jack had challenged the management of the program that if they made the schedule, he would walk on water; so when they did, he walked on water across that fountain [having had a wooden platform constructed just under the water's level] .

Wood: <laughs> There's some interesting stories associated with the fountain — It was a reflecting pond when Hitachi took over. They filled it full of koi because that was sort of a very Japanese thing to do. But all the local blue herons there decided they liked koi so within just a few days the thing was totally empty of koi. It was sad. As I say, that building still exists. It's surrounded with a big fence. It's totally overgrown. And the mobile sculpture thing there is just lying on the ground. I don't know what the plans are

[12] https://en.wikipedia.org/wiki/Charles_Denis_Mee
[13] https://en.wikipedia.org/wiki/Jack_Harker
[14] http://archive.computerhistory.org/resources/text/Oral_History/Harker_Jack_1/102658172.05.01.acc.pdf
[15] Based upon correspondence between T. Gardner and Kenneth Haughton, IBM's development manager responsible for its 3340 (Winchester) disk drive program

for that. But everything else has been turned into condominiums or apartments around there. And so I don't know what the plans are for this building.

Gardner: You said earlier you didn't enjoy IBM for the first few years.

Wood: <laughs>

Gardner: Would you like to elaborate on that?

Wood: <laughs> I didn't — you know, I like to think I was at least a medium-sized fish at Ampex and I think I was making useful contributions. And then you come to IBM and it's a huge organization and you're sort of starting at the bottom again. And I was making suggestions and nobody was really picking them up and in retrospect, they were probably tape-oriented suggestions rather than disk-oriented suggestions.

[01:15:00] And I guess a lot of the spin stand testing which I was very much involved with, they'd been doing things a certain way for so many years that they weren't about to change based on something I said. So it was kind of frustrating in the beginning and you're totally surrounded by experts on every conceivable topic, so it's difficult to make much impression.

Gardner: Yes. <laughs> You, can you share with us some of your suggestions coming from your tape background that were ignored?

Wood: <laughs> I don't know. Probably one of them is the bottom half of this picture. I wanted everybody to be doing this non-linear dipulse extraction and they weren't really about to do that, at least not immediately.

Gardner: Actually, offline we've chatted about the difference between a disk drive orientation in recording channel design and a tape drive orientation in recording channel design as the experts come at it from a very different perspective.

Wood: Yes.

Gardner: And I guess that's what you were talking about.

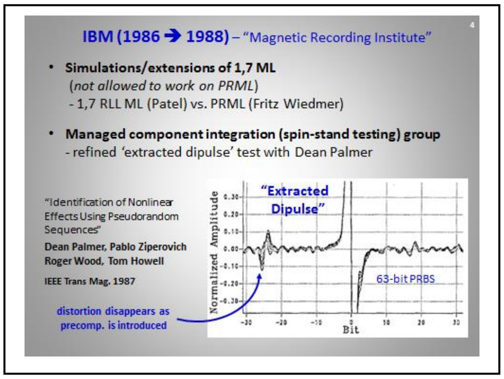

Wood: We mentioned this before during the Ampex portion. But Ampex and the tape world and particularly the audio tape world, they want to reproduce the magnetization on the tape. And so a bit is a little tiny magnet, a piece of magnetized material on the tape; whereas the disk drive world, and I guess the computer tape world, a bit is a transition, it's a change in magnetization. And certainly in the beginning, the using the same name for two totally different things causes a lot of confusion and it extends… I think we talked before about the equalization. That was totally different between the two regimes.

Gardner: So this is one of your successes, I think, the 1,7 maximum likelihood code.

Wood: Yes, I guess IBM was very cautious in those days and they wouldn't let me work on either PRML or MR heads because I'd been associated with them at Ampex. I worked a little bit with Arvind Patel He was doing a maximum likelihood version with a 1,7 code. And so I worked with him and there was a competition going on at the time between Arvind and his 1,7 ML and Fritz Weimer and his PRML team and I think Arvind got into one product and then all the following products were PRML. But it was an interesting sort of competition there. Arvind I most admired because he was so single-minded. Everything about 1,7 ML was wonderful. It had no drawbacks whatsoever. He was a wonderful salesman and advocate for his technique.

And then the bottom half of the picture we just talked about, I was eventually able to persuade people to do this and Dean Palmer in Rochester was very helpful and very interested in doing this. And I remember in particular Tom Howell was involved and he was very, very skeptical that these funny little bumps and wiggles in this picture were real and that precompensation on an NRZ signal was useful. And this picture sort of dispelled that doubt that he had. And as I've said before, this has continued to be a very useful technique.

Gardner: You were talking about that section of the pulse where the blue arrow points to?

Wood: Yes, labeled "Extracted Dipulse".

Gardner: You can eliminate that. That's not noise, that is actually something that could be eliminated by compensation.

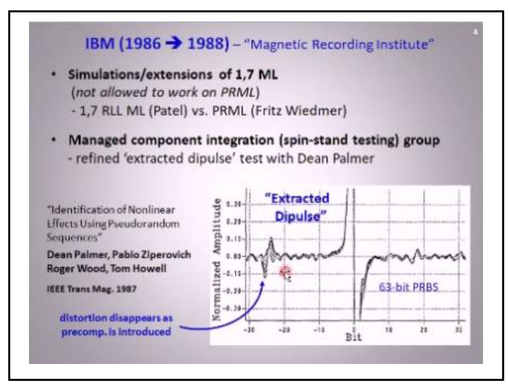

Wood: Yes, that is in nonlinear distortion, nonlinear transition-shift. So the transition or the position of the transition depends on the presence or the absence of the preceding transition. And all of these things, little bumps and wiggles here, they all have a particular meaning. They all relate to some nonlinear phenomenon on the recording channel.

Gardner: No, I was going to say, that's pretty much been adopted as a channel technique in the industry.

Wood: Yes. It's a necessity, really, and things don't work nearly as well if you allow that level of distortion.

Gardner: You mentioned channel wars to me earlier,

Wood: <laughs>

[01:20:00] **Gardner:** Tell me about the wars. Who were the victors and who were defeated?

Wood: Oh, dear. This is a difficult one to talk to. But you can see the cast of characters or at least some of the characters. Let me preface a little bit. IBM always used to run, not always, but often used to run parallel groups in competition and in this case there were groups, I think two groups in San Jose main site, one group in Almaden, one group in Zurich and one group in Rochester. And they collaborated or competed to varying degrees. But this was a horrible time. I especially dislike any kind of unpleasantness or contention in the working environment and this got quite bad. And the first mistake I made probably was hiring Lyle Frederickson. He came highly recommended by Jack Wolf[16] at UCSD. A brilliant fellow, absolutely, but if I'd talked to some of his peers at UCSD I would have got a different impression. So the main sort of competition was between Lyle Frederickson and Paul Siegel. And they both were paranoid about getting credit for their work and every week or every month— every week, I think— I had to do a monthly [weekly] report and try and capture what these guys had been doing. And inevitably, I'd get complaints from one or both of them that I hadn't represented their work properly. Hemant [Thapar] was wonderful, actually, but he deserved a promotion. And I bowed to my upper management, they were worried about their quotas, and I didn't promote him and obviously, he was upset about that because he'd been doing nice work and that was a mistake on my part.

Gardner: So these were people all reporting to you or some reporting to you and some in other groups?

Wood: No. Paul Siegal was in Research. He and Hemant collaborated. Hemant and Frederickson. Hemant Thapar and Frederickson reported to me. And then the final thing I did which was probably the worst thing a manager can do, that hairy guy on the right-hand side there decided that he'd got a better version of what they were doing, so he was promoting his own idea, which is probably about the worst thing that the manager can do to get into competition with his employees. So overall, there was a horrible mess in the end and nobody really won because in the end it came to pass that Rochester and Zurich working together had come up with a simpler scheme, a single parity check. So this Trellis code, it's called matched spectral null Trellis code. It was never used and I think they started some silicon and it was abandoned. I don't remember exactly now.

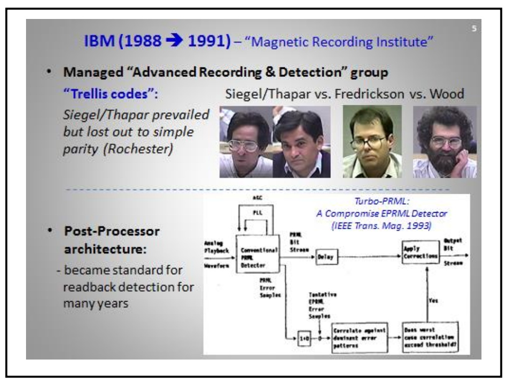

Gardner: So I'm confused. Which PRML Trellis code came out to be adopted in the first IBM product? [Editor's Note: the first PRML product didn't involve any of these distance-enhancing measures. These were introduced in later products]

[16] https://en.wikipedia.org/wiki/Jack_Wolf

Wood: It's difficult to call it a Trellis code. There was this matched spectral null code which probably was correctly called a Trellis code. The thing that finally came out was a single parity check and the detector was designed to take advantage of that parity check.

Gardner: And then that came out of Rochester?

Wood: That came out of Zurich and Rochester, yes. I think Zurich tended to come up with the theory and Rochester with the implementation.

Gardner: And who were the folks that led that?

Wood: Francois Dolivo[17] from Zurich who passed away some time ago, sadly.

Gardner: That's a great name.

Wood: Yes, he is. And I'm not sure who [was] involved — I assume it was the usual team there at Rochester. And we'll come across the names later. - Rick Galbraith was certainly involved

Gardner: Yes, I was struck by the continuation of IBM's management process of competing teams as late as the late eighties. It was well-known in the fifties and sixties and seventies that that was an IBM process, they'd have at least two teams competing to do anything and the way you succeeded in IBM was being on the winning team.

[01:25:00] At least that's a rumor I had heard looking at IBM from the outside and it sounds like at least in recording channel that process was continuing really into the late eighties.

Wood: Right. I think that's right. And maybe we were supposed to be doing different things and if we were, they were subtly different. But yes, it seemed to me they were competing teams.

Gardner: So ultimately you bailed on this job.

Wood: <laughs> Let me mention this, the bottom half of this picture there first. One useful thing did come out of that period and it was this so-called "post processor architecture". And this is the simplest form here, the little bottom diagram there. Basically, the left half of that is a conventional detector. You want to really build a much better detector but it's too complicated to build, so the idea was you build a simpler conventional detector. That gets 99 percent of the bits correct, so you've just got to worry about which bits are wrong. So you look at the output from that simple conventional detector from the perspective of the more complicated detector you want to build and because there's only a few relatively isolated mistakes, you can correct them more easily with this post-

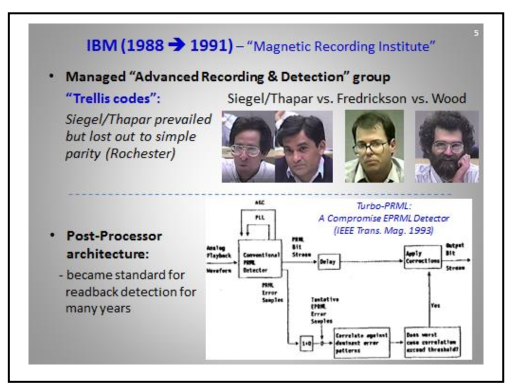

[17] https://www.zurich.ibm.com/news/05/rhein.html

processor. And that became kind of a standard architecture for quite a while and it proved useful in conjunction with this parity check code that we were using. So I was quite pleased with that.

Gardner: And then you bailed.

Wood: And then I bailed, yes. <laughs> I bailed because it was such a mess. Now, I wanted to get Hemant Thapar because I hadn't promoted him and he was upset about that, obviously. I wanted to give Hemant an opportunity to manage that group, which he did for a short time before he left the company and started his own business. But I decided I wanted to do something a little bit different, so we organized an HDD prototyping group, putting together focus on some sort of future disk drive product. We picked out a 2.5 inch product. We nicknamed it Viper. And here are some writing and sketches from that time. Tom Glaser was a manager of that. I forget what we called it. I think the name changed from MRI [Magnetic Recording Institute] to AMRL [Advanced Magnetic Recording Lab] to something else in the end. But Tom Glaser was in charge of it. And it doesn't seem like much of a deal now, but there was a $40,000 dollar incentive award offered, spread among all the people involved, which was about 20 people, so it wasn't a huge deal. We didn't make it. We failed the deadline. I think we got some sort of other compensation, but we didn't achieve the goals we were supposed to achieve.

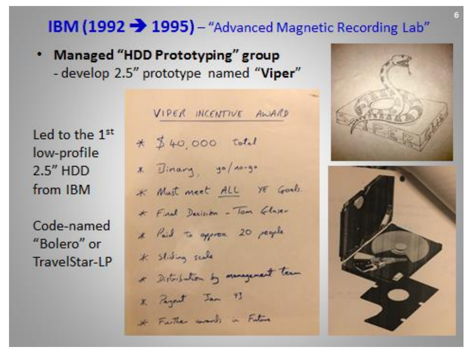

Gardner: You have a cartoon sketch that you made at that time.

Wood: I do. It's interesting to look back on these things. I've been a bit of a pack rat over the years. I've got all sorts of funny little things like this. I did a similar one of Ampex, which I think we showed before, but this is the Viper team. A little 2.5 inch drive, the whole thing held up by Dave Albrecht, whom I'll mention in a minute is the key designer. I won't go through all those names, but I've got just about the whole team there one way or another, all by first names. Top right-hand corner, Bill Hayter was very useful and very powerful, actually.

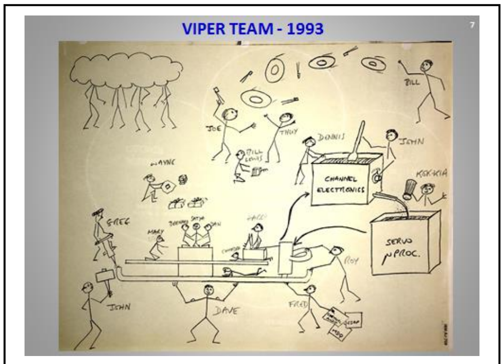

Gardner: <laughs> May I suggest you use your pointer when you talk about a name.

Wood: <laughs> Yes.

Gardner: Fill us in.

Wood: I could spend all day on this, Tom. I don't want to spend too much time. Dave Albrecht, on the bottom here, a mechanical designer. Bill Hayter was the one I was starting to talk about. He was very valuable in obtaining heads and media for the project and he made himself totally unpopular by doing that, but he was very effective. I used to get complaints frequently. But he would go, he would dress up [in clean room gear] and go down on the line and make sure that heads were being processed when they were supposed to be processed, et cetera, and the project wouldn't have been anywhere nearly as successful without him.

[01:30:00] **Gardner:** And he passed those to Joe?

Wood: Yes. This is Joe Feng, which is another name that's probably familiar to the Computer History Museum. He was one of the people who revamped the RAMAC, in that project, with John Best and Mason Williams and the others.

Let me see who else — John Hong, I'll mention in a moment. Superb read-write engineer, designer. And Greg Frees, another I think he's made quite a name for himself, external vibration and TMR testing on this thing. And servo – Kokia Chew.

Gardner: I see those guys on the left there are exciting the mechanics of the systems.

Wood: <laughs>

Gardner: I like your photo.

Wood: Right.

Gardner: Now a lot of people involved. But there are a couple of people on this cartoon that are important to you.

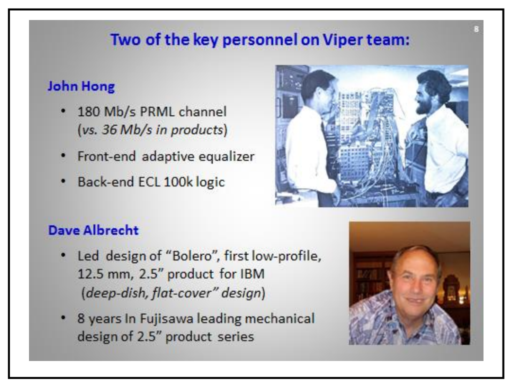

Wood: Right. These two people in particular. John Hong was a superb analog and digital designer. And because of the linear densities we were running, we needed a very high data rate channel. And it takes ages, of course, to make a data rate, a high data rate channel in silicon. And I don't think it was very obvious how to do it in an integrated circuit in those days. So John built the whole thing: analog front end, 100K ECL logic back end and it was a huge board. You can see the board there between the two of us in this photograph, far from being a single chip, obviously, but very valuable in terms of exploring those very high data rates and high linear densities on the spin-stand and in the drive. The other person down below, Dave Albrecht, taught me everything I know about mechanical design, which isn't a lot, <laughs> but he was a wonderful mentor as well. And he basically led the design of this Bolero product. And you probably want to correct me, Tom, but I thought it was the first low profile IBM product and I also thought it was the first time we'd done a deep dish flat cover as opposed to what they used to call the turkey roaster design. But Tom, do you have any comments on that?

Gardner: Well, I believe it truly was IBM's first 12.5 millimeter height. I think IBM had a 12.7 millimeter, a silly .2 of a millimeter difference, a couple years prior to that.

Wood: <laughs> Okay.

Gardner: The whole industry in that period went through a war between 12.5 and 12.7 millimeters and the idea, of course, was if the laptop or notebook manufacturer only had space for 12.5, your 12.7 was excluded. Deep dish, flat cover was well established in other companies. In IBM, it probably was IBM's first. The product shipped was the Travelstar LP, which shipped in October 1994.

I guess you went off into a lecturing mode after this Viper program?

Wood: <laughs> Yes. This was a very enjoyable experience. IEEE Magnetic Society chooses one or more people each year, they're still doing it, to be their Distinguished Lecturer as they call it and there is a certain amount of money allocated to pay expenses. And you're expected to go around the country [world, typically to the local chapters of the Magnetic Society and give talks. And I was selected to be the distinguished lecturer in 1994 and this is a list of my itineraries. So there were three or four trips, I guess. I highlighted in different colors where I hopped from place to place around the country. And this was the first time— move this a little bit— this was the first time I'd been to the Far East, I think, apart from the InterMag in Tokyo, but I was just amazed by the what I call the mind blowing hospitality that I received at Tohoku University. In Fukuoka, there was a symposium there.

[01:35:00] In Singapore in what they called the Magnetic Technology Center at that time. And then Korea as well. And each place outdid themselves with the most amazing dinners and gifts. Not huge gifts, but little souvenirs and tokens which was really nice. And I mean, that's true of all the places I went to. Everybody was — it was wonderful. But in Japan and Korea it was absolutely amazing. So that was quite an eye opener.

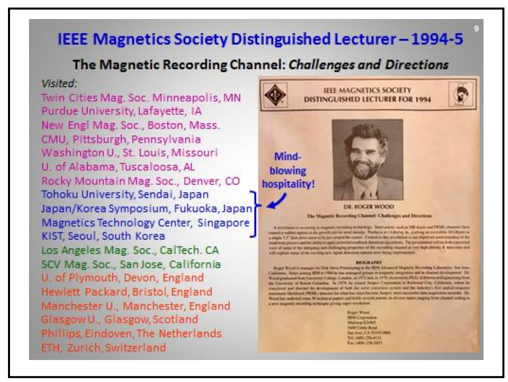

Gardner: By my count, that's 19 cities on 3 continents, 4 different trips.

Wood: Right. <laughs>

Gardner: They kept you quite busy.

Wood: <laughs> Yes, it was a lot of fun. And I think in every one of those I've probably kept in contact with the key people. Actually looking at the list, several of the people have passed away. But, yes, I had a really good time and the hosts in each location were just wonderful.

Gardner: And I gather that visit to Singapore led to something else.

Wood: It did. As I say, I was impressed by hospitality in the Far East, including in Singapore and so I decided I wanted to do something different. I was single again at that time, so and the children were pretty much grown up so I felt free to go on these little adventures. So I talked to the gentleman, the Director of Data Storage Institute in Singapore. It was a gentleman called Tek-Sang Low and he was to become a good friend as well. But I suggested to him that I might want to come out. He apparently checked with Denis Mee to find out if I was a suitable <laughs> candidate to come out and help them. And so I went out there for the calendar year of 1996 on unpaid leave from IBM. And I was employed by the University. IBM didn't want me working for the Data Storage Institute but they didn't mind me working for the University and there was a subtle difference in the way they were organized. I think DSI was a government institute. And there were a couple of projects, both of which were interesting. The other thing I did was to teach a graduate course on magnetic recording and disk drive technology and that gave me a new appreciation. Being an academic and teaching courses is a lot of work and I really hadn't appreciated it. I thought somehow it was the easy life, but especially I suppose preparing new lectures on a new topic is a lot of work. And so that was interesting.

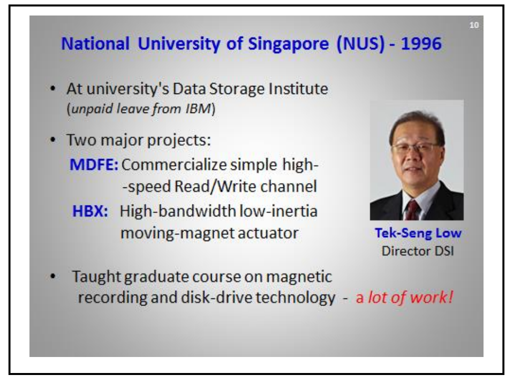

I got a couple of charts next I think elaborating on each project.

Gardner: It helps if you teach the same course several times and build upon—

Wood: Yes. <laughs>

Gardner: <laughs> Your work.

Wood: Yes.

Gardner: The first time as I understand it is rather difficult. Do you want to say a bit about the two projects that you worked on?

Wood: Let me just make a comment about the teaching as well. It was the students in Singapore are incredibly well behaved. They generally sit there fairly quietly absorbing what you have to tell them. But I remember they had a class foreman and after a few weeks the class— Because I would prepare these, you know, I'd be up to 3:00 or 4:00 in the morning preparing the charts for the next day. And then finally, the class foreman came up and said to me quite apologetically that. "The students expected the lecture notes a week in advance, please, if you wouldn't mind." <laughs> So that put me in my place.

<laughter>

Wood: But, yes, there were a couple of projects. I'll say a little bit more about those. The first one was this MDFE project. What did it stand for? Oh, multilevel decision feedback equalization. And it was a concept that came from Jack Kenney, a student at CMU, a part of his Ph.D. thesis. And it was an idea for a very simple detector, again, focused on a 1,7 code. There's always been this sort of stress between— that's not the right word— but between 1,7 and NRZ PRML-based. So it's 1,7 keeps the transitions further apart, which is nice, but it loses distance in that the differences between the waveforms are less than they are with NRZ waveforms.

[01:40:00] But, you know, there's pros and cons between the two. But I think everything now is it's simply NRZ recording, not the 1,7 code. But this was a simple detector that took advantage of the 1,7 properties. It was a decision feedback detector with certain changes to the way the thresholds were set. And there was a nice team working on it. All those names, I think pretty much I've kept in contact with all of them, which is nice. And it was organized at— DSI was very much along the lines of yes, let's do this research, but let's make a product out of it. They wanted it commercialized. So they formed a little consortium including I think it's called TMI, I think that still exists, doing the silicon. Fujitsu was involved and Motorola and I think HP was involved as well. So it was quite a big deal at the time, but it was eventually eclipsed. The idea with the MDFE was that it was simple enough to implement easily. And the progress in silicon was such that you could implement things which were much more complicated and more powerful than this. So there was a window of opportunity but it closed pretty quickly unfortunately, but it was a fun project. And I think the students there learned a lot, which is true of the next project as well. As we move on to the next project here. I shouldn't spend too much time on these. I thought by the time we'd been looking after this prototyping group and I was an expert on mechanical engineering, which I, in retrospect, I obviously was not. But this was my idea that they dutifully implemented and the idea was that the first major mode in a conventional actuator is what's sometimes called the butterfly mode or maybe the rocking mode— butterfly mode, I think's the best. It's basically the entire structure bouncing up and down [in the plane of the disks] on the pivot stiffness, on the pivot bearings from side to side and that limits the server bandwidth that you can achieve.

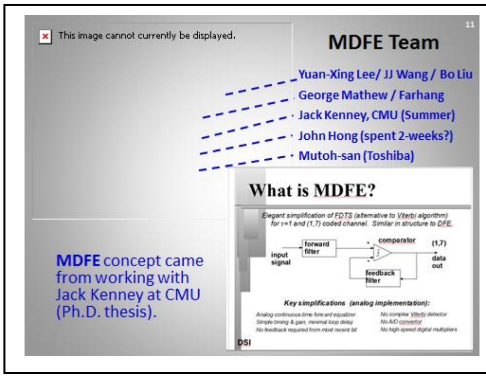

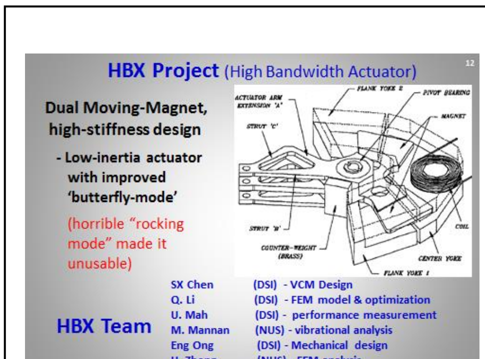

Gardner: You mean something like this? [flapping hands like a butterfly's wings]

Wood: Yes. <laughs> I'm not sure I can illustrate it very well but it's, yes, it's this mode sort of bouncing up and down on the pivot bearing. <laughs>

Gardner: Could you use your pointer to show that on the sketch?

Wood: Yes. Okay. So, yes, normally the lowest mode is this whole stretch of bouncing back and forth and into and out of the picture in this direction and the whole thing bending. The idea here was to make it very stiff in that direction and also to reduce the excitation. Instead of having a coil back here, we had a not of pair of coils which produces torque but a pair of magnets, it was a moving magnet design. There was a pair of magnets, so it was mainly torque it was producing rather than a force. And also the idea was to make a very stiff structure for the arm itself and a sort of truss structure. And it worked great in terms of pushing up that so-called butterfly mode. But in my ignorance, I guess I assumed that other modes were not important. And anybody could have told me that it was. This thing has the most horrendous rocking mode. So this magnet and that magnet are quite heavy. They go up and down in opposite directions and then the whole thing rocks back and forth on this pivot bearing. And that frequency is much lower than it is in a conventional design. And on the face of it it's not excited, but of course in reality there is enough asymmetry in the system that it's excited quite strongly. So it was a total disaster. But again, there was a certainly a good learning process for me and I hope for the team that was working on this. It was a fun time.

Gardner: I take it your team still gets together from time to time.

Wood: Yes. I have this lovely picture. Usually at the end of the year. There's a lot of ex-Singapore, ex-DSI people in the industry in the U.S. in the Bay Area so we have a little get together pretty much every year. I think we skipped this year. I'm not sure why — Oh, COVID probably. <laughs>

Gardner: Yes, COVID.

Wood: Probably, yes, because it's sometimes it's early in the year, New Year. And, but, yes, this is the team.

[01:45:00] The key people I worked with, Yuan-Xing Lee headed the channel read-write team, the MDFE project. He's now a big cheese at Broadcom. And let's see. Servo engineer, Sri Jayantha, I might be mixing the names up there, [wrong, It was] Siri Weerasooriya sorry, Siri <laughs> was a servo engineer. J.J. Wang did the phase-lock loops. George Matthew did the read-write channel. And I think those are the main folks I kept in touch with. But yes, it was a great team and as I say, it's nice to have these reunions every year.

Gardner: You were on leave from IBM. What did you come back to?

Wood: Yes, that was something again, Denis Mee, I keep mentioning his name, but he's been a very important figure in my career and my life and he made sure that I had something to come back to. He was quite adamant about that, that I should know exactly what I'm coming back to. And I didn't, really. I should have perhaps paid more attention. But this I guess worked out okay, but we put together a group focused on high track densities. So it sort of covered, it spanned everything from, you know, the position error signals, the actuator, the mechanics, the magnetics. And it was a good team of people. It was not very successful, though. We did quite a few patents but nothing was picked up by Fujisawa. And the reason I remember being quite upset, I think there was some sort of meeting where little groups were being ranked and Asano-san was the gentleman's name. He was the customer I think in Fujisawa. And there was another team in Yorktown that was helping and somehow or other he gave them an A grade but my little team or I, I only got a C grade, which quite upset me. And as I say, nothing was really picked up. Again, I had a wonderful mentor, Craig Fukushima. I looked for a picture of Craig, but I couldn't find a picture. But I learned a huge amount from Craig. There's something called a TMR budget. What does TMR stand for? Track misregistration which basically tells you how closely you're following the desired track.

Gardner: Use the laser and point your way through it?

Wood: <laughs>

Gardner: We could spend a day in this chart.

Wood: <laughs> Okay. This picture is too complicated to spend time explaining. The bottom line is the write-to-write TMR. This is the difference between the trajectory that the write head takes on the disk while it's writing one track versus writing the next track in the same overlapping. Overlapping's the wrong word. Versus when it writes the next track on top of the previous track. Obviously, you want one track to be laid exactly on top of the previous track to obliterate it completely[18]. There's no separate erase cycle on the disk drive. You have to make sure that you overwrite it properly. And these are the various inputs to that. There's a budget. I mean, each of these contributes a certain amount to that overall TMR. External vibration is a big input. There's some steady state track following issues because there's turbulence inside the drive due to the rotating disks and that can cause the disks to vibrate and the actuator to vibrate. Also if you insist on doing rapid seeks, which you do rapid access which you do in a high performance disk drive, you have to worry about how quickly the head settles on to the new track, so that has to be budgeted in.

[01:50:00] And then the tracks are supposed to follow the servo tracks laid down by the servo writer in prescribed positions and they may or may not do that very closely as well. So all these items go into that budget and you can break down each one into smaller and smaller pieces. But Craig was the maestro at doing this. He understood this very thoroughly and built up these budgets from scratch. And as I say, he was a great mentor that I really enjoyed and appreciated.

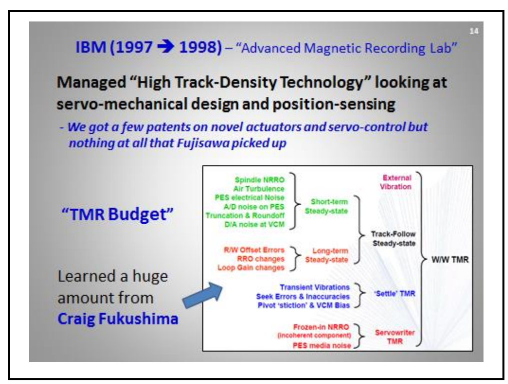

Gardner: My observation is that IBM is the only company in the industry that actually calculated a detailed track misregistration budget and published for its products a track misregistration report. In fact, did this lead to some sort of a budget and a report which was then validated in tests?

[18] A better example is the importance of adjacent tracks not encroaching on each other

Wood: Yes. There was a budget laid out. It was - the various components were measured and yes, when the drive was put together there was a lot of work to try and bring each component into line and make sure it behaved as it was supposed to and met the budget requirements. Yes, it was a big deal all together. And I think we talked earlier, Tom. IBM was probably fairly unique in doing this.

Gardner: I think they were not fairly unique, IBM was unique.

Wood: Yes.

Gardner: By definition, the only one. I could be wrong.

Wood: I think as things progressed through Hitachi and then at WD, this has sort of fallen by the wayside. I'm obviously two or three years out of date now. I think to some extent this has sort of fallen by the wayside a little bit and I don't think there's anything like this or anything to this degree anymore.

Gardner: That's my experience and I've worked at some companies and I've consulted for most and I've not encountered anything quite like the TMR report that I know existed at IBM in the seventies and your experience is it went into the nineties and into this century. IBM had a degree of rigor that no one else seems to have, yet, you know, the other folks make it work.

Wood: Yes. <laughs> Yes, IBM, I guess it was a little bit like Ampex. Things tended to be overdesigned. They were looking for perfection very often and maybe should be paying more attention to the bottom line.

Gardner: You know, my experience again suggests that IBM, now WD, gets it back in higher yields. Certainly HGST designed products appear to be the most reliable in the industry today, that is, as reported by BackBlaze for the last several years[19].

Wood: Yes. I'm very proud of those products. I think a lot of the care that IBM used to take over these products spilled over into Hitachi. Hitachi I think used to try and emulate IBM in some ways and so there wasn't a lot of difference when we went from IBM to Hitachi. And I'm very proud of those drives, the HGST drives. As you point out, they're still the most reliable in the industry.

Gardner: Apparently.

Wood: Yes, apparently. Well, according to the measurements that we see, so presumably they are.

Gardner: Yes. I'm relying on the Backblaze reports which come out periodically and they're based upon statistically valid samples — HGST's failure rate is just better than everybody else.

[19] BackBlaze Inc., San Mateo CA, www.blackblaze.com. See e.g. "Backblaze Hard Drive Stats Q1 2020," May 12, 2020, https://www.backblaze.com/blog/backblaze-hard-drive-stats-q1-2020/

Wood: Yes. And they're large samples, aren't they. They're in the thousands of drives, many thousands.

Gardner: Well, tens of thousands.

Wood: Yes.

Gardner: You know, it's enough to be suggestive that there's something fundamentally different about those products designed by HGST and then produced by facilities that were controlled by HGST management.

Wood: Yes.

Gardner: It would be interesting to see as WD further consolidates whether that excellence in engineering and manufacturing that historically came out of HGST will continue in products no longer HGST products but still the IBM family names, Ultrastar, Deskstar and Travelstar.

[01:55:00] **Wood:** Yes.

Gardner: Who knows? We shall see.

Wood: Yes. I don't know how things are organized now, but it was the heritage IBM HGST, so that's IBM and Hitachi was that heritage portion that produced the capacity drives and not the 2.5 that eventually went over to the WD side. So WD does all the desktop drives and mobile drives.

Gardner: Oh, really?

Wood: Yes, I think that's true, yes.

Gardner: So they're WD design and manufactured but sold under the Travelstar or Deskstar label?

Wood: No, I am confusing you. The sort of heritage WD which does the mobile and the desktop and heritage IBM/Hitachi which is doing the capacity drives.

Gardner: But you can buy a 3.5 inch desktop drive today from WD with I think a color name. A WD Blue PC Desktop Hard Drive for example.

Wood: Oh.

Gardner: And you could still buy a Deskstar until recently, a WD Deskstar now but it is presumably a drive designed out of Fujisawa and produced at a facility under the Deskstar brand, whether that's a facility that was an IBM facility or a WD facility we don’t know.

Wood: Yes. I don't know. I'm trying to remember how things work. Because when WD took over, HGST had to get rid of its 3.5 inch business and sold it to Toshiba. So, yes, I'm confused over the names as well. I'm not sure what Travelstar or Deskstar refer to exactly. I think the capacity drives are called the Ultrastar if I remember rightly.

Gardner: That's true. The enterprise or capacity-level drives are the higher performing drives and are the WD Ultrastars.

Wood: Yes. <laughs> We should probably move along, Tom. <laughs>

Gardner: Right. <laughs> Let's move along. You had a joint appointment in research at Almaden and in San Jose.

Wood: Yes. This topic of offices versus cubicles has come up again, hasn't it, in the context of COVID-19?

Gardner: Oh, yes.

Wood: But all through my career, I have had pretty nice offices everywhere. And everything's moved to cubicles now, I gather, at WD. But at this point, in particular, I had a nice office on the main site in Building 28, and a nice office up at Almaden Research Center. So just a bit of nostalgia there. Yes, I had a dual appointment working both in San Jose on the main site on Cottle Road, and in the Almaden Research Center, which is a beautiful location up on a hilltop. And the work at the main site was on conventional recording at pretty much a very high data rate. And I had to admit in the end that I kind of neglected that portion, because the other group was sort of doing more exciting work. I think perpendicular recording, it's sort of been bubbling away in the background for a long time, and but I think partly because of the terabit per square inch paper that we'll talk about later, there was sort of resurgence in interest in perpendicular recording, so we had the group in Almaden focused on perpendicular recording. And that was a lot of fun. We started off, we had a disk in the beginning and nothing else. It's fairly straightforward to make a disk, because just looking at the coercivity you're making with bulk measurements, if you maximize the coercivity, you pretty much know that— how can I phrase this? If the grains are too big, or if the grains are too small, or if the grains are not aligned well, or if the grains are too strongly coupled to each other, all of those things reduce the coercivity.

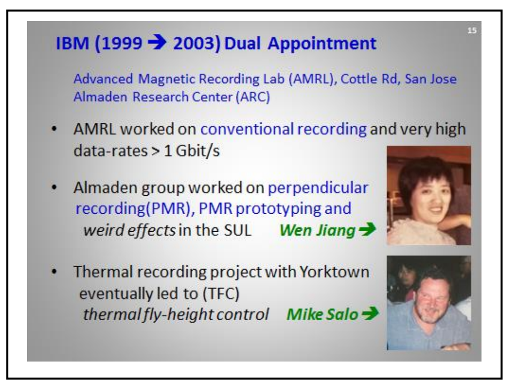

[02:00:00] So if you just aim at maximizing the coercivity on the disk, you're doing a good job. Anyway, we had disks early on, fairly early on. And I think some of the key people were [Yoshihiro] Ikeda-san Ken Takano. I should remember some of the other names, but I don't. But we had disks fairly early on, but we had no heads, no perpendicular heads, we had no perpendicular channel. So we tried writing with a conventional ring head on the perpendicular media, and you get kind of a mess when you do that. And but it became apparent, at least to me, that we'd do much better if flew the head backwards. And the nice thing about Guzik spinstands is that the disks can rotate in either direction. So sure enough, we reversed the direction of the disk, and if you adjust the speed carefully, which you can do on a spinstand, you can get sort of a reasonable flying height. I think even at one point, we had an air bearing designed to fly backwards, and it was a bit complicated. I think they had shift the pivot point if I remember rightly. But anyway, we did quite a few experiments with the thing flying backwards, and it was because of the structure of the head, the way you make a thin film head, you put— you build a slightly wider pole tip than you need on the bottom pole. So that when you put the top pole on, the upper pole on, it doesn't fall off the edge of the lower pole. And what— if you run that backwards, it's horrible, but if you run it forwards, sorry— if you run it forwards, it looks horrible. If you run it backwards, it looks okay.

Anyway, that was fun — and let me mention Wen Jiang , she worked on a bunch of this stuff. And in particular looking at some of the strange effects in the SUL [Soft Under Layer], we were very worried about the SUL. We were worried about "spike noise" and things like that, because of domains in the SUL. As it's turned out, and as we've gone into the future, this soft under layer is so far away, it doesn't really do anything on readback at all. But Wen did a bunch of work on this and published some papers on this. That was a fun time. And then there was another thing going on, and that… — Wen left the company, and Mike Salo took over the perpendicular recording. But the thing I want to mention with Mike is there was a project with Yorktown, IBM Yorktown Research, and the idea was there was a heater placed on the air bearing surface, and the idea was that because if you got close enough separation, a few nanometers, you got extraordinarily high— or you were supposed to get extraordinarily high heat transfer from the heater to the disk. So you put a heating area on the disk, a large spot. It was called Large Spot TAR. Large Spot because it was an area much bigger than the track width, and TAR was [Thermally Assisted Recording] — we refused to call it HAMR, because that was a Seagate thing, so we called it TAR, which doesn't have quite the same ring. And anyway, it didn't work. But Mike wrestled this problem of, you know, when you heat up the— when you apply power into the heater, it distorts the slider. And typically the slider expands, and that portion of the slider gets closer to the disk. So he recognized that this would be useful actually, and tried to sell it to the Head Disk Interface Group in San Jose and failed, and eventually he was able to persuade Mike Suk that it would be a good idea to try this as a means of controlling the flying height, which is the magnetic spacing very exactly. And so he really started this whole endeavor at, I guess it was HGST at that time, he started this endeavor to produce what we call TFC, Thermal Flying Control. I don't think we were the first to do it, but certainly one of the first. It wouldn't have happened so quickly if Mike hadn't taken the initiative on this. So these are two people that have kept as good friends as well.

[02:05:25] **Gardner:** So in 2003, Hitachi bought the IBM Hard Disk Drive Division. What happened next?

Wood: Yes! So it was Hitachi, and it was a Japanese company, and Japan was always very interested in perpendicular recording. And they were doing their own— they had their own project on perpendicular recording, which is the acronym is PMR, Perpendicular Magnetic Recording. So they were busy working on this. And it was looking a little bit more promising by this time because of these three innovations that were happening. And there was a little bit of interaction with Hitachi before the merger, or before the takeover. So I was familiar with some of the people involved. So I took it upon myself to suggest that I might like a nice assignment in Japan to work on perpendicular recording. And the key person in Japan at that time was Hisashi Takano, who was heading that project, and one way or another he enabled me to do that [actually, Currie Munce was probably the main enabler]. So I spent a very nice 18 months in Odawara, Japan, and made some, again, some very good friends there. And I've got a little picture here showing some of the key folks. I don't remember all the names. But Hosoe-san was doing the media, Okada-san the heads, Nishida-san, the integration and Tagawa-san led the integration group. And it was a fun time. I was going to mention that, I almost forgot that Proto-6. I'm very proud of my Proto-6! I'm going to wave it at the camera. Here is my Proto-6![20] This is still a working drive as far as I know. I used to use it for backup, but I've run out of space on it. One of the first perpendicular drives made. Well, I guess there were obviously five iterations before Proto-6. This is the sixth iteration. It was interesting when we joined forces between IBM and Hitachi, that we'd been doing very similar prototyping projects. I think we were on the same number, like Number 3 or Number 4 prototype. And so when we joined forces, we continued this. These drives were I think built in Odawara, I'm pretty sure, which was an Hitachi heritage site. And Proto-6 was the first one we distributed in 2004/'05. And I came across, which I showed you, Tom, I came across a list of people, these certain friendly individuals that we distributed these drives to. We wanted people who would view these drives favorably, and the opportunity to look at these drives and test them favorably. And I found a spreadsheet. We ranked the drives, we tested each drive, ranked them in terms of how many defects, I think, I'm not sure what the ranking was based on, but it was probably the number of defects on the drive. And so we had these drives, right, one, two, three, four, etcetera.[21] So Drive Number 1 went to a Professor Iwasaki[22] at Tohoku University. I wrote myself a list here. Drive Number 2 went to Paul Frank, who headed the INSIC organization at that time. Drive Number 3 went to a gentleman called Kevin Nguyen at Apple. Four went to Nakamura-sensei at the University. Five went to Mike Mitoma at Dell. And then there were others, including one to Jim Porter. And one went to Barry Middleton in the University of Manchester. And we got feedback on each of these. And I think they all worked, and we get feedback, interestingly, on how many hours each drive was being used for, and that was quite insightful seeing how many hours each— some of the executives were using these drives for.

Gardner: The higher the rank, the fewer the hours? <laughter>

Wood: It was something like that, yes, it was. But yes, that was interesting.
[02:10:00] So I have a nice souvenir from that time.

Gardner: Now was Proto-6 announced as a product?

Wood: No, it never made it to market, but it was— there was publicity when we did it— I think the very first drive went to [Shoyama, Etsuhiko] the head of Hitachi — there was a big presentation there. I think Takano-san presented the drive to Shoyama-san. And the same thing with Iwasaki-sensei, there was a

[20] See file named Wood 20200521 Drive Images .pdf in Wood donation for photos of Proto-6
[21] See file named "HGST 20050421 Proto-6 Allocation.xls" in Wood donation
[22] https://en.wikipedia.org/wiki/Shun-ichi_Iwasaki

big presentation made. Lots of publicity around this. But in the end, we were not first to ship the real product. I think that's probably the next chart, isn't it?

Gardner: Back up a chart, please, and if you don't mind put the pointer at the SiOx Segregant and the other three key contributions, and tell us a little bit about it.

Wood: Yes, perpendicular recording had been around for a long time, and the real person who'd pushed this was this Professor Iwasaki, Tohoku University. He wasn't necessarily the first. I think Al Hoagland had used perpendicular recording very early on.

Gardner: Yes, that was a rusty disk and a bent nail head.

Wood: Very different, yes. But Tohoku University—Iwasaki-sensei has really been— I call him the father of perpendicular recording, because he was so persistent.

Gardner: That's true.

Wood: And Tohoku University made some important innovations. They came up with the cobalt chrome medium, which it's still more or less what they use today. But there were three things that sort of really made it work. And I don't think any of them came, to be honest, I don't think any of them came from Hitachi or IBM. Maybe the Cap media did, I'm not sure. That's debatable what that [cap media] means exactly. The first one, you have to keep the grains apart. If they touch each other, then the coercivity drops, and you don't get the well-defined writing process, because you can't separate the grains easily. The segregant is what keeps them apart, and the grains are metallic, the segregant is an oxide material that prevents exchange coupling between the grains. And the trailing shield head was Mike Mallary's invention, and that— there as a big argument as to whether you should use a pole head, or a trailing shield head. The pole head is much easier to build. The trailing shield head gives you higher gradients and more advantageous field angle as well. So you do get better results with a trailing shield head. And then the last one was the cap media, which I think was Sonobe-san, at IBM. And there's various versions of this. And nobody I don't think is really quite sure how the media really works even today. I mean, today it's like five or six layers of material, but at that time the idea was use the silicon oxide segregant to keep the grains apart as much as you possibly can. But it was shown very early on Jimmy Zhu and Neal Bertram and others that you did want a certain amount of exchange coupling. But the question was how to control it and how to make it uniform. And that's where the cap came in. So you kept the grains underneath as best segregated as you could. And you put a thin cap layer on it to provide a very controlled amount of coupling between the grains. And those three things more or less came together and really made the media and the head look very promising all of a sudden.

Gardner: So when did HGST ship the first perpendicular product?

Wood: I think it was 2005, wasn't it? That's the next chart. Yes, so that was a very exciting time, really. This was actually made into a product. Suruga was the internal code name for the drive, which was a two-and-a-half inch drive. It was the— that was the first HGST drive, but unfortunately, again, we always seemed to be beaten to the post with this. Toshiba shipped first with, I believe, a not terribly successful product that suffered from corrosion issues. And then Seagate shipped in bulk, with a successful drive. But I think our drive was designed with low-cost and it did have high reliability as well.

[02:15:00] So I think it's easy for me to say this, but I think we had the best, most reliable drive, lowest cost drive probably out of the three vendors. And some of the key characters there, there's a photograph, a Sugura celebration party held in San Jose. Terashima-san was development manager, Mohammed Mirzamani did media and Ogasawara-san did the program. And unfortunately, you can't really see Yimin Hsu, but he looked after the heads in San Jose. The heads eventually were built in San Jose. And yes, that was a great time. It was the beginning of a very successful series of products, many of which were actually shipped ahead of the prescribed schedule, which was very unusual in the business in that time. And then my main job when I came back to San Jose was to coordinate the interaction between San Jose and Japan on the perpendicular product. And that was fine except that the worse part of the job was typically I ended up writing the meeting minutes, and trying to transcribe from my scribbled notes during the meeting into a coherent story every week ruined many of my weekends. Because it used to take me ages to agonize about how exactly to phrase these things. Because very often, they were action items for various people in there, and you want to make sure you're ascribing it to the right person, and the right action item exactly. But nevertheless, it was a very successful time for the company.

Gardner: Yes, my understanding is that product was the Travelstar 5K160.

Wood: <laughs> You're better at these things than I am, Tom.

Gardner: Yes, I have this need to have these simple facts correct.

Wood: Yes, good.

Gardner: You've always spent time giving talks and organizing conferences even while you were pushing back technology. What drew you to that?

Wood: Yes, I've been lucky actually in that. Generally, the companies I work for have allowed me to do this, and I should express that appreciation. But I always enjoyed that. I think there's nothing nicer than being able to talk about a subject that perhaps, modestly saying, talk about some subject that you have some expertise in, or you think you have some expertise in. It's often fun, you can often try and make the talk as interesting and enjoyable as possible, as well as conveying the technical information. So I did a lot of presentations, and short courses and things like that. I wrote a lot of technical papers, and again, I was grateful for the companies allowing me to do that. Because all the companies aren't quite so open with stuff like that. I was in the Magnetics Society— almost all these papers are published in IEEE TransMag, Transactions on Magnetics. And I was the general chairman of these two conferences. So this is the InterMag in Nagoya in Japan. And this one, two of my good friends, was TMRC in Tokyo in 2013. And TMRC stands for The Magnetic Recording Conference, and it's an annual conference with a small group of about 200 people or so. Specialized in magnetic recording as you can imagine. And these were two of the good friends. Shiroshi-san, I think was Co-Chair with me, the Japanese and American Chairs. And Muraoka-sensei I imagine was probably the Program Chairman. And I should mention Muraoka-sensei. I ran into trouble with this conference, the InterMag in Japan that one of the publication chairmen wasn't able to perform. I think he got quite ill. And at the last minute, things were really falling behind, and the last minute, Muraoka-sensei

[02:20:00] stepped in and took over the job, and did a wonderful job, as he always does. So both of these gentlemen have remained good friends. So yes, I've enjoyed doing that, it's been one of the highlights during my career.

Gardner: Now there's something called TDMR[23], that's been deployed in drives for the past few years. You've been involved with it. Tell me about it.

Wood: <laughs> Yes. There was an organization called INSIC, which stands for, International Storage Industry Consortium. It was a combination of universities and industry. Basically a way of funneling money from industry into universities in such a way that it would be employed usefully towards doing research on magnetic recording. And there was, I think, even at that time, there was HAMR[24], and probably MAMR, Microwave Assisted Recording. HAMR is Heat Assisted Magnetic Recording. MAMR is Microwave Assisted Magnetic Recording from Jimmy Zhu at CMU. And there was a little group set up to look at what other alternatives might there be to HAMR and MAMR, and we looked at various things, but we came up with something called TDMR, Two-Dimensional Magnetic Recording. And there's a little illustration here. Basically, it's shingled writing[25], which, neither of these concepts was really new, but it was sort of put together in a package. Shingled recording, which I'll explain in a minute. And then two-dimensional readback, which is basically reading that with an array of heads on an array of passes. And some pretty fancy two-dimensional signal processing to get the signals back. And we made presentations. This is one of the INSIC meetings, which was, again, it's a fun topic [meeting], because it's so focused on magnetic recording of disk drives. So very enjoyable meetings. Typically held locally in Santa Clara, or Monterey on some occasions.

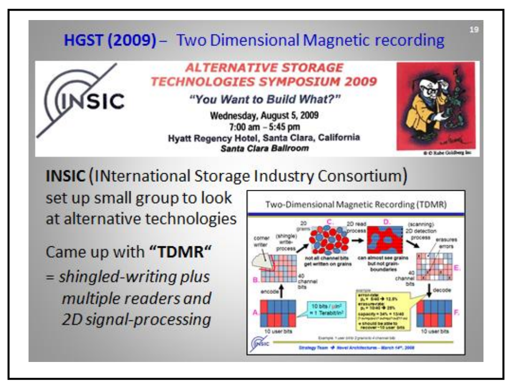

Gardner: <laughs> Why didn't it work out?

23 https://en.wikipedia.org/wiki/Two-dimensional_magnetic_recording
24 https://en.wikipedia.org/wiki/Heat-assisted_magnetic_recording
25 https://en.wikipedia.org/wiki/Shingled_magnetic_recording

Wood: Well, it did work. Yes, it's been implemented in drives, I think probably all the high-end drives. But it didn't work out. The two papers, one was called "The Feasibility of Magnetic Recording at 1 Terabit per Square Inch" [26]. And we'll talk about that, I suppose in a moment. So I guess buoyed by the success of that paper, I wrote another paper, with some collaborators in this case, on Ten Terabits per Square Inch. And using TDMR. So this is the paper, the seminal paper on TDMR. And it's about as far from the truth as you can possibly imagine. And shingle recording doesn't work as well as you'd hope. Largely because you don't particularly get improved field gradients and you end up with a lot of noise between the tracks. So it's not much better than conventional recording. It is a little bit better, but not much better than conventional writing.

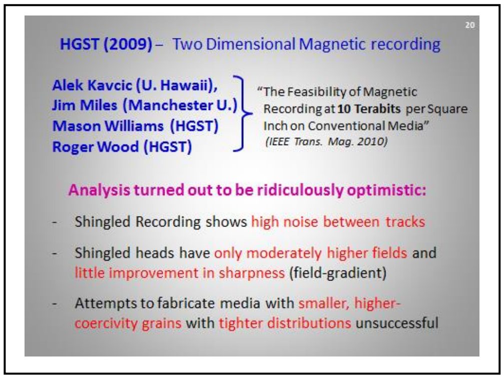

Gardner: That was sort of my impression and when I look at your three points at the bottom part of your slide, it's the one at the bottom that has tended to dominate everything.

Wood: Yes, I suppose the idea of a shingle recording or the shingle writing head was that you could get higher fields, and therefore could make smaller higher coercivity grains in the recording medium, and I think they're still trying to do that, basically. I think the grain size is sort of settled down at eight nanometer grain pitch, and it's been like that probably for the last ten years, and it certainly hasn't changed very much. So as I say, it's, again, it's a surprisingly a highly cited paper, both of those papers are well over 500 citations each. But that second paper was kind of a joke in the end. Sadly.

Gardner: Yes, your first paper, which we'll talk about in a bit, turned out to be quite seminal, and I think this one, TDMR, that allowed higher coercivity, why didn't the bits get less than eight nanometers?

[02:25:00] **Wood:** I don't think anybody fully understands that there was a very intense effort at HGST and WD as well to create media with smaller grains, using some very sophisticated techniques to do that. But I think inevitably, as you go to these smaller grains, it's difficult to control the properties, and I'm not saying that correctly, but the sigmas, the variation from grain to grain [of the properties] tends to get worse and worse. You want all the grains to be uniform in coercivity and magnetization and orientation and everything else, and coupling to any adjacent grains. And that seems to get worse as you make the grains smaller. So eight nanometers seems to be about the best size.

Gardner: Well, we'll talk more about that when we get to your One Terabit per Square Inch paper. Grain size is the current bogeyman of magnetic recording.

Wood: Yes, it is.

Gardner: <laughs> But anyhow, how'd things change when in 2013 Hitachi sold HGST to Western Digital?

---

[26] See file named "2000_00824422_1_Terabit.pdf" in Wood donation.

Wood: <laughs> Not a lot of change. At least from my viewpoint. And this was because of the Chinese government, amazingly. They [WD] had to get permission for this merger from the various governments that they sold products into, or manufactured at, because it reduced the number of people making disk drives. So the Chinese government in its wisdom said, "No, you can't combine forces," and eventually they negotiated that, "Yes, you can combine forces, but the two halves of the company have to remain separate." So for a long time there was a “hold separate” decree from the Chinese government, which we followed, and I think that lasted for several years, during which time HGST was pretty much intact. We weren't allowed to talk to the other half of the company. And I think that continued, it was relaxed a bit, that in the end I think the two halves could talk to each other, except for the marketing departments, which had to be separate. And then finally the two companies merged, obviously, completely and that's the situation today.

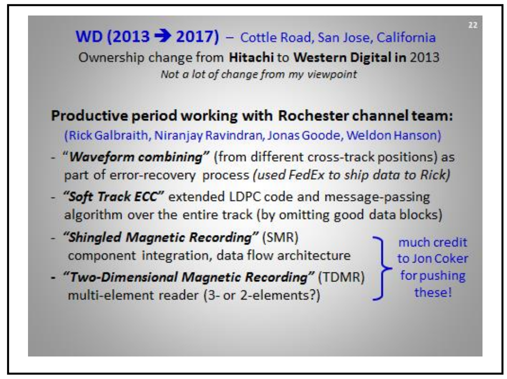

Gardner: And that's quite recent. That's a removal of the constraint is like within the last year or two.

Wood: Two, probably, yes. So actually, it was just about when I left, so it was two-and-a-half years ago. [27]

Gardner: Okay, now I got us a little bit out of order, because I think there's something that you're quite a bit of proud of, and you have to back up one slide.

Wood: I will, yes! Coming up with this! <laughs> Yes, I'm very proud of having received the IEEE Magnetic Society Achievement Award. So if this isn't an opportunity to blow my own trumpet, I don't know what is, so I'm going to blow my own trumpet at this point. So I was delighted, yes. I received the award at InterMag in Sacramento, and I gave a little talk, thanking Mallinson and Neal Bertram, in particular, for being such wonderful mentors. Unfortunately, Neal wasn't there. He'd walked out before I gave my little speech, but never mind.

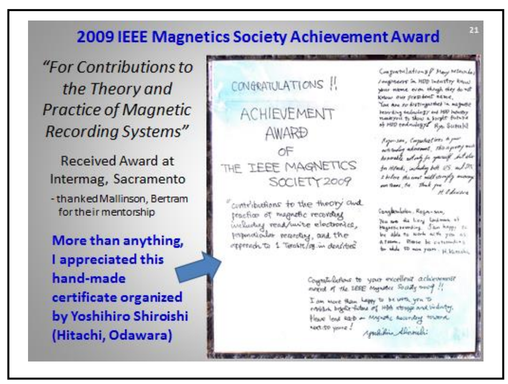

But one thing I really appreciated, Shiroshi-san[28], who we mentioned as Co-Chair on that TMRC conference, he organized this little handwritten certificate from the Odawara team. And I was so touched by that. I thought it was a really nice thing to do, and there's a little gift alongside it as well, which I’ve got. But that personal touch, I think, makes all the difference. It's all very well having a nice award and standing up at a big conference and receiving the award. But little things like this make all the difference, don't they? I count Shiroishi-san as one of my good friends [using pointer on graphic]. So thank you, Shiroishi-san if you ever see this!

---

[27] The removal of the constraint was announced on October 19, 2015, with a 24 month integration period.
[28] https://en.wikipedia.org/wiki/Yoshihiro_Shiroishi

Gardner: Now that goes all the ways back to your Ampex folks, doesn't it?

Wood: Oh, yes! I kind of think of John Mallinson and Dennis Mee, in some ways in the same vein. [02:30:00] I'm sure they wouldn't like either of them to be lumped together, because they're very, very different characters. John has passed away, sadly. But they've both made— had a huge influence on my life in various ways. Both very good mentors, and both had key roles in the sort of direction of my career as well. So yes <laughs>, very different personalities between the two of them.

Gardner: And now we're going to talk about shingled magnetic recording?

Wood: Yes. Just briefly on this picture, I had a very good last few years with working with the Rochester Channel team, very productive. So it started with waveform combining and Soft Track ECC, shingled magnetic recording, two-dimensional magnetic recording. So waveform combining was really the precursor to two-dimensional magnetic recording. And again, it was something that was greeted with a certain amount of skepticism. That's this first bullet here. The idea that you could take waveforms from the read head, placed in two different positions cross-track, and usefully combine them to get a stronger, better signal wasn't that obvious. And there was a lot of work done with that taking waveforms very carefully on a spinstand at various cross-track positions. And it's so funny, really, because we had— there were a couple of technicians. Jana Jarrel , I remember, and Helen [Dang]— Helen's last name I've forgotten— but they would sit at the spin-stand for hour after hour incrementing the read head across the width of the track and capturing waveforms, and it was gigabytes of waveforms. And to actually transmit them to Rochester where the processing was to be done. We found in the end it was quicker … so they had to finish making the measurements by 4:00 p.m. at which time I would pick up the flash stick from them, I would zoom across to the Fed Ex office close-by, give it to them, they would ship it overnight and by 8:00 a.m. the next morning, Rick Galbraith at his home would receive this Fed Ex package with all the data in it. And I guess at that time it was the quickest way of transmitting gigabytes of data. But that went on for many months, this process. And so this waveform combining was included into the error-recovery procedure, that's where it first appeared. But of course, it was an important part of the two-dimensional magnetic recording subsequently as well. And then Soft-Track ECC, so in retrospect, it's an obvious idea, but with LDPC[29], you know whether or not a data block is good or bad. So you can build a much bigger code across all the data blocks around an entire revolution, and use message-passing[30] to try and recover that as you would for the smaller blocks. And that's totally impractical, of course, unless you recognize that most of the data blocks are correct, and you don't

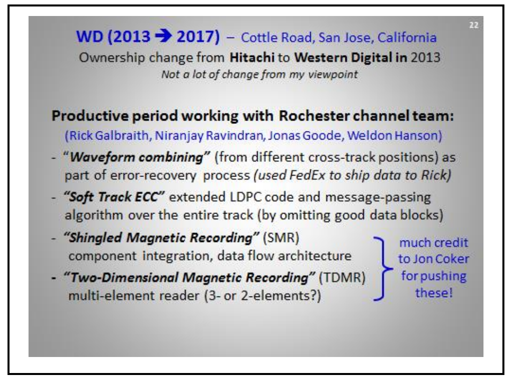

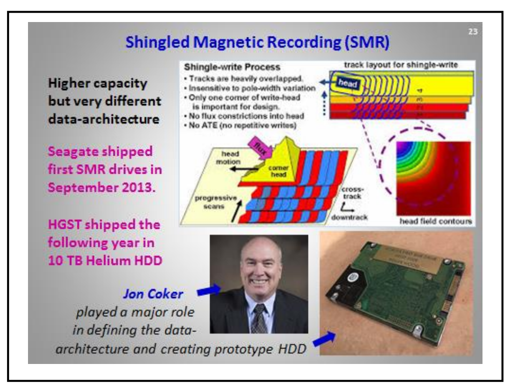

29 https://en.wikipedia.org/wiki/Low-density_parity-check_code
30 https://en.wikipedia.org/wiki/Belief_propagation

need to consider them, so there's a way you can ignore those data blocks and just process the data blocks that have problems. So it's like having a very large data block consuming the entire track. And this sort of thing, it's not very nice in that you have to do read/modify/writes and things like that. But if you're talking about very large blocks, which I think you necessarily are with shingled recording, if you're talking with very, very large chunks of data, then that's fine, and that's exactly the sort of thing you want. And then shingled recording, we've talked about a little bit more, and two-dimensional recording we've talked about. And I have to give credit to John Coker and I'll do that again on the next chart.

And it's a picture of a happy smiling John Coker, and he played a really major role in various aspects, both shingled recording and two-dimensional magnetic recording. He was the driving force behind a lot of the activities.

[02:35:00] He really promoted these ideas and flew from place to place talking to people and making sure things actually happened, and he played a major role in defining the data architecture and created a prototype, and then he was behind this, another [show-and-tell] — this is the SMR drive that he created [lower right hand corner of image], and he gave me a nice souvenir of that. And I think I've just got one more thing to wave in front of the camera in a moment. But yes, unfortunately, again, Seagate shipped ahead of those, I think, and Tom, you had some comments on that again?

Gardner: Well, yes, we have a difficult time describing what constitutes a first shipment for measurement basis. If you go back to RAMAC, for example, IBM shipped its “first” RAMAC disk drive in June of 1956, but that was not a production model, they didn't actually ship production until November of 1957, so what do you count? Shipping for prototype, your Proto-6, if it'd been a product, what do you count those that shipped out to your beta sites or do you wait for an actual production unit shipped to a paying customer? Most companies don't tell us exactly when that happened. So in this particular case, Seagate had a September 2013 announcement.

Wood: I think that September date you did say was correct. And then we'll talk about TDMR in a minute. It was the same sort of disappointment with TDMR, because I think we're not quite sure whether Seagate shipped early or not.

Gardner: That's the point, in September 2013 they announced they'd already shipped a million drives. So we really don't know when the first unit shipped. Right? That gets to if you're going to mark a measurement, yes, by September they had shipped a million drives. Now, you know, that could be several months, or even a year of shipments, we just don't know.

Wood: Yes, and SMR is out there now, but it's taken a long time to gain customer acceptance. And I don't know, we talked about this briefly, Tom, but there's been a bit of a scandal recently over SMR. Scandal isn't the right word, but apparently all three companies, WD, Seagate and Toshiba have been shipping shingled recording in their products without telling the customer. And it has certain performance attributes, some of which are not very desirable, this shingled recording. And this is— it becomes apparent in certain applications that the performance is not what it should be. But that seemed very underhand to me to not specify that it's shingled architecture, because you had to take special provision of how the data is handled on and off the drive. So that's happened just within the last week or so.

Gardner: Yes, the shingled magnetic recording effects access time, and so it depends upon how they specify the access time. The problem is shingled magnetic recording drive really works more like a tape drive than a disk drive.

Wood: Well, it's well-hidden. <laughs>

Gardner: And disk drives, in theory, you can update a single 4K block, 4,096 bytes, whereas a tape drive, you tend to write megabytes or gigabytes at a time, and in an SMR drive, if you're writing very large blocks, you won't notice it, but if you're writing a small block, it's going to be a horrendous impact on access time.

Wood: There's two ways of doing shingled recording. One is to have the host look after it, the computer. And the other is to have the drive look after it.

Gardner: Of course.

Wood: And that's, apparently, what's happening now that they've hidden the SMR inside the drive, so it looks like a normal drive. And in most applications, it works just fine. All the data is staged as you write it onto the drive.

[02:40:00] So it's written immediately onto the drive, and it's just later on that it has to be shuffled around. And it's the situation where the drive is almost full, and being used heavily, that's when it sorts of falls apart. I think from what we've learned over the last week or so, is that all the desktop drives and the mobile drives are shingled recording now, so it really has become the way that these drives are made now.

Gardner: Well, again, it gets to access time, and specification of access time as to whether there's any misrepresentation going on, but you know, with the drive on your laptop and on my desktop as we speak, access time doesn't matter, each drive has lots of time right now to do lots of things. But if you're going to put it into an access time intensive application where you're writing many small blocks, reasonably small blocks, shingled magnetic recording, well, if it's drive-managed, your access time is going to be very erratic. If it's host-managed, presumably, the host will block things up and only write very large blocks, in which case the access time is irrelevant, because the time to transfer a large block is just very, very long. And that's the gist of the problem.

Wood: Sorry to interrupt, Tom. Access time is usually fine, except when the drive gets full, and when it's being heavily used. That's when it sort of starts to fall apart. But otherwise, I mean, you can write the data immediately on the disk. It's always staged. And you can get the data off the disk just the same as you can with a conventional drive. But when you have to start shuffling the data around and there's no time to do it, and you have to do a lot of shuffling, because the drives may be full, then it becomes a problem.

Gardner: I think we're saying the same thing in two different ways, where access time is important shingled magnetic recording will fall apart.

Wood: Yes.

Gardner: You particularly can have problems with writes. What's happened is that disk drives used to be simply addressed. One wrote a block, which had a very specific address, and unless it was a defect, you could pretty much find the data at that location. But when you get into an SMR drive-managed drive, your data are placed wherever there happens to be an availability, particularly as you do update writes. Data just gets moved all around, that means you have a file system on the disk drive, not just a simple access method, and as your accesses go up, the performance drops dramatically. As you and I today sitting here, it doesn't matter. But if I were doing a graphics design on my computer and making lots of changes, the performance might fall down completely, and if this drive was plugging into a server, it would really go down, so it is how they represent it. I haven't read the complaints yet, but there may be an argument for misrepresentation.

The more interesting thing about shingled magnetic recording is again it gave us somewhat of a boost in capacity, but that's it. It was a one-time boost, as opposed to other technology changes which have allowed an ongoing increase. We got what about 20 percent?

Wood: Yes, shingled recording tends to be more valuable at higher densities. So it buys you very little at low densities, but as we've been pushing densities, it's become more valuable. So I mean, they claim numbers like 25 percent, but I think in reality, it's more like ten or fifteen percent.

Gardner: Yes, that's what I mean. It's been a ten to twenty-five percent boost in density. And you know, that's it! You know, it's now a more or less standard process, and it doesn't— it's not going to give us any more capacity to real density increase.

Wood: Right. We should move along.

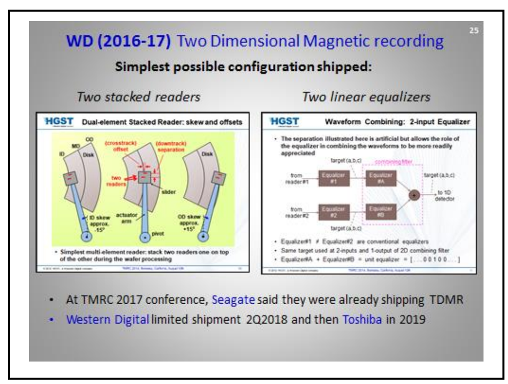

Gardner: Oh, yes, we're going to talk about your collaborating with the Rochester Group.

Wood: Yes.

Gardner: But this image is in San Jose, isn't it?

Wood: This is in San Jose at one of the internal conferences held at the SanDisk Auditorium. But these are actually the three key folks that I collaborated with. All three of them brilliant.

[02:45:00] This gentleman here, Rick Galbraith, is now a Western Digital Fellow. And he's got like 100 at least patents or more. And has been one of the key architects of the channels. I guess, IBM, Hitachi, WD, don't actually do channels as much as such, as they rely on vendors. But pretty much— maybe I'm exaggerating just a little bit, but a lot of the key features in the vendor channels come from this gentleman

[Rick], basically. Niranjay Ravindan here, I think is one of the students from University of Minnesota but, again, he's done some amazing work. And then the gentleman here you can't see terribly well at the back, Jonas Goode, again, outstanding person both theoretically but also very hands on. He actually designs algorithms and brings up the drives and tests them out, and the other attribute is that he is a champion wrestler. There's some particular category of wrestling that he's famous for. He wins competitions. He travels all around the world and very often he's the winner of the competition so he's an interesting personality, as well. So I was very lucky to be able to work with those folks. It says Rochester, too, but actually Jonas is based in Los Angeles, not just Rochester.

Gardner: This looks like the Almaden Cafeteria.

Wood: No, it's not. It's sitting outside the SanDisk auditorium and it's funny I notice it now, but that black blob at the front in the foreground is my trusty briefcase which I still use. I don't have it in this room, but I've been carrying that thing around for the last 20 years probably and had it repaired numerous times because things keep breaking on it. But yes, that's a funny memory. Sorry. That's a distraction.

Gardner: You said TDMR did eventually ship, but was Western Digital first or not first?

Wood: Well, I don't know. Let me first say, it was nothing like we imagined, what finally shipped. It's very difficult to make an array of heads immediately side by side, so we stacked them on top of each other, as you can see here, down-track. So you get some lateral separation across-track just by virtue of it changing skew, and the channel, there's no two-dimensional detection. It's just a pair of equalizers combining the signals from the two readers and feeding it into a conventional detector, and that sort of understates the complexity and the architecture a bit, but very different to what we originally imagined on this, I think one of the reasons why the gain is not anything like what we hoped for. But did we ship first? No. I don't know. You can correct me in a minute but I remember being so disappointed at TMRC 2017. Fatih Erden was the Seagate gentleman we'd been collaborating with. Not collaborating, but we were interacting through the INSIC or whatever the successor was, but Fatih Erden stood up at the beginning of the TMRC 2017 and said, "You can take TDMR off your list of technologies, when are they going to ship question on the questionnaire because Seagate, we have been shipping it for months now. We have shipped thousands or millions of them." So I was very disappointed to hear that. We didn't ship until a little bit later there, but Tom, you had some other comments about this.

Gardner: Again, getting back to how do you measure shipment, indeed they did announce something at the conference in 2017. But Seagate did not announce general availability of the product until third quarter of the next year, more than a year later, when four products were announced as readily available,
[02:50:05] that is, the IronWolf, IronWolf Pro, BarraCuda Pro, SkyHawk, and the Exos TM, products, each product targeted at a different market segment. So that's third quarter 2018 before you could buy one either as an OEM customer or a retail customer. Western Digital, on the other hand, announced limited production

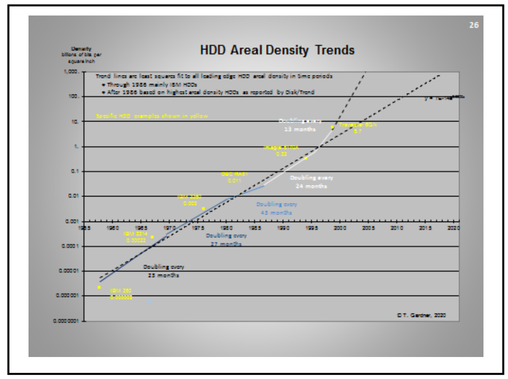

about six months earlier. In second quarter 2018, Western Digital was in shipment. So what was the first customer shipment as we would ordinarily measure it from a production line to a paying customer? Who knows? But if you weren't first, you were very close.

Wood: Okay, that makes me a lot better. Thank you very much. That was good.

Gardner: So in 2017, you decided to retire, but before we talk about your postretirement work, I'd like to jump back to 1998 when you published a paper which predicted an ultimate limit of about one terabit per square inch.

As a way of setting context for discussing the paper, I prepared this chart which shows areal density trend lines and two things. First of all, let me tell you about this chart. It really looks at the entire industry and for the first four decades, it's basically IBM leading areal density. So the trend line is connecting the dots between the IBM product shipments starting with the 350 and going to a variety of products. The first three decades from 1957 the IBM 350 in 1957 going up into the early 1990s, the industry was more or less following Moore's law, doubling every 24 months or so. But beginning about the early '80s, the density rate of increase had fallen off to almost doubling every two years and then there's a kink in the curve about 1980 at which point it is looking beyond IBM HDDs to the whole industry. And beginning in the late '80s and up until the mid-'90s, the rate increased at a higher rate until by the time you wrote the paper, 1998, the areal density was doubling every 13 months, almost once a year.

A few years later, Gordon Moore, the author of Moore's law[31], called this extraordinary. So here you are in 1998 and if you look at the lower dotted line, historically, back to 1957, sometimes faster, sometimes slower, the industry was doubling its areal density every two years. And if you follow that lower dotted line, you'd see you'd hit a terabit per inch squared sometime before 2018. On the other hand, if you looked at the more recent areal density increases occurring in the late 1990s, you follow the dotted line up, you might hit a terabit per square inch around 2004. So Roger, at this time in 1998, looking at the history, looking at the technology, you write this paper that says one terabit per square inch, that's the end of this trend. How was the paper received?

Wood: Looking at this chart, Tom, the top end here is one terabit per square inch but we didn't know that. We didn't know how far it would go.

Gardner: But that's what you said it was.

[31] https://en.wikipedia.org/wiki/Moore%27s_law

Wood: And so the task was to figure out what was the limit, and it was, again, a lot of this was associated with the INSIC organization. Some of the key names there, Mike Mallory, Mason Williams, Randy Victora, Dave Thompson, there was actually a host of people involved in these discussions, and I got very interested in this whole topic of what's the physics, what's the physical limits of recording. And so I set about trying to figure that out. There's noise from the head, there's magnetic noise in the head, there's thermal noise in the head, there's noise in the media, and what are the limitations?

[02:55:20] How far can you really push things?

And I think the first thing that came out of this was that you probably could not get very far with longitudinal recording. So it had to be perpendicular recording, and I put a story together with perpendicular recording with limits to the grain size due to the thermal stability of the recording medium and also data rate, primarily because of gyromagnetic effects. And remarkably, doing this in 1998, both of these have come true. The areal density has stuck at one terabit per square inch and the data rate has stuck at three gigabits per second. So more by luck than judgment but it's worked out very well.

The following year, I tried to put together a strawman product with these at one terabit per square inch, and it's interesting to see. I won't go through this item by item but some of the things are totally off base and some of them are really spot-on.

So the things that are totally off-base, and there's some excuse for this, the disk drive was a terabyte drive. It was, I think, a one-inch disk. Can't remember now but I think it was a one-inch diameter [1.3-inch] spinning at 23,000 RPM because at that point in time, we had no idea about flash memory. We wanted to build a super-high-performance drive, so we wanted extremely low access times. So we wanted the very high RPM, the very fast actuator which arises from miniaturization, as well, but that of course is totally off base because flash memory has come in and taken all that part of the market. We don't even make 15K RPM drives anymore. So that was way off, but a lot of the other things were spot-on. The grain diameter is pretty much spot-on and one thing that is also very close, when it came to spelling out the magnetic spacing, which is really a very critical parameter in figuring out what linear density you can achieve, I talked to Frank Talke[32] at CMRR UCSD. And he thought a bit and came back with a number and he budgeted it quite carefully, what's the limit of carbon thicknesses, what's the limit of physical spacing, air separation, and the answer was 6 1/2 nanometers and that is almost spot-on, as well. It may have changed a little bit, but that's pretty much where it was when I retired in 2017, 6 1/2 nanometers of magnetic spacing and the physical clearance is about a nanometer or less, 6 1/2 with all the carbon overcoats and everything. So that was spot-on, thank you very much to Frank Talke!

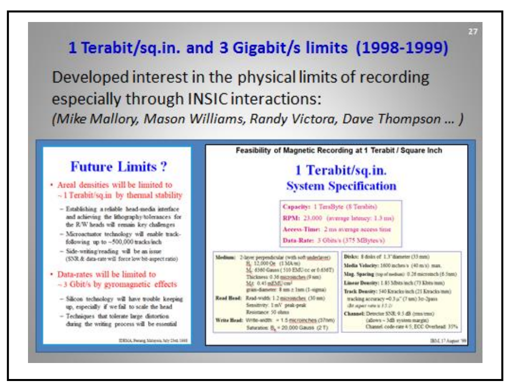

Gardner: So which of those go back to your limits?

Wood: So this is mag-spacing here.

Gardner: Not to your drive but to your future limits? They're still there, right? I mean, no one has come up with technologies that will allow us to continue increasing areal density at a high rate. They're all there.

32 https://cmrr.ucsd.edu/research/faculty-profiles/talke.html

Wood: Yes, this was stated as being conventional recording. I hedged my words. I said that the limit will be “of the order of” one terabit per square inch, which gives you a lot of latitude but it's come out pretty much exactly and shingled recording and TDMR add just a little bit to that and there's debate over whether you call it conventional recording still.

[03:00:00] The data rate, I think I underestimated that. I think we could push that higher but fortuitously, perhaps, the linear density is limited and the velocity hasn't really been pushed because we haven't tried to do 10K RPM or 15K RPM 3 1/2-inch drives. The media velocity has stayed quite moderate at 30 to 40 meters per second, so there has not really been a great need to push the data rate. But yes, they're both about right.

Gardner: When you say the future limits, when you say thermal stability, there you're talking about grain size.

Wood: Yes, and the head, as well, we were looking at. You have to make sure that the magnetic noise in the head, which is also thermally excited, doesn't become excessive, and that's very close to being a problem. And so that will limit future density — if we manage to improve the media by doing HAMR or something like that, the read head is going to be a problem.

Gardner: Well, it used to be, before we got down to these low limits, that system design was about 60 percent media noise, 10 or 15 percent electronic noise, a few percentage for the head. That was your noise budget. It was a media-dominated noise budget.

Wood: It's a little bit different to look at after equalization. The head noise, the white noise becomes not predominant but much more significant after equalization.

Gardner: And I notice on your future limit chart there, micro-actuator technology, the other side of that is how narrow can you make the head. I mean it's sort of an issue there, right? Is it going to be the photolithography is well beyond the radial dimensions, much smaller than the positioning? Do you think it's going to be the mechanical limits that are going to hold us even if we can find a better media and a better head?

Wood: No, I don't see any mechanical limitations, not from a physics point of view, and we talked about this a little bit. I guess I made the claim, probably erroneously, that it's not surprising that HDD areal densities go up 40 percent per year because that's Moore's limit [Law] and that's the semiconductor technology rate. So the lithography goes at that rate so you can make heads and silicon sort of increasing at that rate and it's only when you run into roadblocks and you pointed out the transition from particulate media to thin-film media. But I think one of the big transitions was from inductive readback to magneto-resistive readback and then GMR and then TMR. I think that's been a big factor in allowing us to continue.

Wood: And then, there's a host of things, isn't there, Tom? If you think about the writing process, as well, I mean you have to be able to write high coercivities. You need high-moment pole tips, et cetera. So there's a whole host of things, but yes, it's difficult to pin down exactly what the limits are to each stage.

Gardner: And in our conversations, I have suggested that the electronics has also been, at times, barriers that have been removed which were precluding rapid advances in technology. I think we agree the lithography really hasn't been an issue almost ever, right? The one nice thing as a colleague of ours pointed out, about lithography in the head, is you can make a bunch of heads and throw the ones out that don't work if your lithography is noisy. On a DRAM chip, every bit has to be accurate but if you're making a wafer of heads and your lithography is a little uncertain, fine, you make some bad, some good. It'll increase your cost, but you can finally select the head that will work in a disk drive at the areal density you want.

[03:05:00] But if your lithography is good enough, you'll get a very high yield, and so I would argue that lithography has not been ever a limiting factor. The technology's always been way beyond what we've needed.

Wood: I don't think the head team would agree with that, at all, Tom.

Gardner: I would say they would not. As the grain size is today, we've hit a limit. Right? I mean the argument for thin-film heads was lithography would be cheaper than machined ferrite with metal in gap and it may be true that they were cheaper but you could still machine a head for a long time before the head industry went to lithography.

Wood: Yes, I think it's probably true to say now with the heads that there are lithographic processes available that could create smaller structures. It's just that we can't take advantage of them because we can't do more one than one terabit per square inch.

Gardner: That's exactly my point. We can not only not make the bit cells smaller, we can't make it smaller in any dimension given the mechanical problems of the servo. If we could make a better Servo and a smaller disk, lithography would not be the limiting factor. It might, in fact, be the magnetization of the head, the ability to generate the field or the signal out of the reader. It wouldn't be lithography.

Wood: Yes. The only dimension that has been pushed successfully recently has been the disks per inch stacked vertically on the spindle.

Gardner: That's true.

Wood: That's gone up a lot.

Gardner: And actually the NAND flash guys are in the same boat because the way they're increasing their density is stacking the NAND, the gates higher. They've gone from essentially 4 layers to 128 layers. They're going vertical, too. The question is how high can they go.

Wood: Makes me very jealous, just that. It's not something you can do with magnetic recording, and we'll talk about that in a moment, perhaps.

Gardner: Well, I mean, if things keep going the way we're going, you might resurrect the 5 1/4-inch disks. We could go radially, also.

Wood: There have been attempts to do that, believe me. As I understand it, both Seagate and WD or HGST have had 5 1/4-inch projects, but neither succeeded.

Gardner: Well, Quantum did. It actually came out with a product 20 years ago, ‘Bigfoot’, and it never went anyplace mainly because, again, against the 40-percent-per-year improvement, Bigfoot didn't offer very much and it didn't fit into the racks and stuff like that.

That's actually probably a good introduction to the next slide, with areal density having stopped going at Moore's law rate, so here we have all the highest-production areal density basically starting with Jim Porter's data and then just taking the announcements as various companies announce them and picking the leading, if not always the highest, as companies have announced various products and you can see it almost looks like an exponential decay. Right? I mean, we were doubling every 18 months towards the end of the last decade. This decade, it slowed down to 28 months, and now we're talking 8 to 10 years between the doubling of areal density. That dotted line is the one terabit per square inch that you predicted would be the limit in 1998, and your prediction looks spot-on.

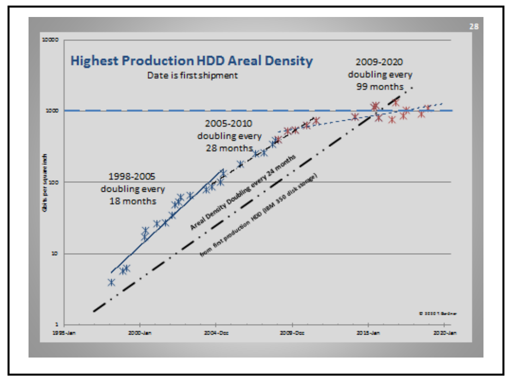

Wood: Those that are above the line, I can argue, probably are shingled recording and TDMR so not conventional recording.

Gardner: Still, I mean, and we went from MFM to 2,7 to PRML. We don't care about the code or the recording process. As I look at it, shingled recording is just another way of getting more areal density, and there have been extraordinary changes in technology from the RAMAC to the disk drives along that line approaching one-terabit-per-square-inch, shingled recording being one of the many technologies. But if you're going 40 percent per year, 20 percent improvement from shingled recording is 9 months.

[03:10:00] **Wood:** Used to be, yes.

Gardner: Yes. Well, that gave us nine months. Now how are we going to keep going? We're not. I don't think there's been a new high-density announcement in several months now. It used to be you got a new product every three or four months. Somebody pushed the limit.

Wood: Just to put this in context, I suppose probably most people are aware that there are two contenders for increasing areal density beyond this limit. One is so-called HAMR, and that's a heat-assisted one, and MAMR is the microwave-assisted one. In both cases, the idea is to make available extra energy, if you like. They're sometimes called energy-assisted magnetic recording [EAMR], but it's to allow more energy to be available during the writing process so that you can write higher-coercivity, higher-energy particles which can be made smaller. And so it, first of all, depends on being able to make a HAMR head, which isn't maybe that easy, or a HAMR system or a MAMR head, but the other thing is of course you have to make a small-grain media successfully at those higher coercivities. And that, well, it depends which system, but typically involves a huge change in the material system on the disk and controlling the grain size and getting the grain size to where it needs to be. So both those things have to happen and I think the story at the moment is that systems are being put together, HAMR systems have been put together and media has been made. HAMR media has been made, but I think they're still struggling with reliability issues in the HAMR head, lifetime issues in the HAMR head and the ability to control, to make good, small-grain HAMR media is not straightforward. MAMR may be a little bit easier, but it doesn't maybe offer such large gains as HAMR would.

So I think initially HAMR was thought to have an extra order of magnitude which would take you up to 10 terabits per square inch, but I think those estimates, that speculation has drawn back to maybe a factor of two.

Gardner: Well, the square root of 10 would say that's how much smaller the grain size has to get because I think your analysis starts with, from a signal-noise ratio, eight grains per bit if that dimensional is right. That's sort of where you are. Maybe you get to seven, but ultimately you get to one, but we don't. In theory, you could get to a bit pattern media, but we don't know how to do that, either. So if you're going to continue to have a bit cell having some number of grains, you're probably going to hit a limit of six or seven. Would you agree with that? Eight, somewhere in that number.

Wood: Yes, it's interesting, and I think we've talked about this before, but the read head, of course, is only about half the width of the track. So if you look at what's happening underneath the read head and if you look at the channel bits rather than the customer bits, then each channel bit, as seen by the read head, has about four grains in it. And that was a number that was used in that 10-terabit-per-square-inch paper, as well.

Gardner: And in your one terabit, it was more like eight?

Wood: Yes, yes, it would have been because it was an 8-nanometer grain pitch, so it would have been 8 or 10 grains per bit.

Gardner: So the point is we're near some fundamental level, and if you're going to increase by an order of magnitude— but an order of magnitude is, what, if you're doubling every two years that's six years.

Wood: I'm very skeptical of whether either HAMR or MAMR will do anything useful, but we will see. It'll be interesting to follow this over...

Gardner: Well, then the only way we can lower the cost per bit is by either going vertically or horizontally, and it's tough to lower the cost per bit. As the industry shrinks, the fixed cost has got to be absorbed over a smaller number of units, so that makes it even more difficult.

[03:15:05] **Wood:** Yes, that counts heavily against HAMR, as well, because it's more expensive.

Gardner: Where do you think the hard disk drive industry is going, just slow growth in areal density and consequential marketing issues?

Wood: No, I think it's pretty much stopped. I don't think it is going to increase. I'm very skeptical that either or HAMR or MAMR will come to a reality.

Gardner: So basically then, the cost per megabyte is going to decline sort of along a learning curve number.

Wood: Yes, not very fast.

Gardner: Not very fast, 4, 5, 10 percent a year maybe as we get better at making drives.

Wood: Yes, but volumes are probably going to shrink, as well, which complicates things. I don't know. I've always been a pessimist, though. That way, I'm pleasantly surprised when things work out better than I expected.

Gardner: How about tape? What's your forecast for there? Areal density is not near ours, but of course, their media is much worse than ours.

Wood: Yes, it seems to me that they're still a long way from being at any physical limit and they're still using particulate media so, yes, there's probably a lot more growth available on tape.

Gardner: Well, I'm not sure. Their limit may be in the tape-head interface. It's just not a very stable medium and so flying height has a huge impact on magnetic recording and the thinner you make the recording layer to make the higher-performance media, the more problems you have with wear. So I think they may be fighting flying — even though, theoretically, the areal density is not near what could be achieved from the particle limitations, the head-tape interface problems, wear and flying height may keep them at a much lower areal density, meaning their cost will also not keep lowering at the same rate it's been lowering in the past.

Wood: Yes, possibly, maybe even probably, but they're engineering limits. They're not physical limits.

Gardner: True. So leaving the industry, you've now retired. I understand you had a nice retirement party.

Wood: The last of my props, let me put this up. There's been a tradition in IBM, I guess, and it was carried over. This probably won't show up at all well, but there's been a tradition. Can this be seen?

Gardner: Oh, yes.

Wood: There we go.

Gardner: Round and brown.

Wood: There's been a tradition. One Terry Whittier, a technician at IBM has got a stack of 14-inch brown disks, and the tradition is that— you probably couldn't see that when I held it up, but there's names on these things. People sign them with a scribing instrument, so I have probably 100 or more names on this thing of the folks I've worked with. So that's a really nice souvenir to have, and again, it emphasizes how far we've come, doesn't it, from 14-inch brown disks to what we do today. It's amazing. But yes, I retired. What does this chart say? Mainly being lazy, yes, that's about right, although Tom has been making me work very hard putting my oral history together.

Gardner: No, I would disagree with that. You volunteered.

Wood: I've given this talk, "Zooming in on Data Storage"[33]. I've given that talk quite a few times, and that's interesting to do. I've enjoyed doing that. It's not a very technical talk. It talks about— tries to involve data storage as you go from the extremes of very large distances to very small distances.

I've got a university project. It relates to what we said earlier, trying to move into that third dimension, put multiple layers of magnetic data on top of each other, and that's on going at Washington State University and Nanjing Institute of Technology.

[03:20:15] And it's really going nowhere, to be honest, but it's a good training ground for students. But the reality is you can't get anything useful from anything but the top layer. The other layers are too far away to either write properly on or read properly on. So I mean, at the moment, they have a paper they've put together which I think says 20 percent gain or 10 percent gain, but that assumes you can build two layers of perfect recording medium with exactly the right characteristics on top of each other and that's never going to happen.

Gardner: And then you focus the magnetic reading? You try to somehow focus the magnetic reading? The optical guys tried to do that. They can focus light.

Wood: Yes, you need to be in x-rays or gamma rays to get into the ballpark where we need to be. Yes, the only thing is, and I did look into this quite seriously at one point — you've heard about racetrack memory?

Gardner: Yes.

Wood: I forget the gentleman's name now, but is there some way you could write...

Gardner: Stuart Parkin[34].

---

[33] https://hgst.app.box.com/s/dzm9mx2qng1hh31v607n26fgn51yng66
[34] https://en.wikipedia.org/wiki/Stuart_Parkin

Wood: Yes, Stuart. Yes, thank you. Is there some way you can write on one layer, on the top layer presumably, and transfer it down into lower layers and there are various physical mechanism you can imagine trying to do that with and then sort of bring it back up to the surface when you want to read it back.

Gardner: I see.

Wood: But there's nothing practical that you can do. And then the other thing there is the Computer History Museum and I've had fun with that involvement. I've written a bunch of Wikipedia pages that we mentioned before. I've written a couple of milestones, and the big project has been this project with Tom. And is this the point, Tom, where I should thank you very graciously for all the work you've put into doing this? It's been quite a remarkable experience, and you must have spent hours and hours on this because I know you've been looking through all the boxes of stuff that I gave you, as well. So thank you very much, Tom, and thanks to the Computer History Museum, too. Much appreciated.

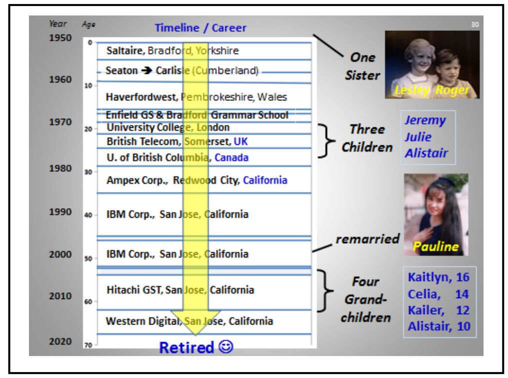

Gardner: Well Roger, I have to thank you. This has been an unusual oral history. I have not had the privilege and really pleasure of working with someone who's been so thorough in preparing for his oral history, the 30-some-odd slides we've collaborated on, but it is unusual. I've not experienced that before. I've had other people prepare exhibits but not to the extent that you've been and even made it very easy for me to moderate, and I've had fun fact-checking you from time to time.

Wood: <laughs> Yes.

Gardner: For the record, I note that Roger has offered the museum three boxes of physical documents going back to his Ampex days through his Western Digital days. And in so far, I have 218 soft-copy documents that Roger has offered, soft copies of his many presentations, many of which don't exist in any other repository other than Roger's garage, and a bunch of copies of his many papers, most of which are online or available in hard copy form but all of which, along with the recordings of these interviews and the slides used will be offered to the museum and they will be researchable under the code A2020.7557. That's the code number that the museum will use to record Roger's offering and most of which, I would expect, would wind up in the museum's permanent collection. So I thank you, Roger, and looking, what we've talked about, the next slide kind of summarizes the discussion from Saltaire in Yorkshire, UK, to retirement here in Gilroy.

[03:25:00] **Wood:** And along the way, I picked up three nice children, Jeremy, Julie, and Alistair, and I remarried, Pauline, and by now I have four lovely grandchildren and my sister, of course, who's still in England. And the next slide, I think is, yes, my nice Thanksgiving from last year with all the people mentioned there, my four grandchildren, my sister. There's five miscellaneous spouses, not all mine, of course, but between the group. But yes, it's been a wonderful adventure, I guess, my career and my family and everything, and hopefully I've got a few years left. We'll see what the COVID-19 has to say about that. Shouldn't joke about that.

Gardner: Well, sheltering in place, that's one way to avoid it. We're a little out of sync. I was going to ask you about comparing IBM, US versus Japan and...

Wood: Oh, yes. To be honest, I haven't really taken that much notice of my surroundings. I've always been focused on what I've been working on, what I've been enjoying. IBM was very IBM. I remember when I joined, I was quite amazed at how much everything was organized, probably well organized, but it was very organized, and people were very IBM-ish. I remember when I joined, there were three people. I kept getting confused as to which was which because they seemed somehow so similar and so IBM-ish, and they were, Jim Belleson was one of them, Roy Jensen, and Larry Rosier. And it was funny because they all seemed so similar to me. It was a very strange experience, but I think that's true of IBM that people were there for an entire career, typically, and they became somehow very IBM-ish. They all had the same sort of outlook, same ideas, followed the same rules, et cetera. It was interesting, and Hitachi, things didn't change a lot in Hitachi. I think I enjoyed Hitachi most out of all these companies. No, maybe Ampex most, but Hitachi, as well. Hitachi tried to emulate, I think, IBM in many ways, and a lot of the structure was the same. One of the most interesting things I noticed was the difference between IBM research and Hitachi research and I'm exaggerating but I tried to encapsulate it this way. IBM Research, Almaden Research Center, there's a whole bunch of researchers, individual contributors, and they would all be heading in different directions. They would all be researching different things in different directions, and if you looked at the research in Odawara, the research crew there, they were all following the same direction. They were all in lockstep as a team, heading in one direction, and I don't know which is best. I mean, if that had been the wrong direction that they were all heading in, but there's certainly a difference.

Wood: And then WD, there was this funny, hold-separate thing, so things didn't change very much with WD. But it got quite unpleasant in the end in the sense that they were getting rid of all the nice offices that we had, and they were merging the groups. I think the media team kind of— quite unhappy, in some ways. The HGST media team got merged with the WD media team and changed locations and stuff, and I know there was quite a bit of upset about that.

[03:30:00] No, I'm not sure. As I say, I haven't noticed much difference, really. I haven’t been very observant as I've gone through life, perhaps.

Gardner: So you really didn't see the WD side, then, of how they did their development and research compared to the way you found Hitachi and IBM were similar. WD, I guess you weren't there long enough?

Wood: Yes, there wasn't really a research effort at WD, not what I would call research. There was development and maybe advanced development. **B**ut now there's a trend. I mean, the Hitachi research department, the HGST research department was disbanded, as well.

Gardner: So tell us a little bit about the history museum and the storage SIG. How has that been going for you?

Wood: Oh, it's been fun. We have a monthly meeting, and the main emphasis has been on creating oral histories. So I've done my part here, and I'm supposed to do one in the future on the Ampex PRML plus the IBM PRML, so that will be fun, as well. But yes, it's been an enjoyable interaction, that, and there are several people, of course, I knew prior to joining that little group. Dennis Mee is a regular participant, Chris Bajorek and others, so that's been fun. So it's kept me occupied.

Gardner: Have you taken your grandchildren to the museum?

Wood: Have I? No, I've taken my sister. My sister was delighted because her boss [Dame Stephanie Shirley, 2018 Museum Fellow] was there, or a picture of her boss about 30 feet high was there because her boss was giving one of these sort of invited lectures on something or other [CHM Fellow!]. Yes, so she wasn't expecting that. So we walked in there and she was confronted with this enormous mural of her former boss.

Gardner: Okay, I'd like to thank you for the time and effort you've put in, too. It's been, as I said a few seconds ago really for me, an enjoyable experience, a level of detail that I've not experienced many times in these oral histories and actually, frankly, it's going to make my editing of the transcript a lot easier because we've gone to a lot of effort to make things consistent, coherent, and clear. So thanks, Roger. It's been a lot of fun. Stay safe in Gilroy. I'll stay safe here in Los Altos Hills.

Wood: Likewise, Tom. Thank you so much for organizing all this and putting all the work in that you have and, as you say, I really enjoyed the interaction, as well. It's been wonderful. Thank you very much.

Gardner: You're welcome, and good bye.

Wood: Goodbye.

END OF THE INTERVIEW

