Oral History Panel on Advanced Ferrite Disk Heads

Interviewed by:
Ian Croll

Recorded: November 3, 2005
Mountain View, California

CHM Reference number: X3331.2006

Ian Croll: Well, we're here to discuss the early and late developments in ferrite heads; my experience is in film heads. So on my left, Paul Frank, Mike Warner, and Erik Solyst, who actually participated in the business over several years. So perhaps we can start off with really the beginning end of what it replaced and what ferrite heads are.

Michael Warner: I think we also introduce ourselves.

Croll: We could do that, yes.

Paul Frank: Okay, I'm Paul Frank. I'm pleased to be here. Thank you for inviting me in. I spent thirty years associated with Applied Magnetics Corporation which was at various times one of the world's largest manufacturers of hard disk drive heads as well as tape heads and certainly was a major manufacturer of ferrite core, ferrite based, and what we call monolithic ferrite heads as well over the years. I started in 1969 just at the tail end of the generations previous to ferrite and I must say that most of my career was spent developing thin film AMR and GMR heads, but I did certainly participate earlier in my career in the ferrite head development as well.

Warner: Oh, I'm Mike Warner. I started with IBM in 1965. We were still making at that time laminated heads so I participated in the transition from laminated heads to ferrite heads and then on to thin film heads so I've spent about twenty years in the technology of magnetic recording in IBM. Erik?

Erik Solyst: I'm Erik Solyst. I started working for IBM in 1961. That was before the ferrite heads, as Warner just mentioned. There was laminated new metal heads, so I went through the whole transition from laminated heads to ferrite heads and the sort of early beginning of thin film heads.

Warner: Ian suggested that we at least give a little history or chronology how we got to the ferrite heads.

Frank: Yeah, courtesy of Denis Mee and some of my own samples, I've brought samples of virtually all the generations from 1301 up to -- I may even have some of the early thin film heads in here. Yes, I do. So it basically encompasses the ferrite core generation of heads. Now this is the one sample I have that actually predates me and I think probably predates Mike. This is a 1301 head assembly from an IBM 1301 drive. My guess, based on my experience with later generations, is that it's a laminated Mumetal core so it predates the use of ferrite for cores. So I'll pass that down and let you guys comment on that if you care to.

Warner: Erik, you were part of that.

Solyst: No, actually this was just prior to when I started at IBM, but it is a 1301 head. 1301 was a twenty-four inch disk, two stacks of twenty-five disks and humungous hydraulic actuator to move this whole stack.

Warner: With a drip pan underneath to catch the oil that leaks.

Solyst : These sliders were the first cylindrical hydrodynamic sliders introduced in magnetic heads and they were flying at two hundred and fifty microinches, if I remember correctly.

Warner: I think that's about right.

Solyst: It was at least in…

Warner: Two hundred and fifty, yeah, correct.

Solyst: 1960…

Frank: '62 according to the…

Solyst: '62.

Frank: And twenty-six kilobits per square inch.

Warner: Yeah. Now if you notice, in this case there's a slider or a magnetic head that's facing up and also one that's facing down so this mechanism went in between the disks, and then had a torsion rod that would turn and load these heads against the disk so that was the way that it provided pressure to load the head. And then the air bearing, cylindrical air bearing that Erik mentioned, provided the opposing force and the head would ride on this air bearing over the surface. So this is kind of the basic technology of flying heads on large disks. And then the next thing we moved to…

Solyst: And the loading force was three hundred and fifty grams.

Frank: Right, it was very high.

Solyst: And that has pretty much been repeated in the next three or four levels of head that we'll show you.

Frank: Yeah. The next generation sample I brought is this one. This is a 2311 head. This is the first generation I actually worked on. I was a student in college and I started working summers. Just to tell a little story about my first introduction to the head business, we were making these for IBM and my job was to adjust the attitude on these. That was before the current connotation of an attitude adjustment, but in some ways similar. I would come in on Saturday, primarily I believe so that the IBM inspectors couldn't see what I was doing, and we had an open disk pack and we'd run these things in and I would eyeball sort of roughly the angle that they came into the disk and I would pull them out and then I would use my highly calibrated pair of pliers to twist the suspension and then run 'em back in and see if it was a little better. And, you know, I had one to five degrees, something like that by eyeball and I would do that until I got them right and then I would move onto the next head. And we didn't do that during the week, just on Saturdays. So similar to the last one, this is a stainless steel air bearing and it is a laminated Mumetal core. I think unlike the last one, it has a slightly more sophisticated air bearing with the two holes in it to achieve the proper attitude. Another thing that's interesting about this one to me is that we sometimes think about technologies being reinvented, this is ramp loaded. This is cammed up and runs off the cam and then this comes crashing down with several hundred grams of load into the air bearing space, I believe typically bounced once, say we're on oxide disks, and then established an air bearing and flew after that. So hopefully flew.

Warner: And if you notice on these guys on the face of the bearing there are two small round holes and that was used to bleed off and lower the air pressure on a portion of the air bearing so that the head could fly at a lower height. Do you remember what the spacing was for these guys?

Solyst: Yeah, a hundred and twenty-five.

Frank: Right, a hundred and twenty-five microinches.

Warner: A hundred and twenty-five microinches.

Solyst: I remember. I worked on that.

Frank: We weren't into nanometers then.

Warner: Yes, actually within this structure, this laminated core is held in with -- there were three- was it three holes? I think three holes with a little ball bearing kind of thing that pressed in and pinched the head, and a ball staking technique that Erik invented for clamping the read/write elements in place. Okay.

Frank: So the next generation sample we have is a 2314 head which we always referred to and I think most people did a monkey face head 'cause it's round, it has the two holes for the eye, it has the ferrite core which looks like a nose, and then it has a crosspiece which actually contains a Mumetal straddle erase core which is wound separately that looks like the mouth. So, you know, this was the next generation sophistication. The stainless steel was replaced by -- I'm thinking it's alumina.

Warner: It is alumina.

Frank: Yeah, it's a solid alumina which is polished and it's actually contoured. It has a crown on it slightly to the side and more to the front and back. It isn't flat. What else can I say about this? It's the first ferrite core generation that I'm familiar with, which of course is what we're primarily here to talk about. Just one little anecdote about these; at Applied Magnetics, we made some really good money making these heads well into the 1980s even though it was introduced I think in 1966, and that's because some of these 2314 drives were installed on Polaris submarines by the navy and they were never willing to upgrade the drives so they just kept ripping the heads out and we sent them new heads to replace so that they could continue running those drives. And the upshot was we had to continue making these things until they decommissioned the final Polaris submarine. And, of course, the prices just kept going up over the years as we continued to make them.

Warner: Yeah, just as far as the ferrite is concerned, the ferrite cores, at least for IBM, were made in Poughkeepsie, New York, in a group under Miles Cook and the cores were made as kind of a unit like this with a flat piece and the gap was a glass gap and a process used for determining the throat heights was bevel grinding. And then these would be processed in these bars and then sliced up and lapped to thickness and then cut in two. So instead of having -- you'd dice this in two so that you'd have half a core and then these would be bonded into this monkey face with epoxy and then a coil was slipped over the top and then a sidebar put on so the magnetic path was made by actually gluing on a piece of ferrite to form that. And then also, as Paul mentioned, there was a yoke for having a tunnel race on either side. Now this particular product had a fatal flaw in it. It had a -- the core was held in with epoxy and as you changed the humidity, then the core would walk in and out.

Frank: Actually, it would swell and the core would move.

Warner: That's right. And the manufacturing people quickly learned this so that they would -- you know, if they tested the parts and they tested bad, they would wait for whatever it was, a dry day or a wet day and then go back and retest 'em so that they would test good. You know, fortunately there was enough, you know, margin in the product, you know, that it worked fine in the field anyway. But an enormous amount of energy was made. Every head had a tally trace over the top to measure the pole tip protrusion and recession.

Frank: I remember that we used to cycle those things through temperature and humidity cycles and then do a final polish in hopes of stabilizing the epoxy for actual use in the drive and it never quite worked right. I think it improved things, but it never quite..

Solyst: The reason of course we didn't use the ball staking which had solved an earlier problem with similar heads is that the ferrite core is very brittle and if we had staked it in as we did in 2311, you know, it would just have crushed. So we were back to an old problem that we had actually experienced before.

Warner: Yeah, see these were epoxied in, right?

Solyst: Yeah.

Warner: So went from epoxy, to the staking process that locked it in, back to the epoxy. But the ceramic was very good for handling the loads and it was more forgiving from a wear standpoint.

Croll: What were the advantages of ferrite over the Mumetal?

Solyst: Well, there were a number of reasons for why we changed over to ferrite. There was an ever increasing demand for higher densities, higher yield densities, lower flying heights, tighter control on the gap width because of the higher bit densities and we could just not achieve those goals in the high permaloy heads within the cost objectives we had. We also had, of course, the frequency, as we increased both the disk speed and the linear bit density the frequency increased and the eddy current losses in these laminated cores increased. So for all of those reasons, we decided to switch to ferrite. And I should point out Mike mentioned that the ferrite material itself, at least for IBM, was developed in our Poughkeepsie laboratory, and they also developed this glass bonding technique which was very significant. They were people that had been doing ferrite magnetic cores before that using epoxy bonding which have always -- some of the early problems that we had when we epoxied the core in. But the glass provided a very, very good surface to maintain the sharp edges on the magnetic core in the lapping process and, of course, the glass was thermally matched to the ferrite so it didn't crack. That was a very significant improvement and that permitted us to achieve the kind of both track densities and linear bit densities that we achieved with the 2314 heads and later on with the 3330.

Warner: Yeah, I would say that we had to do a pole tip inspection on the Mumetal heads to check for the gaps, the alignments because as they were put in there may be some shifts or things of that nature that take place. But with a ferrite glass bonded together, once you had a good core, it stayed good. It was a very stable device.

Frank: Yeah, I agree. I think the glass bonded gap definition where it was -- you could do it very precisely and do it consistently was a big step forward. I think that was probably one of the reasons that these heads were epoxied however is because if you tried to do a second glassing process you might undo the gap, and that was one of I think the advances in the next generation which was the 3330 head. It is a ferrite core and I believe that's -- is that calcium titanate? It's some calcium compound, rather than alumina.

Warner: Was it barium?

Frank: Barium.

Solyst: It was barium titanate.

Frank: Barium titanate. You're right. So it's a little darker. But I think that one of the big advancements here was that it has a two-stage glassing process. The core is made in a glass bonded fashion using a very high temperature glass or a relatively high temperature glass and then the core is put into the air bearing housing here with the lower temperature glass which didn't disturb the higher temperature glass hopefully so you could preserve the gap on the core. So it's a slightly different air bearing. It's a two rail, but it's crowned front to back and not side to side, at least not intentionally, although they always -- it seemed to have some side to side contour. And I think this was the last of the real heavily loaded generations with one hundred grams or more, although this was a little bit high or low.

Solyst: Well, not quite.

Frank: You guys can fill in the date there.

Warner: Within IBM, within the system there were actually -- the 3330 had a high performance drive need that was done in parallel with this and in the IBM labs development took place on two different products at the same time. So in the same building thirteen we had two engineering teams, one of those doing a fixed head which I will let Erik describe because he's the father of that and mother of that particular product which is an all-ferrite head. So this was the last of disk packs, removable pack loading and was the culmination of the ferrite technology in this stream of products, at least for IBM. Erik?

Solyst: Let me just say one more thing about the 3330 head. One of the challenges in that head was the fact that we had four materials that needed to be thermally matched; that is the slider material needed to be thermally matched at thermal expansion to the ferrite material and then we had two bonding glasses, the high temperature and the low temperature bonding glasses. And all these four materials needed to be thermally matched. If they're not, it'll crack, the assembly will crack during the cooling off period after the bonding process. And that was -- one of the things that was developed in San Jose was the material which we mentioned was barium titanate, right?

Frank: Barium titanate, right.

Solyst: Barium titanate. That happened -- in the mixture that we came up with happens to have the thermal expansion characteristics of the ferrite. But it was very labor-intensive work to assemble this. It was time-consuming, it was expensive. The yields that we were running were fairly good. We were running magnetic test yields up in the eighty percent or so which we felt was pretty good, probably better than any of the other previous heads we had had. We were in an overall yield at about seventy-five percent, I believe, on that so we were happy with it at the time.

Warner: You just reminded me of one other area on that. Within this -- the ferrite core needs to be positioned very accurately with respect to the air bearing, meaning that there's a throat height control for the core, so the core had to be positioned relative to the air bearing surface and then glass bonded or held in place, so that's part of what contributed to the processing cost of this style of head and it still had a manual coil that was placed over it.

Frank: Just to put it in time perspective, this generation of technology was introduced in 1971 and it was still a little less than a megabit per square inch. It was about seventy in the first embodiment and then there was a second generation that actually went over a megabit per square inch. Did you want to see this?

Solyst: Yeah, that's right. It was. Linear bit density was four thousand and forty?

Frank: Four thousand and forty, right.

Warner: Is that right?

Frank: That's exactly right.

Solyst: And the track density in the first model was about two hundred, a little less than two hundred and there was -- that increased to three hundred tracks per inch in the second one.

Frank: The second generation, right.

Solyst: Yeah.

Frank: So the next technology we have here is what we always called monolithic ferrite heads which were essentially glass bonded, but where the air bearing and the core were all made out of ferrite glass bonded together and the air bearing was ferrite to the machine. And the very first one of those that I'm aware of is that device there, the 2305 head per track head. So that's one we never did and I have little experience with so I would leave it to Erik 'cause he's intimately familiar with that one I think.

Solyst: I'll pick up on that one. Yes, actually I am and I remember that with fondness. I was very proud of that.

Warner: And you should be.

Solyst: I was proud of it. I got to tell this true story. I was proud of it because in the early development phase manufacturing engineering had heard about this development of a nine element head and they had been struggling with single elements, you know, for years. And all of a sudden during this early phase of development a deputation of manufacturing engineers came to my office and said, "We want to tell you that it is not possible to make a nine element head." And they had picked up the yield numbers from I think it was 2314 at the time and then they had, you know, raised that number to the ninth power and demonstrated to me that we would wind up with a two and a half percent yield or something. I pointed out the flaw in the analysis, and I'm not sure that I convinced them at that time. But anyway, we talked about, you know, the time-consuming effort and operator dexterity necessary to meet all the requirements of the 3330 heads. And so at the time when we got to this nine element head we were striving for design simplicity, especially a simplicity in the material selection that goes into the head. And we developed the concept of integrating the air bearing slider into the magnetic core structure which meant that we had only two materials to work with, the ferrite and the glass. They were a very simple material system and glass was available to us so we eliminated all the other considerations. This was the first monolithic head and because we had excellent control of the core reluctance, we had eliminated the variability that we had in 3330 in the back gap reluctance. So we had extremely good control and a very favorable ratio of core to gap reluctance. And it turned out that, of course, it worked to our advantage. We had essentially one hundred percent yield on these heads. And I do remember that I -- when we got the first performance analysis of our heads, I went back to manufacturing engineering and said, "Look, this is how it came out." Now another thing that was unique on this was that once we had the basic structure which was a very simple structure, since the U-shaped core and the very large side leg that also performed as the slider, once we had developed that concept, the next concept that we had to fight was the air bearing design because up to this point, all the air bearings had been cylindrical and the minimum point in cylindrical air bearings is not at the trailing edge where we would want it to be for this head. So we had to look for some other air bearing shapes than cylindrical and fortunately we had a very efficient computer modeling program for air bearings and we quickly figured out that, you know, many other shaped air bearings that will work and came up very quickly were the taper flat slider that was implemented in this 2305 head. The taper flat slider is such that it would always assure that the minimum gap, minimum head to disk spacing is at the trailing edge of the slider where in this head the cores are located or the gaps are located. It turned out to be an excellent air bearing. It performed as well as cylindrical air bearings and in some respects even better than cylindrical air bearings. And the taper flat air bearing has been used I think in all subsequent recording heads, even today, I believe.

Warner: Variants of it, certainly yeah.

Croll: Did that go into the Winchester product?

Warner: That taper flat? Yeah, I'll kind of explain, because this is the…

Croll: That was the first implementation.

Solyst: Now in the 2305 file or I should actually say buffer, it was a buffer, not a file, there were ninety-six of these heads. Now that was one reason that we had to, you know, do the very best we could to get the cost down and the yields up. There were ninety-six of these heads on four disks I believe, on eight surfaces, I think that comes out right. And the disk was spinning at six thousand RPM, faster than any other disk file at the time. And because of that high velocity and the very relatively large air bearing surface, the load to get this head to fly at fifty microinches was twelve hundred grams. That's a lot. We got the twelve hundred gram loading force from a pneumatic system that we had actually copied off an early magnetic drum assembly that had been developed in Kingston. It was a closed system, no possibility for leaking dirty air into the file and this just worked fine. We developed with this head a batch process. Our manufacturing process was the batch process in which we got eight heads from two bars of ferrite that were glass bonded together and then the next operations was operations of grinding slots to separate the elements and that was essentially it; a fairly simple manufacturing process, but one also that required a very precise grinding process. And manufacturing engineering was actually very effective in developing instrumentation, computer controlled instrumentation to control the location of the grinding equipment. It was totally automated, did a super job, and once we learned how to control the grinding process, that was which kind of grinding wheels to use and it was all diamond grinding wheels we used, once we got the right cool in it and the right speeds, you know, it was a very, very successful operation. And since these grinders were automated, one operator could run four machines at a time so that was very good. I think that's -- no, you worked on that.

Warner: Oh, absolutely, absolutely. There are a couple of other things to mention. This was loaded dynamically with the air so the disk is spinning and so this was plunged onto the disk. And internal to this assembly is a little pocket where the push rod would push and the location of that load point was very critical as far as the flight was concerned. So there was a bar structure that goes through the center with a hole in it that we could push. And then because the head would be gimbaling what you wanted to do is reduce the amount of corrosion, fretting corrosion that would take place on that load point, so that was designed to match the loading system and the load point pushing with the air bearing forces that took place on it. Also to make the taper flat, Erik came up with a sign bar kind of lapping technique that would allow you to lap these to a controlled throat height, you could examine the ends of the part, because you could look at the glass gap throat and make sure that the throats would be even across the nine elements. And then after it's gotten to the right throat height, then a tilt would take place on the lapping plate and that would then lap the taper portion of it. So even though this is a very shallow taper, the location of that taper was quite critical so we're talking about microinch controls on that. And yet the process was quite elegant and cheap, relatively inexpensive for performing that operation. Did you guys ever make any of the Zeus heads?

Frank: We did not. I mean our first..

Solyst: Let me also mention another feature. Mike was talking about that bar that goes across the middle. I don't know if you can pick up this picture here. This is the nine element head assembly and there's a slot ground transverse to the element in which the suspension and this is part of the suspension, this is the suspension attachment and it has four clips that clips in. It's a snap-in type suspension which was easily mechanized, just slip it in, whereas in other assemblies we had used welding techniques and epoxy. And I remember that. Again, I was very proud when I could demonstrate an assembly without epoxies, no epoxies. And I was- at the time I had had very, very bad experience with epoxies in previous development heads and I was at the point where I would not permit epoxy in my laboratory.

Warner: A very interesting point. As Erik brought up, when that bar fits in and clipped in, I had forgotten about that, in the very early versions we did epoxy and you can see they're -- I mean here's a magnetic head, there's a slot cut in it and a bar put in and when you put epoxy in, as the humidity changed, the epoxy would swell and as it swelled it would change the curvature of the air bearing. And so this clip-in suspension allowed a constant force that didn't vary and the bearing stayed the same over all kinds of humidity conditions.

Croll: Paul, why didn't you get involved?

Frank: The market for head per track drives which were essentially buffers was very small.

Solyst: It was not a big market.

Frank: You know, but that led directly to what I think was one of the major breakthroughs in technology for hard drive heads which was the 3340 generation. This is the monolithic ferrite single-track in the center, taper flat air bearing that landed and took off from the disk. This is -- oh, okay. I wasn't sure what the signal meant, but this is a 3340 arm which has four heads on it and as I said, they actually are loaded and then these pins are pulled out and they rest on the disk. And when the disk spins up they take off and when the disk stops they land again. And they're very lightly loaded compared to the earlier generations which were loaded with, you know, as we saw several hundred grams. And this was just a major breakthrough. It took those of us trying to follow this in the industry several years to catch up with this technology, but ultimately we kept pursuing this technology and generations of it up to the mid '90s and this is where Applied Magnetics made most of its profits over the years.

Warner: Yeah. Let me give you just from an IBM perspective, we had a program that was going to start off that was going to be much higher density and a removable pack. And in order to do that a decision was made to make the heads reside with the disks. And so with these kinds of products, packs were interchangeable and so you had the tolerances from pack to pack. And in this first application, the pack and the heads were gonna stay together. And then we're just about the end of this tape so what I'm gonna do is stop here and I'll explain then.

<Break in tape>

Warner: Yeah, at IBM, we started a program that was called the Winchester program. We had these- we were to make a disk pack with the heads residing within the pack. It was determined that, or at least it was suggested, that we needed to fly at much lower flying heights. I think these were what twenty -- the first ones were twenty microinches?

Solyst: Yeah, eighteen microinches.

Warner: Eighteen microinches and much lower loads. And a technology had been worked on for doing starts..

Solyst: May I interrupt you here because I think it's interesting that when we first started on the Winchester program it was supposed to be a recording head in contact with the disk. And some of our engineers had designed a head that was supposed to stay in contact with the disk and could change a single element. And this head was loaded with just two grams and that very low load was necessary to eliminate wear on the disk. Well when we got started and made, you know, a couple of dozen of these heads, it turned out that they operated very erratically and we, you know, nailed it down immediately that it was because it did indeed not stay in contact. Once the disk got up to speed, some head would be flying as high as thirty-five microinches and some would be hopping and skipping on the disk, which was not a very good situation. Well, it turned out that these in contact had some very small flats running on the disk, and once in awhile the edges would be worn down a little bit or rounded a little bit during the lapping process and then would actually result in a sort of a taper flat air bearing and therefore some heads were flying up to thirty-five microinches. Once we discovered that, we said well, let's not fool with that. Let's design a flying head to fly in between these two values. And that's how we came up with eighteen microinches. And I believe that the load proper was what, sixteen, fifteen grams?

Frank: It was on that order, yeah, 'cause it's not in the chart here.

Warner: Well remember, Erik, when we first started out, we had a tripad head?

Solyst: Yeah, well that's the one I was talking about.

Warner: Yeah. Okay, we had a tripad head with two rails and an element that was glassed in, you know, all of this technology in the back so it had a discrete core. And one of the problems with that is making these internal taper flat bearings was extremely difficult. It was not a very cost-effective technique. And when we first went to Menlo Park that's actually what we were building.

Solyst: That's right.

Warner: These tripad with glassed in heads.

Solyst: But we were building those, but we knew they were very difficult to build, and they were absolutely not a manufacturable head. But we were building them for our systems development people so they had some to work with. But that was tough to build.

Warner: Yeah. Actually, it was I think fortuitous that the team that was assigned to work on this and in the IBM lab there were actually two magnetic head teams, two disk drive teams, but this came out of the group that was doing the Zeus or a monolithic head because we knew that manufacturing had this very sophisticated, we used to call it a stiff finger, meaning that they would just poke the button on the machine and it would go through and do the grinding and the slotting and things like that. They had either laser interferometers or glass scales built into the bodies of the machines so they could make very accurate indexing. As we started working on this, the challenge then was how do we get something to fly very low without these, you know, twelve hundred grams that these parts took. Actually, I was the engineer assigned to work on that problem. And in this design you notice that we put a slot across the middle to make it a dual rail. And as I started doing more and more analysis, and in those days I'll tell you how things were. The analysis was done by punch cards. I would just punch up the deck of cards, submit it to the research computer and run the air bearing..

Frank: I had that same job.

Warner: Run the air bearing analysis. And as we started working with that we discovered that as we made these rails smaller and smaller they became kind of like skis with the taper flat being, you know, the bearing with the taper and a flat and the element, read/write element at the rear, that choosing the right taper angle and widths that it would build up a pressure in the front, then the air would leak out the sides, and as it became closer and closer to the disk, because it was flying at an angle, pressure would build up again. So on these two rails you sort of had four posts of pressure holding the part up. And then the idea of making a three rail head, meaning two rails on the outside to make it fly and the center one was defining the track width really came from the manufacturing process that we had that we could make these angles. It had the advantage of being able to vary the track widths, literally with a computer program and we could make single element heads. And we'll talk about that. We made a fixed head version that had three read/write elements on it. I think that we -- what was our- our target was a dollar a head, something like that?

Solyst: No, it was higher than that. I think it was five dollars a head. And very early in the manufacturing program I remember we'd go down to four dollars and fifty cents and I'm sure it came down further than that, but I don't know exactly what the cost was.

Warner: Do you know what you guys built them for?

Frank: I don't recall. I don't recall. I'm sure it was more than that when we started.

Warner: Well, I think that these- if I remember right, you know, later on after I left IBM and we were buying these heads and using them that I think that they were down in the dollar a head, dollar and a quarter per slider kind of area.

Frank: It was probably in that ballpark, yeah.

Solyst: That sounds right. If it cost us four dollars, it would probably..

Warner: Cost somebody else. The initial versions of this used the same clip-in technique that was developed for the Zeus and a suspension gimbaling assembly with a load beam built in. So these parts, one of the challenges, because they were gonna go four parts per arm, is that we could individually test these and then be able to assemble 'em onto an arm and with a very high yield. I think also these were -- the technique for attaching to the arm was done with a swaging operation.

Solyst: Swaging, right.

Warner: Yeah, that was developed by George Paal. And the swaging was stainless steel into aluminum and there again, we took advantage of the coefficient of thermal expansion mismatch and the difference in modules so that the stainless steel would swage out into the aluminum and the aluminum would act as a constant spring to hold this. So I think -- and also you notice in this guy there's a little we call it a pickle fork. And the pickle fork went onto a cam on these assemblies and held them down, so they could slide between the disk, then you'd pull this pickle fork out, and then as you pulled it out then you can see the heads would spring out and load against the disk. And so here's the other assembly.

Frank: Yeah. I'll let you put 'em back in for me. So, you know, as I said before, this is one of the most major transitions technology-wise in hard drive heads, just because it encompassed so many new technologies from the air bearings to other things and very much influenced everything that came after in terms of the taper flat air bearing, the contact start/stop. It took us a couple of years to respond to that technology and produce something similar. And as I was telling Erik, it wasn't so much the technology of the ferrite; it was the fact that these things had to be made in clean rooms and that was a whole different culture for us. We had to learn how to build clean rooms, we had to build clean rooms, we had to learn how to put people in clean rooms, how to get them to work in clean rooms, because all of this, as Mike mentioned earlier, had to go into a sealed assembly which had to be kept clean until it was sealed up and shipped off to the customer. Now I did have one question for you because this is only part of what goes into a 3340. This is the other part and if that gave us grief, this gave us no end of grief. There are ten heads, three tracks each on the trailing edge so there are thirty tracks. This is the head per track piece that goes on the outside, outer surface.

Warner: Actually the under..

Frank: Yeah, the bottom surface of the 3340 pack. And I always wondered why you did this and not this which you could have done with maybe three heads. So maybe someone can explain that because those things were a devil to build.

Warner: Well, the reason for doing this was just because if you look at the machining and the process that came off, this made, you know, a single rail and we said well, if we just make each rail narrower it would fly at the right height and put a coil on then we would wind up with three elements per head and it was made on the same line and it literally was a change in the machine..

Solyst: Computer program.

Warner: Yeah, computer program that was beveling the edges and resulted in a head. And then we didn't have to cope with the large loads that were involved here. We could use the same suspension and the same ball stake operation so it was very easy for us to manufacture. Now Paul mentioned about clean rooms. The disks that these were flying on were lubricated disks and so lubricated disks need to be very, very clean because we're flying at eighteen microinches and later on later versions of these even flew lower, and so very fine particles could get between the head and the disk or contaminate the heads. Actually, the heads as they go along are very good collectors of contamination. They actually, you know, if you wanted to sweep the disk and get it really clean, a good way is to take a magnetic head over the top and it collects those debris. So cleanliness was very important for these kinds of applications.

Croll: Was this the first stop/start in contact with the lubricated disk?

Warner: There was a – what -- the file the Buslik did was a tripad start/stop. What was that, a ram buffer or something like that? I can't remember what that was.

Solyst: Well, I think you're talking about the ram kit, but the ram kit had a 2311.

Frank: Yeah, that certainly was the first volume product.

Warner: Oh, yes.

Frank: With the taper flat contact start/stop.

Solyst: I think it was the first one, yeah.

Croll: Well now with the fixed heads, just a question, on the bottom, was there a stiction problem with that?

Warner: A stiction problem, we had stiction problems with the head-media interface all the time. It was something that we fought, to cleanliness, the right amount of lubes. You may know that, you know, if we had a good case of stiction taking place then it could actually -- when you started to either stall the motor or rip the heads off the suspensions as the drive started up. It's one of the I would say disadvantages of a taper flat bearing is that you now have a flat and so if you have a lube and a flat then you can get some form of stiction. But the real guys that stuck, was when contamination would gather, salts or things like that and the drive went through a humidity change like in a product test lab and those salts would crystallize and then you had a very severe stiction problem. So having clean disks, well-controlled environment, and well-controlled lube process was extremely important for getting these guys to work.

These early guys were -- you notice are four heads per arm and this -- what happened outside of IBM very quickly was this technology was adopted for eight inch drives and five and a quarter inch drives.

Frank: And three and a half inch drives.

Warner: And using the same technology also with oxide media, but also with thin film media.

Frank: Ultimately, yeah.

Warner: Yeah, with thin film media.

Solyst: Except IBM.

Warner: Except -- yes, IBM was slow to pick up on that. A long story, but the interesting thing is these worked extremely reliably. The smaller the disk size goes, the reliability goes up and the sensitivity to contamination actually drops down. And I remember in development labs at Micropolis and Maxtor where the drive development took place out in the open lab and they had a sandwich bag, you know, the plastic sandwich bag and that was opened up, the mouth was opened, the drive was inside, the heads were on it over a thin film disk and they would operate for months and months and months. Completely different environment when you had small disks, and I think it had to do with the amount of energy that was transmitted whenever you had a contact or a particle between the head media, that with small disks there wasn't enough energy to destroy it; with a large disk you had enough energy that you -- the heads would bounce around or start scratching or destroying the element. There's another area that I might mention on these guys. These are all on actuators where they're moving quite rapidly so the suspension designs to go with the ferrite heads was always a major effort. And so these designs that you see here are all linear in nature and if you notice, we have these little cams on the suspensions for holding them down for loading them in a pack. We supported a project out of Hursley, England that used a rotary actuator. And I remember going over and there was the guys having Bodie plots and looking at the suspensions and saying -- I mean the performance, and saying, "You know, we're flying along, everything's looking great and then all of a sudden we'd get these big spikes that our data gets corrupted." And they said, "We have a resonance at this particular frequency." So we looked at this and I got out a pair of snippers and I cut off these little cams. And they said, "You can't do that." I said, "Why not? We designed it. We can take 'em off." And it turns out those little cams were acting as a device to torsion the head and change the spacing. So the design of the head suspension assembly, magnetics, all get interrelated with the mechanical and the magnetics portion of that. Go ahead.

Frank: Speaking of suspensions, this is an interesting specimen. It looks like a prototype of a 3350 head and one of the big innovations for 3350 was I believe it was the first drive where the preamp and right drivers were put on the arm in a chip and was integrated into the arm assembly rather than being back somewhere else further from the head. But this is not a 3350 assembly because I don't recognize these suspensions. I don't believe these were ever produced so it must be some sort of prototype. I don't know if you recognize them or not. I got that from Denis Mee again.

Warner: Yes. This absolutely is a prototype and if you looked at these you'll see that it has a cam assembly on one side only and not on the other side. Well it turns out dynamically that that isn't any good. If you don't have a balanced assembly then as it seeks then you're getting torsion on the part.

Frank: So that's a fairly unique sample I think. Make sure Denis hangs onto that.

Solyst: Oh yeah, that's…

Warner: A very early one.

Solyst: That's an early one, yeah, a very early one.

Warner: Yeah.

Frank: And 3350 came out in 1976. I think it was an even bigger product than the 3340. It was huge. It was about three megabits per square inch as opposed to today we're at a hundred and thirty-three gigabits per square inch so we've come a long way since then.

Warner: Amazing.

Frank: Now beyond this, I mean obviously ferrite heads lasted a lot longer. This came out in 1976, as I mentioned, and in 1979 IBM came out with the first thin film head drive, the 3370. And of course there was a lot of consternation that it had taken us almost until 1976 to be in production with these things and start to make any money and all of a sudden here was a major change in technology and of course we and the other independent suppliers went chasing that. But it turns out that that was a difficult move to follow, particularly all the semiconductor style processing technology. And in reality, most of the rest of the industry continued to use ferrite heads for the bulk of their requirements for many years and, in fact, Applied Magnetics and I think the other independent suppliers were still making ferrite heads well into the 1990s. I can remember it must have been about 1993 we had this -- finally made the call to stop making ferrite heads and put all the effort into continuing to build thin film head volumes. And I think for several years after that that was widely regarded as a mistake. But between 1976 and 1992, the reason we were able to keep ferrite heads going is that a ton of innovations were made.

Warner: Yes.

Frank: Now the heads were made smaller. First we stepped them down to roughly seventy percent in each direction, what we called the mini monolithic head. We changed the materials from nickel zinc ferrite to manganese zinc ferrite which had higher frequency permeability characteristics, because the manganese zinc is largely nonconductive where the nickel zinc is conductive. And even in nickel zinc you would get eddy currents. A big kicker was the development of metal in gap heads where rather than make it just ferrite and a glass gap, we put thin film type metal layer on each side of the glass gap to increase the permeability, to increase the field strength. Those were probably the most profitable products in the history of Applied Magnetics would be my guess. Maybe rivaling later we made a thin film metal andin gap head if you will with permaloy, but a higher permeability metal lining both sides of the gap using a similar strategy. I can remember that when -- this is actually a 3380. I don't have a sample of a 3370, but when it came out of course we were -- a bit of consternation about the size of that slider, it was so much smaller. The day we got one of these I sent it to the lab to be measured and then I sat down with a piece of graph paper and designed a ferrite head, a monolithic ferrite head with the same footprint and a little bit taller in this direction to accommodate the winding window and the core. That became the micro monolithic head and in ferrite and later MIG versions that was one of the things that, you know, carried the company and a lot of the industry, as I said, well into the 1990s. One of our biggest customers for that turned out to be IBM for the drives made in Rochester, the smaller drives. I don't remember whether they were -- maybe started out being eight inch drives and went to five and a quarter inch drives, but we sold those heads to IBM Rochester well into the 1990s. Another innovation we made along the way which doesn't show up in any sample here, but kind of harkens back to the 3330 style head is we made what were called composite heads where we took an in that case calcium type MIG ceramic housing, ground a slot in it, and glass bonded a ferrite core into the slot. That cuts down the amount of core material, gives you a different material that you can use for an air bearing. We sold a lot of those heads, particularly to Digital Equipment. They liked the composite head. It allowed them to put the gap on one side in one rail rather than in the center as with the monolithic heads. Those went through metal in gap versions. Data General bought a lot of those as well. And as I said, we were still doing quite well in ferrite heads up until roughly 1993 when thin film heads had been in production for nearly ten years.

Warner: I'd like to comment a little bit on a couple of things. With the Winchester slider, with three rails, as you fly over a disk, if the disk is curved one way or the other as it's flying over, the air bearings on the two outer rails are supporting the part. So let's say the disk is curved away, then that means the center rail to disk spacing would change. So when we first worked on this design we did a lot of work at looking at the curvatures of the disks that were being made or going into packs to make sure that that variation in flying height could be accommodated. The taper flat bearing is interesting because the trailing edge stays virtually constant and the front is that part that goes up and down so there's very little spacing. But we did have the spacing change as you're flying over the disk, but we did have this spacing due to disk curvature. When you go to a two rail head then the support structure is there where the core is going. So when we had thin film heads..

Frank: Yeah. This is a mockup of a thin film head and you could see the core is on the outer rail on this one.

Warner: Yes.

Frank: This is a 3370.

Warner: So you actually gained some head -- you improved the tolerance on head to disk spacing by having it on one rail. Another kind of a maybe not thought of, but when you lay out a disk, the bit density gets higher as you go in if you write in the same frequency. So if you have a tri-rail head and you're trying to use the whole disk area, the head's in the middle where if you can move the element to the outside you actually move a little further out on the disk and so you..

Frank: It doesn't stress the linear density.

Warner: So it doesn't stress the linear density. So having the element on a rail, one rail or the other and placing that rail on the outboard side helped in the magnetic recording process. And so I think that your composite heads..

Frank: Composite heads did that. It allowed them to put the rail on the outside so we could keep it a little further out of the landing zone.

Warner: Yeah. And I think you guys learned how to put cores in and it was probably cheaper than trying to figure out how to make thin film heads.

Frank: Yeah, we were actually doing both, but we were having much more success at the ferrite head approaches for many years in fact until the mid '90s and then we finally hit our stride in thin film heads. And the biggest product there was this metal in gap like thin film head which we sold a lot of to people that were trying to compete against the early MR heads and didn't have availability of MR heads.

Croll: You know, I must say in this particular head the elements were actually made on both rails.

Frank: Yeah, indeed.

Croll: And you just chose one.

Warner: You chose which one or the other, right.

Frank: Yeah. Well, it gave you a little bit of -- it was called the Gemini, right, twins, there was one on each rail and you could pick the one that worked better.

Croll: And a little bit of history too just from my experience, in terms of the development efforts of film heads which I was managing, the original objective early on and a presumed advantage of the film heads was to be able to make separate read and write elements in a piggyback fashion, so you could get optimum gaps in both.

Warner: Write wide and read narrow.

Croll: Write wide and read narrow. The only trouble is doing two of them at one time from a process point of view looked to be pretty horrific.

Frank: Yeah. We actually made a bunch of those.

Croll: Yeah, well we did too, but the yield was not great.

Frank: Well, the other problem was they tended- since they were both inductive there was feed through between the two cores and we couldn't figure out how to get rid of that when you try to read.

Croll: So that ended up with the single element head, but reduced, you know, the performance characteristics over the original objectives. So actually the film heads and, as you say, the improved ferrite heads coexisted for over ten years.

Frank: For many years, yeah. More than a dozen years.

Croll: And as a little aside, as the manager of that program, I kept getting visitations from our corporate headquarters in Armonk saying, "Why are you fooling around with film heads when they have all these good things?" But the thing that really I think did the ferrite heads in was the magnetoresistive ones which definitely got the advantages of the two.

Frank: Yeah, MR heads.

Warner: Well I think early, you know, like the 3380 product could have been built with a ferrite head.

Croll: Yes, definitely.

Frank: Oh yeah.

Warner: And it was done by other people.

Croll: No, it definitely could have been.

Frank: Definitely.

Warner: So if we kind of look at the timeframe, ferrite heads, at least in hard disk drives, started out in..

Solyst: ’73.

Warner: ’73?

Frank: Well, if you include the ferrite cores, 1966.

Warner: This would have been…

Solyst: Oh, yeah.

Frank: Yeah, I mean if you include ferrite cores, they started in 1966. This came out, what in about 1973, so I imagine this predated it by a little bit, early ‘70s.

Warner: Yeah, so I think- yeah, this was done in the late ‘60s, right, built for the Zeus program.

Solyst: Well, it was actually shipped at the same time as the 3330.

Frank: Okay, so ’71.

Solyst: ’71, yeah.

Warner: Okay, and then they were in production. You guys made them long after IBM did so you were in production with ferrite heads until when?

Frank: I think 1993, as I recall.

Croll: It was a good twenty-five year run, I think.

Frank: Yeah, I would say, yeah.

Warner: In terms of manufacturing and development, somewhere in that range. And I don’t think we’re ever gonna see ferrites in hard disk drives again.

Frank: Probably not, but there are probably a few of ‘em still out there.

Warner: Oh yeah. Does anybody have an estimate of how many were made?

Frank: Oh, not a clue, not a clue.

Warner: There’s got to be, you know, tens of millions of them.

Frank: Oh easily, easily, yeah. Yeah, you know, the volumes we got up to certainly weren’t like thin film head volumes today which are, you know, millions a day, but, you know, I think there were probably years when we made, well several hundred million heads I would guess probably.

Solyst: Yeah, it’s got to be at least that.

Warner: And I remember that, Erik, when you and I don’t know who else, Ken Mishado and others had to go -- I think the projected number for Winchester heads was gonna be like seven hundred thousand heads.

Solyst: I don’t recall that number.

Warner: And they were thinking about should we proceed on with this technology? Can we afford to do that? And so the projected numbers were so low at the time. And obviously, you know, those were way, way understated values. It would have been a shame if it had all been killed because of that.

Frank: Yeah. Let me mention one other thing related to this that for a couple of old air bearing designers I think it's kind of near and dear to our heart. This is a 3370 mockup, but all of the heads up till then were used on linear actuators. They sort of went in and out of the disk like this. But the small disk drives, the innovation that came out there was a rotary actuator and so this had to be taken and turned ninety degrees and the technology used on all of this was robust enough that that worked and it worked well. It led to a host of air bearing innovations with side pressure contours and negative pressure loads and things that are still used to this day that are very innovative, but it worked. I mean we turned it sideways and it went like this to the inside and the outside, caused some problems with skew and things like that, but we were able to design certainly air bearings that flew fairly level all the way in.

Warner: Well I think that these -- the base technology of the taper flat bearing's narrow, you know, ski-like, as you say, has turned out to be extremely robust. And now I guess we do exotic shaped bearings with self-loading..

Frank: Self-loading, various levels of air cavities to create pressure balances and things like that. It's all done today, but it's all fundamentally based on that technology when all is said and done and fundamentally based on this technology. Even thin film, you know, and MR and GMR all utilized things that came out in that.

Croll: And the same bearings.

Warner: Just kind of another historical note, I think we mentioned that original cores were made in Poughkeepsie, New York, but they transferred the whole group out, I guess under -- Miles Cook had the machining portion of that, brought his capabilities out. But what started…

Solyst: And also the material.

Warner: The materials, yes.

Solyst: The material manufacturing was done there.

Warner: Was transferred out so that actually on our factory floor were the ovens, you know, for firing the ferrite material and the mixing of the powders and so forth. But did the ferrite process -- was that an outgrowth of the core memories? Those were made in Poughkeepsie.

Solyst: Yeah. Well, I think so, but I'm not really sure. But I do think that, you know, by the time that ferrite came out, magnetic cores had been out of business for a number of years and I think perhaps they kind of oriented that capability towards ferrite. Do you know that?

Croll: Historically, it's a question of skills because the core memories had -- to do it there were many ceramists and chemists that were involved. And when semiconductor memories came in that was all obsolete and so the skills came out to address the ferrite heads. It's sort of interesting that the same thing happened really with film heads because there was a large effort in the East in developing film memories to replace cores, core memory, and then semiconductor memories won out. So the skills that had been applied to the film memories were then addressed on film heads. So in some ways these were what, jobs programs that were -- I mean very often you think of a project and you have to go get the skills, but in these cases there were skills which could be applied to the projects.

Warner: Yeah, there were guys like Duane Secrist and Hal Turk and Walt…

Solyst: Nystrom.

Warner: Nystrom.

Frank: Dave Thompson. Dave was working on film memories.

Warner: He was doing that back east, but these were the guys who came out to do ferrite; Miles Cook, all came out to do that work.

Frank: Actually, I think in reality that kind of happened twice because when the film memories died, some of the people went to film heads, some went to bubble memories. And later when that died, a lot of those people came around back to the film and MR head business.

Croll: Yes. Well, thank you very much. I think this has been an interesting session. I've learned a great deal and I hope it'll be interesting to others. Thank you.

Frank: Hope so.

Warner: Hope so, yeah. Thank you, Ian.

END OF INTERVIEW

