Oral History Panel on IBM 3340 and 3350 Disk Drives

Interviewed by:
Jim Porter

Recorded: April 6, 2004
Mountain View, California

CHM Reference number: X2718.2004

Abstract: *IBM's 3340 disk drive and 3348 removable data module which first shipped in 1973 were developed with the project code name "Winchester". The 3340 was first to use "low mass heads, closed environment and lubricated disks", a concept which came to be known as Winchester technology and was copied in all subsequent disk drive designs. Panel members devised the Winchester method over an extended development period, and also applied it to the next generation disk drive, the IBM 3350, which shipped in 1976 and was widely used in mainframe applications.*

James Porter (JP): We're here today to talk about two products introduced by IBM in the 1970s, starting with the 3340 disk drive and the 3348 data module, a sample of which you see on the desk right in front of us, and then we'll go on to discuss the 3350 which was introduced later in the decade. But before we get started in why these things were done, how they were done, and their impact on the industry, let's talk about who each of these people are here at the table, all of them who were with IBM during that era, who made important contributions to these products; starting with the gentleman to my left, Ken Haughton.

Ken Haughton (KH): I'm Ken Haughton; I joined IBM in 1957 after having been with a couple of universities and General Electric, and joined IBM in San Jose.

Jack Harker (JH): I'm Jack Harker; I came to IBM in 1952 when they first started a lab in San Jose, straight out of Berkeley, where I had a masters degree, and I stayed with the disk file business from the beginning until I retired in '87.

Chris Coolures (CC): I'm Chris Coolures and I graduated from Berkeley also with my masters degree in mechanical engineering, and joined IBM in 1960 and worked on disk files and printers for 32 years, and retired.

Mike Warner(MW): I'm Mike Warner; I joined IBM in 1965. I graduated in mechanical engineering from San Jose State University, and I joined as a junior quality engineer working on 2311 manufacturing areas.

Bob Friesen(BF): I'm Bob Friesen; I joined IBM in 1964. I was a graduate of the University of Colorado in electrical engineering. I worked for IBM for 32 years almost exclusively in all kinds of different storage, and now retired.

JP: Well, let's talk about why this happened, and what it was all about. But first, let's get one thing out of the way: the 3340 and the 3348 had a code name. All IBM projects apparently had to have code names, and this one was known as the "Winchester" by its code name within IBM, and when it was introduced, that Winchester name became one of the most frequently used names in the computer industry for decades and decades afterward. So, my first question would be with the gentleman who managed that project, Ken Haughton. Ken, why Winchester?

KH: Well, it's an interesting question in the sense that we were defining the product, trying to end up with what would best fit the market requirements for the low end systems, and I went off on vacation, and when I came back, why Chris and one of the marketeers, I guess – I've forgotten, was that Otis?

CC: Yeah, Otis.

KH: Otis Page had been down negotiating with a small systems house, and they came back with a definition that said we would make a device that had two spindles with 30 megabytes on each, and so they were calling it a 30-30, and to me, the 30-30 was a Winchester. So that's how it got its name.

CC: In fact, we used to hunt together, and I had a 30-30 rifle.

JP: So it became codenamed Winchester?

KH: Uh huh.

JP: But gentlemen, why was this project started in the first place? What was the reason for it?

JH: Let me take it back a little ways. In the mid-60s, IBM had pretty much had the whole disk file business to itself and it was very profitable, and found we were inviting competition. We also had a number of events like the dirty dozen leaving IBM. And the company was persuaded to invest in the technology. And so we set up a technology group with an advanced development function, and we had made a lot of progress in things we thought were significant improvements over extensions of the product lines that were presently being done. And we had everything but a product. And so, when we saw the need, and there was a competitor proposal – and Ken can talk about that – just to take a few disks out of a high end product, but we thought there was a better way. And so I think it was in '69, Eisenhower's funeral day, the plant closed, and a small group of us got together and we defined that we would like – we'll come up with our proposal and start appropriate. And Ken, take it from there.

KH: Okay. The need was basically for low end. IBM had recently announced the product, which was the 3330 – which, by the way, also had a code name which was "Merlin"; anybody had to be a magician to make it work! – but it was a high-end device. In those days, 100 megabytes was a lot of data, and so the small systems needed a smaller device, and at the same time we had done some work with low load, low flying heads, and we felt that combination should work out very, very well. In hindsight, it did. We have – we spent, as Jack just said, quite a bit of time putting together a proposal trying to convince the users and their corporation as a whole that this was the right way to go as opposed to stripping down a version of the Merlin to solve that problem. So we initiated the program with that in mind.

JP: Well, one question that would come up is – and it may be out of order, but we'll do a lot of things out of order, I'm sure, as we discuss this today – for a period of time, for at least 10 years before that, the industry had been using removable disk packs, starting with the 1311 in 1963, so at the time this program was introduced in '73, we went a decade. But the disk pack with separate disks and heads, of course, was exposed to particulate contamination, if that were a problem, from airborne particles, etc. The Winchester introduced the first combination of what the industry came to know as low mass heads, lubricated disk, enclosed environment, so the outside world was kept away from those heads and disks.

KH: Actually, you know, there were two major reasons for doing it; probably these fellows will think of some others – one of them was the contamination issue that you mentioned. We did a closed circuit system to filter the air that went around and around and cleaned things out, but one of the really strong driving things I felt was that if you could get the heads uniquely with one disk, ones in disk set, that you would eliminate a whole lot of tolerances involved in the removeability issue. And once you get the heads so that they always read data that they wrote, and didn't have to read data that some other head wrote, why it eliminated an enormous amount of tolerances –

JH: There was another reason. We could see if you looked at the sales figures – I mean, 1311, 2311, you sold lots of disk packs with each drive, because they had a very small capacity. The 2314, which had a larger capacity, you sold a lot fewer disk packs per drive, and it became clear that, as the capacity of the drive went up, and I think we were beginning to understand the impact of the operating systems where, once a system was up and operating, you didn't want people changing things. So, we referred to it as a graceful degradation of the disk pack.

CC: The other things too, we had the system houses for the intermediate to small systems that wanted a very low entry and they couldn't really decide on exactly what they wanted; in the early stages we designed the data module that held five disks maximum, and so we could have a configuration of either one disk to five, depending on what the system houses wanted. You know, we started out with that variety, then it went to the 30-30 for a while, and then we ended up with a 35 and a 70-megabyte, which was two disks and four disks. So we tried to keep that flexibility, and it was easier to do in this kind of a configuration, because you just always kept the same number of arms because you put dummy arms in it, so dynamically with a server, then the problem's been changed. So it really had a lot of advantages.

JP: So, is it fair to say that with the development of this product, which was intended for smaller than the large mainframe in terms of application, and in terms of the higher recording densities which we were gonna be able to use, the need for the save and restore capabilities of the disk pack to save a function, then to take it off then restore it later and use multiple packs, that need was reduced because of the higher capacities and because of going to smaller applications – those two combined?

BF: Although they didn't take advantage of that in this particular timeframe because these were smaller systems and they still continued to use the data module in the same fashion they would have used a small disk frame.

JP: Except that, of course, the data module cost more than an individual disk pack, didn't it, because it was a more complicated device with heads included?

KH: Yeah, that's true, but the cost of the drive –- you saved money in the drive as well, so the cost of a drive and a half a dozen disk packs or data modules was cheaper than the previous drive and disk packs.

JP: I guess I should mention that I was the only person here that didn't work for IBM and for that, during the period of the disk pack, I worked for a company called Memorex and managed the product management on their disk pack product line. So I was watching that market very closely up through that period, and was delighted to sell those additional disk packs that people needed. And, of course, I got out of that business and left Memorex before this came out, so I didn't participate in being cut off at the pass with this product.

JH: But we made trade-offs to minimize the drive cost, as Ken said. And there was their two heads per arm, and that meant the voice coil could be much shorter, and if you looked at the equivalent 3330, the voice coil was massive, and we were able to save drive money by putting a little more in that. And we also had a lower cost head.

JP: Okay, but going to the question of what were the technology challenges in doing this? What were the most important challenges in deciding to do this product, and what had to be overcome?

CC: Well, it would start with the heads.

MW: I remember seeing Chris Coolures the first time when he made a presentation. Dr. Haughton introduced the program – I was in the technology shop, and Chris went through all the things that needed to be done, and the challenge we had from the technology side was, because the heads were gonna be residing with the disk, they had to be very inexpensive. Our target was to make a head suspension assembly for less than a dollar. And at that time, a head suspension assembly was then more like $15 to $18. So we had a dramatic cost reduction that had to take place. Another area is that you couldn't afford to have all the mechanism for loading and unloading heads, so we needed to do start and stop on data, I think, as Ken mentioned, with light load lubricated. It was gonna be the highest aerial density by quite a bit. So we had things like oriented particles in the media, and I'm sure you'll go through the server problems that we went through. So there were a whole lot of technology challenges, and as well as I know Chris from the module side that you had a lot of –

CC: Well, we had, you know, I mean, you have the problems with any development, but we had – well, you know, it's obvious the compromise with the technology and the machine group and how you allocate the tolerances.

MW: But share the poverty!

CC: Right, but having the heads made to the disk and, you know, writing and reading back with the same head added a lot of advantages. Of course, we had problems with the design of the actuator itself, because you have to couple the coil to the carriage, and there was the technical challenges there. And when you did that, you didn't want the coil to be rubbing in the gap of the magnets. Those were all challenges, but with a good team we got through them. I think we had some problems like, the server would lose control and we hit the crash box and bang…

MW: …the head and the disk…and because the head had some short tapers, we would knock really divots out of the disk and it crashed. You know, we had to really work hard to fix that problem, but we did, and it was those kind of things that you get.

JP: Well, let's take a few of those subjects, like the heads and the stories about the development of the heads for this involves numerous entities in the industry, and a number of thresholds that had to be crossed. How did all that happen?

MW: Since I'm the head guy, let me start off. I think the vision we started off with, with this low mass head that had been used on some other programs before it –- I think…

JH: Yes, there was a buffer that we built for Rand Corporation. This model went back to the origin, like, the first person that we saw try this was a fellow at Data Disk, R.J. Miller, had a tri-pad head. It had three points that would fly, but with very low load. He thought he was writing a contact. We did some work and found out, no, it was about a micron or so away from the disk, and we used that design and it worked, but it was not either a dynamically stable design or a manufacturable design.

KH: And it would continue that because the server loop around it…

JP: Well, I know that IBM bought the rights to the design from Data Disk, but then went on to pass the tri-pad into some interesting refinements.

MW: Yeah, let me – the problem with the tri-pad –

BF: Try this head (pointing to sample).

MW: Yeah, I'll grab that in just a second. –- with the tri-pad is you have two bearings in front, a bearing in the back that had the element. It was, I think, a barium titanate or ceramic…

BF: Yes, barium titanate.

MW: Barium titanate, but the element was glassed in separate – the ferrite element was glassed in, and it had a taper that was on this inside surface, and tapers in the front. Those turned out to be real bearers to manufacture, and if you've got to manufacture this, as Bob was about to say, when you try to push on the head, one side is much heavier than the other, and so dynamically it wanted to rotate and wobble instead of going in a straight line. In this technology area, we had a team of guys that had been working on it, our Zeus or 2305 program.

JP: We should mention that 2305 came out about the same time as the 3330 but had fixed heads.

MW: Had fixed heads. As a matter of fact, I think we've got a sample here. This is a little module that has a taper flat bearing on it, has nine elements.

JP: Excuse me, but why don't we break, and why don't you do a close-up on that? Go right ahead with your discussion.

MW: Okay. It's a taper head, a taper flat design; has nine elements. It was pneumatically loaded with a couple of kilograms with a force against the disk. And this had been set up in manufacturing, and it had some very interesting processors. A guy by the name of Eric Solast (sp?) designed the process in this head, A guy named Elliott Flukow (sp?) did the tapered bearing I think with a program that Tom Tolang [sp?] was the author of the software for doing that. Quite a good program, and the interesting thing from a manufacturing standpoint is, we had ferrite cores made in Poughkeepsie, and they also made the ferrite head elements that were cores, since the ferrite cores from memory were going out of business, they needed a place to house these guys that were making the ferrite for the magnetic heads. So Miles Cook and his gang, Hal Turk, Dwayne Seekerest [sp?] and a few other guys came out to San Jose and set up the ferrite business. We had a lab, and so we could make these bricks of ferrite, whatever size we wanted. So this batch fabricated Zeus program had a couple of new things. One, we developed machines that do slicing, so initially they were laser controlled, and then we used glass lights to control, so we could do very accurate slicing. And also inside of a head, there is an apex for the throat height, so there's an edge of ferrite, there's a gap, and we developed a -- remember the sputtered film shims that went in the gaps -- and then there's the top of the throat is a -- we call it an apex, and that had to be very finely ground at 45 degrees, so we had a cup wheel process that we could index across, all generated for making this device. And when we started working on the Winchester, we said, "Gee, here's a team of guys who had been working, I think, myself and Eric and George Powell went up to Menlo Park, ACS location, and joined a bunch of guys from the ACS and brought the people from manufacturing, and I think Jack had a proclamation or something…

JH: The "We" program.

MW: The "We" program.

JH: Background on that. I'd been in -- we had a laboratory up in Menlo Park to develop a supercomputer, sort of, and the program was well funded, had a lot of people, and went belly-up, and San Jose was given the job of absorbing the people. And one of the groups up there was a group that had been doing some work on semi-conductor packaging. And so we just took that group and melded it with the expertise in the heads, and put together with manufacturing what we called the "We" program, which meant they started in development and then when it went to manufacturing, everybody went with it.

MW: That's right. For half the program, I worked for Rex Sellars and the other half I worked for Bob Howard, and on the manufacturing side. Well, anyway, we were making these tri-pad heads, not having very good success with them, and so we started working on using some of the ideas we had for the 2305 or Zeus head. Now, you can see this took a couple of kilograms and we needed to get that down to something like 10 grams or so. So we started playing around with all kinds of varying designs. That was my job, and at the time we rammed out I think it's a model 91 that was in research. I think I had 1 meg of memory – and the programming was done with decks of cards.

BF: These are simulation programming?

MW: These are simulation programs for doing their varying analysis. So I got -- it was gonna take us forever to get this analysis done, but I talked some people at the computer lab and research and said, "If I was there at night from, like, 2 to 4 in the morning, they would let me run my programs." So I would stack -- you know, key up a whole bunch of decks, put them in, and then plot the data that came up, put it – key up some more and put those in, and try to figure out an air-bearing design that would work.

JP: All that between 2 and 4 in the morning?

MW: That's it, only because I could get –

JP: You wanted to tell them you could get up.

MW: I could get batch runs, because I could just go around and around and around. And through that work – obviously, you say, well, one of the things you want to do if you have a big flat surface that will take all this space is you start making it smaller, and cutting grooves in it and doing things like that. Well, through that work, we figured out that, if you make these long narrow bearings, they were kind of like skis, and so the taper flap would build up pressure in the front and then it would bleed off as it got closer to the back, and then you'd have another high pressure. So it actually came up with a double peak on the bearing. And the interesting thing is that, when the head would go forward, pressure would build up here and pull it back. If it started to tip the other way, it would reduce the pressure, and so the pivot point was around the gap. Now, so that was getting the varying part fixed. And then the idea of doing a center track with this cup wheel, we could do 45 degrees either side; we could form the bearing rails, and then the core could be machined to whatever width that we needed for the program. And then it had another rail. Now, we thought this would be a – we did our cost analysis, we did the process analysis, it looked like it would do our job well for us. When we talked with the machine group, with Chris Coolures, Chris said, "You have three surfaces and the disk has curvature, and isn't that gonna do is change the spacing in the middle or, in the worst case, to do – wear it away." So we took hundreds and hundreds and hundreds of traces on disks in order to do is satisfy Chris and give him a database that we could – he'd allow us to use this tri-ped in.

KH: There's a little side story right about this time, my friend Mike, and that is that we discovered that the 30-30 wasn't big enough to be competitive, and we decided that we needed to have more capacity in a data module than 30 megabytes. So I had a meeting and called in the guys from the disk shop and the guys from the head shop and the guys from the machine shop and the guys from – those circuit data recovery guys, and said, "We're gonna have to increase the recording density, and I want you guys to go think about it and come back and tell me what we can do." So they all went off and they came back, and the head guy sits down and says, "Well, I can make that a lot narrower, but I don't think they can make a disk that will work with it." And the disk guy says, "Well, no, I can make a disk that will do very well, but I don't think they can do the read-write circuits" and so on around the table with everybody – and so out of the air I picked a number, 300 tracks per inch –

BF: And I was sitting there at the time.

KH: And everybody gasped and they went out and started working on it, and I've kicked myself ever since for not saying 500!

BF: You probably would have got it!

JP: And I think we should mention that the final drive, when it was finally introduced, did have 300 tracks per inch, and that was a final spec. So this became known as the tri-rail?

MW: This was a tri-rail head assembly. Now, interestingly enough, when you -- I mean, that's what's doing the air bearing and forming the read-write portion of it, and on the Zeus program we worked out how we could control the throat heights, and Eric had developed a very -- Eric Solast -- had developed a very interesting technique where we used what's called a stop ring, so we could measure the throat heights on all these heads, set the ring to where we want to lap, put it on the lapping machine, and it would lap down exactly where we wanted it, within a fraction of a mil, and then it also had a sign plate which would then tilt and then lap the taper on it. So it was a very cost effective way to get a very accurate part. Then the next challenge for this thing was, okay, we've got a head, we have to attach a suspension to it. So, again, suspension needed to be close to the center of gravity, which this is – it has a notch in the back that the suspension fits it, has a dimple on it and a load beam. And a guy by the name of Dick Wilkinson actually did the work, the design up of the suspension assembly, and we would go build them -- we could turn around a suspension design in just a couple of days. They were etched over -- where was that etching was done? I don't know -- we had an etching facility, anyway. We could lay it out, have them etched, form these parts, and then we would go over to Ray AbuZayyad's lab – they had a glass disk with a shaker taper on it, and we'd mount the suspensions, sweep frequencies, put a strobe light on it, and you could see the resonances, the suspensions dancing away, and so forth. And Bob Freeson and his guys would say, "Okay, you've got to keep that resonant away from, what, 1100 hertz…"

BF: 1100 hertz was the bandwidth of the servo, that we had to…

CC: Yes, you'd want to be near that now!

MW: So we had to do is basically tailor through experimentation and some finite element modeling, but actually experimentation is really hard…

BF: Primarily experimentation.

MW: Yeah, that we got a suspension assembly that would be compatible with the…

JH: It had to be very rigid this way, and very soft.

MW: Yes, it had to be extremely soft so it could gimbal but be very stiff as you pushed on it…

BF: In the direction of motion…

MW: Have the direction measured.

BF: The actuator was an interesting thing. When you actually look back on this, the mechanical complexity of this device is something that would make you wonder how we ever got it done.

JP: I was wondering as you were talking!

BF: To a servo guy when you hand him a carriage, a device you're gonna push, a moving mass, and the first day is gonna come in two parts, and is gonna be hooked together with a bunch of springs, and rolling around on a rail, you say, "Well, maybe this will finally get <inaudible> but we did make it work. The servo was all confined to single transducer, namely a head. It was a head that was mounted in the middle of the arm, a single one, on the bottom of those disks so it always saw its own data as well, and all of the inputs for the servo were taken from that disk. So there was no external transducer like a velocity transducer or anything like that. All the position information was taken by reading each track, and as you crossed it, we developed a signal that would allow us to generate both position and velocity and thereby control the servo. We used a different pattern for encoding that servo data than had been used in the past. Die bits had been the Merlin 3330 form, and we used a new form called a tri-bit which was three pulses that allowed us a little more resolution and more integrity in the track crossing ability to count. The servo itself, as like we've alluded to, the very first ones that we closed around these tri-pad heads, they turned to dust in a hurry. They used a servo, it actually wasn't optimized and it started resonating and the head went and then we got that fixed, and then we tried to push on them and they broke, so we were kind of in the dark for a while until we got the tri-rail head. Once we had that, things started to move very well. We had the usual problems of trying to extend the bandwidth of the servo to where it needed to be, but after we got that going, I guess we did have a few problems with the motor itself. We had to linearize it. And one of the bigger problems we had, we had to develop the servo writer as well that created the tracks on the bottom disk, and that was an extension of what had been done for 3330, but because of the way these heads read the track, we found that they varied a lot in the way – because we used the same kind of head to write the servo patterns – as a different head would be put into the module, it would read side rail signals in a different way each time we put one into the pack. And so one of the inventions that came out of this was a simple resistor that we plugged into the data module that equalized all of the different head characteristics. That took us a lot of time to get through that. Jack Cheng (sp?) was the one that got the patent down there. But after we got through those kinds of things, things moved pretty smoothly.

JH: And now you've got to talk some about -- along the way there were some problems.

BF: Oh! Yeah, there were, like the suspension.

JH: Memorial Day massacre.

BF: Oh, yeah, well, that comes later, that comes more of a 3350 problem. I don't think we ever really recognized that we had those problems here. We'll get to that, Jack, soon.

JP: What kind of problems were there in getting this into production, this head assembly?

MW: Actually, my recollection is that it went quite smooth. The design of the head suspension was such that we could test each of the heads – we could do a finite test and an electrical test on a head assembly before they were mounted to the arms. If you've looked at these guys closely, you'll see that this is actually swedged in place, so there's a stainless steel mounting block that the suspension is welded to, and that goes into an aluminum arm, and if you put a tool through or a ball through or expanding – what you're using the different elasticities of the material, so the stainless steel expands out into the aluminum, but the aluminum, when the tool comes up, the aluminum comes back and grabs it. So, it turned out to be a way that it would not change the flying height or the load on the head by mounting. If we tried screws and things like that, it was built more expensive, but it also would distort this plate. And so George Powell came up with this wedging technique for mounting the arm.

JH: Part of the answer to your question is because the group that was put together brought with them the manufacturing approach to the tooling of the 2305 head, and that was always - the design was based on what can we do with the processes so when it moved to manufacturing, it wasn't a new game.

CC: Yeah, and the process was the one used in development.

JH: Yeah, there was no pass off.

CC: It wasn't like, you know, in the machine program sometimes, we developed the product, built in the laboratory, go through the first test, released our drawings to manufacturing and let them try to do it. Well, with the head, they didn't do it that way.

JP: And the development and manufacturing were all in the same group?

CC: Except that they were all part of the same…and just the management structure on the top just changed, but not the people down below that were really doing the work.

CC: And actually, as from a program point of view, the biggest challenge once you got the head working was to convince the head area that the suspension had to be changed. I don't know if you remember the few discussions we had on that, but once that started, and you iterated on the suspensions, then after that we were in good shape.

JP: These heads have to fly over a disk. Now the disk had one interesting change I'm aware of -- the oriented magnetic particle -- that was new with this, wasn't it?

BF: That was. It was introduced in this -- I verified that, Jack. It was, yeah, with a common disk, it was introduced. This device, when you look back at it, both the drive and the data module incorporated so many new things, I'm not sure we were aware of what we were doing in that context…

JH: Probably good!

BF: …which was probably good. I mean, the head technology was very new, the disk had been redesigned to provide us a common platform throughout all of our products, and the orientation was…

JH: Yeah, we came up with the idea that we could build a common disk and just test, because different products had different requirements, and so you'd take your yield from off the most critical and then you'd pass the disk down to a less and less…

BF: Uh huh, like the chip guys do today. Then all of the mechanics of the data module and drive coupling, then the attachment electronics was all different and new – Jack can talk about that. I remember one problem we had getting it into manufacturing. One day Chris called a meeting and said, "We've got to talk about a problem." We were getting ready to go to the final stage, and we'd had all these cost targets to meet, and this and that, and we thought we were doing well. And he comes in and he says, "I've got the cost estimate here, and it says - blah, blah – but the problem is we forgot to include the electronics! And so we got a little problem!" While that was a low-tech problem in a sense, it did keep us up for quite a while until we got it solved.

MW: Actually, speaking of electronics, one thing to mention is that this was the first program where we had then put the electronics out on the arm, so this had –

BF: We had a diode matrix and a pre-amp out there too.

MW: But at that point, pre-amp and diode matrix, and then later on we put more components on the next generation beyond that, so that was another challenge to – and you can see we're using IBM package devices, so we had to figure out ways to have it survive crashed ops and…

CC: Well, getting better flexibility, I mean, on the cable you used, because before we were so noise sensitive at whatever we did, and that really helped a lot.

BF: We haven't talked about the attachment strategy.

JH: I don't know if this is the right session. That's almost a separate item. During this time, and really this is really the result of the 3330. Migrating customers from the 2314 to the 3330 turned out to be a difficult problem, because you had eight spindles of about 25, 30 megabytes on the 2314; on the 3330 you had eight spindles of 100 megabytes, and so once a customer required a lot of storage in those days, they didn't want to make this transition, and it wasn't because you had to buy a controller, and then you add the drives to it. It really wasn't economic to buy a controller and one or two, or two or three drives, you really weren't saving any money. And so we looked at, how could we bring down the entry cost, and we did that by splitting the function of the controller. Some of the things that the controller was doing were device oriented; some of them were system oriented. So we sort of took a cleaver and we separated those. We had what we called a director which was in the low end systems packaged in the system, the high end systems it was either a separate or integrated box, and then just in front of a line of drives you have what we called a day box which just had the drive control function, meaning the clock and data handling, that would sort of associate into that particular technology. And that was part of what – that transition then reflected through our whole product line.

BF: And we were the first to get to do it!

JP: So with the 3340…

JH: Yeah, that was the first product that came out initially.

(Gap in recording tape)

JP: Well, let's go on to that question, which has come up on the disk. The disk, in addition to the oriented magnetic particles, introduced a few other things. What were they?

CC: The lubrication of the disk, which had its advantages, but also you had to worry about stiction which they were -- the head stuck to the disk, but -- that was one of the new things that was done. And the real challenge was to get the right amount of lube and also trying to understand what happens to that lube after a year or so of spinning. Does it migrate off the disk, or does it really stay where you put it?

MW: Well, actually, it required the development of a lube to <inaudible>

CC: Well, right.

MW: Which was a very, very thin cut on a molecular basis.

CC: That's right, yeah.

JH: But it took a ton of testing to get us all as comfortable with the life quality of the lube on the disk.

JP: I would have answered, because we've all had cars, which have run out of lubrication over time. What happens when one of these things -- two to three years in the future, because the chief financial officer that gets the proposal from a company information officer for a five-year amortization on a piece of hardware. I don't know if you'll raise the question, but some of the people involved might raise the question. And how do you test five years without going five years.

MW: Actually, we did a whole series of accelerated wear tests. And some of those would be, you know, going at slow RPM and having heads rub on the disks, so to speak. Another where we've just start-stop-start-stop-start-stop. During the whole program, there was kind of a -- well, there's dozens, but there were certainly multiple spindles that were running the whole time.

CC: Yeah. The telegrade we'd set up a lot, get about like 8 or 10 spindles. We'd have start -- we'd have to do so many start-stop cycles, and, of course, we had three phases of product test, and they would do independent testing of that legally, and then they'd…

MW: But the technology guys were doing tests.

CC: Right. Through environmental swings, and then see if the head stick to the disk, and so a lot of very extensive testing to make sure everything was all right. And…

MW: Even -- things were tested. Everything was tested, you know. Flux side tests were done on the cable assemblies. On the suspension. (cellphone bell rings) Oh, shoot. <laughter> Yeah, right. You didn't have to worry about that. The load was -- there's a little dimple with the load beam on it, so there's a little dimple sitting there to (inaudible) down. (cellphone bell rings) I thought I turned that off.

JP: But it has overpowering capability.

MW: Yeah. And so we set up a structure so that it would actually gimble this, so we'd load it and one of the machinists, Tom Kerr [sp?] worked up this oscillating thing so it would just gimble- wobble plate. And so to make sure that we didn't get frutting [sp?] corrosion. Flex fatigues. All the suspension assemblies were gone through to make sure there would be no flexing on it. And you guys had bearing tests going on.

CC: Oh, yeah, right <inaudible> Well, just about every component you had to do live testing on em.

MW: This lubricant remains a problem. Not just for this product, but for others. Generations of products afterwards, when we, as we got closer and closer, you found we had something, I think, Western Electric had it only- the only one of em discovered, which was the plumerization [sp?] of lubricant.

JP: Well, I know that as a change in the lubricant's…

CC: A form of a hard plastic.

MW: Yeah.

JP: You changed the lubricants you used from time to time, didn't you?

MW: Uh.. Only in the 3380, I think.

MW: Yeah. We- we looked at a couple of different lubricants during this time, but, anyway.

JP: OK.

CC: A kind of a trade off, but, because some lubricants showed less stiction, but in other ways it didn't perform as well, and that's always a debate of how -- which is the best lubricant to use.

BF: Oh, one of our suppliers quit on us, too, as I recall.

MW: Yeah, and we had to scramble and come up with a new one.

JP: Course, IBM was relying on themselves for head development, for disk, for probably most of the electronics...

CC: Yeah, all of the electronics.

JP: And all the manufacturing was all done in San Jose, wasn't it?

CC: Right.

MW: Right.

JP: But, one of the questions is just time. What was the total amount of time lapse from the time this program was started, until it went into production? How long did that take?

MW: Well, we started it, as Jack said earlier, in '69 conceptually. And we shipped the machine, which meant we'd made a few in manufacturing by then. I believe that was in November of '72.

JH: Three.

MW: Three?

JH: Shipped November '73.

MW: '73, right.

JH: Was the formal first customer ship.

JP: And when did you have a staff ready to go to work on it? I mean a full blown…

MW: We started in early June of ’69. That’s when…

CC: Yeah. No, not really June of ’69.

MW: Yeah, that’s when you gave me the job to…

CC: That’s right.

MW: Go to work on it.

CC: You’re right.

MW: Well I started- now.

CC: It was close to four-

MW: And then we started staffing, and probably, you came…

CC: I came at the end of ’70.

MW: ’70. And so by the end of ’69 we were pretty well staffed, I think.

CC: Yeah.

MW: Cause there was one other program that was assigned there was phased out, so we got a lot of resource in late ’69.

JP: So regardless, in ’69 this was still over a four years old program.

MW: That’s right.

CC: It was a four-year.

JP: Do you recall any other disk drive programs that went four years?

CC: For development? Yeah. I have. Always got a little <inaudible>

JH: We had some that got through that soon.

CC: But, typically, what I'll call a first generation product took like three to five years

CC**:** And the second generation, of course, comes, you know, a little quicker because you…

JP: So this was a mid-range development.

CC: This was a new generation product, and so it…

JP: But there were more new innovations in technology in this program than in most others, it sounds like.

CC: I think-

MW: Well, I don't know. There's…

CC: Well, no.

MW: Because the…

JH: The 1301 had a…

MW: This all new -- first year, crossed over the flying head, yeah.

<all talking>

MW: It was like -- it was like a self-acting bearing. The uh…

KH: One other thing I think we ought to mention in this program that was new and that I thought was, well, very clever, was the ability to skip defects on the disk.

JP: Um hm.

MW: Yeah, that's…

KH: We had a scheme where in the home address you could say, so many bites down stream there's a flaw on the disk, don't record there. And we skipped over them all the time. And this was a real blessing on the yield of disk. It did turn on occasions to be a problem, because if you didn't watch how fast those were building up, all of a sudden you'd be on the other side of how many you could stand. And you'd say, my God, we've forgotten how to make disks again. You probably had that experience at Memorex, too.

JP: We also had bad products from time to time. But as the overall manager of this project, what were you most worried about during this whole period?

KH: Well, all these technical problems…

JH: And then how to sell a disk pack.

KH: Well, that's right. Well, just to say all these technical problems really interest me and I'm an eternal optimist. I believe that we can do it. Whatever it is we set out to do. And it turned out that the biggest challenges I had with facing the other way with the marketing people, and as Jack said, how to sell a disk pack. It turns out that a disk pack was sold in IBM by the IRD. Auto- automation records division of the…

KH: Oh, one day a memo came across my desk that said from now on, disk packs, including the Winchester data module, which was before that announcement, are going to be sold by IRD instead of the data processing division.

JP: Oh, OK.

KH: And, holy Toledo, what a mess that's gonna be when we're just moved all the technology into the data module. I call that my million dollar mile memo. Cause I spent about a million miles on airplanes trying to find a solution to that particular problem. And we finally ended up with a compromise that both the DP salesman and the IRD salesman would get a commission when we sold a data module. But interestingly enough, shortly after we started shipping, I went up to Eugene, and went out with a salesman. First place we went, they had Merlins on order. And I looked at the system, and looked at what they were gonna do, and said, why are you ordering 3330s? And the guy says, I make more money. The salesman!

JP: You found the real reason.

KH: And I said, well, you mean, I spent a million miles on an airplane getting you the commission, and you don't even know it? And he said, you did? And we changed that order while I was in the shop that day, from one device to the other. And I encountered that I don't know how many times when I went out into the field. So it's a very powerful issue -- it was who's selling what.

JP: But nevertheless, when it was introduced, it was a successful product.

KH: Oh, yeah. Yeah.

JH: Well actually a 3340 wasn't that big a seller. It wasn't expected to be. It wasn't near the revenue-producer of the 3380 or 3330.

JP: Well, no. Those, of course, were pure mainframe products.

JH: That was the mainframe products, but as we've said, from time to time, if we thought of it in advance, it was pretty good strategy that we got a low production product comparatively, into production with the 3340, and all its new technology, and ramped up, getting good yields, and then we moved on to the 3350, which was a high-volume product, using the same technology.

JP: And I think now is a good time to discuss the 3350.

MW: OK.

JP: Well, I've followed everything from a…

MW: Well, I think there's another one and…

JP: The big study. Which I think you had a business reason for.

JH: Yeah, we were trying to get some extra revenue. And the systems people were always telling us that they needed a high performance, fixed head drive. Even low-end systems wanted it. So we thought about it, and we figured we could give it to em with what we call topalino. Go ahead.

MW: So, the- the topalino actually fits in the base of a data module, and it has a whole series of fixed heads. Now, here's a case where we took the tri-rail head, instead of having just an active center rail, we machined all three rails slightly narrower on the two outer ones, and wider on the middle one, so we then had three elements, and three coils on it. So each one of these head suspension assemblies on here, has three elements. And that fits into the base of this one.

CC: Yeah. And if it was on the bottom disk, and it was using the surface, cause the server only used half of a disk, and so we used the fixed heads on the other.

BF: The key challenge with that besides making all these heads work at once, is we had to have them stop talking to the server head.

JP: The server still in there?

MW: Well, that's true. Right. That's why there's little candle wrappers <inaudible>

JP: You said there's little cans underneath each one of those with magnetic shields,

MW: Yes. So when it's down flying that it doesn't talk to the servo head.

JP: So this was an option available on the 3340.

MW: Yes.

JP: Could you, do you have any idea…

MW: And it was profitable, although tropic [sp?] modules of...

JP: Do you have any idea what percentage of the 3340s were purchased with this option? Was it a big percentage?

JH: No, I don't recall.

JP: It was kind of special, I see.

JH: I think it was small.

MW: It was not mind boggling, I can…

KH: No. That's right.

JP: But profitable?

CC: For the people that wanted it was…

JP: Wanted performance.

JH: Wanted performance. And it gave them a…

MW: Fixed their file.

JP: Well it also enabled them to not bother getting a 2305, which was the-

MW: Oh, yeah.

JP: Fixed head- not <inaudible>

MW: Right.

JH: The cost of getting this feature

JP: Yes.

JH: At a very reasonable price.

JP: Compared to the large one.

JH: Yes, right.

JP: Compared to any sort of fixed head

BF: Yeah. Cause the 2305 was volumetrically probably twice as big as this whole machine.

JP: OK, So much for the 3340 for the moment, which introduced a lot of technology which then went into production on the next generation. A lot of that technology for the 3350. Let's stop now and let's change the props.

<crew talk>

JP: OK, on to the 3350, which uh.. moved up to 317.5 megabytes, all fixed, no removable data module. What you see on the table is the head disk assembly from a 3350. Of course, the entire drive was as big as your refrigerator. But this was one of the data modules. One head disk assembly from inside the 3350, as I said, was all fixed disks. It moved from 300 tracks-per-inch up to 478 tracks-per-inch, and moved up to a somewhat higher spin speed of 3600 rpm on the disk. But it really used a lot of the technology developed with a 3340, did it not?

JH: Yes.

MW: Yep.

JP: Well, any significant changes in that technology? Or just refinements?

CC: No, well one of the main changes as far as the actuator's concerned -- instead of having a two-piece actuator which bound the- <inaudible> the two-coil and the carriage as one unit, and left the magnet assembly and the gap as part of the drive, but then married the coil to the actuator -- we still had to close the loop air system that made it a little bit cleaner to have a good air system in there, because now we didn't have the mechanisms that had to connect the actuator and coil together.

BF: Yeah. Basically, you could take this actuator and all of the coupling mechanisms that went in the middle, came out. The carriage guide system and bearing system was all the same. We did change some of the cabling and electronics. The motor remained the same, it was attached directly to the carriage. The base castings were pretty much the same.

MW: Yeah. You added a few more arms to it.

BF: Yeah. We added more disks and more arms. But the arms were essentially the same. But we'd added some main actuator refinements by that time. I think the important thing is the fact that this technology was so durable, and we had built so much into it in the first place, that it was flexible to be able to change the head dimensions so that we could get the higher track density. The flying heights remained virtually the same, I guess. But the fact that we were spinning it faster caused different loading. The density changed, also to help increase the capacity. And all of those things were built into the original technology. We did have refinements in the read-write electronics, that actually made our detection systems more sensitive. The servo, again, was just an evolution of refinements. Same principles, so the technology was really there from the beginning to be able to extend dramatically for that matter into this kind of thing.

MW: Yeah, I was gonna say from a- <inaudible> if you notice on these arms, there's a bevel in the front, because higher rpm -- we were getting some…

CC: That's right. Yeah.

MW: We were getting some air buffeting, and so the front edge of the arms were tapered for that. Also, remember Jack Schwartz, and his…

CC: A-E module.

MW: Yeah. And his function, put more function in the modules, so there was more. And we went to an all-flexible poly unit -- flex cable as opposed to the ones with the springs that they were on. And still used the basic -- also in it, down at the base area you'll see there's a piece of the ferrite used as a shield for the servo head again. But this…

CC: Yeah, and I think that too, in this program, it was because the same team that did the 3340 both from a technology point of view, and a hardware point of view, was the same team, so that, you know, whenever you get new players, you gotta have a new design, right? And when you have the same team, you use some of that past, you know, history that moves a lot faster, I think. And, well…

JH: Also, the industry was ready for this.

MW: Yeah. Right.

JP: Well, it shared

JH: The experience of the 3340, it convinced people that the disk pack didn't have to be removable.

JP: Well, it shipped only…

MW: Actually it was the end game of the degradation.

JP: It shipped only three years after the 3340. It was shipped in '76.

MW: Yeah. And now…

JP: And of course the 3330 had gone to a second generation, the 3330 Mod 11, a few years after the original one, the double capacity, the 200 megabytes. But to go up to this 317.5 megabytes, in the big market of the day, which was for mainframes, was suddenly pointed at a much larger market than the other 3340 had been.

JH: Oh, yes. Yes. This was now the mainstream product.

JP: Yeah. And…

CC: And we had a version of this, the 3344 which would let you map the 3340 into a 3350.

MW: Yeah. Four of em, yeah.

JP: And apparently there was no resistance at all to the lack of a removable disk pack with data module at this generation.

CC: None.

BF: That was probably one of the two things that we spent more time on in development, subsequently in the market place. It wasn't a problem. But, a lot of our marketing people thought it was going to be, and so they put us through hell, to be quite honest, trying to convince them that (a) we needed to do it, and (b) it wasn't going to foul up the operation of a data center. And we were able to do that, fortunately. Although we did put in an awful lot of things in the attachment electronics in order to be able to have things like dual past all of the spindles, so that if one went down, you could get at it a different way. And because you couldn't take the pack off and move it easily, at least too easily, so we did all sorts of things in attachment electronics to be able to provide multiple paths to data, and that sort of thing, which affected some of the design of the machine.

JP: And that's an interesting observation for the people today, when they're used to RAID -- redundant array of independent disks -- which of course, has a way of recovering. When you use several disk spindles together you can recover the spindles that remain OK. You can recover in operation, and not loose data if one of those spindles fails for any reason. That capability did not exist during this generation, and not till later, so your point about being recoverable, and a way of diverting data, would be very important.

BF: Yeah, actually, oddly enough, this is kind of an aside, but several years after we introduced the 3350, Jack Schwartz, if you'll recall, did IBM's first RAID with these ties. So you can imagine this great big, long line of disk drives, 3350s, and he's scattering data across them in the RAID fashion, so it lived on a little.

JP: No predictions in the ramp-up at all, in production of this, I guess?

BF: Oh, well, Jack alluded earlier so you -- the other problem when we were…

JP: Yes. Now you want to talk about it.

(Inaudible general discussion)

MW: Ed didn't have a problem that I know of.

BF: Yeah. I didn't have a problem with that. We had a Memorial Day holiday, and we had been in, what were we. I guess we had just about announced it. We were just right in front of announcement, of the facts that we were doing the requisite tests, and one of the set of tests that we were doing just before that holiday had been a set of temperature cycle, humidity kinds of tests which were standard on all of our disk drives. And we had left for the weekend, turned off the test chamber. We came back in on Monday morning and about 8:30 I got a call from the test lab manager, who said "Something really strange has happened over here. I think you ought to come over and see it." And I said,"Well, what is it?" He said, "We tried to start the disk drives, and they won't start." "Oh, well, do you know what you're doing?" "Yes. It won't start." And I said, "Well, so what do you want me to do?" And we kind of -- I just didn't think this was a big problem, and so we quit. And he called back a little bit later and said, "No, Bob, this is it. You've got to come over and see this." Well, I came over there and they were stuck. The heads were glued to the disk, in essence, and we couldn't get the thing started. What we found out in that particular test was that we had oil sprayed all over the inside of the disk from some porous castings that had not been cleaned right. We got all that fixed up, rebuilt all of the test equipment, put it back in to test, and started through the same set of cycles. And I think it was 4th of July, if I recall, when we ran into the same thing again. Guess what? Stuck. This time there was no oil. So, this led to a whole, in fact, it was a site-wide and long set of task forces to try to figure out what happened. That's when I got to know intimately some of the people in the materials lab. Leo Phipps (sp?) for example, with his strange concoctions of things. To make a long story short, what it finally turned out to be, was forms of salt, sodium, in nanogram quantities left from various disk washing processes that we had, if I recall, right?

MW: Uh. huh.

BF: And it took an enormous amount of time to get those processes refined. Enormous amount of time bearing in mind we were gonna ship this thing down the road. So a lot of work was put into cleaning up the disk process and in the meantime, Jack was helpful in some of these things. Helpful by telling you, you gotta have a backup plan. You gotta find a way to never let these heads stick. If a disk drive can't fix it, so there were some odd ball things put into the machine, like heaters and coolers and head extractors.

MW: Yeah.

BF: Subsequently, though, we never did have that problem with this drive after we got it into shipment, because we did get the disk process straightened out. However, this problem endured in various forms even till today. It's really <inaudible>

JP: I was going to point out to you that IBM is not the only company that had that problem. And the industry eventually called this stiction.

MW: With the head sticking to the disk it is stiction -- that was the word.

BF: Call if down was the railings, yeah.

JH: Well, it was really the first time that the interface we'd gotten close enough to the…

BF: To the surface finish.

JH: Yeah, the surface finishes were so good that materials that just exuded from even any plastic material you put into the enclosure will out-gas. And it led to a whole new sort of technology.

BF: Yes, it did.

MW: Where: of examining components.

JH: To get components that are stable over years, even not in a closed environment.

JP: So we're representing that this was the second generation of a closed environment.

MW: Yes.

JP: Where gases coming from the interior. Before that they weren't a problem.

JH: Look at any disk drive today, it's all sealed up. And whatever you got in there, you're gonna live with.

MW: But, as the air is pumped around on those surfaces -- the heads very close -- there's a lot of turbulence in there, and it has a way of just backing it up, and getting any particles, anywhere in that system will get caught on the head.

BF: Just stick on the head.

JP: The head's a perfect collector of…

MW: Yes.

JP: But does it give you a feeling of pride to have pioneered a problem?

MW: Right. Well -- it was..

JH: The solution was worth something.

BF: Yeah. It created a lot of legends around San Jose. The code name for the problem -- we gave the problem a code name. I should say the materials lab did. They called it Sirocco, and we all looked at em. Now why do you call it that? And -- well, Sirocco is a hot wind off the desert. And that's how they defined what was happening in there, that this temperature in certain humidity situations would cause the sodium salts to dissolve under the head, and then when the temperature and humidity changed, it became a glue, and held the head. I mean, and a 1.5 horse power motor could not turn this thing, unless it ripped all the heads off, which we did see once happen, too. And we had all the heads spinning on, glued to the disk.

JP: It must have been an interesting feeling to have pioneered through all those other problems and getting the 3340 technology ready, and then transitioning it to this, and then finding a completely new one.

MW: Oh, yeah. Well, I re…

JH: We always do.

MW: There was never…

<inaudible, all talking>

BF: We found a demand for this device became so huge that we actually had to build a new disk plant. We built more space for boxes. I mean, it was an overwhelming success.

JP: It was huge success. Was all the production -- did it all remain in San Jose?

BF: Yes.

MW: Yes.

JH: Well, yes.

MW: Well, as far as the U.S. goes.

JH: We had foreign plants -- we had…

JP: You had foreign plants, also?

MW: Yes. In Japan…Fujisawa.

JP: Meiningen in Germany?

<inaudible, all talking>

JP: Did they know who was up?

MW: Produce all of em.

JP: Yeah. And produce all of em. Yes.

JH: We had been sourcing... But this was the first product where you would see these huge floors.

MW: Right.

JH: Of nothing but disk storage. This is where your data bases came to be monumental.

JP: Well, it was a big success and I guess we should point out that along with an earlier generation, especially in the middle of the '60s, the so-called plug compatible disk drive manufacturers who were making imitations of IBM's drives that were compatible at the interface, and in those days compatible with the same disk packs, transitioned the plug compatible people into making when IBM went back to the fixed disk, they transitioned into this, also. And some of them had better luck in doing the 3350 than others. It became a huge hit for companies like Storage Technology. And a few others. It was a big problem for Memorex. They didn't make the transition as quickly. But there were a few other companies that made all sorts of drives that were exact imCHMOralHistIBM3340.docPhil Glau Page 32 9/27/2006ages at the interface of a 3350. So IBM pioneered with setting a standard for the industry, which was not only a standard for IBM, which was a big success. It also made some other companies very successful in doing exact imitations. And of course, by studying what IBM had done on the technology innovations, they probably saved themselves a lot of trouble, I would imagine.

MW: Um hm. Yeah, yeah, sure.

CC: Yeah, that's true. I'm sure that's true.

JP: Yeah.

KH: Well, the only way is: You gotta keep innovating and moving ahead, or else you get your lunch eaten.

CC: Yep.

JP: I guess that's the nature of the disk drive industry, isn't it?

MW: Yeah.

BF: The interesting thing is that this form of the head, subsequently added with the thin film processes for the active coil, became just the basis for everything that's come after. I mean, all of your disk drives now, use that fundamental form of technology, modified slightly, and then added with thin film,and then conductive, and then thin film magneto resistive elements, to the batch.

JH: Although the…

CC: Yeah. It did do things like eliminate third rail in the middle.

BF: Yep.

MW: Yeah. Slightly, yeah.

KH: And because it got so skinny you couldn't keep from breaking it. Happened to…

BF: So that's why the word Winchester is so pervasive, is because it started here and became essentially the prototype for everything after.

KH: In fact, when I went to the university, I found my faculty in the computer science department called all disk drives like that, Winchesters. That's the only word they used. Winchester.

JP: You see what you did to the industry?

BF: And it was a fun time.

MW: Yeah.

JP: Well, anything else to add to this? No, I think not, and I'd like to thank everyone, and I'm sure the industry thanks you all for sharing your thoughts about how all this happened, and why, and all the interesting problems in getting it started. Thank you, very much.

END OF INTERVIEW

