



Oral History of IBM 1311 and 2311 Disk Drive Panel

Interviewed by:
Jim Porter

Recorded: February 22, 2005
Mountain View, California

CHM Reference number: X3111.2005

© 2006 Computer History Museum

Jim Porter: Well, we're here today to discuss the development of one of the most important types of disk drives that was ever developed. The idea was that the 1311, which we're going to start with today, transformed the industry by going to a smaller disk than had ever been used before in a disk drive and it came up with a removable disk pack, which made possible many applications for these products that were never used before. And, by the way, it probably contributed quite a bit to the disappearance, gradually, of the tabulating card, which had been the method of entering data for the computer industry for many decades before; in fact, back to the previous century. So the 1311 disk drive, which we have a sample of sitting next to the table, was originally announced in October of 1962. It was shipped in July of 1963 and it became, as I pointed out, a very important change in the disk drive industry. And we have here today three of the gentlemen who made all that possible. I'd like to start by asking each of the gentlemen to give us a short background on each of themselves and let's start in the middle with Jack, who was a leader of the project that started all this.

Jack Harker: Jack Harker. I joined IBM in 1952 in a brand new laboratory in San Jose that had just been started under the direction of Rey Johnson. We worked on a lot of things to start with but, pretty soon, our development was pretty much concentrated on the first disk drive, the IBM 350. That was underway in development. Then I worked on the succeeding product, the one that would succeed it, we called it an Advanced Disk File, which became 1301. And then, as that development was proceeding, Lou Stevens, who was the development lab director then, asked three of us to look at a reduced size RAMAC. In other words, half the size and half the performance, half the cost. And we started doing that and we came up with some designs and one of the parts of the design was a small disk file. The other parts of the project got subsumed as part of an effort which you'll hear about later, which is the eventual project that the 1311 came out with, the IBM 1440, but the disk drive took a life of its own. And I was working on it and came up with a design for a half size disk. And, at that time, there was also a project that had been started in the research laboratory, which had separated out under Al Hoagland called the single disk file. And what he was looking at was a removable disk that could store the contents of a nominal tape drive which, in those days, was about five megabytes. The tapes had much more capacity but it wasn't really practical to use them because you could only put a single data set effectively on a tape reel. There was a study done by a guy, John Nolan, in our product planning that suggested that such a file, if you were operating sequentially, as all the data processing was done in those days, where you matched a transaction file against the master file, if there was a relatively low percentage of the records that would be updated, because the disk drive could skip over records that weren't to be updated whereas the tapes had to read and copy each and every record whether it was or wasn't. When I saw that study, I said, "Ah hah!" And, at that point, my low cost file became a removable, low cost file and that's where it started. We actually started designing things and commissioning some hardware in early '59. At that time, I started out, I was a one-man project doing all design and then, later, I don't know, I think it was...

Jim Carothers: In 1959, I was <inaudible>

Harker: ...'59, Jim Carothers joined me.

Porter: Okay, Jim.

Carothers: A little background on my part. I joined -- graduated from college, the University of Kansas, and went to Poughkeepsie, New York, to work on the 701 computer. It was, at that time, called the defense calculator, the first one. And, in aiding and assisting on working that, they needed people to go to the field because the field engineers were not felt to be capable of maintaining a machine of that complexity. So, in any event, I went to Lockheed, California, where I spent approximately a year and a half in aiding the servicing while customer engineers were being educated and brought up as to how to maintain the system. Returned to Poughkeepsie and worked on a -- was to be a forerunner of the 703. That was going to be a tape sorter type of unit that was to be used by the Census Bureau. Well, that was not overly successful. I was able to manage to get transferred to San Jose and became an assistant to Lou Stevens when I first came into the laboratory. Did not know, really, what they were working about but I quickly found out when I was assigned to write the manuals of the operational nature of the 350 RAMAC and so it gave me a great understanding. Then, when they decided that the original version, the (A) version of the RAMAC system was not capable of being sent to the field, that it needed to be upgraded in terms of reliability, I became the electrical engineering manager working for Dennis Willard on that file system. And we improved the thing and got it transferred into product engineering but still the system, at that moment, seemed to be not as reliable as we wanted it so we undertook a file that would duplicate the 350 and it was to a hydraulically driven unit using transistorized circuits instead of the tubes that we were using. When it was found that our system, in going to the field, it's the RAMAC system, had improved the reliability, wasn't as bad as we thought, we stopped the program I was on, Jack offered me a job to come over and begin work on the low-cost file that he had just described.

Alex Solyst: And I'm Alex Solyst. I joined IBM in 1961 and was assigned to a <inaudible> for file technology and spent probably the first year looking at different air bearing sliders and some instrumentation and developed an optical measuring system to measure the spacing between the slider and the disk, the glass disk it was that was implemented in manufacturing. I was later assigned to do the 2311 slider and suspension modifications. There were modifications to the 1301. I had spent some work on 1301 trying to identify and solve some technical problems we had on there and, when I got the 2311 development job, I incorporated some of those fixes. I then later developed the 2305, nine-element head assembly. That was the first slider and magnetic core integrated structure and that was later implemented in the Winchester heads, where I also worked. I was later assigned to inkjet technology and I came back to the file technology again for the last eight years before I retired in 1991.

Porter: Thank you, gentlemen. We come back to the 1311 and I know it went through, as Harker pointed out, a number of stages of development before it actually turned into a tangible product. There were, as I gather, Jack, there were a lot of discussions, maybe arguments, as to how many disks it should have and how many -- what the diameter of the disks should be and all those things?

Harker: Yeah. Well, as I said, it was to be a half-sized RAMAC so I started out saying it should be a 12" disk and I laid out all the <laughter> the data area and I neglected to think that the slider needed at least another half inch outboard of the last data track so then it became a 13" disk, which is how it remained from -- through a lot of the development. Jim can cover how that changed. We built an early model that was sort of a mockup and you'll note this has a five-disk disk pack but the original one I started with was a 10-disk and you'll see that there are part of that structure that appears to be unused. And that is a story that goes with it, too. The reason it went to five disks was these were steel disks that had been oxidized to provide a hard, magnetic layer and -- using a vertical probe head, which is the technology that the Advanced Disk File became in the 1301 was using and I was writing on that technology. But 10 of those steel disks was pretty heavy so if it's a removable disk pack, I was persuaded that you really ought to find a way to do it with half as many disks and that's what we did. And that was the original structure and we built this as a test bed and it actually survived a number of years. I think it wound up being attached to a system in development.

Carothers: It was the forerunner of the 14LC that was being...

Harker: Okay.

Carothers: ...progressed.

Harker: And the program started. It was an interesting selling experience, selling the corporation, et al, on this. Fortunately, I had a very creative fellow in our product planning, Chuck Hester, and Chuck amazed all his compatriots by saying, "If we go ahead with this, we'll build 20,000 of them." And, at that time, no IBM product had ever <laughter> computer product had ever had that kind of volume and he had a very hard time convincing people. And actually, he and I, we built up a wooden mockup of this which we took around the corporation, the two of us on a sales trip. <laughs> Fortunately, we got the support of the company's real mana -- technical management, John Haanstra, Scott, who was the president of the division. They thought this was a neat little device and so we were off and running. After we'd just gotten started, there was an event that came along which was that the file that I was modeling this technology on, the Advanced Disk File, got in deep trouble. And so they pulled a number of people off other programs, including me, and, at that point, Jim took the project.

Carothers: Right, right.

Porter: So, Jim, when you took over the project, what did you think were the most important challenges at that time?

Carothers: Well, we saw the fact that the ADF, which we were using their technology, basically, the esteemed home of this disk probe head and so forth was running into difficulty. That we wanted to achieve the goals of two million characters per disk pack but we needed -- for that size disk, we needed a double bit density. We could keep the track density the same as the ADF. But, with them running into trouble, we said -- and with no commitment as to the technology group to develop a head that would meet our double density requirements, bit density, we said, hey, we better back down. We better make it so we can at least achieve the capacity of two million characters. And the ADF then was going through a conversion to a ring head and then aluminum became a basis instead of steel so we said, let's double the size of the pack. We'll have our two million characters. It'll be manageable. Little heavier than we wanted as a ten pound 6 or 10 surface kind of disk pack and so that's the way we went. And that meant a redesign at that point in time and we got into that. And then, of course, what happens is the ADF starts clearing their problems and then we go back to the technologists and say, "Hey, can you achieve 1,000 bits per inch? Can we return to our original form factor that we were after?" And they looked at it and it required, on the inner track, for us to have 1,100 bits per inch. They would not commit to anything but 1,025. <laughter> Well, what does that mean? <laughter> Well, you got to get this information on there so the disk suddenly grew one inch and so that's where we came up with the 14" disk.

Porter: And the final spec on the product as it was later introduced was 1,025 bits per inch.

Carothers: Yes. That's correct. So that's where that density came through. And it, really, it only came about after they had made one head that achieved those densities. So that's why were delayed and we were out of synch, really, with our low-cost system, the 14LC, as we called it, because they were already developing it. Here we were switching gears, reengineering the base plate, the unit, and so we were under pressure to try to make up that difference in scheduling, which we never did. We came in six weeks later than them in going into the announcement kind of a test.

Porter: Well, how about the issue of the fact that this was going to use removable disk packs and you would have disk packs move between different drives so that the head assembly was not always going back to the same set of tracks, it might have a new set of tracks in front of it the next time?

Carothers: See, the word "removable" was in the original spec. It's still there. And removability, at that time, we didn't know how many spindles were going to be on a system, maybe you could just have one spindle and just a lot of packs in coming on and off. But, as we got into it and we saw the requirement for multiplicity of drives, we said, hey, we're close to making such a thing called interchangeability. This meant that, if we'd accepted that goal, then we would have packs that were -- whatever 1311 disk drive they were written on, they could be read by any other ones, whether in United States or 50-cycle machines over in Europe and so forth. So, in accepting that challenge, there were not a lot of great differences except how to make certain that all files that left manufacturing were aligned to a standard and that, in the field, if they had to replace a head, they had a standard that was created in the factory to adjust that head to so that we could maintain that accuracy. Then, of course, by the -- all these factors of

trying to accomplish all of these things, it meant a real close control on tolerances in every manner possible, from disk coating thicknesses to speed of rotation to bit shift, mechanical tolerances, everything had to be recognized, had to be measured and had to be related between the head, disk, and the electronics and the file itself.

Porter: Of course, one of the most critical parts of this whole thing were the heads and we do have a gentleman with us who's had a significant part of that.

Solyst: Yeah. I mentioned earlier that I joined IBM in 1961 and this assembly is the head assembly for the 1311 where it's just getting out of development at the time and going into product test. And I worked on some problems that were identified in product test and some problems later on I identified in the field and didn't really find good solution to any of the problems before we got to the next level of development, the 2311. But the problems that had been identified were basically two major ones. This 1311 head, the stainless steel slider, is suspended in this yoke with two very small pivot bearings at each end of the slider and the load then is transmitted from this arm assembly through the yoke and into the slider, at each end of the slider. So the load was about 350 grams on the slider and, with that load, at least was tiny, bitty, the pivot bearings, the screws would loosen or something would change, maybe the yoke would give a little bit but they loosened up in the field after some time of use. And one of the maintenance procedures that field engineering had been instructed in was to, at regular maintenance intervals, to go in and check the looseness of these pivots and, if they were loose, tighten them, oil them a little bit. It was very critical, very crucial operation.

Porter: You mentioned oil them?

Solyst: Yeah. These little pivot bearings...

Porter: So you didn't use 3-in-1 oil, I don't assume?

Solyst: No. <laughter> No. The field engineering had a little vial of oil, it wasn't bigger than that. It was enough for 10,000 packs <laughter> and a little oiler, which was a handle with a little thin wire stick and you would dip it in the oil and shake off most of the oil on that little wire and then very, very carefully, on the -- a loop go in and find that little pivot head there and hopefully apply the oil in the right spot. And sometimes it didn't get there, sometimes it got on the disk later on. So it was a problem. It caused a lot of maintenance, difficult maintenance. Another one we had identified was the pull tip. You can see the pull tip in there, which is that black little thing in the middle there. That's a laminated core that's extending through the stainless steel slider and it's epoxied in place. The core consists of a bunch of laminates of new metal and stainless steel slide shields and that assembly, that core assembly was potted in place in the stainless steel slider. And then the final operation was to carefully lap cylindrical surface on this slider. It was found that, with time and with temperature going up and down, that that little magnetic core

that was sitting in a big blob of epoxy was moving up and down with temperature and humidity. And, at elevated temperature, it would extend out from the slider and sometimes extend far enough where it could damage the disk. If it didn't damage the disk, it would wear off the pole tips.

Carothers: You couldn't maintain the space and constant...

Solyst: Right.

Carothers: ...with respect to the magnetic surface.

Solyst: Right. So we had some erratic performances of files in some of these cases that were identified to the movement on the screw. So, when I got the assignment to modify the 1311 head, my manager, Huey Lewinson, he called me into the office and gave me about a five minute spiel as to what he wanted. Fix the problems, two major problems, fix them and the other part he wanted was there was too much operator dependence, very careful adjustments and very difficult assembly operation and the operators needed to have a lot of dexterity. And he said I want to get some of that stuff out of there and "tool it in place" was the term he used. "Tool everything in place." So, with those marching orders, I set out to fix the problems with the pivot ring and pivot bearings and it was a fairly easy solution and someone had already started with a flexure suspension to replace the gimbal yoke. But the first attempt was unsuccessful because I mentioned earlier that the 350 gram load is transmitted to the slider through that yoke and the pivot bearings. And, in that first flexure design, an attempt was also made to transmit the slider load through the flexure. And now the requirements of the flexion being rigid enough to transmit 350 gram load to the slider was not very compatible with the other requirement, which was to have a very soft roll and pitch moment exerted by the suspension onto the slider in order to maintain constant flying height. The solution was fairly easy. It was to separate the load transmittal. Instead of going through the yoke or flexure, step directly on the slider at the exact point where the load needed to be applied, which was at the design pitch and roll axis in a section. And so, in this 2311 head, one can see a little belwin button stepping directly on the slider. And the load arm, which is not shown here or here, was stepping on a dimple up here. But, anyway, that required -- that removed the requirement of transmitting through the flexure. And now the flexure can be optimized for the other purpose, which was to provide a rigid suspension in the plane of the disk but a roll and pitch operation that would afford the slider to follow the contours of the disk. Flexure is made of thin stainless steel and it is spot welded onto the slider at two places here and onto the arm assembly at these two places here. That has been very successful. Just about every slider has followed those same principles in later developments.

Porter: So it worked?

Solyst: It worked well. The other problem was that pole tip movement and we tried some different epoxies and we were just not successful at all with that. And I realized that, if we could implement a mechanical holding system instead that, rather than have to develop, to...

Carothers: Instead of the bonding kind of thing.

Solyst: Yeah. It rather depend on the plastic, the epoxy, you know, have a mechanical holding place. And the first thing that came to mind was a simple staking process. We reduced the slot in the slider for the suspension -- for the magnetic element stack just wide enough to accommodate the tolerances of the magnetic element stack. As you can see, it's a very small slot in there. And then we attempted to stake the slider with a little sharp staking tool on the back side and it sort of worked but, when we started looking into the potential tolerances we would run into and also recognized that this laminated stack of new metal and stainless steel was laminated together with a thin, insulating, bonding film, it was kind of a mushy assembly. So, in order to get reasonable good holding power, you know, we had to overcome that, also. So simple staking didn't work and I devised this system here called four ball staking in which the core is inserted into the slot and holes are drilled adjacent to the slot, small holes, blind holes into the slider from the backs of it. And then oversized balls, one millimeter balls is what we used, oversized balls are pressed into these small holes, squeezing the material out enough to overcome tolerances and mushiness in the stack. And we had some initial problems with it in that, as we squeezed some of the metal, squeezed too hard because, at high tolerances, the gap would open up, which was detrimental to the operation of the file. So we experimented with locations and number of staking balls and found that the best arrangement was to have a system of three staking balls, one at each end of the magnetic stack and one in the middle of the stack. And we further found that, if we stake the outer ones first, they compress enough and sort of hold the elements, the flexible elements, the magnetic element in place so, when the last ball comes in that really does most of the holding, the gap would not open up. So that's what we wound up with staking process.

Porter: Well, that's an impressive list of innovations that were made on the fly as you were working, trying to make all this work for the heads. Now, for the disks...

Carothers: Before you leave that, though, let me have the head here. One of the key things that all this is leading to is to maintain the spacing with respect to the disk of 125 microinches. Now, we get lost oftentimes when you say microinches. We forget that 10 to the minus six power so sometimes it's difficult to interpret what that really means. I have here in my hand, of course, an IBM card. It's got three dimensions. The narrowest dimension is this one. It's about seven and a half thousandths thick. If you divide that into 60 parts, you then have 125 micro inches. Gives you some appreciation to how close we were running with respect to the disk. Of course, the disk had to be flat, could not have asperities over 60 microinches in height. And so that is part of the technology that went into this. Maintaining the disk coating was also very valuable. We had a nominal thickness there, as I recall, of something like 200 microinches and it did vary. And one of the interesting things in -- when we put the -- at one time, put the

1311 into product test, in the initial ones, we thought everything was going great until we put it into a humidity chamber at high temperatures. The coating that had been developed that I think they called it SP-1 in those days, it was a special one that they thought would be better, had gone through what we used to do to show the resiliency of the coating to head landings and so forth. Was a tap, tap tester which tried to break through to a disk surface and, as such, if it did, you counted how many times it did it. In any event, when we got it into product test, we never did break through but the surface and the coating softened. We hit a spot on it, it would pick it up on the head then later it would deposit it some place else. And so we were having head crashes. We often joked after the fact that we never lost a bit, we just redistributed the bits around the track. <laughter> But, from that, though, then we said, no, no, that's the wrong coating. So the people in the disk technology had to go back and reformulate, get a different binder to hold the particulates in the surfaces of that coating.

Porter: You mentioned you'd gone from steel to aluminum. Were these the first disks using aluminum?

Carothers: Say that again?

Porter: Was the substrate?

Carothers: The substrate was aluminum and that's what was going to be, finally, used on the ADF file.

Porter: Uh.

Harker: Well, the first RAMAC...

Carothers: And the first RAMAC was aluminum.

Harker: ...was aluminum.

Carothers: But it had a honeycomb structure...

Porter: Right.

Carothers: ...for support on that one.

Harker: Yes.

Porter: Why was the -- why was steel used on the early prototype work?

Carothers: The steel, as I understood, was because you had a probe head with a coil around it. Think of a pointed nail, a rusty nail, as somebody said, and you needed to form a magnetic flux field that went from the probe in through the oxide coating, through the disk, back through the spindle and have a return path. So you've -- that's why it was steel all the way through and, when we went to the ring head, we didn't need that. That's why the ring head just lets the extension of the flux around the gap penetrate the surface of the coating to record the information.

Porter: Okay. And the coating on the disk was refined such that it had a sufficient consistency and...

Carothers: Of oxide.

Porter: ...lack of hot spots?

Carothers: That's correct.

Harker: Well, the original reason to go to the steel disk was it looked like it was a lower cost recording system because you took a steel disk blank and you put it with a -- clamped with another number of its mates into a furnace with a reducing atmosphere and you formed a ferric oxide coating on the surface. And we'd had enough trouble with that business of coating disks with paint that it looked like a very attractive alternative and it looked like a probe head might be simpler in structure and cheaper to build than a ring head. And that was the technology that it was proposed for and, through most of the program, was used in the ADF became the 1301 and it was only when we realized you just couldn't get steel that was defect-free enough so that you have an adequate error rate in a recording that forced us to go back and switch back to aluminum disks with their paint coating. That is interesting because, now, a lot of the very advanced work in magnetic recording is going to vertical recording <laughter> but this has been said for many years and I don't know yet we've seen real products with it.

Porter: Okay. Stop here.

[video off then on]

Porter: The idea that it used a very interesting head positioning system so that, after each seek to a given track, the head system went back to its starting point, didn't it?

Solyst: Yes.

Harker: Correct. The original -- actually, the other thing that was trying to achieve there was finding a way to get cost out because the actuator on the 1301 was a very complex piece of hydraulic equipment...

Solyst: <inaudible>

Harker: ...that required rifle barrel drilling techniques to get a bore that would support a series of pistons that did coarse positioning and then there was a fine adjustment. And the cost we were proposing for this file was less than the cost of the head assembly on the 1301. So we had to come up with something simple so we just took -- hydraulic actuator was still a simple thing and it just was, you know, used a piston and, in the exit and when you wanted to start moving, you'd open an exit valve and that would drive the actuator forward at a relatively high speed. And then you started, as you crossed the tracks, we counted. There was an optical disk that you count your track crossings and, when you got to 10 tracks, because that was an easy carry point for an adder <laughs> you then, you closed a second valve so it restricted the flow of oil out and that slowed you down and then you just let a decant go free into a cog wheel and that stopped you. And that was about the simplest system we could think of...

Porter: But after each...

Harker: But, after each time, you had to go back and start over again.

Solyst: That's right.

Harker: But, at first, we were happy with that because it was cheap.

Carothers: It was cheap and reliable and it had the registration with the detent coming up against a detent wheel that was always positioned very accurately. It meant, of course, maintaining those detent wheels very accurately for all of the population but, like Jack said, when you wanted to make an access, first when you loaded the heads, when you put a new disk back on, the carriage or the design was such that you pushed the actuator all the way forward. That cammed the heads into a locked position so now, as they came onto the surface, they were being cammed down in preparation for reading and writing. Then, after it got in, then the signal brought it out to track zero. Now new addresses always took the form of what track do you want to go to? Do you want to go track 80? Set it into the register as 80 and, every time that the actuator made a movement, the detent wheel as well as the slotted sensing device of light sensing across slots down-counted that counter. So it came down when you got within 10 tracks, like Jack says, you switched. So you slowed it down so you had more accuracy. When you finally got to that zero, bump! In came the detent and you were locked into a registered position.

Porter: Okay. So the specification then was that the average seek time from track to track was 250 ms?

Carothers: Correct, because you always went back through.

Porter: Right.

Carothers: Zero...

Porter: But you came up with an improvement to make it go down to 150 ms?

Carothers: Yes. You can well recognize the problem of when you, let's say you were on track -- doesn't make any difference, let's say 80 and you wanted to go to track 81. In the old method, back out, you only had one track to move forward to it but back out and in you went. So you were consuming time. So now then what we did is, to make that direct access work, we always took -- knew the address we were on, knew the new address and took a difference. And certain days, it would say negative five or negative one or plus five. If it was plus, that meant you were going to pull your detent and go straight on in five tracks. Now, then, if you were going to back out, have to back out, then you pulled the detent and it went in reverse and you would pass the place you wanted to detent then start moving forward at the pre-selected time, put in the detent and you had it. <laughs>

Porter: Okay. The net result of all of this was a project that was much more successful than anybody had thought it would be. According to the numbers, by the end of 1964, now, mind you, this product first shipped in July of '63. By the end of 1964, over 6,000 disk drives had been shipped and 41,000 disk packs. Now, in those days, IBM leased the hardware mostly, of course, and they were getting \$375 a month for the rental of those disk drives and, for each disk pack, they were charging \$490 on a selling price for each disk pack of 41,000 disk packs.

Carothers: You could rent the disk pack for \$15 a month.

Porter: You could but a lot of people, buying them was a pattern, too.

Carothers: And do you understand where the \$490 came from? No, probably not. <laughs> It wasn't just a discount price. They had gone and, I guess, checked with customers and, if you bought anything over \$500, it was capital equipment. You stayed under it, they could order the department and have multiple drives. <laughter> So that's where that came in but there was also another key factor that I really hadn't expressed and that was in making interchangeability, did we take into account the problems associated with removal of the disk pack and storing it and making certain it was protected when you took

it out and kept it in a protected environment. And that's one of the things that we recognized in the program finally and we said, oh, yeah, we got this thing. It works and everything, you know, with this thing but, well, what's going to happen? How do you get that thing out? Isn't there any cover? Can you get Coca-cola on it? What's the problem? So then I remember John Haanstra had been out and he was kind of querying some of those problems and so I remember Bob Pattison and I and we sat down and briefly discussed this thing and said, "Hey, we need something like a cake cover, you know, something to come down over the thing to protect it." So he and the others, then, undertook an investigation of how could we do something to protect it? And they cleverly came up with the concept of having that like cake cover but it was designed in such a manner that if, let's say, you have a disk pack on the file, took the cover, you dropped it down over the top, it had guides, and, as you turned it off, it locked the cover onto the disk pack.

Harker: And unlocked it from the...

Carothers: And unlocked it from the spindle. And then you took this and you had another cover, you put it on. Now you had your pack that you could put up on a shelf and kept protected.

Harker: Actually, the patent on that unlocking was one of the most valuable patents that IBM had because all of the competitive disk pack manufacturers had to make it so they fit. And, to do that, they had to use the patent.

Porter: Well, I guess I should mention, in passing here, that, seven years after this, I went to work for the company that became the leading maker of disk packs called Memorex and I took over the product management function on all those disk packs. <laughter> I was very familiar with all of that and, fortunately, IBM had a cross-licensing arrangement which made it possible for all those people to just sit down, without much legal work, and arrive at a reasonable license back and forth. So IBM was really quite generous with us and was not creating a monopoly out of their patent structure but creating a source of income as well as allowing other people to be competitors. And, out of that whole thing, of course, came the whole disk drive industry which followed. But the 1311, as we've started out by saying as everybody has said, had become a major, major change in direction for the disk drive industry. It created, as we've talked about, the removable disk pack. It went down to 14" from the 24" which had preceded it. Many other introductions of good ideas in the heads, the disks, et cetera, followed. And, of course, now, we're dealing here with the 1311 with a drive which had 2.68 MB total capacity per pack. In '64, in that year I mentioned a moment ago, a year and a half after the 1311 was shipped with those very large, unexpectedly high shipment levels, IBM introduced the 2311 in April of '64. They didn't actually get it shipped until June of '65.

Harker: With the systems.

Porter: But the 2311 used the same removable disk pack as the 1311, the disk pack became known as the 1316 disk pack. It used the same disk pack, but it increased the capacity from the 2.68 MB of the 1311 up to 7.25 MB for the 2311. And, of course, it did that by increasing recording densities and it made the seek time better and so forth. I would assume that it must have been quite a monkey on the back of the engineers to get this project done right after the 1311, to go up to this higher density product and get it ready to ship.

Carothers: Well, the basic configuration, mechanically, was basically the same, the access mechanisms and things like that. Where we had to pay special attention to, though, was the increased track density and then the tolerances associated with that to maintain our registration over that half/half width track, so to speak. And, in doing that, one of the things we did is really understood, even to greater detail, our temperature variations. We could almost walk in with a disk pack. It had to be stored in the same room on the 1311 and put it on and it would be up to speed and the temperature variations and operation didn't seem to affect it. But then what we did, we said, hey, we better have a little longer delay for further stabilization. That was one of the key things. And when you talk about the capacity that was on the disk file, it was, in MB, 2.68. In characters, and we weren't using the EBSIDIC code on those first files of the 1311, it was 2.98. And each character on there consisted of six informational bits, one parity and two synchronizing bits in that. So we had a code of nine bits that was accompanying it that you didn't first recognize when you look at numbers like this. But the reason for the two synching bits was our oscillator and the speed variations that we might encompass might mean we would mistakenly miss a pulse or erroneously pick up something. And so the resynching always said, when we read back, we always switched when we saw a bit, to another -- we had two oscillators and we switched from one oscillator to the other. But, when there was no bits and the one oscillator continued to run so it's missed registration or its tolerance was getting out of line with -- perhaps, with the information. So that's why we had those multiple ones. So that gave us more flexibility in terms of the capacity on the inside track, which always determines the total capacity of the pack. And so that's why you'll see that kind of a diskrepancy.

Porter: And the 2311 also went down from the -- you mentioned the direct seek feature which -- on the 1311 which achieved an average seek time of 150 ms. The 2311 took it down 85 ms average seek time. You had to tweak something to do that, I assume.

Carothers: I have forgotten what that is, I'll be honest with you.

Harker: Well, was it the not return to zero?

Carothers: No, that was used in the direct access.

Harker: That was used.

Porter: Well, I would assume it was just a matter of fine tuning what you'd done before.

Carothers: It was probably encompassed, to some degree. As I say, I don't recall that right now. The fact the track density is a little higher so you don't have to move as far to get to the tracks.

Porter: Oh, yes.

Harker: Well, actually, you know, this basic mechanical structure, including that head assembly that could take an extra five disks, was carried over to the 2314.

Porter: Right.

Harker: Really, without any substantial change.

Solyst: Yes, that's correct.

Harker: They did find that, as they upped the densities, that particles on the surface were more of a problem and had to put brushes in to clean the disks. <laughter>

Carothers: Yeah, contamination was certainly a problem, what with riding that close to the disk surface. And one of the things we recognized in the 1311 was the fact that our head structure, we had a solid hub disk, solid hub disk pack and so air that was knocked out, you might say, with the heads coming in, was drawn in on the reverse side of the heads and any particles that were in there were pulled in and across the surface. Air always flowed across the surface from the inside to the outside. So what we decided we needed to make a pass so that the air could enter from the hub. And so we created what was called a vented hub with spacer rings which had cuts in them to let the air out to come across the surface. So that was a step forward, then, to preventing any air entering the pack except through a filter in the base of the file and even a coarse filter on the bottom of the pack so that, if you were handling the disk pack outside of the thing, you didn't suck dirt up into there. And so that was recognized as a problem. Also, we had found that our hinges that we used for the lid to the file, this portion here, they were generating some iron filings and, of course, naturally, where do they go? <laughter> Right into the disk pack. So we had bits in there so we changed the hinges on there so that that could -- changed it to a different type and limited that from occurring. Then, at one point, we recognized, even particles, if they got onto the surface, you could run them at 3,600 rpms, way beyond our speed of rotation, they wouldn't come off. They'd stick kind of so they needed to be dislodged. So then what did we do? We created a brush system where the brushes on the startup of a file would come in and go across and come back out, dislodging any particles that might be on there. And of interesting fact, we found that the best brush material was from hogs in China. <laughter> It was an interesting thing. Boy, we were working with all sorts of material.

Porter: It might also be noted that the disk revolution speed went up. In the 1311, it was 1,500 rpm and on the 2311, it went up to 2,400 rpm.

Carothers: Right. Right.

Harker: That probably contributed to the reduction access time, actually.

Porter: Could be.

Harker: Because part of the access time is rotation.

Carothers: Yeah.

Porter: It's latency, yeah.

Harker: Yeah.

Porter: Now, Jack, there were some improvements in the heads, weren't there?

Solyst: Yeah. We should mention that, from 1311 going to 2311, the magnetic element was changed. In the 1311, there were actually two elements. There's a read/write head and an erase head in 1311. And in 2311, we again had two heads, a read/write head and a unique erase head that's called a tunnel erase head that had been invented by Lee Dawson, who was one of our engineers there. So that is probably one of the major contributors to being able to go from 50 tracks to 100 tracks.

Porter: Per inch?

Solyst: Yeah.

Porter: Very good. Well, these improvements also made it possible for IBM to do something else very interesting, to bring out the first single disk cartridge drive, which was known under the group name of a RAM kit. And using the same technology on heads and disks that was used in the 2311, they brought out the first of the single disk cartridges and those drives were then used as computers kept getting smaller and smaller. In those days, what was being small was to build a computer the size of a desk, which was something that was regarded as quite revolutionary. So, on small business computers like the 1130 and

the IBM 1800, they started to use, in 1965, IBM started to use a single disk cartridge using exactly the same disks, heads and basically the same technology as the 2311, all of which was a byproduct which conveniently became available to IBM because of all this great work that had been done on these disk pack drives. This work, was it done in the same group or was it separated?

Carothers: Separate group was working on that. Person by the name of Walk Buzlick[sp?] was in charge of that program.

Porter: And separated in the same geographical area or to a different location?

Carothers: He was in the same building, I think on the same floor but it was a separate program and they were using the components, as you expressed, that came from the 2311 and packaging in a different manner.

Porter: Well, the bottom line, again, was that this -- these programs, in general, going back to the individual, removable disk pack, the 1311 and the 2311, in both cases shipped incredibly more product than IBM's planners, I think, had ever expected. Was that your impression, too?

Carothers: Absolutely. That's happened, it seemed like. And everyone, even the 701 computer which I mentioned I had started on, they thought they were going to have maybe a dozen of these things and double quantity that went out to the aerospace, the Atomic Energy Commission and so forth. But the -- you mentioned the -- well, the 1311 and, really, as recall the numbers, it reached over 12,000 that were really shipped.

Porter: Of the drives.

Carothers: And then the 2311 had something like 25,000 drives and I've lost track of the 2314, which far exceeded that, so...

Porter: Well, that came along just a bit later and -- but -- with a larger disk pack.

Carothers: Right.

Porter: But the 2311 also was very successful in the industry for quite another reason, probably not regarded by IBM as a good thing but the plug compatible drive makers, which were looking at all of this huge success out of nowhere, suddenly decided, gee, I could make a disk drive like that and probably

use that same disk pack, too. And Memorex was the first, the Memorex 630 drive, in 1968, had, if you will, a carbon copy of the 2311, plug compatible version of the 2311. Became the first of what became known in the industry as the plug compatible drives.

Carothers: Correct.

Porter: And they brought out another one, you mentioned the 2314 later that year but the 2311 became the first of the major disk products, disk drive products that developed a pattern of imitators. And, out of that, came basically the whole disk drive industry, which followed, as we knew later. And also the imitators of the single disk cartridge drives, also, became really quite profuse in what they were doing. So what we're talking about here has been the creation of basically much of the disk drive industry and it's been very interesting.

Carothers: That's right.

Harker: Well, it was a -- the basic pieces of the drive were fairly simple. The technology of the heads not so but, basically, the hydraulic actuator, if you saw one and took it apart, it wouldn't be hard to build one. So it led itself, if you will -- I never saw anybody try and build a IBM 407 printer. That would be a total bear and, in fact, most -- even the punch card equipment tended to be a lot more sophisticated mechanically. So we gave the base. <laughs>

Porter: Developed the pattern and, as we have said, a great amount of the whole disk drive industry followed as a result of that.

Harker: And there has been tremendous advancements.

Carothers: Right. Oh, it was well engineered, I'll say that. I mean, when we got into it and started discovering all these parameters associated with interchangeability and everything and we were working as a real group of people to try to understand all those variations and every element of the thing, tie it together, make certain that those things would work that way, and then we ran it through tests. I mean, we have a product test organization that was tough on us. They didn't give us any leeway. They really made you live up to your spec. When we got in this problem with contamination, what did that lead to? That led to a room in which they generated dirt and they put it into the air. People went in that room with gas masks <laughter> to protect their lungs. And they were trying to make the doggone thing fail. <laughter> And they couldn't do it, of course. But we really got tested. On one occasion in the earlier coating that I was talking about that transferred itself around the disk, they had introduced this high humidity hot temperature within the chamber and, when we discovered that we did have a disk failure, I mean, this one, now, I mean, there were other ones were already transferring this kind of a problem, but

they pulled the disk and we looked at the track and, my gosh, under a microscope, you saw the oxide coating but it was raised and it was in the crystalline form. Crystalline form in our coating? Yeah. <laughter> But what had happened is it'd been making contact and, because -- and what had happened is that the chamber had to have soft water and, in generating a softener, they got the wrong, improper amount of salt that was introduced into the air. Well, how does that head work? It compresses the gas in front and it cools on the outtake of the head. It deposited a salt crystal. <laughter> Sodium chloride was right there in crystal form. <laughter> Just an interesting factor of how we looked and how we had to detect when we had a crash. We went back and looked for particles. We had an analysis on each and every one of these things.

Porter: But the bottom line was that you did do sufficient correction so that, even though a few may have failed in the field...

Carothers: Exactly.

Porter: ...it had a low enough failure rate so that it was...

Carothers: Absolutely.

Porter: ...<inaudible>

Carothers: And we also, I mean, looked and identified our reliability. We had to because we were concerned about the quantities of machines going to the field. Field engineers had to have the right amount of people to maintain these in the field. When you're a success, you just don't send out a machine and hope the customer's going to take care of it. It was field engineered. So things like this pig tail that you see here on the head? That -- one end is stationary and the head keeps going back and forth. We ran special tests to make certain that the pig tail would last under all those operations, hundreds of thousands. We identified the range of power on hours and so forth. So our components that we suspected were really put through tests. We even had a pack changing test whereby it was constantly changing packs and you'd run it and stop it, another pack change and we went through thousands of pack changes to make certain we could do that without damaging the reliability of the product.

Porter: Well, the result was it worked and the result of all of that was that it was so successful that IBM had such an influx of orders in the middle '60s, they couldn't keep up and the customers were faced then with several months' delay in trying to order these products because they were so successful. It developed quite a backlog. I think we've covered the ground today. Gentlemen, I would like to...

Carothers: Thank you.

Porter: ...thank all of you for my ability to discuss this with you and I am sure that, looking backward on it, you're glad that it's all in the past but it's a past that you like, I'm sure.

Harker: Yep.

Porter: Thank you.

END OF INTERVIEW