Oral History of IBM 2314 and 3330 Disk Drive Panel

Interviewed by:
Jim Porter

Recorded: June 16, 2004
Mountain View, California

CHM Reference number: X2800.2004

Jim Porter: The objective for this discussion is to discuss the products which became known as the IBM 2314, and along the way, the 2319, the 3330, the controllers that went along with all of this and, oh, by the way, the creation of the whole plug compatible disk drive industry that started during all of this with IBM's competitors. The 2314 happened at the time that IBM changed the whole computer industry by introducing the System 360. It became an industry that was known by a lot of people outside of IBM as "IBM and the seven dwarfs". This is not to say that IBM had as clean an existence as Snow White and the Seven Dwarfs did but, with the introduction of the System 360 and its ability for large users to have everything from small through mid-size systems, the ability to expand, the ability to move things back and forth throughout that whole system, a completely new idea among computer systems. This made IBM the giant and the storage that went along with it, there were tape drives, of course, and disk drives and IBM needed a larger capacity disk drive than the ones they had in the past. They had already moved to 14" disk drives early in the 1960s and, when the System 360 was coming out in the middle of the 1960s, they realized they needed something with more capacity. So along came the 2314, which, in many ways, was a product which evolved from the drive they had earlier, the 2311. So we'll get into those discussions in just a minute but first let's talk to each of these people that are here with us today, find out why they're here. Starting with Hal.

Hal Yang: I'm Hal Yang and I joined IBM shortly after graduating from MIT in 1964. I was with IBM for three years, during which time I worked on the 2841 control unit, which, with slight modification, became the IBM 2314 control unit. I started doing hardware logic design on the 2841 and evolved to doing micro programming on the 2314. I left IBM in 1967 to cofound a disk drive company called Information Storage Systems and I was with Information Storage Systems until 1980. ISS was one of the more successful plug compatible manufacturers and, in 1980, I assessed -- had grown to 4,000 employees and was shipping a billion dollars' worth of disk drives at 75% gross margin.

Jack Clemens: I'm Jack Clemens. I came to California in 1957 to join IBM. Spent 13 years there in various capacities. The last position I held was the 3330 program manager. I left IBM to join Telex, a competitor, to develop a plug compatible version of the 3330, a dream that was never fulfilled. I moved on to doing consulting work with Control Data and several other companies and then, later on, joined Memorex to -- in their storage products area and to be around when they shipped their first unit of their 3330 product. I left Memorex, again did some consulting work and, subsequently, worked for several companies, including Datapoint, LaPine Technology and a company we started called ComPort. I finished my career as a consultant for about ten years and I retired two years ago, three years ago.

Porter: Bob

Bob Pattison: My name is Bob Pattison and I joined IBM in 1951 and I worked in several different assignments. In 1958, I joined the group headed by Jack Harker designing a 14 -- first 14" diameter file. And that was called the 1311 and I worked on that file program and I worked on the 2311 file program

and managed the file assemblies that go into the 2314. And, let's see, after that, I had quite a few other different assignments in the company and I ended up as an IBM Fellow.

Porter: George?

George Santana: My name is George Santana. I joined IBM in 1959, after completing a few degrees from Stanford University in electrical engineering. My initial assignment was in doing several mechanisms work; however, very shortly, I was reassigned to work on a problem upon the ADF file, the recording electronics problem on the ADF file and, at that time, designed a automatic gain control system for it. The rest of my career, then I didn't work with servos. I just continued with the recording electronics and later became a manager in that group and worked on developing a electronics recording system for the 2314 and, later on, the 3330. Then I did add the head and disk and became manager of file technology, which had responsibility for all of the recording technologies that were associated with disk files at San Jose. I basically continued with that work for a number of years. Also, had a staff assignment at headquarters for a couple of years and eventually retired in 1993 as senior technical staff.

Don Massaro: I'm Don Massaro and I went to IBM in 1967 out of Berkeley, and my initial assignment was to work on a simulation program for the airbearing system on the 3330, which I did. In '69, I left with Al Shugart to go to Memorex, to put Memorex in the disk drive business and I ended up being the program manager for the Memorex 3670, which was the plug compatible equivalent of the IBM 3330. Got that product to the marketplace and then, in 1973, Al and myself and a couple other engineers from Memorex left and formed Shugart Associates, which created the floppy disk drive market. And ended up selling that to Xerox in 1978, and that was the last I worked on disk drives.

Porter: Thank you. I guess I should mention that, by way of a little background on myself, I went to work for that company that's been mentioned a few times, Memorex, in 1968 and I took over product management on disk packs. Now, a disk back is similar to what you see here. This is a 2316 disk pack, which was removable, and was used on the IBM 2314 disk drive. And the one alongside is a 3336 disk pack, which was used a few years later on the IBM 3330. I used to know that market quite well in that I was very interested in selling all of the competitive disk packs to this one made by Memorex and I think we had the leading market share, actually. We were delighted by what was going on in that era in the computer rooms, which was called “save and restore”, which is a term you hardly ever hear in the computer industry today. But, because the 2314 had a capacity of 29 megabytes, it was very large compared to what had preceded it. But there were still some applications which there was not enough total capacity on a computer system to do, so that, if you wanted, for example, to post all your employee records on Thursday, you could then write all the checks for employees on Friday and then take the disk pack off which had all that data, put another application on, and that was called save and restore. The next Friday, if you had to issue paychecks, you could put your disk pack back on the drive, which had all that employment data, post the employee records again, do the checks again on Friday and take the data off. Save and restore. And -- which meant that there was a lot more than one disk pack used per drive.

So, as the guy who was selling a lot of those in competition with IBM, we thoroughly enjoyed that phenomenon, and we were watching it very closely. We mentioned, of course, that, in this whole era, the independent disk pack, the independent peripherals manufacturers, the plug compatible manufacturer, wave occurred. Memorex was the first company to make plug compatible disk drives to IBM's. The first was a copy of IBM's 2311 and then the 2314 and these, depending on who you believe, shipped in the-- were announced in the '67 timeframe and shipped in '68, up through '69, et cetera and then on. But they were followed by a wave. Hal mentioned that he was one of the founders of ISS. They, shortly thereafter, were also shipping 2314 equivalent drives which were then sold by the company that Jack worked for, went to work for later, and they resold those drives. I don't think ISS was much of a seller of disk drives in the initial...

Yang: No.

Porter: They were resold.

Yang: Yeah.

Porter: But a wave of companies started, not just Memorex and ISS, but Control Data and Century Data and Marshall and Potter, et cetera, and, out of that, came what became known as the whole disk drive industry. But let's go to the 2314 and talk about how it occurred and, as I think we indicated earlier, it was essentially an offshoot, an elaboration of what had been done on the 2311 and perhaps, Bob, since you worked on both, if you could talk a little bit about that transition and what was going on.

Pattison: Well, it was, as you said, it was because we needed to have increased capacity for the large systems for -- on files, disk files and we -- in order to do that, somebody had to come up with a concept of a 2314, which had nine disk files on it and eight of which were available to the control unit and one of which was a spare. And we could put the control unit, the spare disk into -- we could exchange -- the customer could exchange the spare disk to the data disk if they had a file -- need for the file changes and we needed to get increased capacity. And, as you said, we had 29 megabytes per drive and that gave us-- the total nine file program gave us a lot of extra capacity, total capacity for the drives.

Porter: Any other comments? George?

Santana: Well, I think, coming from the recording technology point of view, one of the innovations that came when this technology was going to a ferrite core, the read-write core, and we needed that because data rate was very high, at least we thought it was very high in those days. I think the data rate was 300 kilobytes per second, something in the order of...

Yang: 312.

Porter: 312.

Santana: ...which translated to about 2.4 megabits per second. And people in the recording channel area and the head area were worried about the serial data rate rather than the parallel byte rate. So that they, in the past, the 2311 and the 1311 files had used a stack of high view metal to form the core. Now, we needed the higher frequency capability of ferrite and so we had a ferrite head and also, to go with the expansion and contraction characteristics of that head, we needed a ceramic material for the slider rather than a metal material, which we had previously used. So the 2314 head was a ceramic material and we epoxied the ferrite core into that and I believe we had a glass gap, the gap dimensions for the head weren't glass but the potting was still done with epoxy. So that was-- from a head point of view, that was changing technology. The disks became thinner. That was as we advanced the aerial density, bit densities, why, we generally made gaps smaller, made spacing smaller and also made the thickness of the disks smaller. And so those, I think, were part of the main areas of the technology challenge.

Clemens: What about the thing I think you're going to get into later on, Jim, was the 1311, 2311 were some of the most profitable products that IBM ever produced. I tell everybody I've been in the disk drive industry to remember when it was profitable. They were attached, in various ways, to systems. The 2841 controller was one of those. There were some native attachments to other systems. The concern, when the 2314 came out, was that it would wipe out the 2311 business and that's why it was bundled into this nine-drive package with a controller. The only way you could buy it was nine drives at a time. As Bob pointed out, there were some plugs where you could reassign drives. If one, for some reason, failed, you could reassign the drive to another position by moving this bull's eye and move the disk pack and continue on until a repairman could come out and fix your disk drive. So that the bundled package for the 2314 and the individual small package for the 2311 were pretty key in the-- and very set in the IBM planning department and business planning department.

Porter: Well, let's just point out the facts on some of this that there were actually nine drives in each of what was called Direct Access Storage Facility, DASF, but only eight were active. One was a spare.

Yang: Well, there was a unique innovation from the controller. Now, the controller, in and of itself, was actually a powerful computer. As a matter of fact, its core was the same as that of the mod 30, model 30 processor, which was designed at Endicott. And so the controller could actually be time-shared to run diagnostics on that spare module and these diagnostics were called inline diagnostics and you could perform rudimentary functions in terms of poling status. You could access, make the actuator seek back and forth. You could actually do reading and writing on a special designated diagnostic track. So you could actually perform some rudimentary diagnosis from the control panel of the controller on that spare drive while the facility was actually running, doing its main business.

Porter: Well, again, let's point out that, in the first year or two, the only way you could buy the 2014 was in a set of the nine drives with the controller.

Yang: Correct.

Porter: And that the selling price-- not too many were sold in those days, they were mostly leased, but the price for that set of controller plus the nine drives was $252,000 in the beginning and the monthly rental, which was the more common way of doing this, was $5,250 per month rental. And the disk pack itself, in the beginning, IBM would rent a 2316 disc pack for $20 a month. If you bought it, it was $650. In other words, they'd found a way to make a little money on all this, hadn't they? In the first few years, I don't know the numbers, but it would appear that IBM had a very profitable business on the 2314.

Clemens: IBM's pricing philosophy on systems at that time was give away the CPU, give away the software and sock it to the customer on the peripherals. The peripherals were extremely profitable. The CPUs themselves were break even. Software was still, I believe, given away free at that time.

Porter: And since this started shipping, let's see, it was announced, along with—part way through the System 360, it was announced on April 22, 1965, the first shipment by IBM was in December of 1966. But the plug compatible drives didn't start shipping for two or three years later. So IBM had two or three years of a clear slate to charge whatever they felt like charging and, in the meantime, the System 360 was becoming the computer used in all the Fortune 1000 companies. So it was quite a landscape they'd moved into with all of this, isn't it? Actually, competition did arrive, though, with all of those plug compatible companies, including ISS and Telex and Memorex and, eventually, IBM did decide to do some competitive things along the way, didn't they?

Clemens: Yeah.

Porter: I think the number 2319 sometimes comes out in this discussion?

Clemens: There was nothing more passionate in IBM executive suites in the 1967, '68, '69 period than the 2319. Technically, there was nothing to it. I believe, I don't think there was any difference.

Yang: Well, there was the dreaded price reduction where they basically started shifting-- raising the price of the processors and cutting the price on the peripherals to cut into the PCM manufacturer's margins.

Clemens: But, technically, the 2841 business and 2314 controller were...

Yang: They're exactly the same.

Clemens: ...easily isolated. And the ability to buy a 2314, basically, with as many spindles as you want was where it finally wound up.

Porter: Let's just point out that IBM...

Clemens: The more profound impact was the pricing change that occurred at that same time.

Yang: Right.

Clemens: When they reduced the price of the controller.

Porter: I think we should point out that, as compared with the 2314 DASF, which was the nine spindles altogether, with the 2319, of course, you could purchase various numbers of spindles along with that controller. And...

Clemens: The drives were physically packaged -- the one isolated drive was a spare but the others were all in pairs over and under configurations so you'd buy them in one, three, or I don't know if one is a possibility, three, five or seven or nine were the package configurations. Whereas the ISS product, you could buy one, two, three, four, five.

Yang: Right. Individually.

Clemens: They were individual cabinets.

Porter: And some of the plug compatible drive companies didn't seem to have a very good response to this except to hire lawyers, did they? <laughter>

Clemens: Well, I don't remember -- I don't know the legal machinations. Turned out that the day I left IBM, I was invited to go back to New York and visit with some of the executives and chasing me through the halls of White Plains was an announcement package for the 2319. One of the questions I had asked the executives at Telex, as a matter of fact, when I took the job, was are you aware that IBM is going to, you know, break the bank and break up the 2314? They assured me they were. They assured me that their business plans would accommodate it and so on. And the week after I got there and the

announcement came out, they just about died. <laughter> There wasn't a one of them that really believed it.

Porter: Well, by that time, with production levels way up there, IBM's internal manufacturing costs must have shrunk down quite a bit per unit and they were then able to bring in the 2319 at levels which the plug compatible drive makers found it very difficult to compete with, didn't they?

Clemens: The plug compatibles, with the exception of ISS, had no controller so the controller was still the IBM box went in, even though they had Memorex drives or Telex -- Potter or one of the other.

Porter: Potter, Marshall.

Yang: Marshall.

Clemens: Yeah. The Marshall.

Porter: Control Data.

Clemens: And so the 2319 would mean that ISS was the one who really had pushed them over the edge as far as the controller went. None of the disk drive companies knew anything at all about controllers. I don't know what Memorex's capability was but I don't think it was very much.

Massaro: We <inaudible>. At Memorex...

Clemens: But, for the 2314 level product, I don't think there was any controller.

Massaro: 660. Yeah, we'd be -- right now, we did not but, you see, back in 1969, Memorex's biggest customer is RCA so we were selling the 660 disk drive, which was the Memorex equivalent to the 2314, to RCA. Of course, RCA was in the mainframe business. Actually, IBM was very concerned about RCA because RCA had a fairly good machine back then. So that was really where the bulk of products were. As a matter of fact, my first year and a half at Memorex, even though I went there to put Memorex into the 3330 business, was… <inaudible> …and try to fix that product so it would work on the RCA machine.

Porter: So tell me what it was like, internally, at these plug compatible drive manufacturers when this sweat was going on. Did they think they had any hope of competing against the 2319 era with IBM? Of course, I should point out that, with the movement of all of these IBM employees, as IBM came to call the

12 people who left to start ISS, we should mention, I don't think it's been mentioned here, that the in-house IBM term, of course, was the “dirty dozen”. And it became known universally throughout IBM and the rest of the industry as the dirty dozen who started ISS. So did the dirty dozen at ISS feel they could compete against the 2319 and the folks at Memorex and the folks at Telex? The 2319 era from IBM, when those prices hit the street? How did they feel about competing?

Yang: Well, I know that the impact it had at ISS was that the management team there decided to run for cover. <laughter> And the way they ran for cover was they went to affiliate themselves with someone that they believed, at the time, had greater management experience, which was Itel Corporation. And the merger between ISS or the acquisition of ISS by Itel, was a direct consequence of ISS's management deciding that they didn't know how to steer the ship, given that the waters had suddenly gotten very, very choppy and murky with IBM's price cut.

Clemens: I don't know. I was new with Telex and didn't know the people there well enough to really say what the pulse was. My impression was that they were-- the expression I used at the time was most people who tried to get into the computer business were sort of like backing into a propeller blade. Most people got their fannies chewed off. Telex got caught on the propeller, was catapulted into the stars. They really never understood the computer business. They really never understood what it was they were doing and it was very hard to judge their reaction. They basically continued, you know, after the shock and everything else, continued just doing what they had done before. That's when Telex went over the hill and really went out of business.

Massaro: If I remember, at Memorex, this was not a big deal at Memorex at the time because most of the effort was on the 3670 program, which was the controller for the 3330. We were building our own controller which was going to hang on the channel so we didn't have to worry about IBM locking us out of that marketplace. You see, it really wasn't a pricing issue because all they actually did with that product was kind of brought the price per spindle down to where the plug compatible guys were. The problem was, if you didn't have any controllers for market entry, you were kind of locked out of that marketplace. So, since we were building our own controller and our own disk drive and we were basically selling product on a per spindle basis for what IBM dropped their price to, it was not a big deal for Memorex. As a matter of fact, I can see many ways it helped Memorex because it scared the hell out of a lot of the other plug compatible guys.

Porter: Except that Memorex did initiate a lawsuit.

Massaro: Well, yeah, but that's just marketing. <laughter>

Yang: So did Telex.

Massaro: Everybody did, right? Well, what they did is they tried to lock it out of the market. It wasn't because of pricing. It's they simply trying to keep you off of their integrated -- I think it was called integrated file adapter, at the time, your IFA.

Clemens: That was one version, yeah.

Massaro: Yeah.

Porter: Well, with that, let's move on. The integrated file adapter and the whole controller area becomes very important as we move on to discuss what IBM was developing during that same period, of course, which was -- using the IBM code name and, of course, these IBM projects mostly had code names and the 3330 was the “Merlin”, to use IBM's code name. The Merlin, which was announced in June 30th of 1970 and finally shipped in August of '71 and what became known as the 3330 mod 1, which was 100 megabytes per disk pack, later was followed with another IBM project, code named “Iceberg”, which was the 3330 mod 11, which had 200 megabytes per disk pack. And that wasn't announced until July of '73 and it was shipped in March of '74. But, in the period that we've been discussing here, which was basically the late '60s, IBM was deeply involved in the preparation of the Merlin and it was, I can tell you, because of all this IBM employee movements over to those other companies, the existence of the Merlin was certainly no secret in those other companies. I'm sure ISS knew about it, we knew all about it at Memorex, et cetera. When, as you mentioned, Al Shugart moved over from IBM in '69 to Memorex and brought whatever number of people it was, whether it was dozens or hundreds has always been an area of contention. The amount of information about the Merlin which accompanied them was really quite large. Subject to, by the way, to some of the lawsuits. And so this higher capacity drive was on its way but equally important with the drive was what was happening on the controller side and also on the-- some of the business side. I should mention that, in '71, about the same time that the Merlin shipped, just before it, IBM had come up with what they called the fixed term plan. Now, prior to this, with the 2314, those we called the direct access storage facilities, the DASF, the 2314s, didn't go out on a fixed term, typically. They went out on so much per month. But IBM changed when they went to some of these other, newer configurations and changed it to a fixed term plan. I think two years was one of the versions so that you could get a lower rate but -- as a user but you had to agree to a fixed term, which was a new thing. And that evolved into then the extended term plan which came along in '72, more of the same. So the environment was changing but the ability of the independents, because of the change of the controllers, and this included moving all or some of the controller functions off a separate controller also onto some of the CPU processors so that, in some situations, you could get along without a controller. You were deeply involved with this controller, Hal, what was going on on all those things during that period, weren't you?

Yang: I was but mostly, at the time, ISS had been acquired by Sperry Univac and most of my energy was focused -- we had a separate group that was doing the cloning of the 3830, which was, from a generic standpoint, a slightly faster 2314 controller. But, from an architectural standpoint, wasn't

significantly different and most of my work was going, at that time, towards developing a 3830 equivalent for Sperry, which we hoped had better system performance. And so I don't really have much recollection in terms of how we were responding to the IBM compatible-type issues because I don't seem to recall that there were very many at the time, at least not from an engineering standpoint.

Clemens: I think one of the biggest beneficiaries was Control Data. Some of the work we started -- when I was hired at Telex, we established the Telex direct access storage division out here on the west coast, hired some people to do a disk drive and a control unit. The disk drive program was terminated, either before or after I left to go to Telex. I was only there about six months. But the control unit work went on and then, eventually, wound up in the hands of Control Data, who really -- Control Data, at that time, had excellent disk drive capability but no plug compatible capability, and it was that controller that gave them the bridge and Control Data really emerged as a major plug compatible competitor because of that.

Porter: Other comments about the controller world?

Massaro: You know, I think the controller was key, though. That was the only way you could guarantee that you weren't going to get locked out in the marketplace by IBM if you developed your own controller. So, as I had mentioned before, we basically bought and were the very first 3330 systems for IBM, which was a controller and the eight drives. And we split it apart and I took the controller into my group because I was developing the disk drive and the gentleman, Larry Pyle, who was developing the controller, he took the disk drive over there. And we each took the other manual so I ended up with a controller and a disk drive manuals. He ended up with a -- these were the published manuals and he ended up with the controller manuals and the disk drive. And, basically, for a year, year and a half, worked alone, my group developing the disk drive, his group developing the controller. <inaudible> communicate with each other. And then, when we were done, we disconnected the IBM controller and the IBM disk drives and plugged 3670 controller into the 3670 disk drive and, within a couple of weeks, we got it to work. And so that's a way of speeding up getting the product to the marketplace to compete against IBM and also to make sure that they don't lock you out because it would have been very difficult for them to now take the controller and integrate into a blocked channel because now this gets really complicated. So we actually were in a good position to compete against IBM because we had that controller expertise.

Clemens: One other comment, I guess, related to the controller was that ISS, for all their technical genius, I think signed a real lousy contract with Telex as their marketing division and gave Telex an exclusive right, not only to the 2311, 2314 products but to future products or at least the 3330 equivalent. So when ISS tried to take some of their technology to the marketplace directly through Itel, I guess, Telex jumped onto the lawsuit and pretty much stopped that and it was about that time, I think, that you guys went to Univac because Univac had their own systems capability and they sort of redirected their activity away from the plug compatible business because of Telex's interference.

Porter: Okay. But the disk drive itself should not be ignored, the 3330, the Merlin. It incorporated some things that were new at the time. It had a track following servo, rotational position sensing. Used a voice coil actuator in combination with all of this and, tell me, what were the important technology innovations that came along with the 3330?

Porter: George?

Santana: The -- well, the 3330 did contain a voice coil actuator. It used a separate servo surface. In other words, you still-- they're not closing the loop on individual data tracks. You are closing the loop on one given surface, which, as I recall, was in the middle of the packs to minimize tolerances. And so that was certainly a great innovation over what we had in the past. We went from a hydraulic system to an electric system which was a much more computer compatible type of technology. The heads and disks in this product were newly developed, our disk for this product. We went both to a higher bit, double the bit density and double the track density over the 2314. In the head area, they continued with the use of ceramic material. They now got -- and ferrite core. They now used also glass bonding and the importance of that is that the epoxy that they'd use in previous heads still had a certain amount of (inaudible), which tended to -- allows the signal to degrade and so, by using glass bonding, we were able to keep the head flush to the surface of the slider and, therefore, as close as possible to the surface of the disk. The newly developed disk was made and we had a program called the Merlin Disk Company that worked on developing that. It was about that time that we really began to realize the importance of development people and manufacturing people working together. And so the Merlin Disk Company was a group of engineers from the manufacturing engineering as well as development engineering working under a single management. And it was, actually, it was a dotted line arrangement but it was a pretty solid dotted line arrangement. And they worked together developing the cross Fs, chemistry, going through the chemistry. There was a lot learned about the coding process, about the importance of alumina in the certain end, the coating of the disk, to give it as much strength as possible. We did -- although these products used heads that took off, they had loading and unload systems, there were still periods where, when you were loading the head, you had some contact. So you wanted to improve the durability of the disk so that you didn't have crashes, you just had instantaneous contact. And, in the recording channel area, we developed clocking circuits and amplifiers to go with this higher data rate. Now, it was up to 800 kilobytes per second, which translated then into 64 megabits per second. So the clocking windows were getting smaller and data rates were getting higher. Those all kept us very busy on computers.

Porter: Going to throw a few numbers at it to compare. For example, the 2314 had used 20 data surfaces there for data and the 3330 used 19 and one surface was strictly for servo.

Santana: Right.

Porter: So all this was done on 19 surfaces compared to the 20 used on the previous one and, as you pointed out, the bits per inch and tracks per inch had basically doubled. And the performance also had changed quite a lot. The 2314, of course, spun disks at 2400 rpm, and the 3330 at 3600 rpm, so it went up 50% on spin rate. And the seek time, originally, the average head positioning time for the 2314 started out at 75 milliseconds average seek. Later versions of it went to 60 milliseconds but then, on the 3330, the average seek time went down to 30 milliseconds. So, in addition to having four times the capacity, basically, 100 gigabytes, performance was also greatly improved.

(Gap in original tape)

Porter: Let's talk about the general summary. What are the most important things that we've learned and what has the world learned from the development projects on the 2314 and the 3330? Going back to the 2314, what did we really learn from all of that, that was really useful to later times? For example, that business about airborne contamination.

Pattison: That was a problem, yes. The problems that we had in developing the program mechanically was the fact that, when we got into the product test cycle, we started having head and disk problems and it was difficult to find out what on earth was causing these problems. And finally they figured out that it was the cigarette smoke building up on products, building up on the slider that then caused the thing to crash. And the only thing you could do -- and what's done from that point on is we had to put an absolute air filter in every drive and we had to put a fan on there and we had to put cowling around it enough so that you could keep the air very, very clean all around the head and disk areas. And, from that point on, all the rest of the group of files had to take the same kind of activity to prevent the contamination.

Santana: That smoke that Bob referred to, smoke contamination, we then developed a test for that to assure ourselves that the head and disk would work properly and we had a system in which several cigarettes were puffed into this head/disk combination and we had to fly through so many hours, I forget what the numbers were, to satisfactorily pass the test. And I think anybody who saw the mess that that made, where they were actually doing the testing, the buildup of brown gum and so forth would have stopped smoking right then <laughter> because it looked terrible.

Porter: Did you have a machine which automatically smoked these cigarettes...

Santana: Yes. Yes, there was-- it was, you know, it was a jury-rigged deal but, yeah, we had a machine that would do it and would have to do a certain number of cycles in order to pass the test.

Clemens: There was a story circulated about a purchase order that went out of George's area. There was a programmer looking at possibly a drum file and somebody found out that beer cans were very

symmetrical and they were going to try to coat beer cans or something like that and make them into a drum. Another activity was ferrite production and they use a condom to hold the powder while it was being-- put it in the furnace and heated...

Massaro: That's making the ferrite.

Clemens: Yeah. Supposedly somebody wrote up a purchase order for cigarettes for the smoking machine, some condoms for the ferrite thing and beer cans for the drum program and this purchase order bounced, you know, what the hell are you guys doing... <laughter>

Porter: And you knew you were having a lot of fun doing disk drives, didn't you? <laughter>

Santana: But George is right. This contamination chamber was awful to have to go into. The smell and what have you. Ugh. I was a smoker at the time. I can't believe what we did, you know? Work on a disk drive with a cigarette hanging out of your mouth, every conference room had, what, six ashtrays heaped with cigarette butts. Unbelievable.

Porter: What other lessons did it-- were derived outside of that cigarettes were bad?

Clemens: Hygiene.

Massaro: I think one of the lessons we learned is that, when you move from one generation to the next, you just don't advance one area, you advance all the areas. So you go from 2400 to 3600 rpm, okay? Now, all of a sudden, you've got more energy in the system, now you have to worry about damping the vibration products, okay? And then, in order to record at the higher density, you got to go to a lower flying height. To go to a lower flying height, you need to have more loading on the head. Now, all of a sudden, with the high loading on the head, when you launch the head onto the disks, now you've got other considerations in terms of those things that create head crashes. So, you know, everything goes up by a factor of 10, just because you double the aerial density, okay? It's not that you double the complexity, you might have gone up by a factor of 10 in the complexity because of how you doubled that. So, you know, you just learn to conquer those things so, every time we double the density, we're not taking on a two times problem, we're probably taking on a 4X or and 8X or a 10X problem.

Santana: I think, you know, the use of a closed-loop servo system also added a lot to these resonance products because you were dealing now with an energy source that was putting energy into the system, the servo system. You were trying to close the loop over a complex structure and so any vibrations that you had, there was an opportunity for the servo to amplify them and so it was very important to control that whole frequency response, recognizing that you had a closed loop system.

Clemens: I think, from an engineering standpoint, too, we first had to wrestle with the fact there was no place to hang your hat. You couldn't make a measurement from anything. Everything -- the servo track wasn't exactly any place and the head wasn't exactly any place. Everything was moving and some of that movement was valid, some of it was invalid. It became a lot of techniques and stuff. The resonance work you guys did and the finite analysis work they did was just marvelous and we'd never done anything like that. The company had no facility at all or that.

Massaro: See, the 2314 was an open loop servo and you had a detent, so you set it to a track and the thing would fall into place and catch the rack and it'd stop there. So you had a reference point, okay? 3330 didn't have that. You had these tracks written on a disk and you had -- the servo head would go track to track and have the track follow it and it was really a -- now, once we solved that technology problem, moving from the 2314 to the 3330 and now we had some track following servos, then it was relatively straightforward. But that was a major, major transition point. I think closed loop servos was clearly a disrupting technology.

Clemens: Big step. Big step.

Santana: But, of course, it was the future.

Massaro: Exactly. If nothing else, you get that hydraulic fluid out of there.

Porter: And, while some of this was going on, another group of people were working on later products such as the Winchester, the 3340, leapfrogging and taking advantage of the work that had been done on these products.

Pattison: That's right.

Porter: Which paid off in later generations.

Massaro: You know, I almost think, though, that going from the 3330 to the 3340, when we put the heads in the pack, actually made it easier than going from the 2314 to the 3330, okay? You took a lot of tolerance -- it created some other problems but you took a lot of tolerances out. The biggest problem was interchangeability, running these interchangeability tests was key because you got to pay us 10 to the ninth, 10 to the tenth, which would take forever to find out if you have any interchangeability problems.

Porter: For the pack going out (inaudible), yeah.

Massaro: <inaudible> …another machine and so forth and so on. And by putting the heads in a pack, it's almost like a quasi-fixed disk drive where you didn't quite have to worry about interchangeability. So I think that actually helped a lot, going to the Winchester, from the mechanical standpoint. It created a lot of problems for the head guys but, for the mechanical engineers that were designing these guys, boy, at least that pressure's off my back. Now it's their problem.

Santana: Yeah, with interchangeability, you had to be concerned about mechanical tolerance. You had to be concerned about the adjustment of the head. When we had this head on assembly was put into the actuator, you had to adjust that so that it was some nominal position because you were going to have interchangeability and another head from another drive, another disk pack would be in there looking for the same track. So Don is right. It really -- the switchover to Winchester, where all those things were just left alone, you didn't have to worry about adjusting the heads. Just put them in there and, if they were off a little bit, it didn't make a difference because the whole track would be off.

Massaro: You realize that, if you look at the tolerances on a disk drive, and you do a worst case analysis, the product would not work. What you depend upon is that not all the tolerances are going to go worst case at the same time, if you do it statistically. And that was fine when you were building 1,000 or 2,000 or even 10,000 drives a month. But now, when you started building 10,000 drives a day, which we did when we got into the floppy, you literally had to throw products away. They were called these hanger (inaudible) and these were the ones that kept going through the line over and over again and you could never get enough and you literally had to throw the disk drive away, because you were building so many disk drives, you actually got to the point where, statistically, you could have one or two of these "worst case disk drives" in a batch of 1,000 or 10,000.

Porter: And I would point out that, while we're still talking about the things that were great advances in connection with these programs, what did the industry get out of it? One of the things they got out of it was a diskette. We haven't really talked about this but, to load that microcode into the controller for the Merlin, IBM did the first usage of an eight inch diskette to load that microcode and it wasn't until about three years later, in 1973, that they brought out the 3740 “key to diskette” system, which set the standard for the whole industry. That they first actually made commercial use of an eight inch floppy which proved out the technical feasibility was to load microcode in that controller and -- for a couple of the processors on the system model 370. So the diskette was another thing which came out of the 3330. And production of diskettes got up into more than five billion a year eventually, and...

Santana: Probably smaller than that.

Porter: ...couple of hundred -- they evolved into smaller sizes as all the hard disk drives did, too. They all evolved into smaller sizes.

Yang: As have computers.

Porter: As have computers.

Massaro: There was no technology relationship between that diskette that was designed as an IPL device, for program mode, and the diskette drives and the diskettes that ultimately were put into mini computers. You would design that to maybe do, in its lifetime, a million revolutions.

Porter: Oh, yeah.

Massaro: We have to do a million revolutions in a day on this stuff or better. Although the form factor looked the same and you could see how it would go back -- the technology investment by taking a program load device and turning that into a system resonant device where you had the operating system on it and you were doing millions and millions of seeks and millions and millions of rotations, totally different.

Porter: But it was the start.

Massaro: The form factor, it was the form factor.

Porter: And it made it happen.

Massaro: Yeah.

Porter: So, out of the 3330 and its era and its controller came a lot of things. Anything else of significance there?

Clemens: Yes. One of the things we haven't talked about is error correction codes, which were very significant, still in use today and really some pioneering work. Error correction codes were known in the communications industry but the requirements for a disk drive and storage were considerably different. We found that out very late in the game. When we were in product testing the 3330, we suddenly discovered that certain data patterns gave an error rate not one in 10 to the 27^{th}, or whatever they were supposed to be theoretically, but they were, like, one in 10 to the 500th. It showed up quite rapidly and we had to redesign the error correction code at the last minute to solve that. And it has to do with the patterning of data and the fixed length of the information. You were asking before about what did we learn from all that. I mentioned, during the break, about the fact that the prize was some market that was

some $2 billion, which made the 3330 the largest commercial program in the history of civilization. <laughter>

Porter: Would you repeat that? <laughter>

Clemens: Well, its, in all humility, yeah. No, it was a huge, huge program. It was valued at more than the entire 747 program to Boeing Aircraft. IBM had, let's see, on the 3330, at the peak, I think we had 250 engineers that had worked for me and all these technology groups, the manufacturing engineering and so on, that were involved and so on. So the number was probably closer to, you know, 500, 600 people who were involved with this. We had a celebration party after the announcement and invited 300 people and there were just equal in number of people at IBM angry because they hadn't been invited, because they were involved in some way or another. One of the things I learned from this and the other program was -- oh, very quickly, when I was at Telex, months after I left IBM, before the first 3330 shipped, a guy from a southern California machine shop walked into my office at Telex, brought in this big box, opened it up and there is a damn actuator, wires off the voice coil, carriage, you know, heads, the whole thing. The 3330 kit. Not like IBM's but a mechanical version of it. And I realized that, you know, the guy that builds the second airplane has got a lot going for him as opposed to the Wright brothers, because you know it can be done. And how quickly the technology can be replicated, copied possibly but duplicated in a useable form. Knowing that it can be done, or perhaps suspecting it can't be done, is a great impediment to the engineers who were working on it and I think that was part of the reason why this program lasted five years. But the industry grabbed the 3330. Same thing happened with the floppy disk. The first one these guys struggled to get out at Shugart and so on and, what, three years later, anybody can build a floppy disk drive. It's rather humbling and it's, you know, we're talking about 100 megabytes and a disk pack, this five-year struggle, a $2 billion program, it's almost-- it's so trivial in terms of today where I can go down to the local electronics store and buy 250 gigabytes for almost chump change. I never thought I'd ever own a disk drive when I worked at IBM. That was just -- it never occurred to me that that would ever happen.

Massaro: I was showing you my watch at lunch, it says 128 megabytes of flash memory on it. 28% more than this disk drive and it's a lot faster. No resonance problems. <laughter>

Porter: Well, it turned out to be quite an industry, didn't it?

Everyone: Sure. Absolutely.

Massaro: It's amazing, though, it's not more profitable today. You mentioned you were getting 75% gross margin at ISS. Memorex was making a fortune on their stuff. And selling these things for $15..., $20,000 a pop and it was costing $3,000 to make them.

Porter: Well, we might point out that, you know, after all this was done, a whole lot of disk drive manufacturers flooded into the industry to take advantage of what you just described, for example, and, at the peak, in the mid-'80s, they peaked out at 76 companies in the world making hard disk drives at one point and now, depending on how you want to count it, you can count eight or nine, depending on what kind of a count you want to make today. One of the problems was that a lot of these companies had -- some of these companies had managements which were able to manage in an environment of keeping up with all these changes that we're talking about, but most of those companies' managements were not able to keep up with the rate of change. They would get very satisfied. You could say they got fat and happy with having a successful product and, unless they were working on the next one to replace it, as IBM was doing when they were working on the-- when they brought out the 2314, they were busy working on what became the 3330. And when they were doing the 3330, they had another group busy working on what became the 3340, et cetera. And, for the independents who came along, some of them had that ability to understand that they had to keep developing to stay up with the pace of short product lifecycles which became shorter as time went on. But the majority did not and the industry has settled down into-- its more and more difficult, with all these technology changes, to achieve those each those each year so the rate of change today is a lot less than it was in those days. To answer your question, Don, why aren't they more profitable? Well, they've settled down into fairly short rates of change and the folks that are involved, all are busy, fairly competent people today, working on those rates of change. It's more difficult to achieve them.

Massaro: I also think they learned to live with 15% gross margin. At Shugart Associates, we were anything below 40% and that was low cost. But now people are happy with 15%, you know? 10% is low, 15 you're happy. I think it's expectation. I mean, there's no reason, with the consolidation we have in the disk drive industry, these people should be living at 15% gross margin.

Porter: Well, it's a competitive industry and there it is. Let me just ask for one round of comments from each of the people involved. Let's take a round around the table. What are they key thoughts that you took away from the disk drive industry, having participated in all of this. Hal, what's your...

Yang: Oh, it was a great ride. It was a -- you know, the disk drive is still, to me, one of the engineering marvels of modern man and if you can engineer, in a disk drive industry, you can disk drive anywhere -- you can engineer anywhere.

Porter: Jack?

Clemens: Yeah, I think that's a good point. The blend of technologies is just intriguing for something as, you know, visually as simple as moving in and out and going round and round is, the technologies that are buried in that are fantastic. At some point or another in my career, I decided, you know, probably what we're doing is building garbage cans. Maybe that's the reason the gross margins aren't any better

than they are but the-- people's ability to use storage, you know, just depends on how much storage is available since it's infinitely of a consumer. The more garbage cans you build, the more garbage people find to put in them and it's amazing the petabytes or whatever it is we're up to now as far as total storage capability in the world. Nobody dreamed, in 1970, this vintage, that there would ever be the requirement for that kind of storage worldwide.

Porter: Bob?

Pattison: I just think that the advance of the technology over the period of years has just been fantastic and it has gone well beyond the 3330 at this point in time and it's been very enjoyable to be part of the ride from beginning to end.

Porter: George?

Santana: Well, I would go along with those comments from others about going on the ride. The thing is, is the ride seems to still be going on, which is great. But I think, you know, it's amazing when you talk about the skills that were used here. We had, initially, we had hydraulics engineers, we had mechanical engineers, we had electrical engineers, we had material scientists, we had computer scientists, we had all of these skills that we brought to bear in solving these problems and it really was a very exciting time for me. I never felt that there weren't plenty of challenging problems to work on.

Porter: Don?

Massaro: Well, the disk drive industry has been very good to me, professionally and financially. I was lucky enough to get started in the computer industry in disk drives and I found it very challenging and there's a lot of room for innovation, so you can go out and start companies, as all of us have. I think the most fascinating thing is we're probably, I don't know, 250, 300 million disk drives a year now and this is the most complicated consumer product by a long shot built anywhere. It's amazing that we can go do that and we can do it for $50 or $60 a unit. And this laid the technology, this was the technology basis for it.

Porter: Well, I would comment, in closing, that, for the 23 years I published the DISK/TREND report, my method of gathering information was to go see the drive producers all over the world, when they were all over the world. Now, given that a disk drive is a weird combination of chemistry, physics, mechanical engineering, electrical engineering and a lot of math, it's fair to say that any place in the world that's advanced enough to design and build a disk drive has good beer, frequently good wine and good food. I found that to be very, very advantageous as I went around and gathered all that information. So I agree. It's been one heck of an industry, gentlemen. Thank you, very much for your time.

END OF INTERVIEW

