Oral History of IBM Data Cell Drive Panel

Interviewed by:
Jim Porter

Recorded: December 8, 2004
Mountain View, California

CHM Reference number: X3047.2005

Jim Porter: Good afternoon. We're here today to do an oral history on a very interesting product that had a very interesting development cycle, a very interesting lifetime and, for its lifetime, it was used in interesting ways. We're talking about the IBM Data CellDrive known as the IBM 2321. It was really something of an offshoot from a paper that was written by John Haanstra in 1957and it actually didn't actually get shipped into the marketplace until 1965. So it had a long development life, a very difficult one, but it was needed, as it was seen in the industry in those days, because it provided something that wasn't otherwise available. It had random access, it was a removable storage medium and, at the time, it had relatively high speed of access, of transfer rates and it had a very high capacity. It, at the time, seemed to be quite a bargain compared to the other methods of storage which were available, especially those that approached being random access. So, at this session today, we have several people who were involved in that program, and, to get started, I'd like to ask each of them to give us a little bit of their own background and a little explanation as to what they were doing at the time, and let's start right here.

Harold Hester: My name's Harold Hester. I joined IBM in 1956. New graduate out of college. I had never been involved or any experience with the equipment that I was to face in the future. I started out a test equipment engineer, eventually migrated into engineering and was involved in the early deliveries of the first machines to the field and had responsibility for field support during that period of time, which was a challenge.

Clarke Carey: I'm Chuck Carey. I joined IBM in 1957 or late '56 and worked on RAMAC for awhile, then on the ADF, which turned into the 1301 and then on the 2321 for a number of years. My last project at IBM was on the 2305, the drum that looked like a disk or the disk that looked like a drum. Then I went to work for Memorex for about 20 years. It turned into Burroughs, which turned into Unisys but I spent most of the time from 1956 to when I retired in 1989 in the disk drive and related computer peripherals business, mostly in design but some in manufacturing also.

Herbert Thompson: My name is Herb Thompson. I first joined IBM in August of 1957, just out of college at the University of Colorado. When I come to IBM, first task was to work on cleaning up problems with the RAMAC drive, 130 -- what's the number of it?

Carey: 650.

Thompson: 650.

Carey: The RAMAC 650.

Thompson: RAMAC.

Bennet: 650 RAMAC, 305 drive.

Thompson: 305. And, after that was finished, I went to work on a new program called the VLCM, Very Large Capacity Memory, which become the 2321 ultimately. So I was there pretty nearly in the first stages of the actual design of that thing. And my background was mechanical engineering so we had this, I call it a mechanical engineer's nightmare because it's so complex, it was really on the leading edge of technology for mechanics, even worse than the typewriter, which had many, many years of development. And all of the functions that the thing had to do make it extremely complex and it was very difficult for me, being a relatively newcomer to the engineering business. But, anyway, it was a good way to learn and I learned a heck of a lot of things about how to do mechanical design. So I stayed with that program until 1965 and that was about the date that it was shipped so I was there for seven years. And, from there, I went on to other miscellaneous tasks at IBM, one being the floppy disk drive. We built the very first floppy disk drive as a program load for the bigger computers. And then I worked on a drum drive for awhile. They were having problems so I had to go over and help them there. And, finally, I left IBM in 1969 and branched out off on my own. So that's pretty well my history up to the time I left IBM.

Porter: Let's wind up with Dave.

Bennet: Hi, my name is Bennet. I joined IBM in 1959 right out of college. I was an electrical engineer, working in Owego, New York, on bomb nav systems for the B-52 bomber. I transferred, in 1963, to San Jose, worked in the plant for a year and, in the fall of 1964, I joined the 2321 program as the first product engineer that they hired and I was specifically hired to be a product field engineer and I spent the first year or year and a half working with the machines in test, trying to make it work, doing some lab work, working with Herb and others. And then, after we shipped to the field, I had an assignment. I worked for Hal Hester most of the time on the Make It Work crew and, when the machines would -- some of them inevitably got into difficulty in the field, I would go out and take trouble calls and make them work. Had an assignment for a couple of months in France with a test for the French government. Actually, joined the 2321 program to see the world but I'm an electrical engineer but I've always been intrigued with mechanical things and I really enjoyed working on the machine. I was on it for slightly over four years and then I went on to a variety of development programs, eventually becoming a manager, a product engineering manager. I managed compliance engineering for seven or eight years, EMC and product safety and a number of things like that and then I went back to my product engineering groups where I stayed until I retired earlier this year, 2004. And, right now, I'm the keeper of the IBM history collection. We have a large collection of old machines in archives and, for the next few months, I have an assignment to, as a contractor, to try and get those machines placed with museums and other places where they'll be on display.

Porter: Well, here we have a Data Cell which consisted of 200 of these magnetic strips, which were, I think, 2-1/4" by 13" long, a piece of Mylar coated with magnetic media on one side and a anti-static coating on the other side. My goodness. I can't imagine how that would work. Gentlemen, can you please explain how this thing was actually made into a system. How did this work as a system, as a product? How do these disks get taken out? How did they get read and write on and how did all this work? Can anybody give us an explanation of that?

Hester: Well, I guess either Chuck or myself. <laughter> Let's grab this thing. We'll do it jointly.

Carey: Okay.

Hester: The Data Cell, this is the Data Cell, is one of ten. You can see it's wedge-shaped and so it would be in a circular -- the other nine would form a circle and the circle, which was called the bin, was hydraulically powered so that it would rotate. And the reason you wanted it to rotate is because you had strips here and you wanted to get them -- the proper strip under the picking station where a device would pick it and lift it up and read it. So the first order of business was to get the bin to rotate into the proper position which was, as I say, done hydraulically. And then, once it was in the right position, then the mechanics took over. What happened was that you needed a device that would locate the strip, there are 200 strips here, and it has to pick out the very one that you want. Well, each strip has a pair of tabs on it which we should probably you later in a close-up, and the tabs are staggered so that, in any one sub-cell of ten strips, there were ten -- if you looked at it, you could see ten different tabs. Actually, you could see 20 because there was an inboard and outboard tab. So the picking device then would be positioned so that some little fingers would come down and enter, just barely enter the array of cells, of strips and then separate from the pack, separate it into a left and a right-hand section and the strip that you wanted, then, would be the one that was not separated out. It would be standing up straight. And so all this had to happen very quickly, in just a few milliseconds. And, again, we'll show you this later in the close-up, I presume. We had little fingers that would move back and forth and, when they reached the location that you wanted them to, these little fingers would come down and separate the strips out, leaving the selected strip standing up, at which point a head would come down and are you the expert on this?

Thompson: Well, I'll give it a try.

Carey: Yeah, okay.

Thompson: Okay. This is what we call the drum. Now, we want to pick that strip out of the cell and wrap that around the drum and we have a magnetic head located on the side here. Let me unfasten this. Yes, there we go. So now this is inside of a strip housing, which we don't have an example of, but it's in a position similar to this when it's in the machine. So now, when you want to pick that strip that Chuck had just separated and left standing, the drum has to pop down and then it -- the head is normally in this position. Now, when you start the thing, as you rotate it, the head goes down. It goes down far enough to where it can pick the strip and then it comes up as the drum continues to rotate and it wraps the strip around the drum now. And the thing was turning, I think, at 1,200 rpm.

Carey: Yeah.

Thompson: And that was sufficient centrifugal force on the strip to load it against the magnetic head. So that was our means of complying that strip to the head. And the head was -- it had two gaps in it. One was a read gap, one being a write gap. And so the head was really a V shape to comply to the radius here, since there was a spacing between the two gaps. And so now the thing would sit here and rotate and you could record, then, with that head on the tracks we had here. And we have an actuator now, an actuator being here, which would move that magnetic head back and forth across that track. And I think there were 20 cores per head. 20 cores for a read and 20 for a write. And that, in turn, would move to one of five positions and we have a mechanical adder here that would move the head back and forth to one of those five positions. So now you could access any one of 100 tracks. Are my numbers correct?

Carey: Yes.

Porter: Let me interject with a piece of information I have here. The average access time to do all of this was -- to mount the strip on a drum, was, in theory, between 175 and 600 milliseconds, all that happened, to mount on a drum and then it was 95 milliseconds average from track to track while on the drum. So the things you just described were really happening very quickly.

Thompson: Yes.

Bennet: Let me point out that the 600 milliseconds included a movement of the array, the hydraulic motion of the -- so that was part of the access. If it was just a matter of picking another strip out of the same sub-cell, it was significantly quicker. I don't remember the number but it was a lot less than 600 milliseconds.

Thompson: Let's see. I would say, we're 1,200 rpm. What's that revolutions per?

Carey: 50.

Porter: Per minute.

Thompson: Per second?

Carey: 50 per second.

Thompson: 50 per second. So I'd say that it was probably doing it in 100 milliseconds.

Bennet: Sounds right.

Thompson: Or thereabouts. As I recall. So it's faster than the speed of light for mechanics. <laughter> Okay. So now we're recording. When we want to restore the strip, we have to drive the drum to a certain spot and then reverse it and it has to be such that the tail of the strip will go back down the little groove in the housing and guide it back down into this opening that we created here such as this. It would drive it back down into that little opening and release this strip from the head. The head has little latches on it that would grab a hold of the strip and hold it there, two holes -- or a single hole in this one, and that little latch on this head, as we call it, would latch onto that and then we had to release that when it got down to the bottom and the drum would continue on a little bit and pull up out of the pack and the cell here could then move on to a different position, to a different cell. And I think that's just about the story behind this in a nutshell. And we have a latch on this I didn't mention that, when the head comes in, this latch flies through and holds it down while it's rotating so it doesn't fly off and hit the housing.

Carey: Yeah. Can you zoom in on this strip or should we do that later? Okay. The strip -- these are the two tabs. Each strip has two tabs and they are staggered along the top. This is probably one of the outer strips. The latching hole is this rectangular hole in the middle and, in essence, the head mechanism's really a four-bar linkage and, at bottom dead center, the latching and the unlatching is done. And then, in the rest position, the head is pulled up past bottom dead center so it doesn't interfere with all the stuff that's going on underneath it, which would be the rotation of the bin. The strip itself [tape glitch] five mils in thickness. This side is the oxide side, the reading and writing. This side is coated with an antistatic carbon coating so that, when you're pulling it out of its sub-cell, it doesn't rub against all its neighbors and create a charge. It also tends to lubricate it. One of the subtle things about this was that, at the very edge here, about an eighth of an inch, you might be able to see what looks like a couple of stripes on the border. Those were actually milled down chamfers on the strip because, when the strip was bent like this, it was subject to what is called anticlastic curvature and, actually, the edges, the left edge and the right edge, would have a slightly larger diameter than the center. You can see this if you have -- I should have brought an eraser. If you bend an eraser...

Thompson: I have one for you.

Carey: Ah! My prop man. Here is an example of anticlastic curvature. If you bend this eraser, you can see that the edges are high. There's a case of anticlastic curvature. Okay. Not only that but it's called a mars eraser, how appropriate. <laughter> And so that would, of course, with the anticlastic curvature, you could not get good compliance with the head. You have too much separation in one place and too little in another. So the chamfers were added to counteract, to compensate for that and that was Dr. Fred Hercter's contribution to a very simple although somewhat expensive solution to a very nagging problem.

Thompson: Yeah, my dad, he first searched the literature to see if he could find anything about anticlastic curvature and he found out that it was analyzed in the 15th century by some mathematician, I've forgotten his name, but nothing since. This sort of relatively old phenomenon that someone had observed but this is the first time we had an application of it.

Carey: You can see it on wrought iron work.

Thompson: Yes.

Carey: Where the wrought iron is bent into arcs, it's quite evident on old wrought iron banisters. I have one at my house, in fact.

Porter: Another thought along the way. You mentioned mars. I don't think we discussed that so far. We should explain that. The project, I think it was mentioned earlier, back in 1958, picked up the project name within IBM -- of course, everything had to have an acronym name, right? So it was the VLCM project, which stood for Very Large Capacity Memory and, in 1960, that was changed to the MARS or MARS project which stood for Modular Access Random Storage. Something about the computer industry, without an acronym, nothing counts, right? <laughter> The acronym name for the project within IBM was MARS until it was introduced as the 2321. Let me ask. The things you described, gee, you wonder how all got done in a millisecond or two -- well, actually, several hundred milliseconds, on the average, but less than a second. So how did all this get done so quickly and were there a lot of problems in doing this?

Thompson: Oh, yes. <laughter>

Carey: There were.

Thompson: Yeah, you have to remember, in those days, we didn't have high performance servo motors so we looked at using a servo motor to drive the drum and it was just about impossible for existing servo motors to drive that drum fast enough for us, to accelerate it. And so we had to depend on energy storage. So we had a system of clutches and flywheels and springs such that, when we'd stop the drum, in its home position, we had wound up a spring. So now -- and then we latched it there. So now, when we start the thing, we can pull the latch and the spring kicks it out and the clutch picks and this then causes the thing to rotate. So that's how we could get the high speed out of it. And that's one example of why we had to do it mechanically. Also, in those days, we didn't have stepper motors and stepper motors would have been a marvelous technology to do the positioning that we have to use to position these little fingers across the disks to select the right one. But stepper motors hadn't been invented yet. And they also have -- you could have put a servo in that but now here you have a big servo motor driving these. So what we've had to do here is to use mechanical adders. And mechanical adders were very common in those days, a good example being the Selectric typewriter. So that's the direction we went there and that was quite fast. That's -- the pick time of a solenoid was maybe 10 or 15 milliseconds so we could move it pretty rapidly. So that was what we call a binary adder. In other words, one solenoid would pick at one increment of disk, one -- it would move it from one tab to the next. You picked the two solenoid, it would move it two positions. You picked the four, it would move it four or you could use any of that combination of that. So we could get it moved to the right spot with the three solenoids.

Carey: These are the solenoids.

Thompson: Yes.

Carey: The things you were talking about.

Thompson: And each solenoid had a different stroke on it, of course, a stroke of one unit, two units or four units. And then we had a detent on it as a final precision positioning which required another solenoid.

Carey: And there is the rack and this is the detent solenoid.

Thompson: And also, as I recall, this is a viscous damper here, is it not?

Carey: Yes.

Thompson: And, to slow these things down when they're moving and stop them properly, we had to put a viscous damper on it now. And also we have to put optical sensors on this to sense position, to be assured that we're in the right position after we move.

Bennet: I think the viscous damper, if I remember correctly, was added after the machine was in production, in the field.

Thompson: That's right.

Bennet: It was in response to a problem.

Thompson: Mm hm.

Bennet: We couldn't maintain accurate positioning and we'd sometimes pick the wrong strip and so that viscous damper was added to the adder on both the finger carriage and the head positioner.

Thompson: Yes.

Porter: What were some of the other significant engineering challenges that had to be taken up to make this thing go?

Thompson: Well, it's just full of those things. <laughter> And every one was a schedule killer.

Carey: Yeah. Where to start? Are you zoomed in on this mechanism so, if I move some parts, you'll be able to see them move? You see the little red guys down here? Okay. I'm going to move them a little bit. Can you see that? Good. Now, this is a pick or rather a separation. When the sub-cell is in the proper location underneath this, that's the down direction, this would be actuated, what my left index finger is doing, and those four little fingers would drop down and then move outward and they would then engage these tabs and all the neighbors of the strip that you wanted selected. And it would hold them like that. In the meantime, we had a couple of wedges here, one there you can see and I can't see it but there it is there, which came down here and opened up this spring steel, or was it beryllium copper? I think it was beryllium copper.

Thompson: Yes.

Carey: It was beryllium copper leaf springs. These springs are such that they clamp the sub-cells together so they're not vibrating but you have to release that spring tension when you want to separate them. So the wedges drive down and the fingers come down all at the same time. You can see the -- maybe. You can see the wedge come down and the fingers come down and separate them. And these red finger carriers moved back and forth in ten different positions through the action of the adder. And what else? Well, yes. There were lubrication problems. One of the lubrication problems was solved by the carbon coating here. At least it was ameliorated. But, still, we needed to have air lubrication so we had -- and I guess they were fixed in the upper housing, we had air nozzles that, when the sub-cell was separated and the strip was standing up, we had a couple of air nozzles that blew down alongside the selected strip to blow compressed air down there to give it a little air bearing on the way up and, just as important, on the way back because, when it's going back, you're pushing -- this is a pretty flimsy thing and you're pushing it backwards down into kind of a tight fit. So, one, you wanted to lubricate it as well as possible mechanically, also with the air. And you wanted to make sure that it didn't crumble and, occasionally, they did in the first -- you know, it was a very serious problem at first. But it got better as we learned how to do it. One of the things that was working in our favor was, about, I don't know, maybe a third of the way back into the stack here, actually, the strip actually got physically stiffer because it was decelerating and so it was, in effect, being pulled up by the head, even though it was -- the velocity was down, the acceleration was up and, as it was being pulled up, of course, it was like a human being in a rack, right? You get stiff. And so it helped us in reinserting the selected strip into that tight location that we actually had Newton working for us during those last critical six inches or so.

Porter: Wonderful. Hal, do you have any comments about what it took to keep all this alive and going?

Hester: I have many hours of <laughs> one of the responsibilities I had, in addition to the logic of the machine, the electrical logic, which was the simple part of the machine from an engineering point of view, the challenging part that I had is I was responsible for field support, which meant that, when it went to the customer and the customer had a problem, my group was responsible for fixing it and Dave was the lead man in that group. And one of the biggest problems we had was training the field engineers on such a new device and Dave and I both spent a lot of hours on the road and in customers' offices trying to explain to the customer engineers and to the customer that things are going to get better. <laughter>

Porter: Did they get better?

Hester: Yes, they got quite good. After, I don't know, six, nine months or a year or something like that?

Thompson: Yeah, first six months were...

Hester: First six months were...

Thompson: ...rocky.

Hester: ...very tedious.

Porter: Of the field experience, the first six months?

Hester: In the field, yes.

Porter: It was a mistake to be a customer in the first six months?

Hester: Not a mistake but you had to be quite flexible. <laughter>

Thompson: I remember getting up in the morning and shaving and I'd have the radio on and I'd hear this advertisement, it was a 2321 operating ker-chunk, ker-chunk, ker-chunk, ker-chunk, ker-chunk. I don't know if you ever heard that.

Carey: I don't know.

Bennet: Absolutely not. <laughter>

Porter: So IBM had radio advertisements for this?

Thompson: Yes.

Porter: Oh, really?

Thompson: Well, I don't know if it was IBM. I don't recall that. Might have been another company that had one of their products. I think that might have been the case. I don't think IBM would -- but the noise was the sound of progress, right? We're on the cutting edge <laughter>...

Porter: We should interject the thought that IBM was not the only company that made a storage system with a magnetic strip card. A company who actually shipped before IBM did was NCR with their very well-known CRAM card system. CRAM stood for Card Read and Access Memory and they actually shipped with a computer in 1962, about three years before IBM shipped the Data Cell. And another company that was well-known at the time in the computer business, called RCA, who later dropped out of computers, but RCA had the MCF file, which was the acronym for the Magnetic Card File, which also shipped in the early 1960s. And I believe there were a few other companies also doing magnetic card files, some of which may even have been a resale on one or more of those. So it was not the only pioneer -- IBM was not the only pioneer into the market with a magnetic strip card file and, as you can imagine, those engineers working for those other companies must have gone through a lot of the same difficulties that you did in making all this happen. But it may have been that you heard a commercial for somebody selling a magnetic card strip file going ker-chunk, ker-chunk. It may not have been IBM. <laughter>

Bennet: We used to claim, when the subject of the CRAM or the RCA race came up, we used to say that we had one significant advantage over those two machines, which would send the strips out and then try to capture them at various parts in their process. In our machine, we never let go of the strip. We always had a positive grip on the strip, no matter what it was doing, which we felt helped us in the reliability mode and I think ultimately made the difference and I claim that our machine was significantly more successful than theirs, the early days notwithstanding, but that's the way it ended up.

Porter: Well, regardless of the difficulties, the computer industry needed this kind of device from IBM and these other companies because, as we said before, it did provide random access to the data and it got to the data one heck of a lot faster, on the average, than using a tape drive, which was the only other cost-effective storage device at the time because the disk drives at the time had started in 1956 with the IBM's disk drive with the RAMAC, all of five megabytes on a big 24" disk set and, when it came to the time when this drive was coming out, it came out and was announced at the same time as the IBM 360 computer, which changed the whole computer industry because of the interchangeability and readability up and down in that system, on the IBM 360 system. So the disk drive IBM brought out at that time, the 2314, and the one before it, the 2311, 14" disk drives with disk packs, removable disk packs, were very big in the market. They provided very fast random access, faster than this but, when you compared the 400 megabytes of a Data Cell drive with the 10 data cells, providing a total of 400 megabytes, this was [tape glitch] all of the 2314s behind one controller, eight drives active at one time. And it was certainly a lot cheaper than the 2314s so you had the choice of using tape, which was serial access if you were using a computer, or you could use the disk drive like the 2314 with very fast random access but the middle ground providing random access and removability. The only good choice was really the Data Cell. So that middle ground found a very large number of customers, didn't it? Do we have any idea how many of these things were actually sold or used?

Bennet: There about 1,300 of them built.

Porter: I use the word "sold". IBM didn't sell most of their equipment in this case. They leased it.

Bennet: No, they were mostly on rent. Very few sold, if any.

Porter: But about 1,300 then were actually manufactured and most of those ended up in use?

Bennet: Absolutely.

Carey: The DMVs and insurance companies were the two biggest users because they had massive...

Porter: Very popular with the people who had to just upgrade files from time to time then, I gather?

Carey: Yes.

Bennet: Actually, some of our more successful installations were the DMVs and that was because they tended to use the machine more than the other customers. We had a projected monthly usage rate of 120,000 pick and restores a month and that was what we figured the average customer would do. But many of the DMVs, and Oregon DMV was an early example and there were others: Dade County in Florida and there were others around the country, I don't remember all of them, but Oregon, for example, had two machines: one with the driver's licenses and one with the registrations for the State of Oregon. And they tended to do 120,000 pick and restores in less than a week. So they used the machines more than four times as much as the actual usage. Well, now -- the projected usage. And, of course, when they first started doing that, it was a disaster. They had a lot of trouble. And, like a lot of accounts, the problem was, as Hal pointed out, had to do with the customer engineer and his state of training. We found out that, in these high usage accounts, that we would have to get the customer engineers up to a level of knowledge -- now, this is early, before a lot of the changes that were made in the machine but back when we were struggling, we had to take the customer engineer by the hand, basically, and take him through the machine, adjustment by adjustment, point by point to the point where they would get comfortable in servicing the machine. And, after we got to that point, say, at Oregon DMV, we had a crew of four or five guys that were in that account so that they could come in at any hour of the day or night and they knew what to do and the Oregon DMV became one of our early shining success stories. So that was the kind of thing -- we had a lot of individual little but major problems along the way but, in the final analysis, successful operation in the early days, when the machine and everything was critical, depended on having a guy in the account that knew what to do and knew how to fix it and didn't put one problem on as he was taking another problem off. So it was interesting.

<crew talk>

[video off then on]

Porter: And one of the things we haven't talked about is we've only got four people here who were actually involved in this program but there were quite a few, weren't there? What was the size of the program and who were some of the key people that we don't have here today? One, for example, a fellow named Al Shugart managed the program at one point, I know, while it was in the MARS status.

Thompson: Okay. I have a timeline here and we should briefly go through this.

Porter: Some of the key people involved, yeah.

Thompson: Yes. Okay. It first started out sometime prior to July 1960 and, in the early days, it was handled by two people: Walker and Ty Cowan and Jake Eddies worked for him and I worked for him then at that point in time. Okay. Then we moved over to a different location. We were at 99 Notre Dame in San Jose and we moved out to a motel on Willow Street in San Jose and a Dr. Peter Lazarus was the program manager at that point in time. And I was just talking with the fellows here about we had built up a model of this and this array of cartridges, it was probably about two feet in diameter, and you had this big, humongous hydraulic motor now that's going to drive this thing very rapidly. So what happens, this thing remains stationary and the whole machine underneath it turns. And we had a heck of a time with that. So we actually built a model whereby we had two of these arrays that fit together like gear's teeth, the teeth on two gears in remission, and that way you'd counterbalance everything out. But we were able to overcome the problems by a lot of engineering and we didn't have to do that in the final end but we did model the thing. Okay. Then, at that point in time -- oh, we first started out with 17 people and then we moved up -- as we progressed along, we went to 29 and then the program manager was Don Johnson. And then we changed the name from VLCM Technology to MARS file. This was now under Dr. Peter Lazarus and now our staff increased to 30. And we had people, some of the peoples' names were Rusty Nagakura, Ty Cowan, and Walker -- what's his first name? I've forgotten.

Carey: He's gone.

Thompson: Yeah. And then we increased staff to 45 people and now here we get Dr. Hurtrig, he was personally mentioned that did the work on this little beveling on the strip that we had to do. And Yang Hu Tong, Ty Cowan, Rusty Nagakura, Nick Stan, Geddes and Herb Thompson, myself. So then, in 1962, we moved to 49 people, just about the same group of engineers. And then, in December of '62, we were still at 49 people. Yes. And we're still at 49 in the middle of 1963. In the middle of 1964, we moved up to 71 people and then, at later part of '64, we're at 76 people. So that's about what it was all that period of time. Now, our first location was a place in San Jose called 99 Notre Dame, which was -- currently, it's a legal facility.

Bennet: Yeah, it's a courtroom.

Thompson: Courtroom.

Bennet: Yeah. Santa Clara County.

Thompson: And then we moved out of that into this motel and IBM still hadn't completed the development lab buildings. And, finally, when they completed those, we moved out to the IBM site on

Monterey and Cottlel Road. So that's pretty well the timeline on the thing. The first machine was shipped in 1965. So this timeline shows out to previous -- this is pre-shipment.

Porter: Well, all those people had to be doing something like solving problems in getting all of this going, didn't they?

Thompson: Oh, yeah.

Porter: You talked about a number of those developments but, looking back at it now, what were the other key things that it took to get this product really into the marketplace and make it work?

Carey: I'll start with this.

Thompson: Okay.

Carey: I mentioned the fingers that we used to separate the strips. I didn't show you that the fingers moved. They had to move to ten positions because there were ten sets of tabs on each strip. And I finally found a place where I could get my thumb in here to work the solenoid. Now, you can see that the fingers will move if I -- there, withdraw the detent. Can you see this in your -- you can see it, all right. Can you see over here where there's this little arc-shaped rack?

<crew talk>

Carey: Okay. And the trick was, while the bin containing all the cells was rotating, the controller also sent the strip address to the drive and the mechanical adder would -- well, first, the detent solenoid would be actuated and the detent withdrawn, permitting movement. And then the fingers would go to one of ten pre-selected locations. When it got to that location, then the solenoid was de-energized and the detent snapped into that rack, locking it firmly into a very precise location. Remember, we're speaking of milliseconds here and all that activity caused vibrations and the fingertips would vibrate and cause damage to the strips. And so we had to add this cylinder to the bottom. This is on the same shaft as that detent arc and this is a viscous damper, which damped out the vibrations of the system when the detent slammed home. And it worked quite well. Also, we should talk about these things. Is this a good time to do that?

Thompson: Yeah.

Carey: All right. I'm holding in my hand a little incandescent lamp like the ones you'd have in your flashlight. We had, I think, ten of these, maybe 12, I'm not sure, in all this mechanism here because we had to know, at all times, where certain key pieces of hardware were and what they were doing. For instance, we had to know whether the fingers were down or up, whether the detents were fired or not fired, whether the bin was rotating or not and a number of other conditions. And, to do that, we used a light and a sensor. And then we had various masks that would go in front of the light. So if you wanted to disable the sensor, you'd have a mask in front of it and then you'd pull it back, sort of saying yes, I'm in position or whatever. And that worked, except we could never find lamps that were of such quality, consistent quality, that the filaments wouldn't vibrate. Now, if you're using this lamp in a flashlight, it doesn't matter if the filament's vibrating. You'll never know it. But, in a machine like this, the filaments vibrating generated a small bit of electrical noise which just drove the machine crazy because it would seriously affect the timing that you were supposed to be controlling. And so we had to finally -- and this was a change that was made after the first customer shipped so we had to retrofit to field. That was a very expensive move. But we ended up with a very powerful -- one very powerful light. It was an Ozram lamp of 70 watts or something, I don't remember, but it was a very powerful, reliable light which didn't shake. And then we have two separate fiber optics harnesses that went out to the 10 or 12 locations, however many there were, that formerly were all occupied by individual lamps. So we had this big octopus of fiber optics that went out there. And so we used the fiber optics just as a light pipe. We didn't take advantage of any of the characteristics that are currently being used for data transmission, just as a light pipe, as a very clean, good light pipe. But, as I say, that was a running change and it necessitated changing machines in the field and it was very, very expensive indeed. But it worked and it solved that vibration problem in the filament.

Porter: Hal, you overviewed the field response to all of this, I guess, didn't you?

Hester: Yes.

Porter: And how did you characterize the attitude of the customers toward all of these problems?

Hester: It was varied, depending on their success, depending on the success of the machine that they had. Initially, they were very enthusiastic. It was something that they had been waiting for for a long time. In the early stages, we had two problems in the field. One was being such a complex machine, we had our own problems, mechanical mainly, because it's a mechanical wonder. Electronics were relatively straightforward, considering the mechanics part. So, mechanically, it was something that the field was not used to and we had a problem of not being able to educate and convince the field that they needed a different CE, Customer Engineer, with a different training than they were used to. And so that was one of the major problems that we had in the beginning. And Dave was the interface to the field and he took all the blows, all the calls, all the long hours <laughter> along with a couple of other guys that we had. After we got a Customer Engineer group up to speed, and I say group because you needed more than one because they tried to use the machine 24 hours a day in some of the major installations and so you had to have highly trained people on board 24 hours a day, seven days a week and that was a little bit of a stretch for the field in those days. Once we got them convinced and got them trained, it settled down and things went reasonably well and according to plan but there were many, many weeks and several months in the first year or two that were nip and tuck all the way.

Porter: Well, looking at it from the point of view of the customer, when his system was suddenly down because of the failure like this, how long did it take, typically, to get him going again?

Bennet: Sometimes it took quite awhile. I would cite the case of the, I believe, the first customer to actually use a 2321 to do work was a small life insurance company in Jackson, Mississippi. Now, Allstate had machines earlier than that but Allstate spent a lot of time in programming preparation and they didn't really use them. This little life insurance company had one 2321 on a model 30. He had a room full of operators sitting at display terminals and he couldn't keep the 2321 running long enough to get the work transferred over from the 1401 that he had used previously. And I was asked to go down and try and get that situation straightened out and it was kind of a do or die thing. And it took -- we went completely through the machine and found a number of things that needed to be done. In those days, there were a lot of things that were unstable. For example, we had the noisy lamp problem, we had a problem that the clutches were inconsistent, we had to develop a technique for stabilizing the clutches and all those things. We actually would lock the shaft and energize the clutch with the motor turning and actually burn the clutch in. We figured out how to do that and then that became a manufacturing process later. But, to get back to the story about how long it took, I think I probably spent the first four or five days actually going through everything on the machine. And, once we got the machine right, then we had programming problems and it went on. And I was in that account for three weeks. I just promised them that I wouldn't leave until they were happy with it, that it was running and I got their work caught up and it took about a week with the machine and a couple weeks with the systems engineers and the programming types and, you know, to get all the kinks out of everything and get it running. Now, that's probably longer than typical but a machine that was in trouble and the customer behind typically was at least a week before we got them up and running and confident in the machine. A lot of it was confidence. If you went into the account and got the machine running, well, that was okay and, if you were to leave and the machine had one other hiccup after you left, well, that would -- generally, that was the kiss of death. They'd say no, I don't want that machine any more. So the strategy was to stay in the customer's account until everybody was happy and everybody agreed that the machine was running. So we did have some long calls in the early days.

Porter: Considering that these were leased machines, not purchased, you could get thrown out on your ear, couldn't you?

Bennet: Oh, yeah. Oh, in a heartbeat.

Hester: Yes.

Porter: Well, again, looking at it from the point of view of the customer, these must have been some very difficult times for a company like you mentioned, that small insurance company to be down for several days.

Bennet: Well, had a room full of girls sitting around, you know, with nothing to do until we got the system up and running. That was part of it. That was a high level thing. We had the president of the company in the computer room, you know, several times a day wanting to know how we were doing.

Porter: When you walked away finally after getting all that finished, were they smiling at you or growling?

Bennet: Oh, yeah, the machine was there for several years. They were very happy with it. Yeah, it was a success story.

Porter: Good.

Bennet: First one.

Porter: Good. I can see how the life of a field engineer in those days would be a little difficult on some of these, especially if you didn't know all of the intricacies of this mechanical monster.

Carey: They learned very fast.

Hester: I think two things that we did wrong as a company. One, we underestimated [tape glitch] of failures that we'd have in the early days and the second one, we underestimated how long it would take to train the typical customer engineer, in those days, that were used to, you know, everything spinning and just the heads moving in and out or something like that, We grossly underestimated the time it would take to train those people and get them up to speed where they had confidence in themselves and the machine and then they could give the customer that confidence. For a long time, we had installations where IBM didn't have confidence. In other words, the IBM customer engineer and the IBM local manager, they were sort of antsy about it. So it took us nine months to a year, I would guess -- at least that...

Bennet: At least that.

Hester: ...optimistically.

Porter: Let me ask a nasty question. Did IBM have enough confidence in this system to use it internally in their own business requirements?

Thompson: I don't know.

Hester: I don't know. Were we using them?

Porter: Did you have any consumer...

Bennet: I think most of the big IBM installations had some. Whether it was mainstream or just so that they'd have the whole product line. You know, in those days, you had a room full of equipment and they wanted to have something of everything and so we typically, at the big accounts, they'd have them. The airlines, which were a big, you know, always big DASDI farms, they always had a few of them around but they weren't mainstream. So I think that's a yes and no kind of a thing.

Porter: Well, according to the information, the 2321, after being introduced -- announced in '64, it was shipped in September of '65 but it was "withdrawn" in January of '75 and, apparently, there was never a follow on product done to the original 2321. Was that because the original 2321 kept getting fine tuned instead of a follow on product?

Carey: I think it was because, by that time, disk drives had advanced to the point where they were so much faster and so much cheaper that it usurped the market that this was intended for. That's my personal opinion.

Porter: By '75, of course, the Merlin, the 3330, had been shipped in '71 and then the Winchester, the...

Carey: The small heads.

Porter: ...for smaller and midrange applications was shipped in '73 so, yes, these were much more cost-effective in terms of price per megabyte and in terms of capacity available at the need to use the capacity.

Bennet: I think that's right.

Hester: And the other thing, the big thing, too, was reliability. The reliability of the disk drive, you know, always exceeded the reliability of the 2321.

Bennet: And maintenance requirements.

Hester: Maintenance requirements, keeping it going.

Bennet: Merlin actually made the 2321 obsolete. As you pointed out earlier, Jim, the 2314, it took two 2314 eight spindle facilities or 16 spindles to have the capacity of one 2321. But, by the time Merlin came along, a Merlin facility, if I remember correctly, actually had more capacity than a 2321. And so that was really the, you know, and, without a follow on product, then the 2321 really didn't have a frontline place in the marketplace any more.

Porter: Yes. The Merlin, the 3330 drive, had a capacity, when it first shipped, of 100 megabytes per spindle. It had eight active spindles so it had twice the capacity...

Bennet: By two to one, yeah.

Porter: ...of the 2321, yes. And...

Bennet: That was it.

Porter: ...I didn't mention this but I had been active in the industry in those late '60s. I'd gone to work for the old Memorex in '68. I think, in traveling with Memorex salesmen around the United States, I was probably in 500 or 600, as they used to say, glass houses, computer rooms, in that era and it was pretty unusual to see only one string of drives behind a controller. Of course, we talk about disk farms, as much as you could see in a large room, the computer room, on a raised floor with disk drives. This was typical of a large installation. So, yes, there was a huge amount of disk capacity in place during that period and it was becoming more cost-effective with each new generation. So it's easy to see why cost-effectiveness probably did away with the drive but, by that time, by the end of the '60s, had the maintenance problems become just routine and not a problem on this product?

Bennet: Let me try that. Yes, they became more routine but they were not zero. This machine required, for its entire life in the field, required preventive maintenance and it was specified as a monthly PM and it was specified in terms of pick and restore cycles. So every 120,000 pick and restores, the machine required preventive maintenance which was, as I recall, at least two hours of customer engineer time. It could have been a little longer than that but it was quite significant and it required a highly trained guy and so that -- so maintenance and -- periodic maintenance, preventive maintenance was a factor for the entire life of this machine. You never got away from that. It was because there were a number of adjustments that had to be checked. They did occasionally slip or you'd have a wear factor. There were a number of points that had to be checked on a monthly PM and that never went away completely.

Porter: Well, it was quite a product. I think we've covered most of the main points. Are there any other significant points you'd like to mention?

Carey: The head.

Thompson: Oh, perhaps I...

Porter: The head.

Thompson: ...should talk a little bit about this head because...

Porter: Go right ahead.

Thompson: This is where I learned how to build magnetic heads, on this guy right here, and this was a good place to start. And I had the technician, Chuck Blackly, who was really a watch maker but he would build these things for us in the early days and, of course, it went into production later on. But this is rather a complicated head because, in it, we have 20 -- I think it was 20, as I mentioned previously, 20 read cores and 20 write cores. So there's a lot of pins on the back of it. And now, since the strip is rotating in a cylindrical fashion, for it to comply on two different gaps, this has to be a V-shaped head. When we first started out this program, we were using foam rubber wrapped around this drum and then we'd grind that off to make it as accurate as we could. And this was very similar to a tape drive, a little tape cassette, right? They have a little piece of felt in there that holds the tape against the head. And so we're going along here for a couple of years trying to struggle and trying to make that thing work and we couldn't do it. We finally got stumped to the point of recruiting people who were able to read and write on the thing and everyone felt, well, we're going to overcome all those problems. But we could never do that. So we had to stop and change things and that's the reason this schedule is stretched out. Because you come up to a problem like that and you just can't solve it. And so we finally then developed the centrifugal loading because, as the drum is rotating, why, the centrifugal force of the strip is sufficient to hold this in compliance on the two gaps in the head. And so that was all sort of a struggle and, even to -- you know, you were talking about six months to a year development time just to do that one thing. If it's new, it's revolutionary, no one's ever done it before. So that's one of the reasons all these things stretched out. Then you run into this problem here of the anticlastic curvature. That takes you six months to solve. Then you put it into production and -- so, all of these new problems come along that really eat into this thing. So, in any event, that's what the head looks like and I guess the last head they ever built looked just like this one, didn't it?

Bennet: Yeah, except it didn't have the shield in it. That's the shield that was on the model one where we had to read the clock track while we were writing?

Thompson: Oh, yes.

Bennet: So it's a little bit unique but, other than that, the slots in the head and the position and everything never changed from the ones that Chuck built.

Thompson: Yeah.

Bennet: That's a genuine Chuck Blackly head, though.

Thompson: Yeah. I don't know what these wound up costing but I'm sure they're worth their weight in gold, literally.

Porter: Speaking of costs, do you think that IBM made money on this whole project and the program afterward compared with all the time it took to get it out to the marketplace?

Thompson: Well, you know, I don't think they could have made money on the 1300 models.

Bennet: I was told that, at one point, that IBM did make some money but they didn't make an awful lot. They did break -- they did better than break even but they certainly didn't make the profit that IBM normally expects.

Porter: Because of the extended amount of time for development? Because of the amount of field support time required for the product and because of the limited time in the market? Those problems, huh?

Thompson: Well, I think it's going to pay off for IBM to have it because now you've got a system 360 that likes to have this device. And so it's going to -- the system is going to have to help pay for it, I would think.

Carey: One of the reasons that the first customer ship was delayed, although not the only reason, was that, initially, there wasn't sufficient programming, internal IBM programming support. That lagged. For us, back on the machine, it was a kind of a fig leaf because it meant that we weren't the only ones holding up the first customer ship but there were definitely program problems in the operating system DOS, OS, what do we call it? Big OS.

Bennet: Yeah, Big OS. The DOS programs were available earlier and all the early accounts that actually...

Carey: Used DOS.

Bennet: ...used the machine were DOS customer but there were a large body of customers that cancelled their orders or delayed them significantly because the Big OS, as you call it, which is what we did call it, was delayed for some time. The 2321 support within Big OS was delayed. Big OS got out for other product types.

Carey: Oh, sure. It got out there and they couldn't operate the system 360 without it but...

Thompson: <inaudible> another point, too. In these days, IBM was deciding do we go with tubes or do we go with transistors? The first equipment we were using was tubes, as I recall. Is that -- am I correct?

Hester: Yeah, RAMAC was tubes.

Carey: By the time I got there in 1962 -- oh, RAMAC was tubes. Yes.

Hester: RAMAC was tubes.

Carey: But, in this project, I don't know. I got to the project in June of '62 and it was solid state.

Thompson: Well, I remember Rusty coming out and showing me the first solid state board. Rusty Nagakura. He held it up and said, "Here's all the transistors on it." It was amazing how you get that many transistors on it.

Carey: Yeah.

Bennet: For technology called SLT, Solid Logic Technology, that was then used in this, there was a prior technology called SMS, Standard Modular System, that was discrete transistors on a phenolic card and that was used on, for example, the 1301 file. But this machine was primarily SLT except for we had a few SMS cards of the earlier discreet type in the power supplies, if I remember, the mini watt supplies.

Thompson: Right. That decision was made about 1962 or '63, I think. That, yes, we will go with that technology.

Porter: Okay. Any other key points before we summarize?

Carey: Yeah, one other thing that Tom mentioned that we didn't talk too much about was the positioning of a cell under the picking station in the upper housing here or the lower housing. There were, as I said before, ten of these guys wrapped around in a circle and -- pretty massive. Weighed probably more than 70 pounds altogether and they sat on a turntable, which also, you know, the mechanism weighed a lot. And so it had to be driven from one position to another, never more than 180 degrees because it always had enough smarts to select the shorter path. But it had to be able to go 180 degrees or less in just a handful of milliseconds. And so, to start that and to stop it, big rotating mass, is a very, very complex thing. And so the hydraulics or the control of the hydraulics and the mechanical valves and the electronic controls of those valves gave us fits for years. Hal Lowe and Brett Nelson and those people worked on that for years, to try to smooth this piece down. Because, if you get this big 80-pound rotating thing and then you have to stop it, you want to stop it very smoothly but very deliberately and it's got to stop in the right place in the shortest amount of time and it was a very tough problem. The hydraulic system was selected because it was the only thing available at the time that could deliver that much power in a short period of time. In fact, we had a pump and then we had a big nitrogen-filled accumulator just to give it that instantaneous boost, kind of like the springs in the drum. So that we could get a lot of torque in a little while. Just milliseconds.

Porter: Well, I was amazed at the time and I continue to be amazed that you guys somehow made this thing work.

Carey: I'm amazed we're all still alive. <laughter>

Hester: I think, from the description that he just gave, it was something that the field and the customer was totally unused to and, when you moved that into an installation with people that have been doing something for the last 30 years a certain way and, all of a sudden, you've got this humongous thing, really something they want and they've sold to the customer, and yet they didn't really appreciate the impact it was going to have on their installation from both the customer engineer's point of view and the customer's point of view, both. And that was one of the, I think, places where we sort of dropped the ball in planning and could have done a better job in that area. But we survived.

Porter: You survived and you still survive.

Hester: Yes.

Porter: I'd like to do a brief summary from each of you, if you wouldn't mind, just if I could go around the table and come up with just a few thoughts on, as you look back on this adventure, at that time, gosh, it's sometimes 45 years ago, I guess, on the average, a lot of time involved. A little more than that, actually, in some -- many cases. What are your key thoughts about all this and how do you feel about it? Let's start with you, Dave. What was your reaction to the whole thing?

Bennet: Well, in a lot of ways, the 2321, for me, was a grand adventure. I've always been kind of a tinkerer and I enjoyed the challenge of making it work and I actually -- and I really enjoyed the job of saving one of these guys in a customer installation, you know? Usually, when I got there, they'd be unhappy [tape glitch] later on or you and the machine are out of here, that kind of thing. And I pretty generally saved them and that was a great experience. And the machine was capable of running very reliably and making customers happy and so it was quite wonderful to me. The development part of it that I was involved in was also an interesting experience because there were some very talented folks that were involved in the development of the machine and making it work. And to participate in that development process, tuning, if you like, which is -- it was tuning when I came along, the machine basically worked but it didn't always work and it didn't work always when you wanted it to and we had to get it through test and we had to do other things but we had some smart guys that always had another idea and, you know, work our way through it one problem at a time. It was fun.

Porter: If it had worked perfectly, you wouldn't have had nearly as much fun.

Bennet: That's right. That's absolutely true.

Porter: Herb?

Thompson: Well, I think this was one great adventure for me, also, because there are very few companies in the United States that would invest that sort of money and time into doing this so I have to give IBM a lot of credit for that. So it was certainly a great education for me and I learned how to be, I thought, a pretty good design engineer. And also, here, you're really on the leading edge of things and it's fertile grounds for new thoughts, new ideas, and patents and that sort of thing. And that's what makes it very, very interesting. So all in all, I'd say it was a real good period of my life. I enjoyed working on this and we had a good group of people, too, a very close knit group that worked very well together.

Porter: Chuck?

Carey: Well, I agree with what these gentlemen have said. Of course, 40 years ago, we all felt happier. One of the things I was amazed at today -- I haven't seen one of these in years and years and years. I'd forgotten how complicated it was. <laughter> And I'd forgotten an awful lot and I still have forgotten an awful lot. But -- there's little adjustments and things here that I'd completely forgotten about but I haven't thought about it in decades so -- one other thing I'd like to say is that the people like Hal and Dave here, on my left and right, really saved us. Their performance in the field was superlative, medal of honor type work and they actually made it work. We were very lucky to have them.

Porter: Hal?

Hester: I think the thing I got off the program more than anything else, I got to see more of IBM and more of the country, of the United States, Europe and most other countries and I got to understand more of why -- what you needed in a product to satisfy a customer and why the customer needed it. I mean, I really began to appreciate the fact that, when a machine was down, the customer really had a problem. So I learned, from that point of view, a side of IBM, having been in the design end all the time and had the expert engineers and technicians and people to support and fix it, you sort of begin -- you get sort of blasé about how hard a problem is to fix until you turn it over to somebody who had a few weeks' training and has got a customer breathing down his neck. So it gave me a much broader understanding of IBM and what a product needed to do.

Porter: Well, I'd like to thank all of you. I knew of this product during its -- not the early stages but the very latter stages of its life and I heard about it in the field all the time and I regarded it always as an amazing product and I now understand why it was so amazing: the people who put it together were amazing, otherwise, it wouldn't have existed. Thank you all very much.

END OF INTERVIEW

