

The Disk Drive Story

Chapter 1: IBM's RAMAC

Transcript #4

Transcript of Videotaping on January 21, 2002

Participants: **Louis D. Stevens (Lou)**
 Joined RAMAC Project - May 23, 1952

William Y. Crooks (Bill)
 Joined RAMAC Project – April 1, 1952

Norman A. Vogel (Norm)
 Joined RAMAC Project – Early 1954

Moderator: **James N. Porter (Jim)**
 President, DISK/TREND, Inc.

IBM Project
Sponsor: **Dr. Robert J. T. Morris**
 Director, IBM Almaden Research Center
 Vice President, Personal Systems & Storage

IBM Facilitator: **C. Denis Mee**

**IBM Assistant/
Transcribed By:** **Merridy Howell**

**IBM Camera/Audio
Visual Coordinator:** **Farukh Basrai**

Location: **IBM Almaden Research Center**
 650 Harry Road
 San Jose, CA 95120

Transcript:

Jim: Today the subject of this session is the inside of the magic disk drive, the first disk drive in the early RAMAC – how it was created and what it took to make a disk and a head somehow do the reading and writing that had to be done. I wonder if we could start by each of us introducing ourselves, and Lou, let's start with you.

Lou: I'm Lou Stevens and I was one of the early members of the 99 Notre Dame Laboratory team. I came from Poughkeepsie as Rey Johnson's Technical Assistant – kind of his arms and legs to make things happen in the laboratory. I am a graduate of The University of California and my wife wanted to come back to California. I talked Rey into giving me a job.

Bill: My name is Bill Crooks. I was the second engineer hired in San Jose at 99 Notre Dame. I joined the company on April 1, 1952 and I don't know if it's significant or not, but it was April Fools Day. We had many projects, and when the RAMAC was first conceived, I got involved in trying to select the disk materials and coating processes and did various other jobs in the lab – we all did a lot of jobs in the lab. It was a lot of fun.

NOTE: In 1956 William Y. Crooks was promoted to project manager of Industrial Engineering. Mr. Crooks started with IBM in 1951 as a design engineer in the Research Laboratory. He was promoted to associate engineer in 1953, and to technical assistant to the laboratory administrator in 1954. Early in 1955 he was transferred to the Industrial Engineering department. He served with the U.S. Army for two and a half years, receiving the Purple Heart while in the European theater. After his discharge from the Army in 1945, he attended UCLA for two years, graduating from the University of Arizona in 1950 with a B.S. degree.

Norm: I'm Norm Vogel. I joined IBM in early 1954. I spent the first four days on a printer project, but I saw the disk file working in the corner. I talked to Rey Johnson and got assigned to the disk file project because I knew it would be a success! I am a graduate of Santa Clara University. I worked for General Motors for about two years before coming to IBM and then had a wonderful 33 years with IBM.

Jim: So let's take you back fifty years and ask Lou – Lou, what was the challenge in deciding what had to be done to develop a suitable approach to the heads and disks to be used in this drive?

Lou: The start, of course, was the hydrostatic air bearing – the airhead off of the very first disk drive file-to-card machine, which we've talked about in the past. It was not a production model sort of a thing you could visualize using in a long pull. It needed a lot of things. One problem was it had to be loaded with a spring that provided resistance to the air escaping out of the holes in the front to provide the bearing. Norm Vogel was asked to take on the task of the airheads in the arm assembly and try to put them in place, in some sensible way. What's in Norm's hand (*airhead assembly*) was the end result of his work -- a very ingenious design, which he'll tell you about a little later.

The disk was the other key ingredient that had to be created. The substrate and the coating material, as well as the means of coating and means of application, were topics that Bill Crooks became very much involved in. He'll also spend some time talking about his experience with the spin coating of the disk and its creation.

It's an interesting combination that these two guys came together on, because the head and the disk work together. Norm Vogel had a brother, Marcel, who he's sitting in for in absentia today (his brother has passed away), who was a paint expert. Norm brought him in to help solve some of the problems with the disk and the disk paint. The paint that Marcel Vogel came up with, and Norm Vogel was instrumental in helping on, was the paint used on all disk drives for many years after the 350. As long as we had particulate media, we basically used the epoxy-based paint that Marcel, Don Johnson and Ralph Flores configured. I'll let Norm talk about that when he gets around to it.

Jim: Let's get started by addressing the question, when you started out on this project relative to the heads, what did you try to do first?

Norm: Actually, I worked on several projects at the same time. I knew I had to end up working on producing an air head kind of system, but there were challenges. At that time, April of 1954, the final speed of the disk was not ascertained. We were trying to run them from 700rpm up to 3,000 or 4,000rpm. The faster the better in reducing access availability. I would

run the disks with a single thickness of .051 aluminum; a very common material – flattened aluminum. I'd find as you went up in speed there would be flutter and squirming. The disk is a very very unstable kind of a thing to deal with. So, I felt that you would either have to go very slowly with disk rotation or work up a kind of disk that would survive holding stable while active mechanisms dealt with it. Toward that end, I imagine Bill might describe what took place in trying to come to a disk material that was stable and usable.

Bill: I was kind of a pioneer because I was doing things in 1953, which was probably a year before you *(Norm)* got onto the project. At first we did all kinds of experiments. We did dip coating. We did spray coating. Finally, we resolved to use spin coating, which is the use of centrifugal force. And to use Lou's term, we had to be 'expedient'. I had to catch up. The project was really moving and we didn't have any disks. So we had to spin *(coat)* and we used 3M magnetic oxide as a coating – right off the shelf. This was the same paint used on the Golden Gate Bridge. It seemed suitable. We had to filter it and dilute it with paint thinner, acetone, that sort of thing, to get the right viscosity to go with the right rotation speed. We couldn't get the magnetic oxide through normal filter paper; so we came up with the silk hosiery idea with help from one of the young ladies in the laboratory. From then on, we used nylon stockings to filter the magnetic oxide. I used a Dixie Cup containing different amounts with different viscosities and used different speeds to apply it to the disks – spin coating. Then I cooked them under infrared lamps to try to harden the surface. Later on we were very ingenious; we put Carnauba wax on top and baked that too. This was April or May of 1953. We tried to get the thing stabilized, but we continued to have problems with rippling. We were 'expedient,' though, weren't we, Lou?

Lou: You bet!

Jim: It probably should be noted that there is a famous picture that IBM has used many times over the years to illustrate the early development process. It shows a hand pouring paint from a Dixie Cup. As I understand it, that was your hand, Bill?

Bill: That's right. I was nicknamed and known in the lab thereafter as 'The Hand'. We all worked together and we tried to help everybody. I also worked on test stands, the first horizontal disks, and we had a lot of fun. Doing track-to-track was our first experience. I think I worked with you, Norm.

Jim: Well Norm, what did you make out of those disks with those heads?

Norm: There is a little more to this story. A variety of disk substrate materials were tried and considered, including glass and magnesium. There was a pretty good Zomag disk of the right thickness. It was a candidate for a while. But the aluminum disk was finally used in the product, and it's really two .051 disks glued together carefully, while bearing against the tombstone as a means of keeping it very very flat. Flatness was super important for any kind of an air head that you might end up with. And so, out came this disk and a kind of coating that had been produced. I might add, the coating was done subsequently. Primarily, Jake Hagopian did the coating and there was a time when we joined him in looking into these things. We went up to San Francisco to my brother's factory and laboratory. He had a firm called Vogel Luminescent Corporation. He built specialty paints – fluorescent, luminescent, radioactive, etc. He was very well acquainted with the constituents of paint. I suggested that he join IBM. He had a business, so he said, "well I'll consult." So he found himself going from Redwood City up to San Francisco in the morning, spend the day there. Then he would drive the 60 miles or so down to the San Jose Lab. I'd join him in the evening. We made the machine for paint making and the testing gear to test whatever paint might come out. Out of the blue, he came up with the epoxy resin as a potential base that had great durability and great sticktivity. Those elements were supremely important for the kind of structure we ended up with. Marcel worked with his associates, Don Johnson and Ralph Flores. They collaborated on a patent, which went through properly. And as Lou said, this disk coating survived for at least thirty years, until we advanced things. So we had good disk coating and the laminated disks were potentially very good. That is sort of where the disk story went.

Jim: Well, you had disks – how about heads?

Norm: Heads – Super interesting challenge! You’ve got to realize that the head is a component that would carry a magnetic element within it. The head had to be moved to a track, applied to a track after the information was gathered, then lifted, retracted beyond the disks and vertically up or down to the next disk position, detented and then into the next track position, and then applied to the surface at that exact position. These were movements that had to take place that were **not trivial**. They were challenging to do. The head, of course, as Lou said, was a hydrostatic head. This means, when air was applied to it (for the bearing) the head would elevate. It was that kind of a thing. We also wanted the head to be serviceable, so we wanted to be able to manually pull the whole thing aside, let the head out, plug it, take the hose off and then service it, if need be with little wires that would go through the orifices. I must point out that the head we are looking at here and the one that went out with the product is a magnetic air bearing head that had six orifices in the front face six thousandths of an inch in diameter, roughly two and one-half times the thickness of the human hair. And on the backside, through the same manifold that the air came through, we put in three 1/16-inch diameter pistons. Those three then, when confronted with the manifold, would push the head to the surface at the same time the air was bleeding out of the orifices and that would all come together and stabilize at some level. That level had to be approximately 4½ to 5 ten thousandths of an inch. To us at that time, it is nothing now compared to the current magnetic devices, it was even difficult to imagine. We didn’t have indicators that spelled out the ten thousandths of an inch. In any event, we found this head and it did work.

Historically, this appeared as a useable thing in about the spring of 1955, at the time the 305A machines were put together and presented to the population. In the prior year, Lou was very much in on things, and a very good manager by the way, who would meet with us frequently and give us assignments -- he asked us to try various things. I found myself first working on disks and disk mounting structures. An interesting disk, in a whole pack of things, was one I had spun in a local shop to a three-degree conical taper for stiffening. It worked. Thank heavens we didn’t have to use it though.

Lou asked Don Johnson and me to work on the heads. That seemed to be a major challenge. Toward that end, there was a very clever head I’ll describe in a moment, called the Bernoulli head, which, when I got into the act, was

alive and working. So Don was going to continue with it and I was supposed to try something else, which was a predecessor to the piston head, except I used a ringed piston – an ‘O’ ring, a Teflon ring. So we were both supposed to work on that and come out with a good usable thing.

The Bernoulli head is something like this. It is a device that had a conical orifice ring in the front face, tipped at a 30-degree angle. With the head about 40 thousandths of an inch off the surface, where it would land when it was moved to a new position, you turn the air on, the orifice ring would spill out the air and in the process it would create a vacuum area in the bonnet region, the middle region, of the head. That vacuum region coming from the vena contracta of the Venturi would then suck the air head to the surface and air exiting that ring would be used as an air bearing and the head would balance out at about the right spacing. Now that was a very clever sort of development. It worked when everything was right.

Lou then challenged Don Johnson and I to have a competition in which we would both build the best head we could, including magnetic element, test them out, and see how they worked. They both ended up working, but because the Bernoulli head sucked itself to the surface, it did not create a warp on the arm, which could strike the upper disk from the arm. So we agreed that the best thing to attack was the Bernoulli head. That’s in mid-1954.

Shortly after that, Lou pulled a cute stunt. He gave me the responsibility to build the head, and he put Don Johnson and Trig Noyce onto the design of the new structure. I spent a pile of time optimizing that Bernoulli head against very nice surfaces. I got it to work, I thought, very very well. But low and behold, after the project was properly funded and put under a very tough deadline (and I was the head guy), I properly equipped the new disks with the Bernoulli heads and they rattled and banged on the disks because the disks had bubbles from the lamination – little wrinkles and irregularities. The poor head did not have enough thrust to follow the surface of those disks.

Bill: They didn’t use the Dixie Cup, that was the problem.

Norm: I thought, Holy Mackerel, we’ve only got about two or three months and we’ve got to come up with a new head that will work in this situation. I

pulled out my old drawings on the earlier version of the ring piston head that I mentioned. I worked up some sketches involving three thirty-seconds or 1/16 inch on the pistons, quickly went out to the shop and built these things, and they seemed to work very well. They worked better when I went from bronze to aluminum heads. Then came the problem of wear on the head from the aluminum. I chased down a fairly new process called 'hard anodized.' So the aluminum head was then hard anodized and they could collide with the surface and not suffer badly. By the way, with this kind of mechanism, every time a head is applied after arriving at track position, the head reaches down and, whether you like it or not, strikes the running surface at an angle, straightens up, levels off and is okay. So we had an awful lot of collisions of heads and disk material. I found myself doing a lot of testing of heads, like 300 or 400 thousand cycles on and off the disks to see if the burnish spot would end up being where we wanted it. Would it still be magnetic the way we liked it? – yeah. And would the head hold together? – yeah – because of the hard anodizing. And so that's sort of the story about heads and the like.

Jim: Lou, there must have been a little bit of concern at the management level of this program about where all of this work was going?

Lou: We had a lot of concerns. Like I said, Rey had created an environment in the lab where everybody did everything. It sounds like chaos. It turns out that it wasn't. It was really a very creative environment to participate in. Everybody who had an idea got a chance to talk to somebody who considered that idea. If you had a good idea, somebody would sit down and listen to you. If you couldn't do it yourself, you could probably find somebody who could do it. Norm happened to be one of the guys who could do everything himself. He said in a kind of quiet tone earlier, "I went out and built them," and he did. I can't remember exactly what magic you used to chose this particular size – 650 thousands or something like that? There was a magic thing.

Norm: Well actually, when you get down to the drawing board and look at this, you want the head to be as large in diameter as you can fit in.

Lou: It doesn't have to be that big.

Norm: I guess we had to decide that the access arm would be a manageable sized unit. We accepted an inch as an okay thickness of an access arm and it had to be as stiff as possible. It had four webs on the access arm to give it a vertical stiffness. Then within that dimensional situation, we would lay out the diameter of the head and it came out to 5/8-inch, which was fine. Everything then fit that we wanted to encapsulate. We also wanted to be able to remove the head in place without requiring tools. It could be done with just fingers and properly repositioned within one thousandth of an inch, as intended. It ended up going out like this, six inches from the inner most track. It had to be able to pop out and push to the surface, at which time it was bending the arm up because it wasn't the Bernoulli head. That was acceptable for the amount of loading we put on it, and there you are.

Lou: This all sounds kind of easy, but it really wasn't. **It really wasn't easy!** We were under terrible time pressures at the time. As Norm said, we had a commitment to deliver to a customer within a few months. We needed to have an answer. It was rather amazing that we were able to come up with something that could survive the first rigors of product testing.

Jim: So as I recall, a decision was made to build fourteen drives to test out the operating principles of everything.

Lou: That's right.

Jim: Was there much of a difficulty to suddenly manufacture parts?

Norm: Yes, just the magnetic element alone – a couple of guys in the shop and I got into etching laminations, stacking them, dressing them with a stone, putting it all back together again and then positioning it on a manual, which then had epoxy placed all around it. Then we would put it in a lathe and turn that, bearing in mind that we would have to perfectly center the pull tip and make it an external diameter that would neatly fit. Then use another epoxy, to epoxy another magnetic element into various heads. It was a real struggle in manual work – not easy. So we found ourselves at a point where I had a bunch of the Bernoulli heads, about 16, at a time when they started putting together the first of the 305A machines. They were not going to work – period. I had to get to something different quickly. I was able to convert a Bernoulli head to a piston and orifice head by carefully machining out the center section and putting in some holes for the pistons. Then put a

new faceplate on and drill 6 mil holes for bearings. I was able to salvage a lot of those elements.

Bill: He saved the program, he really did.

Jim: It must have been a challenge to make them all the same?

Norm: Well the etching pretty clearly made sure they were the same size. You had to be careful not to damage the magnetic characteristics of the mu-metal. It wasn't too hard, with the proper glasses, to look at the pole tip and make sure it was centered in the lathe when you went to turn the exterior then pot it in here.

Lou: Thank God for Chuck Blackley.

Bill: That was the model maker.

Norm: This was a 50-year ago model maker, can you imagine?

Lou: We had a good team of model makers. Couldn't have done it without those guys.

Norm: Another interesting aspect of the design was detenting the access arm in position and the carriage in position. We had a little struggle finding what mechanical approach to take. Lou and I tussled for a moment on this, because I favored a pneumatic detenting system. At General Motors, just before the IBM job, I did a lot of solenoid valves and such type of machine building. I went ahead and carefully measured the latency time of a pneumatic access mechanism and it turned out to be okay.

Bill: Disk-to-disk or track-to-track?

Norm: Both. Disk-to-disk, we had a single plunger that would go into mechanical holes in the back of the carriage and that would then snug the carriage up when it applied. It would be held by the air pressure bearing down on the thing, which was great. Then when the arm went in there were two odd/even detents because they were very close together. One or the other could be pulled, and then detent the arm position. There had to be coordinate action of the pneumatic cylinders and the like withdrawing the

head by cutting off pressure to the manifold, backing the whole arm out to null position and then down and the whole process again. By the way, among patents, I don't have many, but one is the patent that covers the assurance that if an arm is within a disk pack it will not go up and down while in the disk pack. It had to be fully retracted. There is some pretty good evidence of that being a useful thing. As a matter of fact, the server went wild at one point and ended up on the floor before I got into the action. So the carriage was very integral with the arm and the detent system. We didn't have much machinery on the carriage, but that of course, was an essential part of what we had.

Lou: But you know, Norm has pointed out something – the use of compressed air implies the existence of an air compressor. Norm spent a good bit of time working on that air compressor as a practical thing.

Jim: What about reliability? That was a challenge.

Lou: It was wonderful.

Norm: The air compressor I bought was a commercial unit, a Worthington, and it was a well-known compressor in the outside world. After it ran for a while, an awful lot of water and oil came out. Oil is pure dynamite on a disk system like this! We then had to fuss around with devices that would try to separate out the water – although water isn't too bad. Then we picked up a diaphragm compressor that worked, but was noisy. However, it flew apart after about 200 hours of operation. Then we went with a DeVilbus compressor, which was a very well known name, and it worked well and we found the right kind of separation devices and the like. It was all sort of a challenge to make sure that the air was clean enough that was taken in from the orifices and wasn't contaminating the surface of the disk. Oil on the surface of the disk could raise hell.

Bill: We used a lot of air in those days.

Lou: You did some work on valves. We didn't have all the infrastructure on the pneumatics that we needed.

Norm: No we didn't, so we called in Valvair and other people and finally settled on a pretty neat little Skinner solenoid valve for \$9 bucks that had the

speed we desired. We checked that all out and sure enough came through with a very good operating system. Throughout all of this, always bearing on our minds, was a concern about the disk file failing and tearing up a disk surface – crashing. We were always testing one way or the other to be assured that we had done all we could to prevent that from happening.

Jim: Is it true that when you did your detent system for positioning, you actually built in sufficient fineness for the system so it was all ready to go up to a higher track density than was initially used?

Norm: We could go to double track density and we did that on the 1405 machine. A quadruple detent system.

Jim: It sounds like all the problems got solved very promptly and you must have gone home feeling very comfortable each night, huh?

Bill: We had a few nightmares along the way. It was a lot of fun.

Norm: All the while you're working on it, you know darn well that it was an important program, because the whole era of random access systems and the new business that would bring was lurking there. You knew if you could successfully achieve it, it would be very worthwhile.

Jim: That's a key point because the only other random access memory at that point was really a drum.

Norm: Correct.

Jim: They were very expensive.

Norm: Density of storage, it had a single surface. The other things, of course, were card stacks.

Bill: Core Storage.

Lou: No, that was later.

Jim: But that would be classed as main memory rather than auxiliary memory.

Norm: Cards and then magnetic tape.

Jim: Tape is a serial device.

Norm: Serial, you had to sift through a whole pile of time and space to find what you're looking for. Go directly to it in half a second normally and that would really expand the business. The Sabre System is a case in point; a department store is a case in a point for inventory.

Bill: Wasn't a lot of this tube technology? We didn't have transistors.

Norm: It was all tubes. It was later in that decade that the transistor became available.

Bill: So we were really skipping along.

Jim: So Lou, somehow it all got done and were there any problems getting it out?

Lou: I think they were the things that made the difference. If you didn't have a disk and you didn't have a head and you didn't have an access mechanism to move it back and forth, you didn't have a product. In my opinion, Norm's air head was the key ingredient, the disk coating and the disk itself were the second key ingredients, and without those two we wouldn't have had anything. We could have all gone home and let somebody else do it.

Bill: He *(Norm)* made the program.

Lou: The two of you, and the other people who contributed, made the program to give us these two components.

Norm: Another thing that came about in a subsequent project was something we called the Integrated RAMAC – putting everything in one box in order to compact the whole disk file within the remaining space in the 370 printer box. They had to get rid of that hump on the bottom, which was a motor. I wanted to lay it sideways so it could fit under the access arm, so I designed an inside out motor. We would clamp hold on the shaft ends of the

motor and have the wires go through the core of the motor through the shaft and have the outside shell rotate and have it be that upon which all the disks and disk rings would fit.

Jim: Positioned inside the id of the disk.

Norm: Correct. When it all came together, we had a couple of models made in the small shop that we had and it worked fine. Lou came over one day to see one of the motors we had rolling. He walked into the room and didn't realize that the darn thing was even running.

Bill: Success -- It was quiet.

Norm: I think further on, all the motor drives in the gears and the fields as well as the fluid chambers were knocked off and it improved the serviceability of the system. Everywhere we turned we were worried about reliability and serviceability of the system. That was part of the key to making a useful product.

Jim: Actually that's a very interesting example, because decades later the disk drive industry did go to all in-hub motors, as they called them, for the smaller drives. It's interesting that the first drive, the RAMAC, was a pioneer in that area also.

Lou: The guy that Norm was working with, was a fellow by the name of Bill Woodbury (and Greg Tobin). Bill Woodbury was a real maverick sort of a guy. He ended up as a cab driver in San Francisco and that's the last I heard of him. He had this idea of using a plugboard. He was a real plugboard bigot. The reason the plugboard is on the RAMAC is because of Bill Woodbury.

Norm: He was a real maverick. I sort of appreciated that. As a matter of fact, it was kind of entertaining. Along with that motor and the disk drive, which was kind of a glob-adder kind of drive (and it worked splendidly), we also built for the front face of the 370 a punch and a card reader. And that was a swinging die punch at 100 columns per second, which is super fast. We designed the cams to drive all this stuff, and then you would see the three of us out in the shop cutting those cams. So from the ground up, we would design, build and operate. It was exciting.

Jim: Great. Any other key problems in getting this project started?

Lou: Lots more key projects, but these were the things that allowed us to get started.

Jim: Thank you very much. I think that describes how the first hydrostatic heads got developed to work on this disk.

Bill: Thank you for inviting us.

Lou: It was a real pleasure to have you.