

**The IBM
350 RAMAC
Disk File**

Designated An
International
Historic Landmark by
The American Society
of Mechanical Engineers

February 27, 1984



Introduction

May 6, 1955... International Business Machines Corporation made an announcement that went largely unnoticed outside the computer community and other technical circles. The company reported that a team of engineers working in a small research and development laboratory in San Jose, California, had developed a new magnetic disk storage technology.

Few could have guessed in 1955 that the computer industry's first magnetic disk file, the IBM 350 RAMAC (Random Access Method of Accounting and Control), would one day prove to be of worldwide significance.

What the team of IBM engineers had developed was a technology that significantly affected information processing in the worlds of science, agriculture, health, education, government, finance, insurance, transportation and distribution. This technology ushered in a new era of interactive computer applications such as airline reservation systems, inventory management, automated banking, space flights, word processing and personal computing.

It is only through hindsight that the true dimensions of the 350 RAMAC disk file development can be appreciated for what it was -- a historic technological milestone.

A technology that spawned an industry

The first 350 RAMAC disk file became a commercially available product on September 4, 1956, and was a key component of the IBM 305 RAMAC system, which also included a central processor, card reader and printer. (In its early development, the file itself was called the 305. It became the 350 when the 305 system was announced.)

It is difficult to overstate the impact the 350's disk technology has had upon the world in the years since its announcement.

Making information directly available for computer processing on demand meant that no longer would processors stand idle while searches were made through reels of magnetic tape or data was punched into cards and sorted for processing.

Removing these obstacles helped turn the promise of the computer into reality and set the stage for what has come to be called the Information Age.

The 350 was the first step in the evolution of many direct access storage devices. It launched a technology that has been improved and refined during the past three decades, but never superseded. In strictly economic terms, the 350 launched a direct access disk industry whose size goes far beyond disk drive hardware sales and into the worlds of high-speed processors, programming and computer services. There is no way to estimate, for example, other hardware, software, and service revenues which would not have been generated if low-cost direct access disk devices did not exist.

Even discounting such conjecture, it is sufficient to say that fixed and flexible disk drives alone - all derivatives of the basic 350 technology - generated an estimated \$12.5 billion in sales worldwide in 1983 for the 72 manufacturers of fixed disk drives and the 52 manufacturers of flexible disk drives.

Prior to the development of RAMAC in the early 1950's, IBM employed a small number of people at its San Jose, California, punched card facility which had been established in 1943. Today IBM's General Products Division employs some 15,000 people, with two major sites in the San Jose area and a third in Tucson, Arizona.

The economic contribution of disk storage products to California's economy is estimated to be approximately as great as California's more famous semiconductor industry.



Development of the 350 RAMAC file at San Jose led to the company's decision to build its first computer manufacturing plant in California. Above, engineers test units of the 350 RAMAC System before customer shipment from the new facility in 1957.

Getting Started...

The RAMAC file development could not have had more humble beginnings.

It all began with one man, Reynold B. Johnson.

In mid-January 1952, Johnson, a former high school science teacher from Ironwood, Michigan, who had been hired to help develop the IBM 805 test scoring machine, received a visit in his office at the IBM Endicott, New York, development laboratory from W. Wallace McDowell, then IBM director of engineering. McDowell had come with an unusual offer.

The company had decided to set up a small research laboratory on the West Coast, and McDowell wanted Johnson to head the project. His job: to find a suitable site, do his own recruiting and establish a laboratory of no more than 50 people.

The new lab was to work on technologies not being pursued in the East. Non-impact printing was

one of the areas suggested, another was data reduction. These could make up about 50 percent of the new lab's work. The rest was up to Johnson.

One of the most remarkable aspects of the 350 RAMAC development effort was the rapidity with which Johnson put his new laboratory into full operation.

In February 1952, he:
1) signed a five-year lease on a vacant stuccoed cement-block building that had previously housed a printing plant at 99 Notre Dame Avenue in San Jose;
2) began to renovate the building;
3) placed ads recruiting engineers in area newspapers, and
4) started interviewing applicants with the help of Louis D. Stevens, a member of the Defense Calculator Design crew at Poughkeepsie, New York. Stevens, whose assignment was first considered temporary, returned permanently in May 1952 as Johnson's technical assistant.

By July 1952, IBM's new San Jose Research and Development Laboratory was a functioning organization of some 30 people, many hired after only one interview, working on a number of projects, with each engineer usually working on more than one project at a time. Johnson set forth three guiding principles to everyone hired, which added to the vitality of the lab. They were:

- It is essential that each engineer be familiar with the purpose, function and environment of the machine or machine component on which he is working to the degree that his work affects the proper performance of the function in the ultimate environment.
- It is the responsibility of every engineer to be conversant with all other projects going on in the laboratory.
- It is the most important assignment of every engineer in this laboratory to give assistance, in the form of consultation, experimentation or suggestions, when asked to by another engineer; and the second most important assignment is that of carrying forward the project to which he is assigned.

Looking back today, it is hard to believe that within three years one of the computer industry's most important technologies would grow from the efforts of this tiny fledgling operation. Who could have thought so then?



The original building that housed the IBM Research and Development Laboratory. The first unit, on the corner, was 10,000 square feet in 1952, and was enlarged to about 18,000 square feet in 1953.



A meeting of the early lab privy council.
Left to right are: R. Manning Hermes,
William A. Goddard, Reynold B.
Johnson, Louis D. Stevens, Arthur J.
Critchlow, and the late John W.
Haanstra.

Identifying the Goal

The project that was to lead to development of the RAMAC disk file was called "Source Recording," which was defined as encompassing "all processes which take alphanumeric data from any source and transcribe it in a way so that the resulting document may be handled by machine methods." The punched card was the most widely used such document at the time, and early work on the project focused on eliminating the use of cards altogether. As the problem became better understood, however, the focus turned to minimizing or eliminating the key punching task itself. With many shifting goals to come, the journey of discovery had begun.

It is helpful to remember that in 1952 there were only three ways of storing information for use by data processing equipment: the punched card, magnetic tape and, to a lesser extent, magnetic drums.

Each of these methods had significant limitations. Punched cards and magnetic tape essentially limited the user to batch or serial

processing, which meant that expensive central processing systems were often idle while information was being accumulated and sorted for serial processing.

Information stored on drums was also randomly accessible and moderately fast, but the low volume-to-density ratio of the drum technology also made it costly.

From today's vantage point, one can see that what was needed was some form of storage device that was randomly accessible with a greater surface-to-area volume ratio much higher than drums.

In September 1952, however, the focus was on automating the so-called punched card "tub file." Tub files had come into widespread use at that time as a method for making master information punched on cards more readily accessible for machine processing. Essentially, the tub files were very large rectangular trays containing master cards arrayed in sequence by customer number, item number, size, color, etc. Usually, the files contained several copies of each master card.

In a typical operation, clerks receiving an order would search the tubs, pick out cards containing the needed customer and item order information, and send the cards to the machine room, where the needed documents--shipping room instructions, packing slips, invoices, shipping labels and bills of lading--were produced.

An early crusader for tub file automation and a frequent visitor to the San Jose lab was the late Ed Perkins, a special marketing representative in IBM's San Francisco office.

He addressed the San Jose engineers regularly on the inadequacies of the tub file method and took many of them to visit Bay Area IBM customers so that they could see firsthand what needed to be corrected.

Throughout this period and continuing through the entire RAMAC development project, IBM engineers used the tub file billing and inventory control operations of a large paper company as a "real life" touchstone in defining and refining requirements and specifications for the Source Recording study and for the disk memory system they eventually developed.


By November 1952, the tub file automation alternatives had been narrowed to two approaches. One was an endless belt with master cards attached so that operators could see the ones they wanted and have them electrostatically copied. These master cards would then be electrically sensed into a new punched card. The other--Johnson's favored approach at that time--involved a matrix of parallel vertical wires of one-foot length that would sense a card's information and record it magnetically.

It was at about this time that Jacob Rabinow, of the National Bureau of Standards, described a Notched Disk Memory Array in which each disk was rotated independently. His paper was widely circulated at the San Jose lab and led a shift in interest back to disks.

Disks had already been considered earlier and discarded because maintaining the necessarily minute spacing between a recording head and a disk surface (about 1/1,000th of an inch) was considered an insurmountable problem. Other alternatives were magnetic cards, plates, strips, bands, wires and rods. All were finally discarded.



In computing's earliest days, punched cards were the basic means of storing information and were often organized in "tub files". This file was typical of the 1940's and early 1950's. The RAMAC file development grew out of an effort aimed at automating the tub file.



A "File-to-Card" machine first made operable on February 10, 1954, consisted of an 026 key punch modified for input-output service, the prototype disk file, and the control electronics shown against the wall.


Betting on disks

In January 1953, after reevaluating all the alternatives, Johnson came to a decision. The laboratory would focus its attention on disks. A work order for that month reads:

"Magnetic Disks have been selected as the best medium for the Random Access Memory to be used for File Maintenance. Disks will revolve continuously at 16 r.p.s. so that any of 200 columns of a record may be read out in any sequence, from any of 20,000 records."

Those who were there at the time recall that Johnson's disk decision was a very unpopular one. One engineer advised Johnson that he was backing a mechanical folly. The popular name for the disk array was "the baloney slicer."

Despite all disparagement, including cartoons on lab bulletin boards, San Jose engineers went back to work with their focus narrowed to disks. On February 2,



1953, the objectives for a disk storage configuration appeared. The configuration was remarkably close to the final product. While the decision to focus on disks clearly replaced the goal of tub file automation, the general idea at first was that disk memory would be part of a File-to-Card Machine and conceptually the end product would be, in effect, an electronic tub file, if not an "automated" one.

This all changed shortly to a more ambitious concept when the U.S. Air Force asked IBM to submit a proposal to provide an inventory control system for its base supply. The request called for a very-large-capacity, randomly-accessible memory rather similar to what the San Jose researchers were trying to achieve, but with the additional requirement for information processing capability.

The late John Haanstra and other laboratory systems people were assigned to respond to the Air Force's request and on April 1, 1953, submitted a proposal to the U.S. Air Force that called for a memory complex of 10 magnetic drums that might later be replaced by disks.

On April 14, 1953, shortly after submitting the Air Force proposal, Johnson made the decision that the laboratory would abandon all technologies competitive to disk. He assigned three teams to work on various aspects of the project.

William A. Goddard was charged with developing a file model, assisted by Donald D. Johnson, John J. Lynott, Geoffrey Hotham and Warren Gonder.

The late Edward Quade and his group were assigned to design a magnetic read-write head which could be no more than 1/10th of an inch high, and Haanstra and Alton E. Ewing were to concentrate on the File-to-Card Machine while also studying extended systems applications.

It was agreed that one of the first problems to be surmounted was finding a way to maintain constant spacing between magnetic head elements and disks without substantial runout.



Above
Two of the early read-write airheads (circa 1953) positioned as they would be against a magnetic recording surface. Exhaust ports to remove air from between the surfaces are visible. These were necessary to achieve the required spacing between head and disk, then about one-thousandth of an inch.

Top
Rey Johnson's early 1953 decision to proceed with development of a disk file was the subject of much skepticism and disbelief. One of a number of cartoons found on the lab bulletin board at the time refers to a popular song of that year: "How Much Is Dat Doggy In Da Window?"


Goddard recommended an air-bearing approach and asked Don Johnson to try it out. Not only did this approach work, but by June 2, 1953, a third model of an air head with a magnetic read-write element designed for a magnetic drum was used to write 51 bits per inch on an aluminum disk that had been spray-coated with magnetic iron oxide paint.

Thus convinced that an air bearing would "float" a read-write head above the surface of a disk without crashing onto its surface, the development team turned to three other pressing problems: creating disk surfaces flat enough, developing an effective disk coating and devising disk-to-disk and track-to-track accessing mechanisms.

Finding the best method and material for creating flat disk surfaces was a matter of old-fashioned trial and error. Among the materials tried were aluminum, brass, glass, plastics and magnesium. In the sheet aluminum trial, engineers were horrified to find their so-called flat disks had runouts as high as 1/50th of an inch and greater when the disks revolved at 1,200 r.p.m.

The first successful material tried was fabricated from magnesium lithographer's metal, but this was finally replaced by aluminum laminates clamped under pressure and heated in ovens above the annealing temperature.

Jake Hagopian, assigned to the disk coating problem, developed specifications for a coating made from a paint base. Spray coating was tried as was dipping, but neither produced the uniformity of thickness required in the magnetic coating. Hagopian finally settled on a spin method in which the coating was poured onto the inner surface of a rapidly rotating disk and then spread evenly over the surface of the disk.



Bill Crooks produced the first truly successful disks by filtering the dispersion through nylon hosiery and using paper cups to measure out correct amounts for the spin coating. This primitive process was used for a year before it was automated.

Later in the development project, Marcel Vogel, Don Johnson and

Ralph Flores invented the final coating formulation and basically, it is still in use today.

The third major hurdle—development of disk-to-disk and track-to-track access mechanisms—actually posed two challenges.

The first was how to apply the load needed to make the air head function. Initially, the researchers tried to devise a method using springs to apply the load, but this caused problems in moving the arm from disk-to-disk.

It was Norman Vogel who finally solved this problem with a design that retracted the air head into the arm during disk-to-disk motion of the carriage. The design provided a self-loading force using three miniature air pistons on the back of the air head that were activated in the track-to-track mode.

The second challenge was how to provide the disk-to-disk and track-to-track accessing motions. Three



Top

First test model, used with the "File-to-Card" machine, showing heads on and inserted between disks.

Above

Norman Vogel, one of the key members of the 350 development team, at his workbench testing an airhead design.

teams were set up to explore alternatives: Crooks and Haanstra were to work on an electrical-servo drive system; Don Johnson and R. Manning Hermes were to pursue an approach based on the IBM 402 type bar mechanism, and Gonder was assigned a system using cams and spring clutches.

The servo approach was the early favorite for the track-to-track drive but was somewhat suspect for the tougher job of moving the heavy carriage from disk-to-disk. On October 28, Davis was able to locate a previously recorded track with his track-to-track servo, and in November Lynott found a way to use the same electrical servo system in a multiplexing mode to drive the carriage from disk-to-disk.

On November 1, 1953, Rey Johnson turned entire responsibility for the development of the magnetic disk assembly over to Lou Stevens, who was promoted to senior engineer.

Things seemed to be going on nicely until one Saturday evening, during a test, the new servo went out of control and rocketed free of the rail, landing in a heap on the floor.

Within 12 days, Lynott and Davis, with assistance from Trigg Noyes, had completely redesigned a new servo system.

Thus, on February 10, 1954, the San Jose team was able to achieve the first successful transfer of information from cards to disks and back.

This was a huge boost to morale for a brief period, but soon reality began to set in.

Stevens remembers that the first laboratory model looked like a Rube Goldberg arrangement and that "not many people believed we would make such a thing practical..."

"...but it was like a religion to us. We were going to make that thing work for sure...because if we failed...the whole San Jose experiment would fail. None of us were going to allow that to happen."

It soon became obvious, however, that neither the first model of the disk file nor the File-to-Card Machine was performing well enough to be demonstrated. So on March 19, 1954, an in-depth reevaluation of every basic design was undertaken and specifications were drawn up for a revised Model II.

Trigg Noyes had been thinking of a new design for some time, and his first decision was to make the shaft holding the disks upright rather than horizontal, to provide more workable dimensions for the machine, make disk replacement easier and provide more spaces for a number of independent access mechanisms.

Later that year, with construction of the revised Model II well underway, the development team got two tremendous boosts.

The first came in the form of a letter from F. J. Wesley, who had been assigned to review the progress of the work at San Jose and report back on it.


It read, in part: "...We must immediately...attack accounting problems under the philosophy of handling each business transaction as it occurs, rather than under the present condition of batching techniques..."

"We must build storage and peripheral equipment which can spread out into individual accounts every business fact (random access storage) and allow operation of a new concept for handling business information concurrently with its inception..."

"I wish to recommend for your consideration that we double or treble our efforts in this development...I am firmly convinced that (otherwise) we cannot expect to accomplish the real purpose and the real use for electronics in the business world."

The second boost came in November 1954, in the form of official Corporate sanction for RAMAC product development work. The San Jose laboratory was charged with designing, developing and building several field-test models of a machine utilizing a disk random access memory attached to a serial "stick" printer to provide an initial system utilizing random access storage. The machine specified was given the name DRAM (for Direct Reference Accounting Machine), and later the field-test models were known as the 305A.

On January 10, 1955, the first of five Model II 350 RAMAC files was successfully demonstrated—slightly less than three years from the San Jose Laboratory's inception. A new era in information processing had begun.



Wesley Dickinson, associate engineer in 1955 when the photo was taken, points to the reading and writing arm of the redesigned 305 (the disk file's unofficial designation while it was under development; the 350 designation came later).



One of the press release photos for the 305 RAMAC System when it was first announced in 1956. Key component was the 350 disk file. The system also included a central processor, card reader and printer.

Disk technology today

The rapid and widespread acceptance of disk technology is a direct result of the vast improvements in access speed and capacity, as well as dramatic reduction in costs over the years since the 350 was first unveiled.

A comparison between the 350 file and IBM's largest current file, the 3380, tells the story:

The 350's 24-inch disks revolved at a speed of 1,200 revolutions per minute; the resulting data rate was 100,000 bits per second. Today's 3380 disks revolve at 3,600 r.p.m. and the data rate is 24 million bits per second.

The 350's fifty 24-inch disks contained a total capacity of 5 million binary decimal encoded characters (7 bits per character) of storage. Today's 3380, using nine 14-inch disks, offers 1.25 billion

bytes or characters (8 bits per byte) of storage. The complete 3380 has two spindles with 2.5 billion characters.

Today's microminiature read-write heads on the 3380 use thin-film technology. They permit vastly increased density of storage, permitting much faster access time. For example, the original 350 disk file stored 2,000 bits of information per square inch and had an average access time of 600 milliseconds. The IBM 3380 packs 12 million bits of information per square inch and has an access time of 16 milliseconds.

Looked at another way, the 350 technology stored one megabyte of information on an area the size of a pool table. By comparison, the 3380 technology stores one megabyte of information on an area almost the size of a postage stamp.

Meanwhile, the price of disk data storage has tumbled dramatically over the years. Users of the 350 RAMAC file paid \$130 a month to rent a megabyte of storage; today,

the cost of leasing the same amount of storage is about a dollar. This drop in cost was in real dollars, unadjusted for inflation over the more than a quarter of a century in which it took place.

All of these improvements and cost reductions have triggered a huge demand for disk storage devices.

During 1983, for example, it is estimated that the dollar volume of rigid disk drives rose 26 percent, while the dollar volume of flexible disk drives increased some 110 percent during the same period.

Looking back today, it's difficult to believe that a relatively unnoticed announcement on May 6, 1955, about a machine called the 350 RAMAC could have caused such dramatic changes in the pace of technology and in our lives. But that, in fact, is exactly what happened.

