

MAGNETIC DISK STORAGE

A Personal Memoir



By Albert S Hoagland

(c) 2011 A.S. Hoagland. All rights reserved.

ACKNOWLEDGEMENTS

My goal in writing this memoir was to share my unique perspective of the challenges and achievements in the development of magnetic disk storage technology. This perspective is that of a PhD engineering student at U.C. Berkeley (1948 to 1956), a researcher for IBM (1956 to 1984), the director of the Institute for Information Storage Technology (IIST) at Santa Clara University (1984 to 2005), and the founder and director of the Magnetic Disk Heritage Center.

I would like to thank the Computer History Museum, notably Dag Spicer, for their support. More comfortable with technical detail, I found it quite challenging to write for a lay audience. In this, Nicole Hoagland, my daughter, was an invaluable editor in helping me achieve my goal.

I am also grateful for the help and support of many over the years. I would like to thank Kiki Dembrow, Santa Clara University, and Olga Shchennikova for their valuable input and encouragement when I first undertook the writing of this memoir.

Albert S. Hoagland
June 18, 2011

TABLE OF CONTENTS

ABOUT THE AUTHOR	4
INTRODUCTION	5
CHAPTER 1	6
<i>Early History of Data Processing and Data Storage</i>	
CHAPTER 2	9
<i>The Magnetic Drum, Precursor of The Magnetic Disk</i>	
CHAPTER 3	13
<i>The Rotating Magnetic Disk Stack, the RAMAC</i>	
CHAPTER 4	18
<i>The ADF (Advanced Disk File) and Future Generation Disk Drives</i>	
CHAPTER 5	22
<i>Personal perspectives on magnetic data storage during the 1960s and 70s</i>	
CHAPTER 6	26
<i>Importance of Magnetic Disk Technology Recognized</i>	
CHAPTER 7	32
<i>Looking Forward, Remembering the Past</i>	
APPENDIX 1	37
<i>1989 Speech by Rey Johnson</i>	
APPENDIX 2	42
<i>Major Honors and Awards Related to the RAMAC</i>	

ABOUT THE AUTHOR

Albert S Hoagland was born on September 13, 1926 in Berkeley, California. He went on to attend the University of California at Berkeley where he received his PhD in Electrical Engineering in 1954. While a graduate student, he became a consultant to IBM with key magnetic head design and recording responsibilities for the Random Access Method of Accounting and Control (RAMAC) disk drive, the first hard disk data storage device.

Dr. Hoagland joined IBM in 1956, making major contributions to magnetic disk storage technology and the design of magnetic disk drives. During his time with IBM he spent time in the San Jose research laboratory, in the Netherlands as a consultant, in Yorktown NY (Director for Technical Planning for the IBM Research Division) and in Boulder, CO. In 1976 he returned to the San Jose research lab and played a principal role in the formation and leadership of an industry consortium that established the first University Centers to support magnetic disk data storage technology in the early 1980s.

He left IBM to establish the Institute for Information Storage Technology (IIST) at Santa Clara University (SCU) in 1984, where he served as Professor of Electrical Engineering and Director of IIST for 21 years. IIST became the leader in providing courses, seminars, conferences and workshops for the rapidly growing magnetic disk drive industry. In 2002 he persuaded IBM to loan SCU one of the 4 RAMAC disk drives from their archives for an IIST effort to restore this first disk drive to operational status.

Dr. Hoagland was also instrumental in establishing the Magnetic Disk Heritage Center (MDHC), a California non-profit organization whose mission is to preserve the story and historical legacy of magnetic disk storage technology, serving as its director. He also initiated and led the effort to make the original building where the RAMAC was created (99 Notre Dame, San Jose, CA) a San Jose City Landmark.

On his retirement from Santa Clara University in 2005, he relocated both MDHC and the RAMAC restoration project to the Computer History Museum in Mountain View, CA. The project was completed in 2011 and this working RAMAC disk drive is now on display as a major component in their new Exhibit "Revolution," which opened in January 2011.

A Fellow of the Institute of Electrical and Electronics Engineer (IEEE), he is a past president of the IEEE Computer Society, the American Federation of Information Processing Societies (AFIPS), and has served on the IEEE Board and the Board of the Charles Babbage Foundation. He was one of the founders of The Magnetic Recording Conference (TMRC).

Dr. Hoagland is author of the book "Digital Magnetic Recording," the first text published in this field (1963), as well as numerous technical publications in the fields of magnetic recording and data storage.

INTRODUCTION

This book is my personal memoir chronicling the development of an important and vital technology for storing and retrieving digital data: the **magnetic disk drive**. I was extremely fortunate to have been personally involved in the development of this technology and in the larger field of magnetic digital data recording since its inception and over my entire career, covering the last 60 years. With this historical overview I relate my personal perspective of the challenges, successes and attitudes that shaped this technology's growth and development.

The dawn of the electronic computer era in the 1950s created a need for an efficient mechanism to handle the increasing volumes of data being produced. The invention and continued technical improvements of the magnetic disk drive were able to meet this need and the technology became indispensable to the ongoing computer revolution. Today, the rapid storage and retrieval of data (text, video, audio, graphics and pictures, etc.) on magnetic disks using digital techniques is essential to modern life. It is now possible for anyone, anywhere, and at any time, to immediately access or store almost any data desired.

The magnetic disk drive has proven to date to be the only successful economic approach to mass data storage. In the first fifty years following its invention, disk area storage density (the amount of information that can be stored in a given area) increased by a factor of 100 million. The storage capacity of a disk drive is also dependent on the recording surface area, or physical size, of the drive. Technological advances since the 1950s have increased the storage capacity of a magnetic disk drive from a capacity of 5 million bytes on the first disk drive in 1956 (the size of two refrigerators) to a terabyte (10^{12} bytes) today on a pocket-sized disk drive at a cost of only \$60.

I believe the historical importance of magnetic disk storage is still not yet fully appreciated. Current mobile computing devices are moving away from the use of local disk storage while increasing the need for disk-based servers (e.g. the "cloud") where the data is accessed remotely, making possible the benefits we now enjoy from the Internet and other data communications services. Huge disk drive "farms" now exist to meet the needs of a global society. To maintain the safety and integrity of the stored data, both local redundancy and geographic dispersion are required. Continuous availability of the data in the event of hardware or software failures as well as catastrophic environmental events is essential.

Thus, much like the impact the invention of the printing press 500 years ago had on the global dissemination of the printed word, magnetic disk drives began and continue to support another social revolution of great global impact through the almost immediate access and availability of information. In today's age of electronics it is remarkable that the usefulness of such a mechanically based storage device as the magnetic disk drive has survived for over 60 years and remains absolutely essential for both the storage of data for electronic processing and as the repository of an ever increasing source of shared information worldwide.

CHAPTER 1

Early History of Data Processing and Data Storage

In the late 19th century, there was a growing need to mechanize data processing for business purposes. At this time “data processing” consisted of a group of clerks who would manually track all business transactions in a set of ledgers. This method was slow and cumbersome, especially as enterprises expanded. Businesses were therefore looking for a better way to store data in a manner that would allow the data to be processed mechanically. Machines were developed that used punched cards to perform this task. These cards consisted of stiff paper containing information in digital form, represented by the presence or absence of holes in predetermined positions. Machines designed to use punched cards consisted of card sorters, collators, tabulators and a mechanical calculator for computation. Business records could now be organized and updated rapidly but only by mechanically selecting and organizing the cards for sequential processing. The use of punched cards for storing data remained the principal method of business data processing well into the twentieth century, as mechanization brought real efficiencies to the data processing tasks of industry and government (including the organization of the US census database in 1890).




Punched card



Punched card sorting machine

During World War II, the government began to fund efforts to develop machines capable of rapid computation. High-speed computing was important for scientific and military purposes (i.e. improved ballistic missile calculations). This type of high-speed computer required computational rates well beyond what could be achieved by current mechanical devices. Therefore these computers would need to be electronically based. At this time, electronics meant primarily vacuum tubes, an expensive, hot, and power-hungry technology. The first general-purpose high-speed digital computer was called the Electronic Numerical Integrator and Computer (ENIAC). The ENIAC was first publicly announced in 1946, and is regarded as a historical marker that transformed thinking about computers and computation. Computers were now associated with electronics as opposed to mechanical devices and were capable of high-speed calculations. The extensive mathematical calculations required a very high-speed processor but the input and output of data for the ENIAC still depended on the use of punched cards. Thus the development of a high-speed computer processor and memory had no real impact on

business data processing, which required the ability to update, store and access large corporate databases.



ENIAC

The computer design approaches following the ENIAC in the late 1940s were focused on developing a method of data storage and input/output for high-speed computers. To exploit high-speed calculation, it was clear that a much faster method of transferring information to and from the processing unit than punched cards was essential. This need led to a major interest and effort to evaluate the use of magnetic tape, a field of activity with a long history independent of the computer field.

The use of magnetic tape to record *analog* (continuously-varying) signals came into popular and general use only after World War II. It's inventor, the Danish engineer Vladimir Poulson, was awarded a patent for the technology in 1898. He was searching for a method to store voice messages as a new feature for the recently invented telephone. These messages were to be analog recordings of speech for later playback, similar to the function of an answering machine. The first public demonstrations of Poulson's magnetic recording apparatus, which used steel wire for its storage medium, was at the Paris World Exposition in 1889, and generated much excitement. The commercialization of the technology was directed at recording phone messages but was initially unsuccessful due to poor quality.

During World War II, German inventors made major advances in magnetic tape recorders and successfully recorded speeches with a quality that could later be broadcast at chosen times from different locations to confuse enemy intelligence. Their advances in magnetic tape recorders were discovered by the United States intelligence survey of German technology at the end of the war, and stimulated great interest in this country to pursue the technology. Three years later, the commercial success of magnetic tape started to be

realized with radio broadcasting, fifty years after the invention of magnetic recording tape. A small company named Ampex, with the involvement and cooperation of the popular singer Bing Crosby, focused on developing equipment to record his performances for later re-broadcast by radio at suitable times for listeners. Bing Crosby enthusiastically supported and helped underwrite this enterprise. The broadcast industry soon became a major market for magnetic tape recording in the 1950s.

With the success of magnetic tape sound recording for analog information, magnetic tape recording became of major interest for digital data storage and processing as an alternative to punched cards. Magnetic tape was capable of much higher digital data rates (speed of data transfers to and from the computer itself) and higher density, more compact, storage capacities compared to stacks of punched cards. However, both punched cards and magnetic tape only allowed for the sequential or “batch” processing of records since any desired data record had to be accessed by sequentially scanning a long file of data records. Transaction processing, in which individual records can be accessed in any order for immediate processing, was not possible.

The sorting of data on magnetic tape was complicated, time consuming and expensive. The mechanical complexity of these drives, with rapid search and re-read functions--not normally encountered in magnetic tape broadcast applications--led to sophisticated and expensive units. For this reason, magnetic tape drives initially only appeared on high-end computer systems and their adoption still perpetuated the same batch processing methods used by punched cards, albeit at a higher speed. While magnetic tape was suited for archival or off-line storage, its use did not bring a real change in the data processing methods used with punched cards.

In 1947, I had just entered the graduate electrical engineering program at the University of California, Berkeley. I was lucky to begin my graduate studies at a time when many new technological advances from World War II were beginning to bring about major changes in university engineering curricula and programs. The changes marked the beginning of the modern “electronic computer age”, a vital topic for university electrical engineering departments to address, and reinforced by the return of many engineering faculty who were working at government sponsored research laboratories during World War II as well as veterans returning to resume their studies in academia. For a new graduate student these events offered an exciting future in electrical engineering. This time was the beginning of my involvement in the field of digital data storage, an involvement that continued throughout my entire career.

CHAPTER 2

The Magnetic Drum, Precursor of The Magnetic Disk

By the 1950s many companies and new startups began to pursue opportunities in the emerging electronic computer business. The government funded many advanced projects to meet its own needs as well as to stimulate the private sector to expand its activities. Leading companies in high performance computer design, such as IBM and UNIVAC, along with the Massachusetts Institute of Technology (MIT), were all seeking to advance computer technology to address the most challenging scientific computational problems. The key to computational power was to be a high-speed processor associated with an adequate memory device providing high-speed data transfer to keep the processor fully occupied.

The government was also soliciting proposals to fund university computer design activities for a more economical class of electronic computers for academic use as well as to address a broader range of general computational and data processing needs. This funding had the objective of not only creating new technology and design approaches, but also of educating and expanding the base of trained professionals to meet the growing market and business opportunities created by this new field. Moreover, universities possessed the flexibility to pursue research on new design and application approaches without the time pressures associated with developing and marketing new products. Leadership in the new field of electronic computers was seen as vital to the country.


During this time I was looking for a research advisor to pursue a PhD degree. In exploring opportunities within the EE department, I spoke with Professor Paul Morton, who was planning a project to design and build a so-called intermediate-scale electronic computer to serve the expanding needs of the University for mathematical computation. His idea was to seek a government contract to fund a California Digital Computer (CALDIC) and he received wide support from within the Berkeley academic community. He offered me the opportunity to participate in the development of his computer project proposal and I accepted. It was already becoming evident that industry was expanding its recruiting activities in this new field and the opportunity to get involved on the ground floor was exciting. Paul Morton was an exceptional professor who had an ability to keep the project on track while always being willing to let the students pursue their own initiatives in the engineering challenges they faced. Among his students were Lou Stevens, (who later led the IBM RAMAC magnetic disk product development project) and John Haanstra, who was responsible for the systems design of the IBM 305 RAMAC, the first transaction processing electronic computer system (based on the RAMAC disk drive).

An important aspect of the project to me was to be able to become involved in developing the proposal for CALDIC from the very beginning. The planned computer was to be capable of meeting the wide variety of computing needs of departments on the campus, yet be relatively modest in cost. The key factor was identifying a memory design to provide the speed and capacity needed at a moderate cost, which eliminated consideration of all the electronic memory technologies then being pursued for the high-speed

computer market. To design something of affordable cost, the memory would therefore have to be some sort of mechanical device.

Paul Morton had heard of recent successful experiments recording pulses on a magnetically coated drum with a mechanically set separation between the drum and read/write heads at Engineering Research Associates (ERA), a research company founded in 1944 by former engineers from the U.S. Navy's cryptanalysis section. The ERA report greatly excited our group at Berkeley since it strongly suggested we could obtain a very substantial memory capacity on a reasonably sized drum with millisecond access times.

A magnetic drum is a large metal cylinder coated on the outside with ferromagnetic recording material. It could be considered the precursor to the hard disk platter, but in the form of a drum rather than a flat disk. The drum stores information on the magnetic recording surface and uses fixed magnetic read/write heads along the length of the drum to store and access recorded data in the circular recording tracks. Short access times to the data are achieved by the relatively high speed of the spinning drum, which determines the time taken by the drum to rotate wanted data into position to be accessed by the fixed heads. This time could be in milliseconds for any stored data item. Any contact between the rapidly rotating drum and a read/write head could lead to a loss of data, a damaged head or both. Thus, it was essential to always maintain a separation (non-contact) between the heads and medium (drum), which placed limits on bit and track densities and thus the ultimate capacity of the drum.



Magnetic Drum Memory

With a magnetic drum memory, the computational speed possible and input/output data transfer rates were impressive. This made a magnetic drum an attractive option to both our group and other groups as well for the intermediate cost computer we were planning. A magnetic drum is also a non-volatile memory device, unlike the high-speed electronic memory technologies of the time, so no loss of data results from an accidental or deliberate shutdown. Due to cost considerations, punched paper tape would still be the source of input data, and paper tape, printer or punched cards were used for data output. We shared our experiences with other universities as many other groups were also designing computers based on a magnetic drum memory. Even though our team collectively knew little, if anything, about magnetic recording, we wanted to start as quickly as possible on the design and construction of the CALDIC. The project seemed doable, so we immediately proceeded to prepare a formal proposal.

The project received funding from the Office of Naval Research and work began in 1948. It is fair to say that all the students involved were anxious to learn and gain experience in computer logic and computer system design. This was the type of knowledge broadly sought by industry, and certainly offered more career opportunities for the students, most of whom were planning on getting a MS degree and leaving after one year. Unlike them, I expected to be around for several years with the goal to receive a PhD.

The first step was to assign the students areas of responsibility. My longer term “tenure” led to my being chosen to take on the memory drum project. I needed to be on a funded project but suddenly began to wonder about my future since it was anticipated that a magnetic drum memory would be obsolete once electronic memory became available. With electronic memory devices, such as a magnetic core memory, the core is magnetized either clockwise or counter-clockwise and changing the direction and sensing the state is done electronically, not mechanically. Our ignorance of drum technology was not seen as a handicap since there was no extensive base of experience to draw upon.

Once major computer memory design goals for the magnetic drum memory for the CALDIC computer were settled, I still faced the need to define an area for my PhD research. I found no magnetic recording publications that covered the non-contact recording and read back of digital information. I looked forward to seeing what I could learn, although my enthusiasm remained somewhat tempered by the question of whether anyone would be interested given the prevailing attitude that magnetic drum recording would become a dead end for the next generation of electronic computers. My concerns were reinforced by several exploratory interviews I held with company recruiters.

However, while still a student at Berkeley, a visit in 1953 dramatically changed my view of my time spent there and opened up professional opportunities that determined my future career for life. Lou Stevens (who also worked on the CALDIC) was now working at a small new IBM laboratory in San Jose and stopped by to visit me at Berkeley in late 1953. He invited me to come to visit his facility to see what he and his colleagues were doing. He did not mention any particular purpose for my visit except to say that he believed I would be interested in seeing what was going on. I agreed to the visit out of curiosity, but did not relate it to anything in particular. In retrospect making the trip to San Jose to visit the laboratory was certainly the major turning point in my career. For the first time I saw non-contact magnetic data storage being pursued in the form of a rotating

magnetic disk stack. The recording surface area available with many disks in a stack offered a huge advantage over a drum in terms of data storage capacity. Here was the largest computer company in the world pursuing a magnetic disk storage device based on the same mechanical non-contact magnetic recording approach used on magnetic drum memory devices.

From my point of view, the magnetic drum initiated a paradigm shift in magnetic recording, from analog contact recording to digital non-contact recording. The successful development of a magnetic rotating disk drive would allow much higher capacity data storage than a magnetic drum for computer use, ending consideration of magnetic drums as memory devices. However, major technological advances in both mechanics and magnetic recording storage density, yet to be even anticipated, would be necessary to assure a long-term life for digital magnetic recording.

CHAPTER 3

The Rotating Magnetic Disk Stack, the RAMAC

The IBM laboratory in San Jose, where early work on developing a magnetic disk drive was being done, was located in an ordinary one-story commercial building of about 8,000 square feet. It was located at 99 Notre Dame Street in downtown San Jose. During my visit in 1953, I met the lab manager, Rey Johnson, and was immediately struck by his willingness to support new ideas if the individual proposing them was enthusiastic and motivated. I knew that the laboratory had been open for about a year, and I became confused as to its purpose from what I initially witnessed. I toured several small research projects unrelated to computers and it was not evident what the potential value of these studies would be. These projects were described in terms of their inventive, not application aspects.

The early disk drive hardware being developed in San Jose seemed like a “Rube Goldberg” mechanical exhibit. Most of the individuals working on the project were mechanical engineers and they did not seem confident that the hardware mechanisms would actually work. As the nature of the device was explained to me it became clear that the reason for my invitation to visit was to interest me in becoming a consultant to assist with the challenges the project faced in magnetically recording and reading computer data on the disk stack.

I realized that I had been given the opportunity to become involved with the early stages of a new design for a radically different magnetic storage device, the key feature of which was a closely spaced stack of spinning disks. The goal was to develop a random (direct) access storage device that would economically provide a large data capacity with an access time to any record of less than one second. The disk mechanics alone were about the size of a refrigerator. To me, the project size in terms of personnel and facilities seemed smaller than that of the CALDIC. Lou Stevens was thinking at this time of a future commercial product for IBM. In fact, it seemed all the IBM projects at the San Jose laboratory during that time were small and most of the engineers working there were not involved with the disk project. Yet Lou was very positive and I met Rey Johnson, the laboratory manager, whose confidence and enthusiasm for the project impressed me so much I cast aside all my doubts about becoming involved.

I quickly saw that Rey was a visionary, enthusiastic and not subject to self-doubt. In many ways the individuals working in the laboratory had more freedom to explore their ideas than commonly found even at a University. At the time of my visit, I was not yet at the stage of deciding on a full-time job and the factor of most importance to me was the potential opportunity to leverage my on-going studies at Berkeley after graduation since my non-contact digital magnetic recording experience with the magnetic drum was the basis of IBM’s interest in hiring me.

What I witnessed in this laboratory suddenly made clear that my earlier work at Berkeley with the magnetic drum might be relevant in the longer term. In the magnetic disk drive I saw a mechanical design approach providing a very large recording surface area in a relatively small volume, allowing fast mechanical access to tracks on rotating magnetic

disks. The data capacity would certainly be far beyond any electronic implementation then possible. I was invited to become a consultant to the project at the end of my visit and I accepted. While consulting I would also continue to complete my PhD at Berkeley. I began to believe that my earlier research might lead to a future career in computers, if mass data storage was to ultimately become dependent on my area of expertise: non-contact digital magnetic recording.

The IBM laboratory in San Jose was established to give IBM a presence on the West Coast since they found it almost impossible to recruit western engineers for their New York laboratories. They also were aware of major engineering organizations in the Los Angeles and Seattle areas that were active in the new field of electronic computers. San Jose was located more or less in the geographic middle of those two locations. The fact that IBM had a punched card plant in nearby Campbell made the San Jose location a logical choice.

Rey was initially limited to a staff of 50 people, all of whom were from the West Coast. The lab would explore ideas that could be of interest to IBM. Rey personally oversaw the recruiting and selection of the original staff. He approached this challenge as one of heading a small “startup” organization, with funding but no fixed direction or plan. He was to establish close contacts with IBM customers on the West Coast in order to become familiar with, and to understand, their needs in the field of data processing, which mostly related to managing punched card data processing enterprises.

In addition to initiating several small projects, Rey spent the first year after the lab was established (1952) considering a variety of mechanical devices to see what would be the most promising approach for a random or direct access mass storage device. Rey turned the entire laboratory into a task force to evaluate all conceivable devices in terms of their capabilities to randomly store and retrieve data. The task force studied non-volatile storage approaches and considered drums, disks, tape loops, tape strips, wires and other configurations. Rey correctly interpreted his role as searching for a new radical innovation that would be superior to current technology.

Rey had two key experiments performed in the next several months. The first involved spinning a stack of disks that were less than 0.5 inches apart to see if disk flutter would cause them to hit each other. The second test was to demonstrate that directing pressurized air streams against a disk could keep the head from contacting the surface by “floating the head”. These simple experiments convinced him that he had made the right decision in favor of the disk stack, well before the mechanism actually existed. As an ingenious inventor with great vision, he saw promise where others saw problems. He never expressed any doubts about the future success of the magnetic disk drive, unlike his engineers.

The technical challenges of developing a rotating disk stack for data storage were formidable, requiring a significant number of magnetic recording innovations beyond those needed for a magnetic drum memory. An advantage of the disk drive was that the recording area would be far larger than that of a drum of the same volume, thus providing much greater data storage capacity. The limitation of a disk drive design was that the access time would be much slower than with a magnetic drum. However, for many


punched card applications, the access time a disk drive could provide was adequate and the high capacity was key and affordable. He united the group around his choice, in spite of many who felt it was folly to proceed with this design.

In February of 1953, Rey terminated the discussions on the relative merits of all different proposals and made the decision to pursue a magnetic rotating disk stack. He understood the outstanding advantages a rotating disk stack had over other competing concepts and his confidence was such that all the problems and concerns as to its feasibility were not relevant in terms of making his decision. The San Jose team moved ahead with the project. The disk drive would be called the RAMAC, which stood for Random Access Method of Accounting and Control.

With a successful disk drive, the new challenge in moving from prototype to a commercial product was to design a complete computer system able to fully exploit this new large-capacity random access memory device. Of key importance in this endeavor were individuals such as John Haanstra, who were able to see the drive as more than just a storage box. Haanstra's efforts led to the development of a new computer system that was based around the RAMAC disk drive and allowed transaction processing. To design a complete system required not only disk drive specialists, but also an entire team including programmers, computer designers, and experts in computer applications. Now one event, such as a product sale, would simultaneously update inventory status, re-ordering, billing and accounts receivable, all as one business transaction. The computer system that was developed became the IBM 305 RAMAC system. The disk drive itself was called the IBM 350 RAMAC disk drive. The driving factor for IBM was to announce the dawn of a new age in data processing based on magnetic disk storage.

In 1955, the RAMAC 350 disk drive and RAMAC 305 computer system moved to the product development and manufacturing phase under Lou Stevens. A new manufacturing facility was constructed in San Jose to take advantage of the local expertise. At this time IBM was the dominant organization in providing business data processing systems. When Thomas Watson Jr., the president of IBM, became convinced of the promise and opportunity of magnetic disk storage, the company spared no resources in moving forward. Major commitments and decisions were made in spite of difficulties and lack of detailed performance data. With a large customer base, it was clear that IBM would do whatever it would take to lead in introducing transaction processing as the way of the future for business data processing. IBM announced the RAMAC 350 disk drive on September 4, 1956 and the RAMAC 305 system on September 13, 1956. The following day IBM also announced the availability of the RAMAC disk drive on their popular Model 650 Magnetic Drum Data Processing Machine computer system.

I believe Johnson's initiative succeeded for three specific reasons. First, the group was so isolated from the rest of IBM engineering that they were not burdened by reviews and the close oversight that any project of this potential size normally undergoes in a large bureaucratic organization. The project also benefited from the fact that the laboratory was designed to be a hotbed of ideas and not a source of products. Finally, Rey's visionary mind and problem solving attitude associated with his close awareness of customer needs was a major contribution to the final result.



IBM RAMAC 350 Magnetic Disk Drive, 1956



IBM RAMAC 305 Computer System, 1956

Ultimately, about 1,000 RAMAC 305 systems were built and leased. At the time, this was a large number for a specific model of computer system. The RAMAC 350 disk drive was primarily marketed to users of punched card systems in smaller business data processing centers where the performance of the RAMAC could offer real advantages. The RAMAC system was not intended to meet the needs of large data processing centers. At that period IBM only leased equipment. When the lifetime of the machines was over, they were returned to IBM and scrapped. This is why today only three RAMAC 305 systems and one individual RAMAC 350 disk drive have been preserved, as far as is known. The only remaining RAMAC 350 disk drive has now been restored to operational condition and is located in the Computer History Museum in Mountain View, California.

The purpose of the San Jose laboratory was not to encroach upon existing IBM product areas, but to foster innovation and develop new products or technology. Rey Johnson was the right person at the right time with free license to create and lead an innovation that significantly impacted the electronic computer business. The laboratory could be viewed today as a “startup” funded by IBM that was historic in its return on investment.

Besides creating a technological breakthrough, the San Jose research laboratory had a great long-lasting impact on Santa Clara Valley, and the RAMAC today is viewed as a landmark achievement of the area. The IBM San Jose operation became the center of the disk drive industry and held that position for decades. The manufacturing site was a major factor driving industrial and technology growth in San Jose and the Santa Clara Valley. The RAMAC created a pool of expertise in magnetic disk storage and, as the disk drive industry expanded, many companies moved into the Valley to capitalize on the available talent there. For a number of years the disk drive industry was the largest in the region and one proposal was to call this location Iron Oxide Valley, named for the coating on magnetic disks. As we all now know, this name did not catch on and we now have Silicon Valley.

A more detailed and fascinating account of the early environment surrounding the creation of the first magnetic disk storage device was provided by Rey Johnson himself in a speech given in 1989 (Appendix 1).

CHAPTER 4

The ADF (Advanced Disk File) and Future Generation Disk Drives

While the magnetic disk drive clearly demonstrated significant capacity improvements for data storage, it still did not offer either the access time or the capacity needed to implement the kind of rapid transaction processing needed for many computer applications. Companies such as UNIVAC were therefore aggressively developing high performance magnetic drums for mass data storage where access times superior to disks had been demonstrated. Given this environment, it is fair to say that a more advanced successor to the magnetic disk drive was viewed within IBM as absolutely critical to ensure magnetic disk drives were the future direction for data storage and that IBM had chosen the right path by focusing on disk drives and not drums or other technologies.

Even before the announcement of the RAMAC disk drive in 1956, IBM had made a decision in 1955 to move ahead with the design and development of the next generation disk drive, with a project named ADF (Advanced Disk File). The goals of the ADF project were to greatly surpass the capacity of the RAMAC while dramatically shortening access times. One of the first new commercial applications targeted for this next-generation IBM disk drive was the new American Airlines Airline Reservations System (Sabre), which depended on a geographically dispersed real time on-line transaction processing capability. IBM was also developing a new supercomputer, called Stretch, in order to gain the lead in scientific computing. This computer would also require a disk drive with similar advances in capacity and access time.

When the RAMAC project became a formal product development program in 1955, it was placed under Lou Stevens and transferred to a new manufacturing organization associated with a new IBM plant site in San Jose. The overall future of the disk drive became totally vested in the ADF program under Rey Johnson and his original research laboratory (at 99 Notre Dame) was moved to a larger facility in a dilapidated warehouse several blocks away. The ADF project under Rey Johnson began, somewhat like the original RAMAC project, by first choosing advanced performance goals to meet the new market opportunities, since there was no advanced research work to draw upon. Essentially, two performance goals set the technology objectives for the ADF. These were: a faster access to data by almost a factor of 10 and an increase in data capacity 10 times that of the RAMAC with the same size disk stack.

A major issue to address in improving access times for the ADF project was the question of head design. Heads are the mechanisms that read data from or write data to a magnetic disk or tape. The magnetic read/write head is imbedded in a small metallic carrier, called a slider, which is contoured like an airfoil to create and maintain the spacing of the head from the disk. The RAMAC had two heads, one to read/write from the top surface of a disk and one to read/write from the bottom surface. This meant that the heads had to be repositioned up and down the disk stack each time in order to access the correct disk and track and therefore access time was slow. To decrease access time, the solution was to have a head positioned to read/write each disk individually so no repositioning was necessary. To increase capacity it was necessary to decrease the head-to-disk spacing.

With the RAMAC, pressurized air was used to control this spacing and it took a room full of compressors to supply the air needed to float the disks (“floating heads”) to keep them from contact. This problem led the project to investigate using “flying heads” to control the head-to-disk spacing. The term “flying” refers to the fact that the slider carrying the head appears to be “flying” over the moving disk surface. A rotating disk carries a boundary layer of air on its surface when it is spinning. This airflow creates an upward force that is balanced by a slight downward force on the flying head. Thus a flying head does not require an external source of air. There was little understanding of flying head design and behavior at the time. Jake Hagopian was a quiet but prolific inventor in the laboratory and he was able to demonstrate the feasibility of the concept that flying heads significantly reduced the head to disk spacing. He quickly instilled faith in the concept among the engineers. Little did they imagine then how complicated it would be to get the idea to actually work reliably in a disk drive.

Another issue to address was the method of magnetically recording the data on the disk. There were two conventional options for aligning the magnetic pattern with respect to the direction of the recorded track, longitudinal and perpendicular. The longitudinal mode records magnetization patterns parallel to the track. In perpendicular recording, the orientation of the recorded magnetizing field is perpendicular to the track. Both approaches were pursued in the early years of magnetic recording but advances in longitudinal recording meant that this mode became the standard approach for all magnetic recording, including for the RAMAC, until very recently (2000s). Perpendicular recording was an alternative to longitudinal recording but the technology was not considered essential and was pursued only on the basis that it may lead to lower cost heads and disks.

Jake Hagopian introduced the idea of using a type of flying head called a *probe head* which recorded in an up and down fashion and required switching the recording mode to perpendicular. The probe head was viewed as a low cost option compared to the high cost of a ring-type head, which recorded longitudinally. In addition to developing perpendicular recording with a probe head, a new type of disk was needed. This type of head also required a two-layer disk with a new type of coating material. Jake undertook the task of looking for a suitable disk coating. In spite of everyone’s efforts, no suitable disk was produced.

The first fully assembled ADF prototype model was finally available for testing in August 1959. The performance data could hardly have been worse. The most obvious failures were frequent head/disk crashes related to the disk quality. The flying head technology was considered indispensable and there was confidence that it could achieve the target spacing of 250 micro-inches given further studies and testing. Continued advances in magnetic disk storage density would require that the spacing be progressively reduced through technical advances. With anticipated improvements in disk and head surface qualities and more sophisticated head contours, there was no reason not to expect continued reductions in spacing for the foreseeable future.

The ADF program was based on several new technologies, all of which were still not fully understood and none of which had been adequately tested independently. Time lost trying to develop these new technologies caused the ADF project to fall behind schedule

and the entire program faced serious problems, affecting both the IBM 1301 disk drive and the Stretch supercomputer system being developed by IBM.

The IBM 1301 disk drive was the commercial product version of the ADF, with the marketing and system requirements being set by the business community. The Stretch computer system was being designed to meet the needs of the Atomic Energy facility at Los Alamos and serve as an assertion of IBM's dominance in the field of supercomputers. The Stretch disk drive needed to be capable of high-speed data transfers, speeds not required for ordinary business uses. Corporate management gave its greatest attention to the progress of the Stretch disk drive for the Stretch computer system due to the high public visibility of that program. The Stretch version of the ADF was of equal or greater importance in terms of the competitive position of IBM, especially for the commitment IBM had made with American Airlines for the "Sabre" system, one of the most advanced transaction processing systems then under development. The problems with the both the Stretch and commercial version (IBM 1301) were serious enough that IBM began to consider if an outside supplier might be able to produce for them a similar disk drive meeting the 1301 performance objectives. In January 1960, IBM set up a high level task force to conduct a major audit of the ADF project. I was one of three individuals chosen to participate.

Recommendations from the audit suggested several steps be taken. A major research and development program on flying heads would be immediately set up, using the most advanced painted iron oxide disks. Perpendicular recording would be abandoned and advanced development work would begin on modifying RAMAC heads and disks to meet the ADF specifications using conventional longitudinal recording. Finally, the Stretch program would be separated from the ADF project and use already proven RAMAC technology. From the point of view of IBM, its reputation was tied to Stretch and it had to meet the commitment to the Government no matter what the cost.

These decisions not only improved the chances of meeting the earliest date for shipment of the Stretch unit for Los Alamos by 1961, but also bought additional time for the core 1301 group to pursue the basic technology studies required for the commercial version which, while behind schedule, could still be completed well before any competition would exist in the marketplace.

While the audit committee recommendations were accepted and implemented, they were accompanied by major management changes. Vic Witt, the laboratory Director, took over as Director of Storage Products, and Al Shugart became the Engineering Manager for the commercial version (1301). Al Shugart joined IBM in 1951 as a field engineer and subsequently gained experience in several RAMAC development, assembly, and product engineering assignments. Jack Harker rejoined the project and took on the management of the flying head work. I was loaned by Research to handle the change from perpendicular recording back to longitudinal recording. The main advantage of these changes was to move the technology base back to the recording expertise already developed for the RAMAC. This was tremendously important as it allowed a fairly rapid turnaround of the commercial program. The 1301 disk drive was announced in 1961.



1301 Disk Drive (Disk stack of two 2 disk modules independently accessed, 1961)

The ADF effort demonstrated that if you have an already existing technology, in this case non-contact digital longitudinal recording, for which you do not yet know the limits, do not abandon it for another approach you have never pursued and know little about.

The successful turnaround of the ADF program served to convince IBM that their new San Jose facility should formally be given the mission for all disk-drive development within IBM, leading eventually to the San Jose area becoming the center of this industry. The importance of further basic research to advance magnetic disk storage technology was considered essential. In addition, this experience brought about a more conservative approach to the introduction of technological advances in disk drive development for a number of years. It became accepted that only one major technological change at a time should be made for each succeeding generation of disk drive. This approach was reinforced by the need to design disk drives in conjunction with data processing systems. These systems guided disk drive product features, schedules, and manufacturing volumes. One other lesson that may be drawn is that a specialized version of a mainstream product should follow and not precede the final commercial development phase.

Remarkably little remains to document this period. I have not been able to uncover any artifacts at either IBM or the Computer History Museum (CHM) from this era. While IBM was the leader in the early years of disk drives, this period was one of the more turbulent periods in magnetic disk drive history.

CHAPTER 5

Personal perspectives on magnetic data storage during the 1960s and 70s

With the announcement of the IBM 1301 in 1961, I decided to write a book on non-contact digital magnetic recording. A rapidly growing number of engineers were becoming involved in magnetic disk storage within IBM, their computer industry competitors and in newer companies now entering this market place. I saw that a need existed for this information. I was regularly called upon to explain the technology since no published references existed in this new and rapidly expanding field.

Before I began writing I was asked to undertake a two year assignment in the Netherlands, starting in 1962, as a member of an IBM-Philips team set up to support the Dutch government in addressing the conversion of their national banking systems from tape to disk data storage in order to improve the speed of account transfers. Prior to undertaking the assignment, I checked with both IBM San Jose and IBM Netherlands and found that they would fully support me in my work on my book during my time there. I was provided an office and secretarial support in Rotterdam. My book, entitled "Digital Magnetic Recording", was published by Wiley in 1963. This was the first book covering this subject, published at the time disk drives were becoming an increasingly important factor in the electronic computer field.

Following my assignment in the Netherlands, I expected to return to the IBM San Jose research laboratory in 1964, but after completing the book and just before I was to leave, I was offered the position of Manager of Engineering Science at the newly established IBM Research Center in Yorktown, New York. The formation of the Yorktown Research Center was a huge commitment by IBM and the goal for Yorktown was to become the equal of the famed Bell Laboratories.

The opportunity to participate at the beginning of the new IBM Research Center, where major figures in the computer world were being hired, was very attractive. I left the Netherlands in 1964 and moved to Stamford, Connecticut, about a half hour drive from the new Research Center. I also interpreted the recent publication and positive reception of my book as a good omen for the future of magnetic disk storage technology.

At Yorktown, my interactions with professional engineering societies and computer groups increased greatly. My role as Manager of Engineering Science included exploring the potential utility of scientific advances that could have a direct impact on mass data storage. In the expanding engineering science group there was a great desire to explore many data storage alternatives, including holography, electron beam recording and magneto-optics. Work on magnetic cores and magnetic films for memory was also a key area of activity at Yorktown.

Among new PhD hires, all saw the challenges and rewards associated with creating a new electronic (or non-mechanical) technology for mass data storage, and viewed magnetic disk recording technology as already mature and subject to all the limitations normally associated with mechanical devices. IBM divisional responsibility for the development of

magnetic data storage had rested with the San Jose research laboratory, but unfortunately the management of that laboratory also believed magnetic recording would not be the storage technology of the future and that their responsibility was to find its replacement.

Perhaps the biggest obstacle I faced at Yorktown was found in a memo written by Manny Piore, the IBM Chief Scientist and Director of Research in the late 1960s. After careful deliberations with his inner circle, he concluded that magnetic recording technology was now mature, and that Research Division efforts in this field were no longer appropriate. Unfortunately, even many disk drive engineers in San Jose at this time also shared this perspective, hoping only that the technology would survive until they could retire. The majority had mechanical and electro-mechanical skills not easily translated to advanced electronics opportunities. The belief that disk drive technology would soon be obsolete was created by setting ever higher performance goals for new disk drives in response to market pressures with little, if any, major research or advanced technology efforts to draw upon in support of the projected advances. IBM was still the clear leader in magnetic disk devices and competition from other companies was primarily focused on making compatible drives and profiting from the IBM pricing umbrella. In the 1960s and early 1970s, while significant advances were made in disk drive technology, progress was generally viewed as pushing existing designs to their limits.

A new data storage technology, when compared to a magnetic disk, had to offer significantly faster access rates at an equivalent or lower cost to justify its introduction. In the hierarchy of memory storage technology, the magnetic disk occupied the intermediate level, having slower access times than electronic core memory yet offering faster times than magnetic tape. The cost difference between electronic memory (very expensive) and mechanical memory (disk and tape) was so great however that economics dictated that a single computer system would need to use all three-memory technologies. A memory/storage hierarchy had at the top an electronic high-speed memory (working at sub-microsecond speeds) to keep the processor fully utilized. Below that in the hierarchy were magnetic disks, providing a non-volatile mass storage with access times in milliseconds but with very high data transfer rates. The bottom level of this hierarchy was magnetic tape with access times orders of magnitude slower than disk, the most economical method to store backup and archived data.

The magnetic disk therefore created a vital level in this hierarchy and changed the way computers were designed and the applications they could undertake. Clearly, such evaluations and judgments depended not only on the current status of disk technology but also on an understanding of the future potential of magnetic disk storage as well as any other alternative storage technologies that could be identified.

By 1967, I had been promoted to Director of Research Planning for the Research Division. This created a major conflict between my personal interests in magnetic disk and tape technology and my new role in which the Laboratory had committed its research efforts in data storage to electronic methods. My agreement to relocate to Yorktown had been with Art Anderson, then the Assistant Director at Yorktown, who had offered me the original position of Manager of Engineering Science. Being a native of California, one condition I placed on our arrangement was that I would be transferred back to California and the San Jose research laboratory (I regarded Yorktown as a short-term stop

on my way back home to California). It now seemed an appropriate time to review my situation. An agreement was reached with Ralph Gomory, then Director of Research, that I would spend nine months in Boulder, Colorado before returning to San Jose. Boulder was where the IBM magnetic tape product development and manufacturing unit was located. I would initiate and organize a joint Yorktown/Boulder "Benchmark" program to demonstrate that much higher densities were still possible on magnetic tape, a product category that was viewed as having a very long lifeline, at least for archival storage. At this time, the computer industry was still completely dependent on low-cost magnetic tape systems for off-line storage, a very competitive product area. There was no real alternative technology for this purpose then foreseen. I accepted the Boulder assignment with the understanding that Boulder would again be just a stopover on my way back to California.

I arrived in Boulder in September of 1971. All IBM magnetic tape activities had been consolidated at a relatively new IBM plant site in the area. I met with Max Femmer, in charge of product development activities, and we agreed to start the benchmark program by establishing a local task force to assess the limits of known engineering design approaches. On the basis of our analysis and experimental studies the group became convinced that tape storage density could be increased by a factor of 10 or more. An engineering project would be undertaken to attempt to demonstrate the gains predicted.

However, in mid-decade IBM decided to relocate the entire IBM tape organization to Tucson, Arizona. The move came at the cost of major attrition of engineering personnel for whom leaving Boulder was unacceptable. IBM had decided to move the Office Products Division (OPD) from Lexington, Kentucky to Boulder, as that Division needed room to expand. In order to make it easier to convince OPD personnel to relocate, IBM agreed to give this Division site responsibility for the Boulder facility. These changes also triggered a major loss of personnel by the IBM tape division as many would not relocate to Arizona. As for myself, I decided to exercise my option and return to the San Jose research Laboratory, not nine months as originally planned, but after five years.

I relocated to San Jose in June of 1976 and rejoined the San Jose research Laboratory. I was temporarily given the responsibility of setting up a small exploratory magnetic recording group with a starting headcount that included only one professional beside myself. Soon after I arrived my manager requested a meeting. He informed me that I needed to consider changing into another career field since the San Jose laboratory had recently made the decision not to continue working on magnetic recording but focus instead on magnetic bubble technology.

A new department of about 20 or more research scientists had been hired to work on this new technology. Magnetic bubble storage is based on creating small magnetic regions in a continuous film and electronically controlling their movement to a location where they could be accessed. It was expected that this technology would also be non-volatile (which semiconductor memory was not), yet faster than magnetic disk devices and cost competitive. The decision to discontinue work in magnetic recording was consistent with the views at Yorktown and hardly surprised me even though the magnetic bubble project still had not reached the product development stage and a new advanced disk drive product (IBM 3380) was soon to be announced.

Independent of the IBM San Jose facility, interest in magnetic bubbles as an entirely new direction of research for magnetic recording storage was emerging in Japan and being pushed aggressively. Professor Iwasaki of Tohoku University was also drawing attention with the claim that for the highest magnetic recording densities the recording mode needed to be perpendicular rather than the universally used longitudinal mode. In Japan Iwasaki was regarded somewhat like a Nobel Prize winner is regarded in the US, so his statements were taken very seriously there. His papers and talks on this mode of recording raised great interest not only in Japan but also to a lesser degree in the United States.

Iwasaki traveled abroad extensively to share his results with others. Familiar with my work in the field of digital magnetic recording, he stopped by to see me on a visit to the Bay Area soon after my return to San Jose. During our conversation I casually mentioned that the early design of the IBM 1301 disk drive was based on perpendicular recording, a mode that was dropped in favor of the conventional longitudinal recording of the time. Apparently, he was not familiar with this history, viewing perpendicular recording as a uniquely Japanese contribution to the art. I assured him that the perpendicular mode was selected initially for the 1301 design on the basis of offering head and disk cost advantages and was later dropped due to the lack of suitable disks. However, our 1301 experience shed no light on the relative high-density advantages of perpendicular compared to longitudinal recording, as we had not successfully tested the concept due to complications.

Iwasaki was able to provide specific test data that convinced me that perpendicular recording was a research direction IBM needed to seriously investigate. I believed that my exploratory magnetic recording group was the logical group to take on this task. However, perpendicular recording required radical changes in both recording media and magnetic heads before it could even begin to be tested, and so represented only a long-term possibility.

Somewhat surprisingly, perpendicular recording quickly generated widespread interest in the recording community and generated continuous theoretical debates on the fundamental theories and merits of the two recording methods at IEEE Magnetics Society meetings and Conferences for years, in spite of the fact there were very few who believed perpendicular recording would replace longitudinal any time soon. From my point of view, the visit of Iwasaki could not have come at a more opportune time.

CHAPTER 6

Importance of Magnetic Disk Technology Recognized

The late 1970s and early 1980s were a period of great turmoil for magnetic disk data storage. Three major events were to have a major impact on its future. The first event was the failure of IBM to ship the IBM 3380 disk drive on schedule, the first time a shipment date had been missed (delayed from June 1980 to October 1981). The 3380s were the most advanced disk drives the company had yet developed and had the largest market demand ever for such a product. The second event was the growing realization that magnetic bubbles did not appear able to replace magnetic disk storage in the memory/storage hierarchy. The third event of importance was the emergence of a “personal” computer market with the Apple II in 1977 and the IBM PC in August of 1981.

The delay in production of the IBM 3380 was due to the problem of intermittent contacts between heads and disks. A head hitting the disk leads to both head and disk damage and loss of data. This made it very clear within IBM that research on magnetic disk technology had not been adequate to support the product development activities necessary for successive generations of advanced disk drives. The growing complexity and technical sophistication of magnetic disk drives had not been sufficiently appreciated or understood for several years. The 3380 program was in desperate need of more human resources with skills and expertise in magnetics, materials, instrumentation and testing to help “troubleshoot” the causes of failure. San Jose had only a few people in research and product development exploring technical advances, and these were divided between the separate Research and Development organizations. However, an untapped pool of talent with magnetics skills already existed in the magnetic bubble program within the Research Division.

Magnetic bubble technology was proving to be much more expensive than disks and was also not as fast or cost competitive with semiconductor memory. With the necessity to quickly address and solve the problems of the 3380, the Research magnetic bubble group and advanced magnetic recording resources in both the Research and the General Products Division (which was responsible for the 3380 product) were assigned to support the 3380 program. A re-organization combined the magnetic bubble group with the existing advanced recording groups into a new Magnetic Recording Institute. Magnetic disk storage had gone from being viewed as a technology near the end of its life to the main hope for the future of storage.

The third but somewhat independent factor at the time was the emergence of what is now called the PC or personal computer. Hobbyists and small companies were making hardware that suddenly provided electronic computer capabilities for individual use. This development was stimulating great interest and many new projects were undertaken to address this new need. Existing technology was more than adequate to design and make low-cost drives that only required small data capacities and small companies quickly entered this marketplace.

Small hard disk drives (HDD) based on magnetic recording storage first appeared in 1980 when Seagate Associates designed and announced a small 5 MB hard disk in a 5 1/4 inch form factor (the term form factor refers to the physical dimensions of the disk) that could be integrated into the body of a personal computer. This capacity was more than adequate for the computing capabilities of small personal computers. Although IBM introduced the PC, it did not enter the small form factor disk drive market at this time. Disk size for non-PC use remained at 14 inches, which was the standard size for mainframe and minicomputer drives for many years. The story of the impact of disk drive form factors on the emerging disk drive marketplace is well covered in the book “The Innovators Dilemma” by Clayton M. Christensen.

With the market growth in personal computers, business opportunities rapidly expanded in the small form factor disk drive business. The initial absence of IBM in this market opened the door to many small companies. They located around IBM in San Jose and frequently were staffed by ex-IBM employees. The advent of the PC led to the creation of a true disk drive industry. IBM became just one of many companies, although clearly still the technical leader. Innovation and new design approaches came into being. Mechanical tolerances favored smaller form factor drives as the path to higher and higher densities, bringing higher and higher capacities on smaller and smaller disks. The consensus was that several small drives would be a better economic solution than a single large drive. The need for more engineering talent and expertise in magnetic disk storage quickly became an industry-wide problem. Industry would need to greatly increase its funding for research and development. Another way to expand the technical resource base to address the challenges was to more fully involve the academic world, which was already heavily involved in the semiconductor field.

In 1979, I was invited to Harvey Mudd College in Claremont, California to give a banquet speech at the awards ceremony for the conclusion of the annual senior projects day. The topic of this speech was at my discretion. The afternoon of the event I still was not sure what I would say of particular interest to them, as I knew that very few, if any, engineering faculty were familiar with magnetic recording technology. Then it struck me that I could describe how the relationship between the magnetic disk industry and academia was currently the exact opposite of the productive one between the semiconductor industry and academia and how magnetic technology could benefit from a similar alliance.

As a graduate and then a faculty member of UC Berkeley before joining IBM, I was already aware of the growing interaction of the EE Department at Berkeley with the emerging new electronics industry based on semiconductor advances. Thus the EE departments at every academic institution represented a source of collaboration and graduates for the electronics industry. A close bond between industry and faculty was created and maintained by their graduates. So why had we not seen similar relationships develop with the emerging magnetic data storage industry? My belief was that almost all universities are organized into departments around fields of study and disciplines. Magnetic recording storage technology did not fit within the normal University academic structure since it is so interdisciplinary. A disk drive requires expertise in magnetic materials, mechanics, electronics, signal processing and so forth. Design tradeoffs call for a very close interaction between experts in these fields. This difference became the theme

I presented. The point of my speech was that the recent challenges in magnetic recording storage suggested the absolute necessity for the magnetic disk data storage industry to develop a closer partnership with the academic world, where many of the needed areas of expertise existed.

The reaction to my talk was surprisingly favorable and a number of people separately asked me why nothing had been done to change the situation. I found this question intriguing but a clear plan of action was not evident. However, with the 3380 crisis and the internal concerns it elicited, I scheduled a meeting with Art Anderson, President of the IBM General Products Division (which included the San Jose disk drive plant site). Individually we both had been thinking along the same lines and shared our perspectives on the benefits a partnership with academia would bring to IBM as well as to the disk drive industry. He proposed that the industry fund at least one university center for magnetic recording data storage research staffed by faculty members in relevant disciplines. This called for getting both an industrial group to sign on to finance such an endeavor and at least one university to provide a home for such a center.

Art Anderson asked me if I would take on the task of organizing an industry consortium to underwrite a university center dedicated to magnetic recording data storage and locate a university willing to set up such a center. IBM would make a major financial contribution and I was to focus on only a few large corporations to expedite the process. The center would be located on the West Coast to ensure the disk drive industry and their research staffs, primarily clustered in the San Jose area, would have ready access to the Center and vice versa. The center we had in mind was to be for the benefit of all US companies in this field.

I accepted this challenge immediately as I saw this project as essential for the long-term competitiveness of US disk drive companies. I had been involved throughout my career in many presentations and discussions with executives and technical leaders on the importance and exciting future of magnetic disk storage, and relished pushing this cause. I did not know at the time how naïve I was in anticipating the reactions and issues I would encounter both from companies and universities.

The first step was to select and arrange meetings with several major companies with a significant presence in magnetic recording data storage technology. Getting other corporations to join the consortium was expected to be much simpler with IBM already committed. Among my professional contacts was Jim Lemke, the President of a small company in San Diego called Spin Physics that was making magnetic tape heads. Our mutual personal interests included a passion for flying. I discussed my new assignment with him. He liked the idea and agreed to work with me. Initial meetings were set up in Rochester, MN, with Control Data Corporation (CDC), a competitor of IBM, and 3M, a major supplier of magnetic tape among many other products.

In our talks with Control Data Corporation it was clear that they feared Japanese companies would “eat our lunch” in magnetic data storage if nothing were done, and felt this Center should be for US companies only. This proviso had not been discussed previously but was accepted to make it easier to get a Center started. Back in San Jose we approached Memorex, a firm that had a long history in magnetic tape. It did not take long

to enlist the support of these three companies. We did not request specific amounts as contributions but were nevertheless looking to secure several million dollars.

The next challenge was to initiate discussions and seek agreement with a major academic research institution on the West Coast. This step turned out to be a much more arduous task since the academic world had little if any knowledge of magnetic recording technology or the disk drive industry. In the Bay area there were two prime candidates, Stanford and UC Berkeley. Jim Lemke believed UC San Diego would also be receptive and approachable if we were not able to place the new Center in the Bay Area.

Our first formal visit was to Stanford University to talk with Bob White, then Dean of the School of Engineering. The conversation revealed that he had little interest in our proposal. Stanford had just created a new center for microelectronics systems with broad industrial support and that was as much as they wanted to be involved with at the time. However, the key message in his comments was that if they were to set up the Center what would it be doing in five years? In other words, he viewed our efforts as only a response to the short-term crises associated with the 3380, and did not see a long future for magnetic recording technology, just the opposite of our view.

I was more optimistic that UC Berkeley would be interested since it was one of the first universities to get involved in digital magnetic recording back when I was a graduate student there. But creating a new center was considered too complicated and time-consuming a task to consider. Again, as we left, one of the professors I knew personally privately asked me what I thought the Center would do in five or so years, again reflecting the general perception in the academic community that magnetic recording data storage was already mature. In spite of my original enthusiasm, we failed to find a home for the center in the Bay Area.

Now that Stanford and Berkeley had turned the proposal down, UC San Diego was the most obvious next choice. Jim Lemke had close relations with this University and set up a meeting with Lee Rudee, their Dean of Engineering. The school was under growing pressure to expand their school of engineering to meet the needs of the state in coming years due to state population growth. Thus, the University would be receiving additional faculty slots and building funds from additional state funds. Our proposed Center fit within the general long-range growth plan the school had been assigned. We entered into formal discussions on what arrangements we felt were essential and the tradeoffs needed to make a deal possible. After several weeks we came to a general understanding. It was agreed that four new faculty positions would be allocated for the new Center, to be filled by senior professionals with relevant experience in the respective areas of mechanics, signal processing, recording physics and magnetic materials, all expected to be new hires and likely from industry. When final agreement was reached we had created the new Center for Magnetic Recording Research (CMRR). During this period Carnegie Mellon University (CMU) also began discussions with IBM about a research center. CMU had a strong group engaged in magnetic bubble research that could be redirected to research on magnetic recording. Being on the East Coast, the CMU proposal did not relate to our mission but there was mutual support between the two efforts, two centers being far better than none. CMRR was formally established in 1983 and the CMU organization in that same time period.

To nail down our commitment without waiting for all the formalities to take place, I was asked to serve as the acting director (while remaining employed with IBM) and temporary quarters were leased in San Diego. I served in this capacity for about a year and a half. With the formal dedication of the Center looming, I was invited to stay on as the Director by UCSD. I declined this offer seeing it as too remote from my wide interests in magnetic disk storage that ranged over research, development and product-oriented activities, all being centered at the heart of the industry in the Bay Area. In the spring of 1984 I returned to IBM San Jose.

IBM had recently begun an early retirement program for their employees with attractive financial incentives. I now had been away from IBM for a year and a half and no longer had any direct involvement with the San Jose research laboratory. I was eligible for the early retirement option, although I had no thought of actually retiring from my involvement in magnetic disk storage technology. I strongly believed the disk drive activity in the San Jose area would thrive and be accompanied by an increasing need to train employees hired by new and expanding disk drive companies as well as keep current disk drive engineers updated on the state of the art. A new opportunity to remain active in magnetic disk storage presented itself through a university uniquely qualified and ideally located to support these kinds of activities, Santa Clara University. I resigned from IBM in July of 1984 to become a professor of EE and the Director of a new institute at Santa Clara University, the Institute for Information Storage Technology (IIST).

Santa Clara University was the first college chartered in California. While the school of Engineering is small with no major research facilities, education is one role that Santa Clara University (SCU) is uniquely qualified and able to serve well. They had in place an "Early Bird" program that allowed working engineers to pursue a Masters Degree by scheduling graduate engineering courses very early in the morning, permitting students to also continue working on a full time basis. SCU is located in the heart of the disk drive industry and about a mile or so from the San Jose Airport, an ideal location for classes, seminars, conferences, etc.

I put together a one-page charter and scheduled a meeting with Paul Locatelli, then the vice president of the school. He agreed the school would publicize the new institute with a press release. However, the University had completed their budget for the academic year and could not provide any financial support. I did not want to jeopardize the start of the Institute, so took the stance that if the University would not support the Institute then it was not worth doing, and we came to an immediate agreement.

The first challenge was to obtain industrial support. In contrast to CMRR, no company had committed to provide support. Furthermore, having recently secured industrial funding for CMRR amounting to about 6 million dollars, it was not reasonable to go back to that same set of major contributors for more money for similar purposes. I concluded that I would first try to obtain support from the local disk drive startups. These small companies would see an Institute at SCU as of real value as a source of new engineering talent, technical updating for current employees and as host for topical meetings and conferences.

I have been both the President of the IEEE Computer Society and of the American Federation of Information Processing Societies (AFIPS) and one of the most difficult challenges has always been to get the first organization signed on. The first question you are asked is “who has already signed up?” If you say no one, it automatically generates the reply “let us think about it and we will get back to you”.

Years earlier I had worked at IBM with Al Shugart, the current CEO of Seagate, during the IBM 1301 crisis (Chapter 3). I knew he did not hesitate in making decisions or in giving his opinion and decided that he was my best hope. We met and in half an hour and without a formal presentation he simply said that my proposal sounded like a good idea and agreed to sign up Seagate as a member for five years. This decision then made it easier to get other companies to join.

The Institute for Storage Technology, IIST, began operations in 1984 and its first major conference at SCU covered the fundamental technologies underlying disk drives. The interest in learning more of what was now viewed as a new field with exciting opportunities exceeded all expectations. In the second year almost one third of the students in the graduate engineering program at SCU were enrolled in classes relating to the disk drive industry. The teaching faculty for our graduate programs drew heavily on existing experts in the industry. My major responsibility was to market educational services to make IIST financially self-sufficient. In a few years we were no longer dependent on direct funding from industry but rather relied on the registration fees for employees underwritten by their companies.

CHAPTER 7

Looking Forward, Remembering the Past

The RAMAC was a pioneering mass storage device based on non-contact digital magnetic recording of data on a stack of rotating disks, with access through a head-positioning device. Since its introduction in 1956, there have been several major key technological breakthroughs that powered significant gains in magnetic disk drive storage technology. Among these were the successful introduction of flying heads, the development of MR and GMR magnetic reading heads and the transition to perpendicular recording from longitudinal recording.

The introduction of the flying head for the ADF (Chapter 4) provided a means to achieve an increasingly smaller separation between the recording head and recording medium, dramatically reducing access times. This technology was first introduced on the IBM 1301 in 1961 at a spacing of 250 micro-inches or 25,000 nanometers (nm), about the thickness of a human hair. Major improvements to this technology in the ensuing years meant that by 2006 the separation was about 10 nanometers. Advances in air bearing theory initially suggested that a limit of 3 micro-inches spacing could be achieved based on the mean free path of air molecules. However, the technology has been taken well beyond that limit. Major advances have come from more perfect surfaces on both disks and heads, improved flying contour designs, temperature sensitive head loading as well as a more fundamental understanding of the physics of thin film air flows. The reduction in head-to-disk spacing achieved has been a continuing and essential factor driving increases in storage density.

The development of MR (Magneto-Resistance) and GMR (Giant Magneto-Resistance) magnetic reading heads was also significant. Since its invention in the last century, the standard method for reading magnetically recorded signals had been a ring head, a magnetic core with a narrow head gap to sense magnetic flux from the surface produced by a recorded pattern along the recorded surface. The MR head essentially sensed the magnetic field above the surface by a narrow strip and produced significantly larger and higher resolution readback signals as the bits became smaller. In 1989, an even more significant physics breakthrough was announced. A far greater magnetic field sensing phenomena had been discovered that became known as GMR. Initially GMR was not applicable to magnetic recording but the field sensing effect was far superior to the MR head and led to immediate research to find other magnetic materials appropriate for magnetic reading head applications. It is truly amazing that in a short period of ten years this physics discovery of a totally unpredicted effect went from being a curiosity to the commercialization of GMR magnetic heads for magnetic disk storage that greatly increased the storage densities possible. The two physicists who discovered the GMR effect, Albert Fert and Peter Grunberg, won a Nobel Prize for their work.

The long debated technology issue (covering 20 or so years) of longitudinal versus perpendicular recording was eventually resolved. The general view for many years was that the perpendicular mode had no real advantage. In spite of all the advances in longitudinal recording in the 80s and 90s, perpendicular recording, strongly espoused by

Iwasaki, was not forgotten and continued to be an active and emotional item of discussion. The Magnetic Recording Conference (TMRC) did not include perpendicular recording in its program until 2004 on the basis that the purpose of the conference was to address the mainstream technologies being pursued in industry. The earlier interest in perpendicular recording for the ADF (Chapter 4) was related to hoped-for lower costs for heads and disks. This was never achieved due to poor disk quality. There was no other advantage seen at the time over longitudinal recording, so the technology was not pursued any further and longitudinal remained the primary focus. However, as bit density continued to dramatically increase, there was research showing that the instability of the magnetic bits (self-demagnetization) would set the ultimate limits you could achieve in storage density. The advantage of perpendicular recording over longitudinal was that it allowed higher bit densities before self-demagnetization could occur. The conclusion was that to make continuing advances in storage density possible it would be necessary to switch to perpendicular recording. By the early 2000s, all new disk drive products used perpendicular recording.

By 2001, changes were taking place in the disk drive industry that would have a significant impact on both IIST and my career. The industry was undergoing a major consolidation as the number of companies then making disk drives made survival an economic challenge. Being first to market with an advanced drive was key to gaining market share in order to support the ever-increasing investment in research and advanced development needed to remain competitive. These challenges led to only a few companies surviving. As IIST funding was based on broadly supporting a wide range of both large and small competitive companies, the decreasing number of companies presented a challenge to the organization. There was also the uncertainty of being able to attract new students who saw the future as providing fewer job opportunities in disk drive engineering.

Of special interest to me was the approaching 50th anniversary (2002) of the small San Jose laboratory at 99 Notre Dame that Rey Johnson had headed and that had been the birthplace of the first disk drive, the RAMAC. The 50th anniversary of RAMAC would soon follow in 2006. I made the decision to gradually phase out of my IIST role and set up a small organization within IIST called the Magnetic Disk Heritage Center (MDHC) to put in motion steps to get the historical significance of the magnetic disk drive properly recognized as well as to continue the work on restoration of a RAMAC to functional status. The RAMAC restoration project had initially begun as a senior EE design project at SCU under IIST in 1993 and 1994. In this period IBM was in the process of getting out of the disk drive business, providing MDHC a unique opportunity to lead in these history-oriented activities and to be one of the sponsors for the 50th Anniversary RAMAC celebration that was held at the Computer History Museum (CHM) in September 2006. I had retired from SCU in 2005 and relocated MDHC and the RAMAC restoration project to the Computer History Museum. By that time San Jose had made the building at 99 Notre Dame a City Landmark (2002) and the IEEE had made the RAMAC an IEEE Milestone (2005).

Restoration of an IBM RAMAC disk drive

Discussions between IBM and the Institute of Information Storage Technology (IIST) at SCU starting in 2001 led to an agreement that there would be real historical value in restoring a first disk drive to an operational condition for public display. The RAMAC was the only drive where you could easily see the random accessing motions and the access speeds, while quite fast, could be visually followed. There were only four RAMAC drives known to exist. IBM had leased the drives and they were destroyed when no longer in use. In the IBM archives, there was only one RAMAC disk drive, which, while it included all the mechanics, was missing the associated separate unit containing the controller. The proposal was to initiate the restoration project as a sequence of senior engineering design projects, an activity required for graduation. Such a restoration project would generate great student interest in disk drives and their origins as well. Since the senior project requires original design studies using mechanics, the project also provided the students with the challenge of designing an electronics controller for the unit.

With the support of IBM San Jose, SCU submitted a request for the loan of this RAMAC disk drive to the IBM chief archivist, Paul Lasewitz. The good news was that I was only interested in the mechanical hardware and they had one unit in storage without the controller. The request was approved and IBM provided significant assistance in getting the unit located at SCU. IBM, and particularly San Jose “old timers,” provided essential help in refurbishing the hardware. The creative challenge for the students was to design a modern electronic controller for the disk file. A key individual from the very beginning, Dave Bennet of IBM, assisted in this early phase and has been a major contributor to the restoration project since then.

The restoration project was pursued at SCU from June 2003 to June 2005. This was a fun time, not only because I had worked on the RAMAC in the 1950s, but also because the students were excited about the chance to restore such an historical artifact. The engineering school enjoyed the attention it brought. The feasibility of access positioning was demonstrated at reduced speed the first year, and reading and writing data the second year, these two stages establishing confidence that it would be possible to restore the drive to an operational state.

Upon completion of the second senior design project, the RAMAC disk drive was relocated to CHM to continue the restoration project there, where a computer restoration area had already been set up. The restoration at this stage was continued by volunteer professional engineers experienced in disk drive design, to meet the much more demanding requirement that the restored RAMAC be ready to be placed on public exhibit. The leader of this phase of the project was John Best, formerly a vice president of the Hitachi Disk Drive Division and prior to that Director of the IBM Almaden Research Center.

I am very pleased to say that the restoration has been completed in time so that this working RAMAC disk drive is now on permanent exhibit in the Memory & Storage gallery at the new CHM “Revolution” Exhibit that opened in January 2011.

Conclusion

What can we expect for magnetic disk storage in the future? What we do know is that the increasing amount of information to be stored and available on-line will continue to grow at a greater and greater rate. Furthermore, there will be continuing growth in the backup of archived digital data, a function now primarily provided by tape. I believe magnetic disk storage eventually will take over this archival role. With data stored on magnetic tape cartridges, retrieval requires locating and transporting them to be mounted and searched sequentially by tape drives. A concern with magnetic tape is whether the same device used to record the tapes will still be available in future years and whether the tapes will be still readable due to environmental factors after long storage. A disk drive operates only when data is written or retrieved and the read/write electronics is integrated within the drive. Disk drives can be routinely checked and data automatically transferred to another drive on a scheduled basis. Finally, the much higher storage densities achieved with magnetic disks as compared to magnetic tape should eventually make this approach cost-effective.

Billions of disk drives already exist and the number of new drives produced each year is measured in the hundreds of millions. I see no alternative technology that is as economical for storing the huge quantities of data supporting Internet, enterprise and user applications. Today a terabyte of disk storage costs \$60 and can fit in a purse. The cost of developing and manufacturing an alternative would be prohibitive. The economic value to an individual or company of a disk drive actually resides in the content stored and accessible on the unit and not in the initial cost of the device.

Today computers are periodically upgraded and get a new drive with much more capacity, transferring their existing data to this new drive. This situation has many advantages, one of which is that you need not worry a great deal about the product lifetime of the older disk drive. If further advances in storage density are not achieved, product life issues become more important and will further encourage using backup services on remote servers. These factors will accelerate disk drive growth, not through individual user purchases, but through the companies that manage massive amounts of disk storage.

I expect magnetic disk storage to not only survive but also to continually improve for a very long time and well after the current rate of density advances slows down. It is a matter of the economics associated with a huge existing market already providing attractive and very low-cost commodity products. Further additional approaches are already being explored to extend magnetic disk storage density even further, including heat-assisted magnetic recording (using a focused laser to define a smaller bit size) and patterned media, where each bit location is in essence lithographically printed on the disk.

The increasing importance of mobile devices favors a move to solid-state storage because small size, weight and power favor this technology. However, for storage and access to the huge amounts of data required by business, the internet (“the cloud”) and global communications, where size and mobility are not the limiting factor, magnetic disk storage remains the lowest-cost and only practical means to provide this service. Mobile

devices themselves do not need to store large amounts of data because of easy remote access to the information. In my opinion, the continuing growth in sharing information wirelessly and its growing importance to society can only be met by magnetic disk storage for the foreseeable future.

The most unique and singular aspect of a disk drive is that a mechanical feature, the head-to-disk spacing, has been the single most important contributor to the long lifetime of magnetic disk storage and has assured the survival of a mechanical form of data storage in the age of electronics. The disk drive has reached its present state not in spite of being mechanical but rather because it is mechanical. In addition, it is the simplicity of a mechanically based design that makes the magnetic disk drive so successful in meeting the random access data storage needs in electronic computer applications. And all this in 60 years of existence!

APPENDIX 1

Dinner Talk by Rey Johnson at the Data Storage '89 Conference
September 19, 1989, San Jose, California

“Good evening. Thank you for the introduction. I’m delighted to be able to spend the evening with you.

I’ll be talking this evening about some of the highlights in the development of the first random access disk product – the IBM RAMAC 350 file – The why, where, when and how of the first disks.

This is the product that spawned the magnetic disk storage industry – an industry that has since come to generate annual revenue of 23 billion dollars. I think it fair to say that the RAMAC 350 has carved a place in history for itself.

I am sure that many in my audience were not born in 1951 when my story begins. Let me therefore set the stage for the events, most of which took place in San Jose, a few blocks from here.

In 1951, San Jose was a city of 100,000 people. Its economy was based mainly on agriculture. At the same time, IBM was a rapidly growing company, with data processing as its main business. Its revenue in 1951 was about \$250,000,000.

After IBM had developed two large computers that were intended strictly for scientific applications, IBM management saw little evidence that computers would become profitable business products. In fact, in 1951, an extensive survey of all potential computer customers yielded statistics that indicated that 17 or 18 computers would saturate the market.

At that time, Thomas J. Watson, Jr., was succeeding his father, Thomas J. Watson, Sr., as president of IBM. He decided to build 19 scientific computers. The first of them was completed in 1952.

IBM’s business in punched card equipment was growing as rapidly as we would build products. Expansion was in the air at IBM.

In 1951, IBM corporate headquarters in New York decided to establish a research laboratory on the West Coast. Research was a term that at that time covered all types of engineering activity.

The West Coast was chosen because IBM’s customers in the aircraft industry were creating innovations and modifications that were considered to be potential products. For example, the first card-programmed calculator originated with engineers at the Northrop Corporation.

The bay area was chosen as the site of the new IBM research laboratory because it was between Los Angeles and Seattle, where those innovative customers were based.

San Jose was chosen as the specific site because IBM already had a punched card plant here, at 16th and St. John. This plant housed the district manager, an accounting department and a cafeteria. Under the direction of Roger Williams, IBM’s community relations were very good.

“During the first week of January 1952, I was told of my appointment as West Coast laboratory manager. I was told that I would have free rein in hiring a staff of 30 to 50, and that I would be free to choose projects to work on. One-half of my projects were to

be new IBM products and one-half were to be devices in support of customer's special engineering needs. No projects were to be duplicates of work in progress in other IBM laboratories. The laboratory was to be dedicated to innovation.

My first act as manager of the new laboratory was to rent a building, and the second act was to place an ad in all West Coast daily papers, announcing that IBM was opening a laboratory in San Jose. The ad noted that positions were available for scientists, engineers and technicians, and it brought in 400 applications.

The IBM Research and Engineering Laboratory opened its doors at 99 Notre Dame, a few blocks from here, on February 1, 1952.

I was told that my flair for innovative engineering was a major consideration in my selection to manage the new laboratory. During 18 years with the IBM Endicott laboratory, I had had responsibility for numerous IBM products -- test scoring, mark sensing, time clock products, keypunches, matrix and non-impact printers and random card file devices. By 1952, I held over 50 patents, some of them fairly good.

To be given freedom to choose our projects and our staff made the San Jose laboratory an exciting opportunity, especially since funding was guaranteed -- at least for a few years.

The first few months of 1952 were consumed largely in interviewing and hiring a balanced staff of talented and experienced engineers, technicians and administrative personnel.

Except for one person from each of our two New York laboratories, and one engineer from my department in Endicott, New York, we were under orders not to recruit people from the eastern sites of IBM. As a result, our first crew all came from the West.

Among the first projects, undertaken during the start-up were a non-impact printer, a test scoring machine, source recording equipment and a random access replacement for tub¹ files.

It was the search for an automatic random access system to replace tub files that led us to explore magnetic systems.

In 1952 IBM was producing sixteen billion Hollerith cards per year. Each of these cards had to have information entered into it in the form of punched holes before accounting machines could usefully process it. Manual keypunching was one of the most costly items in customer data processing operations. In many applications, most of the information in a card was unchanged from week to week. In a payroll application, for example, only hours worked may be new. An automatic tub file would automatically enter status information and the keypunch operator would be relieved of punching anything but new data.

After deciding that our random access component was to be based on a magnetic recording system, we proceeded to explore the most probable magnetic media. We explored magnetic drums, magnetic tape loops, magnetic plates, magnetic tape strip bins, and even magnetic wires and rods.

Rotating magnetic disks came out on top in our analysis, chiefly because of its rotational dynamics, the potential of multiple accesses and the efficient surface-to-size ratio.

As time went on, our engineers became inspired by the possibility of developing a product that gave essentially instant access to file data, not only when connected to

key punches but also when connected to accounting machines and maybe even to computers.

Two events in 1953 turned out to be fortuitous for our disk project. We ended an automatic data reduction project being done under contract with the McDonald Douglas Aircraft Company. This released half dozen talented electrical and system engineers, who were then available for reassignment to the disk project.

The second fortuitous event was the receipt of a request to bid from the U.S. Air Force Supply Depot in Ohio. They called for a material information flow device. They wanted instant access to each of their 50,000 item inventory records.

We were simultaneously studying file applications in wholesale grocery and wholesale paper supply companies in the bay area.

We pooled these insights and information and prepared a set of specifications for a general-purpose random access memory. These specifications recorded in February 1953, in his notebook by Art Critchlow, turned out to be almost identical to the RAMAC 350 disk file specifications announced two years later.

Going from the wish list provided by the specifications to the reliable operating model required solving many technical problems. With added staff available we proceeded to attack all key problems simultaneously. What kind of disks could we use? How could we get them to run true? How could we best paint them with the iron oxide paint? What kind of magnetic transducer should we use? How could we keep the read/write head close to the disk without having it wear out? And finally, how were we to move the head to any one of 50,000 tracks in less than one second or in one accounting machine cycle?

We proceeded to test out ideas.

We tested the dynamics of rotating disks by mounting 120 aluminum disks two feet in diameter on a shaft with about ¼ inch spacers and rotating this array at 3600 rpm. One test run of this model allayed our fears about problems of excessive wind vibration, power requirements and even excessive disk wobble.

However, one problem that turned out to be quite difficult was coating the disks with iron oxide paint to a uniform thickness and smooth finish. The oxide paint we were using was essentially (except for the addition of magnetic iron oxide particles) the same as was used to paint the Golden Gate Bridge. One of the engineers suggested pouring the paint near the center of a rotating disk and allowing centrifugal force to spread a smooth uniform coat over the disk surface. Another engineer found that filtering the paint through a silk stocking showed that by filling a tray of paper cups with just the right amount of paint, the coating thickness would be the same from disk to disk. This system was used for many years. It was later incorporated into the equipment that automated the process.

Another group of engineers was assigned to develop a small thin head for the record and readback functions. None of our staff knew much about magnetic recording. So we hired a consultant, Al Hoagland, who at the time was a graduate student at UC Berkeley and an expert in magnetic recording. About the only magnetic transducers in use in 1953 were those used with magnetic drum and magnetic tape equipment. Both of these had entirely different space and positioning constraints than we had.

Early in the development of the read/write head we decided to protect the head

against wear by using air pressure with nozzles in the face of the flat head. The airflow spaced the head a uniform distance from the sometimes wobbly disks. Air pressure was also used to force the head toward the disk after it reached its destination.

Stored bit density at the center tracks was made the same as the state of the art density in magnetic drums, 100 bits to the inch and 20 tracks to the inch. This was better than a 4,000% improvement over punched cards in information density and the data was alterable and erasable.

In our first file model two-foot diameter oxide-coated disks were mounted on a horizontal shaft at ½-inch intervals. Fifty-one disks gave 100 inside surfaces. Two opposite facing heads were mounted on one access arm. The access arm was moved so as to place the heads on any of the 100 tracks on each of the 100 disk surfaces at a speed that would match an accounting machine cycle, which was less than one second. We had anticipated that there would be a need for as many as six access stations on each file. Given positioning tolerances, even with two actuators the access positioning of each was limited to the top or bottom half of the disk stack. The maximum travel between addresses was one-twentieth of an inch.

Two access drive systems were designed and modeled, one mechanical and one electronic-servo system. We finally chose an electronic-servo system for the first file model.

The tub file application led us to test out the disk performance by pairing it with a keypunch, because the keypunch could be used for entering information and for recording in punched holes the information read back from the disk file.

On February 10, 1954, this first sentence was fed into and read back from the disk file – “This has been a day of solid achievement.”

By March 1954, tests of components and the card to file machine made us confident of being able to build a product. Lou Stevens was made the manager with full development responsibility.

He and his staff of very capable engineers initiated a program of re-design that started in mid-March 1954. By November this design had matured into a “magnetic disk processing machine.” RAMAC was on its way. Hopes were high that this revolutionary concept would develop into an IBM product.

The potential for large random access memory was attested to be activity among competitors who were using very large drums, drum arrays, tape loops and even the surface of a power station fly wheel as recording surfaces.

The IBM vice-president for marketing, L.H. LaMotte, stationed his long-range planner, F.J. Wesley, in our laboratory mid-summer, 1954. On October 8, 1954, Mr. Wesley sent a memorandum, which he called a “pontifical announcement,” to his boss. In part it started, “we must immediately attack accounting problems under the philosophy of handling each business transaction as it occurs, rather than using batching techniques.” Wesley’s memo was widely circulated among IBM management and, needless to say, in our laboratory. The promise of developing a product for more than a file tub replacement led to a corporate decision in November 1954 to build at least five prototypes of our product to field test.

Initial specifications for the RAMAC were prepared December 17, 1954.

The non-RAMAC projects of the research and engineering laboratory moved to Julian Street to open more spaces for the RAMAC team. At the Julian Street

laboratory work continued on advanced ideas for disk files. Gliding heads and multiple parallel access arms were developed and eventually transferred to Lou Stevens's domain where these features were incorporated into disk file products that followed after the RAMAC 350.

Lou Stevens and his engineers at 99 Notre Dame successfully demonstrated and operated the Model II file on January 16, 1955. Debugging of this machine continued around the clock for months.

Early in 1955, a corporate decision was made to build 14 RAMAC 350 machines for internal use and field-testing.

On May 6, 1955, IBM held a press conference to announce that it was bringing out a new product "that takes information from a stored program using a multi-million character random access memory that makes it possible for the new system to do a whole job automatically without using batch processing." On August 25, 1955, the San Jose Mercury News published a short article which said, "IBM plans a giant new San Jose plant that may employ more than 5,000 here."

The first RAMAC was shipped to Zellerbach Paper Company in June 1956, by an outstanding San Jose manufacturing team. Mr. Porter remembers the event.

T.J. Watson Jr., President of IBM, announced the RAMAC on September 4, 1956. He said in part, "This is the greatest product day in IBM's history and I believe in the office equipment industry."

The RAMAC was featured in the United States exhibit at the World's Fair in Brussels – and in the United States Technical Exhibit in Moscow. Chairman Khrushchev of Russia came to San Jose in 1956 to visit the RAMAC plant.

The American Society of Mechanical Engineers designated the RAMAC an international historic landmark on February 27, 1984.

The first four years of disk file history have been reviewed in my remarks. In the 33 years since the RAMAC 350 left the laboratory, tremendous advances have been made in product development, manufacturing and research, as you all well know."

APPENDIX 2

Major Honors and Award Related to the RAMAC

- **1986**

National Medal of Technology awarded to **Rey Johnson**, “Father” of the RAMAC

Citation: Introduction and development of magnetic disk storage for computers that provided access to virtually unlimited amounts of information in fractions of a second and is the basis for time sharing systems and storage of millions of records. Over \$10 billion in annual sales and over 100,000 jobs arose from his contributions.

- **2008**

Inventors Hall of Fame Louis Stevens: Data Storage Machine Patent # 3,134,097

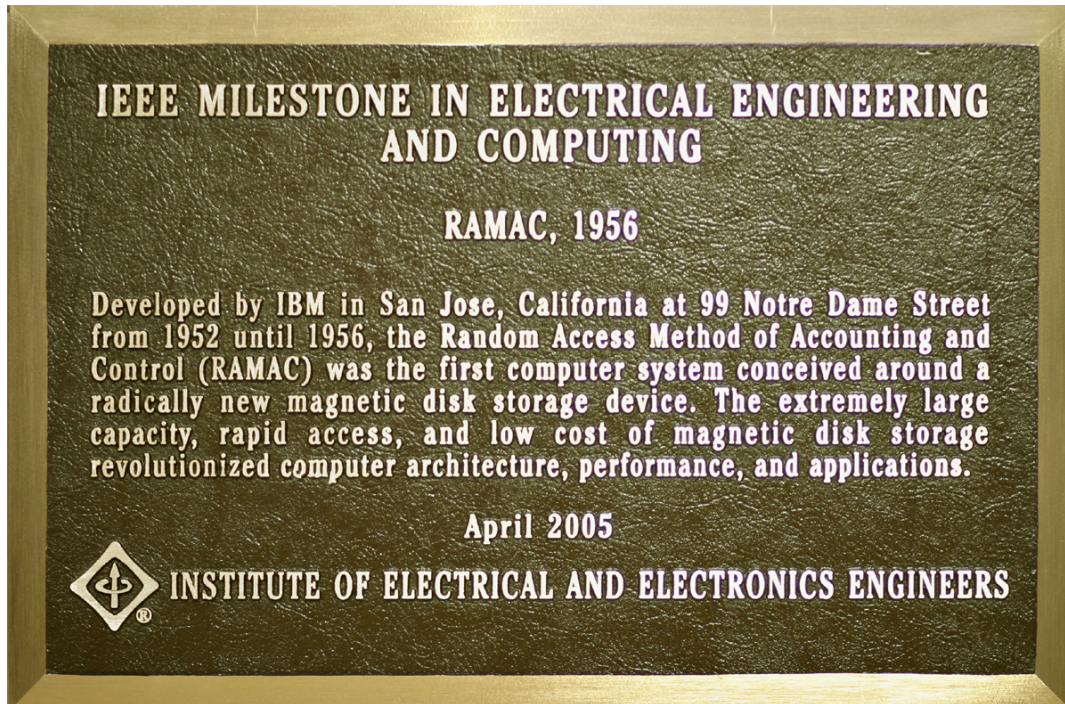
Citation: Louis Stevens, working with William Goddard and John Lynott and a team of engineers, invented a unique magnetic disk storage device. Their disk drive, announced in 1955, was a key component of the IBM 305 RAMAC machine, and allowed users to store and almost instantly access large amounts of data.

- **February 27, 1984**

IBM 350 RAMAC Disk File designated an **International Historic Landmark** by The American Society of Mechanical Engineers



- **May 26, 2005**
RAMAC designated as an IEEE Milestone.



- **January 22, 2002**
Historical Landmark Designation for 99 Notre Dame, San Jose, California

The small IBM laboratory that opened in February 1952 at 99 Notre Dame in San Jose, CA was the birthplace of the magnetic disk drive. The facility was a leased one-story building of 8000 square feet to house a group of only 50 people. The original building had been expanded to 14,000 square feet in the early 1950s to accommodate the growing activity on the disk drive but otherwise preserving the same external façade.

2002 was the 50th anniversary of the founding of this original IBM laboratory. I decided to investigate whether the original building still existed. I found the original building only had two small groups (one city and one religious) in residence. The facility looked long overdue to be torn down. It immediately dawned on me that since this site was of a great historical significance it should be preserved.

An effort was undertaken to have the building designated a City Landmark. The Magnetic Disk Heritage Center spearheaded the campaign. After my visit, the owner of the property renovated the facility and leased the building to the County Court until a planned major new County Courthouse became available to consolidate the many small court facilities in the county. Fortunately, from my point of view, while the interior was totally refurbished, the external appearance was maintained. In addition, upon request, an exception to normal practices was made and the judges unanimously approved the placing of displays and plaques at the site reflecting its historical significance. Our efforts

lobbying the San Jose City Council for a landmark designation were successful and on January 22, 2002, 99 Notre Dame, San Jose, CA was formally approved as a City Landmark.

The building includes large panel displays on the walls of the lobby waiting room describing the early history of the RAMAC. The ASME (American Association of Mechanical Engineers) International Landmark plaque and the IEEE (Institute of Electrical and Electronic Engineers) Milestone plaque, both recognizing the historical significance of the RAMAC, are located on the outside of the building. A special event in recognition of the 50th anniversary of the birthplace of the RAMAC at 99 Notre Dame was held in 2002 at the site.

Note: At the present time the status quo remains since construction on the new Courthouse has yet to begin.



99 Notre Dame, San Jose, CA