Workshop XVI
Keynote Talk

5 decades of magnetic data storage

A. S. Hoagland

1940's -- prologue:

UC Berkeley

Advances in technology during WW II led to tremendous interest in computers as the capabilities of the first electronic computer, the ENIAC, became generally known. The ENIAC used thousands of vacuum tubes so there was a great desire to realize a lower cost design. I started graduate studies at UC Berkeley in 1948 and joined a newly funded ONR computer design project whose goal was to develop a prototype of what now could be called a "departmental" computer. Universities carried out much of the basic research in electronic computers in this period as no significant commercial marketplace existed. We chose a magnetic drum for the memory. At this time data storage, based on digital magnetic recording, was a totally new application of magnetic recording. Up to then the technology focused on the analog tape recording of sound. . New aspects included digital data, update-in-place, short access time and a high data rate. The latter features required the medium move at high speed relative to the head, necessitating a non-contact head/medium interface. In the division of labor, I was given responsibility for the magnetic drum, actually by default, as the desired preference of everyone was to do logic design. The areal density was 800 bits per square inch. But I am not going to talk about that project.

1950's

IBM San Jose

In 1952, Rey Johnson, of IBM, moved to San Jose with the assignment to start and run a new exploratory research laboratory. The purpose of this new facility, deliberately remote from the mainstream IBM development laboratories, was "to go where no man has gone before". His mission was to seek and develop radically new information processing capabilities. Rey was not only a true inventor but had a genius for encouraging and supporting creative ideas by those in his organization. He saw punched card data processing, then the mainstream, as an attractive target. As a point of reference, at that time the use of punched cards as data records forced batch processing, where rather than handle each transaction as it occurred you sorted and processed a large number of transactions sequentially. For applications like inventory control, this was clearly not the method of choice. He decided to design a data processing system for "random" transaction processing, requiring a very high capacity, short access time data storage device.

The laboratory evaluated all sorts of technologies, device configurations and system designs to determine the best approach. Magnetic recording on a stack of disks was selected as the most promising storage component.

Given my activity at Berkeley, I was approached by Lou Stevens (a former student colleague who had joined IBM and now worked for Rey), to do consulting for their disk drive design and development program.

From my viewpoint at Berkeley, the chance to become involved in the RAMAC project was a fantastic opportunity. After all, it was not clear then that expertise in digital magnetic recording offered much of a future. I became so engrossed in disk drive work that in 1956 I joined IBM to work under Rey. By this time the RAMAC was moving into product engineering at the new plant site and Rey was cranking up an advanced disk drive program called the ADF (Advanced Disk File). This device was key to expanding transaction processing into on-line real-time applications, just the sort of new challenge Rey always sought. The ADF technical objectives included 10 times the density of the RAMAC with a flying head per surface (the RAMAC heads required pressurized air). In his usual style, he decided to pursue radical innovations and in addition to flying heads the ADF included steel disks and perpendicular recording

The ADF transferred to the product development organization in 1957. Rey's lab became part of the new IBM research organization and I remained with him and started the SDF (Single Disk File) project. An objective of this program was to design a drive to replace magnetic tape in many applications. Again,in the spirit of his operation, the SDF targeted 10 times the density of the ADF and introduced servoing, using a pattern on the disk, to achieve both a major increase in tpi and disk interchange.

The IBM laboratory provided the bedrock technology and product design directions for disk drive generations to come. I believe the most important contribution during this period was through initiating and accelerating the introduction of transaction processing systems that revolutionized information processing. In this regard, John Haanstra, also a former student at Berkeley, was the key contributor. Rey Johnson deserves all the honors he received and these accomplishments were due to his fundamental goal of bringing new solutions to data processing and not just a storage device.

1960's

IBM (San Jose, Netherlands, Yorktown Research

In early 1960 the ADF engineering problems rose to the level of a major corporate crises, due to the growing need for the advanced performance features of this product. Al Shugart was brought in as manager and I was temporarily assigned to the program to help on the head and disk aspects. Within the year the project was put back on schedule. The changes included dropping perpendicular recording and steel disks and going to advanced RAMAC longitudinal recording.

Since the SDF project was viewed as falling under the mission of the magnetic tape group back East the project was transferred there. However, San Jose was given corporate responsibility for disk drives and their efforts were terminated.

In 1962 I moved to Europe join an IBM/Philips Study defining an on-line transaction oriented replacement for the existing magnetic tape based Post Giro system (the government national banking system). Interest in disk drive systems was rapidly increasing and in addition to this study I was provided specific support, while in Holland, to complete my book “Digital Magnetic Recording” (the remained the only reference work on this topic for years) published in 1963.

I was invited to join the new IBM Research laboratory in Yorktown in 1964 as manager of engineering sciences. I found that the IBM Chief Scientist, Manny Piore, viewed magnetic recording as "mature". Actually, in fact it was over 60 years old at that time! The laboratory enthusiasms were for alternative technologies. The engineering sciences groups were engaged in electron beam recording, holographic recording and different varieties of optical storage. Even when I later became director of research planning and tried to increase magnetic recording storage work, I found my influence was limited. Even the San Jose research group was rapidly losing interest in magnetic recording and shifting into alternative technologies.

I had frequent occasion to visit IBM San Jose, a fair number of times as a member of corporate audit teams trying to understand why magnetic drive programs were constantly in trouble. The problems encountered only reinforced the skepticism within the company about the continuing viability of magnetic recording storage. Disk drive development was still viewed as "grunt" engineering with no perceived long-term future. These perceptions came from the lack of any real "science" and limits that were based only on engineering guesstimates.

However, IBM essentially had an exclusive in disk drives, and all disk drive products were were tied to IBM systems introductions which set price and product life. The business prospered. The success of the American Airlines Sabre Reservation System was of particular importance in establishing confidence in and the potential of disk based on-line real time transaction processing.

1970's

IBM (Boulder, San Jose)

Given my identification with a technology perceived to have a short half life, I was “offered” the opportunity to continue my return to California through an intermediate stop in Boulder (where the IBM tape group was now located) to run a joint Research/Development benchmark program. The objective was to explore the limits of flexible media magnetic recording. Since the Boulder laboratory had a large development group on magnetic tape product programs, the benchmark chose to pursue the potential of flexible magnetic disks. The project first focused on a flexible disk pack and then decided to pursue a single flexible disk drive. In mid decade the entire IBM tape organization was moved to Tucson (suffering great attrition) and I opted to return to San Jose. The major long range impact from this benchmark program came from the migration of key team members from IBM to form with others a new startup called Iomega.

During this decade, within the larger data storage technical community, the major preoccupation was with the so-called "Access Gap". This term referred to the orders of magnitude difference in access time between main memory and mechanical storage and led to innumerable technology proposals and projects on how to bridge this chasm
.

By far the most popular and winning choice within the technical community became magnetic bubbles. In fact because of this, the late 1970's were an extremely precarious time for magnetic disk drives. All the advanced R&D work was focusing on alternative technologies, with the majority on bubbles, in spite of the increasing performance demands being placed on disk drive development for new, more advanced, systems.

These two forces were coming into direct conflict at this time. On one hand, the pool of available magnetics talent was rapidly being channeled to aggressively move ahead with magnetic bubbles, seen as the ultimate successor to magnetic disk drives. On the other hand, the most advanced disk drives had reached a point where the technical sophistication had finally exceeded existing engineering capabilities. Fundamental magnetic recording research in relevant fields did not exist. The dilemma posed by this situation needed to be resolved quickly. Fortunately, this expedited a major reevaluation of the promise of bubbles. On review, the bubble burst, quickly and in dramatic fashion. Essentially, the projections indicated their costs could not compete with magnetic disk or their speed with semiconductor memory devices and there was no defensible ground in between. Within IBM the magnetics expertise in bubbles was quickly reassigned to work on magnetic recording problems, and just in time to make major contributions to the solution of the major problems encountered by the then recently announced 3370/3380.

Thus, the decade was marked by major doubts as to the long-range future of magnetic disk drives with very large efforts to find an alternative technology for the future. It ended when the magnetic disk drive was recognized as still having an untapped potential for storage density advances and ended up the survivor.

1980's

IBM (San Jose, UCSD), SCU

By far, the most significant factor in this decade was the emergence of the PC. The PC, a consumer product, generated a need for low cost disk drive components in volumes that would grow dramatically. The PC led to the creation of a disk drive industry, one that became a high volume source of somewhat generic disk drive storage components to the PC companies. The narrow margins and small size of the initial market lead to a minimal response by existing drive companies, encouraging many new entrants to enter the industry, using the existing technology base but designing and packaging drives into smaller form factors specifically oriented to the PC. The consequence of the increased competition in low-end drives became central to driving future technology advances

In addition, the increasing sophistication and interdisciplinary nature of the technology led, early in this decade, to recognition that a major industry partnership with academia would be vital. (At the time Jack Judy was the only identified academic doing magnetic recording studies in contrast to the situation in semi-conductor electronics where almost all EE departments had programs in semi-conductors).

Having made this suggestion within IBM, I was asked by Art Anderson, then President of SSD (Storage Systems division), to try and create an industrial consortium (involving major corporations that could make significant financial contributions) to support one or more University Centers dedicated to magnetic recording technology

The consensus choice for the ideal location of a Center was the Bay Area and after that in California. However, no major local University (Stanford and California) was willing to respond by making a significant commitment. UCSD did and I spent a year and half in San Diego involved in the startup phase of CMRR (while still with IBM). (A CMU proposal was also supported). Since San Diego was remote from the heartland of the disk drive industry, I then returned to San Jose. However, the fact that the strong wish for a center in the bay area was unfulfilled and I was already heavily involved with university/industry interactions led me to take advantage of this opening and leave IBM to join Santa Clara University and found IIST (Institute of Information Technology).

In terms of growth, excitement, enthusiasm and innovation this was a very exciting decade both in technology and in products. On the technology side the merits of perpendicular versus longitudinal recording generated tremendous interest and controversy, fueled by the rapidly expanding magnetic recording research community.

1990's

SCU

Continued technology advances, driven particularly the introduction of the MR/GMR head supported the achievement of a 60% to 100% CGR in areal density. That said, more impressive in many ways is the fact this sharply higher growth rate was forecast at the beginning of the decade and continued over this whole period.

How could this happen? One explanation. No one can accurately forecast the technical limits or rate of change of the technology due to its complex interdependency on many different disciplines. (There is no such thing as optical wavelength to hang your hat on.) So self-fulfilling prophecies come into play. When a projection is made at a major meeting the audience is put in the following position. If you have no conviction then you must assume your competitor can do what he is projecting and your organization had better act accordingly. The individual who made the projection is stuck with the fact that, even if he does not believe what he said, he is forced to assume his competitors in the audience may and thus act accordingly. This leads to an essentially common technology path and timeline that is accepted industry wide.

The focus on publicizing both a 100% CGR, rapidly escalating competitive benchmarks and while defining limits that are still are only a factor of 10 or so away appears to me to not have served the industry well. At SCU, the interest of students in this field dropped off rapidly in the later part of the nineties. The higher the rate of advances the shorter a possible career foreseen in the field. Even another extension by a factor of 10 of the perceived limits is predicted; this only means an additional 4 more years to the student.

If you can predict your rate of progress accurately 10 years in advance, where is the "breakthrough" to excite a layperson or --- what is so amazing about doing what you said you would do for years?

The more important change in this decade has been emergence of a new and very successful set of companies that design sophisticated storage systems, based on thousands of disk drive components. These storage systems provide capabilities crucial to the new and more demanding applications possible with the Internet and changing business organizations in general.

While Rey Johnson is best remembered for the first disk drive, I suspect his achievements are far greater through the impact of the RAMAC system in accelerating transaction oriented data processing, which was what he really set out to do. The impact of the disk drive far transcends the data storage industry itself.