History of Magnetic Disk Storage Based on Perpendicular Magnetic Recording

Albert S. Hoagland, *Fellow, IEEE*

***Abstract*—This paper reviews a major disk-drive product development effort in IBM during the 1950s, one that was originally based on perpendicular recording. The goal of the project, identified as the Advanced Disk File (ADF), was to develop the successor to the first disk drive, the Random Access Method of Accounting and Control (RAMAC) 350. The ADF eventually became the IBM 1301, the first disk drive to use a flying head per surface. Whereas the RAMAC used longitudinal recording, the ADF project chose perpendicular recording as the mainstream recording technology for magnetic disk data storage and continued on this path for a period of five years. A crisis arose in the later stages of prototype testing due to unacceptable failure rates; one decision made was to change the recording method from perpendicular back to longitudinal. The reasons for the original selection of perpendicular and the subsequent return to longitudinal recording are described and address technology, product, and business issues. General observations on factors influencing the choices made and the final outcome are offered as well as the long-term impact of these events on the disk-drive industry.**

***Index Terms*—Disk-drive product design, IBM 1301, perpendicular recording, RAMAC.**

I. Introduction

IN RECOUNTING the history of perpendicular recording for magnetic disk data storage, I will primarily focus on an IBM development project called the Advanced Disk File (ADF). The disk-drive product that emerged was the IBM 1301, which was the first disk drive to embody a flying head per surface, making it the actual precursor of all subsequent disk drives. The historical importance of this drive is that it was the first to support *real time* online transaction processing, a capability which revolutionized the computer industry. From 1955 until 1961, this was the major disk-drive project at IBM, intended to ensure that IBM, the pioneer in magnetic disk storage, would maintain the leadership position it commanded. It is not widely known that from the inception of this project in 1955 until February 1960 this disk-drive development was predicated on perpendicular recording. The reasons the ADF was based on perpendicular recording instead of pursuing the Random Access Method of Accounting and Control (RAMAC) longitudinal recording technology and why, after almost five years of effort, perpendicular recording was dropped in favor of a return to longitudinal recording, are the main focus of this paper.

Manuscript received October 24, 2002; revised January 24, 2003.

The author is with Santa Clara University, Santa Clara, CA 95053 USA (e-mail: ahoagland@scu.edu).

Digital Object Identifier 10.1109/TMAG.2003.813786

A. First Exposure: Magnetic Drum Memory

In the late 1940s, a new computer design program at the University of California at Berkeley (UC Berkeley) was funded by the Office of Naval Research to pursue an intermediate sized computer based on a magnetic drum memory. A magnetic drum memory provides rapid data access, a feature that necessitated a finite separation between the head and the medium to avoid wear at the high surface speeds required. Noncontact data recording, being a major departure from the typical tape recording of analog signals at that time, generated interest in exploring unique head designs. In the fall of 1949, among the novel head designs tested was "a pointed magnetic recording head with a very wide gap, plus a special composite magnetic surface on the drum. This surface consisted of a layer of Permalloy beneath the usual layer of magnetite, so that in operation the magnetite dipoles are perpendicular to the surface rather than along the surface circumferentially" [1]. The preliminary results from this perpendicular recording investigation showed a higher linear density potential than that obtained with the wide-pole heads then available. The advantages of increasing linear density were twofold: it would simultaneously increase the data rate and reduce the number of heads required for the same capacity. The 1949 status report concluded that more development effort was needed. UC Berkeley's computer design program schedule did not allow for further studies in this direction. However, the potential advantages of the novel longitudinal and perpendicular head designs compared to conventional contact tape heads were apparent.

B. Background: The First Magnetic Disk Drive

In 1952, Rey Johnson was asked to set up a small IBM research laboratory at 99 Notre Dame, San Jose, CA. An outstanding inventor, Johnson was given the freedom to choose his own projects and proceed independently outside of the IBM engineering organization. He quickly selected as his main goal the development of an online transaction processing computer system that could replace the ubiquitous punched card systems of the day, which relied on batch-processing procedures. Such an approach required a low-cost storage device with a large capacity and short access time to any data record. There were no existing devices or technology activities underway that could satisfy this need.

After studying many options, Johnson's final choice for implementation was a rotating magnetic disk stack, based on its volumetric efficiency for recording surface area and head positioning to data tracks to provide direct access to individual records. Perhaps the greatest challenge was the need to maintain a very small spacing between the heads and disks. The solution

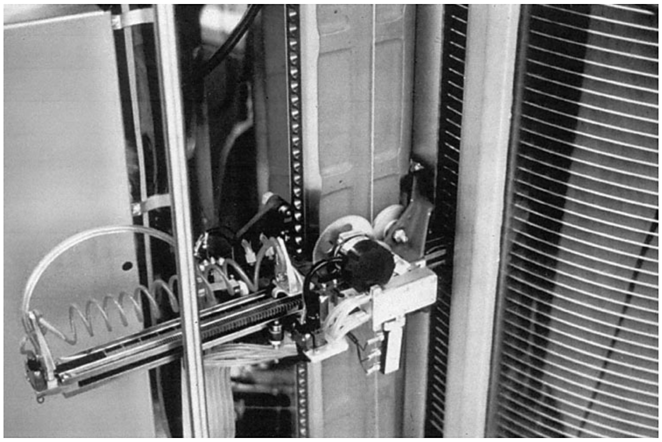

Fig. 1. RAMAC 350 disk-drive access mechanism.

was an ingeniously designed pressurized air-bearing slider. The cost and complexity of this feature led to the choice of an actuator-arm assembly with a head pair—one for the upper surface and one for the lower surface of a disk. Thus, the positioning mechanism transported the head pair up and down the stack to the desired disk and then in and out to a specified track [2]. The magnetic head involved a unique design having very narrow pole tips to maximize density with the head-to-disk spacing of 25 μm. Fifty 61-cm (24-in) aluminum disks coated with iron oxide paint were needed to provide the capacity required for the targeted transaction processing applications. The drive, storing 5 MB at 100 bits/in and 20 tracks/in, had to be able to record or retrieve records spread over 240 square feet of surface area in less than a second (Fig. 1). At first, there was a great deal of skepticism among the engineers that such a mechanical device would ever work [3].

The first transaction-processing computer was the RAMAC 305 system designed around the RAMAC 350 disk drive. From this beginning, considerable expertise was acquired in magnetic heads and disks for noncontact longitudinal recording. While upgrades to the RAMAC disk drive were continued, the mainstream technology efforts for the future were looked for from the advanced technology developments undertaken by the ADF project.

II. ADF: THE SECOND DISK DRIVE

Even before the RAMAC was announced in 1956, a decision was made in 1955 to start on the next generation disk drive, a project named ADF, which had the objectives of obtaining ten times the capacity and 1/10 the access time of the RAMAC. These advanced performance targets for the ADF were essential to expand data processing into major new applications being identified. The first of these was the American Airlines Airline Reservations System (Sabre), which depended on a geographically dispersed *real time* online transaction processing capability. In addition to a commercial version having a single read–write channel for each disk module, a unit with parallel data transfer was crucial to Stretch, the new supercomputer which IBM was counting upon to gain the lead in high-end data processing. The delivery date of the Stretch drive preceded that of commercial version, causing a compression of product development schedules for the underlying technologies. The goals required a radical departure from the RAMAC design as well as major new disk-drive technology advances. The ADF project was viewed as the advanced product development program for the next generation of disk drives and crucial to the data processing systems strategy being developed.

Fig. 2. ADF head-arm assembly module.

A. Design Choices

1) The Flying Head: The only way to dramatically reduce access time was to limit positioning to radial motion only, which called for a head per surface. The pressurized air-bearing slider design used on the RAMAC was complex as well as uneconomical when many heads were needed. A flying (or self-acting air-bearing) slider, on the other hand, is much simpler. Moreover, with better head and disk surfaces, the flying head also offered the best opportunity to reduce head to disk spacing and, therefore, increase linear densities. The head–disk spacing would have to be reduced to about 6.5 μm to meet the linear density increase required. There was little understanding at this time of flying head design and behavior, particularly on a stack of large disks that had significant runout, surface distortions, and uniformity problems. However, there was no alternative to the flying head and so this technological approach was accepted even though the level of air-bearing theory and slider design guidelines necessary to support this decision did not exist.

2) The Recording Method: The next major decision was the choice of recording components. The planned product was to be available with either one or two disk modules of 25 disks each on the same shaft. Thus, there could be 50 disks in the vertical stack or 98 flying heads. (Fig. 2 shows one such

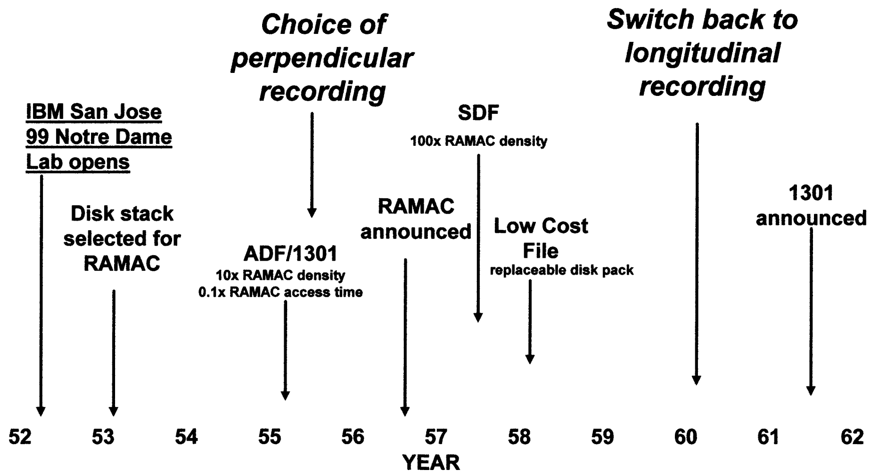

Fig. 3. Time line, magnetic disk-drive projects at IBM San Jose in the 1950s.

module.) In discussions at this time, a novel shielded single lamination probe was proposed as the best solution for a simple and low-cost head structure, a design feature of real importance given the number of heads per drive. The RAMAC experience with head–medium contact causing debris generation and occasional crashes strongly supported the conviction that a harder disk surface was needed for the ADF, especially with a head per surface. Attention turned to find a *compatible* disk for the proposed head that would have a very hard surface. A steel disk was chosen; it was known that oxidizing magnetically soft steel in a steam atmosphere could create a hard thin layer of magnetite on its surface. The additional perceived advantage for the "steam homo" disk was the ability to produce many disks at one time, in contrast to the individual coating of painted aluminum disks. The appeal of the oxidized steel disk lessened interest in looking for other media for perpendicular recording.

The decision was made to pursue flying heads along with probe heads and steel disks, that is perpendicular recording. This choice was primarily influenced by the perceived advantages of a probe head combined with a disk surface harder than that on the painted aluminum disks. The ADF started as an innovative response to the challenge of creating an entirely new disk drive, reflecting the same spirit that prevailed in the original RAMAC program.

This development program incorporated several new technologies, two of which were highly interdependent. Evaluating perpendicular recording was affected by the unpredictability of the flying head behavior and the variable magnetic state of the steel disk itself. The acquisition of performance data and its interpretation resulted in many uncertainties in setting design directions, which were magnified by the lack of sophisticated instrumentation.

B. Risk Summary

With the ADF, IBM became the first and only major company in the computer storage business to make a change to perpendicular recording for its main line of disk-drive products, a business where it totally dominated the market. Whereas there was really no viable technical alternative to flying heads to meet the performance and cost goals, this was not the case with respect to magnetic recording options.

Flying heads were planned for large 61-cm (24-in) disks without any adequate design theory providing criteria for head contours, head and media surface finishes, etc., as well as without adequate measurement instrumentation. Unlike longitudinal recording, where there was considerable experience derived from the development of the RAMAC, there was little knowledge of perpendicular magnetic recording. The project committed to perpendicular recording not for any clear density gains [4], but for the anticipated benefits from a harder disk surface with the expectation of a lower magnetic head cost.

C. Environment and Progress

Two major factors hindering progress were the drain on resources from the need to continue support for the RAMAC and the start of development activities on other new disk drives in order to broaden market penetration as rapidly as possible. In April 1957, ADF efforts were reduced due to the need to lend some engineers to assist in the release to manufacturing of the RAMAC. In the summer of 1957, as part of the newly created IBM Research Division, a separate San Jose research laboratory was formed; its staffing also included transfers from the ADF project. This research laboratory started a new project called the Single Disk File (SDF), which had the objective of achieving ten million characters on a single replaceable disk (100 times the density of the RAMAC). Furthermore, in early 1958, the product development laboratory started work on a low-cost system design based on a small disk pack. Thus, during this critical period, there occurred a major dispersal of key technical resources (Fig. 3).

In February 1958, the first ADF module tests were made. The trials were sufficiently encouraging to justify a year-end target date for an engineering model of the ADF [5].

The first fully assembled prototype model was finally available in August 1959. By this time, pilot production had been started on probe heads and hundreds of units were made. (Fig. 4 shows the final perpendicular head design for the ADF.) In terms of the planned schedule, the test results were

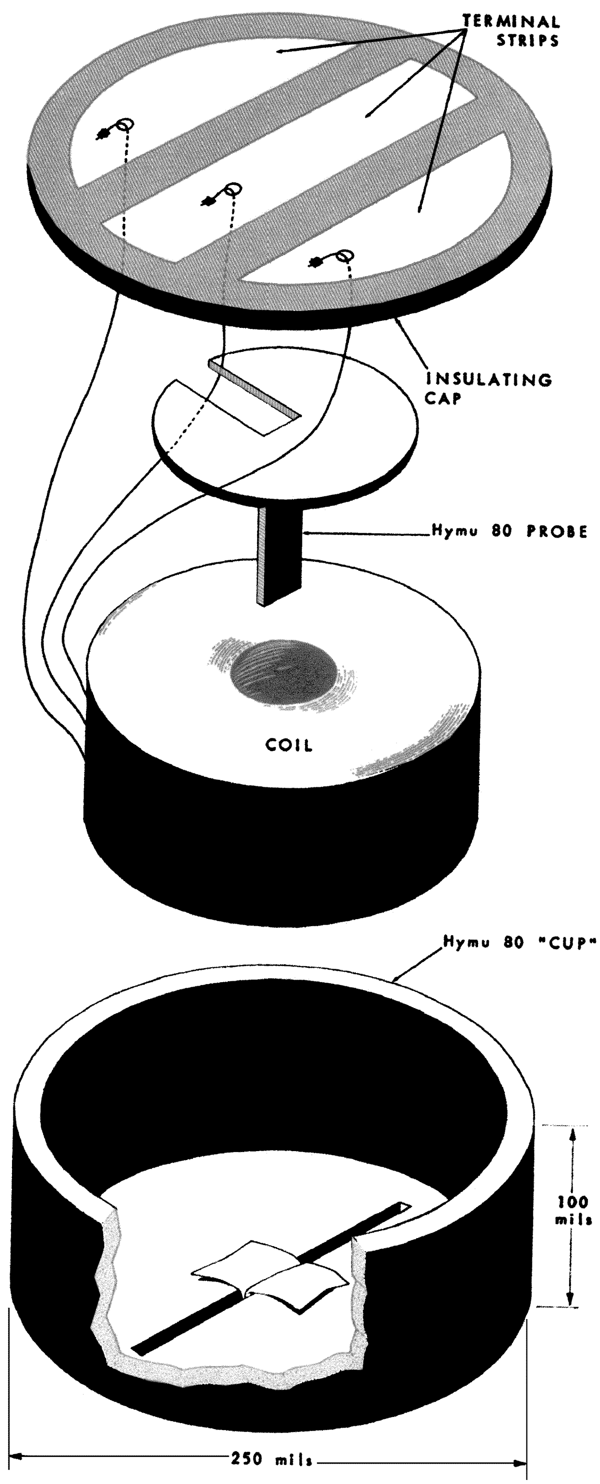

Fig. 4. ADF perpendicular recording head, schematic.

discouraging. Head–disk interference problems persisted and the behavior of the air-bearing slider and the steel disk were both identified as causes. In particular, it was clear that the surface quality of the steel disks was still unacceptable.

D. Crisis and Change

By January 1960, the ADF project reached a major crisis. Continued testing still resulted in numerous failures and the program was so far behind schedule the original plan was meaningless. IBM Corporate Headquarters called for a major review. One obvious conclusion was that the program was based on several new technologies, all of which were still not fully understood and none of which had been adequately tested independently. One measure of the seriousness of the situation was a proposal to look for a possible alternate supplier for the disk drive. Technical recommendations were made leading to the following decisions [6].

1) Immediately set up a major research and development program on hydrodynamic air-bearing sliders.
2) Abandon perpendicular recording and start advanced development work on RAMAC heads and disks to meet the ADF specifications.
3) Use pressurized air-bearing heads and oxide coated aluminum disks for the Stretch file unit.

This last step not only improved the chances of achieving the earliest date for shipment of the Stretch unit (the high costs could be tolerated for this one of a kind device) but also bought additional time for the basic technology studies required.

The Stretch file was shipped from San Jose to IBM Poughkeepsie for systems testing in the fall of 1960. The 1301 disk drive, with flying heads and longitudinal recording, was announced in 1961. In retrospect, it proved fortuitous that upgrades to the RAMAC disk drive and exploratory higher density efforts such as the SDF led to more advanced longitudinal recording components. This situation provided a clear direction for replacing the steel disks and probe heads in an expeditious manner. The ensuing advances made in flying head design were sufficient to alleviate the original concerns over the mechanical surface properties of the oxide-coated aluminum disks.

III. Summary

Over a period of almost five years, the most important disk file project in IBM San Jose and one absolutely crucial to IBM's position in data processing proceeded on the basis of perpendicular recording. One can only speculate on the evolution of future disk-drive technology had the ADF been successful in adopting perpendicular recording. However, my perspective is that there were major consequences from these events on the course of disk-drive development. This experience made it clear basic research was essential to continuing advances in magnetic disk storage, and a major IBM Research Laboratory was started in San Jose. The successful turnaround of the ADF 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 the industry. In addition, this experience led to a more conservative approach to the introduction of technological advances in disk-drive development for a number of years. It came to be accepted that only one major technological change should be made for each succeeding generation disk drive. A factor reinforcing this attitude during these years was that disk drives were designed in conjunction with planned data processing systems. These systems guided disk-drive product features, schedules, and manufacturing volumes; consequently, the financial risk was assessed in terms of major systems markets and sales forecasts. One other lesson that may be drawn is that a specialized version of a mainstream product should follow and not precede the final development phase.

Almost nothing has been written about this perpendicular recording history either within or outside IBM. It is almost unimaginable that IBM would have pursued perpendicular recording as their future direction for so long and then abandon the approach; on the other hand, it is equally unexpected that in spite of this, IBM maintained a commanding lead in magnetic data storage that was not threatened. For perpendicular recording, it may have been an opportunity lost. Recent interest in perpendicular recording makes its earlier history of some possible relevance in the continuing comparisons between longitudinal and perpendicular recording.

References

[1] P. Morton, "Report on the status of contract N70NR-295-XIV, NR-o43-066," Intermediate Electronic Computer Project, Univ. California, Berkeley, Sept. 1–Dec. 1, 1949.
[2] T. Noyes and W. E. Dickinson, "The magnetic disk, random access memory," presented at the Western Joint Computer Conf., San Francisco, CA, Feb. 7–9, 1956.
[3] D. Kean, *IBM San Jose: A Quarter Century of Innovation*. San Jose, CA: International Business Corporation, 1977, p. 23.
[4] A. S. Hoagland, "High-resolution magnetic recording structures," *IBM J. Res. Develop.*, vol. 2, pp. 90–104, Apr. 1958.
[5] C. J. Bashe, L. R. Johnson, J. H. Palmer, and E. W. Pugh, *IBM's Early Computers*. Cambridge, MA: MIT Press, 1986, pp. 300–310.
[6] A. S. Hoagland, H. A. Mussell, and J. M. Norton, "ADF technical audit committee report," unpublished, Feb. 8, 1960.