

Technological impact of magnetic hard disk drives on storage systems

by E. Grochowski
R. D. Halem

Magnetic hard disk drives have undergone vast technological improvements since their introduction as storage devices over 45 years ago, and these improvements have had a marked influence on how disk drives are applied and what they can do. Areal density increases have exceeded the traditional semiconductor development trajectory and have yielded higher-capacity, higher-performance, and smaller-form-factor disk drives, enabling desktop and mobile computers to store multi-gigabytes of data easily. Server systems containing large numbers of drives have achieved unparalleled reliability, performance, and storage capacity. All of these characteristics have been achieved at rapidly declining disk costs. This paper relates advances in disk drives to corresponding trends in storage systems and projects where these trends may lead in the future.

Magnetic hard disk drives are used as the primary storage device for a wide range of applications, including desktop, mobile, and server systems. In 2002, nearly 200 million disk drives were manufactured worldwide,¹ with the total capacity to store more than 10^{19} bytes.

Since the first disk drive was introduced in 1956, drives have undergone a rapid evolution, thanks to the application of new magnetic, electronic, and mechanical technologies. These developments have yielded storage devices with very significant capacity and performance increases. Early disk drives were

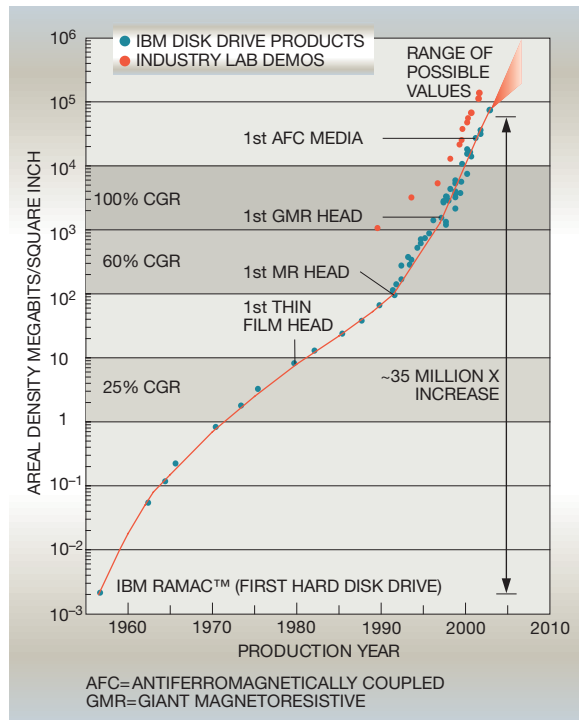
specialized and very expensive, whereas today's drives are nearly commodity items and are universally available. Although the fundamental architecture of disk drives has changed very little in the years since their introduction, the geometric size of drives has been reduced almost to the point of micro-miniaturization, and these smaller sizes have resulted in storage system characteristics that offer new horizons in data retention and availability. The trends characterizing storage systems in which large numbers of drives participate as a single storage unit have followed the evolutionary behavior of their principal component, the hard disk drive (HDD). Storage system characteristics are also influenced somewhat by components other than disk drives, including DRAM (dynamic random-access memory) caches and buffers, cooling systems, frames and cases, and system software.

Areal/volumetric densities

Areal density, a traditional measurement for disk drives, determines capacity, internal (media) data rate, and ultimately price per unit of capacity.² Figure 1 shows the areal density improvement for hard disk drives since 1956. Significant trend changes have occurred when new technologies have been adopted, so that today's CGR (compound growth rate) is essentially 100 percent or doubling every year, and a 35-million-times increase in this parameter has been

©Copyright 2003 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

Figure 1 Hard disk drive areal density trend

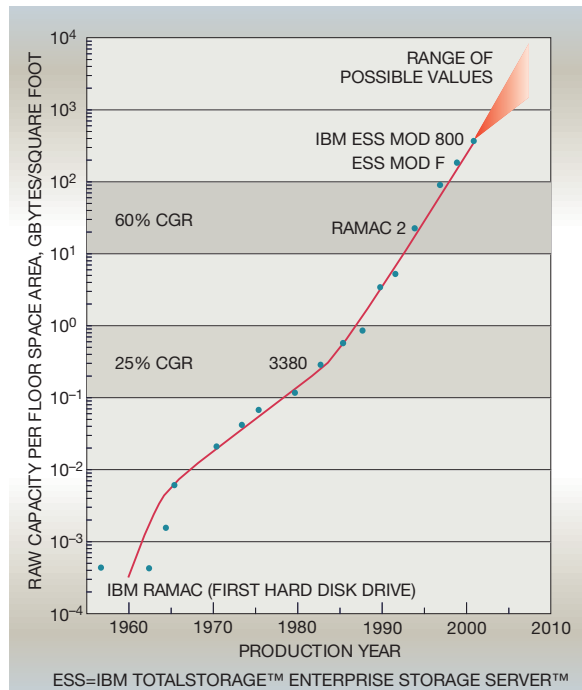


observed since the first disk drive. Although the CGR is expected to eventually decline, based on increased processing difficulties in magnetic head and disk media production, laboratory demonstrations of advanced head designs indicate that a head with a capacity of greater than 100 gigabits/in² is feasible within the near future, representing an increase of two or three times, compared with the areal density of today's production disk drives.^{3,4}

Although there is no direct analogue to areal density at the storage system level, the trend in floor space utilization in terabytes/ft² closely approximates that of areal density at the drive level. Figure 2 illustrates this trend for storage systems and nearly parallels the increases indicated by Figure 1 in disk drive areal density. The enhanced storage capacity per square foot of floor space is the direct result of disk drive increases in areal density and, as this trend continues into the future, it is expected that nearly 10⁴ gigabytes/ft² will be attained before the year 2010.

An alternate method of considering this trend is shown in Figure 3, where the floor space required

Figure 2 Storage floor space utilization trend — IBM storage systems



to store one terabyte of information over the past 45 years is shown to have decreased by more than a factor of 10⁷, similar to increases in server disk drive areal density.

As a further example of the relationship between systems containing multiple disk drives and the evolutionary trends in the drives themselves, the volumetric density of disk drives is shown in Figure 4. This parameter combines disk areal density with the packing efficiency of disks within the drive's frame as well as the packaging of the spindle motor, actuator motor, and electronics. Normally, drives with a smaller form factor (FF) exhibit the highest volumetric density, due to more efficient packing. A marked change in the slope of the volumetric density curve indicates this effect in larger-form-factor (14-inch and 10.8-inch) drives to 3.5-inch and smaller drives. Areal density also influences this slope change.

At the system level, packing large numbers of drives in close proximity combined with high volumetric density of the drives themselves results in the trend shown in Figure 5. The use of smaller-form-factor

Figure 3 Floor space required to store 1 terabyte

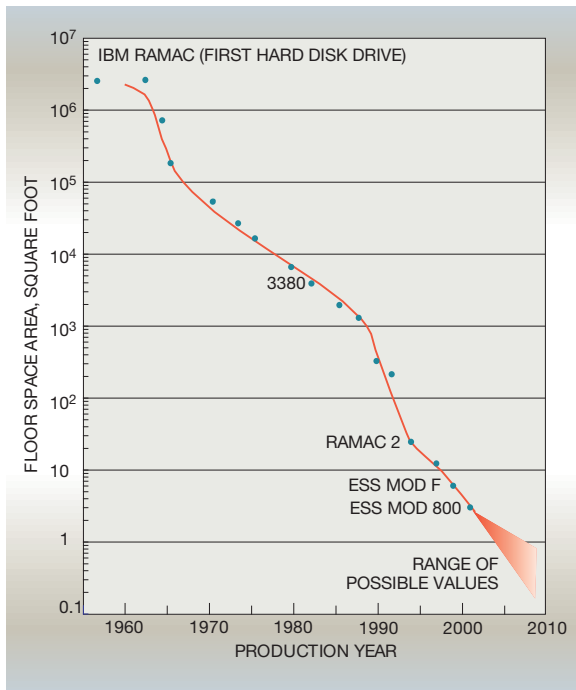
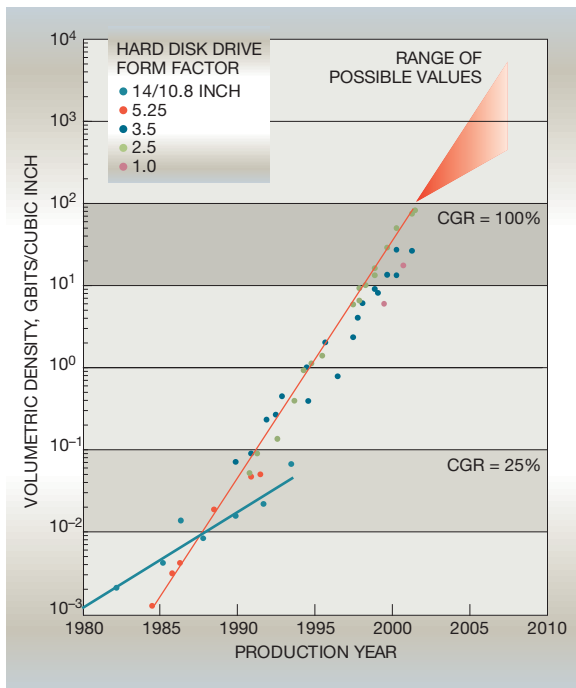


Figure 4 Hard disk drive volumetric density trend



disk drives, in this case drives with a 3.5-inch form factor, directly resulted in the design of storage systems with higher volumetric densities as well as a steeper growth trend.

Cost of storage

The price trends for storage systems would be expected to be dependent on drive costs and to a lesser extent on unit costs of DRAM caches and buffers, other electronics such as controllers, the cooling system, and various types of cabling. On the basis of cost reductions for the latter components, only a modest reduction in system storage costs could be expected. However, as shown in Figure 6, a dramatic drop in cost-per-megabyte is evident, with reduced drive costs as its dominant cause. Since areal density increases within the last 10 to 15 years were 60 percent to 100 percent per year, it would be expected that price declines would average 37 percent to 50 percent, respectively per year, as:

$$\left[\frac{\text{unit price}}{\text{capacity}} \right]_n [1 + r] = \left[\frac{\text{unit price}}{\text{capacity}} \right]_{n+1}$$

where r is the CGR for price-per-capacity change from time period n to time period $n + 1$. The trends depicted in Figure 6 show the 37 percent and 50 percent price declines corresponding to the areal density CGR increases shown in Figure 1. Over the last 15 years, unit price has not changed to the extent that capacity has, for a given drive configuration.

Figure 6 also shows the system-level rate of price decline and indicates that the adoption of small-form-factor drives has significantly increased the rate of price decline, nearly approximating the downward trend for disk drives. Thus, disk drive prices have had a significant influence on system prices. It is projected that any future reduced rate of areal density increase would also slow the rate of price decline for storage systems.

Figure 7 shows the rate of price-per-storage-capacity decline for DRAM (and flash) compared with a similar decline for disk drives.² The significant price decrease in semiconductor memory storage has also influenced the system price trend to a certain extent, as can be seen in the declining price-per-storage-capacity trend. However, it is the disk drive price improvement that is regarded as the principal factor in the trends noted in Figure 6. It can also be observed from Figure 7 that the cost of disk drive stor-

Figure 5 Storage system volumetric density trend

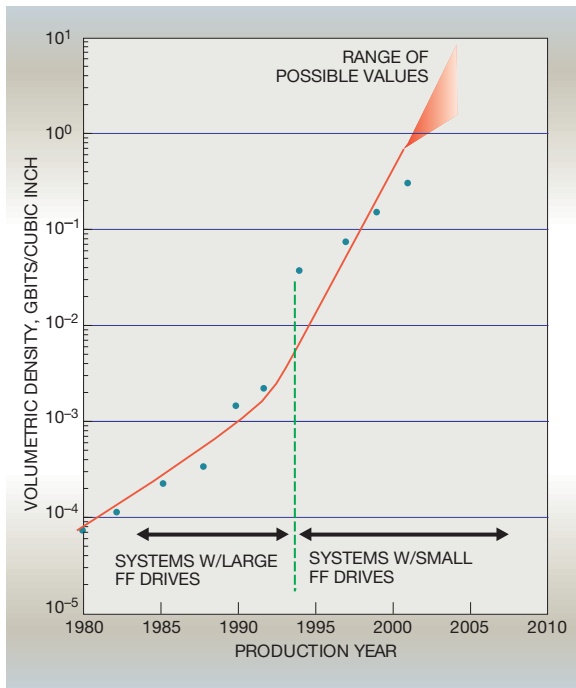


Figure 6 Cost of storage at the disk drive and system level

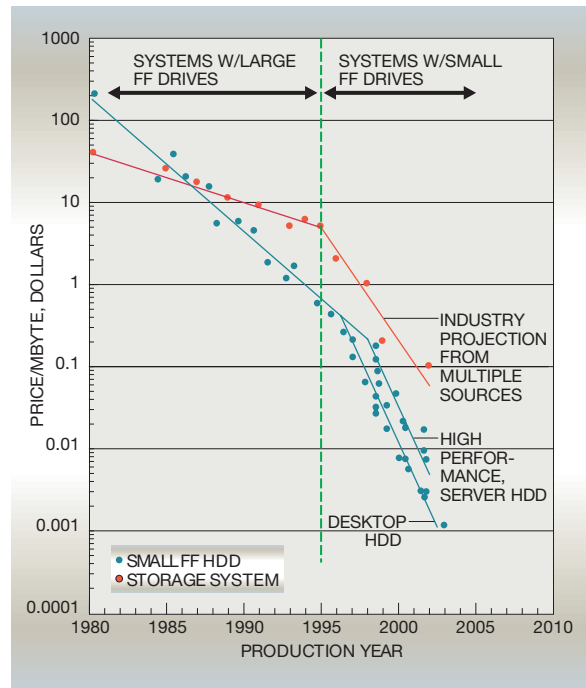


Figure 7 Cost of storage for disk drive, paper, film, and semiconductor memory

age had significantly undercut film and paper costs by the late 1990s. It is obvious that this is a principal advantage to the growth and popularity of on-line (i.e., nonarchival) magnetic storage.

Concurrently, there has been an accompanying increase in the quantity of DRAM found in most storage subsystems. When first introduced in the mid-1980s, a system cache of 8 megabytes was considered large. Today, a system cache of 8 gigabytes is considered relatively small. The corresponding decrease in price-per-storage capacity of DRAM, although prices started at a higher level than magnetic hard disk drives, has enabled larger caches while still maintaining a price reduction per capacity at the system level, as Figure 6 indicates.

Form factor miniaturization

Figure 8 displays form-factor miniaturization for disk drives for the time period 1956–2002. Large form-factor drives, with 14- and 10.8-inch nominal sizes, were replaced by 3.5-inch form-factor drives at about 1995, and it is possible that in the future even smaller-

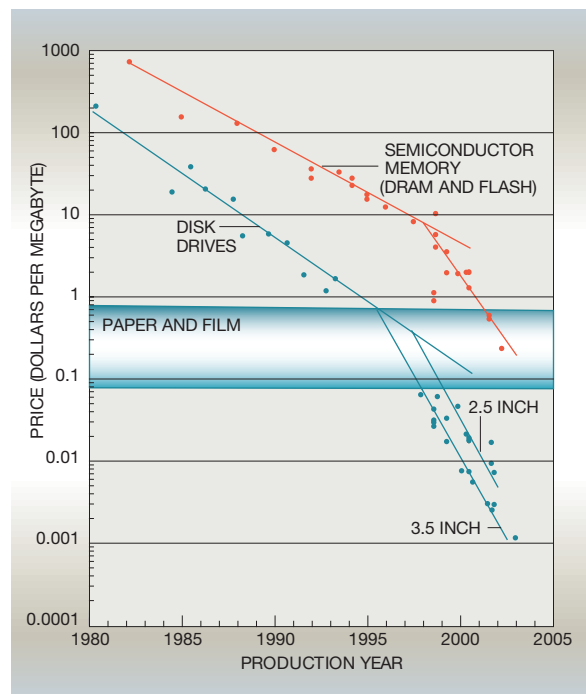
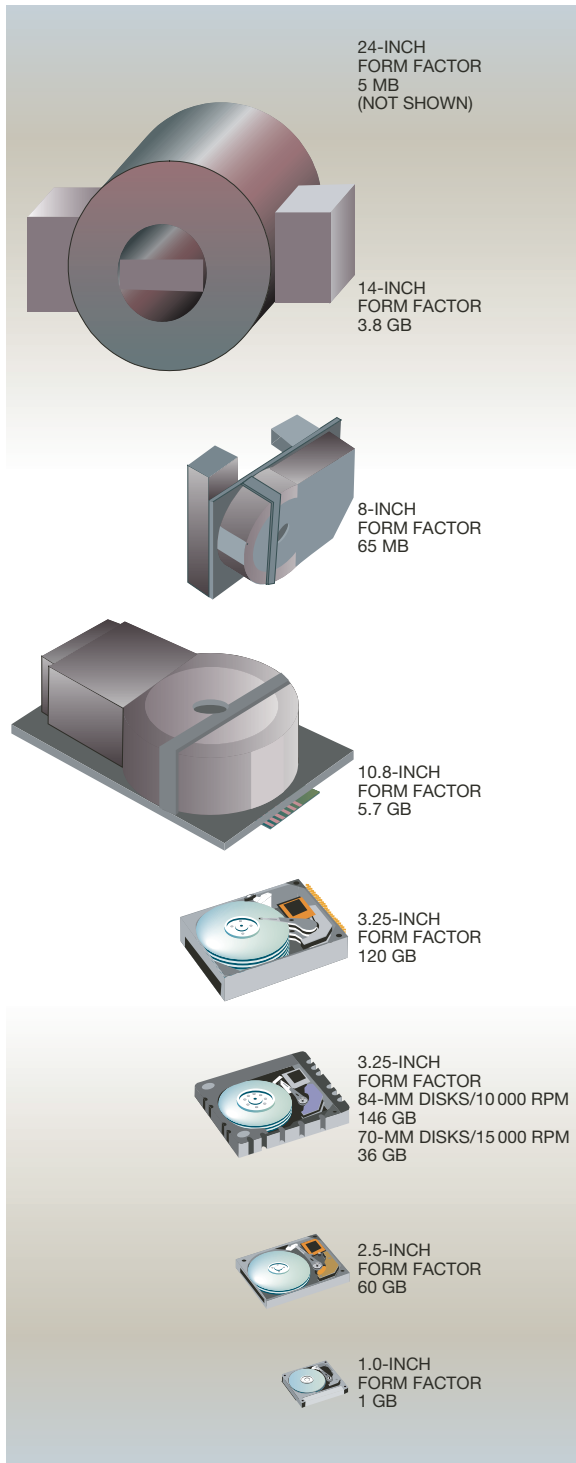


Figure 8 The evolution of disk drive form factors 1956–2002



form-factor drives, such as 2.5-inch drives, will constitute the storage devices used for large storage arrays.⁴ In the past, only eight large-form-factor disk drives could be contained within a storage system frame of reasonable size, whereas today's systems may contain as many as 256 drives (or more) in a RAID (redundant array of independent disks) or equivalent configuration. The availability of small-form-factor drives has been directly responsible for the development and use of system architectures such as RAID, which have advanced storage to new levels of low cost, capacity, performance, and reliability. In addition to the increases in reliability brought about by these enhanced drive technologies, the adoption of RAID architectures has resulted in even higher levels of improvement in system availability. RAID architectures enable the effective uncoupling of hardware reliability from system availability.

The miniaturization trend in disk drives, simultaneous with the trend of increasing drive capacity, is the direct result of the vast technology improvements that have increased areal density over seven orders of magnitude since 1956. These improvements have principally been: the introduction of magnetoresistive (MR) and giant magnetoresistive (GMR) read heads, finer-line-width inductive write head elements, high-signal-amplitude thin film disk materials, lower flying heights, and partial-response-maximum-likelihood (PRML) data channels. These and many more innovations demonstrate the continuing trend toward further enhancements.⁵

Power/performance

Smaller-form-factor drives containing smaller-diameter disks with high areal density permit faster disk rotation rates while maintaining moderate power requirements. Figure 9 shows this trend for disk drives and storage systems, and indicates that both have become significantly more economical with respect to power utilization. The power loss due to air shear for a disk drive is given by

$$P = \text{constant} \times D^{4.6} \times R^{2.8}$$

where D is disk diameter and R is the rotation rate. Reducing disk diameter from 14 inches (355 mm) to 65 mm would allow the drive designer to increase rotation rate from 3600 RPM (revolution per minute) to 15000 RPM in today's higher performance disk drives, while maintaining acceptable power losses

and even contributing significantly to a declining power per capacity trend. Figure 9 indicates that the power per gigabyte for drives is consistently dropping, whereas RPM is increasing and disk diameter is decreasing, in an almost monotonic trend. System power requirements follow a similar trend of reduction. Although system electronics contribute to power requirements, disk drives would again be a major factor, and continuous increases in areal density at the drive level and volumetric density at the system level would be expected to result in further declines in power requirements per storage capacity. This “green effect” is expected to continue with future drives and storage systems. Without the continuous increases in areal density, it would not be possible to reduce disk diameter and increase drive (and system) capacity simultaneously.

The trend in maximum internal data rate for disk drives is shown in Figure 10. This parameter is proportional to linear density, rotation rate, and disk diameter. In 1991, a significant change occurred in drive design and operation, primarily the adoption of MR heads that allowed higher linear density, smaller-form-factor drives, and higher disk rotation rates. Note that at a little over 100 megabytes/s internal clock operation within a data channel circuit is approaching 1 GHz—as fast as today’s microprocessors and also approaching the limit for silicon circuits.

Figure 11 shows the decrease in average seek and accessing times for server-platform disk drives over the past 30 years. Access time is defined as seek time plus latency time, which is inversely proportional to rotation rate. Seek time depends on data band, the difference between outer-disk recording radius and inner-disk recording radius (the region where data are actually recorded). The use of smaller-diameter drives causes a marked change in slope near 1991, similar to trends shown in Figure 10.⁶ The smaller drives are faster, store more information, and consume less power per capacity. Rotation rate continues to be a key parameter in disk drive and system performance, a fact that sometimes is not apparent. Regardless of the inclusion of large DRAM caches, when the server system requests a read of data not contained within the cache, latency becomes the key delaying element. It is the application of small-form-factor drives that has allowed rotation rates to go beyond 10000 RPM and to continually improve system performance while maintaining reasonable power requirements.

Figure 9 Hard disk drive and storage system power trend per Gbytes stored

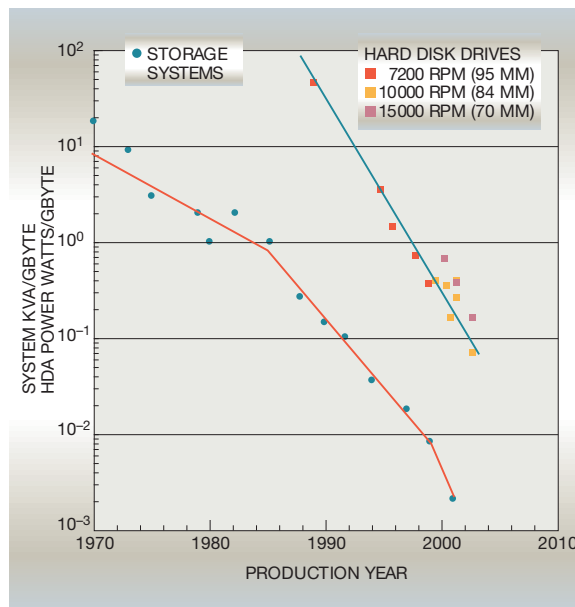
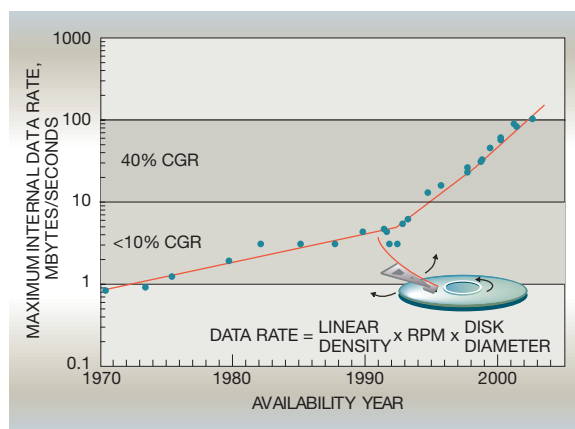


Figure 10 Hard disk drive maximum internal data rate for enterprise/server drives



What impact does drive design have on system performance? Initially, system performance was essentially drive performance. After the introduction of electronic caches, system performance was related to system channel speed, and cache effectiveness became the critical issue. This is shown by

Figure 11 Disk drive access/seek times

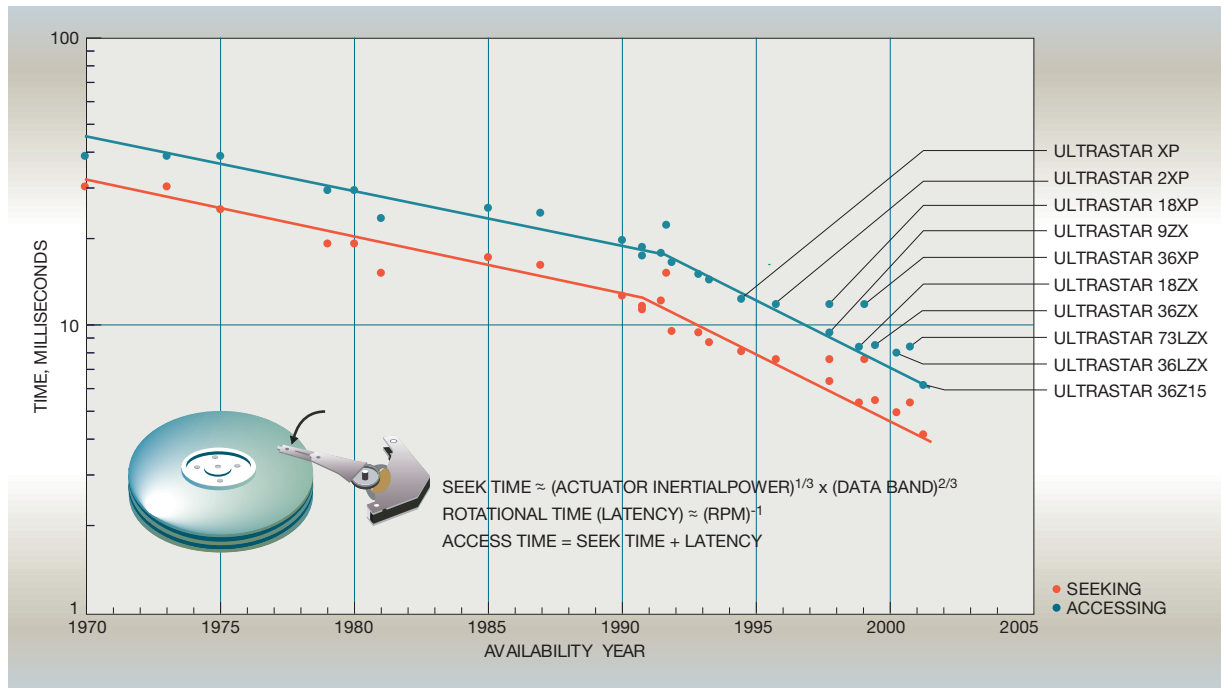
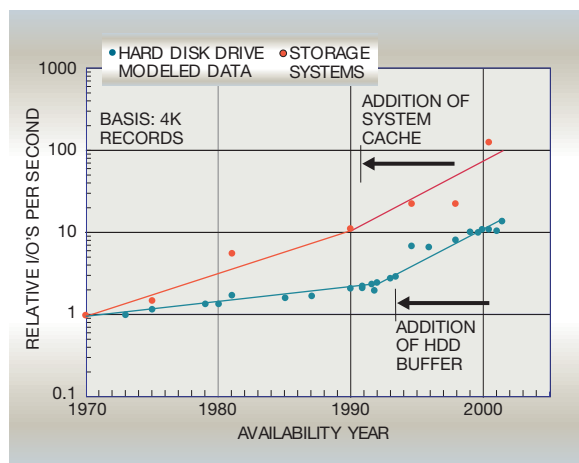


Figure 12 Performance trend for hard disk drives and storage systems



Overall system performance
 = (cache hit ratio × channel performance)
 + (1 - cache hit ratio) × drive performance

System performance is a weighted average of the two performances based on the cache hit ratio. Nearly concurrent with this trend, a disk drive buffer was also included, which enhanced the performance of this component significantly. Figure 12 shows relative performance of both drives and systems over time, and the effect of caching and buffering is clearly visible. The basis for the chart is the storing of 4K records. Although system (and drive) performance would be expected to increase as a result of component technology advances and drive miniaturization, enhancements in electronics have been a significant factor.

The CGR slopes of the relative performance (input/output operations per second) curves in Figure 12 after 1990 are about 25 percent, less than the rate of disk capacity increase based on areal density. There is a hypothetical concern that performance per capacity could decrease to a level where overall system performance would be adversely affected. Presently, this performance trend constitutes essentially no limitation to future disk drive capacity increase, although new design modifications to enhance disk drive and system performance could

be required. One technique to accomplish this would be the addition of dual actuators within the drive to provide concurrent seeks on multiple disk surfaces, as well as within a single disk surface. This design modification, although adding to the drive's cost, could reduce the effective seek time, and in the latter case, a considerable reduction in latency could also be realized. Alternatively, continued miniaturization could result in a reduction of total disk capacity with increasing areal density, a trend which is ongoing today.

Conclusions

The magnetic hard disk drive has been shown to have evolved through the past four or five decades by incorporating successive innovations that have increased storage capacity, performance, and availability, while concurrently allowing a miniaturization in form factor that has also reduced power requirements, particularly in large arrays of drives. The very nature of the disk drive has allowed the design and creation of these large arrays in RAID systems, which today are the cornerstones of the storage and server industry. A major part of the system characteristics have been shown to be directly related to disk drive properties. As disk drive improvements have been introduced, system specifications have reflected these changes. The disk drive is a very significant component in the storage hierarchy, and the future of this device continues to be bright.

The probability of continual increases in areal density is very high, although the rapid rates of advancement must surely slow from today's 100 percent CGR. The result will be a continuous increase in system capacity and performance, particularly with the inclusion of 15 000 RPM drives, a trend that has started today. The ready availability of high-capacity, low-cost storage systems has fueled the application of both SAN (storage area network) and NAS (network-attached storage) architectures and made the Internet a high growth area with almost universal acceptance. Past disk drive areal density limits projected for magnetic recording have all been proven wrong, and capacity and performance could easily continue to increase well into the future.⁵ Although the concept of a maximum capacity per actuator has also been discussed as a limit, miniaturization would be a likely direction for future disk drives with very high areal densities, to allow this capacity per head to level out. The architecture of future storage systems that contain very large numbers of miniature disk drives offers significant design challenges in physical lay-

out, interconnections, cabling, cooling, and finally the electronics and software to fully utilize these drives. The issue of system reliability, as it applies to large numbers of these drives, is an area to be addressed in future investigations.

Although there are alternatives that could replace the present magnetic hard disk drive technology in on-line storage devices in future systems, such as optical/DVD (digital video disk) enhancements and holographic storage products, all are considered years away from complete development and implementation. Each offers some promise of new storage improvements but may lack the wide range of applications enjoyed by today's magnetic disk drive technology.

A major transition from rotating disk to semiconductor flash or DRAM-based storage in storage systems is not likely, since pricing for the latter technology remains prohibitive, as shown in Figure 7. MRAM, in which a magnetic element such as GMR functions as a nonvolatile interconnection between word line/bit line nodes, is also expected to follow this restrictive pricing trend. Future storage systems throughout this decade will, in all likelihood, be based on magnetic disk drives. These will have a small form factor of 2.5 inches or smaller, with high RPM, improved access times, and, of course, lower cost per storage capability.

Cited references

1. M. T. Geenen and J. Kim, "TrendFOCUS 2002," TrendFOCUS (March 2002). See <http://www.trendfocus.com> (subscriber only).
2. E. Grochowski, IBM Storage Technology Web site, 2001. See <http://www.hgst.com/hdd/technolo/grochows/grocho01.htm>.
3. D. A. Thompson and J. S. Best, "The Future of Magnetic Data Storage Technology," *IBM Journal of Research and Development* **44**, No. 3, 311–321 (May 2000).
4. E. Grochowski and R. F. Hoyt, "Future Trends in Hard Disk Drives," *IEEE Transactions on Magnetism* **32**, No. 3, 1850–1854 (1996).
5. E. Grochowski and D. A. Thompson, "Outlook for Maintaining Areal Density Growth in Magnetic Recording," *IEEE Transactions on Magnetism* **30**, No. 6, 3797–3800 (1994).
6. Y. Moribe, "DK711S Disk Drive with Faster Latency," *Hitaichi Review* **37**, No. 5, 283–290 (1988).

General references

- J. M. Harker, D. W. Brede, R. E. Pattison, G. R. Santana, and L. G. Taft, "A Quarter Century of Disk File Innovation," *IBM Journal of Research and Development* **25**, No. 5, 673–689 (1981).
- L. D. Stevens, "The Evolution of Magnetic Storage," *IBM Journal of Research and Development* **25**, No. 5, 663–673 (1981).

Accepted for publication December 18, 2002.

Edward Grochowski *Hitachi Global Storage Technologies, Almaden Research Center, 650 Harry Road, San Jose, California 95120 (grchowski@almaden.ibm.com).* Dr. Grochowski began his career with IBM helping to develop IBM's microelectronic silicon activity and later joined the Almaden Research Center in San Jose, where he was program manager of storage devices. His interests include disk drive component design as well as drive form factor development, both of which have contributed to IBM's latest storage products. Dr. Grochowski holds nine patents. He received a Ph.D. degree from New York University (1971), held the position of Adjunct Professor of Chemical Engineering there, and was also associated with the University of Michigan Semiconductor Research Lab. In addition to being a member of the board of directors of IDEMA, the International Disk Drive Equipment and Materials Association, Dr. Grochowski serves as chairman of the technical program committee for DISKCON Asia Pacific and DISKCON USA, leading storage conferences in the industry. He is also a member of the Institute of Electrical and Electronics Engineers. In 2003, following 41 years with IBM, Dr. Grochowski joined Hitachi Global Storage Technologies at the Almaden Research Center.

Robert D. Halem *IBM Systems Group, San Jose Executive Briefing Center, 5600 Cottle Road, San Jose, California 95193 (sthalem@us.ibm.com).* Mr. Halem is a senior storage consultant on the staff of the IBM TotalStorage™ Executive Briefing Center. He received a B.A. degree in mathematics from UCLA in 1965 and an M.S. degree in computer design from San Jose State University in 1975. He joined IBM in 1965 and has held a variety of positions in hardware and software development and in technical marketing. In 1969, he received an IBM Outstanding Contribution Award for his work on special processor development.