AVOIDING A DATA CRUNCH

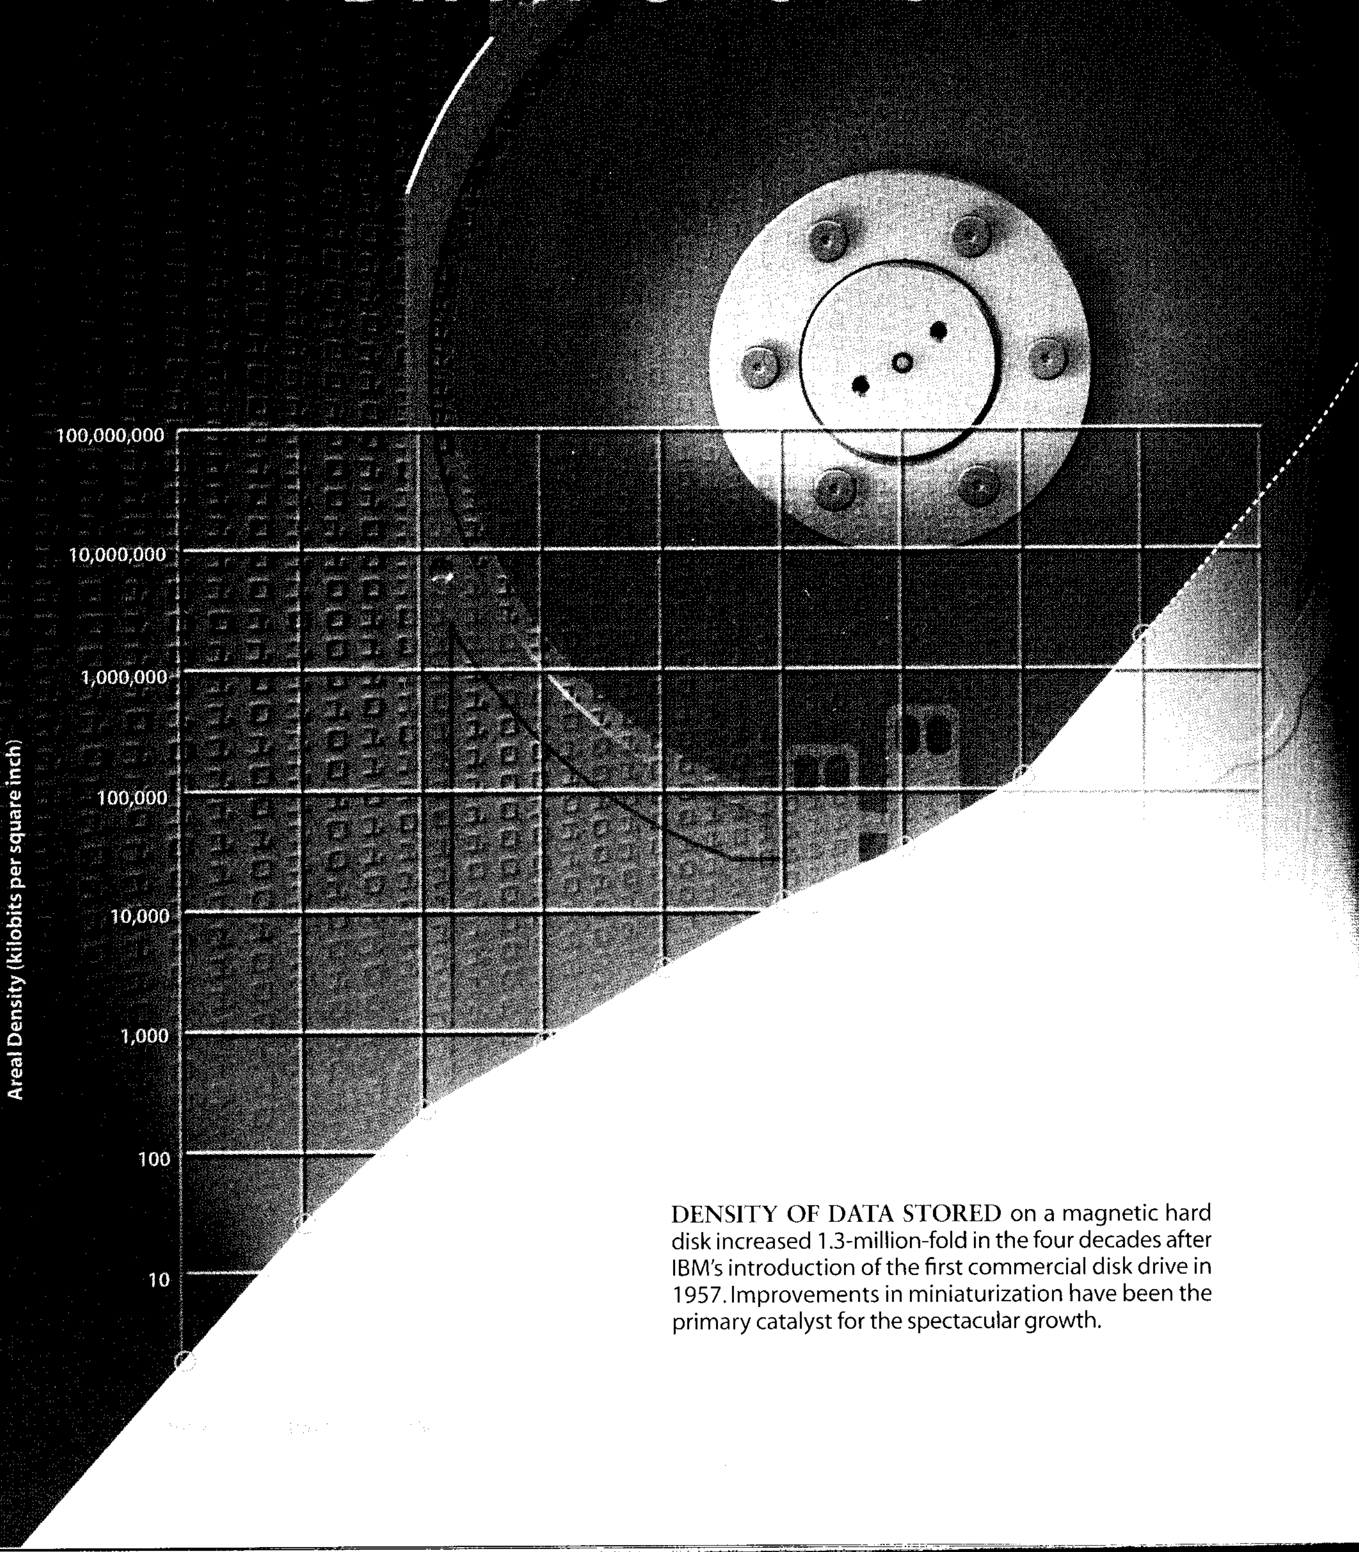

DENSITY OF DATA STORED on a magnetic hard disk increased 1.3-million-fold in the four decades after IBM's introduction of the first commercial disk drive in 1957. Improvements in miniaturization have been the primary catalyst for the spectacular growth.

The technology of computer hard drives is fast approaching a physical barrier imposed by the superparamagnetic effect. Overcoming it will require tricky innovations

by Jon William Toigo

Many corporations find that the volume of data generated by their computers doubles every year. Gargantuan databases containing more than a terabyte—that is, one trillion bytes—are becoming the norm as companies begin to keep more and more of their data on-line, stored on hard-disk drives, where the information can be accessed readily. The benefits of doing so are numerous: with the right software tools to retrieve and analyze the data, companies can quickly identify market trends, provide better customer service, hone manufacturing processes, and so on. Meanwhile individual consumers are using modestly priced PCs to handle a data glut of their own, storing countless e-mails, household accounting spreadsheets, digitized photographs, and software games.

All this has been enabled by the availability of inexpensive, high-capacity magnetic hard-disk drives. Improvement in the technology has been nothing short of legendary: the capacity of hard-disk drives grew about 25 to 30 percent each year through the 1980s and accelerated to an average of 60 percent in the 1990s. By the end of last year the annual increase had reached 130 percent. Today disk capacities are doubling every nine months, fast outpacing advances in computer chips, which obey Moore's Law (doubling every 18 months).

At the same time, the cost of hard-disk drives has plummeted. Disk/Trend, a Mountain View, Calif.–based market research firm that tracks the industry, reports that the average price per megabyte for hard-disk drives plunged from $11.54 in 1988 to $0.04 in 1998, and the estimate for last year is $0.02. James N. Porter, president of Disk/Trend, predicts that by 2002 the price will have fallen to $0.003 per megabyte.

Not surprisingly, this remarkable combination of rising capacity and declining price has resulted in a thriving market. The industry shipped 145 million hard-disk drives in 1998 and nearly 170 million last year. That number is expected to surge to about 250 million in 2002, representing revenues of $50 billion, according to Disk/Trend projections.

But whether the industry can maintain these fantastic economics is highly questionable. In the coming years the technology could reach a limit imposed by the superparamagnetic effect, or SPE. Simply described, SPE is a physical phenomenon that occurs in data storage when the energy that holds the magnetic spin in the atoms making up a bit (either a 0 or 1) becomes comparable to the ambient thermal energy. When that happens, bits become subject to random "flipping" between 0's and 1's, corrupting the information they represent.

In the quest to deliver hard disks with

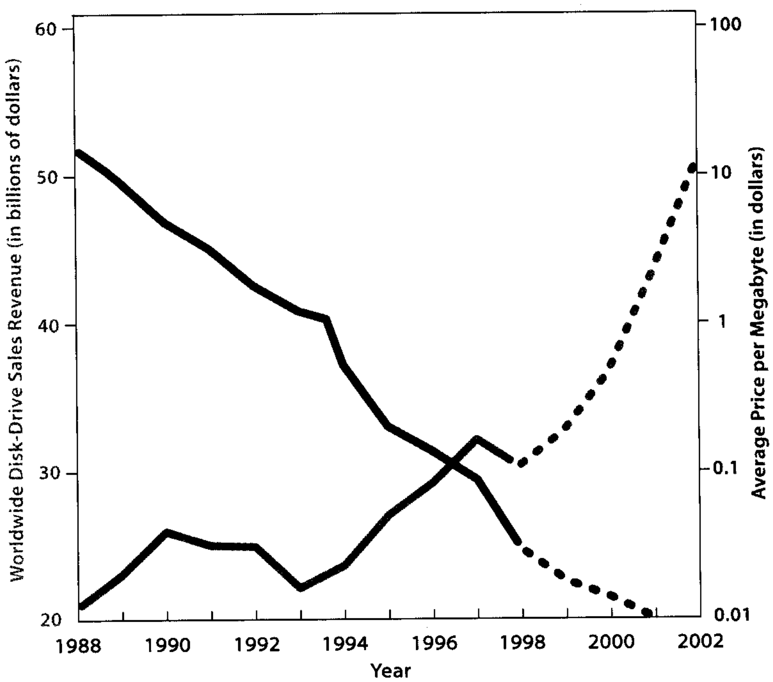

SALES OF HARD-DISK DRIVES have soared as costs per megabyte have plummeted. Sales revenues are expected to grow to $50 billion in 2002.

BRYAN CHRISTIE; SOURCE: DISK/TREND

ever increasing capacities, IBM, Seagate Technology, Quantum Corporation and other manufacturers have continually crammed smaller and smaller bits together, which has made the data more susceptible to SPE. With the current pace of miniaturization, some experts believe the industry could hit the SPE wall as early as 2005. But researchers have been busy devising various strategies for avoiding the SPE barrier. Implementing them in a marketplace characterized by fierce competition, frequent price wars and cost-conscious consumers will take a Herculean feat of engineering.

Magnetic Marvels

The hard-disk drive is a wonder of modern technology, consisting of a stack of disk platters, each one an aluminum alloy or glass substrate coated with a magnetic material and protective layers. Read-write heads, typically located on both sides of each platter, record and retrieve data from circumferential tracks on the magnetic medium. Servomechanical actuator arms position the heads precisely above the tracks, and a hydrodynamic air bearing is used to "fly" the heads above the surface at heights measured in fractions of microinches. A spindle motor rotates the stack at speeds of between 3,600 and 10,000 revolutions per minute.

This basic design traces its origins to the first hard-disk drive—the Random Access Method of Accounting and Control (RAMAC)—which IBM introduced in 1956. The RAMAC drive stored data on 50 aluminum platters, each of which was 24 inches in diameter and coated on both sides with magnetic iron oxide. (The coating was derived from the primer used to paint San Francisco's Golden Gate Bridge.) Capable of storing up to five million characters, RAMAC weighed nearly a ton and occupied the same floor space as two modern refrigerators.

In the more than four decades since then, various innovations have led to dramatic increases in storage capacity and equally amazing decreases in the physical dimensions of the drives themselves. Indeed, storage capacity has jumped multiple orders of magnitude during that time, with the result that some of today's desktop PCs have disk drives containing more than 70 gigabytes. Tom H. Porter, chief technology officer at California-based Seagate Technology's Minneapolis office, explains

HOW A HARD-DISK DRIVE WORKS

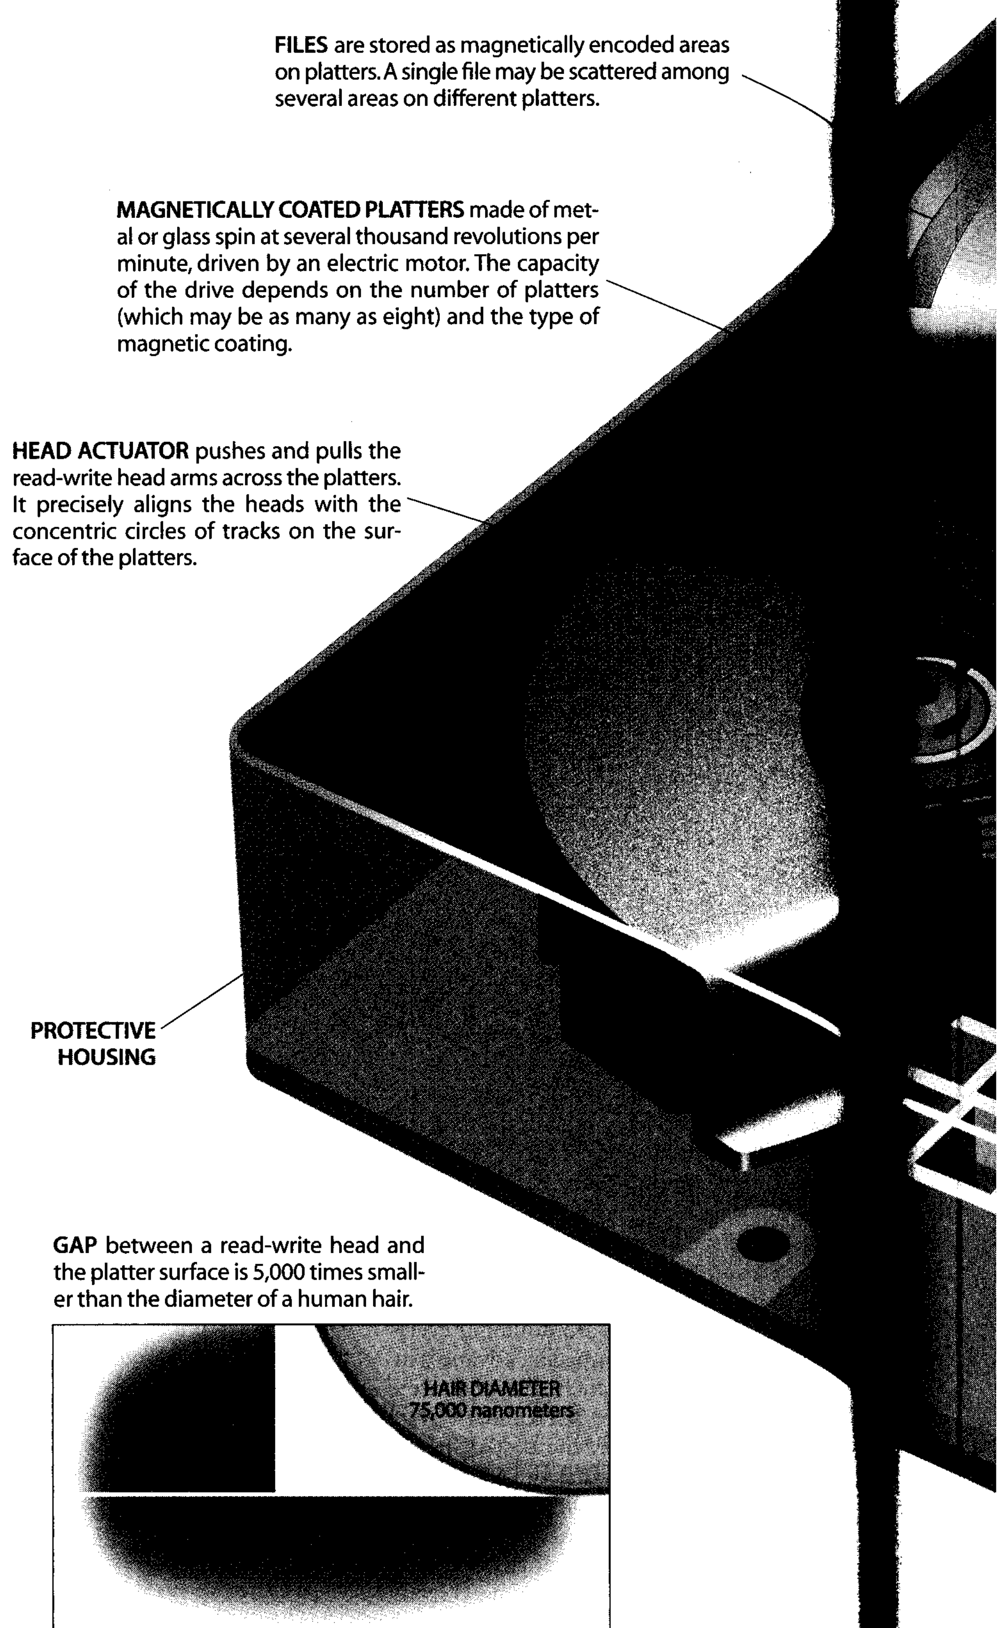

READ-WRITE HEADS, attached to the ends of moving arms, slide across both the top and bottom surfaces of the spinning platters. The heads write the data to the platters by aligning the magnetic fields of particles on the platters' surfaces; they read data by detecting the polarities of particles that have already been aligned.

PRINTED CIRCUIT BOARD receives commands from the drive's controller. The controller is managed by the operating system and the basic input-output system, low-level software that links the operating system to the hardware. The circuit board translates the commands into voltage fluctuations, which force the head actuator to move the read-write heads across the surfaces of the platters. The board also controls the spindle that turns the platters at a constant speed and tells the drive heads when to read from and when to write to the disk.

GEORGE RETSECK; SOURCE: *HOW COMPUTERS WORK*, BY RON WHITE. FOURTH EDITION. QUE CORPORATION, 1998 that the industry has achieved these improvements largely through straightforward miniaturization. "Smaller heads, thinner disks, smaller fly heights [the distance between head and platter]: everything has been about scaling," he notes.

Head Improvements

Many of the past improvements in disk-drive capacity have been a result of advances in the read-write head, which records data by altering the magnetic polarities of tiny areas, called domains (each domain representing one bit), in the storage medium. To retrieve that information, the head is positioned so that the magnetic states of the domains produce an electrical signal that can be interpreted as a string of 0's and 1's.

Early products used heads made of ferrite, but beginning in 1979 silicon chip-building technology enabled the precise fabrication of thin-film heads. This new type of head was able to read and write bits in smaller domains. In the early 1990s thin-film heads themselves were displaced with the introduction of a revolutionary technology from IBM. The innovation, based on the magnetoresistive effect (first observed by Lord Kelvin in 1857), led to a major breakthrough in storage density.

Rather than reading the varying magnetic field in a disk directly, a magnetoresistive head looks for minute changes in the electrical resistance of the overlying read element, which is influenced by that magnetic field. The greater sensitivity that results allows data-storing domains to be shrunk further. Although manufacturers continued to sell thin-film heads through 1996, magnetoresistive drives have come to dominate the market.

In 1997 IBM introduced another innovation—the giant magnetoresistive (GMR) head—in which magnetic and nonmagnetic materials are layered in the read head, roughly doubling or tripling its sensitivity. Layering materials with different quantum-mechanical properties enables developers to engineer a specific head with desired GMR capabilities. Currie Munce, director of storage systems and technology at the IBM Almaden Research Center in San Jose, Calif., says developments with this technology will enable disk drives to store data at a density exceeding 100 gigabits per square inch of platter space.

Interestingly, as recently as 1998 some experts thought that the SPE limit was 30 gigabits per square inch. Today no one seems to know for sure what the exact barrier is, but IBM's achievement has made some assert that the "density demon" lives somewhere past 150 gigabits per square inch.

A Bit about Bit Size

Of course, innovations in read-write heads would be meaningless if the disk platters could not store information more densely. To fit more data onto a disk, says Pat McGarrah, a director of strategic and technical marketing at Quantum Corporation in Milpitas, Calif., many companies are looking for media that will support shorter bits.

The problem, though, is SPE: as one shrinks the size of grains or crystals of magnetic material to make smaller bits, the grains can lose the ability to hold a magnetic field at a given temperature. "It really comes down to the thermal stability of the media," Munce explains. "You can make heads more sensitive, but you ultimately need to consider the properties of the media material, such as the coercivity, or magnetic stability, and how few grains you can use to obtain the desired resistance to thermal erasure."

ADDING OPTICAL TO MAGNETIC

One strategy for extending the life span of the workhorse magnetic-disk drive is to supplement it with optical technology. Such a hybrid approach could push storage densities to well beyond the current range of 10 to 30 gigabits per square inch. In fact, TeraStor in San Jose, Calif., claims that capacities could eventually top 200 gigabits per square inch, surpassing the anticipated limit imposed by the superparamagnetic effect [*see main article*].

The TeraStor disk drive is essentially a variation of magneto-optical technology, in which a laser heats a small spot on the disk so that information can then be written there magnetically. A crucial difference, however, is that TeraStor uses a solid-immersion lens, or SIL, which is a special type of truncated spherical lens.

Invented at Stanford University, SILs rely on the concept of liquid-immersion microscopy, in which both the lens and object being studied are placed in a liquid, typically oil, that greatly boosts the magnification. But SILs apply the technique in reverse

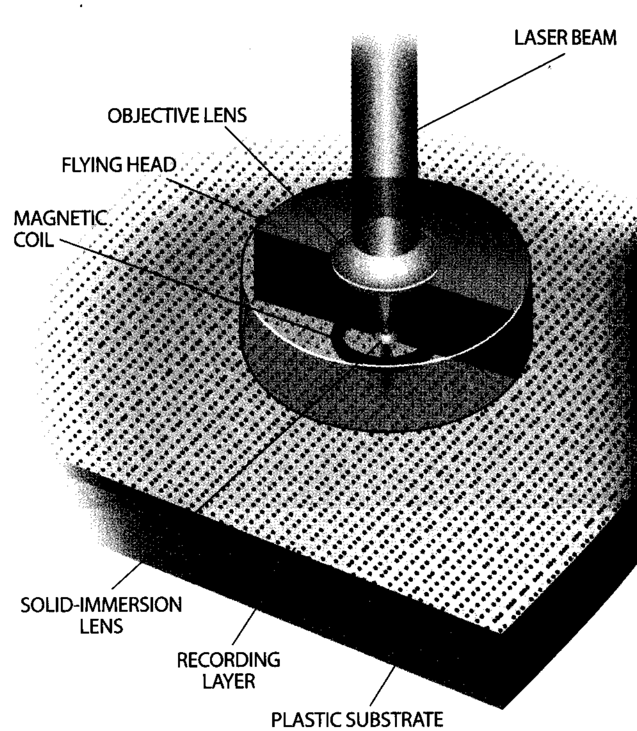

LASER BEAM heats a tiny spot on a disk, which permits the flying head to alter the magnetic properties of the spot so that it stores a binary 1 or 0. Two lenses focus the beam to an extremely fine point, enabling the bits to be written onto the disk at very high density. An objective lens concentrates the beam on a solid-immersion lens—the cornerstone of the system—which in turn focuses the light to a spot smaller than a micron across.

GEORGE RETSECK

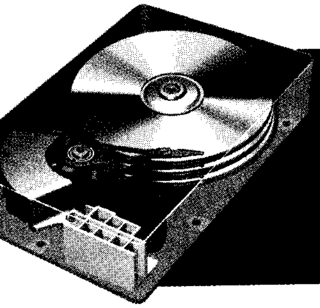

By 2002 the average price per megabyte for hard-disk drives will have fallen to $0.003, predicts James Porter, Disk/Trend

raditionally, Munce says, a mini- m of about 500 to 1,000 grains of gnetic material was required to store it. (In March, however, IBM scien- s announced a process for self- embling magnetic particles into bits t could provide areal densities as h as 150 gigabits per square inch.) rrently researchers are actively look- for improved materials that can ld a detectable magnetic charge and ist SPE with fewer grains. Also, the lustry has been developing better man- icturing processes to decrease the im- rities in the storage medium and there- enable smaller bits.

n lieu of improvements of this type, limit of bits per inch will remain in range of between 500,000 and 0,000, according to Karl A. Belser, a rage technologist for Seagate Tech- logy's research division. But this pa- meter, which is for data stored in a particular track on a platter, is only one determinant of areal density, which is the number of bits per square inch.

Tracking the Tracks

Storage capacity also depends on the narrowness of the tracks, and so far manufacturers have been able to cram up to 20,000 tracks per inch. This number is limited by various factors, such as the ability of the recording head to resolve the different tracks and the accuracy of its position-sensing system. Squeezing in additional tracks will require significant improvements in several areas, including the design of the head and the actuator that controls that head. To achieve an overall density of 100 gigabits per square inch, the industry must somehow figure out a way to fit about 150,000 tracks or more per inch.

With the existing technology, tracks must be separated by gaps of 90 to 100 nanometers, according to Belser. "Most write heads look like a horseshoe that extends across the width of a track," he explains. "They write in a longitudinal direction [that is, along the circular track], but they also generate fringe fields that extend radially." If the tracks are spaced too closely, this effect can cause information on adjacent tracks to be overwritten and lost.

One solution is to fabricate the recording head more precisely to smaller dimensions. "You can use a focused ion beam to trim the write head and to narrow the width of the track that a writer writes," Belser says. But the read head, which is a complex sandwich of elements, poses a harder manufacturing problem. Furthermore, for 150,000 tracks or more per inch to be squeezed in, the tracks will have to be less than about 170 nanometers wide. Such mi-

GEORGE RETSECK

to focus a laser beam on a spot with dimensions of less than a micron [*see illustration at left*]. The TeraStor technology is called "near field" because the read-write head must be extremely close to the storage medium (the separation is less than the wavelength of the laser beam).

The recording medium consists of a layer of magnetic material similar to that used in magneto-optical systems. But rather than being a magnetic layer encased in plastic, the recording layer is placed on top of a plastic substrate, which reduces the production cost and permits data to be written directly onto the recording surface.

As on a conventional magnetic disk, data bits (domains) are laid down one after the other. But the near-field bits are written standing up, or perpendicular to the plane of the disk, and not horizontally along the disk surface. "The magnetic fields of the domains poke out of the media vertically, rather than being laid out longitudinally," explains Gordon R. Knight, chief technology officer for TeraStor. "This configuration means that the magnetic fields of the bits support each other, unlike the fields of horizontally recorded bits, and are not subject to the superparamagnetic effect."

Furthermore, the ultrasmall domains are written in overlapping sequences, creating a series of crescent-shaped bits. This recording method effectively doubles the number of bits that can be written linearly in a track, thus enabling the TeraStor technology to achieve a higher storage capacity. Information is read by exploiting the so-called Kerr effect. A beam of light is bounced off a domain on the disk. Depending on whether the crystals in the domain have been magnetized to represent a 0 or a 1, they will polarize the reflected light in different directions.

The TeraStor technology has been under development for more than five years, and Knight acknowledges that the delivery of products has been delayed several times as the company works out various technical kinks. TeraStor did, however, demonstrate several prototypes at an industry trade show late last year, and the company has already lined up various manufacturing partners, including Maxell and Toso for the storage medium, Olympus for the optical components, Texas Instruments for the supporting electronic chips, and Mitsumi for the drive assembly. Industry heavyweight Quantum Corporation in Milpitas, Calif., which has a financial stake in TeraStor, has provided additional technology and access to its research lab.

If all goes well, TeraStor will be shipping products by the end of 2000. But the drives will contain just 20 gigabytes of storage on a removable CD-size medium. (Current hard drives already boast more than 70 gigabytes.) Knight asserts that the initial products may replace tape and optical storage products in applications in which access speed is important, such as digital video editing. He contends that the technology will ultimately make possible disk drives with much higher capacity—greater than 300 gigabytes—which may enable it to compete more directly with magnetic-disk drives. *—J.W.T.*

croscopically narrow tracks will be difficult for the heads to follow, and thus each head will need a secondary actuator for precise positioning. (In current products, just one actuator controls the entire assembly of heads.)

Last, smaller bits in thinner tracks will generate weaker signals. To separate those signals from background noise, researchers need to develop new algorithms that can retrieve the information accurately. Today's software requires a signal-to-noise ratio of at least 20 decibels. Says Belser, "The industry is at least six decibels short of being able to work with the signal-to-noise ratio that would apply when dealing with the bit sizes entailed in disks with areal densities of 100 to 150 gigabits per square inch."

Nevertheless, such problems are well understood, many industry experts concur. In fact, Munce asserts that the improvements in materials, fabrication techniques and signal processing already being studied at IBM and elsewhere will, over the next few years, enable the manufacture of disk drives with areal densities in the range of 100 to 150 gigabits per square inch.

USING "HARD" MATERIALS

n obvious way to cram more information onto a disk is to shrink the size of the data bits by using fewer or smaller grains of a magnetic material for each bit. The problem, though, is that the tiny bits begin to interfere with one another (think of what happens when two bar magnets are brought close to each other). To prevent such corruption of data, which is caused by the superparamagnetic effect, researchers have been investigating the use of certain rare-earth and transition elements that are very magnetically stable. In the industry lingo, such metals have high coercivity and are called "hard."

But a hard material is difficult to write on, so it may first be "softened" by being heated with a laser. This process lowers the coercivity of the grains so that data can then be written to them. As the material cools, it hardens, protecting the stored information from the vicissitudes of superparamagnetism. The concept sounds simple enough, but it has been difficult to implement: the laser beam must avoid accidentally heating adjacent bits that contain previously stored data.

To that end, Seagate Technology, headquartered in Scotts Valley, Calif., is using a disk that has grooves between the circular tracks of bits (much as a vinyl record does). The grooves block the laser heat from flowing to neighboring tracks. To record information in those narrow tracks, Seagate has been developing a new type of write head that is controlled by a special actuator. Details of these components are being kept under wraps.

The read head also presents certain difficulties. Because even the current experimental devices are three tracks wide instead of one, they have the potential to pick up unwanted noise during the reading process, according to Karl A. Belser, a storage technologist with Seagate. Even if a narrower head were developed, the device would need to be positioned precisely to follow the extremely thin tracks. Solutions include the use of a laser-positioning system, but that would add complexity—and cost—to the overall drive.

An alternative is to make the medium easier to read. This can be done by using a two-layer medium with a permanent storage layer positioned below a readout layer. To read data on the medium, the readout layer would be magnetically erased, and then the appropriate track of the storage layer would be heated with a laser to bring its data to the readout layer through a magnetic coupling process similar to current magneto-optical disk processes. Once the track had been written to the readout layer, its bits could be read in isolation from other tracks. Without the noise of adjacent tracks, even a wide head could read the information in the readout track.

If such a system were workable, the technology could store 1,000 gigabits per square inch, according to Seagate. In contrast, conventional wisdom holds that the superparamagnetic effect limits the storage density of traditional disk drives to a range of 100 to 150 gigabits per square inch. But even Seagate admits that it is at least four years from commercializing its thermally assisted magnetic drives. *—J.W.T.*

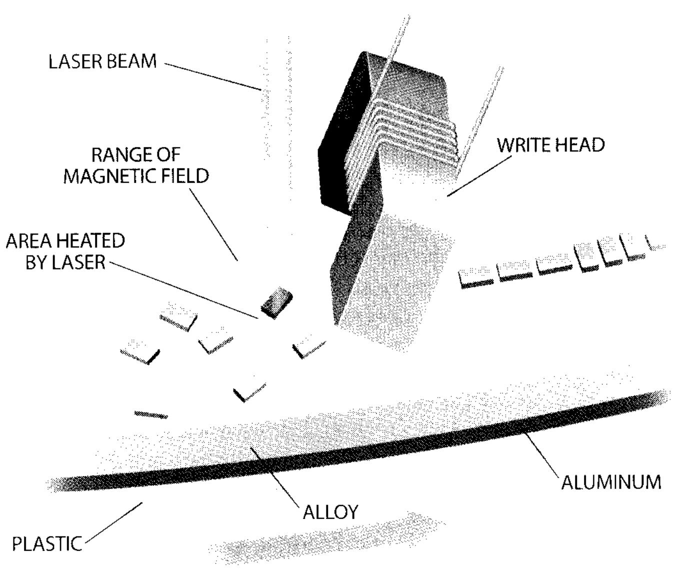

MAGNETO-OPTICAL DRIVE heats a spot with a laser, loosening magnetic crystals so that they can be reoriented with a magnetic field. This basic concept has been difficult to miniaturize because the laser must avoid accidentally heating—and thus possibly destroying—previously stored data. One solution is to manufacture a disk with grooves between concentric tracks of data to block heat from flowing between the tracks. To read the information, Seagate Technology is considering the use of a two-tier system in which the data are stored in tracks on a lower level. When the data are to be read out, a laser heats a section of a track in the lower layer. The heating induces magnetic coupling that transfers the data to the upper level of the disk, where they can be read out in the absence of interfering fields from adjacent tracks.

GEORGE RETSECK; SOURCE: *HOW COMPUTERS WORK*, BY RON WHITE. FOURTH EDITION. QUE CORPORATION, 1998

The introduction of thin-film heads took nearly 10 years. The transition from that to magnetoresistive technology required six more years because of various technical demands, including separate read and write elements for the head, a manufacturing technique called sputter deposition and different servo controls. "Going from thin-film inductive heads to MR heads entailed a number of new processes," Munce remarks. "Delays were bound to happen."

But the switch to giant magnetoresistive drives is occurring much faster, taking just between 12 and 18 months. In fact, IBM and Toshiba began shipping such products before the rest of the industry had fully converted to magnetoresistive heads.

The quick transition was possible because giant magnetoresistive heads have required relatively few changes in the surrounding disk-drive components. According to Munce, the progression to drive capacities of 100 gigabits per square inch will likewise be evolutionary, not revolutionary, requiring only incremental steps.

The Issue of Speed

But storage capacity is not the only issue. Indeed, the rate with which data can be accessed is becoming an important factor that may also determine the useful life span of magnetic disk-drive technology. Although the capacity of hard-disk drives is surging by 130 percent annually, access rates are increasing by a comparatively tame 40 percent.

To improve on this, manufacturers have been working to increase the rotational speed of drives. But as a disk spins more quickly, air turbulence and vibration can cause misregistration of the tracks—a problem that could be corrected by the addition of a secondary actuator for every head. Other possible enhancements include the use of fluid bearings in the motor to replace steel and ceramic ball bearings, which wear and emit noticeably audible noise when platters spin at speeds greater than 10,000 revolutions per minute.

Many industry onlookers foresee a possible bifurcation in the marketplace, with some disk drives optimized for capacity and others for speed. The former might be used for mass storage, such as for backing up a company's historical files. The latter would be necessary for applications such as customer service, in

PATTERNS OF BITS

Magnetic hard drives can store data at incredible densities—more than 10 gigabits per square inch of disk space. But as manufacturers cram information ever more tightly, the tiny bits begin to interfere with one another—the superparamagnetic effect [*see main article*]. One simple solution is to segregate the individual bits by erecting barriers between them. This approach, called patterned media, has been an ongoing area of research at most laboratories doing advanced work in storage technology.

One type of patterned media consists of "mesas" and "valleys" fabricated on the surface of a disk platter, with each mesa storing an individual bit. According to proponents of this approach, one bit of data (either a 0 or 1) could theoretically be stored in a single grain, or crystal, of a magnetic material. In contrast, conventional hard-disk technology requires a minimum of about 500 to 1,000 grains for each bit. Thus, with a grain size of seven to eight nanometers in diameter, this type of storage could achieve a density of more than 10,000 gigabits (or 10 terabits) per square inch.

To fabricate the mesas and valleys, companies have been investigating photolithographic processes used by the chip industry. "Electron beams or lasers would be needed to etch the pattern [onto the storage medium].

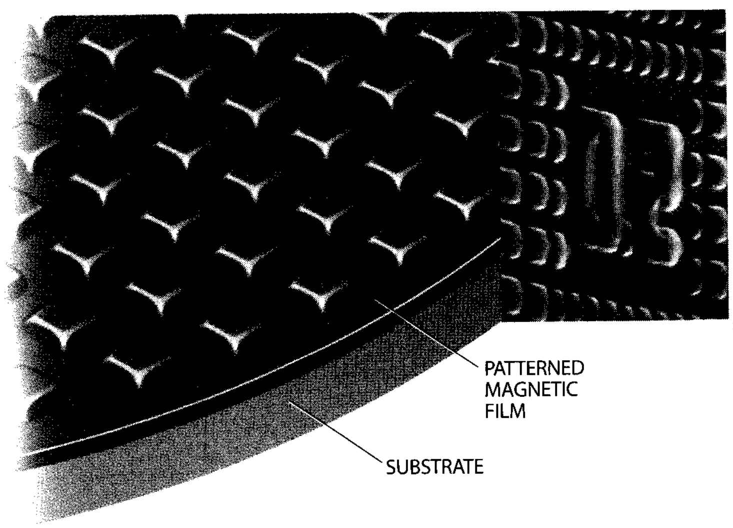

GEORGE RETSECK; IBM (*micrograph*)

"MESAS AND VALLEYS" of a future magnetic disk could help manufacturers avoid the superparamagnetic effect, in which closely packed bits in the magnetic media interfere with one another. In this patterned approach, the problem is circumvented by segregating each bit in its own mesa. The difficulty is in making the mesas small enough: they would have to be around eight nanometers across or smaller in order to achieve the kind of densities that developers are seeking. IBM has been able to build such structures with feature sizes as small as 0.1 and 0.2 micron (*inset*), or 100 and 200 nanometers.

Mesas would then need to be grown on a substrate layer, one bit in diameter," explains Gordon R. Knight, chief technology officer at TeraStor. But this technique needs much refinement. One estimate is that the current lithographic processes can at best make mesas that are about 80 nanometers in diameter—an order of magnitude too large for what is needed.

Even if the industry could obtain sufficiently tiny mesas and valleys, it would still need a revolutionary new type of head to read the data, says Currie Munce, director of storage systems and technology at the IBM Almaden Research Center in San Jose, Calif. According to Munce, various signal-to-noise issues would necessitate a radical departure from current magnetic-disk systems. Consequently, most experts agree that patterned-media technology will take years to become practical. —*J.W.T.*

which the fast retrieval of data is crucial.

In the past, customers typically preferred a bigger drive at the lowest possible cost, even if the product had slower performance. "In our hypercompetitive industry, drives with 30 to 40 percent higher performance sell for only about a 20 percent higher price," Munce notes.

But new applications are demanding faster drives. With electronic commerce over the World Wide Web, for example, companies need to store and retrieve customer data on the fly. In addition, businesses are deploying an increasing number of dedicated file servers for information that needs to be shared and accessed quickly by a number of employees.

The capacity-versus-performance debate could become acute as the industry considers various ways to avoid the SPE barrier. Experts agree that moving beyond areal densities of 150 gigabits per square inch will require a significant

ON THE HORIZON: HOLOGRAPHIC STORAGE

For nearly four decades, holographic memory has been the great white whale of technology research. Despite enormous expenditures, a complete, general-purpose system that could be sold commercially continues to elude industrial and academic researchers. Nevertheless, they continue to pursue the technology aggressively because of its staggering promise.

Theoretical projections suggest that it will eventually be possible to use holographic techniques to store trillions of bytes—an amount of information corresponding to the contents of millions of books—in a piece of crystalline material the size of a sugar cube or a standard CD platter. Moreover, holographic technologies permit retrieval of stored data at speeds not possible with magnetic methods. In short, no other storage technology under development can match holography's capacity and speed potential.

These facts have attracted name-brand players, including IBM, Rockwell, Lucent Technologies and Bayer Corporation. Working both independently and in some cases as part of research consortia organized and co-funded by the U.S. Defense Advanced Research Projects Agency (DARPA), the companies are striving to produce a practical commercial holographic storage system within a decade.

Since the mid-1990s, DARPA has contributed to two groups working on holographic memory technologies: the Holographic Data Storage System (HDSS) consortium and the PhotoRefractive Information Storage Materials (PRISM) consortium. Both bring together companies and academic researchers at such institutions as the California Institute of Technology, Stanford University, the University of Arizona and Carnegie Mellon University. Formed in 1995, HDSS was given a five-year mission to develop a practical holographic memory system, whereas PRISM, formed in 1994, was commissioned to produce advanced storage media for use in holographic memories by the end of this year.

With deadlines for the two projects looming, insiders report some significant recent advances. For example, late last year at Stanford, HDSS consortium members demonstrated a holographic memory from which data could be read out at a rate of a billion bits per second. At about the same time, an HDSS demonstration at Rockwell in Thousand Oaks, Calif., showed how a randomly chosen data element could be accessed in 100 microseconds or less, a figure the developers expect to reduce to tens of microseconds. That figure is superior by several orders of magnitude to the retrieval speed of magnetic-disk drives, which require milliseconds to access a randomly selected item of stored data. Such a fast access time is possible because the laser beams that are central to holographic technologies can be moved rapidly without inertia, unlike the actuators in a conventional disk drive.

Although the 1999 demonstrations differed significantly in terms of storage media and reading techniques, certain fundamental aspects underlie both demonstration systems. An important one is the storage and retrieval of entire pages of data at one time. These pages might contain thousands or even millions of bits.

Each of these pages of data is stored in the form of an optical-interference pattern within a photosensitive crystal or polymer material. The pages are written into the material, one after another, using two laser beams. One of them, known as the object or signal beam, is imprinted with the page of data to be stored when it shines through a liquid-crystal-like screen known as a spatial-light modulator. The screen displays the page of data as a pattern of clear and opaque squares that resembles a crossword puzzle.

A hologram of that page is created when the object beam meets the second beam, known as the reference beam, and the two beams interfere with each other inside the photosensitive recording material. Depending on what the recording material is made of, the optical-interference pattern is imprinted as the result of physical or chemical changes in the material. The pattern is imprinted throughout the material as variations in the refractive index, the light absorption properties or the thickness of the photosensitive material.

When this stored interference pattern is illuminated with either of the two original beams, it diffracts the light so as to reconstruct the other beam used to produce the pattern originally. Thus, illuminating the material with the reference beam re-creates the object beam, with its imprinted page of data. It is then a relatively simple matter to detect the data pattern with a solid-state camera chip, similar to those used in modern digital video cameras. The data from the chip are interpreted and forwarded to the computer as a stream of digital information.

Researchers put many different interference patterns, each corresponding to a different page of data, in the same material. They separate the pages either by varying the angle between the object and reference beams or by changing the laser wavelength.

Rockwell, which is interested in developing holographic memories for applications in defense and aerospace, optimized its demonstration system for fast data access, rather than for large storage capacities. Thus, its system utilized a unique, very high speed acousto-optical-positioning system to steer its laser through a lithium niobate crystal. By contrast, the demonstration at Stanford, including technologies contributed by IBM, Bayer and others, featured a high-capacity polymer disk medium about the size of a CD platter to store larger amounts of data. In addition, the Stanford system emphasized the use of components and materials that could be readily integrated into future commercial holographic storage products.

According to Hans Coufal, who manages IBM's participation in both HDSS and PRISM, the company's strategy is to make use of mass-produced components wherever possible. The lasers, Coufal points out, are similar to those that are found in CD players,

departure from conventional magnetic hard disks. Some of the alternatives boast impressive storage capabilities but mediocre speeds, which would limit their use for certain applications. At present, the main strategies include:

- Change the orientation of the bits on the disk from longitudinal (circumferential) to perpendicular, or vertical, to cram more of them together and to prevent them from flipping [*see box on page 64*].
- Use magnetic materials, such as alloys of iron/platinum or cobalt/samarium, that are more resistant to SPE. If the magnetic "hardness" of the material is a problem for recording data, heat the medium first to "soften" it magnetically before writing on it [*see box on page 66*].
- Imprint a pattern lithographically onto the storage medium to build microscopic barriers between individual bits [*see box on page 67*].
- Use a radically different storage material, such as holographic crystals [*see*

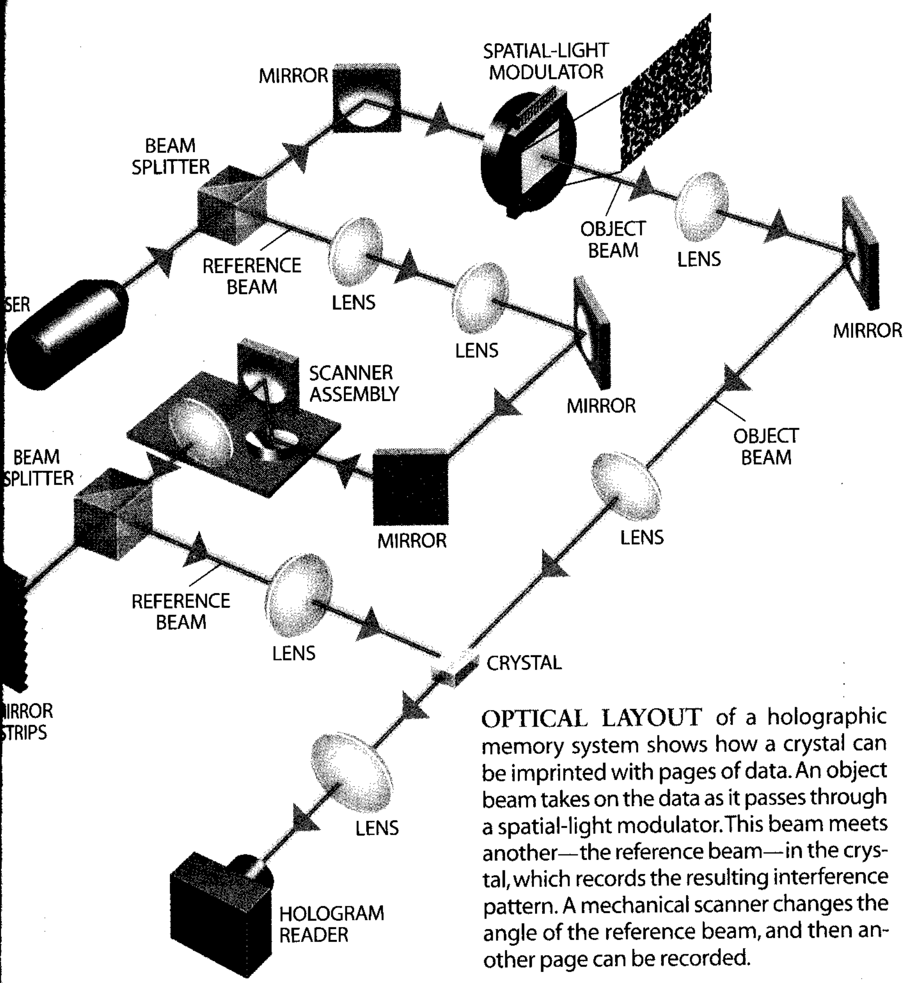

OPTICAL LAYOUT of a holographic memory system shows how a crystal can be imprinted with pages of data. An object beam takes on the data as it passes through a spatial-light modulator. This beam meets another—the reference beam—in the crystal, which records the resulting interference pattern. A mechanical scanner changes the angle of the reference beam, and then another page can be recorded.

and the spatial-light modulators resemble ordinary liquid-crystal displays.

Nevertheless, significant work remains before holographic memory can go commercial, Coufal says. He reports that the image of the data page on the camera chip must be as close to perfect as possible for holographic information storage and retrieval to work. Meeting the exacting requirements for aligning lasers, detectors and spatial-light modulators in a low-cost system presents a significant challenge.

Finding the right storage material is also a persistent challenge, according to Currie Munce, director of storage systems and technology at the IBM Almaden Research Center. IBM has worked with a variety of materials, including crystal cubes made of lithium niobate and other inorganic substances and photorefractive, photochromic and photochemical polymers, which are in development at Bayer and elsewhere. He notes that independent work by Lucent and by Imation Corporation in Oakdale, Minn., is also yielding promising media prospects. No materials that IBM has tested so far, however, have yielded the mix of performance, capacity and price that would support a mainstream commercial storage system.

Both Munce and Coufal say that IBM's long-standing interest in holographic storage intensified in the late 1990s as the associative retrieval properties of the medium became better understood. Coufal notes that past applications for holographic storage targeted the permanent storage of vast libraries of text, audio and video data in a small space. With the growing commercial interest in data mining—essentially, sifting through extremely large warehouses of data to find relationships or patterns that might guide corporate decision making and business process refinements—holographic memory's associative retrieval capabilities seem increasingly attractive.

After data are stored to a holographic medium, a single desired data page can be projected that will reconstruct all reference beams for similarly patterned data stored in the media. The intensity of each reference beam indicates the degree to which the corresponding stored data pattern matches the desired data page.

"Today we search for data on a disk by its sector address, not by the content of the data," Coufal explains. "We go to an address and bring information in and compare it with other patterns. With holographic storage, you could compare data optically without ever having to retrieve it. When searching large databases, you would be immediately directed to the best matches."

While the quest for the ideal storage medium continues, practical applications such as data mining increase the desirability of holographic memories. And with even one business opportunity clearly defined, the future of holographic storage systems is bright indeed. —*J.W.T.*

A DECADE AWAY: ATOMIC RESOLUTION STORAGE

Chuck Morehouse, director of Hewlett-Packard's Information Storage Technology Lab in Palo Alto, Calif., is quick to point out that atomic resolution storage (ARS) will probably never completely replace rotational magnetic storage. Existing hard-disk drives and drive arrays play well in desktops and data centers where device size is not a major issue. But what about the requirements for mass storage on a wristwatch or in a spacecraft, where form factor, mass and power consumption are overriding criteria?

The ARS program at Hewlett-Packard (HP) aims to provide a thumbnail-size device with storage densities greater than one terabit (1,000 gigabits) per square inch. The technology builds on advances in atomic probe microscopy, in which a probe tip as small as a single atom scans the surface of a material to produce images accurate within a few nanometers. Probe storage technology would employ an array of atom-size probe tips to read and write data to spots on the storage medium. A micromover would position the medium relative to the tips.

IBM and other companies are actively developing such probe storage technology, and Morehouse reports that the U.S. Department of Defense has a stake in the work. For example, the Defense Advanced Research Projects Agency (DARPA) is footing the bill for three HP researchers who are working on bringing a device from the test lab to the marketplace.

According to Morehouse, they face four primary challenges. First is the storage medium. The HP group has chosen one consisting of a material with two distinct physical states, or phases, that are stable at room temperature. One phase is amorphous, and the other is crystalline. Bits are set in this "phase-change medium" by heating data spots to change them from one phase to the other.

The second challenge is the probe tip, which must emit a well-directed beam of electrons when voltage is applied. A strong beam flowing from the tip to the medium heats a data spot as needed to write or erase a bit. A weak beam can be used to read data by detecting a spot's resistance or other phase-dependent electrical property. Optical reading techniques may also be possible. HP is looking at a "far-field" approach in which the tip is perhaps 1,000 nanometers from the medium, unlike most probe efforts in which the tip is in contact or almost in contact with the medium.

A third issue is the actuator or micromover that positions the media for reading and writing. HP is developing a micromotor with nanometer-level positioning capabilities.

The final step is packaging. Morehouse explains: "We need to get the ARS device together into a rugged package and develop the system electronics that will allow it to be integrated with other devices." An extra difficulty is that the working elements of the device will probably need to be in a vacuum or at least in a controlled atmosphere to reduce the scattering of electrons from the read-write beam and to reduce the flow of heat between data spots.

Morehouse sees the technology to create the ARS device becoming available within a decade but acknowledges that it may take considerably longer to bring the device to market. The magnetic-disk industry has a significant investment to protect, but he is confident that as applications demand the portability and performance that ARS offers, it will become a significant player in the storage market.

"My imagination for how it can be used is woefully inadequate," Morehouse says. Magnetic-disk drives are being scaled steadily downward in size—witness the example of IBM's Microdrive (340 megabytes in a one-inch form factor). Nevertheless, "ARS may be competitive for many applications," he notes. "One key advantage is low power consumption. When ARS is not being asked to perform an operation, it has no power consumption. Watchmakers may not want a Microdrive with a lot of batteries for a watch."

Morehouse says that the first ARS devices might have a one-gigabyte capacity but that capacities will increase over time: "The ultimate capacity will be determined by how small you can make a spot. Nobody knows the answer to that—100 atoms?" —*J.W.T.*

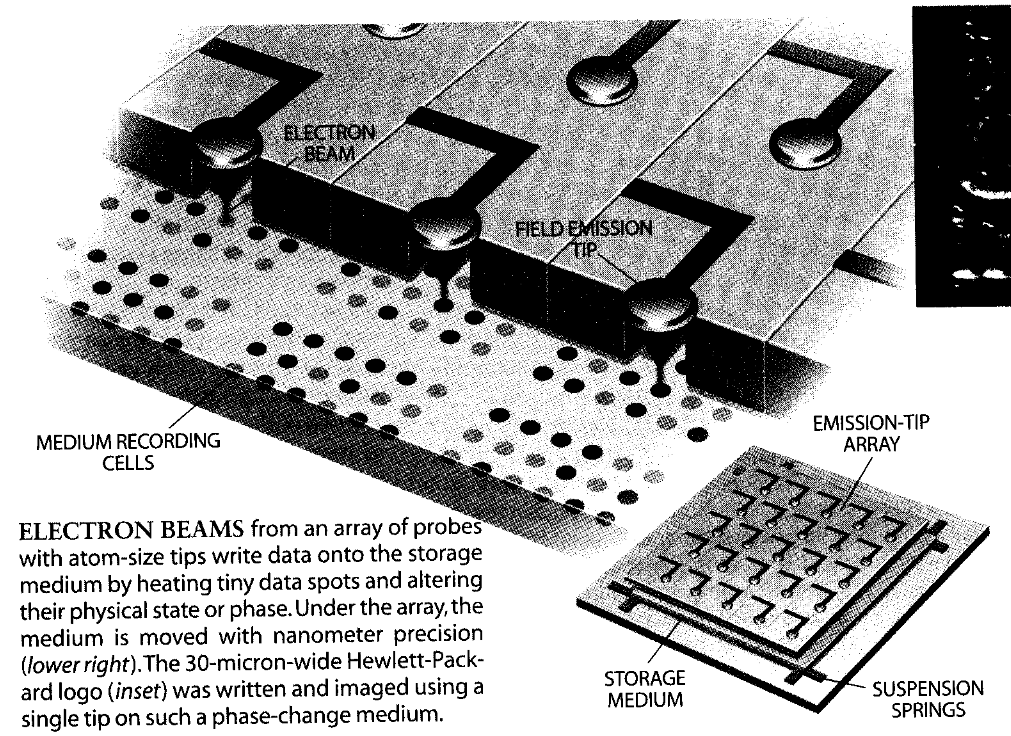

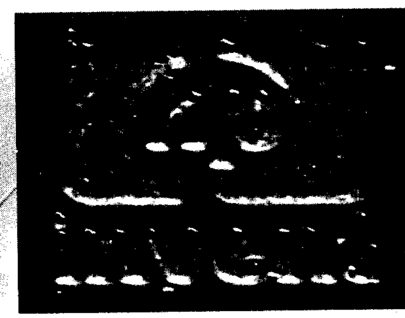

ELECTRON BEAMS from an array of probes with atom-size tips write data onto the storage medium by heating tiny data spots and altering their physical state or phase. Under the array, the medium is moved with nanometer precision (*lower right*). The 30-micron-wide Hewlett-Packard logo (*inset*) was written and imaged using a single tip on such a phase-change medium.

GEORGE RETSECK; HEWLETT-PACKARD (*micrograph*)

"PUNCH CARDS" OF THE FUTURE

ecades ago punch cards were a popular, if clunky, way to store information. Now their mechanical simplicity, shrunk to fantastically small dimensions, is back in vogue at IBM. e company's research group in Zurich has recently developed prototype, dubbed Millipede, that can cram an amazing nount of information into a tiny space—500 gigabits per uare inch—by storing the data as microscopic indentations on lat polymer surface. In comparison, the workhorse magnetic-sk drive is currently limited to no more than 35 gigabits per uare inch, even in the most advanced laboratory prototype. luses are made of silicon and operate simultaneously, making their indentations on a thin layer of plastic coated onto a silicon substrate.

One advantage of Millipede is its anticipated low manufacturing cost. IBM asserts that the arrays can be built economically with processes similar to those used to fabricate computer chips. "You could build hundreds of these arrays on the same wafer," Vettiger predicts.

For increased storage, he says, a future improvement might consist of a larger array (in an earlier incarnation, Millipede had

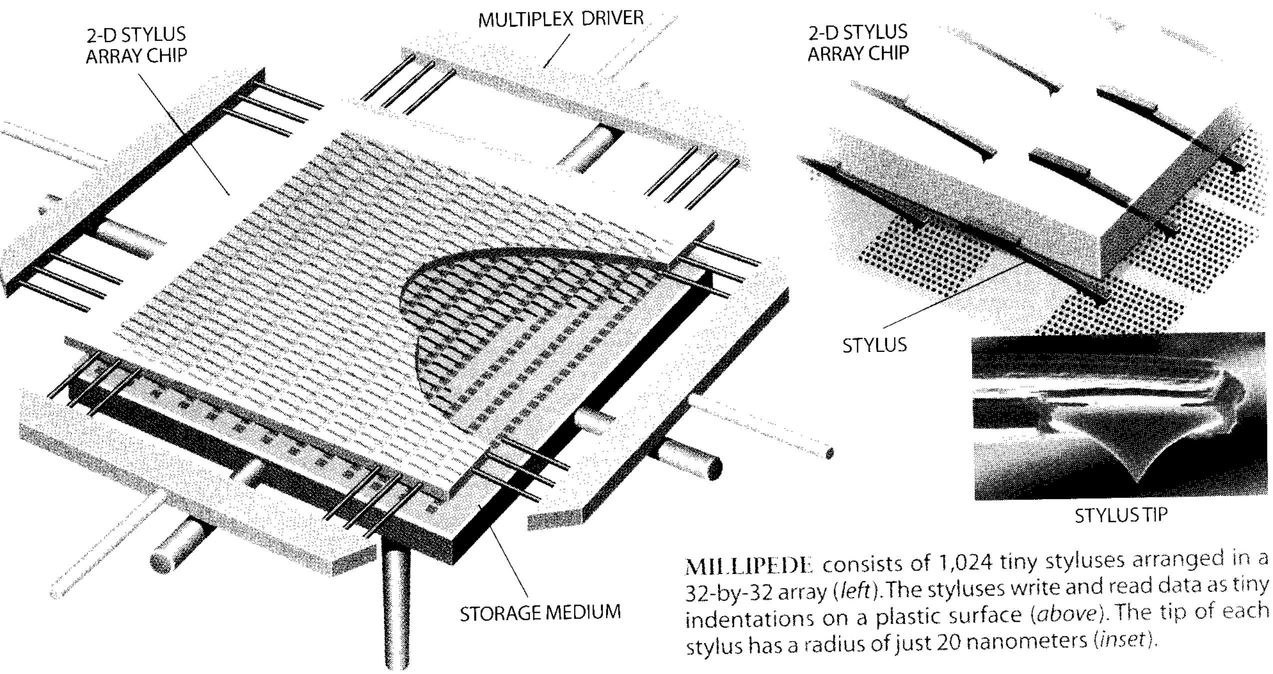

MILLIPEDE consists of 1,024 tiny styluses arranged in a 32-by-32 array (*left*). The styluses write and read data as tiny indentations on a plastic surface (*above*). The tip of each stylus has a radius of just 20 nanometers (*inset*).

'We've reinvented punch cards by using plastics," boasts Peter Vettiger, manager of the project for IBM.

To understand how Millipede works, think of another bygone technology: the phonograph stylus. In IBM's version the tip of the stylus is incredibly sharp (its radius is just 20 nanometers) and rests gently on a smooth, moving plastic surface. To create an indentation, an electric current zaps through the stylus, briefly heating the tip to 400 degrees Celsius, which melts the polymer slightly. A series of such current spikes thus produces a string of indentations, the dips and flats becoming the 0's and 1's of the digital world.

To read this information, the stylus tip is heated to a constant 350 degrees C (below the melting point of the plastic) as it moves across the polymer surface. When the stylus drops into an indentation, heat from the tip dissipates. The resulting temperature drop can be detected by a change in the electrical resistance of the stylus.

IBM has recently turned this basic technology, derived from atomic-force microscopy, into a working prototype. To increase the rate at which data are both written and read, the device consists of 1,024 styluses (hence the name "Millipede") in an area just three by three millimeters. Arranged in a 32-by-32 array, the styjust 25 styluses arranged five by five) or multiple arrays, or a combination of both. "We are completely open at this time. It's a very general concept," Vettiger contends. Another innovation might be the use of carbon nanotubes as stylus tips, which would make them even smaller.

Still, IBM must resolve several core issues before Millipede can become a practical product. One that the company is currently investigating is the long-term durability of the technology. Mechanical wear could dull the sharp tips of the styluses, making their indentations on the plastic surface larger. Eventually the markings might begin to overlap, turning pristine data into digital mush.

To be sure, much work remains before Millipede could even hope to supplant the venerable magnetic-disk drive. A greater opportunity might be in small consumer products. IBM is confident that it can develop a tiny device (10 by 10 by 5 millimeters or considerably smaller) that can store 10 gigabits, making the technology appropriate for cell phones, digital cameras and possibly watches. "This is where we'll concentrate our efforts over the next one to two years," Vettiger says. But he is the first to temper his growing excitement. "This is still an exploratory research project," he admits. "It's not as if tomorrow these products will come on the market." —*Alden M. Hayashi, staff writer*

box on page 70], phase-change metals [see box on page 72], or plastic [see box on page 73].

Although several of these approaches have attracted large investments from the leading manufacturers, most remain in the very early stages of testing. Some of the concepts await research breakthroughs or key advances in supporting technologies before work can begin in earnest on prototypes.

Taking a Breather

Until then, hard-disk manufacturers will continue wringing additional improvements from conventional magnetic technology. "This industry seems to be following an S curve," Porter of Seagate observes. He explains that in the near term, manufacturers will still be riding the wave of advances made in the fields of photolithography and semiconductor-chip manufacturing, which enabled the economical fabrication of magnetoresistive and giant magnetoresistive heads. Consequently, the cost per storage bit has been dropping at a rate of 1.5 percent per week.

But experts agree that the industry will be hard-pressed to sustain that pace beyond the next few years. Gone will be the days when storage capacities grew by 130 percent annually. "Given the technological challenges," Munce says, "we will probably be back to a 60 percent annual growth rate within five years." Adds Porter, "Things are getting more difficult at a faster rate than they have historically. Perhaps a slowdown will give us time to evaluate options."

Interestingly, no one can be sure which of the alternative technologies will bear fruit. IBM itself is hedging its bets, investigating a number of approaches, including the use of patterned media, plastic media and holography. One thing is for certain, though. With the SPE specter looming and the industry confronting core design and materials issues, a window of opportunity might be opening.

But getting through that window will not be easy. "New technologies will need to show a profit quickly and be able to compete on a cost per gigabyte of storage with magnetic disks," Munce asserts. More likely, he adds, the new technologies might first need to carve out niches in the market. (Holographic crystals, for example, might be used for the permanent digitized storage of large libraries.) "The trend is definitely toward cost reduction," Munce points out, "which will make it very difficult for other technologies to get on the tracks."

Most observers agree that, regardless of the uncertainties of future data storage technologies, the Information Age will not come to a grinding halt as a result of the superparamagnetic effect. If delays are encountered in bringing higher-capacity products to market, the void is most likely to be filled in the interim by aggregating many drives into intelligent arrays that represent themselves as individual "virtual" drives to systems and to end users. This approach has already been seen in the case of Redundant Arrays of Inexpensive Disks (RAID) and Just a Bunch of Disks (JBOD) arrays that were deployed during the 1980s and 1990s to meet the storage capacity requirements of companies whose needs were not being satisfied by the single-drive capacities of the day.

Another approach that has garnered significant attention in the industry is to place drives and arrays into storage area networks (SANs) that can be scaled dynamically to meet growing storage appetites. The SAN approach further "virtualizes" storage and will eventually provide a "storage utility" that will automatically deliver to users applications with the size and kind of storage they require from a pool of interconnected storage devices.

These approaches to increasing storage capacity entail greater complexity and, by extension, greater management costs in comparison to single storage devices. But they do provide a means for addressing the needs of businesses and consumers for more elbow room for their data in the short term until new, higher-capacity technologies can be brought to market. Moreover, they will probably be facilitated by declining disk-drive prices, which will result from ongoing improvements in manufacturing processes. ■

The Author

JON WILLIAM TOIGO is an independent information technologies consultant. He is the author of eight books, including *The Holy Grail of Data Storage Management* (Prentice-Hall PTR, 2000), and has written numerous white papers and trade press articles on business automation and storage technology issues. Toigo maintains an information clearinghouse covering data storage and storage management at www.stormgt.org on the World Wide Web.

Further Information

MAGNETO-RESISTIVE HEADS: FUNDAMENTALS AND APPLICATIONS. Electromagnetism Series. John C. Mallinson, contributions by Isaak D. Mayergoyz. Academic Press, 1995.

MAGNETIC DISK DRIVE TECHNOLOGY: HEADS, MEDIA, CHANNEL, INTERFACES, AND INTEGRATION. Kanu G. Ashar. IEEE Press, 1997.

MAGNETIC INFORMATION STORAGE TECHNOLOGY. Electromagnetism Series. Shan X. Wang and A. M. Taratorin. Academic Press, 1999.

MAGNETIC RECORDING: THE FIRST 100 YEARS. Edited by Eric D. Daniel, C. Denis Mee and Mark H. Clark. IEEE Press, 1999.

HOLOGRAPHIC DATA STORAGE. Edited by Hans J. Coufal and Demetri Saltis. Springer-Verlag TELOS, 2000.