this neuron occurred when the sound signal—in this case a sine wave near 1 kHz—reached the contralateral ear (the one farther from the MSO) 200 μs before it reached the ipsilateral or same-side ear.

To examine the role of the inhibitory inputs, which are mediated by the neurotransmitter glycine, the researchers measured the firing rates of this neuron following the injection of strychnine near the electrode used to record the neuron's response. Strychnine blocks the glycine receptors; with the inhibition thus turned off, the peak in the ITD response curve shifted toward zero, as shown in red in figure 2. Similar shifts were seen for the four other ITD-sensitive neurons the researchers examined. Thus, say the researchers, inhibition plays a vital role in determining the ITD response of MSO neurons.

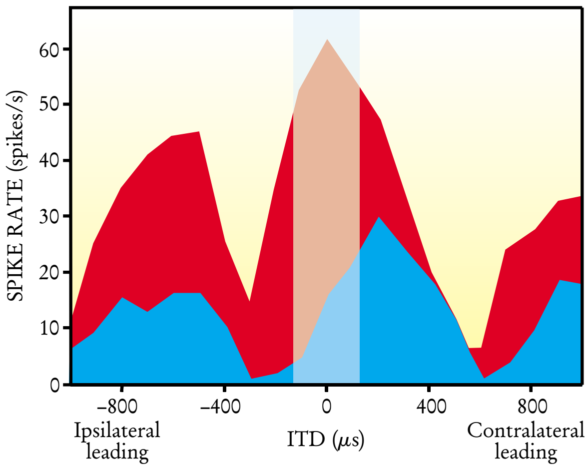

Figure 2. Inhibition is important for establishing the sensitivity to interaural time differences (ITDs). The firing rate, shown in blue, of a sample neuron in the medial superior olive (MSO) of a Mongolian gerbil depends on the difference in sound arrival times at the near (ipsilateral) and far (contralateral) ears. (The side peaks correspond to cross-correlations between different cycles of the input sine wave.) When inhibitory inputs to the neuron are blocked, the cell's response to ITDs shifts to the red curve. The shaded band is the range of physiologically relevant ITDs, and corresponds to the distance between the gerbil's ears. (Adapted from ref. 4.)

The MSO's inhibitory inputs themselves are not sensitive to ITDs. The researchers therefore conclude that the inhibition is precisely timed—that is, phase-locked—to the excitatory inputs. And it is the timing between the inhibitory and excitatory inputs that determines the position of the peak response of the MSO neurons to ITDs. The inhibition is likely dominated by input from the MNTB that reaches the MSO ahead of input from the contralateral cochlear nucleus. In support of that conclusion, Grothe notes that the MNTB receives signals from the contralateral ear through thicker, and hence faster, axons than does the MSO. McAlpine sees advantages to tuning the ITD response through inhibition: It would, for example, allow a mammal's auditory system to adjust to such changes as increasing head size during growth.

Unanswered questions

In addition to demonstrating the role of inhibition in ITD coding, these new results fan an ongoing debate about the nature of ITD processing in the mammalian auditory system. Grothe and McAlpine note that, for most of the neurons they examined, the peak firing rate was found at ITDs outside the so-called physiologically relevant ITD range, given by the spacing between the gerbil's ears divided by the speed of sound and indicated by the shaded band in figure 2. Such tuning seems at odds with the Jeffress model, which holds that the peak firing from coincident inputs encodes the ITD. Instead, the peak location places the steepest part of the response function within the physiological range. And, curiously, the peaks in the responses of MSO neurons are not widely distributed in ITD, as would be expected for a full Jeffress-type map; rather, the peaks all occur at roughly 1/8 of a period of each neuron's so-called best frequency, the frequency at which the neuron generates its largest response. That observation supports similar results found by McAlpine and colleagues in the guinea pig auditory system.[5] With only one peak position for neurons sensitive to a given frequency range, comparisons between the responses from the auditory processing centers on each side of the brain would be required to fully determine the ITD.

Not all hearing researchers are ready to abandon the Jeffress model. Inhibition-mediated ITD tuning could provide an alternative to physical delay lines for realizing a Jeffress-type map. For example, Doug Fitzpatrick (University of North Carolina, Chapel Hill) and his coworkers have shown, in a model, that inhibition can yield sensitivity to large ITDs when only short delay lines are present.[6] Shigeyuki Kuwada of the University of Connecticut Health Center notes that, in most other mammals, particularly larger-headed mammals that might be expected to better exploit ITD cues, peak ITDs do fall within the physiological range. McAlpine counters that in most of the experimental studies to date, the correlation between peak ITD and best frequency, especially at lower frequencies where ITDs dominate, was not determined.

For now, clearer answers to how ITDs are processed in the brain will have to wait for data from more neurons and from more species.

Richard Fitzgerald

References

1. L. A. Jeffress, *J. Comp. Physiol. Psychol.* **41**, 35 (1948).
2. S. R. Young, E. W. Rubel, *J. Neurosci.* **3**, 1373 (1983); C. E. Carr, M. Konishi, *Proc. Natl. Acad. Sci. USA* **85**, 8311 (1988); E. Overholt, E. W. Rubel, R. L. Hyson, *J. Neurosci.* **12**, 1698 (1992).
3. P. H. Smith, P. X. Joris, T. C. T. Yin, *J. Comp. Neurol.* **331**, 245 (1993); G. E. Beckius, R. Batra, D. L. Oliver, *J. Neurosci.* **19**, 3146 (1999).
4. A. Brand, O. Behrend, T. Marquardt, D. McAlpine, B. Grothe, *Nature* **417**, 543 (2002).
5. D. McAlpine, D. Jiang, A. R. Palmer, *Nat. Neurosci.* **4**, 396 (2001).
6. D. C. Fitzpatrick, S. Kuwada, R. Batra, *Hear. Res.* **168**, 79 (2002).

Do Atomic Force Microscope Arrays Have the Write Stuff?

IBM researchers have developed an array of 1024 cantilevers, called Millipede, as a high-density alternative to magnetic recording. Moving across a polymer film, Millipede leaves footprints that encode information.

The information age has been facilitated by the exponentially growing capacity of such storage media as magnetic disks. As demand has soared, the information-storage industry has crammed more and more bits into ever shrinking areas. Significant innovations have already pushed magnetic storage densities well beyond the limits forecast by pundits just a few years ago. Developments now on the horizon promise to raise the densities from values of 30–50 gigabits per square inch (Gb/in^2), which are typical today, to double or triple those values in the near future, and eventually to densities as high as 1 terabit (Tb) per square inch. (See the box on page 16.)

The innovations have not been limited to advances in magnetic-storage technology. Among the new directions being explored is probe storage—that is, using the tips of various types of scanning probe microscopes to write information in the form of bits. Different types of probes can write by inducing different types of surface modification: physical indentations, structural modifications, magnetic alignments, and so forth. To pack 1 Tb/in^2 on a square grid requires a 25-nm spacing between bit centers; scanning probe storage devices should be able to write bits this close together because their tips typically have diameters on the order of ten nanometers. A single probe can't write nearly as fast as today's magnetic writing heads, which record about 1 bit per nanosecond, but large arrays of probes recording in parallel can, in principle, achieve competitive speeds.

Peter Vettiger and his colleagues at IBM's Zürich Research Laboratory have built such a large array, dubbed Millipede. It is made up of 32×32 cantilevers like those of an atomic force microscope (AFM). The researchers recently demonstrated Millipede's ability not only to read and write, with densities of 100–200 Gb/in^2, but also to erase individual bits.[1] They also showed that a single cantilever could write and read with a density of 1 Tb/in^2. For the writing method, the Zürich team built on the work of their colleagues at IBM's Almaden Research Laboratory in San Jose, California, who developed thermomechanical writing—a way of using individual heated cantilever tips reliably to create nanometer-sized indentations on polymer surfaces.[2]

Calvin Quate of Stanford University, a pioneer in AFM technology, called Millipede a "magnificent engineering achievement." He added that, "by operating large numbers of cantilevers in parallel, the IBM Zürich group has advanced probe storage in the direction of integration." Commenting that the IBM Zürich team is several years ahead of other research groups working on probe storage, J. Cock Lodder of the University of Twente in the Netherlands noted that the IBM work "has been an example to us, and one of the main reasons we set out on our project three years ago."

Despite the promise of probe storage, magnetic storage is a well-entrenched technology and one that continues to improve. To compete with magnetic storage—and hence to warrant a major change in industry investment—probe storage would have to perform better and achieve much higher storage densities. "I'm not sanguine about the prospects of any technology to replace magnetic recording in the next several years," commented Thomas Albrecht of IBM Almaden, who has worked in magnetic storage for 12 years and is currently a member of the Millipede team at Zürich. But, he added, "There are already niches where alternative technologies are preferred."

One such niche for Millipede might be handheld devices, for which the cantilever array would be very attractive because of its compact size, low power demand, and low unit cost. Vettiger points out that Millipede offers considerably higher storage capacity than flash memories of the same physical size, such as those used in digital cameras. Magnetic disk storage has a lower cost per gigabit for large memory units, but it does not scale economically to very small sizes. In addition, the motors that spin a magnetic disk during recording eat up a lot more power than the actuators that move the probe-tip arrays.

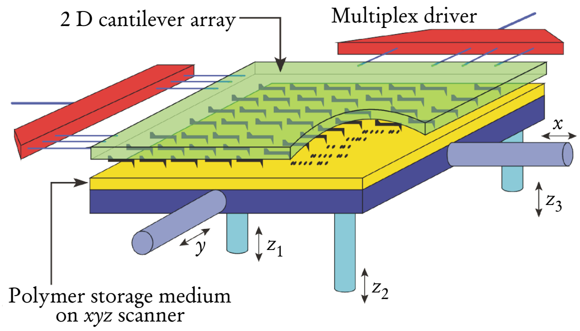

FIGURE 1. MILLIPEDE STORES DATA using an array of 1024 cantilevers (black). The writing surface is a polymer film (yellow) mounted on a silicon substrate (dark blue). Actuators (lighter blue) move the polymer-coated substrate in the *xy* plane beneath the cantilever array, while the individual tips read and write. Other actuators in the *z* direction keep the substrate level. Electrical lines from the multiplex drivers (red) intersect at each cantilever, allowing individual control of each tip. (Adapted from ref. 1.)

Modern punch cards

As illustrated in figure 1, Millipede's two-dimensional cantilever array, which measures 3 mm on a side, sits on a movable storage medium—a thin polymer film that coats a silicon substrate. The cantilevers essentially punch a pattern of bits in the polymer film. Electromagnetic actuators move the medium very precisely in the *xy* plane but no more than the distance between two cantilevers—92 μm. Because the surface moves relative to all 1024 cantilevers in parallel, this small motion enables the entire storage area to be covered with bits. Each cantilever can write about 10 million bits within its local neighborhood, and each is individually addressed by the multiplexed electronics.

The Zürich team fabricated its arrays of cantilevers in batches from crystalline silicon, using surface micromachining techniques. The lever arm for each cantilever was 70 μm long and had a 2-μm downward-pointing tip. When the cantilever array is poised above the writing surface, the height of the tips above the surface must be very uniform. As the array is lowered toward the surface, the tip that hangs down farthest makes contact first; at that moment, no other tip should be more than 500 nm above the surface. This critical criterion minimizes the force necessary to keep the array tips in contact with the surface. When the tip heights are more uniform, they require smaller loading forces and experience less frictional wear as they slide relative to the surface. The Zürich team found that the standard deviation of tip heights for one row of the array was 80 nm, well within the 500-nm upper limit.

To write information on the plastic film, each cantilever can be individually heated to about 400°C. Such hot tips soften the polymer locally and sink into the film, creating tiny encoding pits. As seen in figure 2a, the pits in the Zürich demonstration were as small as 10 nm in diameter. For the reading operation, the cantilevers are also heated, but not to the point of softening the polymer. If a pit (corresponding to a "1") is present, the tip will slide down into the pit, thereby increasing the rate of thermal conduction from

Magnetic Storage: A Moving Target

Like the natural world's evolutionary creatures, the magnetic storage industry has experienced continual innovations to survive—indeed to dominate—in high-density data storage. As the accompanying figure shows, magnetic data-storage density has grown exponentially.

To understand some of the impressive progress that magnetic recording has made, consider first the basic design for storing bits of information on magnetic tapes and disks: A "write head" moves over the magnetic film that coats a recording disk, successively aligning tiny magnetic domains to encode a "1" or a "0." Each bit can be subsequently queried by a tiny sensor in the "read" head, which measures the bit's magnetic field. Typically, the read and write heads are contained in the same unit, which rides on an air cushion above the recording disk as the disk spins at high speeds.

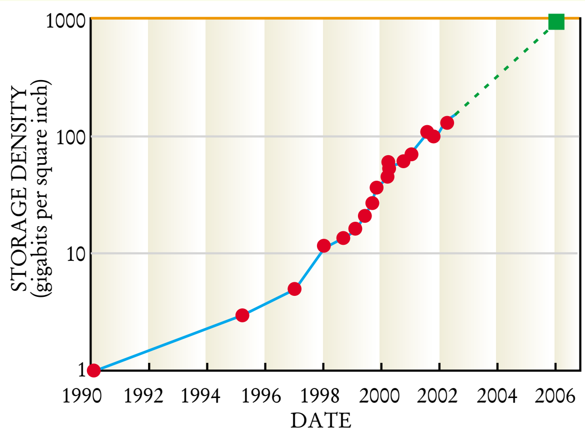

Explosive growth in the storage density for magnetic disk drives since 1990. Red circles mark the announced information density attained in laboratory demonstrations as a function of time. Before 1990, the densities had increased at a slower but still exponential rate; densities in the early 1960s were less than 0.1 megabit per square inch. The Information Storage Industry Consortium (INSIC) has set 1000 gigabits per square inch as a target for 2006 (see green square). (Figure courtesy of INSIC.)

To achieve higher-density storage, industry has made the magnetic films thinner and the domains smaller and more tightly packed. Because smaller domains give weaker signals, industrial developers had to increase the sensitivities of the read and write heads and decrease the gap between each head and the spinning disk. The innovations that have made the greatest impact so far on the growth rate of information-storage densities were the introduction of sensitive magnetoresistive read heads in 1991 and of giant magnetoresistive sensors in 1997.

In the early 1990s, researchers worried that magnetic storage would ultimately be limited by the superparamagnetic effect, in which the magnetic energy of each domain is so small relative to the thermal energy that the magnetic alignment becomes unstable. "We now understand that the superparamagnetic effect is something we must deal with, but not a fundamental limit that prevents all progress beyond a specific density," explained Thomas Albrecht of IBM's Almaden Research Laboratory. One of the first steps that's been taken, he said, is to use an antiferromagnetically coupled medium, which consists of two thin magnetic layers separated by a nonmagnetic spacer a few atoms thick. That approach offers a better combination of the coercivity, grain size, and thermal stability than can be achieved with single-layer media; it thus allows higher-density recording. (Coercivity is the strength of the applied magnetic field needed to permanently change the direction of the magnetization of a material.)

Moving toward 1 terabit per square inch will require further departures from the current way of doing things and will pose considerably greater challenges. For example, one can store the data as perpendicular domains whose magnetization is oriented out of the plane of the disk. Such domains give greater stability to closely spaced bits. Another idea is to avoid the superparamagnetic effect by using a medium with a higher coercivity, but then one must temporarily heat the medium to be able to write to it. Still another avenue to higher density is to use prepatterned media (arrays of isolated magnetic squares, for example), in which each bit consists of only one relatively large grain that's more stable than would be a collection of multiple, small, independent grains. Given its impressive past performance, no one is ready to write off magnetic storage.

the cantilever's heater to the substrate. The cooling rate is determined by measuring the temperature-dependent resistance of the heater.

Daniel Rugar, who helped pioneer this thermomechanical writing process at IBM Almaden, pointed out that the temperature-based readout method introduced by the Zürich group greatly simplifies the device as compared to the piezoelectric readback he and his colleagues had used. For each tip, the Zürich design requires only two leads—for the heaters—and eliminates the additional two required for piezoelectric sensors.

The Zürich group also came up with a way to erase the bits. Their method takes advantage of the lip that piles up like a circular ridge around the edge of each pit. To erase, they wrote a series of offset pits so close to the original pit that the piles from the offset pits fill in the old hole. The results are seen in figure 2.

Wear tests

A natural concern is whether the delicate cantilever tips can stand up to the wear incurred as the surface is constantly dragged beneath them. In the 1990s, the Almaden team conducted a series of wear tests to see if the tips would be subject to unacceptable degradation. They estimated[3] that one could avoid degradation by using a storage medium such as a polymer that is softer than the tip and by keeping the loading force on the cantilevers lower than 5×10^{-8} N.

Vettiger said that cantilevers from Millipede had rewritten the same small area more than 100 000 times without losing the ability to read and write. His team is now operating two computer-controlled testbeds to check the rewritability on a much larger scale. The testbeds will also help them evaluate polymers for their endurance, wear resistance, stability, and power requirements.

The Zürich developers are well aware of the work that still lies ahead before they can claim to have a commercial product. They are currently working to demonstrate a functional storage-system prototype with a standard interface compatible with host devices like digital cameras. Their next goal is to build an array of 64 × 64 cantilevers covering an area 6 mm on a side.

Other types of probe storage

Other development groups are exploring ways to use arrays of probes for data storage, differing primarily in the method of writing bits. One alter-

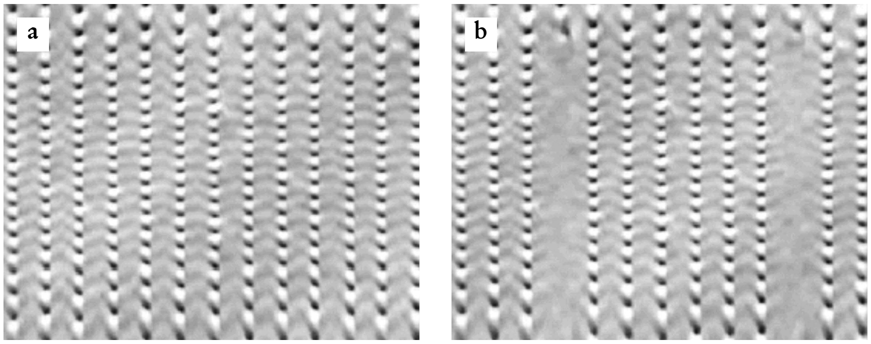

FIGURE 2. WRITTEN AND ERASED BITS are seen in these topographic images, made with Millipede, using the same method by which it reads bits. (**a**) A pattern of pits (bits) having constant spacing in the horizontal direction and variable spacing in the vertical direction. (**b**) Same pattern after most of the bits in two columns have been erased. (Adapted from ref. 1.)

native to making indentations in a plastic film is to record data by changing the structural phase of a small region of material. At Hewlett-Packard's Information Access Lab in Palo Alto, California, for example, researchers are planning to write using the field emission of electrons from a scanning probe tip; the heat imparted by the electrons can change an underlying region from crystalline to amorphous and vice versa. HP's Chuck Morehouse, who heads this "atomic resolution storage" project, said his team has developed a micromotor for the array and is making progress on other components.

Groups at the University of Twente and at Carnegie Mellon University, who cooperate with one another, have chosen a more traditional and proven route to store data: writing magnetic bits. Arthur Davidson, of CMU's center for highly integrated information processing and storage systems, explained that researchers there are designing a cantilever array that's compatible with the standard CMOS fabrication process.[4] They will operate the tips in a noncontact mode, which avoids wear on the tips but requires the control of each tip with a servo and feedback loop. Twente's Cock Lodder said that he and his colleagues are working on the components for what they call a micro scanning probe array. They are exploring probe tips for two modes of reading: one based on magnetic force and the other on magnetoresistance.[5]

Aside from the IBM work, the probe-storage research reported to date is at the individual component level. However, Hideki Kawakatsu and his colleagues from the University of Tokyo have provided a vision of how far probe storage might go: The team used anisotropic etching of silicon to show that one can make an array of millions of fairly uniform cantilevers, with densities of a million per square centimeter.[6] The Tokyo group envisions applications to microscopy and lithography, as well as probe storage.

Barry Schechtman, executive director emeritus of the Information Storage Industry Consortium (INSIC),[7] observed that members of the information-storage community now seem to agree that probe storage is a technology whose time has come. They are starting to focus on the common elements of probe-storage systems and the technical issues that must be solved to bring a product to market. At a November workshop at CMU, jointly hosted with INSIC, one topic of discussion will be the formation of a worldwide probe-storage consortium for the growing number of researchers interested in this area.

BARBARA GOSS LEVI

References

1. P. Vettiger, G. Cross, M. Despont, U. Drechsler, U. Dürig, B. Gotsmann, W. Häberle, M. A. Lantz, H. E. Rothuizen, R. Stutz, G. K. Binnig, *IEEE Trans. Nanotechnology* **1**, 39 (2002). See also http://researchweb.watson.ibm.com/resources/news/20020611_millipede.shtml.
2. H. J. Mamin, D. Rugar, *Appl. Phys. Lett.* **61**, 8 (1992).
3. H. J. Mamin, R. P. Ried, B. D. Terris, D. Rugar, *Proc. IEEE* **87**, 1014 (1999).
4. See http://www.ece.cmu.edu/research/chips/about.
5. See http://utep.el.utwente.nl/smi/content/probe/uspam/uSPAMgeneral.html.
6. H. Kawakatsu, S. Saya, A. Kato, K. Fukushima, H. Toshiyoshi, H. Fujita, *Rev. Sci. Instrum.* **73**, 1188 (2002).
7. See http://www.insic.org.

A Puzzling Increase in Earth's Oblateness

Geophysicists and oceanographers are scrambling to explain why the slight bulge around Earth's equator, which had been slowly shrinking since 1979, abruptly reversed that trend four years ago.

Axial rotation causes Earth to protrude slightly at its midsection. While scientists have long known the magnitude of Earth's oblateness, only since the 1970s have they been able to monitor minute changes in it, thanks to the global view provided by satellites. Such changes in the equatorial bulge reflect large-scale redistributions of mass. For example, the satellite measurements can detect such small effects as the seasonal movements of air masses in the atmosphere and the transport of water among oceans, atmosphere, and land.

On top of seasonal signals, the measured oblateness has shown a slight downward trend over the years, amounting each year to a few tens of parts per billion. Geophysicists attribute the decrease principally to postglacial rebound: Since the polar ice sheets melted away at the end of the last ice age about 10 000 years ago, the underlying mantle has been springing back up in a process that continues today. Postglacial rebound is slowly restoring Earth to a more spherical mass distribution.

The tidy picture has now changed with the recent observation that, sometime around 1998, Earth's oblateness reversed its downward slide and began to increase, as seen in the figure on page 18. "That was something we didn't expect," commented geophysicist Bruce Buffett of the University of British Columbia. Christopher Cox and Benjamin F. Chao of NASA's Goddard Space Flight Center reported the surprising behavior in a recent paper.[1] Several other groups, using different analyses of the same data, have also seen indications of the same effect.

What mass movement could cause so abrupt an effect? That question has sent geophysicists rushing to their computers and databases. In their paper, Cox and Chao surveyed some possibilities: melting of polar ice or mountain glaciers, global sea level