COPYRIGHT 2002 SCIENTIFIC AMERICAN, INC.

THE NANODRIVE PROJECT

INVENTING A NANOTECHNOLOGY DEVICE FOR MASS PRODUCTION AND CONSUMER USE IS TRICKIER THAN IT SOUNDS

By Peter Vettiger and Gerd Binnig

Many engineers have had the thrill of designing a novel

product that then enters mass production and pops up all over the world. We hope—in fact, we would lay better than 50–50 odds on it—that within three years we will experience the rarer pleasure of having launched an entirely new class of machine.

Nanotechnology is much discussed these days as an emerging frontier, a realm in which machines operate at scales of billionths of a meter. Research on microelectromechanical systems (MEMS)—devices that have microscopic moving parts made using the techniques of computer chip manufacture—has sim-

MAKING TRACKS: Arrays of cantilever-mounted tips inscribe millions of digital bits on a plastic surface in an exceedingly small space.

COPYRIGHT 2002 SCIENTIFIC AMERICAN, INC.

ilarly produced a lot of hype and yet relatively few commercial products. But as we can attest, having spent six years so far on one of the first focused projects to create a nanomechanical device suitable for mass production, at such tiny scales, engineering is inextricably melded with scientific research. Unexpected obstacles appear on the road from a proof-of-principle experiment to a working prototype and then on to a product that succeeds in the marketplace.

Here at IBM we call our project Millipede. If we stay on track, by 2005 or so you will be able to buy a postage stampsize memory card for your digital camera or portable MP3 player. It will hold not just a few dozen megabytes of video or audio, as typical flash memory cards do, but several gigabytes—sufficient to store an entire CD collection of music or several feature films. You will be able to erase and rewrite data to the card. It will be quite fast and will require moderate amounts of power. You might call it a nanodrive.

That initial application may seem interesting but hardly earth-shaking. After all, flash memory cards with a gigabyte of capacity are already on the market. The impressive part is that Millipede stores digital data in a completely different way from magnetic hard disks, optical compact discs and transistor-based memory chips. After decades of spectacular progress, those mature technologies have entered the home stretch; imposing physical limitations loom before them.

The first nanomechanical drives, in contrast, will barely scratch the surface of their potential. Decades of refinements will lie ahead. In principle, the digital bits that future generations of Millipede-like devices will write, read and erase could continue to shrink until they are individual molecules or even atoms. As the moving parts of the nanodrives get smaller, they will work faster and use power more efficiently. The first products to use Millipede technology will most likely be highcapacity data storage cards for cameras, mobile phones and other portable devices. The nanodrive cards will function in much the same way as the flash memory cards in these devices today but will offer several-gigabyte capacity for lower cost. The same technology might also be a tremendous boon to materials science research, biotechnology or even applications that are not currently foreseeable.

It was this long-term promise that got us so excited half a dozen years ago. Along the way, we learned that sometimes the only way around a barrier is a serendipitous discovery. Fortunately, besides unexpected obstacles, there are also unexpected gifts. It seems there often is a kind of reward from nature if one dares enter new areas. On the other hand, sometimes nature is not so kind, and you must overcome the difficulty yourself. We have worked hard on such problems but not too hard. If at one stage we had no idea how to address an issue, perhaps a year later we found an answer. Good intuition is required in such cases, in which you expect the problem can be solved, although you do not yet know how.

A Crazy Idea

IN A WAY, Millipede got its start on a soccer field. The two of us played on the soccer team of the IBM Zurich Research Laboratory, where we work. We were introduced by another teammate, Heinrich

<u> Qverview/Millipede Project</u>

- Today's digital storage devices are approaching physical limits that will block additional capacity. The capabilities of the Millipede "nanodrive" a micromechanical device with components in the nano-size range—could take off where current technologies will end.
- Millipede uses grids of tiny cantilevers to read, write and erase data on a polymer media. The cantilever tips poke depressions into the plastic to make digital 1's; the absence of a dent is a digital 0.
- The first Millipede products, most likely postage stamp-size memory cards for portable electronic devices, should be available within three years.

Rohrer. Rohrer had started at the Zurich lab in 1963, the same year as one of us (Vettiger); he had collaborated with the other one (Binnig) on the invention in 1981 of the scanning tunneling microscope (STM), a technology that led to the long-sought ability to see and manipulate individual atoms.

In 1996 we were both looking for a new project in a considerably changed environment. The early 1990s had been a tough time for IBM, and the company had sold off its laser science effort, the technology part of which was managed by Vettiger. Binnig had closed his satellite lab in Munich and moved back to Zurich. Together with Rohrer, we started brainstorming ways to apply STM or other scanning probe techniques, specifically atomic force microscopy (AFM), to the world beyond science.

AFM, invented by Binnig and developed jointly with Christoph Gerber of the Zurich lab and Calvin F. Quate of Stanford University, is the most widely used local probe technique. Like STM, AFM took a radically new approach to microscopy. Rather than magnifying objects by using lenses to guide beams of light or by bouncing electrons off the object, an AFM slowly drags or taps a minuscule cantilever over an object's surface. Perched on the end of the cantilever is a sharp tip tapered to a width of less than 20 nanometers-a few hundred atoms. As the cantilever tip passes over the dips and rises in the surface (either in contact with or in extreme proximity to it), a computer translates the deflection of the lever into an image, revealing, in the best cases, each passing atom.

While Binnig was making the first images of individual silicon atoms in the mid-1980s, he inadvertently kept bumping the tip into the surface, leaving little dents in the silicon. The possibility of using an STM or AFM as an atomic-scale data storage device was obvious: make a dent for a 1, no dent for a 0. But the difficulties were clear, too. The tip has to follow the contours of the medium mechanically, so it must scan very slowly compared with the high-speed rotation of a hard-disk platter or the nanosecond switching time of transistors.

HOW THE NANODRIVE WORKS

Cantilever array

Tracking and

bit-sensing microcircuitry

THE MILLIPEDE NANODRIVE prototype operates like a tiny phonograph, using the sharp tips of minuscule silicon cantilevers to read data inscribed in a polymer medium. An array of 4,096 levers, laid out in rows with their tips pointing upward, is linked to control microcircuitry that converts information encoded in the analog pits into streams of digital bits. The polymer is suspended on a scanning table by silicon leaf springs, which permits tiny magnets (*not shown*) and electromagnetic coils to pan the storage medium across a plane while it is held level over the tips. The tips contact the plastic because the levers flex upward by less than a micron.





WRITING A BIT

Highly doped

Using heat and mechanical force, tips create conical pits in linear tracks that represent series of digital 1's. To produce a pit, electric current travels through the lever, heating a doped region of silicon at the end to 400 degrees Celsius, which allows the prestressed lever structure to flex into the polymer. The absence of a pit is a 0.



Polymer medium

Scanning table

Electromagnetic

coil actuators

ERASING A BIT

The latest Millipede prototype erases an existing bit by heating the tip to 400 degrees C and then forming another pit just adjacent to the previously inscribed pit, thus filling it in (*shown*). An alternative erasure method involves inserting the hot tip into a pit, which causes the plastic to spring back to its original flat shape.

READING A BIT

To read data, the tips are first heated to about 300 degrees C. When a scanning tip encounters and enters a pit (*below*), it transfers heat to the plastic. Its temperature and electrical resistance thus fall, but the latter by only a tiny amount, about one part in a few thousand. A digital signal processor converts this output signal, or its absence, into a data stream (*far right*).





Other pros and cons soon became apparent. Because of the extremely small mass of the cantilevers, AFM operation with the tip in direct contact with the medium is much faster than that of an STM or a noncontact AFM, though still not as fast as magnetic storage. On the other hand, tips of a contact AFM wear quickly when used to scan metal surfaces. And—perhaps most important—once the tip has made a dent, there was no obvious way to "erase" it.

A group led by Dan Rugar at the IBM Almaden Research Center in San Jose, Calif., had tried shooting laser pulses at the tip to heat it; that would in turn soften the plastic so the tip could dent it. The group was able to create compact disclike recordings that stored data more densely than even today's digital video discs (DVDs) do. It also performed extensive wear tests with very promising results. But the system was too slow, and it still lacked a technique to erase and rewrite data.

Our team sketched out a design that

erating a single AFM is sometimes difficult, we were confident that a massively parallel device incorporating many tips would have a realistic chance of functioning reliably.

As a start, we needed at least one way to erase, be it elegant or not. Alternatives, we thought, might pop up later. We developed a scheme of erasing large fields of bits. We heated them above the temperature at which the polymer starts to flow, in much the same way as the surface of wax gets smooth when warmed by a heat gun. Although the technique worked nicely, it was somewhat complicated because, before erasing a field, all the data that were to be retained had to be transferred into another field. (Later on, as we'll explain, nature presented a much better method.)

With these rough concepts in mind, we started our journey into an interdisciplinary project. With the pair of us working in one team, we bridged two IBM departments, physics and devices. (They were eventually merged into a single sci-



99 Percent Perspiration ...

WE WERE NOT MEMS experts, and researchers in the MEMS and scanning probe technology communities had so far dismissed our project as harebrained. Although others, such as Quate's group at Stanford, were working at that time on STM- or AFM-based data storage, ours was the only project committed to largescale integration of many probes. We hoped to achieve a certain vindication by presenting a working prototype in January 1998 at the IEEE 11th International Workshop on Micro-Electro-Mechanical Systems in Heidelberg, Germany. Instead we had a nearly working prototype to show. We presented a five-by-five array of tips in an area of 25 square millimeters.

It was able to demonstrate parallel imaging, but parallel writing failed. We had overlooked a niggling but critical technical detail: the wires leading to the heaters were metallic and too thin to handle the current passing through them. They immediately blew like overloaded fuses because of the phenomenon of electromigration in metal films. Electromigration was well described in the literature, and we should have known about it. This was not our only mistake, but in our group mistakes can be admitted and quickly corrected.

Despite the setbacks, our lab's managers sensed real progress. They allowed us to double the size of our team to eight. We had learned from the 25-tip array that the aluminum wiring had to be re-

Researchers in the MEMS and scanning probe technology communities had dismissed our project as harebrained.

we thought could supply these missing ingredients. Rather than using just one cantilever, why not exploit chipmakers' ability to construct thousands or even millions of identical microscopic parts on a thumbnail-size slice of silicon? Working together in parallel—like the legs of a millipede—an army of slow AFM tips could read or write data quite rapidly.

Here more imagination was required to envision a chance for success than to come up with the idea itself. Although opence and technology department.) We were also joined by Evangelos Eleftheriou and his team, from our laboratory's communication systems department. Today several other groups from within IBM Research and from universities collaborate with us.

When different cultures meet, misunderstandings cannot be prevented, at least not in the beginning. We, however, experienced a huge benefit from mixing disparate viewpoints. placed—which we did with highly doped silicon cantilevers. We also found that it was possible to level out the tip array below the storage medium with high precision in a relatively large area, which made us confident enough to move to a bigger array right away.

Vettiger recognized one serious problem in May 1998 as he was giving an invited talk at the IBM Almaden lab. He was describing how the cantilevers would be arranged in regular rows and columns,

ΙBΜ

MORE TIPS IN A SMALLER SPACE

PROTOTYPE EVOLUTION: Whereas the first-generation Millipede chip contained an array of 25 cantilevers on a square that was five millimeters on a side (*below*), the succeeding prototype (*right*) incorporated 1,024 cantilevers in a smaller, three-millimeter square.





all of them connected to a grid of electrical wires. But as he was explaining how this system would work, he suddenly realized that it wouldn't. Nothing would stop the electric current from going everywhere at once; there would thus be no way to reliably send a signal to an individual cantilever.

The uncontrolled flow of current is actually a well-known phenomenon when units in an array have to be addressed through columns and rows. A common solution is to attach a transistor switch to each unit. But putting transistors on the same chip at the tips was not an option; the tips must be sharpened under intense heat that would destroy tiny transistors. Back at the lab, we tried all kinds of tricks to improve control of the current flownone of which pleased Vettiger. The bigger the array, the more serious this flaw became. A quick calculation and simulation by Urs Dürig of our team showed that for an array of 1,000 units, addressing single cantilevers for writing would still be possible; reading the small signals of individual levers, however, would fail.

Vettiger slept poorly that night, fretting. The team was just about to complete the chip design for a 1,024-tip array. Vettiger told them to wait. For days the team agonized over the problem, until at last Vettiger and Michel Despont saw a practical answer: place a Schottky diode (an electrical one-way street) next to each cantilever. This highly nonlinear device would block the undesired current from flowing into all the other cantilevers. We reworked the design and soon had a 32-by-32-tip array, our second prototype.

This prototype proved that many of our ideas would work. All 1,024 cantilevers in the array came out intact and bent up by just the right amount so that they applied the correct amount of force when mated to a soft polymer medium called PMMA, which is mounted on a separate chip called a scanning table. Copper electromagnetic coils placed behind the scanning table were able to keep the PMMA surface from tilting too much as it panned left, right, back and forward atop the cantilever tips. (A new media scanner designed by Mark Lantz and Hugo Rothuizen has since reduced vibration sensitivity, which was then a problem.) Each 50-micron-long cantilever had a little resistor at its end. An electrical pulse sent through the tip heated it to around 400 degrees Celsius for a few microseconds.

The initial results with our second prototype were encouraging. More than 80 percent of the 1,024 levers worked properly on first pass, and there was only one narrow "dark" zone crossing the center of the storage field, resulting from a twisting of the chip when it was mounted. Not in our wildest dreams did we expect such success at this early stage of the project.

PETER VETTIGER and GERD BINNIG have collaborated extensively to refine technologies for the Millipede nanodrive concept. Vettiger has had a long-standing career as a technologist specializing in microfabrication and nanofabrication. He joined the IBM Zurich Research Laboratory in 1963 and graduated in 1965 with a degree in communications technology and electronics engineering from the Zurich University of Applied Sciences. His academic career culminated in an honorary Ph.D. awarded in 2001 by the University of Basel. Binnig completed his Ph.D. in physics in 1978 at the Johann Wolfgang Goethe University in Frankfurt, Germany, and joined the Zurich lab that same year. His awards for outstanding scientific achievements include the 1986 Nobel Prize for Physics, which he received together with Heinrich Rohrer for the invention of the scanning tunneling microscope.

THE AUTHORS

From R to D

IN THE FIVE-BY-FIVE DEVICE, each lever had at its base a piezoresistive sensor that converted mechanical strain to a change in resistance, allowing the system to detect when the tip had dropped into a pit—a digital 1. We began exploring approaches to detect pits more definitively. We ran tests with Schottky diodes integrated into the cantilevers, hoping that the strain would modify their resistance. Somehow the diodes did not have the expected properties. We nonetheless observed a strong signal when a bit was sensed. After some head-scratching, we found the surprising reason. It turned out same heater on each lever for reading bits as well as writing them. Instead of three or four wires per cantilever, only two would be needed.

We presented this second prototype at the 1999 IEEE MEMS conference. This time the other researchers in attendance were more impressed. But what really excited upper managers at IBM were pictures of regular rows of pits that Millipede had written into the polymer. The pits were spaced just 40 nanometers apart—about 30 times the density of the best hard drives then on the market.

Shortly thereafter, in early 2000, the Millipede project changed character. We



enough to soften the material, surface tension and the springiness of the polymer cause a pit to pop up again. Instead of annealing a larger field using a heater integrated into the storage substrate-as in the block erasure method described earlier-the tip heats the medium locally. Because of electrostatic forces, a certain loading force on the tip cannot be avoided. So when the tip is heated to a high enough temperature and a new indentation is produced, older bits in close proximity are erased at the same time. If a row of pits is written densely, each newly created bit will eliminate the previous one, and only the last bit in the row will remain. This mechanism can even be used to overwrite old data with new code without knowing what the old one was. In a marriage of our experience in physics with Eleftheriou's recording-channel expertise, we developed a special form of constrained coding for such direct overwriting.

At that point it was clear that the team needed to work on the speed and power efficiency of Millipede. We had to start measuring signal-to-noise ratios, bit error rates and other indicators of how well the nanodrive could record digital data. And we had to choose a size and

Although we scientists no longer consider this a high-risk project, we still rejoice when a new prototype works.

to be a thermal phenomenon. If the cantilever is preheated to about 300 degrees C, not quite hot enough to make a dent, its electrical resistance drops significantly whenever the tip falls into a pit [see illustration on page 49]. We never would have thought to use a thermal effect to measure a motion, deflection or position. On macro scales, doing so would be too slow and unreliable because of convection-the circulatory motion that occurs in a fluid medium, in this case air, as heat is transferred between two objects of different temperatures. On the micro scale, however, turbulence does not exist, and hotter and cooler objects reach equilibrium within microseconds.

Although this result was unexpected, it was very useful. Now we could use the

began focusing more on producing a storage system prototype. The team grew to about a dozen workers. We again brought together two departments, with Eleftheriou and his team joining us. They contributed their unique expertise in recording-channel technology, which they had been applying to magnetic recording very successfully. They began developing the electronic part of a fully functional system prototype—from basic signal processing and error-correction coding to complete system architecture and control.

We had just discovered a way to erase a small area, and in cooperation with Eleftheriou, we could even turn it into a system in which no erasing is required before overwriting. In the new, local erasure method, when the tip temperature is high shape for the nanodrive. The "form factor" can be all-important in the consumer electronics marketplace, specifically in the mobile area, which we had chosen to address first.

The Road Ahead

IN THE LAST MONTHS of 2002 our group put the final touches on the thirdgeneration prototype, which has 4,096 cantilevers arranged in a 64-by-64 array that measures 6.4 millimeters on a side. Cramming more levers onto a chip is challenging but doable. Today we could fabricate chips with one million levers, and 250 such arrays could be made from a standard 200-millimeter wafer of silicon. The primary task now is to strike the right balance between two desiderata. First, our design for a complete nanodrive system—not just the array and the scanning table but also the integrated microelectronics that control them—should be inexpensive enough to be immediately competitive and, especially for handheld devices, operable at low power. But it is critical that the system function dependably despite damage that occurs during years of consumer use.

We have found polymers that work even better than PMMA does. In these plastics, pits appear to be stable for at least three years, and a single spot in the array can be written and erased 100,000 times or more. But we are less sure about how the tips will hold up after making perhaps 100 billion dents over several years of operation. Dürig and Bernd Gotsmann of our team are working closely with colleagues at IBM Almaden to modify existing polymers or develop new ones that meet the requirements for our storage application.

And although human eyes scanning an image of the Millipede medium can easily pick out which blocks in the grid contain pits and which do not, it is no trivial matter to design simple electronic circuitry that does the same job with nearperfect accuracy. Detecting which bits represent 0's and which are 1's is much easier if the pits are all the same depth and are evenly spaced along straight tracks. That means that the scanning table must be made flat, held parallel to the tips and panned with steady speed in linear motion-all to within a few nanometers' tolerance. Just recently, we learned that by suspending the scanning table on thin leaf springs made of silicon, we gain much better control of its movement. Even so, we will add an active feedback system that is very sensitive to the relative position of the two parts to meet such nanoscopic tolerances while the device is jostling around on a jogger's waistband.

Any mechanical system such as Millipede that generates heat has to cope with thermal expansion. If the polymer medium and the silicon cantilevers differ by more than about a single degree C, the alignment of the bits will no longer match that of the tips. A feedback system to compensate for misalignment would add complexity and thus cost. We are not yet cer-

High-Density Memory Projects

IBM'S MILLIPEDE PROJECT is only one of several efforts to bring compact, high-capacity computer memories to market.

COMPANY	DEVICE TECHNOLOGY	MEMORY CAPACITY	COMMERCIALIZATION
Hewlett-Packard Palo Alto, Calif.	Thumbnail-size atomic force microscope (AFM) device using electron beams to read and write data onto recording area	At least a gigabyte (GB) at the outset	End of the decade
Hitachi Tokyo	AFM-based device; specifics not disclosed	Has not been revealed	Has not been revealed
Nanochip Oakland, Calif.	AFM-tipped cantilever arrays that store data on a silicon chip	Half a GB at first; potential for 50 GBs	Expected in 2004
Royal Philips Electronics Eindhoven, the Netherlands	Optical system similar to rerecordable CDs using a blue laser to record and read data on a three-centimeter- wide disk	Up to a GB per side, perhaps four GBs in all	Expected in 2004
Seagate Technology Scotts Valley, Calif.	Rewritable system using AFM or other method, operating on a centimeter-size chip	As many as 10 GBs on a chip for portables	Expected in 2006 or later

tain of the best solution to this problem.

Fortunately, nature has helped again. Millipede and the storage substrate carrying the polymer film are both made from silicon and will therefore expand by the same amount if they are at the same temperature. Additionally, the gap between the tip array and the substrate is so small that the air trapped between them acts as an excellent heat conductor, and a temperature difference between them is hardly achievable.

Because the project has now matured to the point that we can begin the first steps toward product development, our team has been joined by Thomas R. Albrecht, a data storage technologist from IBM Almaden who helped to shepherd IBM's Microdrive to market. Bringing the Microdrive from the lab to the customer was a journey similar to what Millipede may face in the next few years.

For the members of our group, this transition to product development means

that we will surrender the Millipede more and more to the hands of others. Stepping back is the most difficult part and, at the same time, the most critical to the success of a project.

Indeed, we cannot yet be certain that the Millipede program will result in a commercial device. Although we scientists no longer consider this a high-risk project, we still rejoice when a new prototype works. If we are lucky, our newest prototypes will reveal problems that our team knows how to fix.

In any case, we are excited that, at a minimum, this nanomechanical technology could allow researchers for the first time to scan a square centimeter of material with near-atomic resolution. So far the project has generated close to 30 relatively basic patents. No one knows whether nanodrives will make it in the market. But they will be a new class of machine that is good for something, and for us that is its own reward.

MORE TO EXPLORE

In Touch with Atoms. Gerd Binnig and Heinrich Rohrer in *Reviews of Modern Physics*, Vol. 71, No. 2, pages S324–S330; March 1999.

The "Millipede"—Nanotechnology Entering Data Storage. P. Vettiger, G. Cross, M. Despont, U. Drechsler, U. Dürig, B. Gotsmann, W. Häberle, M. A. Lantz, H. E. Rothuizen, R. Stutz and G. Binnig in *IEEE Transactions on Nanotechnology*, Vol. 1, No. 1, pages 39–55; March 2002.

For more about nanotechnology in IBM Research and elsewhere, see www.research.ibm.com/pics/nanotech/