A Nanotechnology-based Approach to Data Storage

E. Eleftheriou, P. Bächtold, G. Cherubini, A. Dholakia, C. Hagleitner, T. Loeliger, A. Pantazi, and H. Pozidis,
Advanced Storage Technologies Group
T.R. Albrecht, G.K. Binnig, M. Despont, U. Drechsler, U. Dürig, B. Gotsmann, D. Jubin, W. Häberle, M.A. Lantz, H. Rothuizen, R. Stutz, P. Vettiger, and D. Wiesmann
Micro-/Nanomechanics Group
IBM Research, Zurich Research Laboratory, 8803 Rüschlikon, Switzerland
ele@zurich.ibm.com

Abstract

Ultrahigh storage densities of up to 1 Tb/in.^{2} or more can be achieved by using local-probe techniques to write, read back, and erase data in very thin polymer films. The thermomechanical scanning-probe-based data-storage concept, internally dubbed "millipede", combines ultrahigh density, small form factor, and high data rates. High data rates are achieved by parallel operation of large 2D arrays with thousands micro/nanomechanical cantilevers/tips that can be batch-fabricated by silicon surface-micromachining techniques. The inherent parallelism, the ultrahigh areal densities and the small form factor may open up new perspectives and opportunities for application in areas beyond those envisaged today.

1. Introduction

Data storage is one of the key elements in information technology. The ever increasing demand for more storage capacity in an ever shrinking form factor as well as the pressure to decrease the price per storage unit in $/Gbyte have been a major driving force for substantial worldwide research and development activities to increase storage densities by various means.

For many decades, silicon-based semiconductor memory chips and magnetic hard drives (HDD) have been dominating the data-storage market. So far, both technologies have improved their storage densities by about 60–100% per year, while reducing the cost per gigabyte. However, the areal densities that today's magnetic recording technologies can achieve will eventually reach a limit imposed by the well-known superparamagnetic effect, which today is conjectured to be on the order of 250 Gbit/in.^{2} for longitudinal recording. Several proposals have been formulated to overcome this limit, for example, the adoption of patterned magnetic media, where the biggest challenge remains the patterning of the magnetic disk in a cost-effective manner. In the case of semiconductor memories, such as DRAM, SRAM, Flash etc., the challenges are predominately in lithography to define and fabricate sub-100-nm FET gates as well as very thin gate-oxide materials.

Techniques that use nanometer-sharp tips for imaging and investigating the structure of materials down to the atomic scale, such as the atomic force (AFM) and the scanning tunneling microscope (STM), are suitable for the development of ultrahigh-density storage devices [1-3]. As the simple tip is a very reliable tool for the ultimate local confinement of interaction, tip-based storage technologies can be regarded as natural candidates for extending the physical limits that are being approached by conventional magnetic and semiconductor storage.

Currently a single AFM operates at best on the microsecond time scale. Conventional magnetic storage, however, operates at best on the nanosecond time scale, making it clear that AFM data rates have to be improved by at least three orders of magnitude to be competitive with current and future magnetic-recording technologies. One solution to achieve such a substantial increase in the data rates of tip-based storage devices is to employ micro-electro-mechanical-system (MEMS)-based arrays of cantilevers operating in parallel, with each cantilever performing write/read/erase operations on an individual storage field. We believe that very-large-scale integrated (VLSI) micro/nanomechanics will provide an ideal complement to future micro- and nanoelectronics (integrated

**Proceedings of the 29th VLDB Conference,
Berlin, Germany, 2003**

or hybrid), and may generate hitherto unheard of VLSI-MEMS application opportunities.

Various efforts are under way to develop MEMS-based storage devices. For example, in [4], a MEMS-actuated magnetic-probe-based storage device that should be capable of storing 2 Gbyte of data on 2 cm^2 of die area and whose fabrication is compatible with a standard integrated circuit manufacturing process is proposed. With this approach, a magnetic storage medium is positioned in the *x/y* plane, and writing is achieved magnetically by means of an array of probe tips, each tip being actuated in the *z*-direction. Another approach is the storage concept described in [5], where electron field emitters are employed to change the state of a phase-change medium in a bit-wise fashion from polycrystalline to amorphous and vice versa. Reading is then done with lower currents by detecting either back-scattered electrons or changes in the semiconductor properties in the medium.

The thermomechanical probe-based data-storage concept, our "millipede", combines ultrahigh density, small form factor, and high data rates by means of highly parallel operation of a large number of probes [6-10]. This device stores digital information in a completely different way from magnetic hard disks, optical disks, and transistor-based memory chips. The ultimate locality is provided by a tip, and high data rates result from the massively parallel operation of such tips. As storage medium, polymer films are being considered, although the use of other media, in particular magnetic materials, is not ruled out. Our current effort focuses on demonstrating the concept with areal densities of up to 0.5–1 Tbit/in.^{2} and parallel operation of very large 2D AFM cantilever (up to 64×64) arrays with integrated tips and write/read/erase storage functionality. While a MEMS-based electro-magnetically-activated microscanner moves the polymer medium in the *x/y* directions underneath the array chip, the individual tips can be addressed for parallel write/read operations.

The high areal storage density and small form factor make this concept very attractive as a potential future storage technology in mobile applications, offering gigabytes of capacity and low power consumption at data rates of megabytes per second. Moreover, these features, coupled with the inherent massive parallelism, may open up new perspectives and opportunities for application in areas beyond those envisaged today.

2. Principles of operation

Our AFM cantilever-array storage technique is illustrated in Fig. 1. Information is stored as sequences of indentations and no indentations written in nanometer-thick polymer films using the array of AFM cantilevers. The presence and absence of indentations will also be referred to as logical marks. Each cantilever performs write/read/erase operations within an individual storage field with an area on the order of 100×100 μm^2. Write/read operations

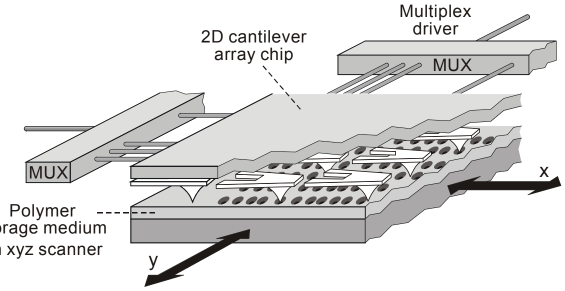

Fig. 1. The "millipede" concept. From [9] © IEEE 2003.

depend on a mechanical *x/y* scanning of either the entire cantilever array chip or the storage medium. The tip-medium spacing can be either controlled globally by a single *z*-actuation system for the entire array, or by simply assembling the device with a well-controlled *z*-position of the components such that the *z*-position of each tip falls within a predetermined range.

Efficient parallel operations of large 2D arrays can be achieved by a row/column time-multiplexed addressing scheme similar to that implemented in DRAMs. In our device, the multiplexing scheme could be used to address the array column by column with full parallel write/read operation within one column. The time between two pulses being applied to the cantilevers of the same column corresponds to the time it takes for a cantilever to move from one logical-mark position to the next. An alternative approach is to access all or a subset of the cantilevers simultaneously without resorting to the row/column multiplexing scheme. Clearly, the latter solution yields higher data rates, whereas the former leads to a lower implementtation complexity of the electronics.

Thermomechanical writing is achieved by applying a local force through the cantilever/tip to the polymer layer and simultaneously softening the polymer layer by local heating. The tip is heated by application of a current pulse to a resistive heater integrated in the cantilever directly above the tip. Initially, the heat transfer from the tip to the polymer through the small contact area is very poor, but it improves as the contact area increases. This means that the tip must be heated to a relatively high temperature of about 400ºC to initiate softening. Once softening has been initiated, the tip is pressed into the polymer, and hence the indentation size is increased.

Imaging and reading are done using a thermomechanical sensing concept. To read the written information, the heater cantilever originally used for writing is given the additional function of a thermal readback sensor by exploiting its temperature-dependent resistance. For readback sensing, the resistor is operated at a temperature in the range of 150–300ºC, which is not high enough to soften the polymer as in the case of writing. The principle of thermal sensing is based on the fact that the thermal conductance between heater platform and storage substrate changes as a function of the distance between them. The medium between the heater platform and the storage substrate, in our case air, transports heat from the cantilever to the substrate. When the distance between cantilever and substrate decreases as the tip moves into a bit indentation, the heat transport through the air becomes more efficient. As a result, the evolution of the heater temperature differs in response to a pulse being applied to the cantilever. In particular, the maximum value achieved by the temperature is higher in the absence of an indentation. As the value of the variable resistance depends on the temperature of the cantilever, the maximum value achieved by the resistance will be lower as the tip moves into an indentation: During the read process, the cantilever resistance reaches different values, depending on whether the tip moves into an indentation (logical bit "1") or over a region without an indentation (logical bit "0"). Under typical operating conditions, the sensitivity of thermomechanical sensing exceeds that of piezoresistive-strain sensing, which is not surprising because in semiconductors thermal effects are stronger than strain effects. The good sensitivity is demonstrated by the images in Fig. 2, which were obtained using the thermal-sensing technique described.

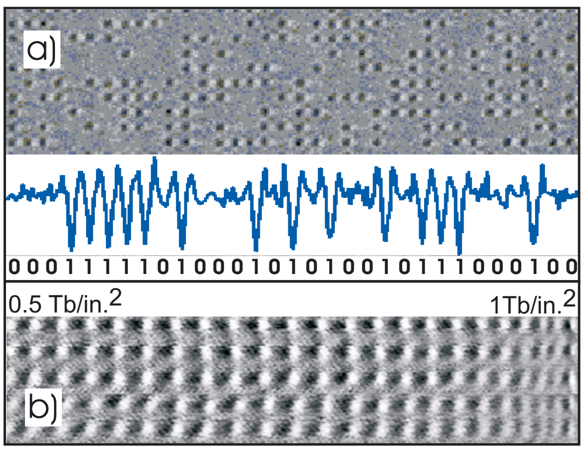

Fig. 2. (a) Data bits section written and readback (160 Gbit/in.^{2}) by a single cantilever using the thermo-mechanical write/read concept as well as the signal line scan read-back signal. (b) Data bits with areal densities approaching 1 Tbit/in^2.

A large 2D array consisting of up to 4096 (64×64) cantilevers with integrated tips, sensors and actuators has been fabricated using silicon micromachining techniques [10]. Figure 3 shows a section of a fabricated chip.

A key issue for our "millipede" concept is the need for a low-cost, miniaturized scanner with *x*/y motion capabilities on the order of 100 μm (i.e. the pitch between adjacent cantilevers in the array). We have developed a microscanner with these properties based on electromagnetic actuation, see Fig. 4. It consists of a mobile platform (that carries the polymer medium) supported by springs and is fabricated from single-crystal silicon.

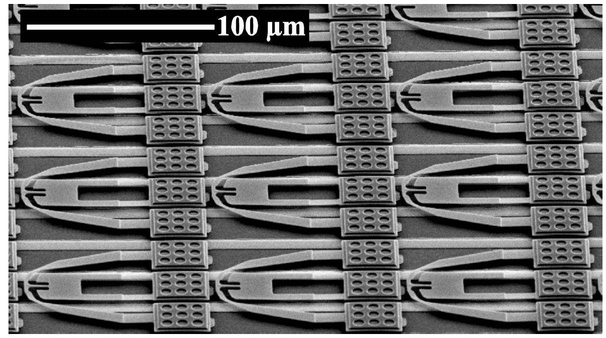

Fig. 3. Scanning electron microscope (SEM) image of a section of the cantilever array transferred and interconnected onto its corresponding carrier wafer.

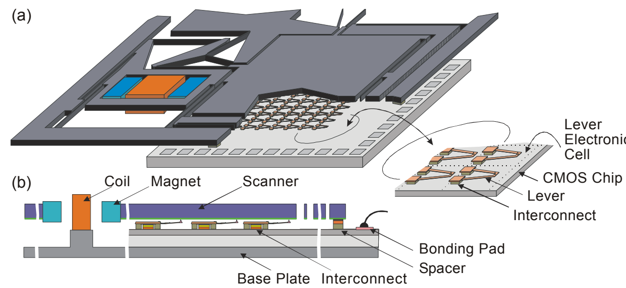

Fig. 4. (a) 3D schematic view of miniaturized scanner and cantilever array chip with integrated electronics. (b) Side-view of scanner, cantilever array, and CMOS electronics.

Actuation in the *x*- and the *y*-direction is achieved by applying a current to a coil positioned between a pair of miniature permanent magnets attached to the silicon scanner. One such coil and its pair of magnets for actuation in one direction are shown in Fig. 4.

3. System aspects

In this section, we describe various aspects of a storage system based on our "millipede" concept including considerations on capacity and data rate. Each cantilever can write data to and read data from a dedicated area of the polymer substrate, called a storage field. As mentioned above, in each storage field the presence (absence) of an indentation corresponds to a logical "1" ("0"). All indentations are nominally of equal depth and size. The logical marks are placed at a fixed horizontal distance from each other along a data track. We refer to this distance, measured from one logical mark center to the next, as the bit pitch (BP). The vertical (cross-track) distance between logical mark centers, the track pitch (TP), is also fixed. To read and write data the polymer medium is moved under the (stationary) cantilever array at a constant velocity in the *x*-direction by the micro-scanner under the control of a servo system.

In general, the servo system in a scanning-probe data-storage device has two functions. First, it locates the track where information is to be written or from which information is to be read, starting from an arbitrary initial scanner position. This is achieved by the so-called seek and settle procedures. During seek, the scanner is rapidly moved with the help of thermal position sensors so that the read/write probes are at a position close to the beginning of the target track. A smaller further move in the cross-track direction from that position to the center of the target track is achieved during the settle mode. As the actuation distances during the seek and settle modes are very small, i.e., on the order of 100 μm, the average data-access time is expected to be on the order of 4 ms. The second function of the servo system is to maintain the position of the read/write probe on the center of the target track during normal read/write operation. This is achieved by the so-called track-follow procedure. Track following controls the fine positioning of the read/write probe in the cross-track direction and is critical for reliable storage and retrieval of user data. It is typically performed in a feedback loop driven by a position-error signal, which indicates the deviation of the current position from the track center line. A robust way to achieve synchronization and servo control in an *x*/*y*-actuated large 2D array is by reserving a small number of storage fields exclusively for timing recovery and servo-control purposes as illustrated in Fig. 5.

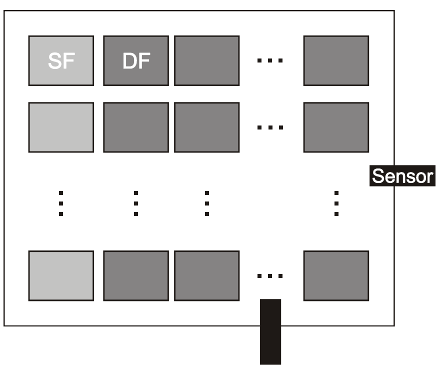

Fig. 5. Layout of data and servo and timing fields. The illustration shows an *x*/*y* top view of the polymer medium on the scanner. Each rectangle outlines the area accessible by a single cantilever tip. Light-grey boxes (SF) indicate servo/timing fields, dark-grey boxes (DF) data fields. The black boxes (sensor) indicate the location of the thermal (position) sensors above the polymer medium.

The two thermal position sensors that are used during the seek and settle modes of operation are also indicated. Because of the large number of levers in our arrays, this solution is advantageous in terms of overhead compared with the alternative of timing and servo information being embedded in all data fields. It has been estimated that the dedicated servo and timing field strategy incurs a very low overhead of less than 5%.

The ultimate locality provided by nanometer-sharp tips represents the pathway to the high areal densities that will be needed in the foreseeable future. The intrinsic nonlinear interactions between closely spaced indentations, however, may limit the minimum distance between successive indentations and hence the areal density. The storage capacity of a "millipede"-based storage device can be further increased by applying (d, k)-constrained codes [11]. The code parameters d and k are nonnegative integers with $k > d$, where d and k indicate the minimum and maximum number of "0"s between two successive "1"s, respectively. For our application, where dedicated clock fields are used, the k-constraint does not really play an important role and accordingly can in principle be set to infinity, thereby facilitating the code-design process. In our code design, where the presence or absence of an indentation represents a "1" or "0", respectively, the d-constraint is instrumental in limiting the interference between successive indentations as well as in increasing the effective areal density of the storage device. In particular, the quantity $(d + 1)R$, where R denotes the rate of the (d, k) code, is a direct measure of the increase in linear recording density. Clearly, the packing density can be increased by increasing d. On the other hand, large values of d lead to codes with very low rate, which implies high recording symbol rates, thus rendering these codes impractical for storage systems that are limited by the clock speed.

Table 1 shows the achievable areal densities and storage capacities for a 64×64 cantilever array with 4096 storage fields, each having an area of 100×100 μm^2, resulting in total storage area of 6.4×6.4 mm^2. For the computation of the storage capacity an overall efficiency of 85% has been assumed, taking into account the redundancy of the outer error-correction coding as well as the presence of dedicated servo and clock fields.

TABLE 1. Total accessible media 0.4096 cm^2, (d=1, ∞) code

pit pitch	track pitch	linear density	track density	areal density	user capacity
25 nm	25 nm	1354.7 kb/in.	1016.0 kt/in.	1376.3 Gb/in^2	10.9 GB
30	30	1128.9	846.7	955.8	6.5
35	35	967.6	725.7	702.2	4.8
40	40	846.7	635.0	537.6	3.7
45	45	752.6	564.4	424.8	2.9
50	50	677.3	508.0	344.1	2.3
55	55	615.8	461.8	284.4	2.0
60	60	564.4	423.3	238.9	1.6

Another important characteristic of a storage device is the sustained data rate for storing or retrieving information. Scanning-probe storage is inherently slow in storing or reading back information with only a single probe or sensor. Figure 6 shows the user data rate as a function of the total number of cantilevers accessed simultaneously. In this diagram, T denotes the time it takes for a probe to move from the center of a logical mark to the center of the next logical mark. Equivalently, $1/T$ represents the symbol rate per probe. In this scenario a (d =1, k = ∞)-constrained

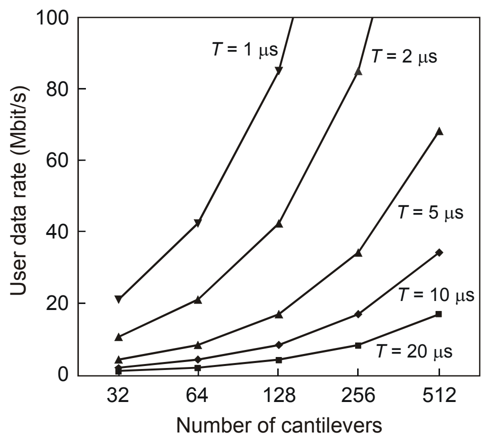

Fig. 6. User data rate versus number of active cantilevers for the ($d = 1$, $k = \infty$) coding scheme. From [9], © 2003 IEEE.

coding scheme is assumed. For example, for a 64×64 cantilever array, a system designed to access a maximum of only 256 cantilevers every $T = 5$ μs yields a user data rate of 34.1 Mb/s.

4. Conclusions

A very large 2D array of AFM probes has been operated for the first time in a multiplexed/parallel fashion, and write/read/erase operations in a thin polymer medium have been successfully demonstrated at densities significantly higher than those achieved with current magnetic storage systems.

The "millipede" array has the potential to achieve ultrahigh areal storage densities on the order of 1 Tbit/in.^{2} or higher. The high areal storage density, small form factor, and low power consumption render the "millipede" concept a very attractive candidate as a future storage technology for mobile applications because it offers several gigabytes of capacity at data rates of several megabytes per second.

Although several of the basic building blocks of "millipede" technology have been demonstrated (including fabrication of large arrays and demonstration of high-density thermomechanical writing and reading), there are a number of issues that need further investigation, such as overall system reliability, including long-term stability of written indentations, tip and media wear, limits of data rates, array and cantilever size as well as tradeoffs between data rate and power consumption.

References

[1] G. Binnig, C.F. Quate and C. Gerber, *Phys. Rev. Lett*. **56**, 930-933 (1986).

[2] H.J. Mamin et al., *Sensors and Actuators A* **48,** 215-219 (1995).

[3] G. Binnig and H. Rohrer, *Rev. Mod. Phys*. **71**, S324-S330 (1999).

[4] L.R. Carley et al., *J. Appl. Phys*. **87**, 6680-6685 (2000).

[5] G. Gibson et al., "Ultra-high Density Storage Device," US Patent 5,557,596 (1996).

[6] P. Vettiger et al., *IBM J. Res. Develop*. **44**, 323-340 (May 2000).

[7] P. Vettiger et al., *IEEE Trans. Nanotechnol*. **1**, 39-55 (2002).

[8] P. Vettiger and G. Binnig, "The Nanodrive Project," *Scientific American*, pp. 34-41 (January 2003).

[9] E. Eleftheriou et al., *IEEE Trans. Magnetics* **39**, 938-945 (2003).

[10] M. Despont et al., *Technical Digest, Transducers'03*, pp. 1907-1910 (IEEE, 2003).

[11] K.A.S. Immink, *Coding Techniques for Digital Recorders* (Prentice Hall International (UK) Ltd., Hemel Hempstead, 1991).