Single-chip computers with microelectromechanical systems-based magnetic memory (invited)

L. Richard Carley,[a] James A. Bain, Gary K. Fedder, David W. Greve, David F. Guillou, Michael S. C. Lu, Tamal Mukherjee, and Suresh Santhanam
Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

Leon Abelmann
Information Storage Technology Group, MESA+ Research Institute, University of Twente, Enschede, The Netherlands

Seungook Min
Department of Materials Science and Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

This article describes an approach for implementing a complete computer system (CPU, RAM, I/O, and nonvolatile mass memory) on a single integrated-circuit substrate (a chip)—hence, the name ''single-chip computer.'' The approach presented combines advances in the field of microelectromechanical systems (MEMS) and micromagnetics with traditional low-cost very-large-scale integrated circuit style parallel lithographic manufacturing. The primary barrier to the creation of a computer on a chip is the incorporation of a high-capacity [many gigabytes (GB)] re-writable nonvolatile memory (in today's terminology, a disk drive) into an integrated circuit (IC) manufacturing process. This article presents the following design example: a MEMS-based magnetic memory that can store over 2 GB of data in 2 cm^2 of die area and whose fabrication is compatible with a standard IC manufacturing process. © *2000 American Institute of Physics.* [S0021-8979(00)90308-6]

I. INTRODUCTION

The purpose of this article is to demonstrate that the underlying technologies required to manufacture nonvolatile rewritable *mass* storage devices using integrated circuit (IC) -compatible fabrication methods are now emerging; and, that these nonvolatile rewritable IC-based mass storage devices have the potential to achieve a quantum decrease in entry cost, access time, volume, mass, power dissipation, failure rate, and shock sensitivity. Such computer systems will dramatically improve performance for today's applications and will offer capabilities that help create many new applications. Specifically, as the technology needed to inexpensively manufacture IC-based mass storage devices is realized, single-chip computer systems integrating several gigabytes (GBs) of IC-based mass storage with processing and traditional random access memories (RAMs) will become a ubiquitous part of our everyday environment—e.g., on our highways, in our homes, in our cell phones, in our bodies—dramatically boosting our quality of life.

There has already been significant progress incorporating nonvolatile memories into silicon ICs. Nonvolatile memories based on charge storage in tunneling oxides (e.g., FLASH EEPROMs) have been in high-volume production for many years.[1] And, there has been a significant body of work focused on development of nonvolatile rewritable memory based on both magnetoresistive materials [magnetoresistive random access memory (MRAM)] (Refs. 2 and 3) and ferroelectric materials [ferroelectric random access memory (FERAM)] (Ref. 4). However, the stated goal of this research is the inclusion of 2–10 GB of nonvolatile rewritable mass data storage on a single IC whose size is generally limited by manufacturing considerations to a maximum of about 4 cm^2 by yield and cost considerations. Allowing roughly 10 bits of raw storage per byte (overhead and redundancy required for error correction codes) and assuming at best 33% of the die area is available for mass storage, the minimum goal of this research must be areal storage densities above 10^{10} bits/cm^2; i.e., bit cells smaller than 100 nm by 100 nm.

FLASH, FERAM, and MRAM all require bit cells that are 3–4 λ on a side, where λ is the smallest feature size that can be defined lithographically. The semiconductor industry is in production today with $\lambda \approx 180$ nm and the Semiconductor Industries Association (SIA) Roadmap for future semiconductor technologies and suggests that we will reach an EEPROM density of 2 GB per chip by 2011 at the earliest.[5,6] The purpose of this article is to explore an alternative approach that does not require advanced lithography to achieve high areal density—using microelectromechanical systems (MEMS) to position magnetic probe tips over a magnetic media with nanometer accuracy.

II. EMERGING TECHNOLOGIES

As we observed in Sec. II, the underlying technologies needed to create MEMS-positioned magnetic probe-based mass storage devices are currently emerging. In this section, we briefly discuss nonmagnetic and then magnetic probe-based technologies. Finally, we discuss the underlying MEMS positioning technologies.

[a]Electronic mail: lrc@ece.cmu.edu

One of the best-developed probe-based mass data storage approaches at the moment is the IBM Millipede.[7,8] Writing in this system is achieved by melting pits in a polymer, and, a readout signal is elegantly obtained by measuring the thermal conductance of the tip to the medium. Since both mechanisms involve only heating and measuring the resistance of the tip, the number of connections per probe is limited to 2. Further, the Millipede employs probe tips that are passively positioned in Z and that are fixed in X and Y. IBM has demonstrated the integration of a large number of probes (32×32) on a small $3\,\text{mm}\times3\,\text{mm}$ surface. Like FLASH RAM, Millipede requires a block erase mechanism. Rewriting can be achieved by heating, which remelts the surface, thus removing the pits. Thus, writing time and storage system power dissipation are both negatively impacted by this choice of writing mechanism. And, the upper limit on the number of allowed rewrite cycles is currently unknown.

Besides melting pits into the polymer, there are many possible nonmagnetic data storage mechanisms. Likely candidates for an IC-based mass storage system include ferroelectric memories,[9] phase-change materials,[10] and charge trapping devices.[11] However, none of these storage mechanisms has demonstrated the combination of a long-term stable memory element and the unlimited rewritability of magnetic recording. Therefore, we will focus on probe-based magnetic recording technologies in this article.

Although there are many scanning microscopy techniques used to detect magnetization patterns in thin films [Kerr-SNOM (e.g., Ref. 12), scanning Hall (e.g., Ref. 13) and scanning superconducting quantum interference device microscopy (e.g., Ref. 14)], only two techniques can be easily adapted for writing magnetic domains: magnetic force microscopy (MFM, e.g., Refs. 15 and 16) and magnetoresistance microscopy (MRM, e.g., Ref. 17).

MFM is capable of detecting details with diameters down to 30 nm.[18] Analysis has shown that it is still the tip that limits the resolution, so higher resolutions should be possible if improved MFM tips become available.[19] Writing can be achieved in the MFM mode simply by tailoring the material properties in such a way that a domain is nucleated when the tip is nearly in contact with the sample, as is shown, for instance, in Refs. 20 and 21. This writing scheme can be modified by applying additional external fields or locally heating the sample.

Although very promising with respect to achievable data density, the MFM mode suffers from a limitation in the data rate set by the cantilever resonant frequency. In principle, this limitation does not have to be serious in a massively parallel system. Assuming a maximum data rate of 100 kbit/s per probe, a 32×32 probe array with all probes working in parallel would still achieve a data rate of over 10 MB/s. Even larger arrays are envisioned in the architecture discussed in the next section.

The MRM mode of reading and writing does not require mechanical displacement of the cantilever: the probes can be run at constant height above the surface. Therefore, the data rate per probe can approach that currently achieved in hard-disk systems; although in practice it will also be limited by the total power dissipation of the die. This suggests that the aggregate data rate of the probe recording system could be one or more orders of magnitude higher than that in state-of-the-art hard disks. Writing in the MRM mode can be achieved just as in a conventional hard-disk writehead by saturating a soft magnetic tip with an external field, locally applied per probe by a micron-size coil. Again the writing cycle can be assisted by external fields applied over the complete probe array. It might be difficult to realize magnetoresistance (MR) sensors on the tip at the very small dimensions which are envisioned. Instead, one could benefit from the tremendous increase in MR sensor sensitivity and guide the flux from the bit in the medium though a soft magnetic tip to a micron-size MR sensor on the cantilever.

Development of MEMS technology for positioning of both media and probe tips has been pursued by several research groups besides our own. As mentioned above, the IBM demonstration of a 32×32 probe array in the Millipede is the largest probe tip demonstration to date. However, their probes are purely passive cantilevers whose resonant frequency is selected in order to keep the probe tips in contact with the media surface at the highest possible bit rates. Another group which has developed arrays of cantilevers with probe tips is at Stanford University; however, their primary focus has been on imaging and patterning. They have demonstrated a parallel array of 50 cantilevers used to generate an atomic-force microscope image spanning a 100 mm^2 area.[22]

The single-crystal reactive etching and metallization (SCREAM)[23,24] process developed at Cornell University achieves high aspect ratio ($>10{:}1$) single-crystal silicon structures, and is ideal for lateral actuation of MEMS. They have demonstrated functional $X-Y$ stages using comb-drive actuators with total sled motions approaching 100 μm and an array of Z-motion actuators for individual probe tips using a torsional parallel-plate electrostatic actuator, with an integrated single-crystal silicon probe tip. Their storage approach proposes a 10 000 element array of torsional Z actuators constructed from single-crystal silicon flipped on top of a large $X-Y$ scanning stage on which the data are encoded. Unfortunately, the perforated nature of the SCREAM media positioner inhibits full utilization of areal density. Furthermore, the SCREAM process inhibits the placement of substantial electronics close to the probe tips, with the resulting additional interconnect reducing the effective utilization of the silicon area for probe tip array.

There has been significant work on developing media acutators other than comb drives. For example, Trimmer and Gabriel[25] demonstrated the idea of an electrostatic linear motor and more recently, Hoen and his colleagues at HP have demonstrated a complete media actuator making use of linear motors to achieve 16 μm of motion at only 4 V in a 3 mm by 3 mm structure.[26] Another alternative actuator structure for long motions is the ''shuffle motor'' (e.g., Ref. 27).

III. EXAMPLE OF MEMS-ACTUATED PROBE STORAGE

In this section, we describe a MEMS-actuated probe-based mass data storage system that has been designed to act as an ultra-high-density, miniature, rewritable data cache for

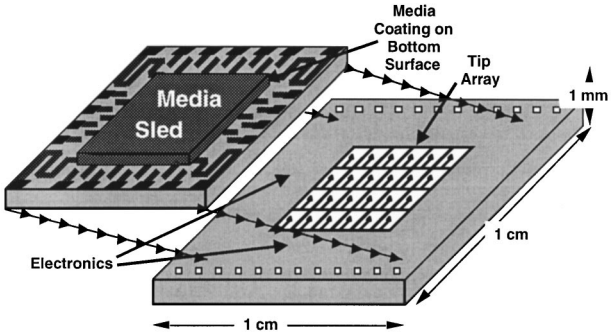

FIG. 1. Conceptual diagram MEMS-actuated data storage devices with $X-Y$ motion of the media and Z motion of a 4×5 array of probe tips.

low-power low-bandwidth communication of bursts of information from sensor arrays. This data cache stores high-bandwidth bursts of information from imagers or other sensors for later transmission via low-power telemetry. In order to carry out this mission, the mass storage device must meet the following performance goals: (i) nonvolatile rewritable magnetic storage; (ii) capacity>2 GB (8 mm$\times$8 mm magnetic media); and aerial density>3 GB/cm^2; (iii) system volume$=2\,500$ mm^3, including package and system footprint$=2$ cm$\times2$ cm; (iv) active power<1 W, and standby power<10 mW; and (v) data transfer rate>20 MB/s and worst-case access time<5 ms.

A. General system-level design constraints

There are several very general design considerations for any MEMS-actuated probe-based data storage system. All such systems must control the position of the probe tip relative to the media in three degrees of freedom: *X*, *Y*, and *Z*. To date the lifetimes of all MEMS bearings are much too short to be useful in a miniaturized disk drive. Alternatively, a piece of media or a probe tip attached to the substrate via a spring can move back and forth. But, the big disadvantage of such reciprocating motion is that the size of the motion is usually a fraction of the length of the springs and actuators along that dimension, typical motions are 10% or less of the suspension/actuator length. Worse yet, moving by only 10% in two dimensions would result in only ''sweeping'' or addressing 1% of the $X-Y$ footprint of the device with a single probe tip. The obvious answer to this problem, as illustrated in Fig. 1, is to have an array of probe tips so that the total motion of the media only needs to match the probe tip pitch in order to completely sweep the media sled (media plate) area. Note, to the extent that the proportions in Fig. 1 are correct, there is a large amount of room to place the rest of the computer system underneath the media acutators and suspension around the outside of the data storage device, making the single-chip computer a cost-effective idea.

There is still an area overhead for the media suspension and actuators. Therefore, increasing the size of the media sled for a given probe tip pitch always results in better overall area efficiency. In fact, only practical factors like the curvature of the MEMS structures and the tolerance to external gravitational forces (i.e., shock tolerance) encourage a smaller size media sled. Although it would be equally possible to move either the media or an array of probe tips, the practical problem of wiring up the probe tips separately makes it far more desirable to have the media undertake the largest motions as it only requires one wire (for magnetic designs). Due to the fact that the media surface is not perfectly flat and the heights of all the probe tips are not perfectly equal, there must also be some independent *Z* motion at every probe tip. And, given that we have chosen to place the media on a sled and move it in *X* and *Y* as shown in Fig. 1, then the easiest way to achieve the necessary *Z* motion is by providing a suspension and actuator for each probe tip.

The specified data transfer rate for this design is quite low; and, in this application it is determined by the need to store digitized video signals in real time. Since we plan to have a large array of tips operating simultaneously, the data rate for an individual probe tip is even lower. For example, if we operate 1000 probe tips simultaneously at 20 MB/s (roughly 200 Mb/s), the required data rate would only be 200 kb/s per tip, which corresponds to about a 100 kHz signal bandwidth. There are three advantages to this relatively low data rate per tip: (1) the noise bandwidth is significantly smaller for the head and electronics, which improves the signal-to-noise ratio; (2) the power required to implement each channel is lower; and (3) the maximum mechanical speed of motion required is much lower, which facilitates precise positioning.

However, the wiring of those probe tips to the servo and channel electronics becomes very difficult unless the electronics can be included directly on the same die as the probe tips. This also greatly improves the bandwidth and sensitivity of capacitive sensors that are integrated into the probe tips to determine their *Z* positions relative to the media. In order to make a highly integrated CMOS+MEMS process, we have developed a series of postprocessing steps following a standard complementary metal–oxide–semiconductor (CMOS) fabrication.[28,29]

B. Magnetic head/media design

A head-medium recording combination that is compatible with the system architecture described above must satisfy a number of design constraints. Examples of *fabrication constraints* include thermal compatibility with CMOS, geometrical compatibility with CMOS-MEMS, and chemical compatibility with the release process. Examples of *operation constraints* include power constraints due to heating, current constraints due to electromigration, small parasitic magnetic forces, sufficient field magnitude for writing, sufficient sensitivity for reading, sufficient field gradient for mark definition, and stability of marks in the medium.

While many possibilities exist for satisfying these constraints, one particular head-medium design, shown schematically in Fig. 2 (not to scale), is featured in this article. It uses a perpendicular magnetic recording scheme with a magnetic tip in proximity to perpendicularly oriented media, either a Pt/Co multilayer film or a rare-earth transition-metal (RE/TM) film, developed for magneto-optic (MO) recording. Unlike films for MO recording, this medium will have a soft underlayer in the manner of perpendicular magnetic record-

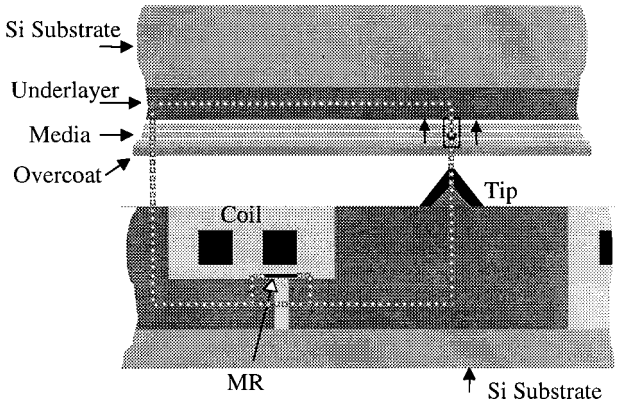

FIG. 2. Head/media structure (cross section).

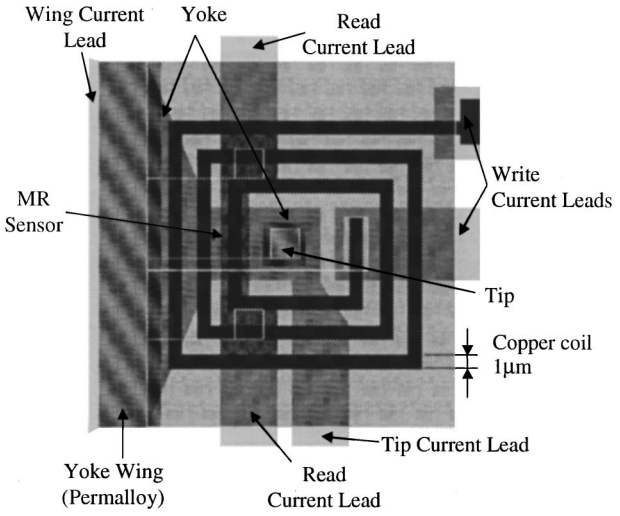

FIG. 3. Top view of head design.

ing. This underlayer will be easy to orient for noise suppression, since the system operates in a rectilinear, rather than circumferential recording pattern. The head is envisioned as a soft tip within a magnetic circuit, in a rather standard perpendicular recording geometry. Detection is accomplished through a giant magnetoresistance (GMR) (or other magnetoresistive) sensor in the back yoke of the system. One interesting design feature is that the writing flux traverses the GMR element. This is possible without premature saturation because the GMR element is so much larger in cross section than the envisioned bit. All of the magnetic processing in this design is done after all CMOS processing and before the MEMS processing.

The envisioned probe-style read–write head will be situated on top of a 25 μm by 25 μm pallet, as part of the CMOS-MEMS actuator discussed later in this section. As a three-metal-layer CMOS process is used, if two anchor points for the pallet are employed, then up to six electrical connections can be made to the probe pallet, allowing a flexible base on which to develop read/write transducers.

The envisioned fabrication process consists of a series of deposition and planarization steps, to produce a magnetically soft NiFe yoke (plated and vapor deposited), a Cu coil (electroplated) and a GMR read sensor. As shown in cross section in Fig. 2, each planarization step buries the previous topography in a dielectric (SiO_2 or Al_2O_3) and then uses chemical–mechanical polishing (CMP) to achieve the planarized surface. This planarization is crucial, as this device requires the fabrication of several thick layers (1–2 μm), while requiring fine linewidth capability (0.5 μm). The sensor is deposited over a planarized gap in the bottom yoke, and separated from the yoke by 0.2 μm of SiO_2. After the sensor is deposited and encapsulated, the coil, yoke wings, and tip are all deposited in succession, with planarization steps after each structure is created. The yoke wings are an attempt to increase the efficiency of the magnetic circuit, bringing the permeable material closer to the medium for a better return path. In manufacturing, deposition of the array of probe tips could be done using the Spindt process[29–32] with minimum effect on the other parts of the structure. Alternatively, a combination of wide-area optical lithography with small-area e-beam lithography can be used to cost effectively manufacture an array of photoresist dots where tips are desired down to 50 nm in diameter.[33] For our analysis, we focus on the latter approach. By using this photoresist mask to pattern a magnetic thin film, 50-nm-diam cylinders with a height of 100–200 nm can be created cost effectively.[33]

A simple magnetic circuit analysis indicates that the device described here will be able to generate adequate field to write media with coercivities of up to 400 kA/m (5000 Oe), while being able to extract enough flux to read. Note, in order to write this high coercivity media, we plan to use tip materials with $M_s > 1600$ kA/m such as iron nitride.[34–36] Micromagnetic models, assuming that the tip is shaped like a brick (45 nm×45 nm×180 nm),[37,38] indicate that achieving a relatively high flux transmission efficiency through the tip will be possible, but at the cost of severe nonlinearity. Under the assumption of perfect head efficiency, it is estimated that the minimum current necessary to generate the required write field is approximately 3 mA for a three-turn-head design. The coil layout of the prototype device is shown in Fig. 3, in plan view, indicating how the coil is routed.

One concern is whether the head will saturate (during writing) at flux levels below that which are needed to generate enough field to switch the medium. However, it is important to remember that in this geometry the bit area is much smaller than the cross section of the MR sensor. In our prototype design we plan to create an MR sensor that has a 10 nm magnetic thickness and a 1 μm width which is four times larger than the proposed 50 nm×50 nm media bit size. Consequently, the tip will saturate before the MR sensor.

It is also appropriate to assess the power consumed by this device during the writing process. In the design shown in Fig. 3, there are about 350 squares in the write coil, using a coil layer that is 0.5 μm thick. This gives a coil resistance of about 15 Ω. At 3 mA, this resistance would consume about 0.1 mW. Estimates of the steady-state heat transport out of the head suspension under these conditions generates a long-term steady-state temperature rise of less than a 50 °C, which is acceptable.

For readback, the yoke design spreads the flux collected by the tip out over the entire length of a 7-μm-long GMR sensor. This design satisfies all of the thermal and SNR con-

siderations that have currently been identified. In the proposed MR sensor, the saturation magnetization is 160 kA/m. This is about 20% that of permalloy, and would be achieved by using a synthetic ferrimagnet as the free layer.[39–41]

If all of the flux from the medium reached the sensor in such a head, the output voltage can be calculated simply from the expression $\Delta V = J\Delta RL$, where J is the current density, ΔR is the change in resistivity of the sensor after exposure to the flux, and L is the length of the sensor exposed to the flux. Analysis indicates that at least half should actually make it through the sensor. Assuming a 20-nm-thick sensor, having a 10-nm-thick free layer with a M_s of 160 kA/m, a sheet resistance of 15 Ω/square, a length L of 7 μm, a 1 μm stripe height, 50% flux guide efficiency, and a 10% magnetoresistance coefficient, a value of 132 μV is obtained for ΔV at 1 mA of current. A 1 μm by 7 μm sensor will have a resistance of about 105 Ω, which yields a Johnson noise-limited SNR of approximately 22 dB in a 1 MHz band width when combined with a preamplifier having an input noise floor of 10 nV/sqr (Hz). The power dissipated in this design at 1 mA of current is only 105 μW and the current density for the sensor is about 5×10^6 A/cm^2, which is acceptable for these transition-metal materials.

The perpendicular magnetic media envisioned for this probe recording system is PtCo multilayer media. These media have been heavily investigated for magneto-optical recording and have several desirable properties. Most important, these media have been demonstrated to support 60 nm magnetic domains (see Betzig *et al.*[12]) and have suitable magnetization, coercivity and Curie temperature for thermomagnetic recording. It also has the desirable feature that the coercivity, magnetization, and Curie temperature can be tuned across a relatively wide range of values by changing the ratio of Pt to Co layer thickness. The impediment to using these media will be the development of a low-noise soft underlayer for flux closure.

C. Probe tip positioning

The job of the probe tip positioning system is to adjust for the difference in Z heights between different probe tips. The source for these differences can be curvature or roughness of the media or the CMOS MEMS die and differences in probe tip height across the die. In addition, there are uncertainties in the exact spacing between the media wafer and the CMOS MEMS wafer during bonding. Finally, external acceleration forces can cause a change in the Z height between the media and the probe tip. The probe tip suspension must be designed so that it has enough range of actuation to be able to servo out all of these variations in Z. One possible prototype that can achieve a quite large range of Z values is shown in Fig. 4. The probe is mounted on a pallet (on this prototype it is smaller than the 25 μm^2 specified above) is at the bottom and just above it is the array of fingers is used to provide a parallel-plate Z actuation upward to the media surface. Note, rather than a square, as suggested in Fig. 1, this tip actuator has a 4:1 aspect ratio. Staggering alternate rows by 100 μm, the probe tips will be in an array with a 100 μm in both X and Y directions.

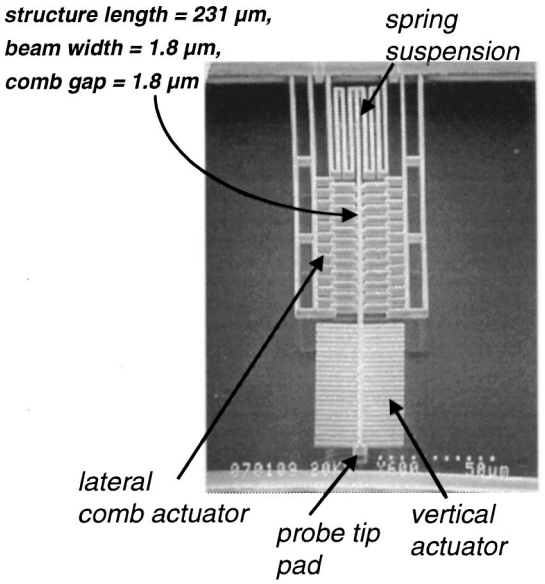

FIG. 4. Prototype probe head positioning system.

The structure in the center at the top of Fig. 4 is a serpentine flexure designed to make the spring softer (smaller spring constant) without making it longer. On either side of the spring, we can see a set of comb fingers. These are used to provide a small, up to 0.5 μm, lateral actuation. This is for use in large media sled designs where a small temperature difference between the media wafer and the CMOS MEMS wafer makes it impossible to position all of the tips simultaneously. For example, with an 8 mm by 8 mm media sled, a 5 °C thermal difference results in a few hundred nanometers of difference that must be taken out by this lateral actuator.

Experimental measurements on the design shown in Fig. 4 indicated a resonant frequency for the beam of about 10 kHz. A static movement of the pallet upward in the Z direction of 0.5 μm required only 8 V on the actuator fingers (near the bottom of Fig. 4). And, a static lateral motion of 0.5 μm required 19 V on the lateral actuator fingers (near the middle of Fig. 4).

D. Media positioning

All of the media area must be accessed by the heads to exploit the ultrahigh data density of probe storage. Therefore, the media actuator must move ±50 μm in both X and Y directions. Bigger actuators are better for large stroke and, hence, swept media area. The stroke of a suspended electrostatic microactuator is equal to the electrostatic force divided by the suspension spring constant and varies as roughly L^4, where L is a characteristic microactuator length. The ratio of swept media area over microactuator footprint increases as L^6. A millimeter-scale actuator achieves the desired 50 μm stroke. Since the media size is also mm scale (about 8 mm by 8 mm), combining the media and actuation into a single device is feasible and appropriate.

A conceptual layout of an initial media actuator design is given in Fig. 5. Our initial design for the media actuator is a decoupled-mode $X-Y$ microstage originally conceived for use as a vibratory-rate gyroscope.[42] A first-generation proto-

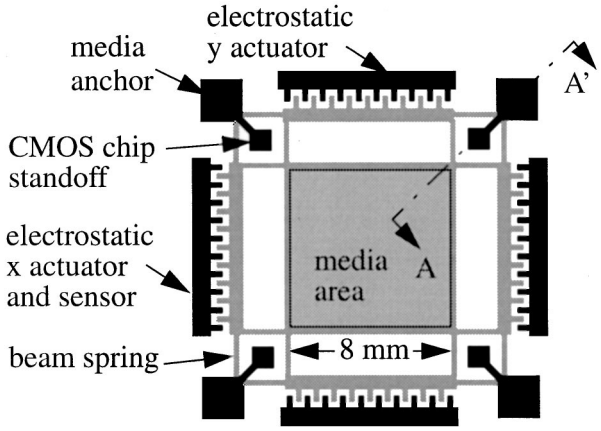

FIG. 5. Conceptual diagram of a 100-μm-stroke media actuator fabricated using deep Si RIE. The foot print is around 14 mm×14 mm. Anchored regions are black, the movable structure is gray.

type actuator has been fabricated using deep-Si reactive-ion-etch (RIE) technology. The high-aspect-ratio silicon structures from the through-wafer process provide media stability in the vertical direction while enabling electrostatic actuation to drive up to ± 50 μm. A boxspring suspension is used to decouple the two lateral directions $(X-Y)$ of actuation, so that comb fingers can be used for $X-Y$ actuation without mechanical interference. Electrostatic actuators with greater than 200 interdigitated fingers with 500 μm thickness and 16 μm gap are designed to achieve the ± 50 μm displacement in each direction. A by-product is that the stage is a sensitive two-axis accelerometer. Electrostatic force in addition to the spring reaction force is required to balance the inertial forces on the stage. In this case, a lateral 10 G force on the device will use up 10% of the actuating force that is available for moving the media sled.

Detailed mechanical finite-element analyses have been carried out on this actuator structure shown in Fig. 5. The lowest mechanical resonant frequency in X and Y was 739 Hz, while the lowest resonant frequency in Z was 2940 Hz. These high resonant frequencies make the positioning system relatively shock tolerant and also make it relatively easy to achieve under 1 ms worst case access time. Note, in order to achieve the ± 50 μm displacement, a large actuator voltage of 120 V is required. However, since there are only X and Y drivers needed, these could be implemented with off-chip high-voltage drivers.

ACKNOWLEDGMENTS

The authors thank the staff of the DSSC/CMU cleanroom and MOSIS. This research was sponsored in part by the NSF under Grant No. ECD-8907068, and in part by DARPA under the AFRL, Air Force Materiel Command, USAF, under Agreement No. F30602-97-2-0323. The United States government has certain rights to this material. The research of Leon Abelmann has been made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences.

[1] T.-S. Jung *et al.*, IEEE J. Solid-State Circuits **32**, 1748 (1997).
[2] J. M. Daughton, Thin Solid Films **216**, 162 (1992).
[3] J.-G. Zhu, Y. Zheng, and G. Prinz, J. Appl. Phys. (to be published).
[4] Y. Shimada *et al.*, Jpn. J. Appl. Phys., Part 1 **36** (1997).
[5] *1997 Technology Roadmap for the Semiconductor Industry* (Semiconductor Industries Association, 1997).
[6] *Update to 1997 Technology Roadmap for the Semiconductor Industry* (Semiconductor Industries Association, 1998).
[7] M. Despont *et al.*, Proceedings of MEMS 1999, p. 564–569, (1999), pp. 564–569.
[8] P. Vettiger *et al.*, Proceedings of 10th International Conference on Scanning Tunneling Microscopy (STM-99), p. 4, Seoul Korea, July (1999).
[9] S. Molitor *et al.*, Proceedings of 10th International Conference on Scanning Tunneling Microscopy (STM-99), Seoul Korea, July (1999), p. 347.
[10] N. Yamada, MRS Bull. 48 (1996).
[11] B. D. Terris and R. C. Barret, IEEE Trans. Electron Devices **42**, 944 (1995).
[12] E. Betzig *et al.*, Appl. Phys. Lett. **61**, 142 (1992).
[13] A. Oral, S. J. Bending, and M. Henini, J. Vac. Sci. Technol. **14**, 1202 (1996).
[14] L. N. Vu and D. J. Harlingen, IEEE Trans. Appl. Supercond. **3**, 1918 (1993).
[15] Y. Martin and H. K. Wickramasinghe, Appl. Phys. Lett. **50**, 1455 (1987).
[16] S. Porthun, L. Abelmann, and J. C. Lodder, J. Magn. Magn. Mater. **182**, 238 (1998).
[17] R. O'Barr, M. Lederman, and S. Schultz, J. Appl. Phys. **79**, 6067 (1996).
[18] L. Abelmann *et al.*, J. Magn. Magn. Mater. **190**, 135 (1998).
[19] L. Abelmann, R. Proksch, and J. C. Lodder, Proceedings STM'99, Seoul Korea (1999), pp. 477–478.
[20] S. Manalis, K. Babcock, J. Massie, V. Elings, and M. Dugas, Appl. Phys. Lett. **66**, 2585 (1995).
[21] T. Ohkubo, J. Kishigami, K. Yanagisawa, and R. Kaneko, IEEE Trans. Magn. **29**, 4086 (1993).
[22] S. C. Minne, J. D. Adams, G. Yaralioglu, S. R. Manalis, A. Atalar, and C. F. Quate, Appl. Phys. Lett. **73**, 1742 (1998).
[23] N. MacDonald, Microelectron. Eng. 49 (1996).
[24] Z. L. Zhang and N. C. MacDonald, J. Micromech. Microeng. **2**, 31 (1992).
[25] W. S. N. Trimmer and K. J. Gabriel, Sens. Actuators **11**, 189 (1987).
[26] S. Hoen, P. Merchant, G. Koke, and J. Williams, Proceedings of the 9th International Conference on Solid-State Sensors and Actuators (Transducers '97), Chicago, II, June 1997, pp. 41–44.
[27] N. Tas, J. Wissink, L. Sander, Th. Lammerink, and M. Elwenspoek, Sens. Actuators A **70**, 171 (1998).
[28] G. K. Fedder *et al.*, Sens. Actuators A **57**, 103 (1996).
[29] G. K. Fedder, Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS '97). Hong Kong, June 9–12, 1997.
[30] C. A. Spindt, J. Appl. Phys. **39**, 3504 (1968).
[31] C. A. Spindt, I. Brodie, L. Humphrey, and E. R. Westerberg, J. Appl. Phys. **47**, 5248 (1976).
[32] C. A. Spindt, C. E. Holland, A. Rosengreen, and I. Brodie, IEEE Trans. Electron Devices **38** (1991).
[33] Y. Kondoh, J. Seeger, and P. Merchant, J. Microelectromech. Syst. **7**, 428 (1998).
[34] J. A. Bain, J. Wolfson, M. H. Kryder, Y. Yip, J. Reid, R. D. Silkensen, and R. H. Dee, IEEE Trans. Magn. **32**, 166 (1996).
[35] H. L. Hu, L. Vo, T. Nguyen, N. Robertson, M. Re, and C. Jahnes, IEEE Trans. Magn. **30**, 3870 (1994).
[36] W. P. Jayasekara, S. Wang, and M. H. Kryder, J. Appl. Phys. **79**, 5880 (1996).
[37] L. Abelmann *et al.*, J Appl. Phys. (to be published).
[38] L. Abelmann, J.-G. Zhu, J. A. Bain, K. Ramstock, C. Lodder, J. Appl. Phys. (to be published).
[39] V. S. Speriosn *et al.*, Intermag Conference '99, Seattle, WA, April (1996).
[40] H. A. M. van den Berg *et al.*, IEEE Trans. Magn. **32**, 4624 (1996).
[41] J. L. Leal and M. H. Kryder, Symposium L of the Spring 1998 MRS Meeting: San Francisco, CA, April 12–16 (1998).
[42] M. S. Kranz and G. K. Fedder, Proceedings of the Symposium Gyro stroke. Since the media size is also mm scale (about 8 mm by Technology, Stuttgart, Germany, September 16 and 17 (1997), pp. 3.0–3.8.