Simulations of magnetic materials with MDGRAPE-2

B. G. Elmegreen
R. H. Koch
M. E. Schabes
T. Crawford
T. Ebisuzaki
H. Furusawa
T. Narumi
R. Susukita
K. Yasuoka

The use of accelerator hardware for micromagnetics simulations is described, along with some initial results. The accelerator calculates the dipole interactions at 400 gigaflops, allowing large simulations to be performed with arbitrary geometries. Two research programs are highlighted, the simulation of a curved MRAM cell and the simulation of the write head in a computer disk drive.

Introduction

Magnetic materials continue to play an important role in digital data storage [1, 2]. Intensive research efforts currently aim to extend the storage capacity of a magnetic disk to areal densities of 1 Tb/in.^{2} or more using perpendicular recordings [3], while magnetic random access memory (MRAM) shows great promise for a large range of innovative applications in information technology and consumer electronics [4, 5].

Perpendicular disk recording is projected to allow for continued reduction in the size of the data bits of a magnetic disk by polarizing the magnetic surface layer of the disk in the vertical direction with the north pole up or down, unlike today's disks, which have in-plane polarity. The perpendicular orientation makes it possible to achieve higher bit densities because the data can be written using closed loops of magnetic flux between the write head and a magnetically soft underlayer. This makes it possible to concentrate the flux and write sharper transitions, facilitating the use of narrower tracks on the disk and closer bit spacings along the tracks.

In the MRAM technology, data is stored on integrated chips as an array of magnetic bits with two stable polarities. One polarity represents a 1 bit, and might, for example, have the north magnetic pole pointed to the left, while the other polarity represents a 0 bit with the north magnetic pole pointed to the right. For an elongated bit analogous to a bar magnet, the polarity direction inside the material is stable, and the memory cell does not need a continuous current to maintain it. This stability contrasts with the situation in the DRAM technology, which is the most commonly used memory chip technology. Without the need for a current, MRAM data stays written even when the computer is turned off. This means, e.g., that through its use, portable PC battery life is extended during cordless use and PC operating system and commonly used software do not need rebooting after a power-down.

Computer simulations of these devices can greatly aid in the design process by optimizing the sizes and shapes of various components and the material properties that are used [6–10]. A typical simulation might divide the device into small portions and approximate each with a cell of uniform magnetization. The interactions between these cells then mimic the behavior of the whole material.

Here we discuss a novel technique for such simulations, pioneered at IBM Research over the last few years, in which the mutual interaction between computational magnetic cells is accelerated with special-purpose hardware linked to a conventional computer. The hardware was developed in the astronomical community in order to accelerate gravitational interactions between stars in a cluster [11]. The acronym for the technique is GRAPE, from the term *GrAvity PipE*. The "pipe" in this expression refers to a pipeline of sequential arithmetic operations that determine the squared distance between a pair of coordinate points, (x, y, z) and (x_1, y_1, z_1),

$$r^2 = (x - x_1)^2 + (y - y_1)^2 + (z - z_1)^2, \quad (1)$$

and then the distance-dependent vector acceleration between them, $(x - x_1)/r^3$ for the x component, $(y - y_1)/r^3$ for the y component, and so on in the case of gravity. The pipeline then adds this acceleration to the running sum of all accelerations between the test point at (x, y, z) and the other points representing pieces of the device at (x_2, y_2, z_2),

0018-8646/04/$5.00 © 2004 IBM

(x_3, y_3, z_3), and so on. Because each cell interacts with every other cell, the number of accelerations that must be calculated equals the square of the number of cells. Pipelining speeds this up tremendously because all of the operations for a single term in the sum are carried out simultaneously for one particle or another down the pipe. Many pipelines can be placed on a single chip, many chips on a computer board, and many boards can be attached to a host computer.

For a problem of modest size having a number of cells less than about a million, the accelerator is faster than conventional techniques by a factor of ~10 to 100 for irregular cell positions and complex geometries. For larger numbers of cells, more GRAPE hardware can be added, but more economical methods should be considered instead, such as grouping the cells into a hierarchical tree structure and using only the branches for the more distant parts of the tree [12, 13]. For regular cell positions, FFT methods are faster [14]. GRAPE accelerators have been used in astrophysics for more than ten years [11, 15].

In 1995, IBM undertook a collaborative project with Dr. Ebisuzaki and his team at The Institute of Physical and Chemical Research (RIKEN) in Japan to design and manufacture a chip that could accelerate the calculation of arbitrary central forces between all pairs of particles in a system. The main application of such a chip is in molecular dynamics (MD) calculations in which long-range forces between charged particles and intermediate-range van der Waals forces must be evaluated at regular time steps to determine the accelerations and motions of the atoms. One method commonly used for these simulations is the Ewald sum method [16], which involves the evaluation of error functions and Gaussian functions of the separations between pairs of particles. There were two previous versions of this general-force hardware, GRAPE-2A [17] and MD-GRAPE [18]. The current version is called MDGRAPE-2. Extensive discussions of MDGRAPE-2 and its applications may be found in [19] and [20].

Several years ago, RIKEN built a 15.4-teraflop (Tflop) machine using these chips and another type of chip for performing the wave-space part of the Ewald sum (called "WINE," for Wave INtegrator for the Ewald method). The combined system ran molecular dynamics software as rapidly and as accurately as a 1.34-Tflop conventional computer. It shared the 2000 Gordon Bell Prize of the IEEE with GRAPE-6, which is another version optimized for calculations involving the inverse square force law [21]. In 2001, RIKEN completed its full Molecular Dynamics Machine (MDM) with 1536 MDGRAPE-2 chips and 2688 WINE chips. The MDM runs at a peak rate of 78 Tflops and an effective rate for calculations involving the Ewald method that is equivalent to 8.61 Tflops on a conventional computer. RIKEN is currently planning a 1-petaflop-peak molecular dynamics machine for 2006.

This paper discusses some of the research done at IBM using MDGRAPE-2 for micromagnetics. The computer consists of an IBM parallel AIX* machine with eight MDGRAPE-2 boards attached to the PCI buses of separate nodes. Each board contains four MDGRAPE-2 chips and storage for one million particle positions and charges, and each chip has four pipelines with a 100-MHz clock rate. The chips were fabricated by the IBM Advanced Microelectronics Division in Burlington, Vermont, and the boards were fabricated by Advanet, Inc., Tokyo, Japan. The distribution of the MDGRAPE-2 boards and the software for the drivers and libraries used in programming them is handled by a Japanese company, Peta Computing, Inc. (*http://www.peta.co.jp/*). Our eight-board system runs at a sustained speed of 400 gigaflops for the evaluation of the dipole interactions, which is the most time-consuming part of the simulation when the number of dipoles exceeds about one thousand.

Magnetic simulations via MDGRAPE-2

The micromagnetics simulations were based on an assumed array of dipoles that could respond individually to the magnetic fields in their neighborhoods. The field that results from all of the other dipoles could be evaluated in several ways with various approximations, but with the accelerator hardware we used, it was most convenient to sum the contribution of each dipole directly. A typical simulation also included three other types of magnetic forces from short-range interactions, which were evaluated on the main processors of the computer. The short-range forces arise from 1) the exchange magnetic field, which tends to co-align near-neighbor dipoles; 2) the unilateral anisotropy field, which tends to give all of the dipoles a preferred direction; and 3) the external field. These other three fields do not use MDGRAPE-2 for evaluation and are discussed in more detail elsewhere [22, 23].

The force on each dipole from the local magnetic field causes it to precess around an equilibrium orientation that is aligned to the field. The mutual precession of all of the dipoles causes all four fields to change with time. The simultaneous evolution of the changing fields and dipole orientations is too complicated to evaluate analytically from the initial conditions, so it must be followed numerically by stepping along each dipole orientation in small increments and then re-evaluating the new fields everywhere for these slightly different orientations. This time-stepping is done on the host computer, following the Landau–Lifshitz–Gilbert formalism for the precessional motions, and using a fourth-order Runge–Kutta integrator with a variable time step [24]. The time step was determined by forcing the rms change in the dipole position to be between 10^{-4} and 10^{-7}; with these limits, the integration over time was stable.

The only purpose of MDGRAPE-2 is to find the magnetic field at each dipole resulting from the summed magnetic fields of all of the other dipoles. This is determined from the curl of the local magnetic vector potential, **A**, which is evaluated from

$$A_x(k) = f_{y,z}(k) - f_{z,y}(k),$$
$$A_y(k) = f_{z,x}(k) - f_{x,z}(k),$$
$$A_z(k) = f_{x,y}(k) - f_{y,x}(k), \quad (2)$$

where

$$f_{x,x}(k) = \sum_j \frac{\mu_x(j)[x(k) - x(j)]}{|r(k) - r(j)|^3},$$
$$f_{x,y}(k) = \sum_j \frac{\mu_x(j)[y(k) - y(j)]}{|r(k) - r(j)|^3},$$
$$f_{x,z}(k) = \sum_j \frac{\mu_x(j)[z(k) - z(j)]}{|r(k) - r(j)|^3}. \quad (3)$$

The position $\mathbf{r}(k)$ is that at which the vector potential is to be measured, and $\mathbf{r}(j)$ is the position of some dipole j in the field of dipoles. The x-component of the dipole moment is μ_x. There are similar equations for $f_{y,x}$, $f_{y,y}$, and $f_{y,z}$, which arise from the y-component of the dipoles, μ_y, and for $f_{z,x}$, $f_{z,y}$, and $f_{z,z}$, which arise from μ_z.

The functions f are evaluated using the MDGRAPE-2 accelerator. At first, the dipole positions $\mathbf{r}(j)$ are sent to all of the MDGRAPE-2 boards along the PCI bus inside the computer. This is done with a subroutine, `M2_SET_POSITIONS`, that is called by the main program. The positions are stored in memory on the boards. In a second step, the x-components of the dipole moments, μ_x, are sent to all of the boards by another subroutine, `M2_SET_CHARGES`, and these are also stored on the board memory. Finally, a subroutine, `M2_CALCULATE_FORCES`, sends the positions $\mathbf{r}(k)$ directly to the pipelines of the MDGRAPE-2 chips. These $\mathbf{r}(k)$ are the positions at which the $f_{x,i}(k)$ functions are to be evaluated. In parallel operations, these $\mathbf{r}(k)$ are divided equally between the boards. With each clock cycle in the pipeline, a new term is added to the f-sum. When all of the field dipoles on the board have had their terms added, the three f components for the particle at $\mathbf{r}(k)$ are returned to the host. Next, the y-component of μ is sent to all of the boards with another `M2_SET_CHARGES` call, and the corresponding functions $f_{y,i}(k)$ are determined with another `M2_CALCULATE_FORCES` subroutine. A third pair of calls sends the z-component and calculates $f_{z,i}(k)$. The host computer then rearranges these f functions to create the vector potentials at the $\mathbf{r}(k)$ positions, as in Equation (1), and the difference of vector potentials determines the dipole field.

The positions $\mathbf{r}(k)$ are displaced from each dipole by half the dipole spacing, aligned in *x, y,* and *z* coordinate directions. If the dipoles are not on a regular array, each dipole must have six distinct positions $\mathbf{r}(k)$. In the parts of the array that are colinear, coplanar, or part of a regular cubical lattice, the intermediate positions between adjacent dipoles can be shared by the same k-indices to minimize the number of points at which f is evaluated.

The dipole approximation assumes that a continuous magnetic material can be simulated by an array of discrete dipoles. This is only an approximation because the nearest-neighbor dipoles do not exert the same force on each other as two smooth cubical lattices of the same magnetic material. The approximation is good if the number of dipoles in the system is large. For system sizes exceeding ~20 dipoles on a side, whether in a cube, layer, or between parallel layers, the difference between the theoretical value for the magnetic field and the value measured by MDGRAPE-2 is around 1%.

Sample results

The effects of layer curvature on the bit-switching properties of an MRAM element were investigated using MDGRAPE-2. A thin rectangular layer of magnetic material is assumed to be initialized with an external field pointing mostly along the long axis, and then the external field is turned off to allow the layer to relax. The magnetization finds an equilibrium that is aligned along the long direction of the rectangle in the center, and it bends near the short ends to be parallel to those edges as well. This is the state of dipole orientations that has the least amount of magnetic flux outside the rectangle. External flux costs the system magnetic energy, so the lowest-energy state, which is the equilibrium, minimizes this cost.

Figure 1 shows an example of the results of a simulation of the equilibrium state. The arrows represent the in-plane components of a planar array of 48 × 96 dipoles that correspond to a single MRAM bit; the colors represent their curl. After this initial setup, a new external field is applied in a different direction. If the new external field is strong enough and aligned mostly in the opposite direction to the setup field, the magnetization in the layer rotates to follow the new field. If the external field is too weak or if it has too small a component in the long direction, the magnetization turns only a little and becomes stuck. The problem is to determine what external field will flip the magnetization from one direction to the other. Each direction is a stable bit of information that can be read with a perpendicular current [4, 5].

The simulations carried out with MDGRAPE-2 utilized its distinctive capability to place dipole tracers anywhere in three-dimensional space without the loss of computational speed. Normally, FFT techniques are used to calculate dipole interactions, and these require the use of regular grids [7]. Without the constraint of a flat geometry for

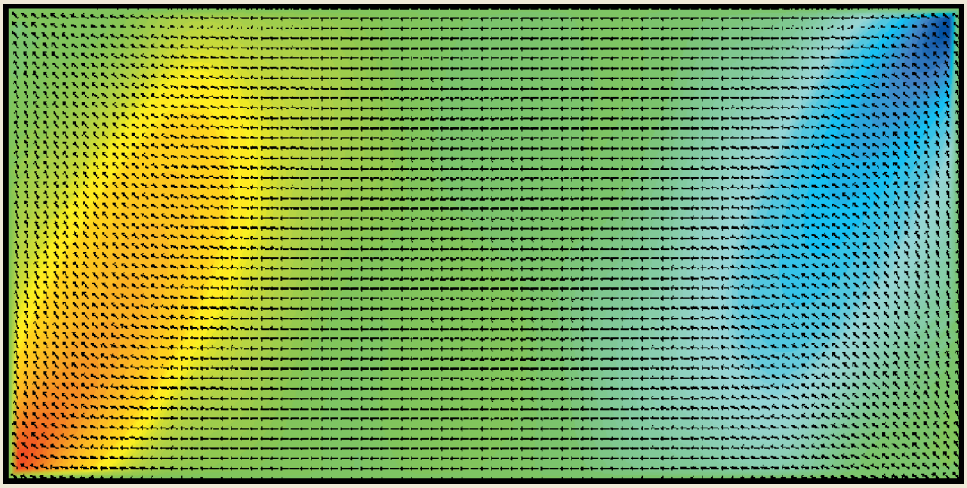

Figure 1

Example of results of a simulation using MDGRAPE-2: Orientations of a planar array of 48 × 96 MRAM dipoles in the equilibrium state, following a 10-ns set-up phase. Tiny arrows represent the in-plane components of the dipoles, and color represents the perpendicular component of the curl of the magnetization, with red representing clockwise and blue counterclockwise directions. (Reproduced from [25], with permission.)

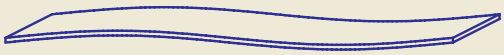

Figure 2

Scale diagram of a sinusoidal irregularity in a thin layer. The amplitude of the wave is 1.5 times the layer thickness (peak-to-peak is three times the thickness). (Adapted from [25], with permission.)

the models, we can investigate the effects of layer curvature and other three-dimensional irregularities on the switching properties of an MRAM bit [25].

The results obtained showed that layer curvature causes the magnetization to relax to an orientation that differs from the usual direction. Instead of aligning its dipoles along the long axis of the overall rectangle, a curved layer aligns its dipoles parallel to the long axis of any locally flat rectangle, even if that part is not parallel to the outside edges. For a sinusoidal curvature along the long dimension of a 2:1 rectangle (see **Figure 2**), the equilibrium alignment in the middle of the rectangle becomes slightly bent toward the short axis (**Figure 3**) because the layer is locally flat there, with a half-sized rectangle perpendicular to the main rectangle. The dipoles actually feel the new "edge" from the crest and trough of the sine function, and they try to align parallel to this new edge just as they did to the long axis in the original flat case. The reason for this alignment is again to minimize the external magnetic flux. When a layer is curved or when it has kinks, bumps, or other vertical protrusions, the field can leak out near the edges of the kinks in the vertical direction if it is not locally aligned to minimize this.

As a result of this curvature-induced anisotropy, layers with slight bends act as if there is a bias field that is applied in a direction parallel to the crests and troughs. The switching fields shift accordingly, in the direction of this bias, as shown in **Figure 4**. The axes of the figure are the two components of the external field. The solid line marks the field components at which a transition is made from no-switching at smaller fields to switching at larger fields when the layer is flat. The other lines indicate critical fields when there is curvature. The amount of curvature needed to influence the critical field can be very small. The cases shown have sine-function amplitudes $\Delta z/L$ that are only 0, 0.5, 1, and 1.5 times the thickness of the layer, which is 1/96 times its length. An analogous shift in critical fields occurs for layers that are curved along the short dimension; then the bias is along the long dimension, and the curves in the switching diagram move closer to the vertical axis. In both cases, the external field that is required to switch a curved MRAM bit is smaller in the direction of the crest or trough than when the layer is flat.

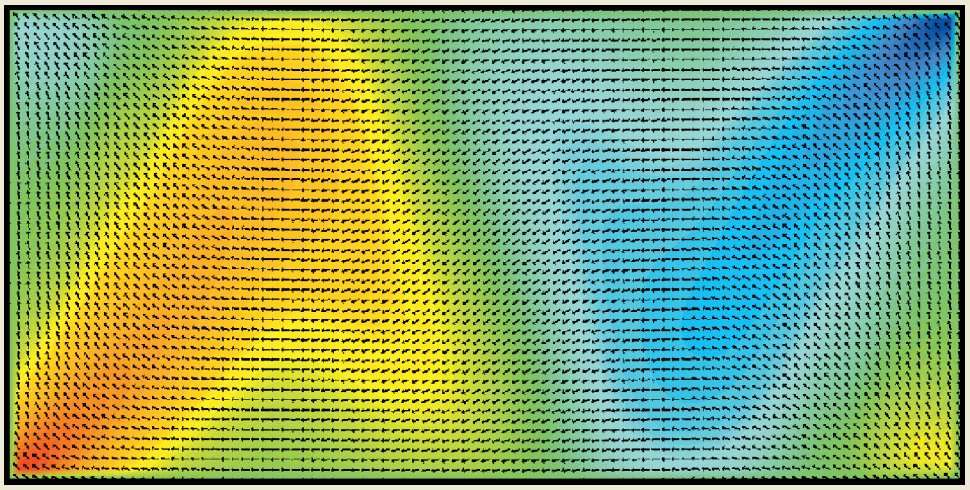

Figure 3

Dipole orientations in the equilibrium state for the layer shown in Figure 2. (Reproduced from [25], with permission.)

A second example of the use of MDGRAPE-2 in micromagnetics is for simulating a perpendicular write head in a computer disk drive. **Figure 5** shows the setup phase for a three-dimensional array of 70,046 dipoles that make a closed magnetic flux loop. The dipoles are shown as colored arrows. The yellow towers are initialized with dipole orientations that point downward on the left and upward on the right. The right tower is thicker than the left tower, so the field strength below the right tower is not high enough to switch the data dipoles below it. This way only the left tower can write a bit. The thin red layer below the towers is part of the moving disk, where the orientation of the towers is to be recorded; this data layer has a vertical anisotropy causing the dipoles to prefer to align up or down. The dipoles are red when they point up and blue when they point down; all of them point up initially. The thick blue layer below this is the soft underlayer, which connects the flux between the two towers. Initially all of the dipoles in it are aligned to the right, which is depicted in blue. Upon their switching to the left, their color would change to red.

Figure 6 shows the dipole orientations in the data layer and soft underlayer after 2.5 ns with the magnetization in the left tower pointing down and the magnetization in the right tower pointing up. The scale has been expanded for clarity. The dipoles below the left tower have flipped down, indicating that a bit has been written, and the soft underlayer points right to close the flux loop generated by the towers. **Figure 7** shows the dipoles after another 2.5 ns, this time with the opposite orientations in the towers followed by a shift in both the data layer and the soft underlayer. The data layer beneath the right tower does

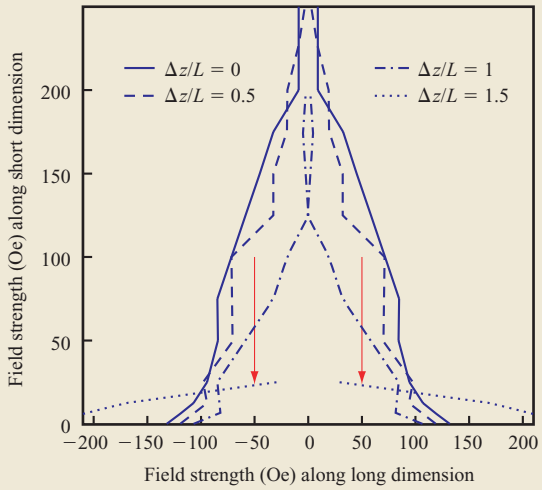

Figure 4

Critical switching field strengths for four cases of a layer having a sinusoidal curvature along the long direction. The amplitude of the curvature, measured in layer thicknesses, is given for each line type. As the curvature increases, the internal magnetization becomes more asymmetric, with a component along the short direction. This lowers the required switching field in this direction, and so lowers the curves in the figure, as indicated by the arrows. (Reproduced from [25], with permission.)

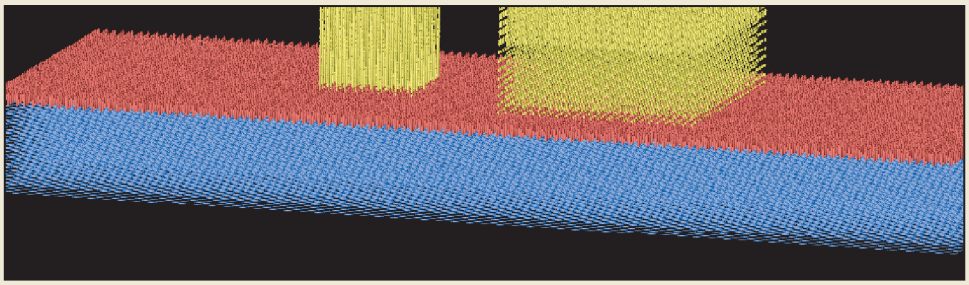

Figure 5

70046 dipoles shown as vectors for a magnetic disk head simulation. This first view shows the initial state: two towers (yellow) with vectors pointed down on the left and up on the right; a data layer with red vectors indicating the upward direction; and the soft underlayer with blue vectors indicating the rightward direction.

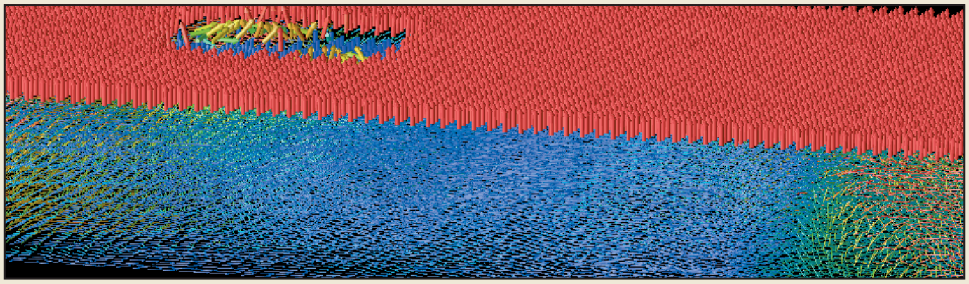

Figure 6

Blow-up of the data layer and soft underlayer after 2.5 ns, with the flux in a counterclockwise direction. The dipoles in the data layer flipped downward under the left tower, indicating that a bit of information was written. The soft underlayer is pointed mostly to the right between the towers.

not change because the magnetic flux concentration there is not large enough to overcome the local anisotropy field.

Simulations of the magnetization distribution around magnetic disk heads with MDGRAPE-2 are relatively easy to set up and modify because the dipoles can be positioned anywhere and each can have a different property. For example, the dipoles can be concentrated into a dense array near the data layer, where physical resolution is important, and placed more sparsely far from the data layer, where the magnetic field is relatively smooth. The uniaxial anisotropy can vary from dipole to dipole, particularly in the data layer, to represent granularity, and each anisotropy field can vary slightly with time to represent noise.

Disk head simulations were performed for a wide variety of material parameters and geometric shapes. Successful operation was found to depend sensitively on the assumed parameters and shapes. For example, if the anisotropy in the head was too weak, the dipoles aligned parallel to the bottom surface of the head and very little flux emerged for writing. If the anisotropy was too strong, the flux orientation in the magnetic loop was hard to switch during bit reversal. Similarly, the properties of the data layer and soft underlayer required fine tuning so that they could be written successfully by the head. The calculations indicated that a properly tuned combination of pole head, perpendicular data layer, and soft underlayer has the potential for writing bits at very high data densities.

Summary

This paper discusses two examples in areas of considerable technological importance that demonstrate the benefits of

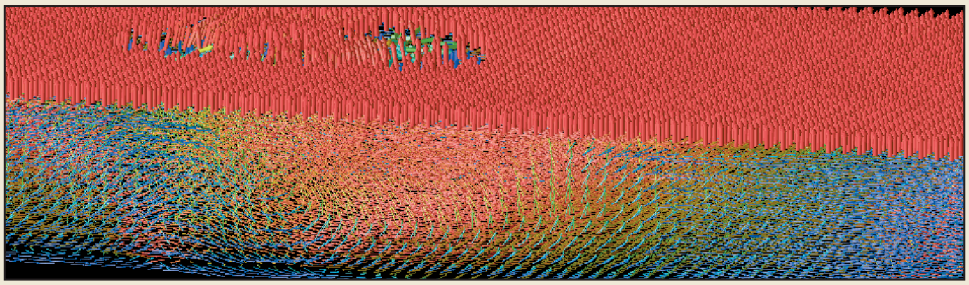

Figure 7

Data layer and soft underlayer after another 2.5 ns, following a switch in polarity of the towers. Now the magnetic flux is in the clockwise direction. The data layer under the left tower has mostly returned to the upright orientation, and the soft underlayer has switched to point leftward between the towers (turning red).

innovative GRAPE hardware and novel computational algorithms. The approach can easily handle the complexity inherent in magnetic materials and devices with long-range interactions. As a result, the spatial and dynamic features of magnetization processes can be simulated at the detail required for the optimization of high-performance magnetic designs. The versatility and ease of MDGRAPE-2 in designing complex physical systems, and the speed, accuracy, and scalability of its computations, make it an important new tool for solving complex micromagnetics problems. GRAPE hardware in general also shows considerable potential for other simulations involving long-range interactions, such as those in molecular biology and astrophysics. These advantages become even more compelling as the computational power of the GRAPE chip continues to increase along the same technology path as general-purpose processors.

Acknowledgment

We are grateful to Mr. Osahiko Ohno for his extensive contributions to the MDGRAPE-2 project.

*Trademark or registered trademark of International Business Machines Corporation.

References

1. E. D. Daniel, C. D. Mee, and M. H. Clark, *Magnetic Recording: The First Hundred Years*, IEEE Press, New York, 1998.
2. D. A. Thompson and J. S. Best, "The Future of Magnetic Data Storage Technology," *IBM J. Res. & Dev.* **44,** 311–322 (2000).
3. R. W. Wood, J. Miles, and T. Olson, "Recording Technologies for Terabit per Square Inch Systems," *IEEE Trans. Magn.* **38,** 1711–1718 (2002).
4. S. S. P. Parkin, K. P. Roche, M. G. Samant, P. M. Rice, R. B. Beyers, R. E. Scheuerlein, E. J. O'Sullivan, S. L. Brown, J. Bucchigano, D. W. Abraham, Y. Lu, M. Rooks, P. L. Trouilloud, R. A. Wanner, and W. J. Gallagher, "Exchange-Biased Magnetic Tunnel Junctions and Application to Nonvolatile Magnetic Random Access Memory," *J. Appl. Phys.* **85,** 5828–5833 (1999).
5. W. Reohr, H. Hoenigschmid, R. Robertazzi, D. Gogl, F. Pesavenot, S. Lammers, K. Lewis, C. Arndt, Y. Lu, H. Viehmann, R. Scheuerlein, L.-K. Wang, P. Trouilloud, S. Parkin, W. Gallagher, and G. Mueller, "Memories of Tomorrow," *IEEE Circuits & Devices* **18,** No. 5, 17–27 (2002).
6. J.-G. Zhu and H. N. Bertram, "Reversal Mechanisms and Domain Structures in Thin-Film Recording Media," *J. Appl. Phys.* **69,** 6084–6089 (1991).
7. R. H. Koch, J. G. Deak, D. W. Abraham, P. J. Trouilloud, R. A. Altman, Y. Lu, W. J. Gallagher, R. E. Scheuerlein, K. P. Roche, and S. S. P. Parkin, "Magnetization Reversal in Micron-Sized Magnetic Thin Films," *Phys. Rev. Lett.* **81,** 4512–4515 (1998).
8. M. E. Schabes, B. Lengsfield, and T. Schrefl, "Micromagnetic Modelling of Soft Underlayer Magnetization Processes and Fields in Perpendicular Magnetic Recording," *IEEE Trans. Magn.* **38,** 1670–1675 (2002).
9. D. Z. Bai and J.-G. Zhu, "Micromagnetics of Perpendicular Write Heads with Extremely Small Pole Tip Dimensions," *J. Appl. Phys.* **91,** 6833–6835 (2002).
10. X. Zhu and J.-G. Zhu, "On the Shape Optimization of Magnetic Random Access Memory Element Design," *J. Appl. Phys.* **93,** 8376–8378 (2003).
11. *Astrophysical Supercomputing Using Particle Simulations* (Proceedings of the 208th Symposium of the International Astronomical Union, Tokyo, Japan, July 10-13, 2001, J. Makino and P. Hut, Eds.); published in 2003 on behalf of IAU by the Astronomical Society of the Pacific, San Francisco.
12. J. Barnes and P. Hut, "A Hierarchical O(NlogN) Force-Calculation Algorithm," *Nature* **324,** 446–449 (1986).
13. J. Makino, "Treecode with a Special-Purpose Processor," *Pub. Astronom. Soc. Pacific* **43,** 621–638 (1991).
14. H. N. Bertram, *Theory of Magnetic Recording*, Cambridge University Press, Cambridge, England, 1994.
15. D. Sugimoto, Y. Chikada, J. Makino, T. Ito, T. Ebisuzaki, and M. Umemura, "A Special-Purpose Computer for Gravitational Many-Body Problems," *Nature* **345,** 33–35 (1990).
16. P. P. Ewald, "Calculation of Optic and Electrostatic Lattice Potential," *Ann. Phys.* **64,** 253–287 (1921).

17. T. Ebisuzaki, J. Makino, T. Fukushige, M. Taiji, D. Sugimoto, T. Ito, and S. K. Okumura, "GRAPE Project: An Overview," *Pub. Astronom. Soc. Jpn.* **45,** 269–278 (1993).
18. T. Fukushige, M. Taiji, J. Makino, T. Ebisuzaki, and D. Sugimoto, "A Highly Parallelized Special-Purpose Computer for Many-Body Simulations with an Arbitrary Central Force: MD-GRAPE," *Astrophys. J.* **468,** 51–61 (1996).
19. T. Narumi, R. Susukita, T. Ebisuzaki, G. McNiven, and B. Elmegreen, "Molecular Dynamics Machine: Special-Purpose Computer for Molecular Dynamics Simulations," *Molec. Sim.* **21,** 401–415 (1999).
20. R. Susukita, T. Ebisuzaki, B. G. Elmegreen, H. Furusawa, K. Kato, A. Kawai, Y. Kobayashi, T. Koishi, G. D. McNiven, T. Narumi, and K. Yasuoka, "Hardware Accelerator for Molecular Dynamics: MDGRAPE-2," *Computer Phys. Commun.* **155,** 115–131 (2003).
21. P. Hut and J. Makino, "Astrophysics on the GRAPE Family of Special-Purpose Computers," *Science* **283,** 501–505 (1999); see *http://grape.astron.s.u-tokyo.ac.jp/grape/*.
22. M. Mansuripur, "Magnetization Reversal Dynamics in the Media of Magneto-Optical Recording," *J. Appl. Phys.* **63,** 5809–5823 (1988).
23. J.-G. Zhu and H. N. Bertram, "Micromagnetic Studies of Thin Metallic Films," *J. Appl. Phys.* **63,** 3248–3253 (1988).
24. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, Cambridge, England, 1989.
25. B. G. Elmegreen, R. Koch, K. Yasuoka, H. Furusawa, T. Narumi, R. Susukita, and T. Ebisuzaki, "Numerical Simulations of Magnetic Materials with MD-GRAPE: Curvature Induced Anisotropy," *J. Magn. Magn. Mater.* **250,** 39–48 (2002).

Received August 5, 2003; accepted for publication November 3, 2003

Bruce G. Elmegreen *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (bge@watson.ibm.com).* Dr. Elmegreen received a B.S. degree from the University of Wisconsin in 1971 and a Ph.D. degree in astrophysics from Princeton University in 1975. He was a Junior Fellow at the Harvard University Society of Fellows until 1978 and then joined the faculty of the astronomy department at Columbia University in New York. In 1984 he joined the IBM Research Division, where he continued to work on star formation and galactic structure. In 1995 his astronomy connections with Dr. Toshikazu Ebisuzaki and others in the GRAPE team at RIKEN, Japan, led to a collaborative project between IBM and RIKEN to design and manufacture the MDGRAPE-2 chip. His subsequent research on micromagnetics began as a project to demonstrate the power of this device in novel fields with high potential impact in information technology.

Roger H. Koch *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (rhkoch@us.ibm.com).* Dr. Koch manages the Novel Physical Systems group in the Physical Sciences Department at the IBM Thomas J. Watson Research Center. He received an undergraduate degree in architecture from the Massachusetts Institute of Technology and a Ph.D. degree in physics from the University of California at Berkeley. In a career spanning twenty years with IBM, he has investigated the fundamental physics of superconductivity and magnetism, and has published more than 176 papers.

Manfred E. Schabes *Hitachi San Jose Research Center, 650 Harry Road, San Jose, California 95120 (manfred.schabes@hgst.com).* Dr. Schabes has been a Research Staff Member at the Research Division of Hitachi Global Storage Technologies since 2002. He received a Ph.D. degree in physics from the University of California at San Diego in 1989, and his research has been focused on the micromagnetic theory of nonuniform magnetization processes in magnetic nanograins and thin films. After assignments at Honeywell and Komag, he joined the IBM Research Division as a Research Staff Member in 1997 at the Almaden Research Center, where he worked on micromagnetic theory and simulations of novel magnetic recording media architectures. In 2002 Dr. Schabes received an IBM Outstanding Technical Achievement Award for his contributions to the invention, development, and implementation of antiferromagnetically coupled magnetic recording media. He has several patents and many publications in the field of micromagnetic theory and magnetic recording technology. Dr. Schabes is a member of the American Physical Society and the Institute of Electrical and Electronics Engineers.

Tiffany Crawford *Department of Electrical Engineering, Cory Hall, University of California, Berkeley, California 94720 (tiffcraw@EECS.Berkeley.EDU).* Ms. Crawford recently received a B.S. degree in electrical engineering from Howard University. She was a summer intern in 2002 at the IBM Thomas J. Watson Research Center, where she worked on modeling the magnetic write head. She is currently pursuing a Ph.D. degree in electrical engineering at the University of California at Berkeley, with an emphasis on digital signal processing. Ms. Crawford is a member of the Institute of Electrical and Electronics Engineers and the National Society of Black Engineers.

Toshikazu Ebisuzaki *Advanced Computing Center, Institute of Physical and Chemical Research (RIKEN), 2-1 Hirosawa, Wako, Japan 351-0198 (ebisu@postman.riken.go.jp).*
Dr. Ebisuzaki is the Director of the Advanced Computing Center at RIKEN. He received a B.S. degree in physics from Osaka University in 1981 and M.S. and Ph.D. degrees in astronomy from the University of Tokyo in 1983 and 1986, respectively. From 1986 to 1988 he was a Research Fellow of the National Research Council at the NASA Marshall Space Flight Center. Dr. Ebisuzaki subsequently joined Kobe University as an Assistant Professor. He and his collaborators began developing GRAPE special-purpose computers for gravitational *n*-body simulations, such as galaxies and star clusters. In 1995 he joined RIKEN and initiated a project to develop the Molecular Dynamics Machine, a special-purpose computer for molecular dynamics simulation. Dr. Ebisuzaki was awarded the 2000 Gordon Bell Prize of the IEEE Computer Society.

Hideaki Furusawa *Computer Information Division, Institute of Physical and Chemical Research (RIKEN), 2-1 Hirosawa, Wako, Japan 351-0198 (furusawa@riken.go.jp).*
Mr. Furusawa was a technical staff member of the Advanced Computer Center of the Computational Science Division at RIKEN during most of the development work on MDGRAPE-2. He graduated from a Tokyo liberal arts university and received a master's degree in education in 1996. He is currently a member of the support staff of the Computer Information Division.

Tetsu Narumi *Genomic Sciences Center, Institute of Physical and Chemical Research, (RIKEN), E503 214 Maeda, Totsuka, Yokohama, Japan 244-0804 (narumi@gsc.riken.go.jp).*
Dr. Narumi is a Research Scientist in the High Performance Biocomputing Research Team of the Bioinformatics Group at the RIKEN Genomic Sciences Center. He received B.S., M.S., and Ph.D. degrees in philosophy from the University of Tokyo in 1993, 1995, and 1998, respectively. He subsequently joined RIKEN, where he worked on a special-purpose computer for molecular dynamics simulations. In 2000, Dr. Narumi received the Gordon Bell Prize of the IEEE Computer Society for molecular dynamics simulation. He is the first author of seven technical papers, and a member of the Institute of Electrical and Electronics Engineers.

Ryutaro Susukita *Computational Science Division, Advanced Computing Center, Institute of Physical and Chemical Research (RIKEN), 2-1 Hirosawa, Wako, Japan 351-0198 (susukita@postman.riken.go.jp).*
Dr. Susukita received a Ph.D. degree in science from Kyoto University in 1997 and then joined RIKEN to work on the MDGRAPE-2 system. He is also working on computational techniques for large-scale simulations in physics and astronomy using MDGRAPE-2. Dr. Susukita is a member of the Japanese Physical Society.

Kenji Yasuoka *Department of Mechanical Engineering, Keio University, Yokohama, Japan 223-8522 (yasuoka@mech.keio.ac.jp).*
Dr. Yasuoka is an Assistant Professor at Keio University. He received a B.S. degree in physics and M.S. and Ph.D. degrees in applied physics from Nagoya University in 1992, 1994, and 1997, respectively. He subsequently joined RIKEN as a postdoctoral researcher, moving to Keio University in 1998 as a faculty member. He is an author or coauthor of 24 technical papers on molecular dynamics simulation of phase-change phenomena, including evaporation and condensation. Dr. Yasuoka is a member of the American Physical Society, the Japanese Physical Society, the Japan Society of Mechanical Engineers, the Heat Transfer Society of Japan, and the Japan Society of Thermophysical Properties.