

NONVOLATILE RAMS BUILT WITH THIN FILMS OF FERROMAGNETIC MATERIAL ARE POISED TO CHALLENGE DYNAMIC AND NONVOLATILE MEMORIES BASED ON CONVENTIONAL SEMICONDUCTORS


Magneto-electronic memories last and last...

MARK JOHNSON
Naval Research
Laboratory


THE EMERGING FIELD OF MAGNETOELECTRONICS promises huge enhancements of the speed, reliability, and power consumption of solid-state memory. At the root of the excitement are electronic devices with a ferromagnetic component that lets them not only switch between two stable states in one clock cycle but also retain the state they are in when power is removed.

The compact design and simple operation of magneto-electronic devices make them ideal storage cells for fast random-access memory (RAM) [Fig. 1]. Prototypes improve by several orders of magnitude on the speed, power, and reliability of semiconducting floating-gate nonvolatile memory and have achieved cell sizes and operating speeds that are competitive with dynamic (volatile) RAM [see "Challenging the establishment," p. 35]. Further off in the future is the intriguing prospect of instantaneously reprogrammable logic.

Magneto-electronic devices may be grouped according to the physics of their operation into three main categories: hybrid ferromagnet-



[1] The 1Mb GMRAM, a magneto-electronic memory chip under development at Honeywell Corp., uses a type of spin valve based on so-called giant magnetoresistance. Production is slated for 2001.



[2] A thin ferromagnetic film [left] has bistable magnetization states, oriented right or left and described by a hysteresis loop [right]. An external magnetic field H , if larger than switching field H_s , sets the orientation to the right, if less than $-H_s$, sets it to the left, or if removed, leaves the magnetization in its last orientation.

In an integrated device, a current pulse through an integrated, contiguous write wire that is directly over the ferromagnetic element generates magnetic field H parallel to the wire's plane and close to its surface. The magnitude of H is αI_{write} , with α the inductive coefficient.

semiconductor structures, magnetic tunnel junctions, and all-metal spin transistors and spin valves (details to come). The process of writing data into a cell is essentially the same for any of these approaches. Reading a cell's contents, however, may utilize any of several physical mechanisms, each in some way dependent on the direction of the magnetization in the ferromagnetic element. In fact, a distinguishing characteristic of the three device categories is their readout mechanisms. Also different for each type are its prospects for becoming high-density, low-power, high-speed nonvolatile RAM, as well as challenging issues in development and manufacturing.

BASIC PRINCIPLES

The bistable orientation of their magnetic state is a defining characteristic of ferromagnetic materials and a natural basis for nonvolatile bit storage. For a properly fabricated thin ferromagnetic film the two possible states of magnetization can be described by a hysteresis loop [Fig. 2, left]. The film may vary in thickness from about 1 nm to 100 nm, and be as narrow across as 10 nm. Typically the material is a transition metal, such as nickel, iron, cobalt, or an alloy of these materials.

With proper preparation, the film will preferentially form a single magnetic domain with a single magnetization axis. An instance would be a rectangular shape with an aspect ratio of 4:1 or 5:1 promoting magnetization along the long axis. Such anisotropy may in addition be induced during film growth by, say, applying an external magnetic field to the substrate to orient the film's crystallites and thereby create a preferred crystallographic magnetic orientation. An appropriate element has two stable magnetization states, pointing in opposite directions, with

saturated magnetization values. A saturated value corresponds to an internal mean magnetic field of about 1 tesla.

The magnetization of the film as a function of magnetic field follows a hysteresis loop [Fig. 2, right]. A magnetic field applied along the axis of magnetization sets the magnetization to the saturated value in the direction of the applied field when the applied field is greater than the switching field—this last being the minimum field required to switch the film's magnetization state. In ferromagnetic films used in memory cells, the switching field is typically 2–6 kA/m, which is adequate for stability in ambient magnetic fields.

The nonvolatility natural to magnetic storage is reflected in the square shape of the hysteresis loop. When the field is removed, the magnetization retains its orientation and will do so for many years afterward.

At the heart of a random-access memory is an integrated method of addressing the element in order to write a bit into it or read a bit out. In magnetoelectronic memories, writing is done by a wire fabricated directly over, and inductively coupled to, the magnetic element. A current pulse traveling down the wire generates a magnetic field parallel to the wire's plane and close to its surface. This write current is chosen so that it couples a field greater than the switching field into the element and is therefore adequate to set its state to a binary 1 or 0. To write the opposite bit, the current is reversed, requiring this scheme to have a bipolar current source.

For a two-dimensional array of memory cells, a 2-D array of rows and columns of write wires is employed, with each cell inductively coupled to a write wire from one row and another from one column. When a write current of half the value needed to switch a cell is applied to one row, and equal

current is applied to one column, the full current addresses a single cell in the array, leaving the others along that row and column unchanged. This half-select process requires each cell's hysteresis loop to be square enough that applying and removing half the switching field leaves the element in its initial state. The possibility of crosstalk during the write process is negligible because the magnetic fields from the write currents are localized to regions under each write wire. The switching time for transition metal ferromagnetic-film elements is a few nanoseconds, and the write current for an element with dimensions smaller than a micron is a few milliamperes.

There are other requirements. If a magnetoelectronic device is to succeed within electronics, it must be capable of being integrated with other systems technologies. In particular, a successful magnetoelectronic memory cell should be based on a device with impedance of about 1–10 k Ω . Readout voltage discrimination should be 20–40 mV for compatibility with existing sense amplifier circuitry. Device fabrication must be compatible with standard CMOS processing steps. The device should be scalable down to a minimum feature size f of 100 nm, and should be amenable to vertical (as well as planar) cell designs. To compete with nonvolatile RAM, the cell size should be no larger than $16 f^2$, and to compete with dynamic RAM, no larger than $8 f^2$.

Consider how well the three types of magnetoelectronic devices fulfill these requirements. Drawbacks of all-metal spin transistors and spin valves are their low-impedance and relatively low levels of readout voltage. The magnetic-tunnel junctions are high-impedance devices with good readout voltage levels but strict fabrication requirements that cause problems with operating margins

Challenging the establishment

A new technology will never win adoption if it improves only slightly on an approach already embodied in an established product. At the very least, it must promise an order of magnitude improvement in one or more valuable performance categories.

The floating-gate approach to nonvolatile semiconductor RAM is based on either the electrically erasable programmable ROM (EEPROM) device or the flash-EEPROM device. Such devices use hot-carrier injection at a relatively high voltage to write a bit into a memory cell, plus Fowler-Nordheim tunneling at a high reverse bias to remove charge from the floating gate, erasing the bit. An EEPROM memory cell has two transistors and an architecture that permits data to be erased a byte at a time. The flash variety has a single-transistor cell with a simpler architecture that requires erasure by sector.

Floating-gate nonvolatile RAM technology is highly refined and deeply entrenched and its US \$5 billion-per-year market is still growing rapidly. All the same, the characteristics of an ideal memory are well known, and by that standard semiconductor nonvolatile RAM fares poorly in many regards.

Nonvolatile RAMs do not excel in some of the categories articulated by W.D. Brown and J.E. Brewer in *Nonvolatile Semiconductor Memory Technology* (IEEE Press, New York, 1997). Examples are fast read/write/erase times, good endurance (high number of write/erase cycles), and low power consumption. Further, though flash-EEPROM memory does well on several counts because its rather high density and bit count come at a fairly low cost, it cannot be altered one bit at a time. EEPROM memory achieves bit alterability, but at the expense of low density and bit count at a rather high cost. And both types of nonvolatile memory are less than rugged, being susceptible to radiation damage of the gate insulator.

Thus, despite the success of semiconductor nonvolatile RAM, it is vulnerable

in roughly half the categories that are characteristic of an ideal memory.

Ferroelectric RAM technology is based on a bistable ferroelectric capacitor and offers an alternative for nonvolatile RAM technology. But after two decades of incremental improvement, ferroelectric RAM still suffers from weak performance. The large cell entails low density and bit count (the largest memory is 256Kb) and high cost. Read times of a few hundred nanoseconds are greater than those of traditional nonvolatile RAM. Some cells have limited retention, and endurance is typically 10^{13} cycles or less. These problems have limited ferroelectric RAM to a mere 1 percent of the nonvolatile RAM market.

The magnetoelectronic approach to integrated nonvolatile RAM is relatively new and represents a paradigm shift for memory. The physical principles underlying the writing, reading, and storing of a bit of information are completely different from the semiconductor approach. The established paradigm uses a capacitance for storage and capacitively coupled voltage pulses for writing and reading. But the magnetoelectronic approach turns to the bistability of a magnetization state for storage, inductively coupled current pulses for writing, and voltage sensing for reading.

The vulnerability of semiconductor nonvolatile RAM stems from the very nature of capacitance. Between the control gate and the floating metal gate, which sets the device state to one of two capacitive values, the coulomb charge repulsion is so strong that the two gates must be isolated by a thick dielectric layer to keep leakage to a minimum. As a consequence, the time required to add or remove charge through the oxide barrier is extremely long and the process gradually damages the barrier. Despite decades of incremental improvement, write times vary from 10 μ s to 10 ms, and erase times vary from 10 ms to 1 s. Writing and erasing require a relatively high bias of 10 V, which of course consumes extra power.

To compare power consumption of different kinds of devices, the power for a read or write operation can be taken as the product of instantaneous voltage

and current required for operation on a single device. Since all relevant operations are performed serially, this also defines the steady-state power for operation on a sector. For a chip with four or eight sectors, operating power is four or eight times as large.

For any nonvolatile memory, the quiescent power is, of course, zero. Floating-gate write and erase power consumption can be several hundred milliwatts, while reading power can be up to 100 mW. The damage caused by these processes prevents floating-gate cells from lasting longer than about 10^5 write cycles. Magnetoelectronic nonvolatile RAM, though still in its infancy, has demonstrated write and rewrite (erase) times of 10 ns or less, read and write powers of 1–10 mW, and infinite endurance—in other words write/rewrite cycles much greater than 10^{15} . It therefore outperforms semiconductor nonvolatile RAM by two or three orders of magnitude in the crucial operating categories of speed, power consumption, and durability.

Of potentially greater importance is the competition with dynamic RAM that may be on the horizon. A highly successful dynamic RAM, when evaluated against an ideal memory, fares extremely well in nearly all categories but fails in the two most important: It loses stored charge in an instant when the power is turned off. When powered on, it must always, even in the quiescent state, draw high power of about 100 mW in order to continuously refresh the data in the cells.

Existing magnetoelectronic prototypes are less dense than dynamic RAM, having cells that are roughly twice as large. But if magnetoelectronic technology becomes established, enhancements in cell design may yield a superior nonvolatile RAM offering nonvolatility plus order of magnitude improvement in write/read times and power consumption with little or no loss of density. Because it holds data fast, magnetoelectronic RAM can combine the functions of dynamic RAM and direct-access nonvolatile storage, eliminate magnetic disk drives, and thereby offer cost and performance benefits for many computing systems. —M.J.

and yields. Both approaches incorporate a MOSFET in the memory cell for isolation.

And here's the rub—fabrication procedures for multilayer magnetic elements and standard CMOS processing are incompatible. Making magnetic tunnel junctions and spin valves involves dry etching by an argon ion mill or else a reactive ion etch, techniques that damage semiconductor structures. In addition, standard CMOS requires

temperatures as high as 800 °C, or at least 400 °C, so as to anneal ion damage. Unfortunately, magnetic multilayers are unstable at temperatures higher than 200 ° or 300 °C.

From the perspective of manufacturability, the most promising choice is the hybrid ferromagnet-semiconductor device. It not only offers good readout voltage levels but, of greater importance, it also involves a single ferromagnetic layer that is electrically


isolated from the rest of the device. During processing of the magnetic layer, the device structure can be protected by insulating layers, say, and the single ferromagnetic layer tolerates the high temperatures of CMOS processing steps.

THE FINALIST


The hybrid ferromagnet-semiconductor device uses the magnetic fringe field of its

[3] A ferromagnetic element generates locally strong magnetic fringe fields [top]. The hybrid ferromagnet-semiconductor gate [bottom] uses the fields to exert a Hall effect on charge-carriers in a semiconductor channel with voltage bias V_b . The effect can be sensed in a transverse arm of the semiconductor as a Hall voltage. When the magnetization of the ferromagnetic element reverses orientation, the signs of the fringe fields and Hall voltage also reverse.

The readout voltage tells a stored 0 from a 1 by picking up a difference of twice the Hall voltage measured on a background of $V_b/2$. The channel is heavily doped [dark blue] between the right transverse arm and ground so that the voltage drop occurs mostly between the bias and readout voltages.



[4] A typical hybrid ferromagnet-semiconductor nonvolatile memory uses a half-select technique to uniquely address each of the cells in its array. It utilizes diagonal bit and word lines as write wires. Each line selected carries half the write current, so that only the selected cell gets the full write current. A Schottky diode isolates the semiconducting channel from the array.



single ferromagnetic element for the readout operation. The mean internal field of a single-domain ferromagnetic element leaks out of the ends [Fig. 3, top]. The direction of the external magnetic field is outward at one end and inward at the opposite end. These external fields, generated by the spontaneous magnetization of the ferromagnetic material, are termed demagnetizing because they can affect the state of the ferromagnetic element.

Contributions to the total energy of the element come from the dipolar energy of the magnetization in its own demagnetizing field and from the domain wall energy where two regions of different magnetization orientation meet. The magnetization will tend to orient itself in such a way as to minimize its total energy. In many cases, it splits up into several magnetic domains, with a different magnetic orientation in each, and in a pattern that minimizes external fields. But a properly designed element has a single magnetic domain.

The mean internal field is about 1 or 2 T, as is the external fringing field at a few nanometers from the element's end surface

(though with increasing distance its strength diminishes rapidly). When one end of the element is directly over a semiconducting channel, the fringe field has a large component perpendicular to the plane of the channel. When the channel is voltage biased, the field causes a Hall effect—the name for the voltage that develops across the channel and perpendicular to the current. The Hall effect is due to the Lorentz force exerted by the magnetic field on the charge-carriers. The force is proportional to the product of the field's perpendicular component and the carrier velocity, and it pushes the carriers sideways, in a transverse deflection [curved arrows in Fig. 3]. This shift can be sensed in a transverse arm as a Hall voltage. When the magnetization of the element has the opposite orientation, the signs of the fringe fields and the Hall voltage are reversed.

In a typical planar memory cell, the bit and word lines also function as the two integrated write wires for the half-select operations [Fig. 4]. The maneuver increases packing density and reduces the number of mask levels.

Research on this type of device is being conducted at Honeywell Corp., Plymouth, Minn., with work ongoing at the U.S. Naval Research Laboratory, Washington D.C. as well. Prototypes fabricated with minimum feature sizes of 0.5 μm and 1.0 μm have demonstrated excellent performance. The write/read cycle time of individual prototypes has been measured to be 10 ns or less, with readout voltage discrimination of 40 mV and readout power of a few milliwatts. Planar twin cells, which share an element such as a contact, have been built with an area of 12 μ^2 , and vertical cells with a ferromagnetic layer perpendicular to the substrate have been designed with an area of 7 μ^2 .

An important benefit of the hybrid ferromagnet-semiconductor device is inverse scalability. As the device is shrunk, the magnitude of the magnetic fringe fields that drive the output voltage increases. At the same time, the switching field can be kept constant. Accordingly, the current per unit width (current density) in the write wire stays the same, but the total current shrinks as the write wire shrinks. Write currents can be reduced

to the order of 1 mA, and write powers (assuming 50- Ω metallized write lines) are reduced to the order of tens of microwatts.

A second factor that favors development is that device parameters can be controlled by doping. Reasonable operating margins with high yields should be achievable.

Several issues about this device are unresolved, mostly because of its early state of development—only a few person-years of research have been invested to date. The device concept was derived from radiation-hard electronics research based on compound semiconductors, so that early prototype work has been restricted to III-V semiconductors such as gallium arsenide. The switch to silicon is not expected to degrade performance significantly. Silicon devices have been designed for fabrication with silicon-on-insulator or bulk silicon wafers using CMOS technology but the actual fabrication and demonstration of acceptable silicon-based prototypes are steps that remain to be taken. Further, the ferromagnetic elements are thick enough—typically 40 nm—for there to be a danger of magnetic cross talk between cells: the fringe field of the element of one cell may alter the magnetization state of the element in a neighboring cell. This has not been a concern in prototypes with 1- μ m minimum features, but is likely to be a problem at feature sizes below 0.5 μ m. The solution is to fabricate a second ferromagnetic element as a flux keeper, but prototypes of this kind have not yet been tested.

MAGNETIC TUNNEL JUNCTIONS

A magnetic tunnel junction (MTJ) is a sandwich of two thin ferromagnetic films separated by a very thin dielectric tunnel barrier [Fig. 5]. The device is magnetoresistive, which means that the tunneling resistance from film to film depends on the magnetic states of the two films. Ideally, each film is a single domain, and when the magnetizations of both films are parallel, the resistance is lower than when the magnetizations are antiparallel. This effect is known as spin-dependent tunneling. The relative difference of resistances—the difference between the two resistance values divided by the lower value—depends on voltage bias across the barrier and is roughly 12 percent for typical values of bias voltage.

The bistable magnetization states of the magnetic tunnel junction, with the magnetizations of the two ferromagnetic layers parallel or antiparallel, form the basis of non-volatile storage, and a variety of designs have been developed for magnetic tunnel junction memory cells. In a typical prototype cell, the word line serves as a write wire for setting the state of the cell and also for interrogating cell contents. The bottom ferromagnetic film is constructed to have a larger value of switching field than the top film. The bit is stored as the orientation of the bottom film, and is written using a write pulse of relatively large amplitude. (Cells in a two-dimensional array have two write wires for half-select processes, and are slightly more complicated.)

Reading out the cell contents is done by interrogating the cell with a sequence of smaller-amplitude write pulses. While the resistance of the cell is monitored by sensing either voltage or current, a sequence consisting of first a positive and then a negative current pulse is transmitted to the word line. The associated magnetic fields are larger than the switching fields of the upper magnetic film, which is set first to the left and then right. For a stored binary 0 this corresponds to a sequence of antiparallel and parallel orientations, sensed as a relatively high voltage followed by a relatively low voltage. For a stored 1, the sensed response is reversed.


IBM Corp. and Motorola Inc. are leading the research on magnetic tunnel junction-based RAM. A typical memory cell includes a single transistor for isolation. IBM recently demonstrated results from a simple prototype array with a cell area of 12 μ^2 and a minimum feature size of 0.5 μ m, a device impedance of about 1 k Ω , read times of 10 ns or less, and write/read cycle times of about 30 ns. Write currents were the order of a few milliamperes, and power consumption for both write and read operations was the order of 1 mW.

Although magnetic tunnel junction prototypes have performed on a par with the hybrid ferromagnet-semiconductor device, the technology must overcome a number of barriers. Tunnel junctions scale poorly, and controlling the margins of device impedance will be crucial. To achieve a device resistance of about 1 k Ω for prototypes with 0.5- μ m

[5] A typical magnetic tunnel junction memory device has a trilayer structure in which two ferromagnets are separated by a non-magnetic metal layer.

Ferromagnet No. 1 has the larger switching field of the two, and its magnetization orientation stores the bit. The contents of the device are read out by manipulating the magnetization of ferromagnet No. 2 while monitoring the resulting change in voltage (or current).

This kind of memory cell generally relies on a single transistor for isolation [not shown] and has two write wires that permit half-select processes.



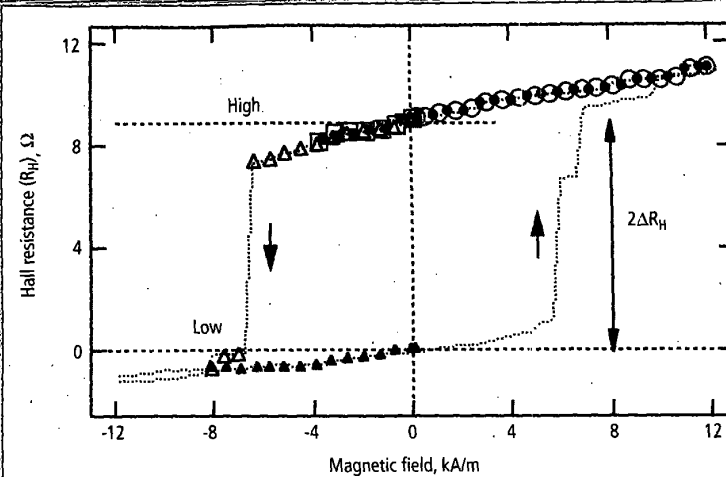
Dynamics of magnetoelectronic logic

The bistability of the magnetization states of a ferromagnetic element has applications besides nonvolatile memory. After all, the half-select process used to address a single cell in a two-dimensional array is a Boolean AND operation. More generally, the hybrid ferromagnet-semiconductor device can be viewed as a latching Boolean gate, with application as instantaneously reprogrammable logic.

Programmable logic is an emerging technology with a market of about US \$2.5 billion per year. Today, most of these chips can be programmed only once, and reprogrammable chips need anywhere from a millisecond to a full second to reconfigure high-level blocks of hard-wired devices for special functions. (In the main, transistor-based logic families, CMOS and TTL, a single Boolean gate is formed by hard-wiring a particular combination of several transistors.) In contrast, the magnetoelectronic latching Boolean gate offers new and unique speed and flexibility: logic operations can be continuously programmed and reprogrammed on the time scale of a single clock cycle and at the most basic cellular level, that of a single Boolean process.

For digital logic, the magnetoelectronic device's possible states must correlate with logic low and high (binary 0 and 1) electronic states, to match the low and high levels of transmitted voltage or current pulses. For an appropriately fabricated hybrid ferromagnet-semiconductor device, the low and high output states take the convenient values of zero and twice the Hall voltage.

The quasistatic output characteristic of such a device, measured with a steady-state bias current and an externally applied field H_x , is shown with dotted lines in the figure above. The switching field of the device is 7.5 kA/m. The nonvolatile, nonpowered condition is at zero



field, where the bistable device states correspond to the two remanent magnetization states with resistance levels 0.0 Ω (low) and 8.5 Ω (high).

To aid an understanding of the loop, individual processes were traced out with limited field sweeps (the open symbols indicate increasing, the closed symbols indicate decreasing field magnitude). Starting at the high state, when the externally applied field is swept to 12 kA/m and back to 0 (green circles), the resistance response retraces the upper portion of the loop and returns to high.

In a similar fashion, a sweep from 0 to -4 kA/m and back (red squares) results in a reversible trace that returns to high. However, when the field is swept from 0 to -8 kA/m and back again (brown triangles), a hysteretic loop is traced and the value of the resistance across the device is set to low.

These magnetic maneuvers can also

be achieved by applying current pulses to an integrated write wire; the result being a Boolean logic operation with schematic opposite (top of figure, right) represents the prototype device, whose hysteresis loop is shown above.

Binary input pulses applied simultaneously to input terminals A and B, or to control terminal C, are transmitted down a write wire. This is demonstrated using individual pulses (green in figure at right). The device is biased with a dc current and the readout (pink lines at right) is recorded as current pulses are applied to the input terminal of the write wire.

The sequence of pulses for a Boolean NAND operation is: apply a reset pulse of amplitude 12 kA/m to terminal C to set the device's initial state to high. A pulse of 4 kA/m is insufficient to switch the state, so it remains high. After another reset pulse, a pulse of -8 kA/m traverses the loop and switches the device to low. If the 4 kA/m

minimum features, the thickness of the aluminum oxide barrier was reduced to 0.7 nm. Tunneling resistance varies exponentially with barrier thickness, and a 0.1-nm variation changes device impedance by a factor of 10. Because of these variations, high yields with reasonable operating margins will be hard to achieve across production-sized wafers and from one wafer to the next. By way of comparison, digital superconducting electronics employs the same aluminum oxide technology to fabricate slightly thicker superconducting tunnel junction switches. After more than two decades of development, the state of the art for the industry is a minimum feature size, limited in part by junction margins, of 2 μm and a maximum

density of 40 000 junctions per chip. It will be a challenge for magnetic tunnel junctions to achieve substantially higher bit-counts—say, three orders of magnitude larger—along with minimum feature sizes an order of magnitude smaller. And the smaller the feature size, the worse the problem. This challenge would be even greater for vertical cells, and no vertical design has yet been proposed.

Also at issue is the reproducibility of the structure's magnetic states. The two planar ferromagnetic layers are fabricated in such close proximity, separated by about 1 nm, that the magnetization of one element can become coupled to that of the second element. Naturally, the magnetization state of the coupled system is more complicated than

that of a single element, and prototype cells have shown a wide spread in the value of the switching field. In some devices the repeated switching of the top ferromagnetic layer has gradually demagnetized the bottom layer, setting a limit on device durability and endurance. Furthermore, as mentioned above, magnetic tunnel junction fabrication and CMOS processing are not compatible. Prototypes have been fabricated with back-end processing, but final annealing to repair damage to CMOS structures frequently damages the magnetic tunnel junctions.

ALL-METAL DEVICES

In the third, all-metal, category, the storage cell is usually based on an all-metal mag-

pulse is identified as the unit write current associated with the binary 1, the sequence is seen to be a NAND.

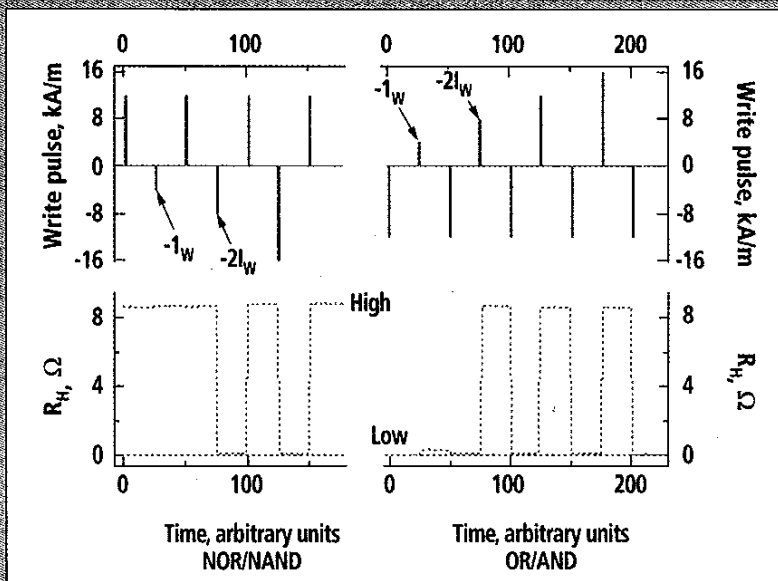
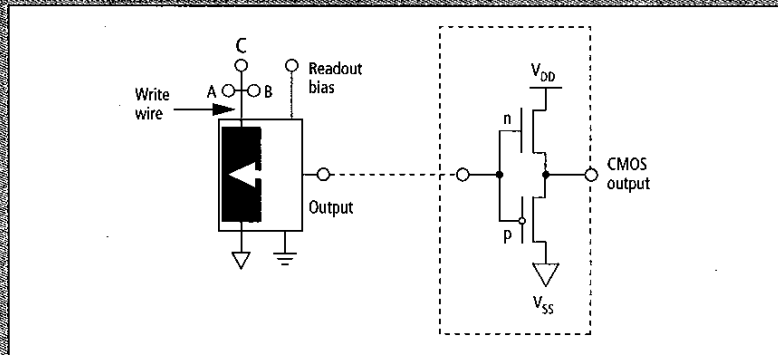
Note that a single unit current pulse (1) applied to either terminal is too weak to switch the device, whose state remains high. But two unit current pulses applied simultaneously to terminals A and B, with net magnitude of -8 kA/m will switch the device state to low.

For a Boolean NOR operation, a reset pulse again sets the device state to high, whereupon a -16 kA/m pulse traverses the loop and the device switches to low, identifying an -8 kA/m pulse as a unit write current reveals this process as a NOR. A current pulse of unit amplitude applied to either terminal, or to both terminals, changes the device state to low. As for the inverse operations, AND and OR result from changing the polarity of the input write pulses.

Remarkably, this single magnetoelectronic device can perform any of the four Boolean operations in two clock cycles, reset and input, with the result latched for readout at any later time. The function to be performed is determined instantaneously by the way the device is addressed: the normalized value of the write current pulse has one of two values and either of two polarities.

An equivalent mode of operation is even simpler. A control pulse applied to terminal C, simultaneously with the input pulses applied to A and B, determines the function.

To amplify, if no pulse is applied to C, unit pulses of write current at A and B are required to switch the state of this device, which operates as an AND gate. If a unit amplitude pulse of write current is applied to C, then a single unit pulse at either A or B is necessary and sufficient to switch the device state. Changing the



polarity of the control pulses at terminal C sets the device function to be a NOR or NAND gate. In this mode, it is possible to program a single pulse to control the function.

Once again, any of the four Boolean functions can be performed in only two

clock cycles. Combine memory and logic capabilities, and large-scale on-chip integration is automatic. The prospect beckons of versatile high-density arrays of high-speed devices, individually programmable for a particular logic or storage function. —M.J.

netoresistor. This device resembles the magnetic tunnel junction prototype sketched in Fig. 5, but instead of a thin insulating layer, a layer of a nonmagnetic metal (such as copper) a few nanometers thick separates the two ferromagnetic layers. The nonmagnetic metal mediates a coupling of the two ferromagnetic films' magnetization states, which are sensitive to variations of a few tenths of nanometer in the thickness of the nonmagnetic layer. Trilayer magnetoresistors with certain forms of magnetic coupling are called spin valves, and this term is often applied to all multilayer magnetoresistors.

The operation of the magnetoresistor and that of the magnetic tunnel junction are closely related. A sense current flows through

the sandwich along the plane of the films, and the structure's resistance is relatively low when the magnetizations are parallel and relatively high when they are antiparallel. With this device, the effects of spin-polarized transport are smaller than in magnetic tunnel junctions, and typical values of relative resistance difference are 6–8 percent.

There are several storage cell designs for magnetoresistors, but typical cell design and operation are analogous with the tunnel junction cell design and operation. A bit is stored as the orientation of the lower ferromagnetic layer, and readout involves interrogating the cell with a sequence of current pulses that modulate the relative magnetization orientations between parallel and

antiparallel while the output is measured. As a consequence of their all-metal composition, magnetoresistors are low-impedance devices. A typical device has an impedance of roughly 100Ω and, given a relative resistance difference of 6 percent, requires a current of 3 mA for readout discrimination levels of 20 mV. Cells with a single isolation transistor are plausible, but fabricating a small transistor that can carry high currents is difficult. A more common approach is to string a number of magnetoresistors on a single line and address all the cells in a row at once. While this leads to higher packing densities, the penalty is diminished signal-to-noise, longer read times, and higher read power.

Motorola's Corporate Research Center in

Tempe, Ariz., and Honeywell Corp., Plymouth, Minn., have been leading the research efforts for this kind of magnetic RAM. Prototype cells have been made with minimum feature sizes of 0.5 and 0.25 μm , an area of 12 μm^2 , and readout discrimination of roughly 10 mV. Write currents were about 10 mA, corresponding to 5 mW for 50- Ω write lines. Currents in the sense lines were about 3 mA, and the power consumption during read operations, a few milliwatts, was dominated by the power dissipated in the write lines while interrogating the cell. The low readout voltage level necessitates the use of a reference resistor in some prototype designs, thereby doubling the cell size. Since the write/read cycle time is slow (about 100 ns), spin valve prototypes have yet to achieve acceptable performance levels.

THE OLDEST APPROACH

All-metal magnetoresistive RAMs have enjoyed the longest history of research and development. Low density (16 Kb) radiation-hard memories based on a simpler cell are manufactured and sold by Honeywell. These cells use a property of individual ferromagnetic layers called anisotropic magnetoresistance, which gives relative resistance differences of 2 to 3 percent for magnetization orientations parallel or perpendicular to the sense current. They are more robust to the conditions of manufacturing processes than spin-valve multilayers. But the lower magnetoresistance ratio means low readout voltages, and the read times of 300 ns are slow. Many of the obstacles to be overcome to achieve high density and performance relate directly to their intrinsically low impedance and hence low output voltage. The sense currents needed to achieve 10-mV readout levels correspond to huge current densities on the order of 10^7 to 10^8 A/cm 2 .

The unfortunate consequence is electromigration, the mass transport of atoms in metal films that results in voids and the diffusion of material across ferromagnetic and nonmagnetic metal interfaces. Since the value of relative resistance difference is sensitive to the detailed nature of these interfaces, electromigration can damage the cell irreversibly. The relatively high levels of operating power can also cause local heating and hence changes in resistivity, which further degrade operating margins.

Another source of problems is that the two ferromagnetic layers of spin valves are magnetically coupled through the thin nonmagnetic layer. The magnetization states of the bilayer, which are sensitive to 0.1-nm variations in thickness, are far more complicated than the states of a single element. Margins for both the switching field and the magnetoresistive coefficient are hard to control, and yields have been low. Although several prototype arrays have shown no degradation for 10^{15} write/read cycles, the

demagnetization of one ferromagnetic layer, resulting from modulation of the magnetization of the other, is an endurance problem for some cell designs. Because the bilayer structures are not thermally stable at temperatures of 200 ° to 250 °C and higher, prototype fabrication by back-end processing has run into annealing problems.

Furthermore, spin valves scale poorly. The thickness of the sandwich structure is determined by magnetic coupling constraints, and the aspect ratio of length to width is set to optimize the magnetic anisotropies that give the films the desired magnetic orientations.

Thus, the ratio of the length to the width of a planar cell stays the same. As a result, device resistance is constant, independent of minimum feature size. Since a given value of sense current is required to achieve a target output level, the current density increases as the inverse of the minimum feature size so that problems with heating and electromigration become worse as that dimension decreases. Vertical cells are not feasible, and prototype structures that use vertical stacks of ferromagnet-nonmagnetic metal layers have not yet succeeded.

Nevertheless, there is reason for optimism. Ongoing research is focused on developing materials systems with larger coefficients of magnetoresistance. Doubling or tripling the relative resistance difference would significantly reduce bias currents and resolve the associated problems.

IN HOT PURSUIT

Advances in magnetic materials research during the past decade have vastly improved the density and capacity of magnetic disks and tapes, but the pursuit of integrated magnetic-based RAM is in its infancy. Apart from a 10-year history of marginal R&D on low-density, radiation-hard magnetic RAM, a reasonable estimate is that the field has seen roughly 200 person-years of research, most of it during the past three years.

Nonetheless, its achievements have already been considerable. Prototype devices from the main device types have been built at competitive dimensions, with cell areas that compare well with those of floating-gate cells, and in a bit-addressable architecture. Several prototypes have shown write/read times on the order of 10 ns, at least three orders of magnitude faster than floating-gate write times and six orders faster than floating-gate erase times. Write and read power usage of a few milliwatts has been shown, smaller by two to three orders of magnitude than typical floating-gate write and read powers. Several prototypes have demonstrated a minimal endurance of 10^{15} write/erase cycles, at least nine orders of magnitude greater than those of floating-gate devices.

Magnetic RAM may establish a foothold in niche markets such as radiation-hard

memories, low-density and high-performance embedded applications that presently use static RAM, or reprogrammable logic [see "Dynamics of magnetoelectronic logic," p. 38]. If continuing R&D is successful in solving issues of manufacturability, integrated magnetoelectronic memories may well replace floating-gate technology.

The ouster of dynamic RAM is less predictable. Magnetoelectronics has demonstrated write/read cycle times that are faster by a factor of 10, alongside operating powers that are smaller by two to three orders of magnitude. But because of its nonvolatility, high-density magnetoelectronic RAM could do away with the memory overlap in many computer architectures by combining the functions of cache, dynamic RAM, and disk storage. Although disks are much less expensive per bit than magnetoelectronic RAM could ever be, disk drives consume high power and require a large footprint—disadvantages for portable or low-power applications.

An integrated magnetoelectronic architecture has its cost and performance rewards, but it is also true that dynamic RAM is a moving target. Magnetoelectronics must move rapidly just to stay in the ballpark, and it must anticipate the achievement, for example, of cells with a minimum feature size of 0.1 μm and an area of 0.06 μm^2 on a time scale of a few years in order to challenge the existing paradigm for main memory. ♦

TO PROBE FURTHER

For a general background in magnetoelectronics, see G.A. Prinz, "Hybrid ferromagnet semiconductor devices," *Science*, Vol. 250, 1990, p. 1092-97; or Mark Johnson, "The all-metal spin transistor," *IEEE Spectrum*, Vol. 31, no. 5, 1994, p. 47-51.

The *Journal of Applied Physics*, Vol. 85, no. 8, 1999, contains two pertinent papers. One is "High density submicron magnetoresistive random access memory," by S. Tehrani and others, p. 5822-27. The other is "Exchange-biased magnetic tunnel junctions and application to nonvolatile magnetic random access memory," by S.S.P. Parkin and others, p. 5828-33.

A background review of semiconductor nonvolatile RAM can be found in Ashok K. Sharma's book *Semiconducting Memories: Technology, Testing and Reliability*, IEEE Press, New York, 1997.

ABOUT THE AUTHOR

Mark Johnson is a research physicist at the Naval Research Laboratory in Washington D.C. A condensed matter physicist, he does theoretical and experimental research in the areas of magnetism, superconductivity, and semiconductor physics.

Spectrum editor: Linda Geppert