

Design considerations for MRAM

T. M. Maffitt
J. K. DeBrosse
J. A. Gabric
E. T. Gow
M. C. Lamorey
J. S. Parenteau
D. R. Willmott
M. A. Wood
W. J. Gallagher

MRAM (magnetic random access memory) technology, based on the use of magnetic tunnel junctions (MTJs) as memory elements, is a potentially fast nonvolatile memory technology with very high write endurance. This paper is an overview of MRAM design considerations. Topics covered include MRAM fundamentals, array architecture, several associated design studies, and scaling challenges. In addition, a 16-Mb MRAM demonstration vehicle is described, and performance results are presented.

Introduction

MRAM may be a cost-effective solution for long-term data retention and rapid on/off applications such as mobile handheld and general consumer electronic systems. In such cases MRAM may effectively replace a battery and SRAM (static random access memory) and/or flash memory to provide fast, low-power, nonvolatile storage. In large-system applications requiring a reduction in system start-up time (boot time) and the protection of memory contents in the event of sudden unexpected power-down events, MRAM may serve as a replacement for various combinations of SRAM, DRAM (dynamic random access memory), and flash memory components.

MRAM fundamentals

MTJ device structure

Figure 1 is an illustrative drawing (not to scale) of the fundamental MTJ device structure used for binary storage; it consists of two ferromagnetic layers separated by a thin tunnel dielectric [1, 2]. The lower layer is “fixed,” implying that its magnetic orientation cannot be changed during operation, whereas the magnetic orientation of the upper, “free” layer can be changed by the application of a sufficiently large magnetic field. The MTJ is shaped to fit into a rectangular box and is patterned in various shapes within the box as a circle, an oval, an ellipse, or some sort of re-entrant “Saturn” shape [3]. The long axis of the free layer is oriented parallel to the uniaxial anisotropy magnetic orientation of the fixed layer, resulting in a magnetic orientation of the free layer in two stable states: in the same direction as the fixed

layer (parallel) or in the opposite direction (anti-parallel). While conceptually simple, the fixed and free layers are in fact multilayer structures constructed to achieve the desired read, write, and thermal stability characteristics.

When a small bias voltage is applied between the fixed and free layers, a tunneling current flows through the thin intervening dielectric layer. The magnitude of the current depends on the state of the free layer, with the parallel state having a higher current. The current–voltage characteristic of the device can be modeled as a nonlinear resistor, with the resistance being dependent upon the state of the free layer. The fractional change in the effective resistance is known as the magnetoresistance (MR), which is defined by

$$R_1 = R_0(1 + MR),$$

where R_1 is the effective resistance of the anti-parallel state and R_0 is that of the parallel state. At this stage in the development of MRAM technology, MR values for integrated devices are typically in the range of 30–50%, while new materials under development provide MR values in excess of 100% [4].

Figure 2 illustrates the percentage of change in the effective resistance (relative to the low-resistance, parallel state) as a function of applied magnetic field. The applied field is assumed to be parallel to the long axis of the MTJ. The MR for this particular example is approximately 65%. Note that the switching between the two states is hysteretic; that is, the transition from the high- to the low-resistance state (1 to 0) does not occur at the same applied magnetic field as in the reverse order (0 to 1). This hysteretic behavior allows the device to be used as a memory element.

©Copyright 2006 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/06/\$5.00 © 2006 IBM

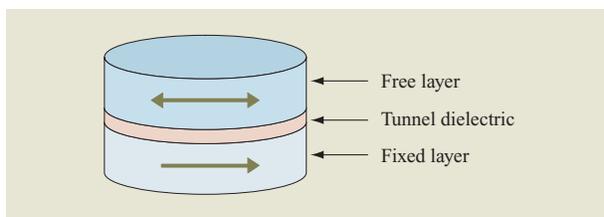


Figure 1

Illustrative drawing (not to scale) of fundamental MTJ device structure, with indicated directions of layer magnetization.

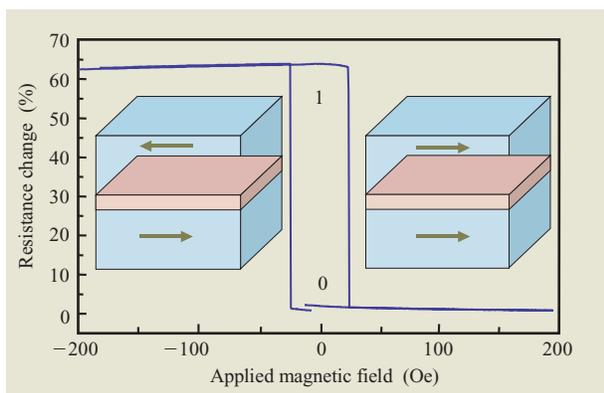


Figure 2

MTJ device switching.

The MTJ structure is integrated into the interconnect portion of an otherwise typical CMOS integrated circuit structure. The CMOS devices allow the integration of circuits to address, read, and write the MTJ memory elements.

MTJ read operation

The MTJ device is read by measuring the effective resistance of the structure, which is a function of the state of the MTJ free layer. This can be achieved by applying a voltage and sensing the current (current sensing) or by applying a current and sensing the voltage (voltage sensing). In either case, the sensed parameter (assumed to be current in the following) is compared to a reference value to determine the state of the device.

The fractional value change in effective resistance or *MR* is not constant but rather decreases with increasing read voltage. Therefore, the relative signal or fractional difference between the data and reference currents decreases with increasing voltage. However, the absolute

signal or absolute difference between the data and reference currents vanishes as the voltage approaches zero. Both relative and absolute signals are critical for a robust, high-performance design. Therefore there exists an optimum value of the read voltage, which appears to be approximately 200 to 300 mV.

Reference method

The reference value must be designed to compensate for process-related variations in MTJ parameters (R_0 and *MR*) and for environmental variations such as voltage and temperature. Three general methods are known for generating the reference value: the twin cell, reference cell and self-referenced methods [5–8]. Current sensing is described in the following examples, although the methods apply to voltage sensing as well.

In the twin cell method, two MTJs are used to store one data bit. The true and complementary MTJs are always written to opposite states. The current associated with the true MTJ is sensed and compared with that of the complementary MTJ to determine the value of the stored data. Use of the twin cell method results in the maximum possible raw signal. However, it has the obvious density disadvantage of requiring two MTJs per bit. In addition, it is sensitive to parameter mismatch between the true and complementary MTJs.

In the reference cell method, the current associated with the data MTJ is sensed and compared with that associated with one or more reference MTJs, which are preprogrammed to known states. If a single reference cell of known state is used, the reference cell current must be multiplied by a certain factor in order to position the reference midway between the 0 and 1 state currents. In another approach, two reference cells are used and are preprogrammed to opposite states. The average of the two reference cell currents is used as the reference. The use of the reference cell method results in only half the raw signal of the twin cell approach, but the MRAM that can be fabricated using this method is much denser, since a reference cell can be shared among many cells. It is also sensitive to parameter mismatch between the data and reference MTJs.

In the self-referenced method, the current associated with the MTJ is sensed and the value is momentarily stored. The same MTJ is then written to a known state and the current is sensed a second time. The original value of the current is compared with the known state current, again multiplied by a certain factor to position it midway between the 0 and 1 state currents. Alternately, the second current value is also momentarily stored, and the MTJ is written to the opposite known state and sensed a third time, permitting the original current to be compared with the average of the 0 and 1 state currents. Assuming that the read cycle is not allowed to disturb the

stored data, the original state of the MTJ must then be restored. The self-referenced method utilizes the same raw signal as the reference cell method, requires no chip area for reference cells, and is insensitive to MTJ parameter mismatch, since MTJ is referenced solely to itself. Unfortunately, the repeated write and sense cycles add considerably to the read access, cycle times, and read active power.

Because of its attractive combination of high density, high performance, low power, and high degree of symmetry, the current-sensing two-reference-cell design appears to be the most popular approach. With this method, the raw signal must be sufficiently large to compensate for parameter mismatch between the data and reference cells as well as offsets within the sense amplifier (SA). This requirement places strict requirements on the MR, MTJ parameter matching, and design of the SA.

MTJ write operation

Figure 3 illustrates the MRAM write operation. The selected MTJ, shown in red, is situated between the selected word line (WL) and the selected bit line (BL), both shown in green, which are orthogonal to each other. During the write, currents (blue arrows) are forced along the selected WL and the selected BL, creating magnetic fields in the vicinity of these wires. The vector sum of the fields at the selected MTJ must be sufficient to switch its state. However, the field generated by the WL or BL alone must be small enough that it never switches the state of the so-called half-selected MTJs that lie along the selected WL and BL.

The process is designed so that the word lines and bit lines are as close as possible to the MTJs for good magnetic coupling to the MTJs. Nonetheless, currents of the order of 5 mA are typically required to switch the state of an MTJ. These currents are considered large by integrated circuit standards and create a variety of challenges for write circuit design. Further, the associated IR drops along the lines limit their allowable lengths, limiting the maximum number of cells in a memory array.

The pulse widths of the WL and BL current pulses are typically approximately equal to or less than 10 ns. However, the two pulses are typically offset by a few ns, with the WL pulse beginning first, so that the free layer can be switched to its new state in a controlled manner.

The magnetic field experienced by WL or BL half-selected MTJs is perpendicular to the wire that generates the field. Further, the field applied to the fully selected MTJ points in a third, somewhat diagonal direction. The hysteresis loop of Figure 2 is insufficient to fully describe these situations, since it is limited to fields in one direction only (along the long axis). The astroid plot, illustrated in **Figure 4**, describes the switching of the free layer in response to both field strength and direction.

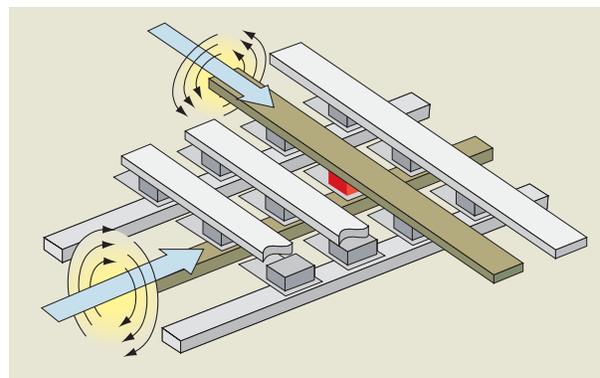


Figure 3

MRAM write operation.

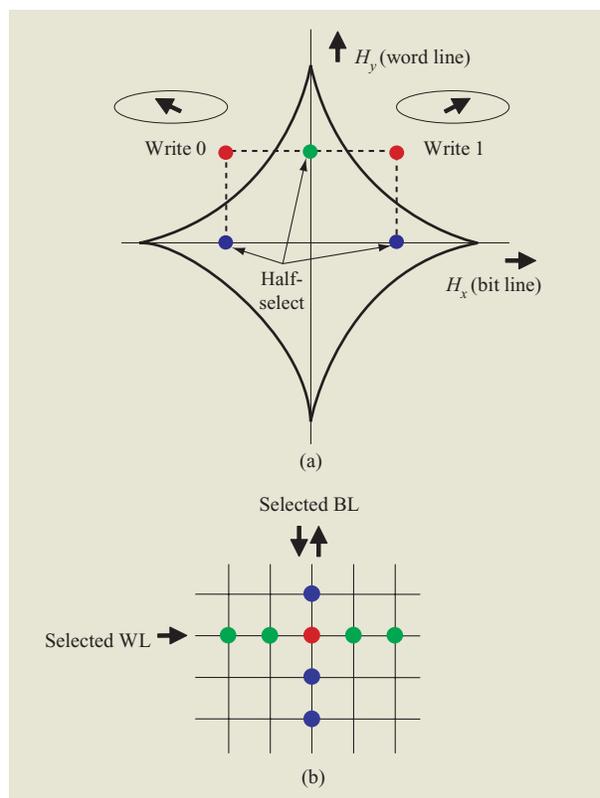


Figure 4

Ideal astroid plot describing the switching of the free layer of an MTJ (a) and the corresponding MRAM array (b).

The x - and y -axes represent the x and y components of the magnetic field applied to the MTJ. In this figure, the long axis of the MTJ and the WL are assumed to be horizontal and the BL to be vertical. Since the WL field

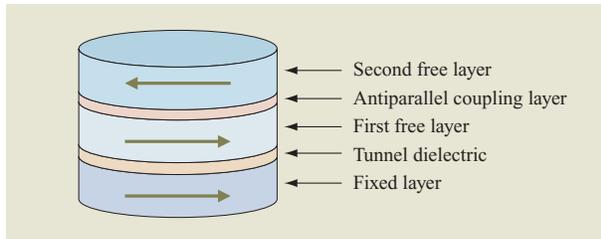


Figure 5

Example of toggle-mode MTJ device structure.

applied to the MTJ is perpendicular to and proportional to the WL current, the y component of the field is proportional to the WL current. Similarly, the x component of the field is proportional to the BL current.

The astroid plot is interpreted in the following manner. If the applied field begins at the origin (no applied field), moves to a point to the right of the y -axis and the diamond-shaped region, or astroid, and returns to the origin, the free layer will point to the right (data state 1). Similarly, if the applied field begins at the origin, moves to a point to the left of the y -axis and the astroid, and returns to the origin, the free layer will point to the left (data state 0). If the applied field remains inside the astroid, the state of the MTJ remains unchanged.

The fully selected MTJ experiences both x and y field components, placing it in the first or second quadrant of the figure depending on the data state to be written. Since the polarity of the x field component or BL current determines the written data state, the BL write circuitry must support a bidirectional current. The y field component is independent of the data state to be written, simplifying the design of the WL write circuitry because bidirectional currents are not required. In order to write successfully, the fully selected field points must always lie outside the astroid.

As indicated in Figure 4, WL half-selected MTJs experience a y field component only, whereas BL half-selected MTJs experience an x field component only. The polarity of the field experienced by a BL half-selected MTJ depends on the state being written to the fully selected MTJ. To avoid half-select disturbs (data loss between an MTJ being written and read), the half-select field points must always lie inside the astroid.

Write margin

Write margin is the ability to reliably write the selected MTJ without disturbing other bits. Write margin requires that the fully selected fields always lie outside the astroid, while the half-selected fields lie within the astroid. Several additional mechanisms degrade the write margin, as described below.

In addition to the half-select field described above, the two WL half-selected MTJs immediately adjacent to the fully selected MTJ experience a small x component field because of the adjacent BL current. Similarly, the two BL half-selected MTJs immediately adjacent to the fully selected MTJ experience a small y component field because of the adjacent WL current. The magnitude of these stray fields depends on the design of the memory cell and is typically less than 10% of the corresponding BL or WL field. Nonetheless, these stray fields further degrade the write margin. The stray field problem becomes more significant as the cell size and hence the distance to the adjacent WL and BL are reduced.

The astroid shown in Figure 4(a) is hypothetical. The shape and size of the astroid are dependent upon the shape, size, and other properties of the MTJ. Correspondingly, the shape and size of each MTJ within a chip vary because of local variations in shape, size, and other properties. The resulting statistical spread of the astroid further degrades the write margin. The write margin challenge is further exacerbated by the finite chance that MTJs operated close to the astroid boundary may undergo undesired thermally activated switching over a vanishingly small potential barrier from one data state to the other [9].

In addition, the applied field varies with variations in circuit parameters (FET parameters, wiring resistance, and supply voltage). The resulting variations in the position of the full and half-select field points [see Figure 4(b)] on the astroid plot degrade the write margin still further. It is the goal of the write circuit design to limit these variations and to compensate for the temperature dependence of the astroid.

Toggle mode

In response to the write margin difficulties associated with the conventional MTJ device, a more complex, “toggle-mode” MTJ device and switching method have been developed [10, 11]. **Figure 5** illustrates the toggle-mode MTJ device structure. The structure is similar to that of the conventional MTJ except that the free layer consists of two weakly anti-parallel coupled ferromagnetic layers. In addition, the long axis of the structure lies at approximately 45 degrees with respect to the WL as opposed to being parallel to the WL. The read operation is essentially unchanged, with the magnetic orientation of the lower free layer determining the effective resistance of the structure.

Whereas the conventional MTJ is written directly into one state or the other depending on the polarity of the BL current, the toggle-mode MTJ toggles its state when exposed to a similar WL and BL current pulse sequence. As illustrated in **Figure 6**, the dipoles of the free layer rotate slightly in the direction of the applied field, and

essentially follow the applied field as it rotates during the WL and BL current pulse sequence. At the end of the sequence, the free-layer dipoles have rotated 180 degrees from the initial state, regardless of the initial state. The criterion for a successful toggle is that the applied field must trace a path in the applied field plane that encloses a particular point in the plane, referred to as the “spin-flop” point.

Unlike a conventional MTJ, a toggle-mode MTJ is largely insensitive to half-select disturbs, regardless of WL and BL field strength, since such disturbs do not trace a path that encloses the spin-flop point. In addition, since the free layer has no net magnetic moment, the field experienced by a particular device is insensitive to the state of adjacent devices. This advantage is of particular importance as cell size and hence the distance to the adjacent devices are reduced. A final advantage of the toggle-mode MTJ is that only one BL write current direction must be supported, simplifying the design of the write circuits.

Because of the toggle nature of the device, the device must be read at the start of the write cycle. The device is then toggled if its current state does not match that of the incoming write data. Although the read can be performed concurrently with preparations for the WL and BL write pulse sequence, the required read represents a write-performance disadvantage compared with that of a conventional MTJ.

Array architecture

There exist two basic architectures for constructing an MRAM array—the cross-point (“XPT”) architecture and the one-transistor, one-MTJ (“1T1MTJ”) architecture [12], as illustrated respectively in **Figures 7** and **8**.

In the XPT architecture, the MTJs lie at the intersection of the WLs and BLs, which connect directly to the fixed and free layers (or vice versa). This arrangement allows for a considerable packing density. Since no contact is made to the silicon within the cell, it is possible to stack such arrays, thus further increasing MRAM density. In addition, it might be possible to place peripheral circuits under the array, increasing the density even further.

However, the XPT architecture involves several significant design challenges. Each MTJ introduces a resistance through which write current may be lost. The only effective way to limit this loss is to increase the effective resistance of the MTJ, which in turn reduces the absolute value of the signal during the read operation. Further, during the read operation, there is no device within the cell to assist in selecting the cell. As a result, current from other cells along the BL interferes with the sensing operation. The result of these effects, described later in greater detail, is very poor read performance.

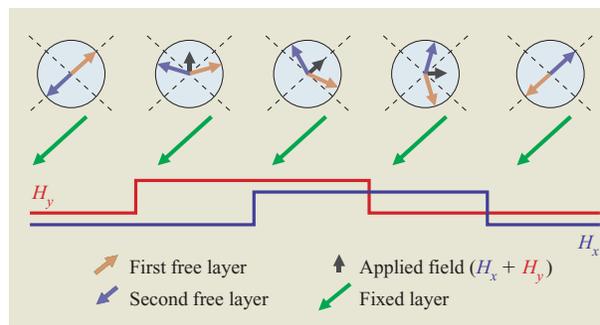


Figure 6

Illustration of toggle-mode switching.

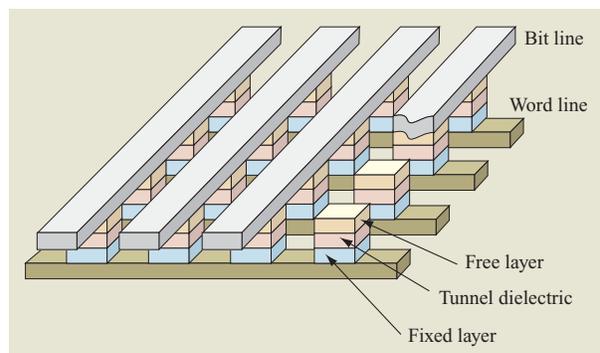


Figure 7

XPT architecture for MRAM array. Adapted from [12], with permission; ©2002 IEEE.

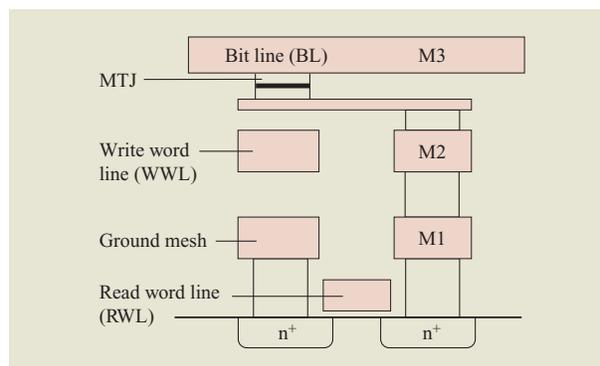


Figure 8

1T1MTJ architecture for MRAM array. From [13], with permission; ©2004 IEEE.

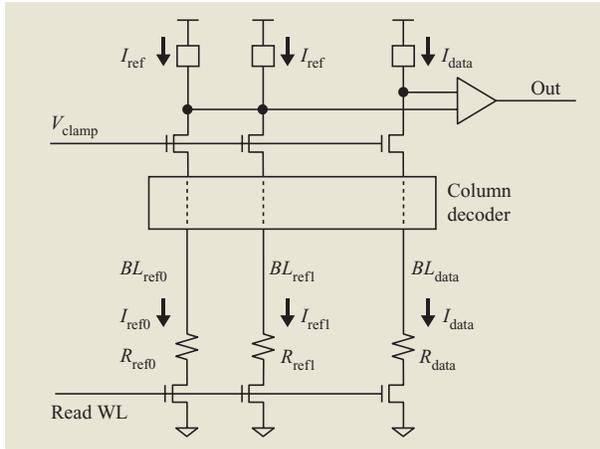


Figure 9

Simplified schematic of a 1T1MTJ read system.

In the 1T1MTJ array architecture, each MTJ is connected in series with an n-type FET, or n-FET. The n-FET, the gate of which is the read word line (RWL), is used to select the cell for the read operation. The write word line (WWL) runs directly below but does not actually contact the MTJ. The RWL and WWL run parallel to each other and perpendicular to the BL, which contacts the free layer of the MTJ. The source of the n-FET is grounded, whereas the drain connects to the fixed layer of the MTJ via a thin local interconnect layer. This layer and the dielectric below it are relatively thin in order to ensure good magnetic coupling from the WWL to the MTJ.

The density of the 1T1MTJ array architecture is less than that of the XPT array architecture for several reasons. The 1T1MTJ cell must include sufficient space for the contact extending down from the thin local interconnect layer, which is typically adjacent to the MTJ since the WWL is directly below the MTJ. The cell size may also be limited by the size of the n-FET and its associated source/drain contacts. In contrast to the XPT architecture, it would be difficult to stack multiple layers of 1T1MTJ arrays, since each cell must make contact with the silicon below it. Similarly, it would not be possible to place the peripheral circuits below the array because the silicon is utilized by the n-FETs of the cells.

Electrically, the 1T1MTJ array architecture has several advantages. During write operation all RWLs are low, eliminating the possibility of losing write current through the MTJs. As a result, the effective resistance of an MTJ may be much lower and the absolute read signal therefore much higher than in the XPT array architecture. Further, only the selected RWL is driven high during a read operation, preventing currents from other MTJs on the

BL from interfering with the sensing operation. For these reasons, the read performance of the 1T1MTJ array architecture is far superior to that of the XPT.

For a variety of reasons including its superior read performance, the 1T1MTJ appears to be preferable. The read and write operations for the two array architectures are described in greater detail below.

1T1MTJ array architecture

Read operation

Figure 9 illustrates a simplified schematic of a 1T1MTJ read system featuring a current-sensing two-reference-cell design. This appears to be the most popular design, although variations on it exist and are described later.

The sense amplifier (SA) is connected to three BLs, one data and two references, by the column decoder system. The selected RWL is driven high, connecting the data cell and two reference cells to their respective BLs. The two reference cells are preprogrammed to opposite states. Three n-FETs, gated by V_{clamp} , operate as source followers, holding the three BLs at the desired voltage, approximately one threshold voltage (V_t) below V_{clamp} . For maximum signal, it is critical that the impedance of the source follower, as viewed from the BL, and the impedances of the column decoder, BL, and cell n-FET all be small relative to the effective resistance of the MTJ (i.e., small enough to permit the latter to determine the current flowing in this path). This generally requires that the effective resistance of the MTJ be at least 5–10 k Ω .

The drain of each source follower is connected to a load device. The load devices are connected in turn to the power supply and serve to convert the current signal into a voltage signal that is sensed by the differential voltage amplifier to create the SA output signal *Out*. However, note that the drains of the two reference source followers are shorted together. By solving Kirchoff's current law at this node, it is easily shown that the current flowing through each of the reference load devices is equal to the average of the two reference currents. This provides the ideal voltage at the reference input of the differential voltage amplifier: exactly midway between the voltages corresponding to the two data states.

The raw signal must be large enough to compensate for 1) signal loss due to random parameter mismatch between data for reference cells and 2) SA offset resulting from random parameter mismatch in the devices within the SA. Careful design of the SA is required in order to minimize the SA offset while maximizing the read performance.

Write operation

Figure 10 illustrates a simplified schematic of a 1T1MTJ BL write system for a conventional (not toggle-mode)

MTJ requiring bidirectional BL write currents. The 1T1MTJ WWL write system is similar, though somewhat simpler, since bidirectional currents are not required. Nonetheless, several design variations exist, as described in a later section.

The selected BL is represented by the resistor in the center of the figure. At either end, the selected BL is connected to the master bit lines (MBLs) by column decoder gated switches. At the left end of each MBL is a current source circuit, and at the right end is a current sink circuit.

The write cycle proceeds as follows. The column decoder and one sink circuit (lower right in this example) are enabled, ensuring that the entire path (selected BL and both MBLs) are discharged to ground. One current source (upper left in this example) is then enabled, creating a current from source to sink and passing through the selected BL as indicated. The timing of the current pulse is controlled by the current source. The polarity of the write current can be reversed by enabling the opposite source and sink circuits in order to write the opposite data state.

The diagonally opposite placement of the current source and sink circuits ensures that the effective resistance of the write current path from source to sink is essentially independent of the position of the selected BL. Since the current source is not ideal, this arrangement reduces the column address dependence of the write current and thus improves the write margin.

XPT array architecture

Read operation

Figure 11 illustrates an equivalent circuit of an XPT array architecture read system. Starting from the left side of the figure, the resistor labeled $R_{\text{unselected}}/(n-1)$ represents the parallel resistance of the $n-1$ unselected MTJs along the selected BL; it is connected to the selected BL and a node representing the unselected WLs, which are driven to the equalization voltage (V_{eq}). The selected MTJ is represented by the resistor labeled R_{selected} , which is connected to the selected BL and the selected WL; the latter is driven to the equalization voltage minus the voltage intended to be applied to the selected MTJ, or $V_{\text{eq}} - V_a$. In order to sense the current through the selected MTJ without interference from the unselected bits, the SA, which consists of an operational amplifier A and feedback element F in a negative feedback configuration, attempts to force the selected BL voltage to V_{eq} and measure the current required to maintain this voltage.

If A were ideal (infinite gain, no offset), the system would achieve equilibrium with the selected BL at V_{eq} and the SA output voltage V_{out} equal to V_{eq} plus the voltage across F corresponding to the current through the selected MTJ. Since both terminals of $R_{\text{unselected}}/(n-1)$

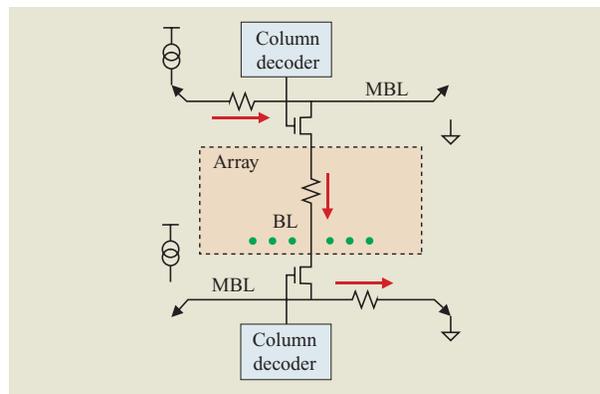


Figure 10

Simplified schematic of a 1T1MTJ bit-line write system.

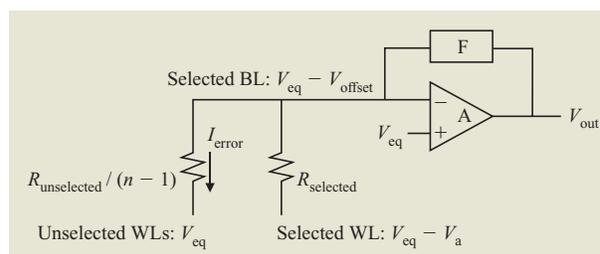


Figure 11

Equivalent circuit of an XPT array architecture read system.

are at V_{eq} , no current flows through the unselected MTJs. The analog output voltage V_{out} is compared with that of one more identical circuit sensing a reference cell of known state in order to determine the state of the data cell.

Unfortunately, despite the use of layout techniques to minimize such effects, there will always exist a certain amount of random mismatch in the parameters of the devices used to construct A. These mismatches arise from random local fluctuations in device dimensions and channel doping, for example, and typically limit the standard deviation of the offset of A to a value of the order of 1 mV. This offset (V_{offset}) causes the system to reach equilibrium at a BL voltage of $V_{\text{eq}} + V_{\text{offset}}$. While V_{offset} is small relative to V_a (perhaps 200–300 mV) and leads to a negligible change of the current flowing through the selected MTJ, it creates a sizable error current through the many unselected MTJs (I_{error}). The SA output reflects the value of the selected MTJ current plus the random I_{error} term.

For robust sensing, the standard deviation of the I_{error} must be much less than that of the raw signal, or I_{signal} :

$$\sigma(I_{\text{error}}) \ll I_{\text{signal}} = \frac{I_0 - I_1}{2}.$$

Substituting and solving for the allowed offset gives

$$\frac{\sigma(V_{\text{offset}})}{\frac{R_0}{n-1}} \ll \frac{1}{2} \left(\frac{V_a}{R_0} - \frac{V_a}{1+MR} \right).$$

Hence,

$$\sigma(V_{\text{offset}}) \ll \frac{V_a MR}{2(n-1)(1+MR)},$$

which is much less than 230 μV for $n=128$, $V_a=250$ mV, and $MR=30\%$.

Such a small value of offset cannot be achieved by layout techniques alone; therefore, it is necessary to resort to offset compensation techniques. With such techniques, the offset of the amplifier is measured and stored away, perhaps as the voltage on a capacitor, during a calibration phase. The selected MTJ is then sensed by using the stored offset to compensate for the offset, ideally creating a zero-offset amplifier. Although the use of compensation results in some improvement, the compensation techniques are not perfect, and a finite amount of random offset will remain.

The SA feedback loop must be sufficiently damped so that the design is stable (no ringing). Both the calibration phase and the measurement phase must be long enough to allow the system to stabilize. For these reasons, the read access time for an XPT design is typically in excess of 100 ns. While several different XPT sensing approaches have been proposed, they typically include a two-phase (calibration and measurement) method of offset compensation in which the length of each phase is limited by a slow-settling negative feedback amplifier [12].

Because of the relatively small voltages involved in the XPT read operation, such a design is expected to be very sensitive to noise. Thus, it is probably not appropriate for an embedded memory, in which power-supply variations from other activity within the chip require a very robust design.

Write operation

In MRAM cross-point array architecture, the write current diminishes as it traverses a WL or BL, since each MTJ represents a resistance through which write current may be lost. The only effective way to limit this loss is to increase the effective resistance of the MTJ. For arrays of reasonable size, this limits the effective resistance to values well in excess of 100 k Ω , or more than ten times higher than the optimal value for the 1T1MTJ array architecture. This in turn reduces the absolute value of the read signal, further complicating

and degrading the performance of the XPT read operation.

Illustrative designs

This section focuses on several designs for the SA and write driver of the 1T1MTJ array architecture. Circuits relevant to that architecture were chosen because it is currently the more feasible option. The SA and write driver were chosen because they represent the most critical and novel circuits involved in the MRAM read and write operations. The WWL write system was chosen over the BL write system for simplicity, though the two circuits are typically very similar.

Three SA designs for 1T1MTJ array architecture

The primary challenges associated with the design of the SA of 1T1MTJ involve minimizing the SA offset and maximizing the read performance. SA offset results from random parameter mismatch in the FET devices within the SA. While sensing an MR of 30% may appear trivial at first glance, the size of the signal relative to the reference current is considerably smaller, as shown by the following:

$$\frac{I_{\text{signal}}}{I_{\text{reference}}} = \frac{(I_0 - I_1)/2}{(I_0 + I_1)/2},$$

where

$$I_0 = \frac{V_a}{R_0}$$

and

$$I_1 = \frac{V_a}{R_1} = \frac{V_a}{R_0(1+MR)}.$$

Substituting and simplifying,

$$\frac{I_{\text{signal}}}{I_{\text{reference}}} = \frac{MR}{MR+2}$$

and

$$\frac{I_{\text{signal}}}{I_{\text{reference}}} = 13\% \text{ for } MR = 30\%.$$

The MTJ resistance mismatch between the data and reference cells degrades this signal further. A typical memory redundancy system is capable of replacing cells which fall outside approximately 4.5 standard deviations from the mean. For example, if the standard deviation of the mismatch is 1%, 4.5% of the signal budget must be allocated to the cells, leaving 13% – 4.5%, or 8.5% for the remaining terms of the signal budget.

Since the number of SAs on a chip is far less than the number of cells, perhaps only 3.5 standard deviations

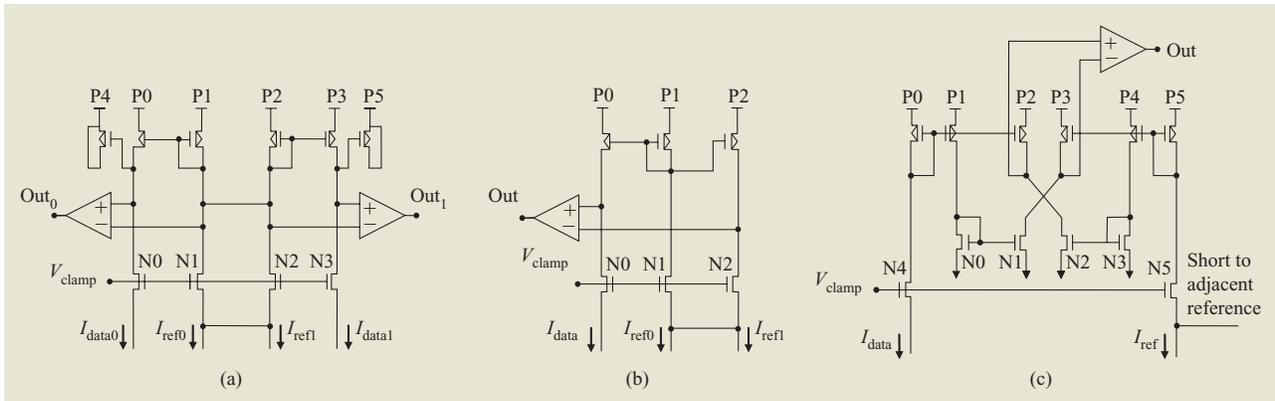


Figure 12

Three SA designs for the 1T1MTJ array architecture. Design (a): Two SAs share two reference BLs. The reference BL currents I_{ref0} and I_{ref1} are averaged to create an ideal reference current through current mirror load devices P1 and P2 and hence an ideal reference voltage at the negative input terminals of the differential amplifiers. The voltages at the positive terminals correspond to the data BL currents I_{data0} and I_{data1} . The differential amplifiers in turn generate the SA output signals Out_0 and Out_1 . Design (b): Each SA requires two reference BLs. The reference BL currents I_{ref0} and I_{ref1} are averaged to create an ideal reference current through current-mirror load devices P1 and P2 and hence an ideal reference voltage at the negative input terminal of the differential amplifier. The voltage at the positive terminal corresponds to the data BL current I_{data} . The differential amplifier in turn generates the SA output signal Out . Design (c): Two SAs share two reference BLs, although only one is shown in the figure. The reference BL current associated with this SA, I_{ref} , is averaged with that of the adjacent SA via the short to the adjacent reference to create an ideal reference current flowing through device P5. The data BL current in turn flows through device P0. Devices P0 through P5 and N0 through N3 create a highly symmetric cross-coupled current-mirror amplifier which generates a voltage signal at the input of the differential amplifier, which in turn generates the SA output signal Out .

of SA offset must be accommodated within the signal budget in order to achieve an acceptable chip yield. If we set an aggressive goal for the standard deviation of the SA offset of 1%, 3.5% of the signal budget must be allocated to the cells, leaving $8.5\% - 3.5\% = 5\%$ for the remaining terms of the system budget. This remaining 5% relative signal must ensure that the SA achieves the correct result within the signal development time. This example illustrates the importance of achieving low offset and high performance in the design of the SA for the 1T1MTJ architecture.

While both current and voltage sensing are possible options, current sensing is generally the higher-performance option. In voltage sensing, the BL is driven with a current source, such that the time constant of the BL equals the relatively large BL capacitance times the resistance of the MTJ. In current sensing, the BL is driven with a voltage source, such that the time constant of the BL equals the relatively large BL capacitance times the effective impedance of the voltage source, which must be much smaller than the MTJ resistance in order to achieve maximum signal. While other delays influence the SA performance, current sensing appears to be the higher-performance option.

The three SA designs described here have a great deal in common. Each utilizes current sensing and averages

the current from two reference cells to create the reference value; an n-FET source follower to drive the BL; a load structure to convert the current signal into a voltage signal; and a differential amplifier to sense this latter signal. The differences in the designs lie primarily in the configuration of the reference and load circuitry.

Figure 12 illustrates the designs [13–15]. In the first design [Figure 12(a)], two SAs share two reference BLs, shown in the center of the figure. A p-type FET (p-FET) current-mirror load structure is used for the load devices within each SA to achieve relatively high current-to-voltage gain. The diode-connected side of the current-mirror load is connected to the reference side of the SA. The reference sides of the two SAs are shorted together both above and below the source-follower n-FETs to ensure accurate averaging of the two reference currents. The data and reference sides of each SA are well balanced, improving performance and noise immunity, with the exception of the gate load associated with the p-FET current-mirror load structure, which appears solely on the reference side. Dummy p-FET capacitors are introduced on the data side to rebalance the design.

The symmetry and simplicity of this design minimize the SA offset distribution. The two reference-side source-follower n-FETs are essentially one device, since all of the connections to the two devices are shared. The resulting

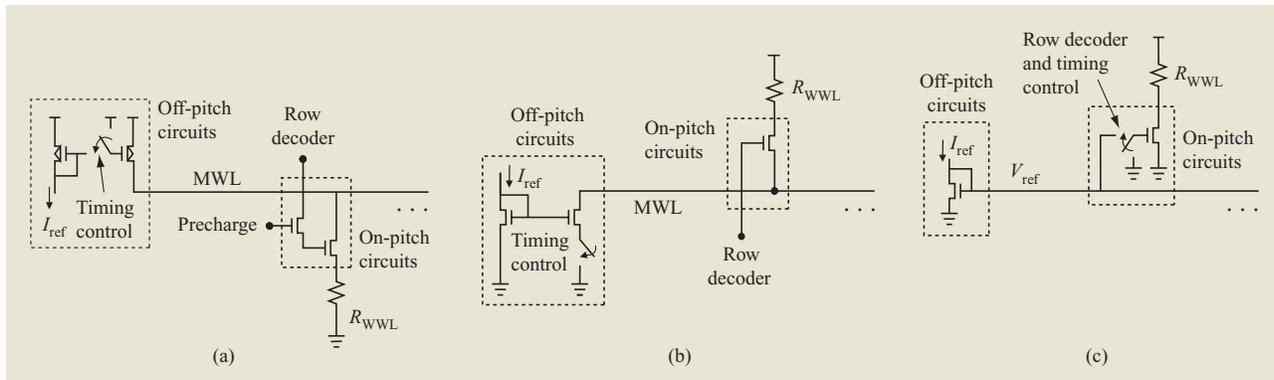


Figure 13

Three WWL write system designs for the 1T1MTJ array architecture. Design (a): An off-pitch p-FET current source drives the MWL. The on-pitch circuitry consists of an n-FET write-driver device, the gate of which is bootstrapped above the supply voltage V_{dd} . Design (b): An off-pitch n-FET current source drives the MWL. The on-pitch circuitry consists of a simple n-FET write-driver device. Design (c): An off-pitch diode-connected n-FET generates a reference voltage V_{ref} which is distributed to the on-pitch circuits. The on-pitch circuitry consists of an n-FET current source which uses V_{ref} as its gate voltage.

device is effectively twice as large and thus has improved parameter-matching characteristics. The same is true of the reference-side p-FET current-mirror load devices. This effect results in a modest but useful improvement in the SA offset distribution.

In the second design [Figure 12(b)], each SA requires two reference BLs. The SA thus has one data and two reference legs. The two reference legs are shorted below the source-follower n-FETs to allow averaging of the reference currents. The load structure is a modified current-mirror load design, with one reference leg being diode-connected and the other two legs current-source-connected. The differential amplifier senses the two current-source-connected legs, one data and one reference. This design develops signal rapidly at the input of the differential amplifier, since these nodes include no gate capacitance associated with the current-mirror load circuitry.

As for the first two designs, in the third [Figure 12(c)] two SAs share two reference BLs. However, the reference legs are shorted below the source-follower n-FETs only. The load structure has been replaced with a highly symmetric cross-coupled current-mirror amplifier. Ideally, this amplifier provides twice the gain of the load structures of the other designs; however, the additional matched pairs in this design degrade the overall SA offset distribution.

Three write-system designs of 1T1MTJ array architecture

The primary challenges involved in the design of the WWL write system for the 1T1MTJ array architecture involve maximizing the write-current uniformity and the layout efficiency of the circuits. Since current variations

due to variations in circuit parameters (FET parameters, wiring resistance, and supply voltage) degrade the write margin, it is essential that the write system be as insensitive to these variations as possible. Furthermore, since the write currents are generally considered large by integrated circuit standards, it is important that the write-system devices be operated very efficiently to minimize the write-system circuit area.

For the three WWL write-system designs for the 1T1MTJ array architecture described here, a similar concept can be applied to the BL write system, though the system may be somewhat more complex to permit the use of bidirectional write currents. Each design employs an off-pitch circuit which is shared for one array and receives a reference current. Each design also employs an on-pitch write driver for each WWL. In the first two designs, an off-pitch current source drives a common line that is connected to the WWL by a switch device. In the third, a current source is included on-pitch for each WWL.

Figure 13 illustrates the designs. In the first design [Figure 13(a)], the off-pitch circuit consists of a p-FET current-mirror current source which drives a shared master WL (MWL) and in turn the WWL, the far end of which is connected to ground. Timing of the write pulse is controlled by the activation of the current source. On-pitch, the MWL is connected to the WWL by a large thin-oxide n-FET write driver, the gate of which is connected to the row decoder by a small thick-oxide n-FET which in turn is gated by the precharge signal.

At the start of the write cycle, the row decoder is activated, and the precharge signal is pulsed from the supply voltage V_{dd} to approximately twice V_{dd} and back again, leaving the gate of the driver device floating at V_{dd} .

The current source is activated, driving current through the MWL and WWL to ground. Initially at ground, the voltages of the WWL and MWL rise in response to the current pulse, coupling the gate of the write driver above V_{dd} and maintaining a gate-to-source voltage on this device that approaches but does not exceed V_{dd} . The design therefore makes efficient yet reliable use of the write-driver device. An advantage of this write-system concept is that it is readily extended to support bidirectional BL write currents by including similar circuitry at the either end of the BL.

In the second WWL write-system design [Figure 13(b)], the off-pitch circuit consists of an n-FET current-mirror current source which drives the MWL and WWL, the far end of which is connected to V_{dd} . As for the first design, the timing of the write pulse is controlled by the activation of the current source. On-pitch, the MWL is connected to the WWL by a simple n-FET write-driver device gated by the row decoder. The simplicity of this concept is very desirable, and the concept can be modified to support bidirectional BL write currents.

In the third design [Figure 13(c)], the off-pitch circuit consists of a diode-connected n-FET, which generates a reference voltage corresponding to the reference current. On-pitch, the gate of an n-FET current source is connected to this reference voltage when active and grounded otherwise, thus controlling the pulse timing. The n-FET current source drives the WWL, the far end of which is connected to V_{dd} . Relative to the other designs, this design removes the row-select switch and MWL impedances from the WWL current path, increasing either the current capacity or the allowable array size. Whereas the current pulse rise and fall times of the earlier designs are limited by a time constant equal to the WWL resistance times the heavily loaded MWL capacitance, this design potentially supports faster current-pulse rise and fall times and is limited by the time constant of the WWL itself. However, device parameter mismatch between the off-pitch diode-connected n-FET and the many on-pitch n-FET current-source devices increases the width of the write-current distribution, degrading the write margin. With some modification, the concept can be extended to support bidirectional write currents.

The scaling challenge

To be viable, an integrated circuit memory technology must be scalable in order to take advantage of the rapid advance of lithography technology and to keep pace with advances in other system components. The most significant challenge facing MRAM in this regard pertains to the write currents. The magnetic fields and hence the write currents required to write the MTJs cannot be chosen arbitrarily because they are related to the thermal stability of the MTJs. The MTJs are designed

to achieve an acceptable soft-error rate, comparable to that of other memory devices, and this in turn defines the magnetic fields and currents required. Unfortunately, as the physical size of the MTJs decreases, the fields and currents required to maintain the same soft-error rate increase.

Increasing write currents present two problems. Since the voltage drops along the WL and BL are limited to a fraction of the supply voltage, increasing currents imply shorter WLs and BLs and hence smaller arrays. Smaller arrays degrade the layout efficiency of the chip, since more decoders, SAs, write drivers, and other peripheral circuits are required, potentially negating the density benefits of a smaller cell. Additionally, the WL and BL write currents and the circuits which generate them dominate the active write power consumption of an MRAM chip, which already exceeds that of chips of competing technologies such as DRAM and SRAM. For certain applications that involve frequent writes but may not value the high performance of MRAM, further increasing the active write power may make MRAM a less attractive choice.

Potential solutions

A variety of approaches for lowering write-current requirements have been proposed. While no single solution appears to solve the problem completely, combinations of these methods are likely to permit the scaling of the MRAM technology at an acceptable rate. Some of these involve modifications of the MTJ device design or materials to reduce the required write field. Further optimization of the vertical dimensions and of the layout of the cell may result in more magnetic field at the MTJ per unit of write current. Also, cladding of the three sides of the write wires that do not face the MTJ with a thin layer of ferromagnetic material has been proposed as a method of increasing the field at the MTJ (Figure 14) [7]. Such ferromagnetic liners have been shown to increase the field at the MTJ per unit of write current by a factor of 2 or more depending on the material and geometry involved. Finally, writing with higher WL currents and lower BL currents has the potential to save power, since the WL current can be shared among many bits and the BL current cannot.

16-Mb demonstration vehicle

Description of design

A 16-Mb MRAM was designed and fabricated for the purpose of demonstrating the potential of the MRAM technology [13]. A photograph of the fabricated chip is shown in Figure 15. The chip was designed in an 0.18- μm CMOS technology with three copper metal levels and

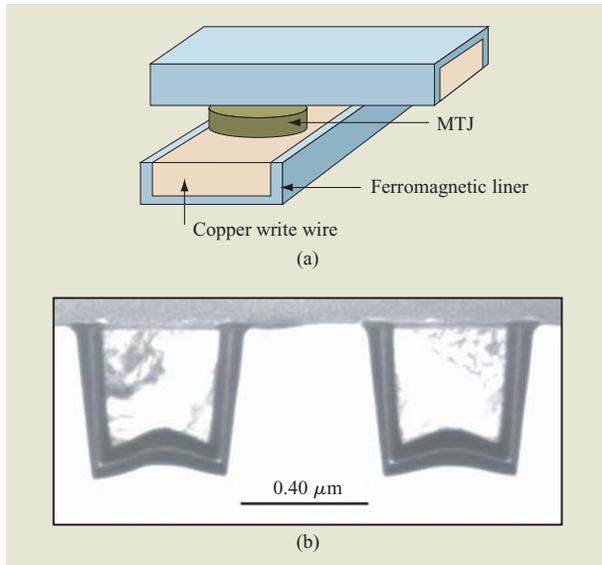


Figure 14

Illustration (a) and micrograph (b) of ferromagnetic liners.

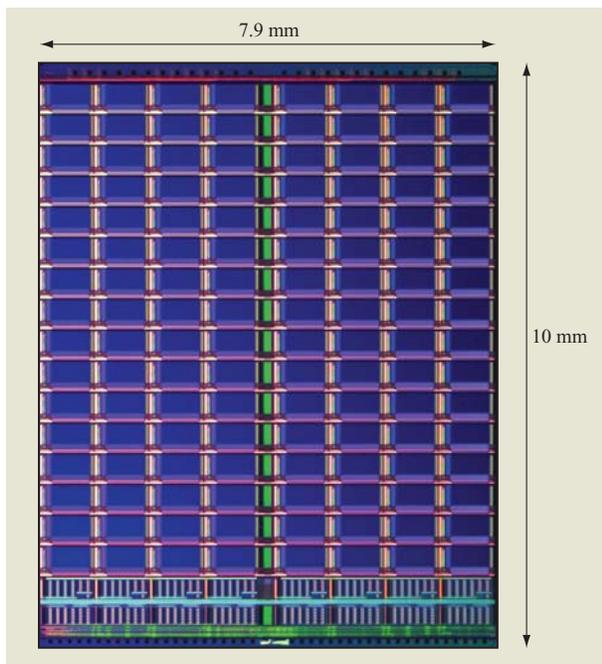


Figure 15

Photograph of 16-Mb MRAM chip. Adapted from [13], with permission; ©2004 IEEE.

three MRAM-specific masks. The chip utilizes a low-power SRAM-like interface with a 16-bit-wide data

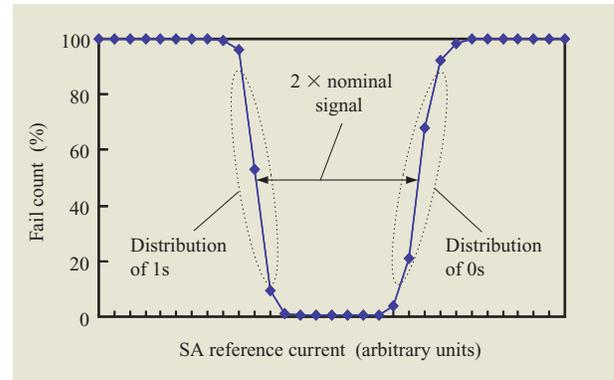


Figure 16

Measured read characteristics of one 32-Kb domain of the 16-Mb MRAM chip. From [13], with permission; ©2004 IEEE.

bus and contains 128 128-Kb arrays, four of which are activated in a given cycle, each contributing 4 bits to the 16-bit data word. The cell area is $1.42 \mu\text{m}^2$ and the chip area is 79mm^2 . The chip operates at an external voltage of 2.3–3.3 V and is regulated to 1.8 V internally. The design supports both conventional and toggle-mode operation.

Read and write characteristics

Figure 16 illustrates the measured read characteristics of one 32-Kb SA domain of the chip. For this measurement, the reference cells were disabled and an externally controlled reference current was provided. The figure illustrates the number of failing bits (fail count) vs. this externally controlled reference current. The current corresponding to 50% fail count on the left side of the chart corresponds to the median high resistance value and that on the right to the median low resistance value of the MTJs within the domain. An ideal reference value would lie midway between the median values, and thus the distance between them represents twice the nominal signal. The adjacent regions represent the cumulative distributions of the corresponding resistance values.

The measured write characteristics of one 128-Kb array of the 16-Mb chip are illustrated in **Figure 17**. **Figure 17(a)** illustrates contours of fail count (0% to 100% in 5% steps) as a function of the BL and WL reference currents for a simple no-disturb pattern (one in which the data is written and immediately read). A series of tightly spaced contours divides the chart into two regions: a switching region, in which sufficient field exists to switch the MTJs, and a non-switching region. The plot resembles one quadrant of the astroid plot discussed earlier. **Figure 17(b)** illustrates a similar

plot for a checkerboard pattern (one in which a number of half-select disturbs occur between the data being written and read). The plot resembles the plot of Figure 17(a), with the addition of some fails (“disturb fails”) near the top of the plot corresponding to high values of WL current and reducing the size of the operating region. More detailed testing has revealed the fails to be WL half-select disturbs in which the adjacent BL is written to the opposite data state. As expected, the fails increase with both WL and BL current.

Read and write cycle times of 30 ns, with read and write active powers of 25 mA and 80 mA, respectively, were achieved. A standby current of 32 μA and a deep power-down current of less than 5 μA were measured at 40°C.

Summary

An overview has been provided of design considerations for MRAM, an emerging nonvolatile memory technology, with emphasis on the challenges faced by the MRAM circuit designer. MTJ device structure and associated write and read operations have been described. Write margin, or the ability to write the selected cell without disturbing others, is a particular MRAM challenge and appears to be greatly improved with the toggle-mode structure.

Two array architectures, the XPT and 1T1MTJ, have been described. While potentially offering higher density, the write and read design challenges posed by the lack of an isolation or select device in the XPT architecture are significant and result in slower read performance. Consequently, at this time it is not surprising that the 1T1MTJ architecture has received more attention.

Several different SA and WWL write system circuits have been described. While they share certain common elements, each represents a unique optimization. The SA systems strive for low offset and high performance and the WWL write systems strive for write current uniformity and layout efficiency.

A memory technology must be scalable to be economically viable. All memory technologies face scaling challenges in one area or another. MRAM faces a particular challenge with respect to write current, which must generally increase with decreasing MTJ size to maintain data stability. A number of technology- and design-related developments were described which may permit the scaling of MRAM write currents for several technology generations.

Finally, a 16-Mb MRAM demonstration vehicle has been described, and illustrative performance results presented. Read and write cycle times of 30 ns were achieved. While it is still too early to predict the long-term success of MRAM as a memory technology, it appears to possess a unique combination of density, performance, and write endurance.

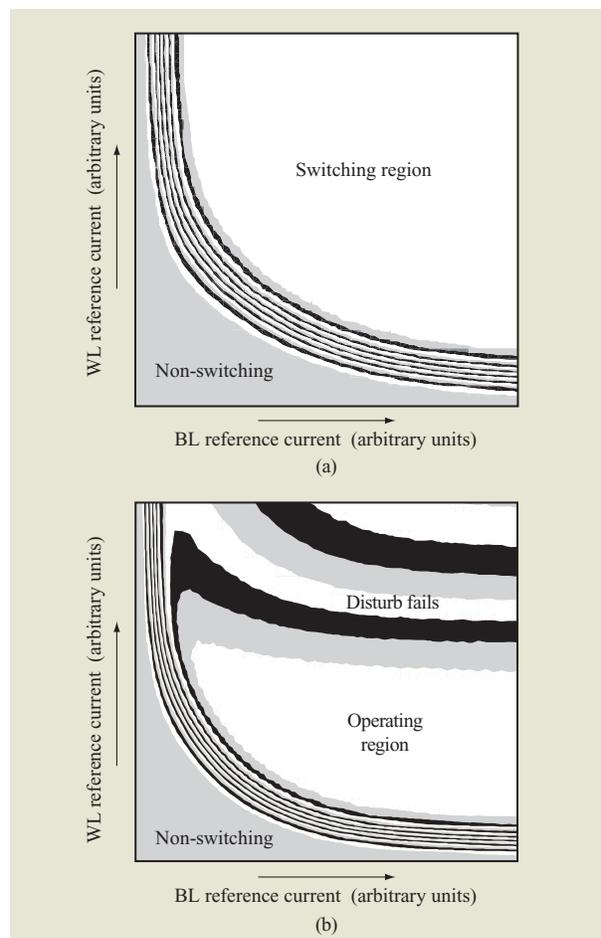


Figure 17

Measured write characteristics of one 128-Kb array of the 16-Mb MRAM chip: (a) Characteristics for a simple no-disturb pattern, from [13], with permission; ©2004 IEEE. (b) Characteristics for a checkerboard pattern.

Acknowledgments

The authors wish to thank the many members of the IBM–Infineon MRAM Development Alliance (MDA) for their many contributions to this work.

References

1. S. Tehrani, J. M. Slaughter, M. Deharrera, B. N. Engel, N. D. Rizzo, J. Salter, M. Durlam, R. W. Dave, J. Janesky, B. Butcher, K. Smith, and G. Grynkeiwich, “Magnetoresistive Random Access Memory Using Magnetic Tunnel Junctions,” *Proc. IEEE*, pp. 703–714 (May 2003).
2. A. R. Sitaram, D. W. Abraham, C. Aloh, D. Braun, S. Brown, G. Costrini, F. Findeis, M. Gaidis, E. Galligan, W. Glashauser, A. Gupta, H. Hoenigschmid, J. Hummel, S. Kanakasabapathy, I. Kasko, W. Kim, U. Klostermann, G. Y. Lee, R. Leuschner, K.-S. Low, Yu Lu, J. Nutz, E. O. Sullivan, C. Park, W. Raberg, R. Robertazzi, C. Sarma, J. Schmid, P. L. Trouilloud, D. Worledge, G. Wright,

- W. J. Gallagher, and G. Muller, "A 0.18 μm Logic-Based MRAM Technology for High Performance Nonvolatile Memory Applications," *Digest of Technical Papers, IEEE Symposium on VLSI Technology*, 2003, p. 15.
3. M. Motoyoshi, I. Yamamura, W. Ohtsuka, M. Shouji, H. Yamagishi, M. Nakamura, H. Yamada, K. Tai, T. Kikutani, T. Sagara, K. Moriyama, H. Mori, C. Fukumoto, M. Watanabe, H. Hachino, H. Kano, K. Bessho, H. Narisawa, M. Hosomi, and N. Okazaki, "A Study for 0.18 μm High-Density MRAM," *Digest of Technical Papers, IEEE Symposium on VLSI Technology*, 2004, pp. 22–23.
 4. S. S. P. Parkin, C. Kaiser, A. Panchula, P. M. Rice, B. Hughes, M. Samant, and S. H. Yang, "Giant Tunneling Magnetoresistance at Room Temperature with MgO(100) tunnel barriers," *Nature Mater.* **3**, 862–867 (2004).
 5. R. Scheuerlein, W. Gallagher, S. Parkin, A. Lee, S. Ray, R. Robertazzi, and W. Reohr, "A 10ns Read and Write Non-Volatile Memory Array Using a Magnetic Tunnel Junction and FET Switch in Each Cell," *ISSCC Digest of Technical Papers*, February 2000, pp. 128–129.
 6. P. K. Naji, M. Durlam, S. Tehrani, J. Calder, and M. F. DeHerrera, "A 256kb 3.0V 1T1MTJ Nonvolatile Magnetoresistive RAM," *ISSCC Digest of Technical Papers*, February 2001, pp. 122–123.
 7. M. Durlam, P. J. Naji, A. Omair, M. DeHerrera, J. Calder, J. M. Slaughter, B. N. Engel, N. D. Rizzo, G. Grynkeiwich, B. Butcher, C. Tracy, K. Smith, K. W. Kyler, J. J. Ren, J. A. Molla, W. A. Feil, R. G. Williams, and S. Tehrani, "A 1-Mbit MRAM Based on 1T1MTJ Bit Cell Integrated with Copper Interconnects," *IEEE J. Solid-State Circuits* **38**, No. 5, 769–773 (2003).
 8. A. Bette, J. DeBrosse, D. Gogl, H. Hoenigschmid, R. Robertazzi, C. Arndt, D. Braun, D. Casarotto, R. Havreluk, S. Lammers, W. Obermaier, W. Reohr, H. Viehmann, W. J. Gallagher, and G. Muller, "A High-Speed 128Kbit MRAM Core for Future Universal Memory Applications," *Digest of Technical Papers, IEEE Symposium on VLSI Circuits*, 2003, p. 217.
 9. J. Z. Sun, J. C. Slonczewski, P. L. Trouilloud, D. Abraham, I. Bacchus, W. J. Gallagher, J. Hummel, L. Yu, G. Wright, S. S. P. Parkin, and R. H. Koch, "Thermal Activation-Induced Sweep-Rate Dependence of Magnetic Switching Astroid," *Appl. Phys. Lett.* **78**, No. 25, 4004–4006 (2001).
 10. M. Durlam, D. Addie, J. Akerman, P. Butcher, J. Chan, M. DeHerrera, B. N. Engel, B. Feil, G. Grynkeiwich, J. Janesky, M. Johnson, K. Kyler, J. Molla, J. Martin, K. Nagel, J. Ren, N. D. Rizzo, T. Rodriguez, L. Savtchenko, J. Salter, J. M. Slaughter, K. Smith, J. J. Sun, M. Lein, K. Papworth, P. Shah, W. Qin, R. Williams, L. Wise, and S. Tehrani, "A 0.18 μm 4Mb Toggling MRAM," *IEDM Tech. Digest*, pp. 995–997 (2003).
 11. C. K. Subramanian, T. W. Andre, J. J. Nahas, B. J. Garni, H. S. Lin, A. Omair, and W. L. Martino, Jr., "Design Aspects of a 4 Mbit 0.18 μm 1T1MTJ Toggle MRAM Memory," *Proceedings of the IEEE International Conference on Integrated Circuit Design and Technology*, 2004, pp. 177–181.
 12. W. Reohr, H. Hoenigschmid, R. Robertazzi, D. Gogl, F. Pesavento, S. Lammers, K. Lewis, C. Arndt, Y. Lu, H. Viehmann, R. Scheuerlein, L.-K. Wang, P. Trouilloud, S. Parkin, W. Gallagher, and G. Muller, "Memories of Tomorrow," *IEEE Circuits & Devices Mag.* **18**, No. 5, 17–27 (September 2002).
 13. J. DeBrosse, C. Arndt, C. Barwin, A. Bette, D. Gogl, E. Gow, H. Hoenigschmid, S. Lammers, M. Lamorey, Y. Lu, T. Maffitt, K. Maloney, W. Obermeyer, A. Sturm, H. Viehmann, D. Willmott, M. Wood, W. J. Gallagher, G. Mueller, and A. R. Sitaram, "A 16Mb MRAM Featuring Bootstrapped Write Drivers," *Digest of Technical Papers, IEEE Symposium on VLSI Circuits*, June 2004, pp. 454–457.
 14. T. W. Andre, J. J. Nahas, C. K. Subramanian, B. J. Garni, H. S. Lin, A. Omair, and W. L. Martino, "A 4-Mb 0.18- μm 1T1MTJ Toggle MRAM with Balanced Three Input Sensing Scheme and Locally Mirrored Unidirectional Write Drives," *IEEE J. Solid-State Circuits* **40**, No. 1, 310–309 (January 2005).
 15. T. Tsuji, H. Tanizaki, M. Ishikawa, J. Otani, Y. Yamaguchi, S. Ueno, T. Oishi, and H. Hidaka, "A 1.2V 1Mbit Embedded MRAM Core with Folded Bit-Line Array Architecture," *Digest of Technical Papers, IEEE Symposium on VLSI Circuits*, June 2004, pp. 450–453.

Received March 29, 2005; accepted for publication May 25, 2005; Internet publication January 5, 2006

Thomas M. Maffitt *IBM Systems and Technology Group, 1000 River Street, Essex Junction, Vermont 05452 (tmaffitt@us.ibm.com)*. Mr. Maffitt is a Senior Engineer. He joined IBM in 1979 after graduating from the University of Notre Dame with a B.S.E.E. degree. He has contributed to DRAM product development and design from the 64-Kb to the 256-Mb level and is the author or co-author of 12 patents and six technical papers. Mr. Maffitt is currently involved in MRAM development.

John K. DeBrosse *IBM Systems and Technology Group, 1000 River Street, Essex Junction, Vermont 05452 (jdebros@us.ibm.com)*. Mr. DeBrosse received a B.S.E.E. degree in 1983 and an M.S.E.E. degree in 1984, both from Purdue University. He joined IBM in 1985, has contributed to four generations of DRAM technology development and product design, from the 4-Mb to the 256-Mb level, and is currently working on MRAM design. Mr. DeBrosse is an author of 24 patents and 22 technical papers.

John A. Gabric *IBM Systems and Technology Group, 1000 River Street, Essex Junction, Vermont 05452 (jgabric@us.ibm.com)*. Mr. Gabric is a Senior Design Manager. He joined IBM in 1970 after graduating from the University of Akron with a B.S.E.E. degree. His contributions include DRAM product design and development from the 64-Kb to the 512-Mb level and SRAM development for standalone and embedded applications. He is the author or co-author of five patents, ten publications, and three technical papers. Mr. Gabric currently manages MRAM, SRAM, and CAM development.

Earl T. Gow *IBM Systems and Technology Group, 1000 River Street, Essex Junction, Vermont 05452 (gowe@us.ibm.com)*. Mr. Gow is a Staff Engineer. He joined IBM in 1997 after graduating from the University of Vermont with a B.S.E.E. degree. His contributions include test and characterization work on multiple generations of DRAM chips from the 64-Mb to the 512-Mb level. Mr. Gow is currently involved in MRAM chip development and EDRAM chip characterization.

Mark C. Lamorey *IBM Systems and Technology Group, 1000 River Street, Essex Junction, Vermont 05452 (mlamorey@us.ibm.com)*. Mr. Lamorey received an A.S.E.E. degree in 1997 from the Vermont Technical College and a B.S.E.E. degree in 2002 from the University of Vermont. He joined IBM in 1997. He has contributed to three generations of DRAM technology development and is currently working in the field of MRAM and PCM design and technology development. Mr. Lamorey is the author of one patent and two technical papers.

John S. Parenteau *IBM Systems and Technology Group, 1000 River Street, Essex Junction, Vermont 05452 (jparente@us.ibm.com)*. Mr. Parenteau is a Senior Laboratory Specialist. He joined IBM in 1982 after graduating from the New England Institute of Technology. He has contributed to the design and support of the MACE memory tester used for product development of memory products including MRAM and embedded memory. Mr. Parenteau is the co-author of one patent and two technical papers. He is currently involved in chip design and layout.

Dennis R. Willmott *IBM Systems and Technology Group, 1000 River Street, Essex Junction, Vermont 05452 (dwillmot@us.ibm.com)*. Mr. Willmott is an Advisory Engineer. He graduated from Iowa State University in 1976 with a B.S.E.E. degree, and then joined Texas Instruments Corporation. He subsequently returned to school and, in 1979, received an M.S.E.E. degree from the University of Illinois. He then joined IBM. Mr. Willmott is the author or co-author of four technical papers; he is currently involved in computer-aided design.

Mark A. Wood *IBM Systems and Technology Group, 1000 River Street, Essex Junction, Vermont 05452 (woodma@us.ibm.com)*. Mr. Wood is a Senior Mask Designer. He joined IBM in Manassas, Virginia, in 1985 after graduating from the United Electronics Institute, Tampa, Florida. He has worked on the physical design of DRAM chips from the 4-Mb to the 256-Mb level, as well as designs for SRAM, PowerPC, and MRAM chips.

William J. Gallagher *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (glgr@us.ibm.com)*. Dr. Gallagher joined the IBM Research Division in 1978 after receiving his B.S. degree in physics (summa cum laude) from Creighton University in 1974 and his Ph.D. degree in physics from MIT. He worked for five years at IBM on scientific and engineering aspects of Josephson computer technology and then, for six years, managed the IBM Exploratory Cryogenics Research Group. In 1989, he participated in the formation of the IBM-AT&T-MIT Consortium for Superconducting Electronics (CSE). He served as a director of the CSE from 1989 until 1995. Since 1995, he has led an effort to explore the use of magnetic tunnel junctions for a nonvolatile random access memory, MRAM, including serving from 2000 to 2004 as the IBM project manager in the MRAM Development Alliance with Infineon. Currently Dr. Gallagher is Senior Manager of the Exploratory Nonvolatile Memories program at the IBM Thomas J. Watson Research Center. He is a Fellow of the American Physical Society and of the Institute of Electrical and Electronics Engineers.