

Global Semiconductor Alliance



SEMICONDUCTOR
FORUM

Vol. 18 No. 2 June 2011
Published by GSA \$60 (U.S.)


SEMICONDUCTOR LIGHT—LET IT SHINE!

AUTOMOTIVE FOUNDRY AND BACKEND FLOW

POWER MANAGEMENT IC TRENDS AND COST-EFFECTIVE TESTING

ELECTRONIC SYSTEM DESIGN CONSIDERATIONS TO MEET EMERGING MARKET NEEDS

*MEMORY SOLUTIONS FOR MEDICAL, SMART GRID, INDUSTRIAL AND OTHER
LONG-LIFECYCLE APPLICATIONS*



*Exploring the Potential of Emerging
Semiconductor Technology*

OVERCOMING THE STABILITY BARRIER TO HIGH-DENSITY MRAM

BARRY HOBERMAN, CHIEF MARKETING OFFICER, CROCUS TECHNOLOGY
KEN MACKAY, DIRECTOR, MAGNETIC CELL ENGINEERING, CROCUS TECHNOLOGY

Successful memory chips must support a wide variety of specific performance and reliability requirements. The physics, technology and design of each memory type dictate the specifications which define the memory's reliability and performance envelope. At the IC level, data retention, write endurance, susceptibility to ionizing radiation, read and write access and cycle time, read and write voltage requirements, and power are sensitive to the physics and characteristics of the core storage element.

For example, information in dynamic random access memory (DRAM) is stored in the form of an electrical charge retained in a small capacitor. The characteristics of the capacitor and surrounding circuits directly affect data retention, susceptibility to ionizing radiation, read voltage, read access time and power. The charge tends to leak out of the capacitor, and so the device requires a regular charge refresh (which requires power) to maintain reliable data retention over extended periods of time. If power to the device is shut down, the charge stored on each capacitor leaks away rapidly and the data is lost, hence DRAM's designation as "volatile memory." Flash memory is a type of "non-volatile memory" (NVM), which characteristically retains its stored information when power is removed from the chip. Stability of stored data in NVM (regardless of whether power is applied or not) is a key performance specification in such devices.


Magnetoresistive RAM (MRAM) is a memory technology that blends the best characteristics of volatile and non-volatile memories. In MRAM, as in Flash memory, data storage is non-volatile, but the write endurance of the memory cell is unlimited, and reading and writing are very fast (similar to DRAM and static RAM (SRAM)). As with other memory technologies, MRAM is constrained by its own physical limits to define the chip performance envelope. Data retention is governed by the fundamental stability of the physical quantity being stored. In a magnetic memory, magnetic polarization is stored rather than electrical charge. This is achieved by controlling the orientation of magnetization in a small magnetic element (analogous to the orientation of a compass needle). As for all memory technologies, data retention is of particular importance and essential

to the future of MRAM as a mainstream memory technology. As such, the stability of the magnetic orientation of the MRAM's bits is crucial.

Thermal Stability in MRAM

At the heart of all memories, a fundamental tension exists between the stability factors that govern data retention and the ability to scale that memory cell to a smaller lithographic feature size. The critical tension in MRAM is that the stability of the magnetic orientation is proportional to the volume of the magnetic particle, so as an MRAM bit is scaled down, its stability generally declines. This destabilizing effect can be countered, however, in several creative ways with differing effectiveness. The various approaches to maintaining MRAM data stability while scaling to competitively small feature sizes (i.e., 90 nanometer and smaller) is one of the distinguishing characteristics of competing approaches to advanced MRAM.

The core storage element in MRAM is a device called a magnetic tunnel junction (MTJ), as depicted in Figure 1. A very thin (1 to 2 nanometer) insulating layer is sandwiched between two thin magnetic layers. In a typical configuration, one of these layers has a fixed magnetic orientation for the life of the device and can be referred to as the reference layer. The second layer, which can be referred to as the storage layer, has a magnetic orientation that can be reversed depending on whether the stored data is a '0' or '1.' The read-back signal is derived from a transverse resistance measurement across both layers and changes from low to high depending on whether the storage layer magnetization is aligned parallel or anti-parallel to the reference layer, respectively.



Stored data follows the magnetization direction (parallel or anti-parallel) of magnetic layers in the MTJ.

The storage layer is designed to have two stable orientations (Figure 1), where the magnetization in the storage layer is either oriented to the left or to the right. Each of the stable orientations has an associated energy level. The energy levels of both orientations are equivalent, but there is an energy barrier to overcome when switching from one orientation to the other. The stability of a given magnetic state over time is intrinsically linked to the magnitude of the energy barrier between the two possible stable orientations of magnetization. It is critical to observe that the ability to write to a new state is also intrinsically linked to the height of this energy barrier. Indeed, during the write process, energy must be supplied to orient the magnetization of the storage layer in the desired direction. The greater the energy barrier between the two orientations, the greater the magnitude of the energy required for the writing process will be. This increase in writing energy tied to an increased stability gives origin to the fundamental tension that exists between the stability factors governing data retention and the ability to scale that memory cell to smaller lithographic feature sizes.

Modeling Memory Chip Stability

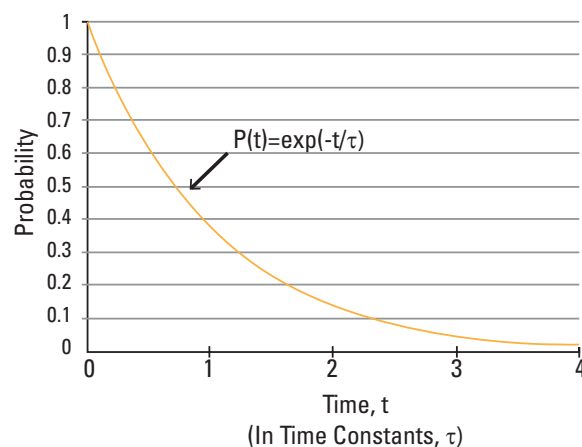
When examining the stability of MRAM, it is important to understand that each individual bit is subject to a potential thermal disturbance of its magnetic orientation, with the previously described energy barrier as the “keeper” of the existing state. The thermal disturbance of an individual bit can be equated to a randomly fluctuating magnetic field that acts to potentially disturb the magnetic state. The magnitude of this equivalent random field is proportional to the device temperature. The effects of such random thermal fluctuations depend on the relative magnitudes of the energy barrier and the available thermal energy, which is proportional to the device temperature. If the energy barrier is high enough, the probability of overcoming it with only random thermal fluctuations is negligibly small. For example, in macroscopic systems such as a compass needle, once magnetized to point north, the needle will continue to do so forever unless it is remagnetized in the opposite direction. If the energy barrier in a MTJ is comparable to the ambient thermal energy available to cause fluctuations, there is a good chance that the orientation can spontaneously flip from one orientation to the other in a random manner within the specified operating life of the device, thus creating a data storage error.

The chance of flipping from one orientation to the other has a mathematical probability profile that is exponential in nature (Figure 2). The characteristic lifetime τ over which flipping occurs depends on the relative magnitudes of the energy barrier and the thermal energy (with the thermal energy quantified as kT , where T is the temperature and k is Boltzmann’s constant). The energy barrier is proportional to the volume of the magnetic dot and can be expressed as KxV , where K is an energy density describing the material and shape of the dot and V is the volume of the magnetic dot. Thus, the characteristic lifetime required to overcome such an energy barrier can be expressed as

$$\tau = \tau_0 e^{(KV/kT)}$$

where τ_0 is the minimum time required to reverse magnetization and is typically on the order of 1 nanosecond. It is clear that the energy barrier and characteristic lifetime increases with increasing volume (i.e., the size of the dot), thereby making the dot more stable, but decreases with increasing temperature, thereby making the dot less stable.

Figure 2. Probability of MTJ Data Retention over Time



This characteristic lifetime determines the reliability statistics of a memory chip relating to its data stability. The probability that an individual bit retains its data follows an exponential decay that depends on the characteristic lifetime:

$$P(t) = e^{-(t/\tau)}$$

(where τ is the characteristic lifetime).

In a memory application, the question is now what τ is required, so over the 20-year life of the device, no bits erroneously change state. For a single 64 Mbit device, this would lead to the requirement that

$$P(20 \text{ years}) = e^{-(20 \text{ years}/\tau)} > 1 - (1/6.4 \times 10^7).$$

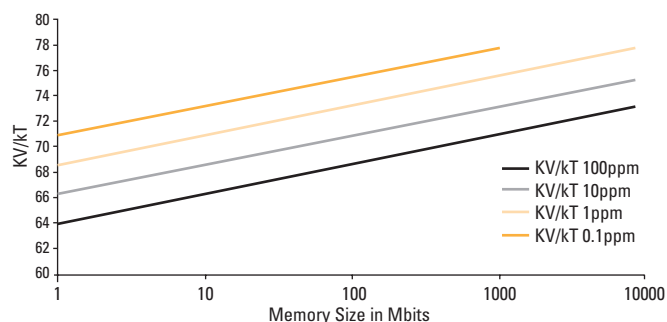
In production, this criterion must apply to all outgoing product, so the requirement becomes even more stringent. Application of the basic rules of probability and straightforward approximations yield the required probability for a given memory density and quality level (typically expressed in parts per million (ppm) of fallout):

$$P(20 \text{ years}, 100 \text{ ppm}) = e^{-(20 \text{ years}/\tau)} > 1 - (\text{ppm_Rate}/\text{Memory_Size})$$

This immediately imposes a lower limit on the ratio KV/kT since $KV/kT = \ln(-20 \text{ years}/(\tau_0 \times \ln(P(20 \text{ years}))))$. For example, for a 64 Mbit memory and a 100 ppm quality level over 20 years, $P(20 \text{ years}) = 1 - 1.6 \times 10^{-12}$, and so KV/kT must be greater than 68.2.

The conventional wisdom within the MRAM community is that production of megabit density memories requires the factor $KV/kT > 60$, which is in line with the above example. Figure 3 specifies this stability requirement by denoting the minimum KV/kT requirement as a function of memory density and acceptable fallout level while assuming a 20-year life.

Figure 3. Minimum KV/kT Required vs. Memory Size



Developers and customers of MRAM technology should evaluate the performance of considered technology against the preceding criteria. The first-generation MRAM technologies, with bit densities up

to about 16 megabit, rely on a KV/kT on the order of 60 for successful chip products. Next-generation MRAM technologies intended for manufacture with 65 nanometer lithography and finer are targeting densities that start at 64 Mbit with a roadmap up to gigabit densities and beyond, requiring KV/kT in excess of 70 (Figure 3).

Achieving High-stability MRAM

The challenge for next-generation MRAM technologies is keeping this stability factor high enough to produce reliable memory while sustaining a writing mechanism that allows for a minimum lithographical size. The problem links back to the previously described energy barrier. The barrier is kept high to maintain stability, but increased barrier height impedes production of a very small memory cell due to the high energy required for writing (i.e., larger currents and voltages are required, thus necessitating larger transistors and other structures in the memory cell).

Thermal-assisted switching (TAS) is an important breakthrough for solving this tension. In TAS, the essential concept is that the KV/kT can be set at different values when “holding” data and when “writing” data. As a result, a very high KV/kT (typically >150) can be engineered for holding data, and a much lower KV/kT (typically <40) applies for writing data. As a result, the maintenance of high data stability ceases to be a constraint in reducing the size of the bit cell, opening a path to 45 nanometer lithography and enabling gigabit memory densities.

To make TAS work, the holding KV/kT applies when the MTJ is at normal operating temperature. During writing, the MTJ is heated in several nanoseconds to about 220 C—the temperature at which the material characteristics of the MTJ change and the KV/kT falls to a much lower level, thereby facilitating a very low-energy write operation. This low-energy write operation, coupled with high stability during the hold, provides for use of a lithographically minimum data bit cell with extremely high stability. As shown in Figure 3, TAS technology enables stability factors that can produce memory chip densities in excess of 10 GB. TAS technology has been proven at the 130 nanometer process node and will enter production in 2011.

Across the industry, there are two writing mechanisms under

development for advanced MRAM technologies, one using locally generated magnetic fields and the other using electron spin polarization (often referred to as spin transfer or spin torque). The TAS concept for stabilization can be used with either of these writing mechanisms, and in both cases resolves the stability barrier looming on the path to very high-density MRAM. Without TAS, each step forward in MRAM density will require innovation in MTJ device structure and writing mechanism to provide acceptable stability and a small bit area. It is not proven that production-worthy commercial MRAM of 64 Mbit and larger densities can be achieved without an advanced stabilization mechanism such as TAS.

Summary

The road to high-density MRAM requires a strong solution for the problem of maintaining data stability while reducing feature size. Most MRAM device approaches are constrained by the challenge of materials and device physics that cannot simultaneously provide stability and a reduced feature size. TAS is a new MRAM technology that decouples the limiting interaction between stability and scaling, thereby opening the door to very high-density MRAM technology that scales reliably to 45 nanometer, 32 nanometer and smaller. ■

About the Authors

Barry Hoberman is the chief marketing officer at Crocus Technology. Mr. Hoberman possesses 30 years of experience within the semiconductor industry. Previously, he was the founder and chief executive officer of inSilicon, a leading semiconductor intellectual property (IP) supplier serving the communications and connectivity markets, which was acquired by Synopsys in 2002. His industry leadership experience also includes CEO positions with Virtual Silicon and TZero Technologies. In addition, Mr. Hoberman has held various product management and development positions at Monolithic Memories and Advanced Micro Devices (AMD). Mr. Hoberman earned Bachelor of Science degrees in electrical engineering and biology from the Massachusetts Institute of Technology and holds 14 U.S. patents.

Ken Mackay is director of magnetic cell engineering for Crocus Technology. With over 20 years combined semiconductor and magnetics experience, including past positions at IBM and Hitachi, Mr. Mackay is responsible for magnetic technology and core memory cell integration at Crocus Technology.