Two-level BEOL processing for rapid iteration in MRAM development

M. C. Gaidis
E. J. O'Sullivan
J. J. Nowak
Y. Lu
S. Kanakasabapathy
P. L. Trouilloud
D. C. Worledge
S. Assefa
K. R. Milkove
G. P. Wright
W. J. Gallagher

The implementation of magnetic random access memory (MRAM) hinges on complex magnetic film stacks and several critical steps in back-end-of-line (BEOL) processing. Although intended for use in conjunction with silicon CMOS front-end device drivers, MRAM performance is not limited by CMOS technology. We report here on a novel test site design and an associated thin-film process integration scheme which permit relatively inexpensive, rapid characterization of the critical elements in MRAM device fabrication. The test site design incorporates circuitry consistent with the use of a large-area planar base electrode to enable a processing scheme with only two photomask levels. The thin-film process integration scheme is a modification of standard BEOL processing to accommodate temperature-sensitive magnetic tunnel junctions (MTJs) and poor-shear-strength magnetic film interfaces. Completed test site wafers are testable with high-speed probing techniques, permitting characterization of large numbers of MTJs for statistically significant analyses. The approach described in this paper provides an inexpensive means for rapidly iterating on MRAM development alternatives to converage on an implementation suitable for a production environment.

Introduction

MRAM may be suitable for replacement of many volatile and nonvolatile memories that are currently in use. Offering performance similar to that of dynamic random access memory (DRAM), density significantly better than that of static random access memory (SRAM), and nonvolatility with greater speed and write endurance than flash memory, MRAM has the potential to dominate the memory market in the near future [1]. **Figure 1** is an example of the most commonly implemented MRAM circuit topology: the so-called one-transistor, one-MJT ("1T1MTJ") "FET cell" [1], wherein one FET is coupled with each single-MTJ memory element. In this cell concept, the read path to measure the resistance of the MTJ is through a transistor, isolating the desired MTJ for relatively high signal-to-noise ratio. Array efficiency is improved by sharing a bit line for both reading and writing, and device isolation is facilitated through the use of two word lines—a "write" word line and a "read" word line. As can be seen in the figure, MRAM offers the advantage of fabrication of the critical magnetic components solely in BEOL structures for reduced cost and flexible integration with CMOS technology. In the figure, a red oval encloses the critical MRAM components: the carefully designed stack of magnetic film elements, the precisely patterned shape for ideal switching behavior, and encapsulation chosen to maximize the magnetoresistance and thermal stability of the devices.

Fabrication of the FET-cell circuit, from the CMOS "front-end" through the MRAM "back-end," can encompass several hundred process steps. The MRAM-critical (red-oval-enclosed) portion of the circuit is a relatively small portion of the entire configuration. After the last standard CMOS step (the M2 wire completion), there remains the need to pattern the shallow vias, the MTJs, the local interconnects, and at least one level of wiring with contact to MTJs and the functional circuitry below. Even for simple functional circuits, five or more photomask levels are required to complete the MRAM-specific portion of the structure, with more than one

0018-8646/06/$5.00 © 2006 IBM

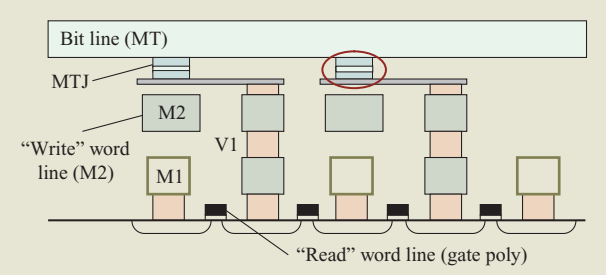

Figure 1

Example of the 1T1MTJ FET cell topology, with two adjacent cells shown atop the silicon CMOS "front-end" structure. The red oval encloses the critical components for MRAM implementation. As cell size is determined primarily by the MTJ and via chain above the via V1, two FETs can be used for each MTJ in order to achieve lower resistance and some redundancy. Thus, the FET gates on either side of a V1 via chain will be connected to the same "read" word line. Wires formed in the first level of metallization (M1) (outlined in green) form a grid at a reference potential. M2 denotes the second level of metallization. The reader is referred to Reohr et al. [1] for a detailed explanation and 3D views of such structures.

hundred process steps between fabrication of the M2 layer and the finished counter-electrode wiring layer (designated here and elsewhere as the "MT wiring layer"). This complexity in processing is needed to fully characterize the functional performance of MRAM memory arrays. However, it would be advantageous to reduce cost and turnaround time when experimenting with the development of improved MRAM devices irrespective of the CMOS portion. For best economy, one would process and test only the following critical steps in MTJ formation: magnetic film deposition, MTJ patterning, and MTJ encapsulation. Ideally, these steps would be performed in a manner essentially identical to that when they are processed in a complete CMOS-functional wafer. In this way, the newly developed steps could be inserted at low risk into the fabrication of a marketable product. There is the potential for significant cost savings in the development of MRAM should such an economical, rapid-turnaround process and test vehicle be implemented.

One possible option for rapid characterization of some critical elements of magnetic stack composition and patterning is through the use of contact atomic force microscopy (C-AFM). In this technique, an electrically conducting tip is used in an AFM; it can be placed directly atop an MTJ to electrically sense the resistance of the MTJ. The details of this technique are described in a companion paper in this journal [2]. Drawbacks of this technique include the manual nature of the data gathering and the need for an MTJ counter-electrode that can easily be electrically contacted by the AFM tip. The lengthy data collection process consequently implies small data sets and relatively poor statistics. The counter-electrode limitations can prohibit the use of certain desirable self-aligned MTJ hard masks and preclude the use of MTJ encapsulation. These restrictions relegate C-AFM to being a research tool, since detailed statistics and realistic processes are a must for the qualification of processes to be utilized in a production environment.

A more suitable alternative for obtaining rapid feedback with test data from production-relevant processing is a "short-loop" subset of the fully integrated functional structure. The short-loop test vehicle contains a complete set of the important process and design elements for magnetic device development, and also includes MTJ counter-electrode wiring to facilitate rapid and automatic characterization with semiconductor-industry high-speed electrical probers. **Figure 2(a)** schematically represents the active elements in the short-loop structure. These elements include a base contact (MA) that is a continuous conducting plane, the MTJs, and the upper MT wiring layer. When it is compared with the fully functional structure of Figure 1, the simplicity of the short loop is apparent. The simplicity comes from removal of extraneous CMOS elements; the most critical elements of MRAM device fabrication are still present.

A series of MTJs fabricated with the short-loop process is shown in **Figure 2(b)**, clearly illustrating the focus on the magnetic devices to the exclusion of the CMOS elements. The patterning method used to create the devices shown in Figure 2(b) utilized a self-aligned counter-electrode (CE) contact between the MTJ and the MT wiring layer.

We present here an overview of the issues involved with the creation of such a short-loop test vehicle (or simply "short loop"), from initial design to final test. The manuscript is organized as follows: first, a discussion of the design issues to explain how one can minimize the complexity of processing while still maintaining useful electrical measurement capability; next, a description of the processing issues involved with the fabrication of the structure; then, the presentation of test data showing the applicability of the short loop; and finally a discussion of future improvements and applications.

Circuit design

As discussed above, to greatly reduce the number of processing steps in the short loop, a continuous base electrode is shared among all devices on a wafer. This imposes requirements on the test site design that limit the flexibility and accuracy with which measurements can be made. For example, the fidelity of a simple four-wire resistance measurement can only be approximated in

the short-loop structure because voltage drops in the conductive sheet can significantly affect the voltage reading. In a conventional design, one minimizes measurement errors by sensing the potential of the base MTJ electrode as close as possible to the MTJ under test. However, in the two-photomask short-loop framework, the only means for contacting the MA plane is through another MTJ. The requirement that MTJ patterning closely emulate that used in fully functional wafers dictates that one maintain a relatively wide MTJ-to-MTJ spacing, as in product arrays. This wide spacing can result in a substantial potential drop across the MA plane and consequent measurement errors. For simpler analysis of test results, a related attribute of the design is that the effect of the MA voltage drop be similar for the wide range of MTJ sizes and shapes being tested in order to facilitate the study of the scaling properties of the MTJs.

To satisfy the aforementioned desires, we have devised a scheme to take advantage of symmetry to define a reference potential with a design illustrated in **Figure 3**. Two identical MTJs are separated by a distance similar to that used in product design. When an equal amount of current is sourced into one MTJ and drained from the other, and no other currents are flowing into or out of the MA plane, the points in the MA plane that are equidistant from the two MTJ devices acquire an electrical potential that is the average of the potential of the two lower MTJ electrodes. Errors in this approximation result from narrow spacing between MTJs and from any variability in MTJ shapes and resistances.

The potential of the symmetry line can be conveniently detected by connecting one voltage-sensing contact to an MTJ that is far distant from the pair of MTJs under test. This is because the potential at infinite distance is equivalent to the average of the two base electrodes under test. This distant contact is implemented in the short loop by connecting the middle pad (#13) to a base electrode through an MTJ, as illustrated in **Figure 4**. Even with the ideal symmetrical case, there will be an error in the MTJ resistance measurement due to the series contribution of MA sheet resistance between the MTJ and the line of symmetry. Since the MTJ cannot be treated as a point source, an analytical solution of the problem becomes very complicated. A numerical simulation was therefore used to estimate an upper limit of the error due to finite MA sheet resistance. For a 0.50-μm-diameter circle centered 1.16 μm from the center of its mirrored counterpart, an example barrier resistance–area (RA) product of 600 Ω-μm^2 and MA sheet resistance of 85 $\Omega/\square$ were used in the simulation. The true resistance of the MTJ was 3.148 kΩ, taking into account the discretization error of the simulation grid. For the perfectly symmetric case, the resistance of the MTJ plus the effect of finite MA resistance would give an approximate answer of

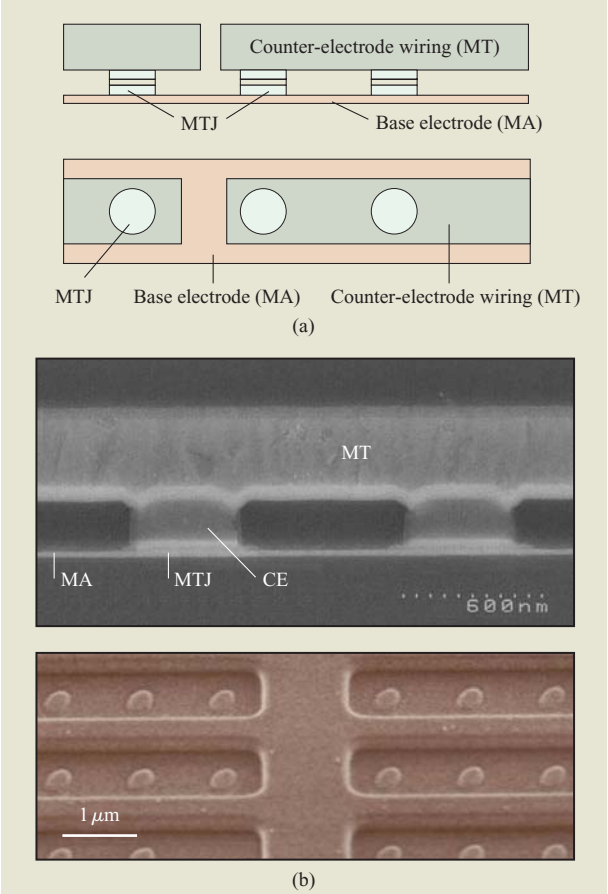

Figure 2

(a) Schematic drawing showing the active elements of the short-loop configuration, with cross-sectional view (top) and top-down view (bottom). (b) SEM images, including (top) a cross section of a completed short-loop wafer, corresponding to the right side of the schematic in part (a), showing the continuous base electrode (MA), and magnetic tunnel junctions (MTJs) with metallic counter-electrodes (CEs) contacting the upper (MT) wiring level. The MTJs are approximately 400 nm wide. A tilted top-down SEM image (bottom) reveals the structure after the MT trench is etched but before filling with MT metal. Elliptical MTJ counter-electrodes are exposed at the bottom of the MT trenches and will be contacted by the ensuing MT copper wiring.

3.172 kΩ—less than 1% error from the finite MA resistance. Similar computations showed that the error is less than 1% for MTJs smaller than 0.5 μm in diameter, and increases to approximately 9% for 2.0-μm-diameter MTJs. The measurement error for such large MTJ devices is inherently large because the finite MA sheet resistance induces a nonuniform current flow below the device [3]. Asymmetry is generally less of a problem than MTJ device size. Considering the case in which one MTJ

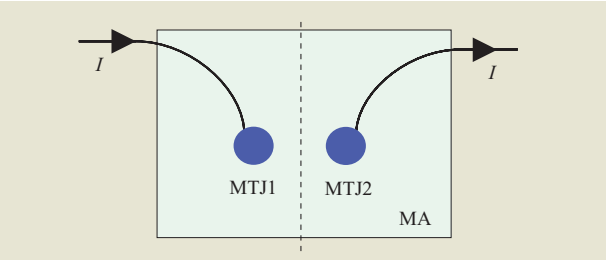

Figure 3

Equipotential line of symmetry between two identical MTJs.

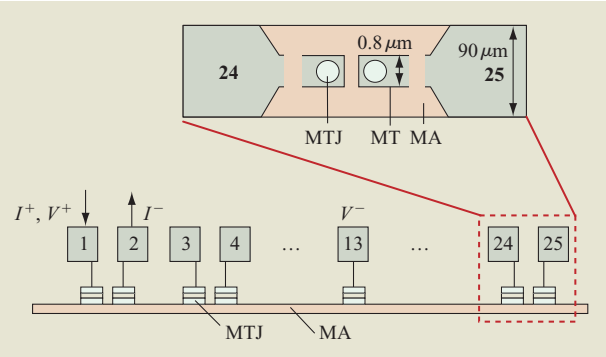

Figure 4

The structural layout for the symmetric measurement of two MTJs utilizes a far-removed contact to sense the potential at the line of symmetry between the two MTJs under test. In this figure, MTJs connected to pads 1 and 2 are under test. Through an MTJ, pad 13 is the distant measurement contact. A standard 25-pin probe card spaces the pads with a 125-μm period, and the paired MTJs are separated by less than 1 μm. The wiring to the MTJ counter-electrodes and the contact pads are formed in the MT metal layer. Because of the vast difference in scale, the dimensions in the figure are compressed for the contact pads relative to the MTJs, as indicated by the dimensioning in the detail view. The entire 25-pad structure is approximately 3 mm long and 90 μm wide.

is a short circuit, the error in measurement of the other MTJ is approximately twice the error computed for the perfectly symmetric case (current in the MA plane must travel approximately twice as far). Simple algorithms can be used to refine the data to correct for the bulk of this error, given inputs of tunnel barrier RA product and MA sheet resistance, which can be accurately measured by test structures on the wafer.

An added benefit of the symmetric circuit design is the need for only n contact pads for approximately n devices. Since no transistors are fabricated in the short-loop wafers, the number of testable devices—and hence the statistical significance—is essentially limited by the number of probe pads that can be fit into each die. The vast majority of electrical tests are performed with industry-standard source-measure units (SMUs) at low frequency (essentially dc). It is therefore desirable to use the symmetric design to minimize the number of probe pads required for each testable MTJ. In principle, one can test more than a million separate MTJs on a single 200-mm-diameter wafer. The symmetric MTJ configuration provides the necessary device density and measurement accuracy for rapid iteration of materials and process experiments needed to refine the process integration for implementation in fully functional wafers.

Process integration of the short loop

Many of the short-loop processes can be translated directly from fully integrated wafer processes, but the inhibition of delamination requires special care. In conjunction with the diagram in **Figure 5**, we present below a list of the basic process steps in short-loop fabrication, with coverage of the important considerations in each step. Although the reduction in number of process steps makes the short loop an attractive development vehicle, the topographical flatness of the short-loop structure makes it prone to delamination. After coverage of the basic process steps, we discuss in more detail methods by which we prevent delamination and generally improve yield of the MTJ devices.

Basic process steps

1. *Preparation of the substrate (Figure 5–1)*. To obtain ideal MTJ device performance, it is critical to achieve atomic-scale flatness over areas of the order of the size of an MTJ. This reduces Néel coupling effects and makes for a well-controlled RA product for the MTJ devices [4]. Substrates for the short-loop process are thus generally prepared either with a careful silicon wafer oxidation/cleaning or with the deposition of a dielectric such as silicon nitride on the silicon substrate, followed by a chemical–mechanical planarization (CMP) step to smooth the surface. The latter option most closely resembles the fully integrated wafer structure, but MTJ performance has not been found to depend on which of the two techniques described above is used.
2. *Magnetic film stack deposition (Figure 5–2)*. The magnetic stack is generally composed of the following layers: a (typically nonmagnetic) seed layer to promote proper polycrystalline growth (e.g., Ta), an antiferromagnet for strong pinning of the reference layer (e.g., PtMn or IrMn), an antiferromagnetically exchange-biased pair of

ferromagnets (e.g., CoFe/Ru/CoFe), the insulating tunnel barrier (e.g., Al_2O_3 or MgO), a switchable free layer (e.g., CoFeB/Ru/CoFeB), and a suitably stable cap and hard mask layer (e.g., Ta, TaN, or TiN). Detailed discussions of the motivation for such a complex stack can be found in [5]. The subsequent patterning of the MTJ introduces device-to-device isolation in the counter-electrode (the conductive portion above the tunnel barrier), but maintains electrical continuity between all devices in the base electrode (the conductive portion below the tunnel barrier). Often negligible in fully integrated wafers, the resistance of the base electrode after MTJ patterning is germane to the short loop. The use of a continuous planar base electrode incurs additional measurement error at final electrical testing if the base electrode possesses a high sheet resistance. Subject to the constraint of emulating the stack used in fully functional wafers, the magnetic stack of the short loop will therefore include thick or low-resistivity films beneath the tunnel barrier. One cannot arbitrarily thicken materials to meet this requirement, as it can further separate the free layers from the write wire in fully integrated wafers, and generally involves a tradeoff against increased Néel coupling from surface roughening in thicker films.

3. *Tunnel junction patterning (Figure 5–3)*. A commonly used, straightforward approach to patterning the MTJs is through the use of a conducting hard mask. The conducting mask is later utilized as a self-aligned stud bridging the conductive MT wiring to the active magnetic films in the device. Such a processing scheme is among the simplest and fastest ways of creating and contacting the MTJs, making it an ideal approach for use in the short loop. Choices for the hard mask are numerous, with necessary characteristics being etchability and a resistance that is negligible when compared with MTJ resistance. Refractory materials commonly used in the semiconductor industry such as Ta, TaN, and TiN are suitable as masks for MTJ patterning. The MTJ shapes are defined in the hard mask by transfer from a first photomask level in a process such as the following: apply resist/expose and develop/RIE through hard mask/strip resist. The pattern is further transferred downward to penetrate to (or through) the tunnel barrier, leaving behind a low-resistance base layer which covers the entire wafer.
4. *Dielectric encapsulation (Figure 5–4)*. The encapsulation of the etched MTJs protects them while at the same time forming the environment in which the MT wiring level will be created. The choice of encapsulation is determined from three requirements: a) it must not damage the MTJs; b) it must adhere well to the substrate; and c) it should closely emulate the interlayer dielectrics (ILDs) that would be used in a fully integrated wafer process. Damage to the MTJs can arise from chemical interactions and thermal stress. Standard semiconductor-industry dielectrics typically are deposited or cured at temperatures around 400°C, whereas degradation in submicron MTJs can set in at temperatures below 350°C. Thus, a major challenge to the integrator of MRAM devices is the development and utilization of suitable low-temperature dielectrics. Adhesion of the dielectric to the substrate can be particularly problematic given the characteristics of the magnetic films being used. Noble-metal-containing antiferromagnets can be particularly difficult to adhere to, and, if exposed by the etching used for MTJ patterning, can require specialized surface-cleaning or surface-preparation techniques to promote adhesion to the encapsulating dielectric. The dielectric thickness is chosen such that

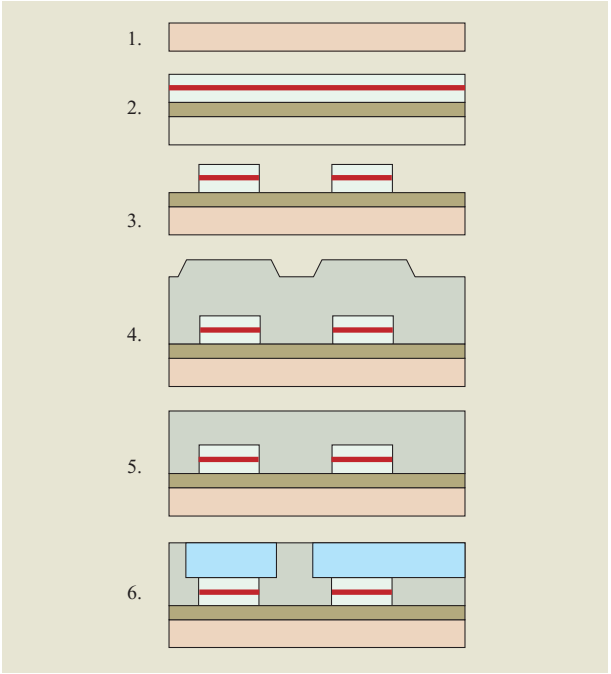

Figure 5

Schematic cross-sectional representation of the basic steps involved in fabricating the short loop. The MTJ tunnel barrier is represented by the red lines.

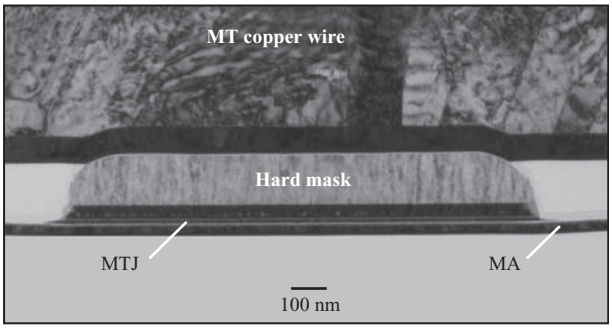

Figure 6

TEM cross-sectional image of a completed MTJ device, showing clearly the continuous bottom (MA) contact, the multilayer magnetic film stack forming the active MTJ region, and a thick hard mask serving as a counter-electrode contact between the MTJ and the MT copper wiring. The tunnel barrier is the bottom-most bright horizontal line, and in this device forms the stop layer for the MTJ etch. (TEM courtesy of Philip M. Rice, IBM Almaden Research Center, San Jose, CA.)

it will be thick enough to provide the environment for the wiring level above the MTJs.

5. *Planarization (Figure 5–5)*. To facilitate industry-standard damascene copper wiring, the wafers generally undergo a gentle dielectric CMP process at this stage. The purpose of the CMP is to remove topography from the surface that is caused by the underlying MTJs. This step is also the first check of the adhesion of the dielectrics to the underlying metal films, as well as the cohesion of the metal films to each other. If the encapsulating dielectrics are suitably planarizing in their deposition, or the upcoming MT trench etching is formulated to aggressively target protruding features above the planar "field" regions, this CMP planarization step can be eliminated for faster turnaround time and potentially higher yield.
6. *Wiring (Figure 5–6)*. After completion of the critical steps required for MRAM development (layer formation, patterning, and encapsulation), the wiring is instituted in the simplest manner consistent with the available tooling. Relying on well-established semiconductor-industry techniques, a photomask-defined trench is etched into the dielectric with RIE, to be filled with a liner and high-conductivity copper. The depth of the trench is sufficient to expose a portion of the conducting hard-mask stud (the counter-electrode), while not so deep as to create a short circuit to the planar base electrode. Endpointing during the trench RIE can facilitate the proper choice of trench depth even for relatively thin hard-mask films. After the trench etching and a suitable cleaning step, the wiring liner film is deposited, along with a thin copper seed layer. This deposition is followed by the electroplating of copper to completely fill the trench and provide enough overburden so that the ensuing CMP step will planarize the metal coincident with the surface of the dielectric. This final CMP step can be aggressive enough to cause shear failure of the films on the wafers, and care must be taken to prevent such delamination. A post-polish cleaning of the wafers is the final preparation step before electrical testing. Illustratively, **Figure 6** shows a TEM cross section of a completed MTJ device after final processing.

Film adhesion

The largest yield detractor in the short loop arises from the relatively low shear strength of the films and interfaces and the lack of substantial topographical "footholds" which could inhibit shear failure. The problem is worsened by the use of antiferromagnetic alloys such as PtMn and IrMn with substantial noble-metal content (offering poor interfacial adhesion). More than the internal film stresses, we find that the considerable shear forces applied during the CMP steps initiate delamination. The weakest region is typically the dielectric/metal interface formed during encapsulation of the patterned MTJs. The failure of this interface is apparent in **Figure 7(a)**, where sample preparation during cleaving for cross-sectional analysis has pulled the dielectric away from the underlying metal layer. In the foreground of the image remain patterned MTJs (as would be seen after the third step discussed above). A portion of the dielectric encapsulation has broken off and been pulled away. In the rear portion of the image, the dielectric/metal interface that remains is seen to be failing due to inadequate adhesion.

The adhesive strength of the interfaces can be increased dramatically through a better choice of encapsulating dielectrics and improved post-MTJ-etch cleaning techniques (details of which are currently proprietary). **Figure 7(b)** shows SEM imaging of a sample similar to that from Figure 7(a), but with improved interfacial adhesion. Aggressive sample preparation was utilized to induce interface failure (i.e., some of the encapsulating dielectric and counter-electrode wiring were ripped from the sample). The remaining structure reveals cracking in the encapsulating dielectric, suggesting that cohesive failure of the dielectric develops under the harsh conditions needed to induce interface failure. As can be seen in the foreground of the image, a substantial portion of the magnetic film stack has been pulled free along

with the dielectric, suggesting that the dielectric/metal interfacial adhesive strength has improved enough to be comparable to the intrinsic interlayer adhesive strength of the magnetic stack.

Improving adhesion through materials choices and cleaning is an effective but restrictive manner of improving yield. Often, dielectric MTJ encapsulants must be chosen on the basis of the way in which they stabilize the MTJ for best performance. For instance, encapsulants that inhibit manganese diffusion from the antiferromagnet toward the tunnel barrier can facilitate higher thermal stability for the devices [6]. Since these encapsulants may not always be the best from an adhesion viewpoint, additional techniques must be employed to prevent yield loss through delamination. One such technique is the reduction of shear forces during CMP processing. CMP techniques have evolved with the semiconductor industry's push toward low-K dielectrics. The low-K dielectrics tend to be weaker than traditional silicon oxides, and they necessitate gentler CMP processing to prevent film failure. These techniques include low-downforce processing and electrochemical CMP (e-CMP) [7]. Such CMP techniques can be utilized effectively to minimize shear forces that delaminate films in the short-loop processing.

Another technique employed to strengthen the short-loop structure is the creation of topography that gives the dielectric a better "toehold" with which to grip the underlying metal. Extensive use of MTJ device fill is used to create such toeholds throughout the wafer, including in particular the kerf regions. A 10% or greater local filling factor, with no unfilled regions greater than 50 μm $\times$ 50 μm in size, is found to inhibit the nucleation of delamination during CMP processes. Care is taken to print the partial chips at the edge of the wafer so as to extend MTJ feature topography to the edge of the wafer. Doing so prevents delamination from nucleating at the wafer edge. The lithography exposure dose is reduced on partial chips at the wafer edge, where MTJ shape and size are not critical. The lower dose ensures the printing of a substantial number of the MTJs even as the wafer flatness degrades near the wafer edge.

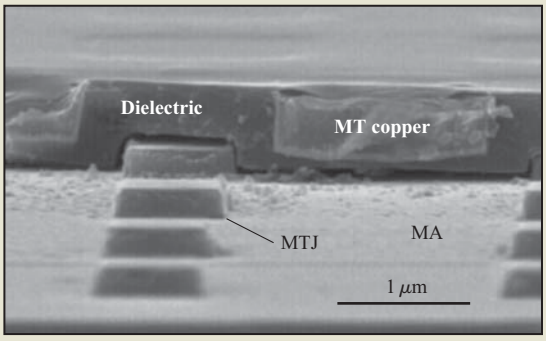

(a)

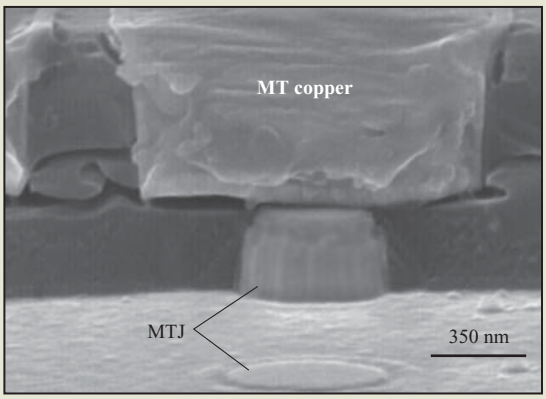

(b)

Figure 7

(a) Tilted-angle SEM image of the surface of a short-loop wafer after final process steps have been completed. Failure of the dielectric/metal interface was induced during sample cleaving in preparation for SEM analysis. Note that the MTJs in this micrograph are inactive "fill" structures inside a probe contact pad, and thus were not configured to make contact with the MT wiring. (b) Illustration of how the use of alternate encapsulation materials and improved cleaning methods can improve dielectric adhesion appreciably. Aggressive sample attack has induced delamination of the dielectric from the metal surface, but not before cohesive failure in the dielectric has occurred.

Dielectric encapsulation

In addition to being a critical element in the performance of the MTJs, the dielectric encapsulation plays a critical role in the integration of the counter-electrode wiring levels. Even before the dielectric is deposited, however, one must pay careful attention to the *in situ* preclean performed with the sensitive edges of the MTJ stacks exposed. As in industry-standard copper damascene processing, the use of hydrogen or NH_3 plasmas to preclean the structure can be very effective at improving the adhesion necessary for successful CMP, but can chemically and physically alter the exposed structures. The precleaning, to some extent, physically etches the MTJ, adding a layer of complexity to the critical procedure of MTJ patterning. Constituents in the cleaning plasma such as oxygen and nitrogen can diffuse into the MTJs, increasing the tunnel barrier resistance at their edges. These in-diffusing elements are also found to alter the magnetic characteristics of the stack films—for instance, changing the coercivity and the direction of intrinsic anisotropy. The *in situ* precleaning and, in a similar manner, the dielectric material itself, are of critical importance to the operation and thermal stability of

MRAM devices. An integration scheme must be found for which the cleaning and dielectric encapsulation are compatible with the requirements of both MTJ performance and functional wafer integration. This is greatly complicated by the large number of choices and the large parameter space related to precleanings and dielectric materials. This situation illuminates one of the strengths of the two-level short loop: It facilitates rapid iteration of material choices and processing parameters at relatively low cost.

Silicon nitride and similar compounds are desirable for their adhesion to the MA and MTJ metal surfaces, and for strong interfacial bonds that inhibit the migration of metal atoms along the dielectric/metal interfaces. Such metal migration is one well-documented cause of MTJ thermal degradation, and can limit processing temperatures in patterned MTJ devices to less than 350°C [6]. Unfortunately, the deposition of silicon nitrides with industry-standard plasma-enhanced chemical vapor deposition (PECVD) techniques requires the use of expensive high-density-plasma tooling and process temperatures of 400°C or higher to achieve a semblance of conformal fill around the protruding MTJ studs. Worse for higher-aspect-ratio structures and lower-temperature deposition, seams and voids are found in the silicon nitride dielectric at the juncture of films growing from two nonplanar surfaces. These seams and voids can degrade reliability and yield, in part because of poor enclosure of the ensuing MT wiring. In addition, the high dielectric constant of silicon nitride films is undesirable for the integration of MRAM into high-speed, low-power CMOS circuitry. For the aforementioned reasons, even though silicon nitride and related compounds can be deposited at low temperatures with good interaction with MTJs, the compounds are not suitable as an environment for MT wiring.

Silicon oxide compounds are desirable because they are more benign with respect to the magnetic behavior of certain stack film choices, and offer the potential of low-temperature deposition of conformal films. The use of TEOS (tetra ethyl ortho silicate) as a precursor in the deposition of silicon oxide films is known [8] to offer the benefits of a relatively inert depositing species which can readily diffuse into spaces adjacent to high-aspect-ratio structures, even at temperatures below 250°C. Low-temperature deposition of TEOS-based silicon oxide in industry-standard tooling is used effectively as a seam/void-free MTJ encapsulant and as an integration-friendly environment for housing the counter-electrode wiring layer. The main drawbacks to the use of silicon oxide encapsulants relate to thermal stability and adhesion problems.

The adhesion problems can generally be overcome through the use of multilayer dielectrics or with more complicated "spacer" techniques. The multilayer approach incorporates a thin layer of well-adhering, MTJ-compatible dielectric (such as silicon nitride) with an overlayer of thick, conformal dielectric (such as TEOS-based silicon oxide). The silicon nitride is kept thin enough so that miniscule seams and voids have no impact on yield, but thick enough to form a suitable MTJ encapsulant. For dielectrics to serve as MTJ encapsulants even with poor dielectric/metal adhesion, one can implement a spacer approach as follows: A conformal coating of the dielectric is deposited atop the MTJ; this deposition is followed by a blanket, anisotropic dry etch to remove the dielectric from the horizontal surfaces on the substrate while leaving behind a wedge of material in the area shadowed by the high-aspect-ratio MTJ stud; next, an adhesive film is deposited and covered with a thick conformal dielectric to serve as the environment for the MT wiring layer.[1] Such spacer and multilayer techniques can also be applied to enable the use of spin-on and other low-K dielectrics which can be leveraged to eliminate a CMP step and enable more direct integration of MRAM with advanced CMOS processing.

Trench RIE

Although not necessarily a process that is critical to the performance of isolated MTJ devices, the formation of the MT wiring layer involves some subtle issues that must be taken into account when fabricating short-loop wafers. A fully functional wafer would utilize the MT wiring layer for the bit-line readout and for field generation used during the "write" operation. Since the strength of the write field depends on the wiring distance from the free layer to the MT, minimizing the operating power of the circuit dictates that the MT wires be positioned close to the MTJ free layer. This is traded off against an increased process window obtained by spacing the MT wires farther from the MTJ free layer through the use of a thick conductor (for example, a self-aligned metallic hard mask or a via between the MT wire and the cap layer of the magnetic stack). The process window is reduced by thickness nonuniformity of the encapsulating dielectric, by polishing rate nonuniformity during the dielectric planarization step, and by etching rate nonuniformity in the reactive ion etching of the MT trench. For ease of process transference, it is advantageous to minimize differences in processing these steps between the short-loop and fully functional wafer methods of integration. Here again, the short loop shows its importance: By rapidly iterating the process parameters in dielectric deposition, polishing, and trench etching, one can converge on a solution that can be inexpensively

[1] S. Kanakasabapathy and I. Kasko, "Decoupling Tunnel Junction Encapsulation and Interlevel Dielectric Deposition in MRAM Crosspoint Cell Architecture," unpublished; see *www.IP.com*, Document ID IPCOM000124658D (2005).

incorporated into fully functional wafers to minimize the power required to write the MTJ devices. The short-loop test vehicle has enabled us to couple (e.g., with wafer rotation speeds and back-pressure) the CMP removal rate inversely to the thinner areas of nonuniform dielectric deposition for minimized dielectric thickness nonuniformity as the wafers enter MT trench RIE. With improved dielectric thickness uniformity, the MT wiring can on average be brought closer to the MTJ free layers, reducing the write currents used in fully functional wafers. Any uniform variation in dielectric thickness is sensed by an endpointing technique used during the MT trench RIE. The detection of certain plasma constituents halts the etching when the top surfaces of the counter-electrode metal are exposed or when the lowest layer of dielectric encapsulation is etched.

Electrical and magnetic characterization

The reduction in complexity of processing the short loop (relative to fully functional wafers) adds some complexity to the interplay between circuit design and the interpretation of test results. There is also an additional burden placed upon the test sector to enable the simplified short-loop processing: the requirement of off-chip magnetic field generation. To switch toggle-mode devices [9], one requires well-controlled timing of magnetic fields along two axes. The simple one-layer (MT) wiring in the short-loop structure is too difficult to use for more than one axis of field generation, so an external (off-chip) field source must be used. No write operations are driven with circuitry on the wafer. This situation results in a significant shift of burden from the processing sector to the test equipment, and comes at a cost: The intricacies of on-chip device writing are not fully tested by this short-loop structure. This wiring is considered a second-tier issue, less difficult to develop when compared with the more critical issues of MTJ stack composition, etching, and encapsulation. Indeed, comparison of device performance between short-loop wafers and similarly processed fully integrated wafers suggests that this is a reasonable tradeoff: On-chip switching is sufficiently comparable to externally generated magnetic field switching. What the short loop and externally generated fields do enable is rapid feedback with automated testing of statistically significant numbers of devices.

External field generation

For rapid testing of the individual MTJs in a short-loop test vehicle, external generation of "write" magnetic fields must be done in an efficient manner so that rapid field ramping and magnet cooling are practical. We have adapted a toroidal magnet as an efficient generator of arbitrarily directed in-plane magnetic fields. Its relatively small size (of the order of 10 cm) allows it to be placed close to the wafer, and within existing industry-standard semiconductor test probers. Referring to **Figure 8**, if the X coils are activated, flux is driven in opposition in the two sides of the core. The convergence and divergence points act as sources (positive and negative magnetic charges) that create a magnetic field at the sample location close to the center of the core. Similar behavior is found from the Y coils. As long as the core is far from saturation, the fields from the X and Y coils add as vectors. As a result, the field at the sample can be set in any in-plane direction by adjusting the currents in the X and Y coils. A key feature of this design is low remanent fields at the sample. If the coils are not activated, the flux from any remanent magnetic charge in the core closes through the low-reluctance core and does not create a field at the sample location. For the most rapid testing, the magnet is cooled with flowing air.

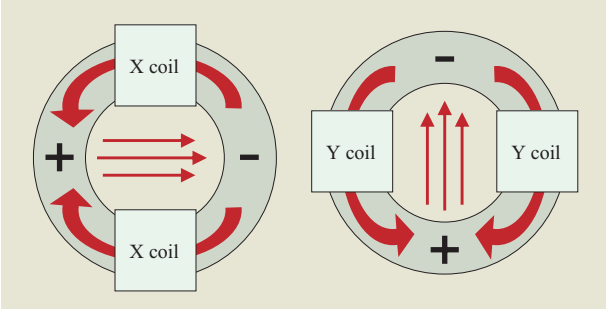

Figure 8

Toroidal magnet configuration for generation of external fields to switch the state of MTJs of a short-loop test vehicle. The inner diameter of the toroid is roughly 5 cm, and the distance from the bottom of the toroid to the wafer surface is approximately 0.5 cm.

Prober operation

One drawback to the lack of transistors to drive the devices on short-loop wafers is the need for approximately one probe pad touchdown for every device tested. This necessitates a very large number of prober touchdowns to gather reliable statistics for a given process being tested. Residuals on the MT copper pads on the wafer tend to contaminate the probe contact needles after hundreds of touchdowns. A particular problem is the copper-passivating benzotriazole (BTA) layer that is formed during MT liner CMP to prevent excessive oxidation of the copper surface [10]. After repeated prober touchdowns, this BTA film coats the probe contact needles and degrades the conductance of the contact interface between the needle and the MT

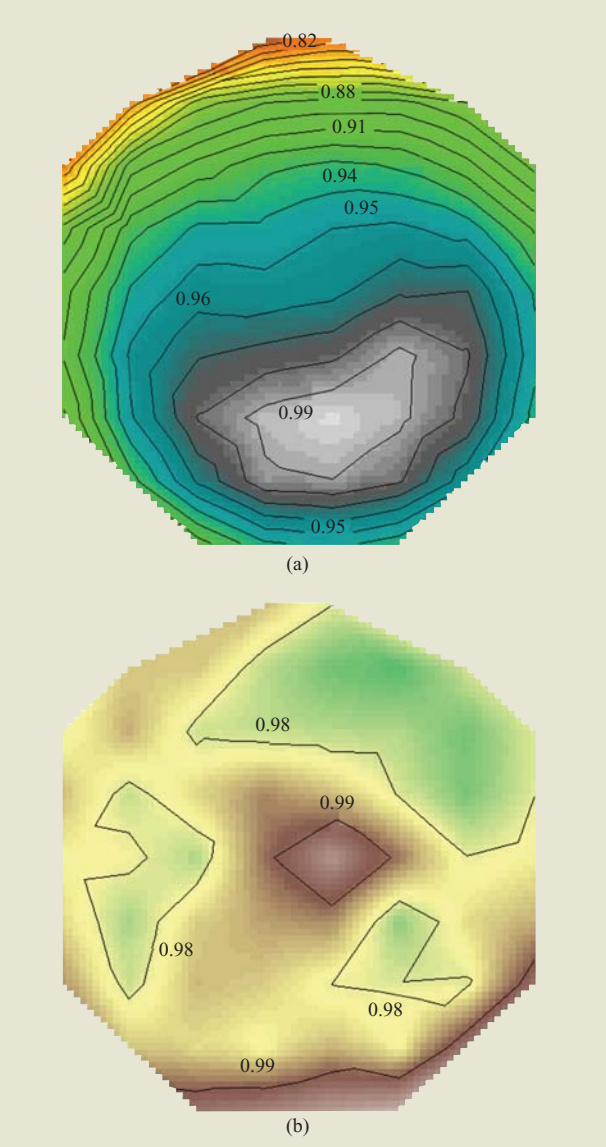

Figure 9

Normalized MA sheet resistance (Ω/□) plotted as a function of wafer position across a 200-mm-diameter wafer, clearly showing (a) a nonuniform etching front from the tool used to pattern the tunnel junctions, and (b) a more uniform result, obtained when use was made of an etch-stop layer developed with the short-loop test vehicle.

copper pads. The large number of prober touchdowns makes simple in-process abrasive cleaning of the needles impractical and rather ineffective. Semiconductor-industry processing often dissociates the BTA layer before testing by subjecting the wafers to a brief 350°C anneal. However, given the thermal sensitivity of the MTJ devices, we have instead employed a wet chemical means for removing the BTA before testing. A two-step AZ–NMP[2] chemical strip has proven to be extremely effective in improving probe contact without harming the MTJ devices. Suitable tuning of prober needle downforce ensures good contact without device destruction or excessive noise.

Illustrative examples of advances from short-loop experiments

The short-loop test vehicle enables rapid process development and characterization without the overhead of expense and delay associated with fully integrated wafer processing. In addition to the complexities of magnetic device operation, the short loop is effectively used to tweak processes for optimal uniformity and yield across entire wafers. The MA sheet resistance wafer maps in **Figure 9** illustrate how the implementation of an etch-stop layer in our magnetic film stack was effective in improving the across-wafer uniformity of a certain MTJ patterning etch. The etch-stop layer was placed between the hard mask and the magnetic free layers, and served to compress nonuniformity originating in the hard-mask etching. Thus, characteristics of the "inactive" layers of the short loop were used to adjust processing for best performance of the active magnetic devices.

One of the more difficult issues to overcome when patterning and encapsulating MTJs is the thickening of the tunnel barrier near the edges of the MTJs. For example, the use of certain etchants and encapsulants can enable oxygen or nitrogen to diffuse into the tunnel barrier from the edge of the device, creating a ring of highly resistive barrier material surrounding the nominal-thickness tunnel barrier at the center of the device. This loss of electrically active material can be gauged by measuring numerous MTJs of different areas and calculating the effective loss of area with a straightforward resistive model. An ideal process would exhibit a resistance that scales with the inverse of the nominal design area of the MTJs, but a process with tunnel-barrier edge degradation would show much higher resistance than expected for smaller-area MTJs. **Figure 10** shows magnetoresistance (MR) and resistance data for a series of device sizes on short-loop wafers processed with MTJ patterning and encapsulation techniques that were developed using short-loop wafers. The resistance data indicate that the MTJs are relatively immune to edge degradation. Each data point is an average of approximately 150 yielding devices from three separate wafers using this same process. The resistance scales roughly as the inverse of the area, and the magnetoresistance is roughly constant for all device sizes. We attribute the outlying point at very small device size

[2]NMP: N-methylpyrrolidone; AZ: a proprietary chemical resist stripper manufactured by the Clariant Corporation, P.O. Box 3700, 70 Meister Ave., Somerville, NJ 08876.

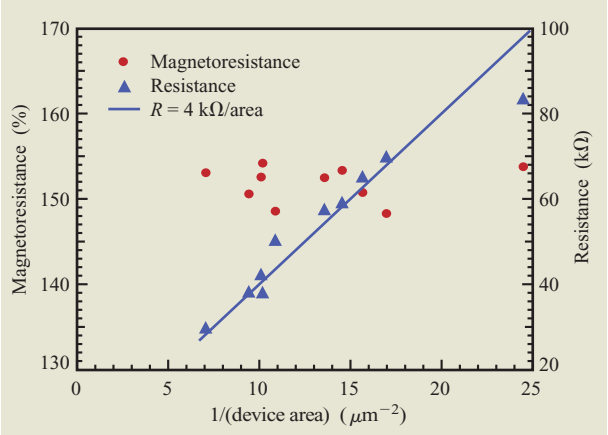

Figure 10

Magnetoresistance and resistance as a function of inverse device area. Highlighted by the solid-line "guide to the eye," the increase of resistance with decreasing device size is consistent with a well-defined MTJ edge. The edge definition is relatively insensitive to encroaching oxidation during dielectric encapsulation of the MTJs. Note that the solid line *does* go through the origin (per discussion in the text).

(equivalent to a 230-nm-diameter circle) to an uncertainty in the measurement of the active device diameter of approximately 20 nm. This arises in part from line edge roughness of the MTJs and is more significant for the smallest devices.

It is known that the edge degradation can be accelerated by annealing an MTJ at moderate temperatures (near 300°C). The short loop was used to converge rapidly on processing techniques that demonstrated much better thermal stability. **Figure 11(a)** shows data from a non-optimal MTJ patterning process before and after annealing. This is in contrast to **Figure 11(b)**, which shows data from an improved MTJ patterning process, exhibiting much improved MTJ stability with the same annealing.

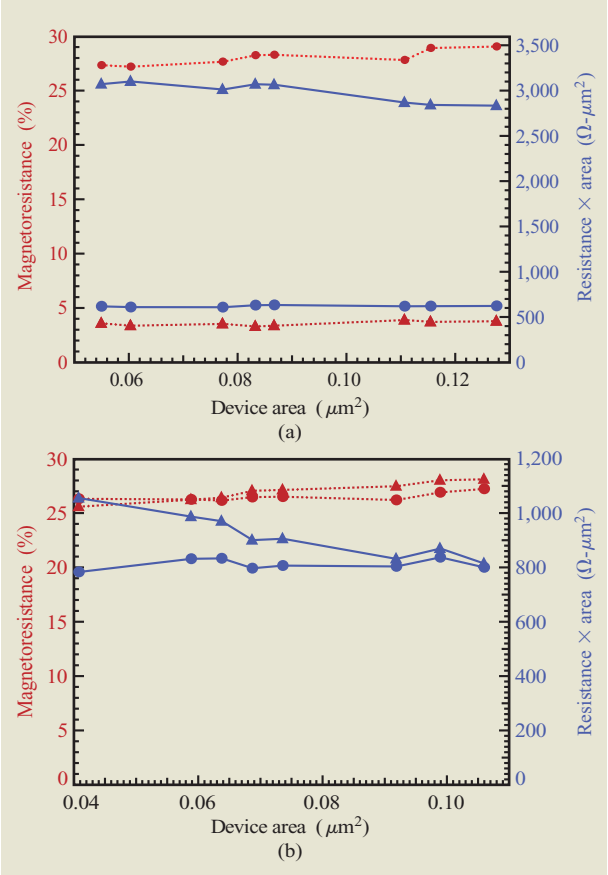

Figure 11

(a) MTJ magnetoresistance (*MR*) before and after thermal stress annealing, showing a dramatic increase in resistance times area (*RA*) and a concurrent decrease in *MR*. *MR* is represented here as dotted red lines, and *RA* as solid blue lines. Pre-annealing data are depicted by the circles and post-annealing data by the triangles. (b) Same data, but for an improved method of MTJ device patterning—*MR* remains essentially unchanged, and *RA* increases slightly, but only for the smallest devices. Note that both are similar *before* annealing for the two processes, suggesting that the observed degradation in part (a) is largely due to the use of different patterning processes rather than a difference in intrinsic film properties.

Concluding remarks

The use of a two-photomask short-loop test vehicle has proven to be invaluable for inexpensive and rapid development of MRAM technology within IBM. The short loop has proven to give a trustworthy and statistically relevant indication of how the critical MRAM processes behave in fully integrated wafers. Rapid turnaround with functional front-end CMOS circuitry is acknowledged as a good alternative for development in a subsequent piloting operation, but is considered overly expensive for the R&D environment to which the present work has been applied. The short-loop test vehicle described here omits extraneous CMOS circuitry in favor of faster turnaround and less expensive process integration.

Modifications of the short-loop design described could be developed to focus on different aspects of the MRAM technology. For instance, more accurate resistance measurement could be implemented with an adaptation to a four-wire configuration in place of the symmetric three-wire design described above. With two pads wired to each MTJ of the pair, accuracy would be improved, but at the expense of halving the number of testable

devices. Designs incorporating large numbers (20–40) of MTJs in parallel could utilize the symmetric three-wire design and be used as the basis for rapidly searching through huge numbers of MTJs to examine MTJ yields at the 99.999% level, thus being applicable to process development of large MRAM memories.

Acknowledgments

The authors wish to thank the following IBM employees for contributions to this work: John Hummel for integration advice and facilitation of tooling, Don Horn and the IBM Almaden Research Center for magnet design, Philip Rice for TEM imaging, Robert Thompkins for SEM imaging, Richard Conti and Michael Smits for dielectrics development, Gerald Gibson and Timothy Dalton for RIE assistance, Michael Lofaro for advances in CMP, Kenneth McCullough for wet chemical etching and cleaning, and Steven Steen and William Enichen for refinement of lithography. We also wish to acknowledge that the work summarized here was built on prior work done with Infineon Technologies within the MRAM Development Alliance and on earlier MRAM work at IBM that was supported in part by DARPA.

References

1. W. Reohr, H. Hoenigschmid, R. Robertazzi, D. Gogl, F. Pesavenot, S. Lammers, K. Lewis, C. Arndt, Y. Lu, H. Viehmann, R. Scheuerlein, L.-K. Wang, P. Trouilloud, S. Parkin, W. Gallagher, and G. Mueller, "Memories of Tomorrow," *IEEE Circuits & Devices* **18**, 17–27 (2002).
2. D. W. Abraham, P. L. Trouilloud, and D. C. Worledge, "Rapid-Turnaround Characterization Methods for MRAM Development," *IBM J. Res. & Dev.* **50**, No. 1, 55–67 (2006, this issue).
3. R. J. M. van de Veerdonk, J. Nowak, R. Meservey, J. S. Moodera, and W. J. M. de Jonge, "Current Distribution Effects in Magnetoresistive Tunnel Junctions," *Appl. Phys. Lett.* **71**, 2839–2841 (1997).
4. B. D. Schrag, A. Anguelouch, S. Ingvarsson, G. Xiao, Y. Lu, P. L. Trouilloud, A. Gupta, R. A. Wanner, W. J. Gallagher, P. M. Rice, and S. S. P. Parkin, "Néel 'Orange-Peel' Coupling in Magnetic Tunneling Junction Devices," *Appl. Phys. Lett.* **77**, 2373–2375 (2000).
5. S. S. P. Parkin and M. G. Samant, "Magnetic Random Access Memory with Thermally Stable Magnetic Tunnel Junction Cells," U.S. Patent 6,518,588, 2003; S. S. P. Parkin, "The Magic of Magnetic Multilayers," *IBM J. Res. & Dev.* **42**, 3–6 (1998); S. S. P. Parkin, N. More, and K. P. Roche, "Oscillations in Exchange Coupling and Magnetoresistance in Metallic Superlattice Structures: Co/Ru, Co/Cr, and Fe/Cr," *Phys. Rev. Lett.* **64**, 2304–2306 (1990); J. M. Slaughter, R. W. Dave, M. DeHerrera, M. Durlam, B. N. Engel, J. Janesky, N. D. Rizzo, and S. Tehrani, "Fundamentals of MRAM Technology," *J. Supercond.—Incorporating Novel Magnetism* **15**, 19–25 (2002); S. S. P. Parkin, K. P. Roche, M. G. Samant, P. M. Rice, R. B. Beyers, R. E. Scheuerlein, E. J. O'Sullivan, S. L. Brown, J. Bucchignano, D. W. Abraham, Y. Lu, M. Rooks, P. L. Trouilloud, R. A. Wanner, and W. J. Gallagher, "Exchange-Biased Magnetic Tunnel Junctions and Application to Nonvolatile Magnetic Random Access Memory," *J. Appl. Phys.* **85**, 5828–5833 (1999).
6. M. G. Samant, J. Luning, J. Stohr, and S. S. P. Parkin, "Thermal Stability of IrMn and MnFe Exchange-Biased Magnetic Tunnel Junctions," *Appl. Phys. Lett.* **76**, 3097–3099 (2000).
7. R. K. Singh and R. Bajaj, "Advances in Chemical–Mechanical Planarization," *Mater. Res. Soc. Bull.* **27**, 743–751 (2002); A. S. Brown, "Flat, Cheap, and Under Control (Electrochemical Mechanical Planarization)," *IEEE Spectrum* **42**, 40–45 (2005); J. Pallinti, P. Burke, W. Barth, W. Catabay, S. Lakshminaryanan, and D. Wang, "Removing Copper Over Low-K Films Using Stress-Free Polishing," *Solid State Technol.* **47**, 69–72 (2004); G. T. Stauf, K. Boggs, P. Wrschka, C. Ragaglia, M. Darsillo, J. F. Roeder, M. King, Jun Liu, and T. Baum, "Copper CMP Formulation for 65nm Device Planarization," *Mater. Res. Soc. Symp. Proc.* **816**, 23–28 (2004); C. L. Borst, S. M. Smith, and M. Eissa, "Challenges and Rewards of Low-Abrasive Copper CMP: Evaluation and Integration for Single-Damascene Cu/Low-*K* Interconnects for the 90nm Node," *Mater. Res. Soc. Symp. Proc.* **816**, 3–14 (2004).
8. J. Crowell, L. Tedder, H. Cho, F. Cascarano, and M. Logan, "Model Studies of Dielectric Thin Film Growth: Chemical Vapor Deposition of SiO_2," *J. Vac. Sci. Technol. A* **8**, 1864–1870 (1990).
9. L. Savtchenko, B. N. Engel, N. D. Rizzo, M. F. DeHerrera, and J. A. Janesky, "Method of Writing to Scalable Magnetoresistance Random Access Memory Element," U.S. Patent 6,545,906, 2003; D. Worledge, "Spin Flop Switching for Magnetic Random Access Memory," *Appl. Phys. Lett.* **84**, 4559–4561 (2004).
10. M. K. Carter and R. Small, "Electrochemical Measurements of Passivation Bilayers on Copper in a CMP System," *J. Electrochem. Soc.* **151**, B563–B571 (2004).

Received May 2, 2005; accepted for publication July 19, 2005; Internet publication January 25, 2006

Michael C. Gaidis *IBM Systems and Technology Group, Technology Development & Manufacturing, Route 52, Hopewell Junction, New York 12533 (gaidis@us.ibm.com)*. Dr. Gaidis is a process integrator for the development of MRAM in the IBM Materials Research Laboratory. In the course of obtaining his B.S. degree in physics and his M.S. degree in electrical engineering from the Massachusetts Institute of Technology (1989), he developed rapid thermal annealing techniques suitable for III–V semiconductor processing. He received a Ph.D. degree in applied physics from Yale University in 1994 for his work with superconducting X-ray detectors. After developing superconducting mixers for far-IR astronomy during a postdoctoral position at the California Institute of Technology, he joined NASA in the Observational Systems Division of the Jet Propulsion Laboratory. There, Dr. Gaidis engineered a terahertz receiver, now orbiting the earth on NASA's AURA spacecraft, to detect traces of atmospheric ozone depletion reactions. He joined IBM in 2001 to help bring MRAM to a state of maturity suitable for production in an advanced semiconductor fabrication facility.

Eugene J. O'Sullivan *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (eosull@us.ibm.com)*. Dr. O'Sullivan is a Research Staff Member in the Advanced Manufacturing Technology Department. He received his Ph.D. degree in electrochemistry from University College, Cork, Ireland, in 1981. From 1980 to 1983 he was a Postdoctoral Fellow at Case Western Reserve University, and he was a Senior Research Associate at Case from 1983 to 1984, when he joined the IBM Research Division. At IBM, Dr. O'Sullivan's research has included fundamental and applied aspects of electroless metal deposition, e.g., diffusion barriers, selectivity of deposition, polyimide adhesion, and new electroless alloy deposition solutions. He has worked on the fabrication of a multilayer magnetic minimotor. Dr. O'Sullivan is currently working on aspects of MRAM device fabrication such as post-dry-etch wet cleaning, chemical etching, and chemical–mechanical planarization (CMP).

Janusz J. Nowak *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (nowakj@us.ibm.com)*. Dr. Nowak is a Research Staff Member in the Magnetoelectronics Department of the Thomas J. Watson Research Center. In 1981 Dr. Nowak received his Ph.D. degree in solid-state physics from the Institute of Physics, Polish Academy of Sciences, Warsaw. From 1978 to 1994 he worked on magnetic thin films, Lorentz electron microscopy, and magnetic tunnel junctions in the Department of Solid State Physics of the Polish Academy of Sciences in Zabrze, Poland. In 1995 he joined Jagadeesh Moodera at the Francis Bitter Magnet Laboratory at MIT to continue research on magnetic tunnel junctions. In 1998 he joined Seagate Technology, Minneapolis, Minnesota, and as a principal engineer worked on developing ultra-thin tunnel barriers for magnetic hard-drive head applications. In 2004 Dr. Nowak joined the IBM Research Division in Yorktown Heights and began work on MRAM.

Yu Lu *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (yulu1@us.ibm.com)*. Dr. Lu received his Ph.D. degree in physics in 1997 from Brown University, subsequently joining the IBM MRAM project at the Thomas J. Watson Research Center. His work has included device characterization, magnetic and physical cell design, design manual development (including design rule ownership), and parametric macro and test site design. Dr. Lu is an author or coauthor of more than 20 papers and patents, most of them related to MRAM.

Sivananda Kanakasabapathy *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (sivakana@us.ibm.com)*. Dr. Kanakasabapathy is an engineer in the Physical Sciences Department at the Thomas J. Watson Research Center. He received a B.Tech. degree in electrical engineering from the Indian Institute of Technology, Madras, in 1995, and M.S. and Ph.D. degrees in electrical engineering from the University of Texas at Dallas in 1997 and 2001, respectively. He subsequently joined IBM at the Thomas J. Watson Research Center, where he has worked on patterning of magnetic tunnel junctions for MRAM applications. He is an author or coauthor of three patents and seven technical papers. Dr. Kanakasabapathy is a member of the American Vacuum Society.

Philip L. Trouilloud *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (plt@us.ibm.com)*. Dr. Trouilloud is a Research Staff Member in the Physical Sciences Department at the IBM Thomas J. Watson Research Center. He is a graduate of the Ècole Polytechnique, France. He completed his Doctorate at the Laboratoire de Physique des Solides in Orsay, France. After postdoctoral studies at the University of Texas at Austin, he joined the IBM Research Division in Yorktown Heights, New York. Dr. Trouilloud's work focuses on MRAM development, specifically on magnetic devices and magnetic test and characterization.

Daniel C. Worledge *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (worledge@us.ibm.com)*. Dr. Worledge received a B.A. degree with a double major in physics and applied mathematics from the University of California at Berkeley in 1995. He received a Ph.D. degree in applied physics from Stanford University in 2000, with a thesis on spin-polarized tunneling in oxide ferromagnets. After joining the IBM Thomas J. Watson Research Center as a postdoctoral researcher, he became a Research Staff Member in 2001, working on fast-turnaround measurement methods for magnetic tunnel junctions. In 2003, Dr. Worledge became the manager of the MRAM Materials and Devices group. His current research interests include magnetic devices and their behavior at small dimensions.

Solomon Assefa *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (sassefa@us.ibm.com)*. Dr. Assefa joined the IBM Research Division in 2004. He received a B.S. degree in physics and B.S. and M.S. degrees in electrical engineering and computer science (EECS) from MIT. In June 2004, he received a Ph.D. degree from MIT for his research on novel passive and active photonic crystal devices on III–V and Si-based platforms. Dr. Assefa is currently integrating novel tunnel-junction barrier materials into MRAM technology for high MR demonstration, investigating MRAM scaling for 45- and 32-nm nodes, and exploring spin-torque-transfer memory devices.

Keith R. Milkove *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (milkove@us.ibm.com)*. Dr. Milkove is a Staff Engineer in the Physical Sciences Department at the Materials Research Laboratory (at the IBM East Fishkill facility and the IBM Thomas J. Watson Research Center). He received S.B. and M.S. degrees in physics from the Massachusetts Institute of Technology in 1975 and 1978, respectively, and M.S. and Ph.D. degrees in materials science and engineering from Cornell University in 1979 and 1984,

respectively. He subsequently joined IBM, where his research interests have included dry-etch technology, physical failure analysis of magnetic materials, thin-film growth, superconducting tunneling phenomena, X-ray diffraction, and semiconductor materials processing. Dr. Milkove is an author or coauthor of 22 publications and three patents. He is a member of the American Physical Society, the American Vacuum Society, and Sigma Xi.

Gwen P. Wright *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (gwenwr@us.ibm.com)*. Ms. Wright is a process engineer in the MRAM group at the Thomas J. Watson Research Center working on RIE and IBE processes. She received a B.S.E.E. degree from Fairleigh Dickinson University and attended graduate school at the Polytechnic Institute of New York. Her previous work at IBM included novel photolithography techniques for silicon technology, fabrication and testing of early GaAs MESFET devices, electron-beam lithography, and scanning electron microscopy.

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.