

Resistive Memory Based Acceleration of Data Intensive Computing

Engin Ipek and Eby G. Friedman

Department of Electrical and Computer Engineering
University of Rochester
Rochester, New York 14627

1 Introduction

Resistive memory technologies hold the promise of replacing mainstream on-chip memory while providing enhanced throughput and capacity in modern compute systems. Demonstrating material, process, and circuit compatibility with existing CMOS infrastructures, resistive memories deliver non-volatility, no static power consumption, and improved density. Application of these technologies, however, requires novel circuits and architectures that exploit these features.

Several approaches are summarized in this article in which resistive memory technologies are leveraged to achieve significant performance enhancement in modern data intensive applications. Two recent results utilizing phase change memory (PCM) and spin torque transfer magnetoresistive RAM (STT-MRAM) based resistive TCAM systems demonstrate significant acceleration and energy reduction over a broad set of data intensive applications. Additionally, recent circuit level enhancements to increase the sensing ratio are described that demonstrate improved read latency for STT-MRAM arrays. Lastly, a magnetic field is applied to MTJ devices to improve both write energy and latency in high performance STT-MRAM on-chip caches.

2 Resistive TCAM Accelerator

A new technique [1] is described that aims at cost effective, modular integration of a high capacity ternary content addressable memory (TCAM) system within a general purpose computing platform. The TCAM density is improved by more than $20\times$ over existing, CMOS-based circuits through a novel resistive TCAM cell and array architecture. High capacity resistive TCAM circuits are placed on a DDR3 compatible DIMM, and accessed through a software library with no modifications to the processor or motherboard. The modularity of the resulting memory system allows the TCAM to be selectively included in systems running workloads that are amenable to TCAM-based acceleration. Moreover, when executing an application or a program that does not benefit from an associative search capability, the TCAM DIMM can be configured to provide standard RAM functionality. By tightly integrating TCAM with conventional virtual memory, and by allowing a large fraction of the physical address space to be made content addressable on demand, the proposed memory system improves the average performance by $4\times$ and average energy consumption by $10\times$, as demonstrated on a set of data intensive applications [1].

2.1 Overview

An example computer system with the proposed resistive TCAM system is depicted in Figure 1. A multicore processor is connected to main memory through an on-chip memory controller. The TCAM DIMM sits side-by-side with DRAM on the DDR3 bus. An on-DIMM TCAM controller serves as the interface to DDR3, and manages the DIMM. The processor communicates with the controller through a set of memory mapped control registers (for configuring the func-




Figure 1: Illustrative example of a computer system with the proposed resistive TCAM DIMM.

tionality) and a memory mapped key store that resides with the controller (for buffering the search key). Each TCAM is composed of eight banks. A bank is comprised of a set of arrays that are searched against the query key, as well as a hierarchical reduction network for counting the number of matches and choosing the highest priority matching row.

2.2 Structure and Operation

A TCAM is a special type of associative memory that supports both storage and search with a wildcard (X) in addition to a logic zero or one. A wildcard matches against both binary states (as well as another wildcard), and can be used in both the search key or the stored data word. Some circuit-level features of this TCAM system are described below. R_{HI} and R_{LO} refer, respectively, to the high and low resistance states of a resistive storage element.

TCAM Cell Structure. The proposed area efficient resistive TCAM cell, which consists of three pairs of resistive storage elements and access transistors, is shown in Figure 2. The first two resistors store the data bit and complement; the third resistor is permanently programmed to R_{HI} . To store a logic 1 or 0, the leftmost resistor is programmed to store the data bit (D), while the resistor in the middle is programmed to store the complement of the bit (\bar{D}).




Figure 2: Resistive TCAM cell.

To search for a logic 0 or 1, SL and $\bar{S}L$ are driven, respectively, with the search bit and complement, turning one of the access transistors on and the other off. A match is decided based on the effective resistance between the matchline and ground. If the resistor in the high resistance state is in series with the on transistor—adding a resistance of R_{HI} between the matchline and ground—the search results in a match; conversely, a resistance of R_{LO} connected to the matchline indicates a mismatch. To search for a wildcard (X), SL and $\bar{S}L$ are disabled and SX is driven high; hence, a resistor in the R_{HI} state is connected to the matchline regardless of the value stored in the cell. The proposed TCAM cell consists of three 1T-1R cells sharing a matchline and contacts to ground. At 22 nm, this cell exhibits an area

of $27F^2$ (three times the 1T-1R PCM cell size projected by ITRS [2]), which is $\frac{1}{20}$ of the area of a CMOS TCAM cell.

3 AC-DIMM

The AC-DIMM [3] is a flexible, high performance associative compute engine built on a DDR3 compatible memory module, which is based on TCAM associative computing. The AC-DIMM addresses the limited flexibility of previous resistive TCAM accelerators by combining two powerful capabilities—associative search and processing-in-memory. Generality is improved by augmenting a TCAM system with a set of integrated, user programmable microcontrollers that operate directly on the search results. A bit-serial TCAM array is proposed which enables the system to exploit STT-MRAM. The AC-DIMM achieves a $4.2\times$ speedup and a $6.5\times$ reduction in energy as compared to a conventional RAM-based system on a set of 13 benchmark applications [3].

3.1 Overview

The organization of an AC-DIMM enabled computer system is illustrated in Figure 3. A multicore processor accesses main memory via an integrated memory controller, which buffers memory requests, schedules memory accesses, and issues DRAM commands over a DDR3 bus. The system supports one or more AC-DIMMs on the memory bus, each comprising an on-DIMM controller and eight associative computing integrated circuits (AC-IC). The DIMM controller consists of control logic, interface logic, and RAM-based storage (the shaded blocks in Figure 3). To reduce the peak power, the AC-DIMM adopts a bit-serial search scheme; only one of the AC-ICs can be searched at a time to ensure that the instantaneous power does not exceed the maximum power rating of a standard DDR3 DIMM (15 watts [4]).

An AC-IC is built from STT-MRAM arrays. A set of specialized microcontrollers, each co-located with a group of four arrays, perform ALU operations on the search results. A reduction tree forwards processed results to the DIMM result store. By mapping part of the physical address space onto the AC-DIMM, data is made content addressable and is processed directly by the memory circuit, which significantly reduces data movement, increasing energy efficiency.

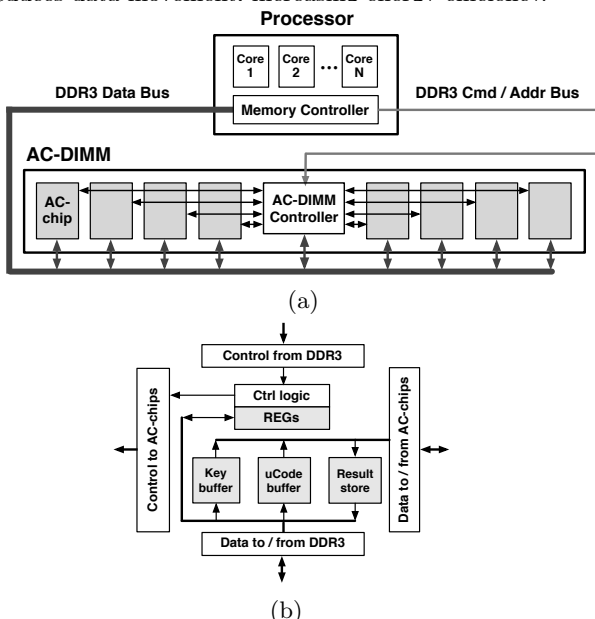


Figure 3: An example computer system with (a) an AC-DIMM, and (b) an on-DIMM controller.

3.2 Summary of Key Results

The proposed resistive TCAM cell and array architecture deliver a $20\times$ density improvement over existing CMOS-based solutions. A modular memory system that places resistive TCAM circuits on a DDR3-compatible DIMM, and accesses the DIMM through a software library with no modifications to the processor or the motherboard has been explored [1]. The evaluation compares a baseline multicore system running eight threads to a single-threaded system with the proposed TCAM accelerator, exhibiting average performance and energy improvements of, respectively, $4\times$ and $10\times$. The efficiency gains are due to two factors. First, a TCAM eliminates off-chip data movement and instruction processing overheads by processing data directly on the IC; second, the faster execution time leads to lower leakage energy.

The AC-DIMM uses a novel 2T-1R STT-MRAM cell, which is $4.4\times$ denser than a CMOS TCAM cell. When implemented in an embedded STT-MRAM process, the cell topology is applicable to any memory technology behaving as a RAM. The AC-DIMM broadens the scope of associative memory systems over existing approaches by allowing a key-value pair to be co-located within the same row, and by employing integrated microcontrollers to execute user defined operations on the search results. A high performance, energy efficient solution is produced that successfully combines associative search and processing-in-memory capabilities.

4 STT-MTJ Memory Cells

STT-MRAM has unique advantages over traditional memory circuits. The Achilles heel of STT-MRAM, however, is the small on/off resistance ratio. This limitation requires sophisticated read circuitry which leads to greater sensitivity to noise. To address these limitations, two memory cells have recently been proposed that significantly improve the output read ratio [5, 6]. These memory cell variants utilize additional CMOS transistors within the cell to enhance the observed on/off resistance ratio of the MTJ device, leading to a shorter read delay. Each cell exhibits an order of magnitude increase in the current ratio as compared to a traditional 1T-1R structure while requiring additional area and delivering comparable energy efficiency under high bias.

4.1 Overview

Three basic cell types are considered for use in STT-MTJ memories: a standard 1T-1R memory cell as well as the proposed 2T-1R cell variants.




Figure 4: Circuit diagram of STT-MTJ memory cells: a) standard 1T-1MTJ, b) 2T-1MTJ diode cell, and c) 2T-1MTJ gate cell.

4.1.1 1T-1R cell

The 1T-1R cell, the standard basic building block of resistive memory arrays (see Fig. 4a), must satisfy several design constraints to operate correctly. At full bias, the internal cell transistor and access circuitry must supply sufficiently high current to ensure that the MTJ switches. For reads, the cell

current must remain sufficiently below the critical current to mitigate the potential for erroneous writes to the device. Moreover, each transistor isolates a selected memory cell from peripheral cells to maintain the required sense margin. For this purpose, the read operation biases the access transistor to operate within the linear region. The sense margin of the device is observed as a voltage or current proportional to the on/off resistance ratio of the device.

4.1.2 2T-1R cells

Alternate 2T-1R memory cell topologies utilizing an additional transistor can produce voltage and current amplification without sacrificing immunity to leakage current within an STT-MRAM array.

Diode connected transistor read port. A diode connected transistor incorporated into a memory cell, as shown in Fig. 4b, amplifies the voltage of the internal node of the memory cell (node B) to produce a current and voltage signal at the transistor output. The maximum amplification occurs when node B is biased to ensure that the R_{on} and R_{off} states produce a voltage, respectively, above and below the threshold of the transistor.

Gate connected transistor read port. A gate connected memory cell, as shown in Fig. 4c, achieves the same amplification as the diode connected transistor and operates at a similar maximum voltage. This topology, however, differs in several key aspects. First, the gate connected transistor is electrically isolated from node B, facilitating the addition of multiple gate connected read ports. Secondly, the source of the transistor is connected to ground, eliminating any source body voltage bias, improving the conductance of the transistor. Thirdly, the output current margin is a function of transistor width which can be increased to improve the sense margin.

4.2 Summary of Key Results

A comparison of an 8T SRAM with STT-MRAM memory cells in terms of the read delay, read energy, and physical area are listed, respectively, in Tables 1, 2, and 3. The SRAM read ports (RP) and write-read ports (WRP) are evaluated for both memory specific high density (HD) and logic process (Logic) design rules [7]. Each memory cell and the associated parasitic impedances are scaled to the 22 nm technology node. The 8T SRAM read port is sensed using a standard single-ended inverter sense amplifier [8]. The SRAM write-read port is sensed using a standard dynamic latch sense amplifier [9]. Each of the STT-MRAM cells is sensed using a clamped bitline sense amplifier [10]. The array sizes are typical of an on-chip cache array.

Table 1: Single bit access delay (ns)

Number of Bits	SRAM 8T HD RP	SRAM 8T HD WRP	SRAM 8T Logic RP	SRAM 8T Logic WRP	1T-1R	2T-1R Gate	2T-1R Diode
2,048	14.879	14.708	25.793	26.927	3.106	4.200	3.762
1,024	4.471	3.716	7.189	6.754	0.718	1.242	0.969
512	1.537	0.960	2.219	1.721	0.265	0.377	0.295
256	0.626	0.273	0.800	0.466	0.127	0.139	0.111
128	0.306	0.094	0.352	0.145	0.078	0.067	0.057

Table 2: Single bit access energy (fJ)

Number of Bits	SRAM 8T HD RP	SRAM 8T HD WRP	SRAM 8T Logic RP	SRAM 8T Logic WRP	1T-1R	2T-1R Gate	2T-1R Diode
2,048	31.182	28.529	31.197	28.561	5.382	50.285	98.113
1,024	19.144	18.243	19.175	18.274	1.081	26.014	39.430
512	12.093	12.014	12.124	12.045	0.568	12.559	17.441
256	6.047	6.235	6.078	6.266	0.370	6.250	7.891
128	2.736	2.973	2.767	3.004	0.284	3.170	3.620

Delay metrics for square array sizes ranging from 128 to 2,048 bits are listed in Table 1. The STT-MRAM arrays

Table 3: Area comparison

	SRAM 8T HD	SRAM 8T Logic	1T-1R	2T-1R Diode	2T-1R Gate
Cell Height (F)	8	8	7	7	7
Cell Width (F)	31.6	45.4	6.65	10.8	14.5
Density (F ²)	252	363.2	46.55	75.6	101.5

exhibit significantly less delay than the SRAM counterparts. At an array size of 2,048 cells, the delay of SRAM and STT-MRAM is dominated by the wordline delay. As compared to the single ended SRAM read port, the delay of each STT-MRAM memory cell type is smaller by a factor of 3.9, 4.6, and 5.37, respectively, for the 1T-1R, 2T-1R gate connected, and 2T-1R diode connected memory cells. Both the gate and diode connected cells exhibit an area overhead larger than the 1T-1R cell but overcome this issue through an improved current ratio which reduces the delay.

The energy consumption of each cell type is listed in Table 2. Each cell type exhibits a significant reduction in energy consumption with a smaller data array. The gate connected and diode connected cells plateau at an energy similar to SRAM arrays at smaller sizes. This behavior is due to the additional bias required to drive the internal node of the cell. The 1T-1R cell does not require an additional bias, enabling more energy efficient reads than the other memory cell types. At larger array sizes, the 2T-1R cell variants require more energy than the other cell types.

Among each of the memory types, the SRAM requires longer delays and greater energy than the STT-MRAM memory. In general, the 1T-1R outperforms SRAM for all array sizes. Both of the 2T-1R cells require more energy at large array sizes, indicating that each topology is better suited to small active on-chip caches where speed is paramount. At these sizes, the 2T-1R topology exhibits the fastest read operation of any memory cell type at an energy consumption comparable to SRAM.

5 Field Enhanced STT Switching

A key issue constraining the use of STT-MRAM is the switching latency of the magnetic tunnel junctions (MTJ) within each memory cell. The long latency causes the switching energy of an MTJ to be much greater than traditional CMOS SRAM.

To address this issue, the first generation MRAM cell topology is utilized with an STT-MTJ where an additional field current is applied to destabilize the MTJ prior to switching, reducing the switching latency. An analytic framework for assessing and optimizing field driven writes in STT-MRAM arrays is described in recent work [11]. Building on this framework, the switching latency and energy can be reduced by amortizing the additional field current over many cells, leading to a significant reduction in energy consumed per bit.

5.1 STT-MRAM Device Switching and Array Structure

Since the spin transfer torque effect was first incorporated into MTJ switching, MRAMs have exclusively used this effect for writing. The STT effect, however, can complement field driven excitation of the magnetic free layer within an MTJ. Classical MRAM approaches use two perpendicular currents, with a single selected MTJ at the intersection, to create a magnetic field that acts on the free layer of an MTJ (see Figure 5a). This approach suffers from several problems: (1) the use of two currents to switch a single bit consumes a large amount of energy as compared to DRAM, (2) MTJs in adjacent cells on the path to the target MTJ are

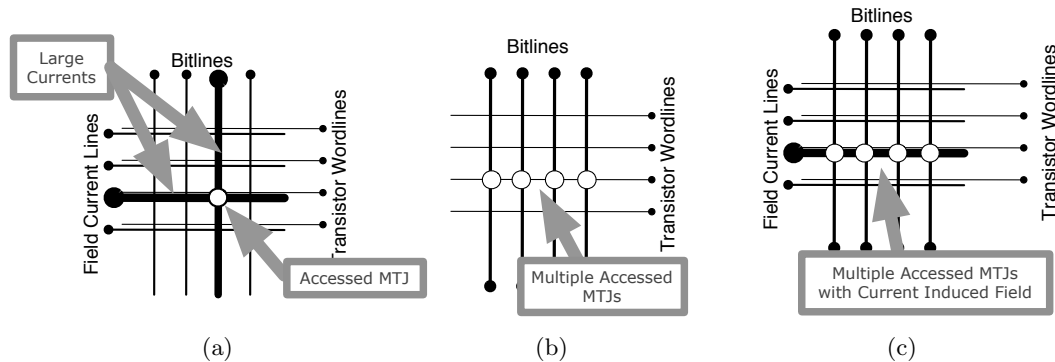


Figure 5: Current biasing scheme for a) classical MRAM, b) standard STT-MRAM, and c) proposed STT-MRAM arrays

half-selected by the high fields of the write currents, potentially inducing erroneous writes, and (3) a checker read operation is required to ensure that the correct state is written into the device. These issues limit the scalability of classical MRAM devices.

The STT effect overcomes these problems by using a single current that passes through the MTJ. This technique enables many MTJs to be written in parallel, as illustrated in Figure 5b. The overall switching current is much lower than in classical MRAM, which removes the half select problem. The write latency, however, remains significantly longer than the read latency, and the switching energy is also significantly higher than DRAM. Supplying a sufficiently large write current requires a large access transistor, which reduces density.

The approach proposed here combines an STT-based current with the field generating current used in classical MRAM circuits, as shown in Figure 5c. In this approach, the field current creates an additional magnetic field that destabilizes the MTJs across the row. Each MTJ is biased with an STT current that controls the switching of the MTJs in each column. The use of a field current in this manner has two beneficial effects: (1) the alignment of the field with respect to the MTJ can destabilize the device, which reduces both the write latency and energy, and (2) the cells in a row share the field current, ensuring that there is no half-select problem [12] and the energy consumption of the field current is amortized across the row. The energy consumption per bit is therefore much less than a standard STT-MRAM.

5.2 Summary of Key Results

The field driven approach utilized in classical MRAM cells reduces the switching latency of an STT-MTJ. An array model is presented that describes the switching energy and energy consumption for different field currents and array sizes. The per bit switching latency can be reduced by a factor of ten. As compared to nominal STT-MRAM, an 82% reduction in switching energy per bit is achieved. The reduction in both switching energy and latency provides a significant enhancement in performance for embedded high performance STT-MRAM based memories and enables the use of STT-MRAM in write latency critical applications.

6 Summary

Several recent innovations are highlighted that leverage resistive memory to improve power and performance in data intensive processing applications. TCAM-DIMM, a TCAM based accelerator, exhibits a 4x reduction in computation time at less than one tenth of the baseline system energy consumption. By combining TCAM functionality with pro-

cessing in memory, AC-DIMM improves performance while increasing the variety of applications that can be accelerated using a TCAM based approach. Circuit level approaches are presented to reduce the latency and energy consumption of STT-MRAM based resistive memories. Two 2T-1R cell topologies are described that increase the on-off ratio of resistive memories and thereby reduce delay. A field driven approach to STT-MRAM switching further reduces the critical write latency of the device, providing a 10x reduction in delay, and an 82% reduction in power consumption. The confluence of circuit and architectural innovations applied to emerging resistive memory technologies provides a significant improvement in the performance of data intensive computational systems.

7 References

- [1] Q. Guo, X. Guo, Y. Bai, and E. İpek, "A Resistive TCAM Accelerator for Data-Intensive Computing," *Proceedings of the IEEE/ACM International Symposium on Microarchitecture*, pp. 339–350, June 2011.
- [2] ITRS, *International Technology Roadmap for Semiconductors: 2010 Update*, <http://www.itrs.net/links/2010itrs/home2010.htm>.
- [3] Q. Guo, X. Guo, R. Patel, E. İpek, and E. G. Friedman, "AC-DIMM: Associative Computing with STT-MRAM," *Proceedings of the IEEE/ACM International Symposium on Computer Architecture*, pp. 189–200, June 2013.
- [4] Micron Technology, Inc., <http://www.micron.com//get-document/?documentId=425>, *1Gb DDR3 SDRAM*, 2006.
- [5] R. Patel, E. İpek, and E. G. Friedman, "STT-MRAM Memory Cells with Enhanced On/Off Ratio," *Proceedings of the IEEE International System-on-Chip Conference*, pp. 148–152, September 2012.
- [6] R. Patel, E. İpek, and E. G. Friedman, "STT-MRAM Memory Cells with Enhanced On/Off Ratio," *Microelectronics Journal* (in press).
- [7] K. Nii, Y. Tsukamoto, T. Yoshizawa, S. Imaolka, and H. Makino, "A 90 nm Dual-Port SRAM with 2.04 μm^2 8T-Thin Cell Using Dynamically-Controlled Column Bias Scheme," *Proceedings of the IEEE Solid-State Circuits Conference*, Vol. 1, pp. 508–543, February 2004.
- [8] S. Cosemans, W. Dehaene, and F. Catthoor, "A Low-Power Embedded SRAM for Wireless Applications," *IEEE Journal of Solid-State Circuits*, Vol. 42, No. 7, pp. 1607–1617, July 2007.
- [9] A. Hajimiri and R. Heald, "Design Issues in Cross-Coupled Inverter Sense Amplifiers," *Proceedings of the IEEE International Symposium on Circuits and Systems*, Vol. 2, pp. 149–152, May 1998.
- [10] T.N Blalock, and R.C. Jaeger, "A High-Speed Clamped Bit-Line Current-Mode Sense Amplifier," *IEEE Journal of Solid-State Circuits*, Vol. 26, No. 4, pp. 542–548, April 1991.
- [11] R. Patel, E. İpek, and E. G. Friedman, "Field Driven STT-MRAM Cell for Reduced Switching Latency and Energy," *Proceedings of the IEEE International Symposium on Circuits and Systems* (in submission).
- [12] S. Tehrani *et al.*, "Magnetoresistive Random Access Memory using Magnetic Tunnel Junctions," *Proceedings of the IEEE*, Vol. 91, No. 5, pp. 703–714, May 2003.