



HAL
open science

High Performance SoC Design Using Magnetic Logic and Memory

Yahya Lakys, Jacques-Olivier Klein, Daniel Etiemble, Dafiné Ravelosona, Claude Chappert, Weisheng Zhao, Lionel Torres, Luís Vitório Cargnini, Raphael Martins Brum, Yue Zhang, et al.

► **To cite this version:**

Yahya Lakys, Jacques-Olivier Klein, Daniel Etiemble, Dafiné Ravelosona, Claude Chappert, et al.. High Performance SoC Design Using Magnetic Logic and Memory. VLSI-SoC: Very Large Scale Integration - System-on-Chip, Oct 2011, Hong Kong, China. pp.10-33, 10.1007/978-3-642-32770-4_2. hal-01519767

HAL Id: hal-01519767

<https://inria.hal.science/hal-01519767>

Submitted on 9 May 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

High Performance SoC Design using Magnetic Logic and Memory

Weisheng Zhao¹, Lionel Torres³, Luís Vitório Cargnini³, Raphael Martins Brum³, Yue Zhang¹, Yoann Guillemenet³, Gilles Sassatelli³, Yahya Lakys¹, Jacques-Olivier Klein¹, Daniel Etiemble², Dafiné Ravelosona¹, and Claude Chappert¹

¹ IEF - Université Paris-Sud 11 / CNRS

² LRI - Université Paris-Sud 11 / CNRS

weisheng.zhao@u-psud.fr

³ LIRMM - Université Montpellier 2 / CNRS

torres@lirmm.fr

Abstract As the technology node shrinks down to 90nm and below, high standby power becomes one of the major critical issues for CMOS highspeed computing circuits (e.g. logic and cache memory) due to the high leakage currents. A number of non-volatile storage technologies, such as FRAM, MRAM, PCRAM and RRAM, are under investigation to bring the non-volatility into the logic circuits and then eliminate completely the standby power issue. Thanks to its infinite endurance, high switching/sensing speed and easy integration on top of CMOS process, MRAM is considered as the most promising one. Numerous logic circuits based on MRAM technology have been proposed and prototyped in the last years. In this paper, we present an overview and current status of these logic circuits and discuss their potential applications in the future from both physical and architectural points of view.

Keywords: MRAM, Non-volatile CPU, Magnetic Logic, Reconfigurable logic.

1 Introduction

Thanks to its fast speed, small size and low power, CMOS is the dominant technology to build high-speed computing circuits (e.g. logic and cache memory) [1, 2]. However, the high standby power due to the increasing leakage currents becomes a more and more critical issue as the fabrication node shrinks down to 90nm or below [3]. A number of non-volatile storage technologies such as Magnetic RAM (MRAM) [4], Ferroelectric RAM (FRAM) [5], Phase-Change RAM (PCRAM) [6] and Resistive RAM (RRAM) [7] are under investigation by both the industries and academics.

They are expected to bring the non-volatility into the CMOS logic circuits and then allow them to be powered off completely. All the data are protected and can be retrieved instantaneously on active state. This approach could overcome definitely the standby power issue and allows the circuits to be further shrunk down.

MRAM is certainly one of the most important applications of Spintronics, which is a very rapidly emerging R&D area (Nobel Prize 2007) and would have a significant

impact on the future of all aspects of electronics beyond CMOS as it did for Hard Disk Drives (HDDs) [8,9].

In 2006, the first MRAM based on Field Induced Magnetic Switching (FIMS) was commercialized [10] and it addresses particularly some specific niche markets such as space, security and aeronautics thanks to its radiation hardness. However, the first MRAM is greatly limited by its FIMS approach, requiring high currents for programming the cell ($>10\text{mA}$) [5] and this leads to low memory density and high switching power. Today, most of R&D efforts in MRAM focus on new switching approaches which are expected to be scalable, energy efficient, reliable (>10 years) and fast. A number of solutions have been investigated, such as Thermally Assisted Switching (TAS-MRAM) [11], Spin Transfer Torque (STT-MRAM) [8,12,13] and Domain Wall Motion (DW) [14]. STT-MRAM is the most promising MRAM technology and numerous companies have promised to commercialize it very soon [15,16].

Table 1: Comparison of different non-volatile memory technologies

Technology	Write endurance (cycles)	Switch time (ns)	Read time (ns)
MRAM	Infinite	~ 10	> 0.5
Flash	10^6	> 100	> 10
FRAM	10^{12}	~ 10	~ 20
RRAM	10^6	> 50	> 50
PCRAM	10^{12}	~ 50	~ 60

Beyond “data storage”, MRAM is the most promising technology for logic applications thanks to its infinite endurance, high switch/sense speed and easy 3D integration with CMOS processes. Flash memory, FRAM, PCRAM and RRAM, characterized by limited endurance $< 10^{12}$ [17] (see Table 1). This makes MRAM a really good candidate for many applications, such as non-volatile configuration FPGAs, design of Non-Volatile Flip Flop (NVFF) and embedded MRAMs. For instance, benefiting from the easy integration of MRAM on top of CMOS process, the final die area could be reduced and the interacting speed between logic and memory circuits can be accelerated over CMOS counterparts (see Figure 1).

MRAM based logic circuits (magnetic logic) were initiated in 2000 [18] and considered as a potential computing paradigm featuring high performances in terms of power, speed and area. Numerous academic and industry research groups joined this field since 2006 [19–22] and some Magnetic logic circuits have been presented and successfully prototyped, such as Magnetic Look-Up-Table (MLUT), Magnetic Flip-Flop (MFF) and embedded MRAM (eMRAM) as different levels of cache memory [15,23,24]. In this chapter, we describe an overview of the magnetic logic circuits and discuss their potential applications from both the physics and architecture points of view. We propose herein to give an overview of several applications that could be realized with MRAM technologies.




Figure 1: (a) 2D structure of CMOS logic (b) 3D structure of MRAM logic, the distance between logic and memory can be greatly reduced and then accelerate the computing speed.

First we briefly introduce Spin-Transfer Torque (STT) switching approach of MRAM and present its fundamental physics to achieve high-speed computing. In the following, we review different designs of magnetic logic circuits. In the fourth section, we focus on the MRAM based FPGA (MFPGA) and at last we discuss the applications of magnetic logic circuits and the future integration of MRAM as cache memory with processors or microcontrollers.

2 Spin Transfer Torque (STT) RAMs

2.1 STT based Magnetic Tunnel Junction (STT-MTJ)

Magnetic Tunnel Junctions (MTJ) is the basic cell of MRAM [8]. It is a nanopillar composed of two ferromagnetic (FM) layers and one oxide thin barrier (see Figure 2(a)). As the magnetization direction of the two FM layers is either in parallel or anti-parallel, a MTJ shows two different resistance values R_P and R_{AP} . For practical applications, the magnetization direction of one FM layer is pinned as reference and that of the other ferromagnetic layer is free to be switched to store binary state [5, 9].

As mentioned, a number of new MRAM switching approaches were investigated in order to replace FIMS and then expand widely its applications. TAS still requires two currents for the switching operation [11]: one passes through and heats the MTJ cell, and the other generates magnetic field to change the magnetization direction of free layer. This approach promises a number of advantages in terms of power, data reliability and scalability etc., which allow it suitable to be embedded as configuration memory in FPGA [25, 26].

Spin Transfer Torque (STT) based devices [12, 13] are able to switch the free layer with only one low spin polarized current I_{switch} , as shown in Figure 2(b)). Thereby it promises naturally low power, small die area (22 nm technology node) and fewer masks for fabrication. Its switching speed is very fast ($< 10ns$) [27, 28] allowing fast logic circuits to be built such as NVFF. Furthermore, most of MRAM prototypes and commercial programs are based on this approach, which accelerate progressively its maturity.

We believe that STT-MRAM could be the mainstream technology to manufacture magnetic logic circuits [15, 29]. Nevertheless, data thermal stability is a critical issue for STT-MRAM due to its planar anisotropy storage principle, which leads to important random sensing errors.

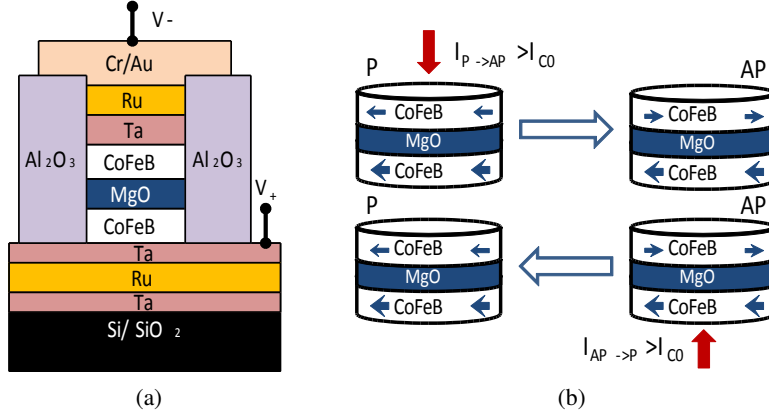


Figure 2: (a) Vertical structure of an MTJ nanopillar composed of CoFeB/MgO/CoFeB thin films. (b) Spin transfer torque switching mechanism: the MTJ state changes from parallel (P) to anti-parallel (AP) as the positive direction current $I_{P \rightarrow AP} > I_{C0}$, on the contrast, its state will return to P state with the negative direction current $I_{AP \rightarrow P} > I_{C0}$.

2.2 Tradeoff between power, speed and data retention

For standalone memory applications, the thermal activation energy factor of STT-MTJ $\xi = E/k_B T$ should be more than 42 according to Equation 3 [30] to allow the data retention Rt longer than 10 years (thin dot line in Figure 2). The acceptable error rate here is as high as 10^{-1} thanks to ECC circuits, which on the contrary cannot be embedded in cache memories and logic circuits to avoid performance degradation [31].

Figure 2 shows the relationship between Rt and ξ with lower error rate, 10^{-3} for L3 cache memory and 10^{-9} for NVFF. We can observe that ξ should be more than 48 and 61 for these two cases. However, for high-speed computing circuits, the Rt is not required to be longer than 10 years as the storage is often the intermediate transition between standalone memory and computing unit. If the needed Rt is 1 ms, ξ more than 35 is enough to ensure the logic operation. The equations governing the MTJ data retention can be written as follows,

$$\tau_1 = \tau_0 \exp(\xi) \quad (1)$$

$$Rt = -\tau_1 * \ln(R_{error}) \quad (2)$$

$$E = \frac{M_s * H_K * V}{2} \quad (3)$$

where τ_0 is characteristic switching latency between two states=1 ns, R_{error} is the acceptable error rate of data storage, k_B is Boltzmann constant, T is the temperature, E is the energy barrier that separates the two magnetization directions, which depends on the saturation magnetic field M_S , the shape anisotropy field H_K and the volume V of free layer of MTJ nanopillar.

From the Equation 3, we can find that the activation energy E depends greatly on the anisotropy field H_K , which is mainly dominated the shape of MTJ nanopillar. An elongated cell can provide relatively high H_K , and then improve the thermal stability of MTJ (see Figure 4). These relationships are described by the equations below,

$$H_K = M_S(N_b - N_a) \quad (4)$$

$$N_a = \frac{4\pi}{m^2 - 1} \left[\frac{m}{\sqrt{m^2 - 1}} \ln \left(m + \sqrt{m^2 - 1} \right) - 1 \right] \quad (5)$$

$$N_b = 2\pi - \frac{N_a}{2} \quad (6)$$

$$m = \frac{a}{b}, \quad (7)$$

where a and b are respectively the length and width of MTJ nanopillar, N_a and N_b are the demagnetization factor along the longer a -axis and the shorter b -axis of the elongated cell, m is the length/width ratio. When $a = b$, $N_a = N_b$, which means that there is no shape anisotropy for in-plane MTJ if it is a circular.

In STT-MRAM based non-volatile logic circuits, we can use different shapes of MTJ nanopillar to address different speed requirements based on the same magnetic process. For instance, we can use $\xi=29, 43$ and 50 (see Figure 4) for respectively NVFF, cache memory and main memory. One potential constraint of this approach is the sensitivity to process variation for the advanced technology node beyond 45 nm. As the shape of MTJ is mainly ellipse, its manufacture needs more precisely lithography techniques than the rectangular of CMOS fabrication.

2.3 Future Trends

With the miniaturization of the fabrication node (e.g. for lateral sizes of $40nm$), data thermal stability becomes a critical issue for STT-MRAM due to its planar anisotropy storage principle, which leads to important random sensing errors. For standalone memory applications, this issue can be overcome through Error Correction Codes (ECC) circuits [31]; however there isn't any efficient solution for logic circuits except for reducing the sensing current value and enlarging the sensing circuit area.

One solution is to use thermally assisted switching for spin transfer torque (STT + TAS) by using exchange bias effect in the MTJ nanopillar [11, 32, 33], which may overcome completely the thermal stability issue of in-plane STT-MTJ. However this approach needs heating and cooling mechanisms for each switching operation, which slower greatly the speed and limit its speed to be around 10 ns. Thereby this solution cannot be suitable for high-speed computing.

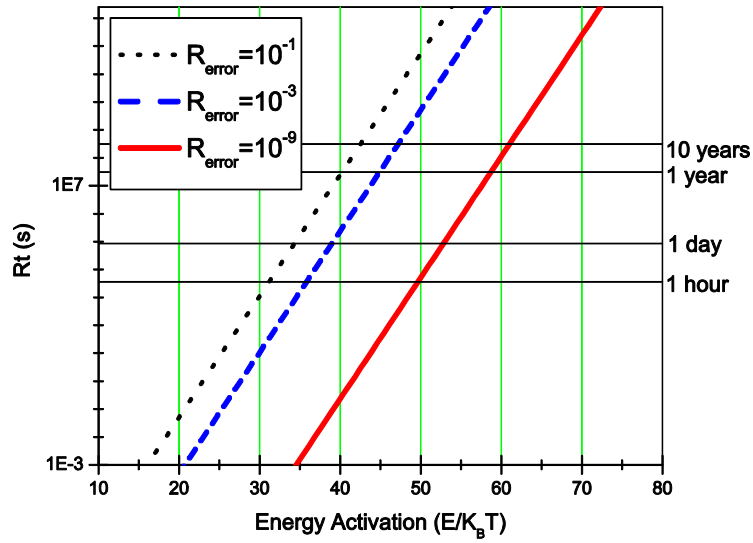


Figure 3: Relationship between data retention duration R_t and ξ . R_t can be reduced to obtain low ξ (e.g. 35) while keeping low error rate.




Figure 4: Influence of the shape anisotropy to the activation energy of in-plane MTJ.

Perpendicular anisotropy (PMA) based MTJ addressing this issue has been rapidly advanced [28], which can significantly decrease the sensing error caused by random magnetization switching while keeping high R_{AP}/R_P ratio for sensing performance [29]. PMA MTJ is becoming the mainstream solution to overcome the thermal stability issue and attracts the most attention from both the academics and industries. The switching duration of PMA STT-MTJ is as well as reduced from 10 ns to about 1 ns, which makes it the best candidate to build non-volatile high-speed computing circuits [34]. Furthermore, PMA can be obtained in a circular shape, which is less sensitive to the process variation unlike that of in-plane MTJ. Recently, it has been found that one can

use the same stack configuration of in-plane MTJ (e.g. CoFeB/MgO/CoFeB, see Figure 2(a)) to obtain PMA by controlling precisely the thickness of free layer, which should be from 1.0 nm to 1.3 nm [28].

Lower ξ can linearly reduce the threshold switching current I_C for STT-MTJ (see Equation 4) [28], which can lead to the switching duration $D_t < 1ns$ while keeping the same I_{switch} value (e.g. $300\mu A$) according to the Equation 5 [27, 30] (see Figure 4). Thus this reduction of ξ implies the reduction of both power dissipation and switch duration. Higher density can be achieved with low I_{switch} keeping the same speed, as the footprint of MRAM depends mainly on the current source for switching current, as shown below,

$$I_c = \frac{2\alpha\lambda e}{\mu_B g} * E \quad (8)$$

$$D_t = \frac{2}{C + \ln\left(\frac{\pi^2\xi}{4}\right)} \frac{\mu_B P_{ref}}{em(1 + P_{ref}P_{free})} (I_{switch} - I_C) \quad (9)$$

where α is the magnetic damping constant, e is the magnitude of the electron charge, γ is the gyromagnetic ratio, g is a function of the spin current polarization and the angle between the magnetization of the free and the reference layers, C is Euler's constant =0.577 for CoFeB/MgO/CoFeB nanopillar, μ_B is the Bohr magneton, P_{ref} and P_{free} are the tunneling spin polarizations of the reference and free layers, and m is the free layer magnetic moment.

This study demonstrates that high computing speed $> GHz$ and low power of embedded MRAM can be achieved at the same time from the physics point of view. For different applications (e.g. three levels of cache memory), its characteristics in terms of power, data retention and area are customizable through specific algorithms, circuits or architectural tricks.

For instance, in [35], the authors explored these ideas by creating a cache hierarchy in which each level is constructed with a different data retention time. This was done by modifying the MTJ shape as described before. A data refresh scheme similar to the ones used normally in DRAM devices was added to this scheme.

In addition to this contribution, they also presented a multi-retention, 16-way L1 cache composed by one fast set, which has a retention time of just $26.5\mu s$, and fifteen slow but highly nonvolatile sets. Write-intensive blocks are allocated in the first way, while the remaining blocks should be allocated in the regular cache sets. For this purpose, their system is able to keep track of the read/write operations done in the past to predict the best location for each cache line, and to migrate data between the low-retention and the high-retention portions of the cache memory. They claimed substantial energy consumption reduction while maintaining the same performance when comparing their implementation to a conventional all-SRAM cache implementation.

3 Magnetic Logic Circuits

A number of logic circuits based on MRAM have been presented and prototyped in the last five years. Some of them have been already under development by the industry




Figure 5: Relationship between switch duration Dt and ξ . Dt can be reduced greatly with low ξ (e.g. 35) while keeping the same $I_{switch} = 300\mu A$. The speed can be further improved by scaling down MTJ size.

towards practical applications. We have seen that from physical point of view we can address high computing performance with STT-MRAM, we propose here an overview and current status of these magnetic VLSI circuits.




Figure 6: (a) MLUT structure (b) MFF and its non-volatile master register.

3.1 Magnetic Look-Up Table (MLUT)

We imagine that Magnetic Look-Up Table (MLUT) could be proposed as replacement of the conventional SRAM based LUT in the FPGA circuits (see Figure 6(a)) [36]. We propose here a new Non-Volatile, dynamic reconfigurable LUT based architecture. A

non-volatile configuration point of MLUTs consists of an SRAM based Sense Amplifier (SA) associated with a couple of complementary MTJs [29]. FPGA circuits can be configured instantaneously and the high-speed SA ensures nearly the same speed as SRAM-LUT [19–24]. Thanks to the small cell area and 3D integration of MRAM, multi-context can be easily implemented, allowing dynamical and run-time reconfiguration methods [26, 29, 37]. The first prototype of MLUT has been presented in 2009 based on 130 nm hybrid MRAM/CMOS process [21].

3.2 Magnetic Flip-Flop (MFF)

MFF stores the intermediate computing data in non-volatile mode. It was firstly proposed in 2006 and the idea is to replace one of the master or slave parts by the circuit similar to a non-volatile configuration point [38].

In 2008, NEC presented the first prototype based on 0.15 μm hybrid process and high performance up to 3.5 GHz was shown. MFF is expected to make low standby power for electrical appliances like LCD TV, PC and portable devices in the next years [23]. MFF is the key element to build non-volatile logic circuits and allows true instant on/off and zero standby power [29].

As the switching of MRAM cell consumes always much more than a SRAM, two special techniques have been developed to reduce the switching power of MFF (see Equation 10), which should be operating at a very high frequency f_{switch} (>MHz) [38, 39]. One is checking point: MFF stores the intermediate data in MTJs with a manageable slow frequency (e.g. 1KHz) [40] and the other is “write before sleep”: MFF stores the intermediate data in MTJs, as there is a request by user or in case of power failure.

$$P_{dynamics} = f_{switch} * \int_0^T V_{dd} * I_d(t) dt \quad (10)$$

3.3 Magnetic NAND/NOR/XOR logic gates

Magnetic non-volatile NAND, NOR and XOR logic circuits have been presented and demonstrated experimentally using different switching techniques [41, 42]. The intermediate logic computing is based only on magnetic signals, which reduces the transition times between magnetic and electrical signals through CMOS switching/sensing circuits. Nevertheless, magnetic signals are difficult to be propagated and reinforced in case of attenuation; thereby these circuits cannot be easily cascaded to build up complex logic circuits. Important breakthroughs are required to explore these logic gates for practical applications.

3.4 Magnetic Full Adder (MFA)

MFA, one of the innovative circuits based on hybrid MTJ/CMOS circuits, has been proposed to build ultra-low power high density ICs. However, as there is no Error Correction Codes (ECC) block for advanced magnetic logic circuits, high reliability becomes one crucial bottleneck following the miniaturization of fabrication node beyond 90 nm.

A new MFA design recently presented, which is based on Pre-Charge Sense Amplifier (PCSA) circuit [29] and logic-in-memory [43, 44] (see Figure 7), has demonstrated a satisfying reliability performance. Thanks to the non-volatile data stored in MTJ, this MFA can also provide an instant ON/OFF to consume nearly zero static power. Moreover, the 3D integration of memory cell above logic circuits reduces significantly both the area cost and the interconnection delay between the power and area efficiency.

Other switching approaches like STT and domain wall motion [8] are also under investigation. Based on this MFA, more complex magnetic logic circuits like nonvolatile arithmetic logic unit (NV ALU) can be built up, which may be used as the elementary computing unit of a non-volatile CPU (NV CPU). This new processor presents great interest for low-power applications as it promises to overcome both the static and dynamic power consumption.

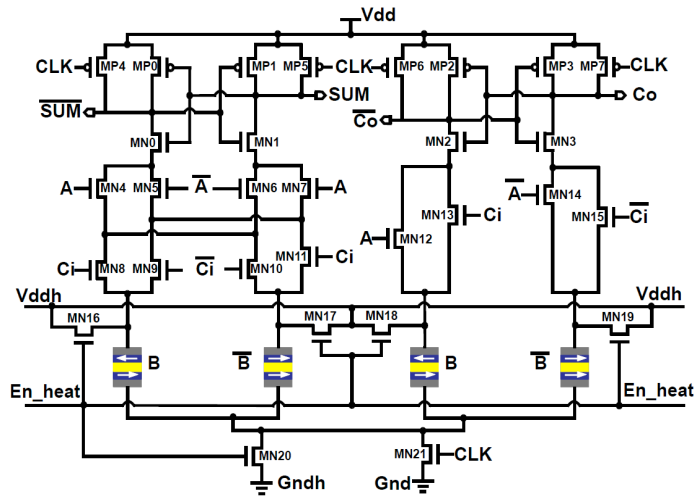


Figure 7: Magnetic Full Adder (MFA) based on high stability, low power Pre-Charged Sense Amplifier (PCSA) and “logic-in-memory” architecture.

3.5 Shift Register Based on Domain Wall

The recent progress demonstrated that the domain wall (DW) motion in a magnetic strip driven by a short spin-polarized current pulse is a promising technology to build up nonvolatile memory device with its low-power, high-speed, high-density performance. A shift register concept based on domain wall, which is the fundamental building block of magnetic racetrack memory, has been recently presented (see Figure 8) [45].

Individual writing and reading MTJs are used to write and sense a series of data bits stored in the nanowire. The moving direction of the DWs depends on the current flow direction. It is the high velocity of DWs exceeding 100 m/s that makes it possible to be




Figure 8: Structure of domain wall propagation shift register, which is composed of write head (MTJ_0), read heads (MTJ_1 and MTJ_2) and storage magnetic stripe.

moved over distances of several micrometers within a few nanoseconds. Furthermore, as lower resistance can reduce the rate of breakdown and higher resistance can improve the sensing performance, the size of the read heads should therefore be smaller than that of the write heads to obtain the best switching and sensing reliability. By using this shift register based on current induced DW motion, some new digital circuits and architectures based on DW can be built up like configuration memory of FPGA circuit.

Besides the magnetic logic circuits shown above, embedded MRAM dedicated to high-speed and low-power applications have been also developed [15, 24]. They are compatible with advanced logic fabrication process and are expected to be used as cache memory in processor and working memory in micro-controller to improve their performances and simplify the architectures [44]. Another potential application of MRAM is into the FPGA domain combining MLUT [29] and MFF [45] shown previously. In the next section we propose a new FPGA based on embedded MRAM.

3.6 Magnetic-RAM-based FPGA (MFPGA)

This section summarizes detailed information about a full Magnetic FPGA design [20, 26, 37] and preliminary results are discussed. It is another typical example of the application driven by MRAM technology. The technology targeted is 130nm for CMOS part and 120nm for the TAS-MRAM cells.

Magnetic Tile Architecture. The MFPGA is an array of magnetic tiles (MTILE), each one containing an SBOX of 16 tracks and 4 Logic Elements (LE). These LEs are based on a LUT-4 and a Flip-Flop. A dedicated interconnection scheme with the SBOX is used to manage input/output of the LEs. Figure 9 describes the general architecture of the Tile, and the TAS-MRAM based logic elements such as runtime reconfigurable LUTs, multi-context LUT and the Magnetic Flip-Flop are further described in [37].




Figure 9: Magnetic Tile Architecture

To control the configuration Tile a specific configuration chain based on a serial register (scan chain) allows configuring directly the MRAM. In this way two configuration modes are available at same time, one in the SRAM part, and the other in the MRAM. At any time, it is possible to dynamically reconfigure each Tile by just configuring the MRAM cells.

The reconfiguration time for a Tile is about 260 clock cycles for the CMOS scan chain, plus the time to configure the MRAM part (to transfer CMOS data into MTJ), which is about 200ns. Once the configuration is loaded in the MRAM part, the read time to apply the configuration to the user logic is done in just 4 ns.

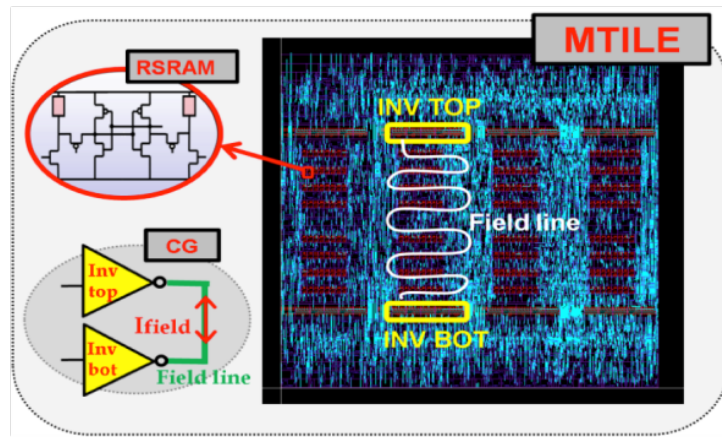


Figure 10: MTile layout, showing both the RSRAM cell and the current generators (Metal-2 only).

Core Architecture. The core architecture of the FPGA is based on a 19x19 regular array of tiles. The general core architecture is presented in Figure 11 and its layout in Figure 12.

The configuration time for the overall FPGA is about 93860 clock cycles plus time to transfer CMOS data to MRAM part. The dynamical reconfiguration granularity is done at the Tile level. In the same picture, the full layout of this FPGA is shown. For further details, refer to [37].

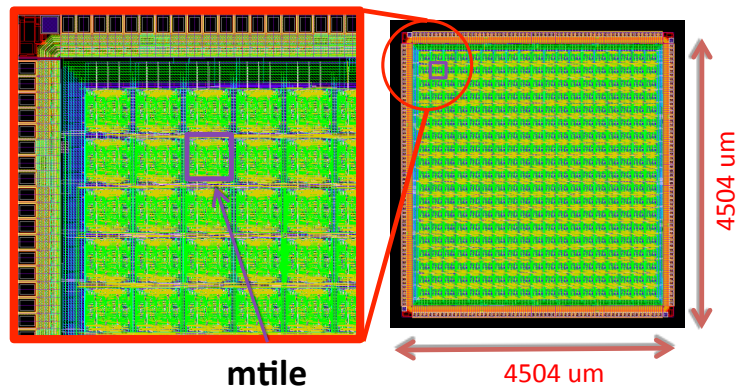


Figure 11: Magnetic FPGA Core Layout

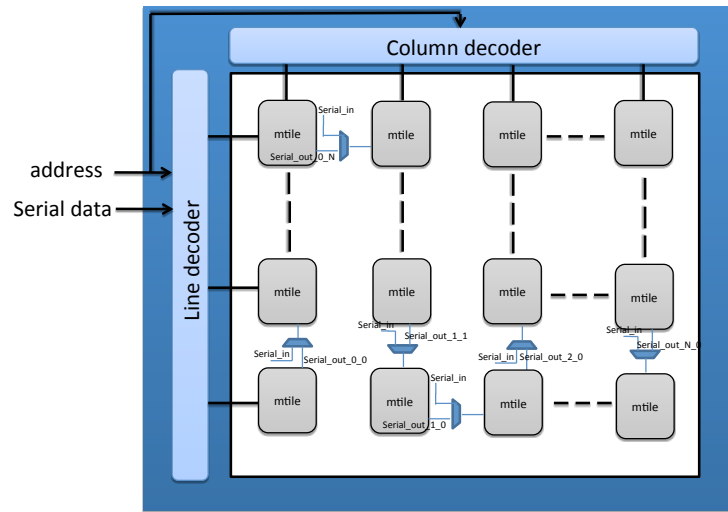


Figure 12: Magnetic FPGA Core Architecture

4 Embedded MRAM for Processor Applications

In this section, potential ideas of MRAM applications for microprocessors are analyzed, both for the processor core and for the memory hierarchy.

4.1 MRAM Applications for Processor Cores

In this context, the most straightforward idea is to replace all the SRAM-based registers by MRAM registers. The processor could then be turned-off while retaining its context. Conversely, a conventional processor would have to remain in sleep mode, consuming energy.

Since current MRAMs cannot perform as fast as SRAMs, a hybrid MRAM/SRAM cell was proposed in [38]. These registers keep their information stored as an electrical charge during their regular operation. In the presence of a power failure or a request for entering the sleep mode, the data can be saved in their MTJs.

Another approach is to provide the operating systems with an infrastructure for context switching. Currently, this feature is available either by providing several register windows or by saving and restoring the register contents in the cache memory. MRAM-based multi-context registers [38] can be used to save multiple contexts locally, avoiding transactions with the cache hierarchy.

The third concept is to provide hardware support for rolling-back mechanisms. State rolling-back on its own is useful for dealing with unexpected system failures. By creating system checkpoints from time to time, one can then return to the last safe state by restoring the last checkpoint and continue from there onwards.

These checkpoints can also be stored in the multi-context register, in the same fashion of the context-switching implementation. As mentioned before, due to the use

of magnetic storage, the snapshots are intrinsically radiation-hardened and protected against power supply failures. Should we provide a way to access the contexts without disrupting the current one, temporal redundancy could also be implemented on top of this infrastructure.

4.2 MRAM applications for the processor memory hierarchy

MRAM can also be used for building the processor memory hierarchy. As mentioned earlier, even though the MRAM density is between two and four times higher than the SRAM's, its access time is between three and ten times higher. When using the same silicon area for implementing both variations, the first might compensate the latter.

In [44], for example, a 2 MB L2 SRAM Cache was replaced with an 8 MB L2 MRAM Cache, using roughly the same silicon fingerprint. In their particular case, the increase on the cache size was not enough to compensate the penalty due to the cache access delay. By employing write buffers and a novel cache access policy, they managed to achieve similar performance while reducing the power consumption by almost 74%. They also present a hybrid MRAM/SRAM cache organization, having 31 sets implemented in MRAM and 1 set implemented in SRAM. The write-intensive data is kept in the SRAM part, in order to mitigate the higher write delay. A method for determining which data is suitable for being placed in the SRAM set is also discussed.

We propose here a preliminary study on the impact of cache delays on the processor's performance, based on the SimpleScalar simulator [46] and the set of benchmarks Mediabench 1 [47]. This simulator implements a MIPS-like architecture, and the memory hierarchy can be fully configured in terms of capacity, access speed and access policies. Our goal was to determine whether replacing L1 SRAM caches by L1 MRAM caches, while keeping the same silicon fingerprint, is worthwhile.

SimpleScalar was configured to mimic a processor designed for embedded applications, such as LEON3 [48]. The baseline configuration is described in Table 2. It consists of a single processor having a single cache level and a large external memory, assumption that can be considered for many systems.

Differently from our previous work in [49], we assumed that the MRAM density is four times the SRAM's [50]. We are then comparing, for instance, a 4 KB SRAM-based cache with a 16 KB MRAM-based cache.

For this set of experiments, we assumed a latency of 3 clock cycles during each cache access. It means that the processor will stall upon each cache request, waiting for the data to become available. We also assumed a latency of 1000 cycles for the external memory to make the first word available, and 10 cycles for each subsequent word while doing burst reading [51].

In Figure 13, we compare a 1KB SRAM cache with a 4KB MRAM-based memory. We can see that the increase in capacity can easily compensate for the delay in such a case.

In the same manner, as shown in Figure 14, where a 128 KB SRAM cache is compared with its 512 KB MRAM counterpart, the latter shows comparable performance to the smaller, yet faster SRAM.

Table 2: SimpleScalar baseline configuration used in all experiments

Option	Value	Meaning	Configuration
-cache:memlat 1st	1000	Ext. memory latency - 1 st word	1000 cycles
-cache:memlat burst	10	Ext. memory - burst	10 cycles per word
-cache:d12	none	L2 Data Cache Parameters	No L2 Cache
-cache:i12	none	L2 Instr. Cache Parameters	No L2 Cache
-res:ialu	1	# of integer ALUs	1 integer ALU
-res:falu	1	# of floating-point ALUs	1 FP ALU
-res:imult	1	# of integer multipliers	1 Multiplier
-fetch:ifqsize	1	Instruction Fetch (IF) Queue Size	1 IF per cycle
-fetch:mplat	1	Branch misprediction latency	1 extra cycle
-cache:d11 linesize	32	L1 Data Cache Line Size	32 bits per cache line
-cache:d11 policy	1	L1 Data Cache Replacement Policy	Least Recently Used (LRU)
-cache:i11 linesize	32	L1 Instr. Cache Line Size	32 bits per cache line
-cache:i11 policy	1	L1 Instr. Replacement Policy	Least Recently Used (LRU)

Table 3: SimpleScalar configuration used in Figures 13-15

Option	Value	Meaning	Configuration
-cache:d11 assoc	2	L1 Data Cache Line Size	2-way associative
-cache:d11lat	3	L1 Data Cache Access Latency	SRAM: 1 cycle, MRAM: 3 cycles
-cache:i11 assoc	2	L1 Instr. Cache Line Size	2-way associative
-cache:i11lat	3	L1 Instr. Cache Access Latency	SRAM: 1 cycle, MRAM: 3 cycles




Figure 13: Overview of the processor performance using low-capacity L1 caches.

In order to generalize this conclusion, let us then define the CPI penalty as the increase in the CPI caused by replacing an SRAM cache with an MRAM cache using the same silicon area, as follows:




Figure 14: Overview of the processor performance using high-capacity L1 caches.

$$CPI_{penalty} = \frac{CPI_{MRAM}}{CPI_{SRAM}} - 1 \quad (11)$$

Based on the $CPI_{penalty}$, in Figure 15, the best-case, the worst-case and the average performance over the benchmark set are shown as a function of the cache capacity. Given our assumptions are valid, MRAM does present a CPI gain rather than a CPI penalty for most cases. Once the cache capacity is large enough to contain the whole benchmark data, the CPI gain turns into a penalty which can no longer be compensated if no specific technique is employed.




Figure 15: Overview of CPI Penalty: best-case, worst-case and average of the Media-bench benchmarks' performance.

4.3 Case Study: CJPEG

Based into the data analysis for the set of benchmarks, one was selected for a in-depth analyses, the CJPEG. The choice for this particular algorithm is because it is a data-driven benchmark performing a substantial number of access to the cache memories.

Also for this particular benchmark, we performed the tests for different associativities 1,2 and 4-way caches. We also adjusted the delays of the cache latency, focusing on three and ten for the MRAM latency. Also is assumed as our reference that the SRAM has one cycle delay [52–57].

The Figure 16 and Figure 17 are depicting the two cases regarding delay of 1000 clock cycles for the first word been retrieved from the external memory [51,58] and one cycle between subsequent words in burst. A latency of 1000 cycles is assumed to be more conservative, regarding the external memory access and the circuitry in the path to access it [59].

This case study assumes that for the same silicon fingerprint the MRAM integration density is x4 in comparison to SRAM (for instance for 1KB of SRAM, meaning it is possible to integrate in the same silicon area a 4KB MRAM memory), considering also as parameter the associativity available to access to the cache memory (1, 2 and 4 way). All the results are normalized by the baseline SRAM-L1 of 1KB, delay of 1 cycle. This way we compare SRAM [1,2,4,8,16,32,64]KB with MRAM [4,16,32,64,128,256]KB. The comparisons are based into the CPI penalty, and CPI penalty is defined as specified in Equation 11.

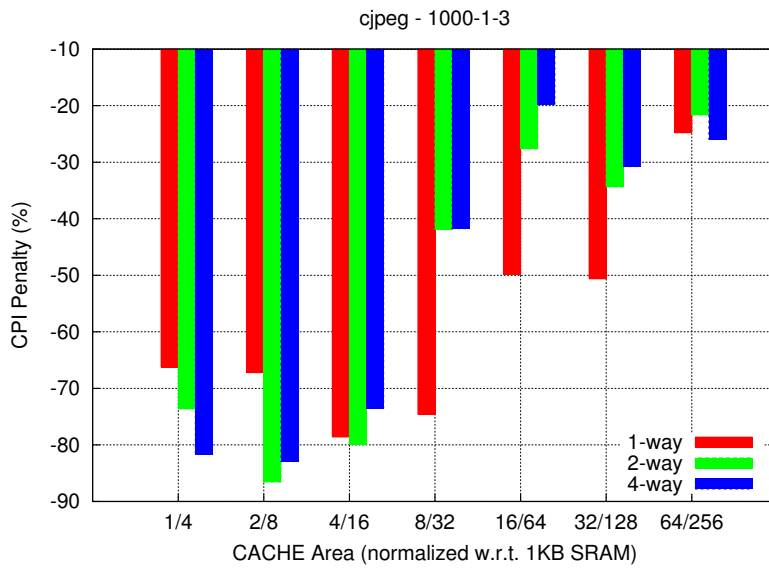
Observing the Figure 16(a) and Figure 16(b), is possible to notice that an increase of seven cycles into the cache latency were necessary to affect the MRAM cache performance for sizes larger than 16 KB. We also note that the 1-way cache kept steady for sizes 16 and 32 KB despite the MRAM cache delays, this result is widely described into [53]. Similar results are observed in Figure 17(a) and Figure 17(b), for a different MRAM cache parameters, also in this case a delay of ten cycles between subsequent words are assumed for read burst from the main memory.

So the strategy of a n-way cache is interesting for caches larger than 16 KB. The SRAM despite the density of 4 times smaller than the MRAM (for the same silicon area) outperforms the MRAM for sizes larger than 32KB.

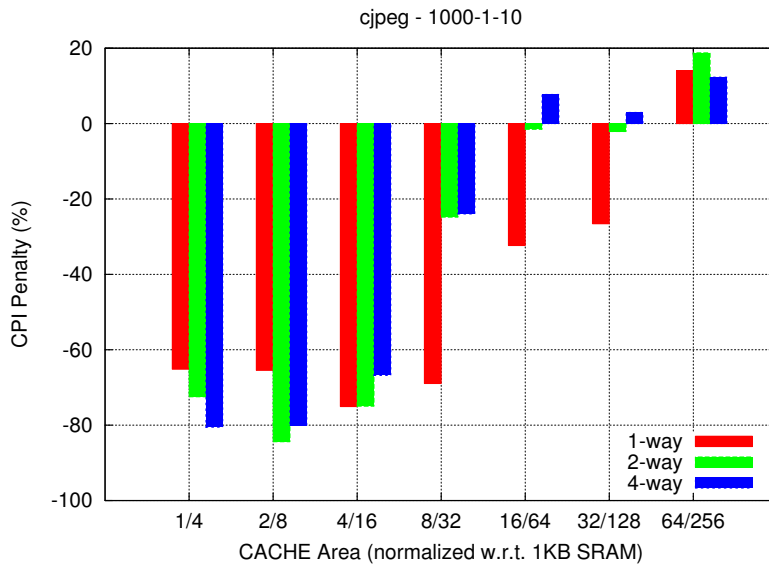
The cache size combined with the associativity have a relevant impact on the cache miss ratio. For this reason, MRAM obtained a better result, despite the higher access delay. But for caches larger than 16 KB we should use an additional mechanism to mitigate the MRAM delay, assuming the worst case scenario of ten cycles of cache latency.

So, when using an MRAM L1 cache in a microprocessor, the higher density mitigates the lower delay up to 64 KB in this particular experiment. For L1 caches beyond this capacity, the MRAM density itself is not enough to mitigate the delay.

For these cases one possible solution could be the usage of write-buffers, or a MRAM cache working in a higher frequency and phase-aligned. Another approach would be a hybrid MSRAM cell, for L1 cache, for L2 caches the access delay of the current state of the technology are not an issue, only adopting a L2 MRAM it already cuts the leakage current in 70% for a SoC as described in [44].




(a) 1000 cycles for the first word, than 1 for the following, 3 cycles cache delay




(b) 1000 cycles for the first word, than 1 for the following, 10 cycles cache delay

Figure 16: Simulation results for 1000 cycles and 1 cycle of delay for each subsequent word in burst mode.



(a) 1000 cycles for the first word, than 10 for the following, 3 cycles cache delay



(b) 1000 cycles for the first word, than 10 for the following, 10 cycles cache delay

Figure 17: Simulation results for 1000 cycles and 10 cycles of delay for each subsequent word in burst mode.

5 Conclusions

This work presented both an overview of the currently available MRAM technologies and their applications for FPGAs and embedded processors. Apart from the innumerable advantages intrinsic to this technology, it was shown that the writing speed is still a concern for all of the currently available MRAM implementations.

Nevertheless, by either taking advantage of the higher MRAM density or by applying compensation techniques such as the “write-before-sleep”, it is possible to achieve competitive performance with this technology.

We discussed the building blocks and demonstrated an implementation of an FPGA built over the MRAM technology. The FPGA has essentially the same architecture of a standard implementation, but for the basic cells, which were replaced by magnetic-based elements.

Finally, a comparative study on MRAM applications for embedded processors was presented. It was shown that, despite of higher latency introduced by the magnetic memory cells, for caches up to a certain value (64 KB in our experiments) this problem is mitigated by the higher density of the MRAM. For some applications, we can even reach comparable-to-better performance with the MRAM counterparts.

Based on these results, we intend to continue our research towards the magnetic microprocessor and FPGA. One possible idea is to combine both on the same silicon die, using the FPGA as a dynamically reconfigurable accelerator. Specifically for processor, there are still improvements in the memory hierarchy to be done, aiming to further counterbalance the latency of nowadays MRAM implementations.

Acknowledgements

The authors wish to acknowledge support from the French National Agency (ANR) projects CILOMAG, NANOINNOV SPIN, NANO2012 and MARS. We thank also G. Prenat and B. Dieny from SPINTEC laboratory, O. Redon from CEA LETI, K. Torki and G. Dipendina from CMP and L. Prejbeanu from Crocus Technology for decisive inputs and scientific discussions.

References

1. Weste, N., Harris, D.: CMOS VLSI Design: A Circuits and Systems Perspective. 4th edn. Addison-Wesley Publishing Company, USA (2010)
2. Kang, S., Leblebici, Y.: CMOS digital integrated circuits: analysis and design. McGraw-Hill series in electrical engineering. McGraw-Hill (1999)
3. Kim, N., Austin, T., Baauw, D., Mudge, T., Flautner, K., Hu, J., Irwin, M., Kandemir, M., Narayanan, V.: Leakage current: Moore’s law meets static power. *Computer* **36**(12) (dec. 2003) 68 – 75
4. Slaughter, J., Dave, R., Durlam, M., Kerszykowski, G., Smith, K., Nagel, K., Feil, B., Calder, J., DeHerrera, M., Garni, B., Tehrani, S.: High speed toggle mram with mgo-based tunnel junctions. In: *Electron Devices Meeting, 2005. IEDM Technical Digest. IEEE International*. (dec. 2005) 873 –876

5. Hoya, K., Takashima, D., Shiratake, S., Ogiwara, R., Miyakawa, T., Shiga, H., Doumae, S., Ohtsuki, S., Kumura, Y., Shuto, S., Ozaki, T., Yamakawa, K., Kunishima, I., Nitayama, A., Fujii, S.: A 64mb chain feram with quad-bl architecture and 200mb/s burst mode. In: Solid-State Circuits Conference, 2006. ISSCC 2006. Digest of Technical Papers. IEEE International. (feb. 2006) 459–466
6. Wong, H., Raoux, S., Kim, S., Liang, J., Reifenberg, J., Rajendran, B., Asheghi, M., Goodson, K.: Phase change memory. *Proceedings of the IEEE* **98**(12) (dec. 2010) 2201–2227
7. Kund, M., Beitel, G., Pinnow, C.U., Rohr, T., Schumann, J., Symanczyk, R., Ufert, K.D., Muller, G.: Conductive bridging ram (cbram): an emerging non-volatile memory technology scalable to sub 20nm. In: Electron Devices Meeting, 2005. IEDM Technical Digest. IEEE International. (dec. 2005) 754–757
8. Chappert, C., Fert, A., Van Dau, F.N.: The emergence of spin electronics in data storage. *Nature Materials* **6**(11) (nov. 2007) 813–823
9. Wolf, S.A., Awschalom, D.D., Buhrman, R.A., Daughton, J.M., Von Molnár, S., Roukes, M.L., Chtchelkanova, A.Y., Treger, D.M.: Spintronics: a spin-based electronics vision for the future. *Science* **294**(5546) (2001) 1488–1495
10. Freescale: Freescale leads industry in commercializing mram technology; 4 mbit mram memory product now in volume production. <http://investors.freescale.com/phoenix.zhtml?c=175261&p=irol-newsArticle&ID=880031> (jul. 2006)
11. Prejbeanu, I.L., Kerekes, M., Sousa, R.C., Sibuet, H., Redon, O., Dieny, B., Nozières, J.P.: Thermally assisted mram. *Journal of Physics: Condensed Matter* **19**(16) (2007) 165218
12. Sun, J.Z.: Spin angular momentum transfer in current-perpendicular nanomagnetic junctions. *IBM Journal of Research and Development* **50**(1) (jan. 2006) 81–100
13. Kawahara, T., Takemura, R., Miura, K., Hayakawa, J., Ikeda, S., Lee, Y., Sasaki, R., Goto, Y., Ito, K., Meguro, I., Matsukura, F., Takahashi, H., Matsuoka, H., Ohno, H.: 2mb spin-transfer torque ram (spram) with bit-by-bit bidirectional current write and parallelizing-direction current read. In: Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International. (feb. 2007) 480–617
14. Parkin, S.S.P., Hayashi, M., Thomas, L.: Magnetic domain-wall racetrack memory. *Science* **320**(5873) (2008) 190–194
15. Lin, C., Kang, S., Wang, Y., Lee, K., Zhu, X., Chen, W., Li, X., Hsu, W., Kao, Y., Liu, M., Lin, Y., Nowak, M., Yu, N., Tran, L.: 45nm low power cmos logic compatible embedded stt mram utilizing a reverse-connection 1t/1mtj cell. In: Electron Devices Meeting (IEDM), 2009 IEEE International. (dec. 2009) 258–259
16. Tsuchida, K., Inaba, T., Fujita, K., Ueda, Y., Shimizu, T., Asao, Y., Kajiyama, T., Iwayama, M., Sugiura, K., Ikegawa, S., Kishi, T., Kai, T., Amano, M., Shimomura, N., Yoda, H., Watanabe, Y.: A 64mb mram with clamped-reference and adequate-reference schemes. In: Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International. (feb. 2010) 258–259
17. ITRS: International roadmap for semiconductors 2007 and 2008 update. <http://www.itrs.net/Links/2007ITRS/Home2007.htm> (2007)
18. Black, W.C., Das, B.: Programmable logic using giant-magnetoresistance and spin-dependent tunneling devices (invited). *Journal of Applied Physics* **87**(9) (may 2000) 6674–6679
19. Zhao, W., Belhaire, E., Chappert, C., Mazoyer, P.: Spin transfer torque (stt)-mram-based runtime reconfiguration fpga circuit. *ACM Trans. Embed. Comput. Syst.* **9** (October 2009) 14:1–14:16
20. Guilleminet, Y., Torres, L., Sassatelli, G., Bruchon, N., Hassoune, I.: A non-volatile runtime fpga using thermally assisted switching mrams. In: Field Programmable Logic and Applications, 2008. FPL 2008. International Conference on. (sept. 2008) 421–426

21. Suzuki, D., Natsui, M., Ikeda, S., Hasegawa, H., Miura, K., Hayakawa, J., Endoh, T., Ohno, H., Hanyu, T.: Fabrication of a nonvolatile lookup-table circuit chip using magneto/semiconductor-hybrid structure for an immediate-power-up field programmable gate array. In: VLSI Circuits, 2009 Symposium on. (june 2009) 80–81
22. Yamamoto, S., Sugahara, S.: Nonvolatile delay flip-flop based on spin-transistor architecture and its power-gating applications. Japanese Journal of Applied Physics **49**(9) (2010) 090204
23. Sakimura, N., Sugibayashi, T., Nebashi, R., Kasai, N.: Nonvolatile magnetic flip-flop for standby-power-free socs. In: Custom Integrated Circuits Conference, 2008. CICC 2008. IEEE. (sept. 2008) 355–358
24. Chaudhuri, S., Zhao, W., Klein, J.O., Chappert, C., Mazoyer, P.: Design of embedded mram macros for memory-in-logic applications. In: Proceedings of the 20th symposium on Great lakes symposium on VLSI. GLSVLSI '10, New York, NY, USA, ACM (2010) 155–158
25. Zhao, W., Belhaire, E., Dieny, B., Prenat, G., Chappert, C.: Tas-mram based non-volatile fpga logic circuit. In: Field-Programmable Technology, 2007. ICFPT 2007. International Conference on. (dec. 2007) 153–160
26. Guilleminet, Y., Torres, L., Sassatelli, G.: Non-volatile run-time field-programmable gate arrays structures using thermally assisted switching magnetic random access memories. Computers Digital Techniques, IET **4** (may 2010) 211–226
27. Devolder, T., Hayakawa, J., Ito, K., Takahashi, H., Ikeda, S., Crozat, P., Zerounian, N., Kim, J.V., Chappert, C., Ohno, H.: Single-shot time-resolved measurements of nanosecond-scale spin-transfer induced switching: Stochastic versus deterministic aspects. Phys. Rev. Lett. **100** (Feb 2008) 057206
28. Ikeda, S., Miura, K., Yamamoto, H., Mizunuma, K., Gan, H.D., Endo, M., Kanai, S., Hayakawa, J., Matsukura, F., Ohno, H.: A perpendicular-anisotropy cofeb–mgo magnetic tunnel junction. Nature Materials **9**(9) (jul. 2010) 721–724
29. Zhao, W., Chappert, C., Javerliac, V., Noziere, J.P.: High speed, high stability and low power sensing amplifier for mtj/cmos hybrid logic circuits. Magnetics, IEEE Transactions on **45**(10) (oct. 2009) 3784–3787
30. Faber, L.B., Zhao, W., Klein, J.O., Devolder, T., Chappert, C.: Dynamic compact model of spin-transfer torque based magnetic tunnel junction (mtj). In: Design Technology of Integrated Systems in Nanoscale Era, 2009. DTIS '09. 4th International Conference on. (april 2009) 130–135
31. Nepal, K., Bahar, R., Mundy, J., Patterson, W., Zaslavsky, A.: Designing mrf based error correcting circuits for memory elements. In: Design, Automation and Test in Europe, 2006. DATE '06. Proceedings. Volume 1. (march 2006) 1–2
32. Dieny, B., Sousa, R.C., Herault, J., Papisoi, C., Prenat, G., Ebels, U., Houssameddine, D., Rodmacq, B., Auffret, S., Prejbeanu, L.D.B., et al.: Spin-transfer effect and its use in spintronic components. International Journal of Nanotechnology **7**(4/5/6/7/8) (2010) 591
33. Zhao, W., Duval, J., Klein, J., Chappert, C.: A compact model for magnetic tunnel junction (mtj) switched by thermally assisted spin transfer torque (tas + stt). Nanoscale Research Letters **6**(1) (2011) 368
34. Worledge, D.C., Hu, G., Abraham, D.W., Sun, J.Z., Trouilloud, P.L., Nowak, J., Brown, S., Gaidis, M.C., O'Sullivan, E.J., Robertazzi, R.P.: Spin torque switching of perpendicular ta|cofeb|mgo-based magnetic tunnel junctions. Applied Physics Letters **98**(2) (2011) 022501
35. Sun, Z., Bi, X., Li, H., Wong, W., Ong, Z., Zhu, X., Wu, W.: Multi-retention level stt-ram cache designs with a dynamic refresh scheme. In: Proceedings of the 44th annual ACM/IEEE international symposium on Microarchitecture. MICRO 44, Porto Alegre, Brazil, IEEE Computer Society (dec. 2011) 329–338
36. Xilinx, I.: Spartan-3 fpga family data sheet. http://www.xilinx.com/support/documentation/data_sheets/ds099.pdf (dec. 2009)

37. Torres, L., Guillemenet, Y., Ahmed, S.Z.: A dynamic reconfigurable mram based fpga. In: *ERSA10 Keynote Paper*. (2010) 10
38. Zhao, W., Belhaire, E., Javerliac, V., Chappert, C., Dieny, B.: A non-volatile flip-flop in magnetic fpga chip. In: *Design and Test of Integrated Systems in Nanoscale Technology, 2006. DTIS 2006. International Conference on*. (sept. 2006) 323–326
39. Kang, S.H.: Embedded stt-mram for mobile applications: Enabling advanced chip architectures. In: *Non-Volatile Memories Workshop, San Diego, CA, USA* (apr. 2010)
40. Kothari, L., Carter, N.P.: Architecture of a self-checkpointing microprocessor that incorporates nanomagnetic devices. *Computers, IEEE Transactions on* **56**(2) (feb. 2007) 161–173
41. Behtash Behin-Aein, Deepanjan Datta, S.S., Datt, S.: Proposal for an all-spin logic device with built-in memory. *Nature Nanotechnology* **5**(4) (feb. 2010) 266–270
42. Allwood, D.A., Xiong, G., Faulkner, C.C., Atkinson, D., Petit, D., Cowburn, R.P.: Magnetic domain-wall logic. *Science* **309**(5741) (2005) 1688–1692
43. Matsunaga, S., Hayakawa, J., Ikeda, S., Miura, K., Hasegawa, H., Endoh, T., Ohno, H., Hanyu, T.: Fabrication of a nonvolatile full adder based on logic-in-memory architecture using magnetic tunnel junctions. *Applied Physics Express* **1**(9) (2008) 091301
44. Sun, G., Dong, X., Xie, Y., Li, J., Chen, Y.: A novel architecture of the 3d stacked mram l2 cache for cmps. In: *High Performance Computer Architecture, 2009. HPCA 2009. IEEE 15th International Symposium on*. (feb. 2009) 239–249
45. Lakys, Y., Zhao, W., Klein, J.O., Chappert, C.: Low power, high reliability magnetic flip-flop. *Electronics Letters* **46**(22) (28 2010) 1493–1494
46. Burger, D., Austin, T.M.: The simpliscalar tool set, version 2.0. *SIGARCH Comput. Archit. News* **25** (June 1997) 13–25
47. Lee, C., Potkonjak, M., Mangione-Smith, W.H.: Mediabench: a tool for evaluating and synthesizing multimedia and communications systems. In: *Proceedings of the 30th annual ACM/IEEE international symposium on Microarchitecture. MICRO 30, Washington, DC, USA, IEEE Computer Society* (1997) 330–335
48. Gaisler, A.: Leon3 multiprocessing cpu core. http://www.gaisler.com/doc/leon3_product_sheet.pdf (feb. 2010)
49. Zhao, W., Zhang, Y., Lakys, Y., Klein, J.O., Etienneble, D., Revelosona, D., Chappert, C., Torres, L., Cagnini, L., Brum, R., Guillemenet, Y., Sassatelli, G.: Embedded mram for high-speed computing. In: *VLSI and System-on-Chip (VLSI-SoC), 2011 IEEE/IFIP 19th International Conference on*. (oct. 2011) 37–42
50. Mackay, K.: Tas, tas+stt-mram and magnetic logic unit (nov. 2011) Property of Crocus Technology. Non authorized publication.
51. JC-42.3: Double data rate (ddr) sdram standard. Standard, JEDEC (2008) <http://www.jedec.org/standards-documents/docs/jesd-79f>.
52. Powell, M., Agarwal, A., Vijaykumar, T., Falsafi, B., Roy, K.: Reducing set-associative cache energy via way-prediction and selective direct-mapping. *Microarchitecture, 2001. MICRO-34. Proceedings. 34th ACM/IEEE International Symposium on* (2001) 54–65
53. Hennessy, J.L., Patterson, D.A.: *Computer architecture: a quantitative approach*. 4 edn. Volume 1. Elsevier - Morgan Kaufmann - Denise E. M. Penrose, 500 Sansome Street, Suite 400, San Francisco, CA 94111 (Jan 2007)
54. Patterson, D.A., Hennessy, J.L.: *Computer organization and design: the hardware/software interface*. (2005)
55. Boschma, B., Burns, D., Chin, R., Fiduccia, N., Hu, C., Reed, M., Rueth, T., Schumacher, F., Shen, V.: A 30 mips vlsi cpu. *Solid-State Circuits Conference, 1989. Digest of Technical Papers. 36th ISSCC., 1989 IEEE International* (1989) 82–83
56. Nambu, H., Kanetani, K., Yamasaki, K., Higeta, K., Usami, M., Fujimura, Y., Ando, K., Kusunoki, T., Yamaguchi, K., Homma, N.: A 1.8-ns access, 550-mhz, 4.5-mb cmos sram. *Solid-State Circuits, IEEE Journal of* **33**(11) (1998) 1650–1658

57. Alvarez, J., Barkin, E., Chao, C.C., Johnson, B., D'Addeo, M., Lassandro, F., Nicoletta, G., Patel, P., Reed, P., Reid, D., Sanchez, H., Siegel, J., Snyder, M., Sullivan, S., Taylor, S., Vo, M.: 450 mhz powerpc601 microprocessor with enhanced instruction set and copper interconnect. Solid-State Circuits Conference, 1999. Digest of Technical Papers. ISSCC. 1999 IEEE International (1999) 96 – 97
58. Gharachorloo, K., Gupta, A., Hennessy, J.: Performance evaluation of memory consistency models for shared-memory multiprocessors. ACM SIGARCH Computer Architecture News **19**(2) (Jan 4) 245–257
59. Gutierrez, A., Dreslinski, R., Wenisich, T., Mudge, T., Saidi, A., Emmons, C., Paver, N.: Full-system analysis and characterization of interactive smartphone applications. Workload Characterization (IISWC), 2011 IEEE International Symposium on (2011) 81 – 90