



MATCH-IN-PLACETM

**A NOVEL WAY TO PERFORM SECURE AND FAST USER'S
AUTHENTICATION**

**IMPLEMENTATION INTO A
MAGNETIC LOGIC UNIT**

**Bertrand Cambou
CROCUS Technology**

SUMMARY

In the last 20 years, secure microcontrollers provided quick and mainly secure methods to authenticate users. These chips store a reference pattern in their embedded memories such as a pin code, a biometric print, or a secure private key. The crypto-processors of the chips simply compare patterns entered for or by the user with stored patterns read during the authentication cycle.

This traditional architecture presents some limitations. Non-volatile memories that are commonly used in secure chips tend to be slow, so the stored keys are potentially exposed to hackers during the process. An attack may not always be detected before the full completion of the authentication cycle, creating potential security breaches.

This paper describes a new architecture, called Match-In-Place™, which Crocus Technology has developed on its “Magnetic Logic Unit™” (MLU™) to authenticate users without exposing any confidential data to hackers. Advantages expected are:

- The stored patterns are never read, and never exposed to potential hackers
- The matching cycles are very quick and could be orders of magnitude faster than existing methods and burn less power
- Match-In-Place engines could act as a hardware accelerator, and simplify the overall chip
- The cost structures of the secure chips are considerably reduced

Each cell of the Match-In-Place™ architecture is a non-volatile memory point combined with a virtual XOR gate (exclusive OR). Multiple cells, typically 16-512, are connected in series to form a NAND chain acting as a linear Match-In-Place™ engine. Multiple NAND chains together form a bi-dimensional “Match-In-Place™” matrix that match simultaneously one input pattern with multiple stored patterns.

Fields of uses of this new architecture are quite wide and include secure microcontrollers, SIM cards, banking cards, biometric authentication chips, NFC, and hardware acceleration. Benefits will be enhanced security, lower cost, and faster response time.

INTRODUCTION

Crocus Technology was founded 5 years ago in Grenoble France out of the CEA (Commissariat a l'Energie Atomique) with the full support of a group of energetic international investors. The corporation located part of its team in Silicon Valley, and developed the fully integrated MRAM (Magnetic Random Access Memory) technology at SVTC a local corporation with semiconductor pilot line. The technology was transferred for volume manufacturing from SVTC to Tower Jazz Semiconductor in Israel. The commercialization from Tower is expected for mid-2012. Crocus Technology has recently announced a partnership with the Russian firm Rusnano to extend its manufacturing capability to 300mm wafer size and 65nm lithography processing.

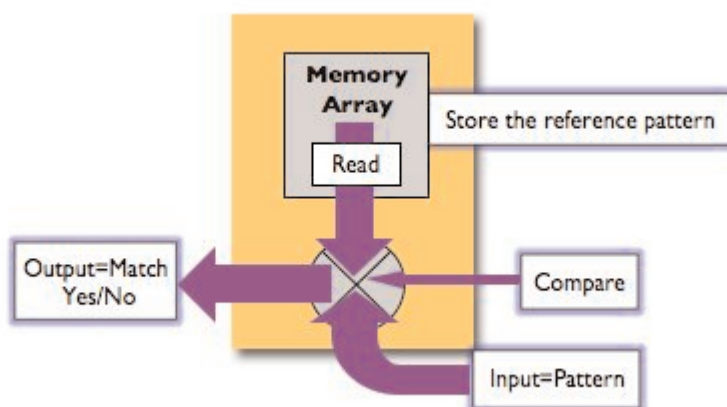
In addition to standard memory products, Crocus has developed an important variation of the MRAM cell, the Magnetic Logic Unit (MLU) that allows design engineers to turn these cells into logic gates and realize multiple logic functions such as an "XOR" at the magnetic level without any transistor level CMOS logic embedded in the arrays. The purpose of this paper is to describe a new architecture, Match-In-Place™, that leverages MLU arrays for the purpose of designing new applications such as secure chips, search engines, pattern recognition, or hardware acceleration. For the readers who wish to understand the basic structure behind the architecture we included appendixes describing the magnetic technology, and experimental measurements done on the silicon produced for Crocus by Tower Jazz semiconductor.

READ AND COMPARE VERSUS "MATCH-IN-PLACE™"

Let us describe first the current methods prevailing in existing secure microcontroller chips: "read then compare".

Figure 1 shows the architecture of the traditional authentication system. The patterns are stored in a mature non-volatile memory, usually a NOR flash or EEPROM, either of which stores arrays of "0's" and "1's" as charges above a thin dielectric on each cell. Data which is stored in such memories can typically be read in between 50ns to 150ns.

Figure 1. Read and Compare

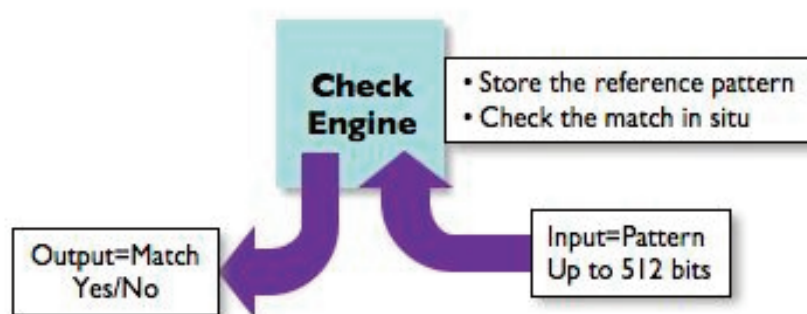


The comparator (or a crypto-processor) reads the stored pattern then compares it to the incoming pattern to authenticate the user. In this approach the stored pattern could be exposed to a hacker during the authentication process because this stored pattern needs to be read for comparison purpose not knowing if the input pattern is correct or not. Hackers have developed sophisticated methods to read the stored pattern, for example by presenting random input patterns posing as passwords. A well-funded hacker with access only to the memory chip can use electron beam detection and other sophisticated techniques to reveal the stored pattern.

Let us now describe Crocus' architecture with its MLU technology and how this fundamentally departs from the traditional method.

Figure 2 illustrates the alternative authentication approach based on a direct Match-In-Place™ to compare the input pattern with the stored pattern. In this method all bits of the input pattern are simultaneously compared with the corresponding bits stored in the Check Engine.

Figure 2. Match-In-Place™ Authentication



The Crocus implementation uses magnetic memories to allow up to 512 bits to be checked in about 15-20ns without ever exposing the stored pattern.

This produces a real “Zero Knowledge Proof” authentication method that never reveals information other than the level of match.

These engines do not have any conventional “read” cycles; they only perform a “matching” function, so hackers can never be in a position to break the system without presenting the correct input pattern. Long and complex patterns combined with standard cryptography methods could offer an unprecedented level of security.

We are now going to describe step by step how the Match-In-Place™ technology is constructed, and why Crocus' MLU is well suited for an effective implementation.

SINGLE CELL MATCH-IN-PLACE™ ENGINE

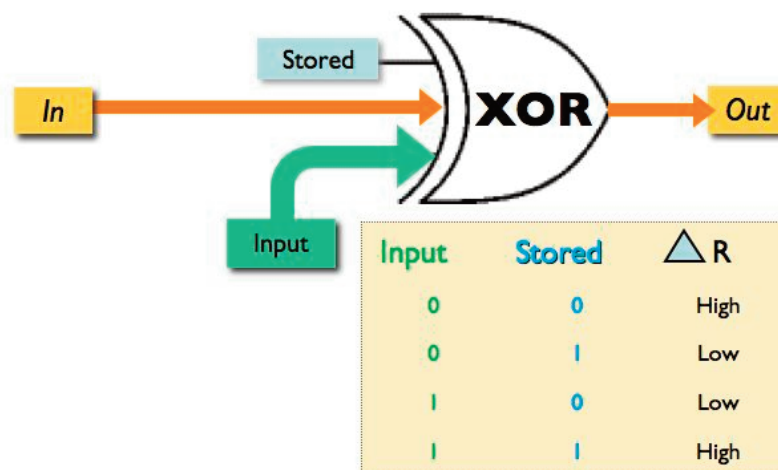
The basic Match-In-Place architecture is built at three levels: The single cell for direct matching at the bit level, NAND chains that combine multiple individual cells in series for linear Match-In-Place™ engines, and a matrix that combine multiple NAND chains in parallel to match one input pattern to a stack of stored patterns.

The foundation of the architecture is a quite sophisticated single cell that is very difficult to build economically with conventional semiconductor elements. These memory cells are described in the Match-In-Place. The single cell has to perform the Match-In-Place™ quickly in a native way without external circuitry

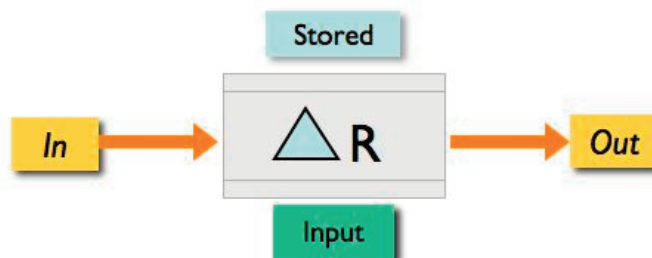
Unlike standard memory cells that are only capable of storing, erasing, and reading individual bits, this cell can also directly compare the input bit to the stored bit, a necessary function for Match-In-Place. The single cell has to perform the Match-In-Place™ quickly in a native way without external circuitry.

Figure 3. Single-Cell Architecture

Logical



Electrical



Similar to a CAM (Content Addressable Memory), common in network engines or fast microprocessors, this cell perform an XOR (Exclusive or) operation.

When the input bit matches a stored bit (a “0” with a “0” or “1” with a “1”) the cell will measure a match. Conversely, if the input data does not match the cell’s contents (a “0” with a “1” or “1” with a “0”) the cell will measure a mismatch.

Unlike a RAM-based CAM, an MLU-based Match-In-Place cell will retain the stored bit when power is removed. This avoids any need to reload the stored information, a step that exposes the information to hacking. An MLU-based Match-In-Place™ engine can handle unlimited write cycles, and a read or match cycle does not disturb the stored information. The read cycle of an individual cell is extremely fast; a compare can occur in as little as 5ns. Program cycles can be implemented through a separate path with a cycle time of about 20ns.


In brief, the MLU cells are highly suitable for the fast Match-In-Place architecture. In addition, these cells are less sensitive to α (alpha) particle radiation than traditional SRAM cells and scales better to smaller process geometries than SRAM. As a result, some SRAM designs are migrating from smaller 6-transistor memory cells to larger 8-transistor cells in new designs simply to overcome issues concerning cell stability at 20-30nm processes. An MLU embodiment of the cell in Match-In-Place. The single cell has to perform the Match-In-Place™ quickly in a native way without external circuitry assumes that the cell exhibits a high resistance for a match, and low resistance for a mismatch.

NAND CHAINS FOR MATCH-IN-PLACE™ ENGINES

Multiple cells connected in series form a NAND chain that can create a linear Match-In-Place™ engine. The sensing elements are located at the end of the chains to extract the level of matching.

Figure 4 describes a NAND-chain. The input pattern $\{I_1, I_2, \dots, I_i, \dots, I_{n-1}, I_n\}$ is compared at with the stored pattern $\{S_1, S_2, \dots, S_i, \dots, S_{n-1}, S_n\}$ in a NAND chain of XOR (exclusive OR) gates. The output of the NAND chain indicates either that all bits in the chain match, or that one or more bits do not match the stored value.

Figure 4. NAND Chain



The circuit in figure 4 uses a self-referenced two phase cycle to sense the level of matching within about 10 to 20ns. During the first phase of the cycle the resistances of all individual cells of the NAND chain are measured at once when all the input lines are idle (not driven) to get a reference to be sampled. Each cell has been set up in such a way that its resistance, when idle, is the same as that of a match. The measurement of this first step is temporarily 'held' as a voltage level in a sample-and-hold circuit (shown as the capacitor in the diagram), a step that takes less than 5ns. During the second phase the resistance of the NAND chain is measured again when the input pattern is compared to the held pattern. The sensing module compares the difference between the held and sensed resistances. Such a self-referenced scheme, which compares the relative values of the two resistances rather than comparing a resistance against an absolute level, is insensitive to manufacturing process variations as well as physical parameters like temperature that might impact each cell's intrinsic resistance.

The sensing circuit creates an output signal to quantify the level of mismatch. Both signals are identical when all bits perfectly match their stored reference; this yields a signal that is close to zero. The output signal will increase when the number of cells with a mismatch increases. This constitutes a linear Match-In-Place™ engine as described in figure 2.

These NAND chains could be designed to render it impossible to read the contents of the information stored in the chain. The integrity of the stored pattern can be checked just after programming by performing a compare with the correct data without reading (and possibly revealing the contents of) any individual cells. An entire NAND chain can be written to at one time.

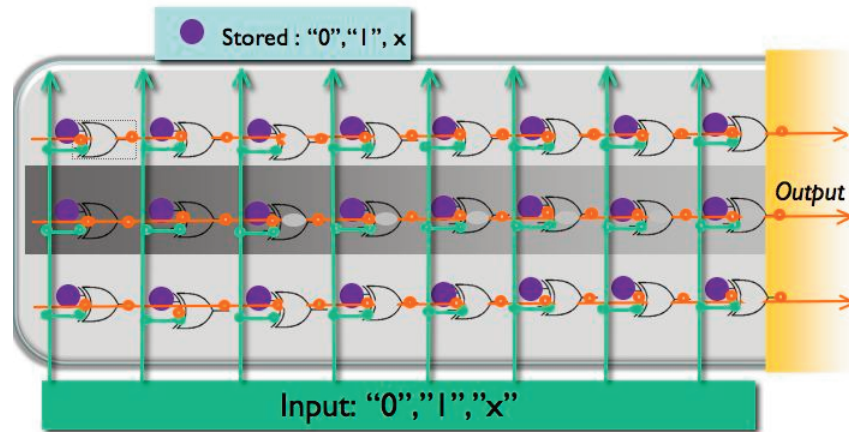
Because the NAND cells are packed together to form the chain with one select transistor per chain it is anticipated that the memory array will be considerably denser than a NOR implementation which has a select transistor for each memory cell. It is also possible to use this same basic concept to design NAND architecture memories. In this case the sensing, unlike a linear Match-In-Place™ engine, reads the resistance difference of each individual cell independently during a read cycle while all other cells in the NAND chain are kept idle.

MATRIX FOR MATCH-IN-PLACE™ ENGINES.

The NAND chains could be stacked in parallel to create a bi-dimensional matrix, sharing the same input line for a simultaneous read Match-In-Place. A design example of an 8x3 Match-In-Place check engine is presented

Figure 5. In this implementation three NAND chains of 8 cells each are compared simultaneously. This allows fast comparison of a particular input pattern against multiple stored patterns.

Figure 5. Example of 8x3 Match-In-Place™ Check Engine



This concept of an 8x3 table could be extended to significantly larger look up tables. For example, a 512x10K look up table with this type of architecture could compare a 512-bit input word with 10K stored entries simultaneously in about 15ns. This simple binary CAM design can be converted into a ternary CAM by adding a “don’t care” state, a straight forward extension of the MLU architecture.

The nonvolatile storage capability of the MLU eliminates the need to download the contents of the lookup table after a power-off cycle, and reduces the standby current to zero. Traditional non-MLU based MRAM (including STT) technology, or flash memory technology, although nonvolatile cannot compare at the bit level for such an application. Standard CAM or TCAM cells are currently implemented using SRAM structures manufactured using leading-edge semiconductor technology nodes in order to be able to operate fast enough; however their high cost and volatility limits their field of use. Conventional MRAM technologies simply cannot address this problem as they cannot perform a CAM function at the cell level (because of the lack of a self-reference capability).

The MLU technology with its Thermally-Assisted-Switching self-referenced configuration store and compare at the single cell level in one cycle.

At this point of the paper we already described the difference between the traditional method of “read and compare” and the Match-In-Place™ architecture, and how operations like authentications become faster. Step by step, the buildup of Match-In-Place™ engines was presented starting from the single bit structure, to NAND chains that can provide linear matching, then to the bi-dimensional matrix connecting multiple NAND chain for parallel matching. Match-In-Place™ is a generic architecture that transcends a particular implementation; however it was stated that Crocus’ MLU technology was particularly suited for it. The readers interested by the analysis will find that information in the appendix.

Now, we are going to pick an application example: biometric based authentication, and describe how the Match-In-Place™ architecture can drastically accelerate the authentication cycles, increase the protection from the hacker, and in addition, reduce unit chip costs.

EXAMPLE OF APPLICATION: BIOMETRIC BASED AUTHENTICATION

Biometric patterns like finger prints, finger vein prints, eye iris prints, or facial shapes are commonly used as reference patterns for secure chips. The word size of each biometric pattern is in the 512 bits to 2,000 bits range. One of the difficulties for quick authentication is the natural variation of the shape of the user's biometric print (pressure, angle, shifted patterns, moisture, and cracks).

Traditional secure chips use Digital Signal Processors (DSP) as well as sophisticated algorithms for authentications. The DSPs need multiple clock cycles to compare the stored print with the measurement, calculate the mismatch, and converge to a conclusion, in an iterative way. When the input signal is significantly shifted, the authentication process could be extremely slow. To add complexity, the stored pattern could be encrypted and need decryption before or during authentication.

Leveraging the MLU, and Match-In-Place™ architecture, it is possible to store multiple configurations of biometric patterns in the engine, then to simultaneously compare the measurement with a database of patterns. As needed, 100 to 1000 configurations could be constructed by software from the original biometric print and stored in a matrix similar to the one described in Figure 5. If for example the size of the biometric patterns is 1,000 bits and if we wish to compare the measured pattern with 200 different configurations, the MLU matrix will have only 200K cells. Such a matrix would be tiny to implement in silicon with an estimated area of 0.2mm² (with 130nm CMOS). Direct Match-In-Place of the measured pattern with the 200 different configurations will be done in the MLU simultaneously in one cycle, in about 15ns. Responding to a known attack from hackers, the MLU could be erased as fast as in 15ns while traditional memories need 1000 times more.

A small CMOS state machine that controls the MLU can replace both the DSP and the Crypto-Processor, be very inexpensive, fast, and unbreakable. The state machine verifies that at least one of the stored patterns is similar enough to the input pattern, with a match statistically valid for the authentication. It is expected that the complexity of such a state machine will be considerably less than a DSP and a crypto-processor; this will reduce the die size and costs.

A back to back comparison between both methods is following:

	Read and Compare	Match-In-Place™
Stored Memory	512 bits to 2000 bits	50Kb to 500Kb
T Authentication	Slow: 1 ms to few seconds	Fast: 15 ns
CMOS Logic/Security	Full DSP/Crypto	Small state machine
Erase Mode (defense)	Slow: 1 ms	Fast 15 ns
CMOS Technology	Need HV for memory	Basic LV
Die Size	Big: Typ 5mm ²	Small: Typ 1mm ²

The Match-In-Place architecture brings additional features:

- Very fast Programming of the MLU to load the biometric print (15ns – 100ns) – this operation is often called personalization.
- In addition to the fast erase feature, to respond to an intrusion by a hostile hacker, as a catastrophic defense mode, the MLU can also be destroyed with low power and low voltage to render sensitive data permanently inaccessible.
- During the Match-In-Place cycle of a NAND chain we can directly measure a percentage of the bits that are matching via an analog measurement. This is very practical considering that biometric patterns can not be 100% similar twice: aging, scratches, dirt, and other variations of the subject may prevent perfect matching. We can simply adjust with the state machine the tolerance or the desired level of matching for a valid authentication. It is interesting to notice that the traditional DSP implementation can do the same job, but at the expense of clock cycles, and more CMOS gates (e.g., die size).

Many of the advantages of this particular utilization of Match-In-Place for biometric authentication are generic and are expected to be directly applicable to multiple fields of uses as presented in this last section, and conclusion.

FUTURE WORK AND CONCLUSION

Crocus is currently a small startup corporation, a technology provider, which does not have the bandwidth to fully take advantage of the full potential of the novel architectures on its own. So, we are concentrating our internal efforts in selected areas, and we are establishing strategic partnerships with industry leaders that will be early adopters of the Match-In-Place™ architecture with MLU implementation.

Applications that are been considered include:

- **Data base search**

Data to be loaded in parallel NAND chains to allow large quantities of comparisons at once.

- **IP Packet routing**

Addresses to be stored in the MLU engine for fast matching cycle and routing.

- **IP Packet filtering**

The MLU can filter the packets and search in real time the status for applications such as content deliveries, and handset security.

- **Image detection**

The image to be broken in smaller pieces to be directly compared with stored patterns in data bases.

- **Hardware acceleration**

The MLU will archive complex mathematical addresses, and extract the already calculated results from the data base.

All of these applications will reuse the same basic architecture, and technology.

The size of the MLU arrays will vary from a simple one dimensional NAND chain to large bi-dimensional matrix. It is anticipated that the effort to deploy this architecture will benefit from the transportability of the basic concept.

This work is just the beginning of a bigger task.

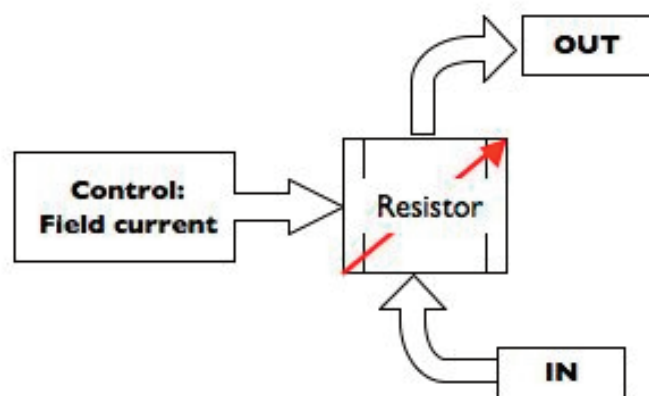
Near term, leveraging the experimental work already done, we are confident in the deployment of the technology for arrays of relatively small sizes, 10Kb to 64Mb, read and write access times in the 15ns to 35ns range, commercial temperature range, and usage of mature CMOS from 150nm to 65nm.

The applications described in this paper represent an extremely large market that currently uses rather mature CMOS technologies. This offers Crocus the ample opportunity to demonstrate the value of the disruptive novel Match-In-Place architecture and its MLU technology without worrying about leading edge lithography. In addition, we do not see obvious show stoppers which prevent the basic concepts described here to scale down to large memory densities, multi Giga-Bytes, with few nanosecond response times, and using sub 20nm CMOS. For example, we characterized successful operation of the magnetic stacks all the way to 10nm. The research work done by Crocus' close research partner Spintech on STT (Spin Torque Transfer) is very synergetic, and could extend the technology even further.

FROM MRAM TO MAGNETIC LOGIC UNITS (MLU)

1. A traditional MRAM cell is based on two ferromagnetic layers: a reference layer and a storage layer. When the magnetizations of the two layers are aligned the resistance of the stack is low, this could be a "0" (or arbitrarily a "1"). When the layers are anti-aligned the resistance is high, this could be a "1" (or vice versa).
2. In Thermally-Assisted-Switching (TAS) MRAM's the storage layers are blocked by an antiferromagnetic material to achieve superior stability in normal operating temperatures. During the programming cycle, the temperature of the cell is momentarily locally raised above the blocking temperature, through resistive heating allowing the value of the memory cell to be changed. At operating temperatures the information stored in TAS-MRAM memory cells is thus not affected by external fields and noise.
3. Self-referenced cells are based on TAS MRAM, the reference layers are not pinned in a permanent way. A current applied in an adjacent field line will create a magnetic field to control the reference layer during the read cycle. A two phase read operation utilizes the natural tendency of an un-driven field line's effect on a selected memory cell to create a momentary reference value that is compared to the cell's value when the field is driven. The stored information is thus read as this field is applied.

MLU cells are designed on the Self-referenced TAS MRAM technology. During logic operations the field lines are acting as controlling gates modulating the resistivity of the magnetic stack. Each MLU cell behaves as a 3 terminal device capable of performing native logical functions. The currents circulating in the field line could be modulated in direction, and intensity.



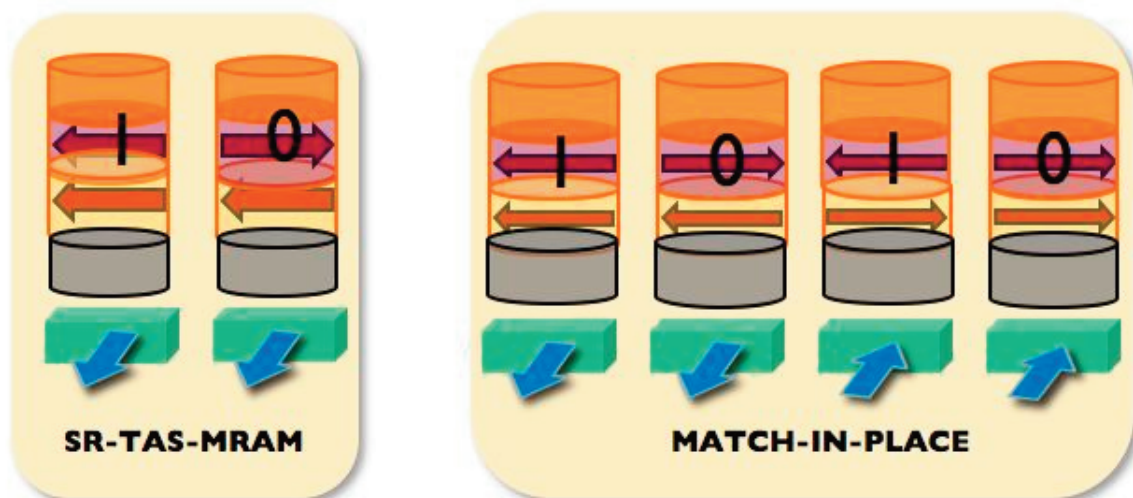
4. MLU cells can be used to design:
 - a. NOR - MRAM. Each self-referenced cell has a select transistor to read the stored bit.
 - b. NAND - MRAM. Multiple cells are connected in series.
 - c. Match-In-Place. The field lines drive the input patterns.
 - d. Table look up. The field lines drive multiple NAND chains in parallel.
 - e. Various logic functions. The field lines are acting as control gates.

More importantly, the tremendous potential of Match-In-Place transcends a particular process implementation. We are initiating research work in cryptography, systems and solutions, and secure architectures with leading institutions like l'Ecole Normale Supérieure de la rue d'Ulm in Paris with Pr. David Naccache. The Match-In-Place architecture is a new tool; how to fully take advantage of it is ahead of us.

APPENDIX : MLU FOR MATCH-IN-PLACE ENGINES

In Match-In-Place engines, as described above, the object is to compare a stored pattern with an input pattern, not always to read the stored pattern. Magnetic Logic Units are suited to implement this architecture (see MLU sidebar).

Figure 6. MRAM cells



A conventional MRAM memory as shown in **Error! Reference source not found.** is designed with all of the sense layers oriented according to the same convention, allowing “0” and “1” to be programmed and read in the same fashion in all cells. In a Match-In-Place implementation the sense layers, rather than following the memory reading convention, are individually aligned by selecting the direction of the current in the field line adjacent to each memory cell to correspond to the input value to be matched

This changes the nature of the cell sensing mechanism: if the sense layer’s orientation matches the orientation of the stored layer (a “0” stored is compared with a “0” input, or a “1” stored is compared with a “1” input) then the cell indicates a match. If, conversely, the two do not match (a “0” stored is compared with a “1” input, or a “1” stored is compared with a “0” input) then the cell indicates a mismatch.

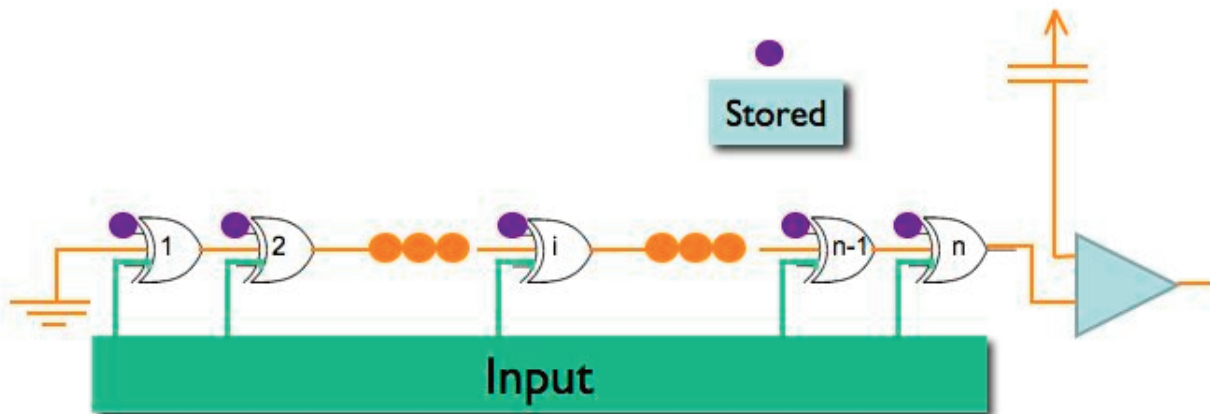
This binary match/no-match implementation with nonvolatile retention of the stored information makes the MLU a very attractive implementation for the Match-In-Place single cell engine described in Match-In-Place. The single cell has to perform the Match-In-Place™ quickly in a native way without external circuitry.

MLU cells perform as 3 terminal logic devices. The resistive stack is modulated by the direction of the current flowing through the field line which is set by the input bit to be matched. A high or low resistance at the output indicates whether the input bit matches the stored bit or not, respectively.

A single MLU technology can produce both traditional MRAM cells and Match-In-Place cells.

Multiple MLU cells (Figure 7) could be tied in series by local interconnect to form a NAND chain. The stored bits $\{S_1, S_2, \dots, S_i, \dots, S_{n-1}, S_n\}$ programmed into the storage layers remain stable at normal operating temperatures. The input bits $\{I_1, I_2, \dots, I_i, \dots, I_{n-1}, I_n\}$ are presented simultaneously to the sense layers through n field lines.

Figure 7. MLU for NAND Chains



When no current flows through the field lines the NAND chains naturally orient the sense layers in the direction that creates a match: At rest the sense layers are always oriented by their proximity to the storage layers regardless of whether a “1” or a “0” was stored. The total resistance of the chain at rest will always be n times the value of the resistance of each cell at rest.

During the active matching mode, all field lines carry currents that orient each field line’s associated sense layer. If all sense layers in a chain are oriented in a way that matches all the bits in the stored word, the total resistance of that chain will be exactly the same as it was when measured at rest, but any mismatched chains will have lower resistance. The differential sensing components at the end of each chain simultaneously measure the level of matching for all stored words.

Crocus' MLU technology directly implements bit-to-bit matching within a NAND-chain without any need for other logic or memory components. Unlike other memory cells, the MLU cell structure conveniently supports "0" and "1" state storage together with the comparison with an external bit function. NAND-chain configuration supports word-level matching.

CHARACTERIZATION OF MLU CELLS

The MLU cells were characterized for applications such as Match-In-Place engines on 0.5Mbit memory wafer lots processed in the last 18 months. The test chip was designed on 130nm CMOS with 200nm magnetic dots. Tower-Jazz Semiconductor manufactured the 200mm CMOS wafers, and the multi-layer CVD deposition of the metallic and magnetic layers was processed by Crocus Technology in Grenoble.

The cell resistance of the "0"s and the "1"s can be adjusted by changing the thickness of the thin magnetic tunnel oxide. With a magnesium oxide 1.4nm thick ($R_a=30 \text{ } \Omega \cdot \text{sq}$ for MgO resistivity), the resistance is around 1k Ω for one state, and 2k Ω for the other state.

The TAS programming cycle was characterized with programming pulse widths varying from 14ns to 77ns. The voltage across the stack necessary to raise the temperature of each cell above blocking temperature is about 0.9V with current of 0.5mA. This current is cut in half when the technology migrates to each new node. Satisfactory programming was observed within the full pulse width range.

Read margins were characterized against read pulse (address access t_{ACC} and chip enable access t_{CE}) width for both "0" and "1". Read margins for both "0"s and "1"s are excellent at pulse lengths of 13ns or longer, and degrade below 11ns.

We also characterized the test chip under permanent magnetic field. Within an operating temperature range of 0-100°C the TAS MRAM is insensitive to parasitic magnetic fields as high as 100Oe. Low cost plastic packages without shielding will be perfectly acceptable for these chips.

Finally accelerated write cycling of the magnetic dots under stress was performed at room temperature. No failures were observed at voltages up to 2.2Volts in the magnetic stack; under normal operation the supply voltage should remain below 1.6Volts. As a result we anticipate that in a TAS-MRAM Match-In-Place application the magnetic cells are capable of withstanding a number of write/erase cycles of 10^{12} under normal operating life of the chip in the 0 to 105°C range.

SYSTEM-ON-CHIP (SOC)

The MLU as characterized above is to be implemented into full SoC.

- **Compatibility with CMOS wafers processing**

Crocus' MLU is compatible with any basic CMOS technologies. The wafer fabrication of MLU adds 3 lithographic steps just before passivation and final processing. Crocus' technology does not require

special packaging. CMOS SOC could be generated with mainstream design technologies.

- **Can replace all embedded memories**

Crocus MLU is non-volatile like FLASH, fast and high endurance like an SRAM, and operates without the need for high voltage transistors.

Sub 3.0 volt writing versus 15-20 volt erase/program voltages required by embedded flash:

- Elimination of complex high-voltage processes
- Simplified periphery, advantageous at low density and NAND architecture

*Non-volatile, fast reads and writes, with unlimited program/erase cycles:
MLU can replace embedded SRAM, with a die area of only 20-30% that of the
same-sized SRAM*

Crocus Technology Sunnyvale

710 Lakeway Drive, Suite 100

Sunnyvale, CA 94085 USA

Phone: +1.408.732.0000

Fax: +1.408.732.8250

Crocus Technology Grenoble

5 Place Robert Schuman

38025 Grenoble Cedex, France

Phone: +33 (0) 4.38.12.10.70

Fax: +33 (0)4.76.21.45.18