Electronic Engineering Times

http://www.eetasia.com/ART_8800429807_499486_eccefe11200608_no.HTM

MRAM ushers in era of new memory tech

Posted : 16 Aug 2006

For the past 11 years, through periods of hyperbole and doubt, Saied Tehrani and the MRAM development team he directs at Freescale Semiconductor Inc.'s Arizona facility have worked on a new type of IC memory—one using magnetic resistance instead of charge storage.

Today, after two years of sampling to some 40 companies, Freescale will begin selling 4Mbit magnetoresistive RAMs, priced at $25 each in small quantities and at "negotiable" prices for high volumes.

"You can put in an order, and we will ship it, starting today," Tehrani said.

While MRAMs have their own cost and power consumption challenges, a handful of companies—notably Renesas Technology, Samsung Electronics, Sony, and partners Toshiba and NEC Electronics—have joined Freescale in the MRAM race as traditional, charge-storage memories run out of steam. SRAM faces current-leakage problems, DRAM trenches now look like long noodles, and embedded flash confronts reliability concerns and charge-pumping costs. Particularly for embedded applications, where memory bits will far outnumber logic transistors going forward, the need for a new form of non-volatile memory has raised interest in MRAM.

"Freescale's entry says we are officially in the era of new memory technology," said Bob Merritt, analyst for emerging memory technology at Semico Research Corp. "By being one of the first guys out there with a viable manufacturing strategy, Freescale will give a boost to the other companies working on MRAM. And once they see MRAMs succeed, it could help the companies developing other new memory types," such as phase-change, metallic and floating-body memories.

The non-volatility and endurance attributes of MRAM could open up new applications in mobile systems, Merritt added.

Freescale's accomplishment is "the most significant memory introduction in a decade," said Will Strauss, president of market research firm Forward Concepts Co. He noted that other new memory advances have all been derivatives of the big three: DRAMs, SRAMs and flash (itself a derivative of EEPROM).

"Everything else is a variant. This is a totally new technology," Strauss said.

Cars are target

Manufactured at Freescale's 200mm Chandler fab, the discrete 4Mbit MRAMs are just a prologue.

Freescale, Tehrani said, "has no intention of becoming a commodity MRAM supplier." Rather, over the next 18 months to two years, Freescale plans to bring MRAM onto its automotive-use MCUs, replacing flash. The latter technology suffers from wear-out mechanisms that tend to worry automakers' reliability engineers. MRAMs, unlike flash or ferroelectric RAM, can run through an effectively infinite (1,015) number of cycles.

Freescale's MRAM cell is a magnetic tunnel junction (MTJ), with a thin (15Å) insulating layer separating two thicker magnetic layers. When a voltage is applied, the polarization of one of the two magnetic layers reverses, and current tunnels through the insulation layer. The state of the memory is non-volatile, endurance is practically infinite and operation is fairly fast (read/write times are about 35ns). Unlike flash, which has very fast reads and relatively slow writes, MRAM is symmetrical in its read and write times, which means that it is fairly simple to switch from embedded SRAM, which loses data when power is turned off, to MRAM.

Nonetheless, the approach has challenges including cell size and write power. To ensure that bits are stored with non-volatility, electrons must overcome an energy barrier. If that barrier is too low, bit cells can be disturbed too easily—too high, and write power becomes excessively high.

The 4Mbit MRAM currently available requires a write current of 150mA, which Tehrani said "is not very different from standalone SRAMs that you find today. Flash takes milliseconds of write time, while our MRAM can do a write in 35ns. And the total energy that MRAM extracts from the battery is orders of magnitude lower than flash."

The jury is still out on costs. Freescale uses low-cost sputtering and etching techniques, avoiding expensive tools such as molecular-beam epitaxy or atomic-layer deposition.

Semico's Merritt noted that MRAMs are compatible with CMOS logic processes: The memory tunnel junction is built in the back-end, after the transistor is created. In Freescale's case, only three levels of metal interconnect deposition must be altered in the back-end of the line to create the MTJ, leaving the two or three other metal levels in the MCU undisturbed.

Freescale now uses six extra mask layers to embed MRAM with logic, an approach that compares favorably with the six to 10 extra masks needed to embed flash or DRAM in a CMOS process, Tehrani said. The MRAM cell size—1.26µm for the 0.18µm process used for the 4Mbit discrete MRAM, shrinking to 0.29µm for Freescale's 90nm process—is larger than embedded DRAM. It is roughly the same as embedded flash, but it's smaller than a six-transistor SRAM, which consumes silicon area at the rate of $1\mu m^2$ per bit at the 90nm node, he said.

To some extent, costs in the memory business depend on volumes, which provide manufacturing and yield learning. Tehrani said Freescale expects to start out making MRAMs at the rate of only a few hundred 200mm wafers per month. The company is targeting a range of customers who now use battery-backed SRAM for high-security network systems, for storing file-structure data in disk-drive networks, and for other applications in which data must be written and retrieved quickly. While SRAMs can be made non-volatile by attaching them to large batteries, those subsystems can fail if batteries leak, and they can be more expensive than MRAM-based solutions, he said.

For customers who do a lot of data writes to memory, such as in RAID storage systems, MRAM appeals because it has "no wear-out mechanism," Tehrani said. "Wear-out is especially bad with flash and, at some level, with ferroelectric memories."

Strauss, who specializes in signal-processing markets, said many portable systems use flash to store code and data, which is then written into SRAM or DRAM during operation. With MRAM replacing those memory types, he said, "you can execute directly out of MRAM at a faster rate than with some types of flash. And power consumption is less because, with flash, you need large charge pumps, which are hard to integrate with CMOS."

While many companies have learned to tolerate the failure mechanisms in flash by developing wear-leveling techniques, MRAM can be more reliable than flash, Strauss said. "MRAM could open new market opportunities, including use in high-performance DSPs."

Reliability at high temp
The market for battery-backed SRAM is fairly small at less than $200 million annually, Tehrani said. But it provides a good ramp to the larger plum: the multibillion-dollar market for automotive MCUs.

The 4Mbit MRAMs now being sold by Freescale are guaranteed for the commercial range of temperatures, roughly 0°C to 70°C. The next step for the company is to offer chips in the industrial temperature range, where they must be guaranteed to operate continually at temperatures ranging from -40°C to 105°C.

Automotive temperature "profiles" demand that engine controllers, for example, be able to withstand very high temperatures for certain periods over the lifetime of the chip. That upper range is moving upward, from 150°C now to a more demanding profile that sets 170°C as the upper range for a limited number of hours. Freescale thus has more work to do, adjusting the design and making the process more robust against current drifts at high temperatures, if it hopes to hit the profile goals set by automakers.

At last year's International Reliability Physics Symposium, Freescale published its early reliability data, showing good 10-year data retention levels. "We are very pleased with the levels of reliability we have been able to see so far, and we continue to make improvements as we drive to industrial and automotive temperature ranges," Tehrani said.

Paul Grimme, senior VP and general manager of the company's transportation and standard products group, said MRAM also can be combined with magnetic sensors—a mix that he said would serve "a fairly large market"—as well as embedded on analog components that require on-board memory. The company is investigating where it might make sense to integrate passives with MRAM.

Longer term, Freescale is researching the use of magnesium oxide for the tunneling layer, replacing the aluminum oxide now in use. That would reduce the power required to read a bit, though it would not improve the write-power consumption.

"MgO gives you a much bigger signal between the two states. That makes reading much faster because the signal size grows," Tehrani said. "We have not made a firm decision; we still have to make sure reliability is acceptable with MgO. But we believe that is just a matter of time."

Spin torque
Further out on the horizon is spin-torque technology, which Tehrani said "would definitely help reduce the write current."

The MTJs in use today employ metal lines on the top and bottom to generate current locally. In the case of spin-torque MTJs, the current flows directly through the stack of magnetic layers to switch the bit. While spin torque is widely seen as where MRAM technology is headed, there are reliability concerns.

"There is a lot of promise there to go to the next generation and scale the memory going forward," Tehrani said, "but there are still a number of areas where we need to make considerable improvement to make spin torque viable from a production perspective. There is current going through the tunnel barrier, so we have to make sure the tunnel junction is reliable at these higher currents."

- **David Lammers**
EE Times