Can MRAM (finally) be a factor?

Jean-Pierre Nozieres
Spintec and *eVaderis, Grenoble, France*

Despite all the fuss about the "more-than-Moore" approach, introduction of new technologies in semiconductor fabs is always a challenge. All is not to blame on conservative operations executives, though, and technologists also carry their share of responsibilities. With never-ending improvements from researchers, who always have the "next-big-thing-that-will-make-it-work-better", displacing existing mature technologies, even scaled beyond reason, is a difficult decision to make. At the same time, an early adopter may displace the market if he takes a timely risk and becomes a leader, while competition is playing catch-up. It is really a game of cat and mouse between the risk of being too early on the technology or too late on the market … and this may be exactly where MRAM stands today.

Since its inception in the late 1990's and despite numerous promising announcements, MRAM has yet to live up to expectations. The recent advent of spin transfer torque (STT), however, has shed a new light on MRAM, with the promise of much improved performance and greater scalability. Semiconductor giants, IDMs, fabless manufacturers and, now, pure play foundries are entering the game with promises of available technologies and products … by the end of the year (as usual).

As exciting a technology as it is, however, STT-MRAM is by no means a champion: It is slower than SRAM, bigger and more expensive than Flash, and does not reach the truly infinite endurance of DRAM and SRAM. What makes STT-MRAM so unique in this landscape is its unique combination of non-volatility, speed, endurance and low active power.

For the past few years, the focus has been on cutting-edge technology, with players shooting for sub-20 nm feature sizes in densely packed, multi-Gb arrays, with DRAM replacement in sight. While major strides have been made, it is still an uphill battle, against a constantly moving target, with the focus now on sub-10 nm technology. All major DRAM players (Samsung, Micron and Hynix) have on-going STT-MRAM development programs, yet the jury is still out whether real products will ever be launched.

At the same time, the same core technology at more relaxed technology nodes may be the ideal embedded memory solution, in particular where power consumption (or battery lifetime) is key. In today's SoC, memory power now accounts for a large fraction of the total dissipated power. The best solution is to embed large amounts of *nonvolatile* memory within the logic chip itself, in a so-called "logic-in-memory" architecture: No more static loss (thanks to non-volatility) and minimized dynamic loss, provided said memory exhibits low enough active power and fast enough read/write … which are precisely the STT-MRAM's attributes! As the icing on the cake, STT-MRAM is compatible with any form factor, so the capacity to turn off power from the processor down to the individual logic gate almost on-demand paves the way for ultra-low power "instant-on/normally-off" computing.

All this of course has a price. Beyond the added processing costs, MRAM demands a total rebuilt of the circuit and system level architectures.

At the same time, because we are scientists, we cannot resist the excitement of "the next one" and the burgeoning SOT (spin orbit torque) and OST (orthogonal spin torque) MRAM technologies – what a lack of imagination in the acronyms! – both promise SRAM-like, ultrafast switching, thereby paving the way for directly-attached microprocessor level-1 cache.

Not all routes pursued in MRAM research may lead to the Promised Land, however this may just be the right time, with STT-MRAM technology nearing maturity and foundries finally entering the game. One way or another, MRAM will bite into the semiconductor memory market and never look back.