CROSSING the CHASM

MRAM is finding ways across the memory chasm and into mainstream applications.By **Chris Edwards**.

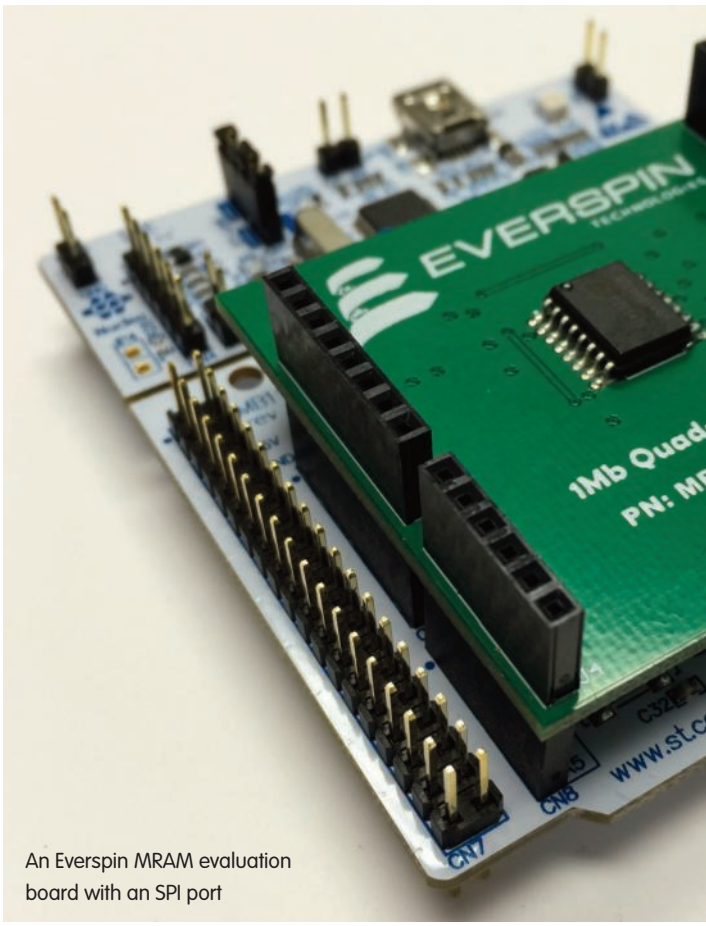

An Everspin MRAM evaluation board with an SPI port

When Freescale started work on magnetic random access memory (MRAM) two decades ago, it looked as though it could provide a fast, low power memory that does not need a constant flow of current to store data. With a bit cell that looked to be competitive with DRAM, but with better storage behaviour than flash, MRAM offered the potential to be the ultimate memory.

While we are still waiting for MRAM to go mainstream, the technology looks as though it is finding a way across the yawning chasm that has prevented many promising memory technologies from gaining commercial success.

The drawback has often been the problem of playing catch-up with DRAM and flash in terms of cost and density. MRAM makers have found a way to chip away at the massive memory market that avoids having to compete directly on price per bit or density.

Hamid Haidari, vice president of sales at MRAM specialist Avalanche Technology, said: "All manufacturers have internal programmes for MRAM. Their primary focus is to replace DRAM because they see challenges in further scaling. But we see an opportunity before that – there is significant demand for lower density DRAM."

Freescale spun out its MRAM interests around 10 years ago into Everspin Technologies, which became the first to sell MRAMs commercially and has also licensed the technology to Cobham and Honeywell, where it brings advantages in resistance to radiation-borne upsets.

"We've been shipping MRAM products since about 2008," says Phillip LoPresti, Everspin's president and CEO. "We have shipped 60 to 65million units that are either standalone or embedded in a processor core. But until now, we've mainly been addressing persistent SRAM needs."

LoPresti says the existing markets span industrial, transportation, medical and server applications. "The core application is fundamentally datalogging. We see a lot of growth moving forward because of the increased interest in automated driving. It requires a form of black-box recorder to keep track of what tasks were completed last before a crash. Insurance companies want to know whether the car failed or if it was driver error."

The key advantage of MRAM over flash in datalogging is that writes are much faster, consume a lot less energy and can be performed on a per-bit rather than a per-block basis. While the power and time taken per bit written is higher than that of SRAM, the static dissipation is much lower.

Vendors expect that a shift to spin transfer torque (STT) – a form of MRAM first explored in the early 2000s – will enable DRAM, rather than SRAM applications, to be targeted. STT MRAM reduces the potential for a write to one bit cell to upset its closest neighbours. That means bit cells can be packed closer together. This, accompanied by a migration towards 28nm manufacture, is bringing parts with hundreds of megabits and more to the market.

According to LoPresti, sales of the datalogging MRAMs go into storage devices, where the memory is used to hold important filesystem metadata. The next step is to push MRAM into the write caches that help maximise throughput in flash – and disk based storage subsystems for servers.

"It's really important that you have a low latency cache that's protected in the event of a power failure. You are collecting writes and holding them before writing through to the media," says LoPresti. "In an enterprise storage subsystem, it's vital that data in flight gets to the media, regardless of the kind of system cache. Today, they use DRAM that backs up to NAND flash or the storage media. But they have to retain the power during a failure. So they use either supercapacitors or a battery tethered to the board."

Haidari adds: "That all takes real estate and adds cost. Our approach is attractive from a cost perspective, as

well as providing a simpler system design. We are not going to compete with DRAM at the level of bit cost, but if you look at the system level, we become more attractive."

handle compute tasks within the storage array. Everspin started sampling a 256Mbit memory in the Spring and expects to have a 1Gbit part based on a 28nm process ready for trials by the end of the year.

"With the 256Mbit and 1Gbit densities, we can cover a wide slice of these write cache applications," says LoPresti. "We see a trend of memory moving closer to the compute side and see that as a market that's addressable by MRAM. It is, maybe, a market that starts getting active with 1Gbit parts, but which becomes more viable with 4Gbit devices. High-performance computer and server companies are very interested in these non-volatile DIMMs that are popping up."

LoPresti sees the technology scaling readily to 4Gbit and 8Gbit densities. "We still have some room to catch up with DRAM. We design today onto logic design rules, but DRAMs are on custom memory processes."

Scaling and write power demands may lead to a switch away from STT MRAM in the future. Technologies such as spin-orbit torque or voltage-controlled magnetic anisotropy (VCMA) do not need high currents to switch

"We see a trend of memory moving closer to the compute side and see that as a market that's addressable by MRAM."
Phillip LoPresti

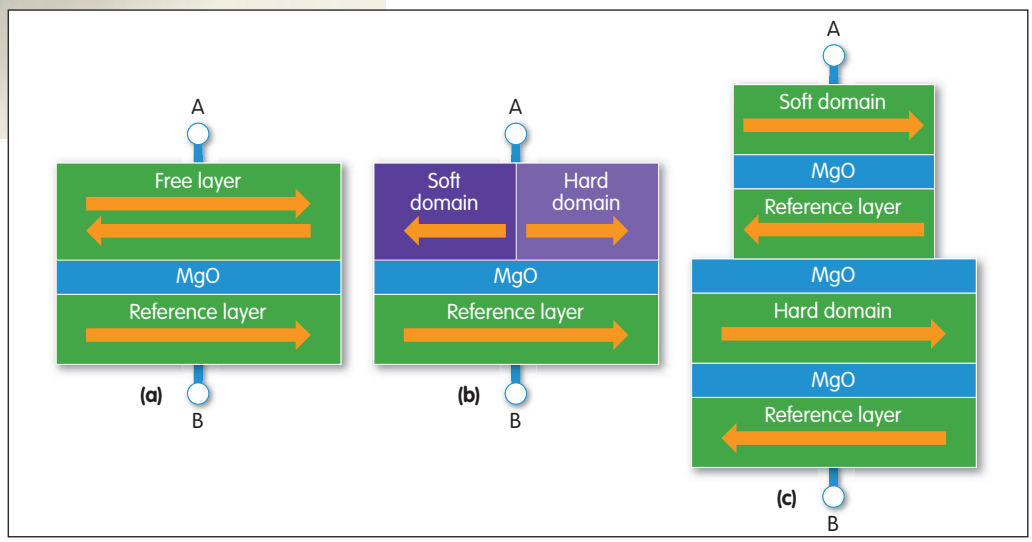

Below: An MRAM cell, left, alongside a parallel MLC cell, centre, and a serial MLC design that stacks two 1bit cells

torque's need for an extra write line, which will take up space in the memory array.

"A lot of optimisation has yet to occur in MRAM," says LoPresti. "Most other technologies are in the early stages of development and hold the promise of lower switching currents, which would also allow a smaller bit cell. Maybe there will be some breakthrough but, for five years at least, it seems there is enough room to expand spin-torque to capture a bigger part of the non-volatile memory needs of servers."

A smaller bit cell may not be entirely necessary for MRAM. The access and manufacturing techniques for MRAM lend themselves to 3D structures, LoPresti says. Controlling the change in resistance more precisely may then enable multilevel cells.

In parallel with the market for discrete MRAM, embedded applications are emerging. "Wearables and the IoT in general represent a huge potential market," Haidari says.

By replacing onchip SRAM with MRAM, it becomes possible to design MCUs where power can be removed from all but a few vital circuits without having to copy cache contents to flash or maintain a trickle charge. MRAM may also prove a better option for non-volatile memory on MCUs built on sub 65nm processes as flash capable processes lag far behind mainstream CMOS. And because MRAM is added in the metal layers, rather than in the base silicon, it should be easier to port to new processes.

MCU makers are cautious, aware that changes to memory technology are risky. Geoff Lees, general manager of NXP's MCU business, sees MRAM as a potential replacement for flash. One way to test the waters for the technology, he notes, is to add a comparatively small amount of MRAM and use that for datalogging before attempting to replace larger flash arrays.

The result is that MRAM can use specialised applications to build a big enough base to finally take on the other technologies and, perhaps, avoid the memory chasm completely.

The availability of higher-density MRAM is letting computer makers experiment with novel storage techniques. IBM has used Everspin's 64Mbit memories in its experimental ConTutto platform, where FPGAs

the spin state and are likely to be more energy-efficient – assuming they translate into volume products. These technologies can introduce other drawbacks, such as VCMA's need for a read before a write and spin-orbit