



## MRAM Overcomes SRAM, DRAM, and Flash Limitations

By *Andy Walker* January 18, 2018

We generate data at an accelerating rate, and we need more space to store it. While memory was once scarce and expensive – meaning every byte was carefully considered, today the mantra is, “Memory is cheap!” And so we blithely throw terabytes of data into memories with scarcely a thought.

That might suggest that we have all the memory solutions we need, but today’s primary memory technologies – as well as they have served us – all have drawbacks. They get the job done, but they leave us wanting more.

- SRAM is used for limited, fast working memory – but it’s expensive and power-hungry
- DRAM is used for high volumes of working memory – but it’s complex
- Flash memory is used for non-volatile storage – but it’s complex and power-hungry

We use these memories both in dedicated chips that can leverage custom processes and as memory blocks embedded on systems-on-chip (SoCs) using standard CMOS processes. None of these three works equally well in both embedded and stand-alone implementations.

Because these memory types have limitations, there has been an ongoing search for new approaches to storage. Of the relative newcomers, magnetoresistive RAM, or MRAM – already in production – does the best job closing the gaps left by the incumbents.

## The Incumbents – Good and Bad

While each memory type has its niche, a closer look reveals where they fall short.

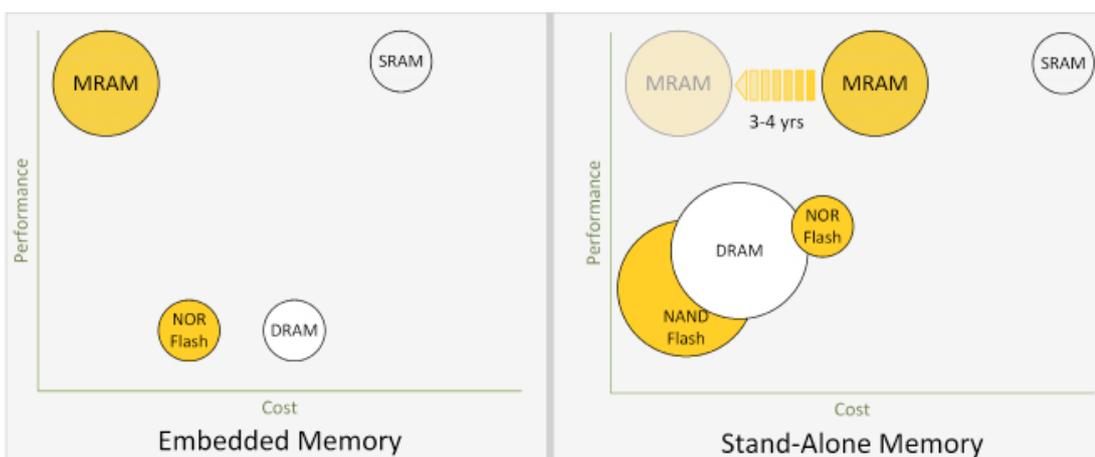
- DRAM provides bulk working memory during software execution. It can store the code being executed or data being generated and consumed during execution.
  - ◇ **Strengths:** it’s inexpensive and of moderate complexity. The memory cell consists of a capacitor that holds the data; it can be built in a way that allows dense packing of cells on dedicated memory chips.
  - ◇ **Weaknesses:** it’s relatively slow and power-hungry; memory cells are subject to corruption by radiation; it’s very hard to embed onto CMOS SoCs; data disappears when power is removed; and, even when power is applied, the capacitors leak data, meaning that cells must be constantly refreshed. Refresh is hidden from users, but it represents extra complexity. In addition, memory modules incorporating newer technologies require exquisitely sensitive timing to strobe signals precisely as they “fly by” the memory chips.
- SRAM provides much faster working memory – at a cost. So it’s primarily used for embedded working data storage or for caching DRAM data to improve performance. SRAM can consume more than half the die area of modern microprocessors and high-performance computing SoCs.
  - ◇ **Strengths:** speed, pure and simple. It’s also easy to embed, since it is built from CMOS transistors, and it has a simple architecture that allows byte addressability.

- ◇ **Weaknesses:** It's expensive. The bit cell consists of six or eight transistors that take an enormous amount of space as compared to other technologies. It's for this reason that it is used as sparingly as possible. It also consumes more power than any other memory; it's volatile; and bits are subject to corruption by radiation.
  - Flash memory will be most familiar in the form of thumb drives and solid-state drives (SSDs). It's used for longer-term storage that can be disconnected from power. Sophisticated engineering has made possible bit cells that can hold more than one bit of data and a 3D structure for dedicated flash chips. There are two configurations: NOR, typically used for code store, and NAND, used for bulk storage and for higher code-store capacities. NAND overwhelmingly predominates. The bit cell consists of floating gates or charge traps for holding electrons.
    - ◇ **Strengths:** inexpensive and available in high capacities. The contents are also impervious to radiation.
    - ◇ **Weaknesses:** First, it's slow, and, while data can survive power-off, it can do so for only a limited amount of time (typically 5-10 years). Its high capacity comes at the expense of high complexity: data cannot be read byte-at-a-time; cells must be erased before being rewritten, and erasure must be done on entire blocks, not individual bytes. And complex wear-leveling algorithms are required to maximize memory life. Finally, it's not easy to embed on CMOS chips. Doing so for NOR requires 9-12 additional process steps. Embedding NAND is not even attempted whether it be 2D or 3D.

So, in summary, we have simple, fast memories that are expensive and power-hungry; cheaper working memories that are slow, moderately complex, and almost impossible to embed; and non-volatile memories that are slow and very complex. What's needed is a simple memory that goes light on the wallet, light on the battery, can hold its contents when power is removed, and can readily be implemented either as a stand-alone chip or as an embedded memory block.

## MRAM Steps Up

The memories discussed so far all rely on storing electrical charge in some fashion. Magnetoresistive memory, by contrast, uses a magnetic layer to store a bit. A single large version of such a cell has been in use for a long time – in the read heads of hard disk drives. That version, however, is not well suited to densely packed arrays of cells, and so further work has yielded more compact cells suitable both for stand-alone chips and for embedding.



Circle diameter indicates capacity potential  
 ● Gold fill indicates rad-hard  
 ○ Cross-hatched indicates byte-addressable

The bit cell consists of two magnetic layers. One has a fixed polarity; the other can have its polarity set either parallel or anti-parallel to the fixed layer. If both layers are polarized in the same direction, then the current tunneling across a barrier will flow relatively easily. If the layers are in opposite directions, then the tunneling current will have a harder time flowing. The cell value is read by measuring that current.

MRAM technology has evolved for denser embedding in SoCs. The magnetic layers, however, require CMOS-friendly materials that aren't used for standard CMOS. Those extra steps can increase the cost of the finished wafer. But because of the small cell size, an MRAM-based SoC will be smaller than an equivalent SRAM-based SoC would be.

For dice with a substantial amount of embedded SRAM, the benefit of the smaller die can easily exceed the impact of the extra wafer cost, resulting in a net less-expensive chip. Future work on multi-level cells can reduce the die size even more. So for caching and local memory, MRAM competes well.

It will also beat embedded flash, since many of the tricks used to lower the cost of stand-alone flash chips can't be used in embedded applications.

For stand-alone MRAM chips, the process can be further optimized, since no compromises are needed on behalf of processors or other logic blocks. Such MRAM chips will beat DRAM chips on cost in 3-4 years.

Beyond cost, there are several other key advantages of MRAM technology.

- MRAM arrays are byte-addressable, in a manner similar to SRAM. That makes them far simpler to access than flash, and, for non-streaming data, simpler also than DRAM. In addition to reducing system complexity, this feature improves both performance and power.
- Non-volatility is an important differentiator for MRAM. Given that MRAM can be used as a working memory that keeps its data when powered off, it may be possible to reduce system cost even more by eliminating – or at least reducing – the amount of external flash memory required.
- MRAM access speed comes close to what SRAM can provide. Reading and writing are dramatically faster than flash memory, and they're far less complex.
- MRAM cells have no leakage when in standby, making them a much lower-power option than SRAM. They also compete well with DRAM, since reads are non-destructive and no refresh is needed. And the simpler read and write mechanisms dramatically reduce power – write power in particular.
- Finally, the magnetic layers can't be perturbed by radiation. MRAM joins flash in its ability to maintain rad-hard storage.

While all four memory technologies – SRAM, DRAM, flash, and MRAM – can provide benefits in some applications, MRAM can excel in many applications that today require combinations of SRAM with DRAM or flash. The result will be simpler, less-expensive, and lower-power systems.

# A Promising New Embedded Memory

In summary, then, we've long been able to leverage benefits from DRAM, SRAM, and flash memories, but those benefits have come with significant limitations. It's been a challenge to find a new memory technology that could compete on the benefits while addressing the limitations. MRAM is now a leading contender for full commercialization of such a memory.

You will soon see MRAM replacing SRAM and embedded flash in their many guises. Embedded DRAM is not so common – largely because it's difficult, so any such applications are particularly likely to embrace MRAM technology. But a few years from now, you should expect to see dedicated MRAM chips that are cheaper than dedicated DRAMs. Existing technologies won't disappear overnight – perhaps not ever. But much of what they do will be better served by MRAM in the future.

Andy Walker, Ph.D. is VP of Products at Spin Transfer Technologies. A well-known semiconductor leader with expertise in 3-D Flash memory, Walker holds over 40 U.S. patents and has authored over 30 academic articles. Prior to joining Spin Transfer Technologies, Walker founded Schiltron Corporation and worked for Cypress Semiconductor, Artisan Components, Matrix Semiconductor and Philips Research Laboratories. Walker holds a B.S. in physics from the University of Dundee, Scotland and a Ph.D. from the Technical University of Eindhoven, The Netherlands.