New non-volatile memory technologies: an engineer's comparison of nvSRAM, MRAM and F-RAM

The technologies most commonly used for storing user data and configuration information when a system loses power – Flash and EEPROM – are now very old, and very well understood. Although these non-volatile memories have their flaws, they will remain in widespread use because they are very cheap, and are effective enough for many applications.

However, the market has recently welcomed new technologies that get much closer to being the perfect non-volatile memory, in that they are completely reliable, instantly programmable, comparable in price to a standard memory, powered by a single supply voltage, and support an infinite number of erase/write cycles. It is a tall order, but the raw performance data for nvSRAM, Magneto-resistive RAM (MRAM) and Ferroelectric Random Access Memory (F-RAM or FeRAM) shows that the latest entrants to the market come remarkably close. Colin Weaving, Technical Director, Future Electronics (EMEA) explains.

Technology	Programming speed	Cost/bit	Data retention	Write erase cycles	Industrial grade
MRAM	35ns	1	20 years	infinite	Y
Non-volatile RAM	15ms	1.7	20 years	200,000	Y
F-RAM	110ns	1.2	>10 years	10 exp 14	Y

Table 1: Comparison chart.

As always, the raw data, (see Table 1) does not tell the whole story. The manufacturers of each of these new technologies have taken very different approaches to solving the same problem. The question is, are there any clear ways for the circuit designer to distinguish between them?

A look under the hood

In order to understand the different properties of each of these three memory types, it is helpful to understand the ways in which they operate.

nvSRAM, a product marketed by Cypress Semiconductor, is a hybrid device containing both a standard Static RAM (SRAM) and a non-volatile (EEPROM) memory (see Figure 1).

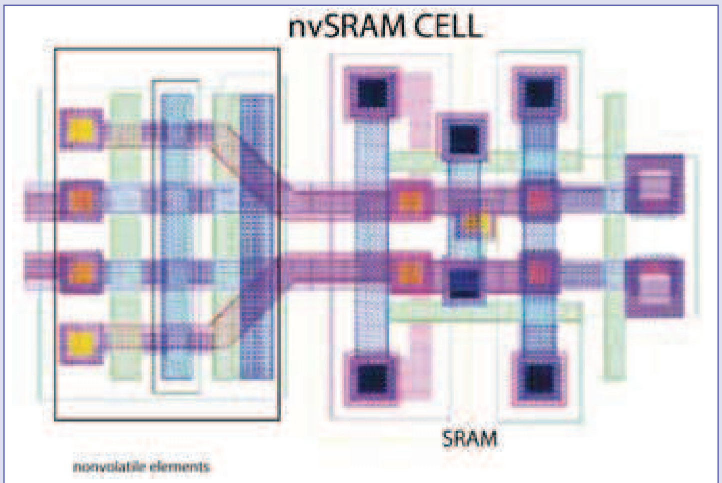

Fig. 1: nvSRAM contains Static RAM and EEPROM memory.

READ THIS TO FIND OUT ABOUT:

- The latest non-volatile memory solutions.
- A technical examination of each technology.
- The key differentiators depending on application and requirements.

It looks to the user like an SRAM with a non-volatile mirror for each bit. In normal operation, data is saved to and retrieved from the SRAM. When power fails, the data is automatically dumped into the non-volatile part of the array. Power for this process is derived from a small capacitor connected to the device, which holds it active while bits are programmed to the EEPROM. Data can be retrieved from the non-volatile array at any time and will overwrite the contents of the RAM.

nvSRAM does not use standard floating-gate EEPROM technology. Instead, it is based on Cypress's Silicon Oxide Nitride Oxide Semiconductor (SONOS) process. This uses Fowler-Nordheim, field-emission tunnelling techniques to store charge, and hence data, on a nitride layer. More on Fowler-Nordheim effects can be found at http://ece-ww.colorado.edu/~bart/book/msfield.htm.

SONOS devices can be manufactured on standard silicon processing equipment, which helps keep costs low. SONOS technology is also scalable to large array sizes: devices up to 4Mbits were available as of June 2008, but future process improvements will enable higher densities to be made available.

While nvSRAM is an innovative way of manufacturing and packaging old memory types, its two rivals are based on far more esoteric technology.

Standard non-volatile technologies use quantum mechanical tunnelling to store data on insulating layers; in F-RAM, of which the leading manufacturer is Ramtron, data is stored in the state of a ferro-electric crystal. The atom in the centre of the crystal has two equal and stable low-energy states (see Figure 2). These states determine the position of the atom. An alternating electric field applied across the crystal will cause the atom to move from the top of the crystal to the bottom and back again. Each transition will produce charge, (Q_S). The position of the atom can be read, thus providing a non-volatile data storage mechanism.

Data is handled as words rather than a single bit at a time, which means that the device looks like a normal CMOS RAM to the user. As each cell may require only a single transistor and capacitor, F-RAM can achieve high densities, reducing manufacturing cost.

There are two important consequences of the physical structure of an F-RAM device. Firstly, the voltages used are extremely low. The absence of the high programming voltages required in many non-volatile memories helps F-RAM offer extended life. F-RAM devices can be rated for 1014 erase operations – 108 times more than a typical Flash device can offer although, since current architectures use a 'destructive read', an endurance cycle is used on reads as well as writes. Secondly, the physical structure of an F-RAM means that programming operations can be accomplished in nanoseconds – far faster than the milliseconds that are required by Flash.

MRAM adopts yet another approach to achieving non-volatile memory storage. Here, magnetic data elements store the information in memory cells. The structure is made up of two insulated ferro-magnetic poles separated by an insulating layer.

Freescale Semiconductor, which has pioneered the commercial implementation of MRAM, uses a Magnetic Tunnel Junction (MTJ) as the storage element. The fixed layer of the MTJ is permanently polarised in one direction. The other layer is free and can be polarised in either direction (see Figure 3). The resistance across the MTJ depends on whether the alignment of the layers is in the opposite or the same direction.

The MRAM reads and writes in just 35ns. As the polarity of the magnetic storage element does not leak away when power is removed, it operates as a non-volatile cell.

Again, as with F-RAM, some of the physical properties of an MRAM device produce real benefits for the circuit designer. For instance, the magnetisation of the cell does not decay with time or temperature, so an MRAM device can offer long storage intervals at relatively high temperatures.

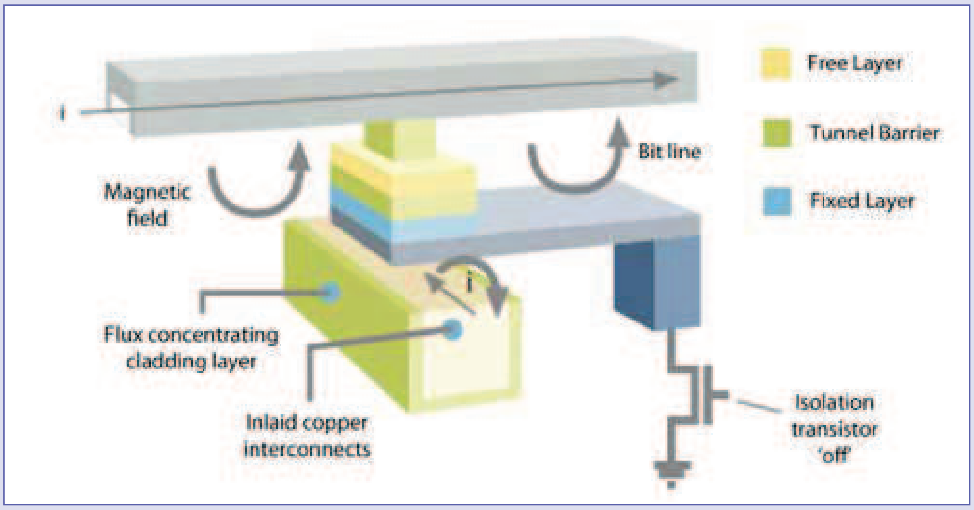

Fig. 3: MRAM utilises a Magnetic Tunnel Junction as a storage element.

Further, the structure is implemented with cells containing just one transistor and one MTJ. This will enable the production of devices offering higher density and lower cost in the future.

Lastly, as no electrons or atoms move during programming, there are no mechanisms that can wear out, which means that MRAM devices are rated for a very high number of write/erase cycles.

How to compare these new technologies

A glance at Table 1 appears to show little that would let the circuit designer distinguish one of these new technologies from another. It is, of course, clear that in terms of read and write speed and write/erase cycles, these new technologies offer far better performance than Flash and EEPROM. All three claim long data retention periods, they all support industrial temperatures, and they all have a similar price per bit (based on a comparison of 4Mbit parts).

The table appears to show that nvSRAM offers a lower read/write cycle capability. It is important to understand, however, that in the case of nvSRAM, it is the non-volatile part of the array that has lower read/write endurance. This part of the memory is not used in most write operations – it is only written to when power goes down. The endurance of the little-used non-volatile part of the array will therefore be sufficient in nearly all applications.

The same argument applies with the apparently longer programming time of nvSRAM. The figure quoted here is for the non-volatile array. In fact, in normal operation the system writes to the SRAM, and achieves read/write times of 25ns, faster than F-RAM or MRAM. This means that the nvSRAM devices might actually better support very high-speed processors in applications in which wait states are unacceptable. The trade-off, however, is that the designer will need to add an extra capacitor on to the board to provide power for the non-volatile array when system power goes down.

One common misconception should also be addressed in any study involving MRAM: MRAM technology is not, in practice, highly affected by magnetic fields. In the datasheet, Freescale specifies that its devices are immune to magnetic fields no higher than 25 Gauss. To put this in context, a loudspeaker will typically have a field of 90 Gauss directly on its surface. This drops, however, to less than 10 Gauss at 5mm distance. At 1cm from its surface, the loudspeaker's magnetic field is negligible. Practically, therefore, interference from magnetic fields should be easy to avoid in the majority of applications.

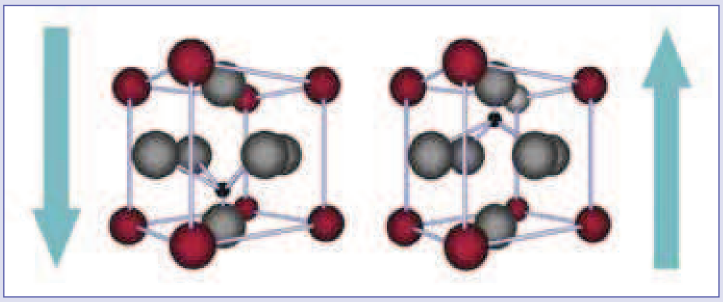

Fig. 2: F-RAM utilises two equal and stable low-energy states.

This kind of misconception is likely to bedevil MRAM more than F-RAM and nvSRAM, simply because it is the newest of the three technologies, and therefore the least widely used. Few designers have learned the reality of using MRAM. However, a large number of both licensees and product introductions have been announced in recent months, so it would appear that a considerable commercial momentum is now building behind the technology.

Indeed, all three technologies are commercially viable; they should definitely not be regarded as experimental. Each system designer will have their own particular requirements in terms of endurance, speed, board space and cost, and it is the combination of these factors that will point them towards one or other of the three technologies.

Conclusion

The three technologies discussed above are relatively new and advanced; their widespread adoption is far from assured. Flash, EEPROM and even traditional battery-backed SRAM have a far greater share of the non-volatile memory market, and they will hold that share for as long as they continue to do an adequate job. The traditional route can often be the safest, and many designers avoid taking risks with new designs.

Looking at the detail in the table, it could be concluded that MRAM wins hands down: it is reasonably fast, has good endurance, and offers the lowest price per bit (based on a 4Mbit part). However, as a newer entrant to the market than F-RAM or SRAM, it is not multi-sourced like the others, and this can be a crucial criterion for an industrial OEM. It is also worth remembering that pricing is dynamic, and new entrants to the market often use pricing policies to try to gain market share. These can sometimes prove to be unsustainable in the long term if volumes do not ramp up.

In the final analysis, these new technologies will find a place if sufficient applications require the very high speed and long endurance that they offer and can accept the extra cost.