SEMICONDUCTOR ENGINEERING

MANUFACTURING & PROCESS TECHNOLOGY

What Is Spin Torque MRAM?

Phill LoPresti drills down into new memory types, why and where they're needed, and why it's so hard to develop them.

AUGUST 17TH, 2017 - BY: MARK LAPEDUS

Mark LaPedus is Executive Editor for manufacturing Semiconductor Engineering

Phill LoPresti led Everspin Technologies, Inc. *for 7 years as its CEO and President during which Spin-Torque was developed and brought to commercialization by the company*

The memory market is going in several different directions at once. On one front, the traditional memory types, such as DRAM and flash, remain the workhorse technologies. Then, several vendors are readying the next-generation memory types.

As part of an ongoing series, Semiconductor Engineering will explore where the new and traditional memory technologies are heading.

For this segment, Phillip LoPresti, president and chief executive of Everspin, sat down with Semiconductor Engineering to discuss the magnetoresistive random access memory (**MRAM**) market. What follows are excerpts of that conversation.

SE: What is MRAM?

LoPresti: MRAM is a nonvolatile memory that provides the aspects of a working memory like **SRAM** and **DRAM**. What I'm referencing there is write speed, write

bandwidth and symmetrical read and writes with high endurance. It combines that with the data retention that you would see in a memory like a NAND flash product. It retains the data after the power is down or the power is inadvertently cut off. For MRAM densities, our target is to have them track closer to DRAM density products. It's not a product that is inclined to be a storage media or compete with **3D NAND**—or NAND, for that matter.

SE: What are the key markets for MRAM?

LoPresti: Today, we service various markets, such as transportation and automotive. We also service enterprise storage and industrial applications.

SE: How does MRAM work?

LoPresti: We use magnetics or the manipulation of electron spin to control the resistance of the bit, which allows us to program 1s and 0s.

SE: There are different types of MRAM, right?

LoPresti: MRAM is the basic category. Under MRAM, there are various types of ways to implement the MRAM bit. Our first-generation product, in which we have shipped over 70 million units, is targeting the SRAM density space. Its densities are 128-Kbits to 16-Mbits. That first-generation technology is referred to as toggle MRAM. It's called a field-switched MRAM. That means you use a magnetic field that is created on the chip to actually change the resistance on the bit from one state to another.

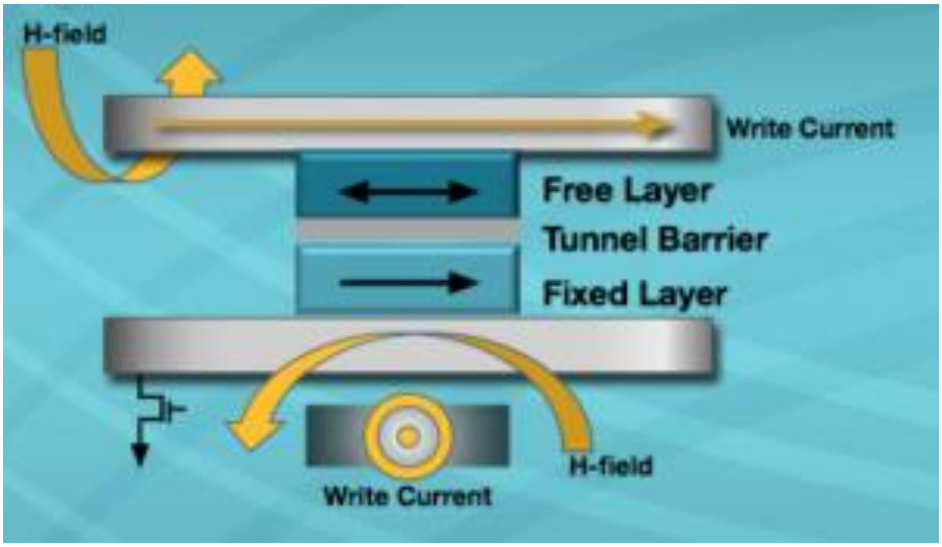

Fig. 1: Toggle MRAM technology. Source: Everspin

SE: Everspin and others are developing a more advanced version called spin torque MRAM (ST-MRAM or STT-MRAM). How does that work?

LoPresti: Instead of using a field to switch the bit, we actually use a spin polarized current that you either drive up or drive down the bit. Depending on the direction of that current, you can flip the spin of the electrons from one direction to the other, thus changing the magnetic polarity and the resistance of that bit. In the spin torque MRAM category, there are at least two flavors. One is called in-plane, which means that the electron spins are horizontal or in the plane of the bit. That was our first spin torque part, which we introduced some time ago. It was a 64-megabit DDR3 part. That was a 90nm process on 200mm wafers.

SE: Now, Everspin and others are focusing on another version called perpendicular ST-MRAM or STT-MRAM. How does that work?

LoPresti: It's basically the same type of concept. But instead of the electrons spinning horizontally and changing in the plane, it changes perpendicular to the plane. Perpendicular spin torque MRAM has the advantage of providing lower switching currents to allow us to make smaller bits, utilize smaller transistors, reduce the power, and increase the density of the parts. We've successfully implemented perpendicular spin torque with our recent 256-megabit DDR3 part, which we put into production back in March. That's the first production chip coming out of GlobalFoundries' production line. That was a partnership announced with GlobalFoundries for 300mm back in 2014. It's a 40nm process.

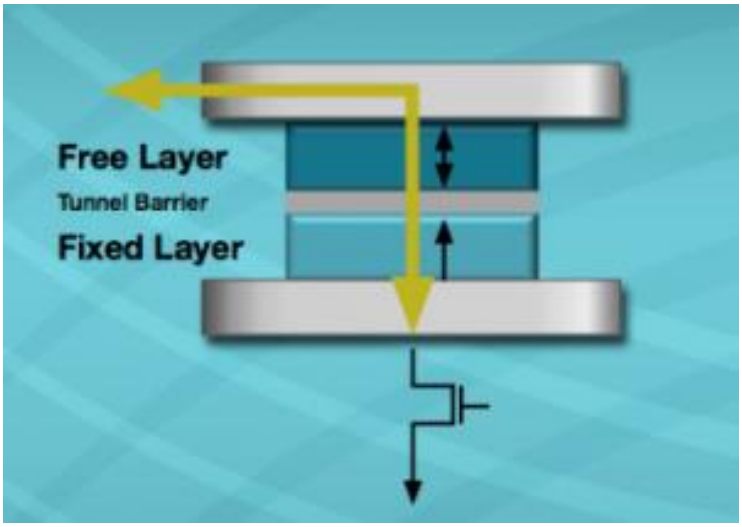

Fig. 2: Spin torque MRAM technology. Source: Everspin

SE: What's the latest from Everspin?

LoPresti: We've announced another perpendicular part, which is a 1-gigabit DDR4. Again, it's using GlobalFoundries, but this time it's on 28nm.

SE: Why move to 300mm?

LoPresti: We realized that spin torque had the potential of driving higher-density parts and getting us into persistent DRAM. Those parts needed to be built on 300mm wafers. You needed process nodes that didn't exist on 200mm wafers.

SE: Recently, GlobalFoundries, Samsung, TSMC and UMC have separately announced plans to develop and offer embedded STT-MRAM for customers. Why are they jumping on the technology?

LoPresti: MRAM has an interesting benefit to the foundry market. It's a very good substitute for embedded memory in projects that their **SoC** customers are going to want and try to do in sub-28nm process nodes. In particular, you can configure MRAM to operate like an embedded flash IP core, or you can configure it to look like embedded DRAM. Or you can configure it to look like embedded SRAM.

SE: Everspin's foundry partner is GlobalFoundries. Everspin also licensed the technology to GlobalFoundries for embedded MRAM, right?

LoPresti: GlobalFoundries captured the concept quickly. They realized that embedded flash would maybe get them through 28nm. But it becomes very challenging to put that on smaller nodes and even more challenging on a **fully-depleted SOI** process. So we licensed it to GlobalFoundries. They can deploy the technology in the embedded memory space. In the meantime, they put it in a production line. We started using that to produce our parts. Then, in September 2016, GlobalFoundries announced that they were rolling out embedded MRAM in their 22nm FD-SOI process. They are considering it for other process nodes. They have a 12nm FD-SOI that they plan in the future, and perhaps their 7nm or 14nm finFET processes.

SE: Where is Everspin targeting its spin torque MRAM devices in the market?

LoPresti: They have been optimized and specifically targeted to service the write caches and write buffers for enterprise-class SSDs and RAID systems. Both SSDs and RAIDs have a write cache. In both cases for the enterprise world, they need to be write-protected on power fail.

SE: For these applications, some use DRAM. What problem does ST-MRAM solve for write-cache apps?

LoPresti: To get the bandwidth and latency, you use DRAM. The problem is that DRAM is volatile. So if it's a write-protected cache and they use DRAM, then they have to connect it to an alternative power source to ensure that there is enough energy in the system or on the card. This is to ensure whatever data is in flight is either protected in some kind of nonvolatile memory. In other words, it has to get scrammed out of the DRAM into perhaps a bank of NAND. Or, it must get to the final media, which could ultimately be SSDs or hard drives. That is where we've targeted using ST-MRAM parts at

256-Mbit and 1-Gbit. We essentially provide similar latency to DRAM. We offer the same bandwidth—DDR3 or DDR4—on our parts. However, as soon as data is written into our parts, you no longer have to connect to supercapacitors or batteries because they are nonvolatile. You no longer need to have a power fail safe system set up, or have charging circuits, or have scram-mode programed into your firmware. It's quite simple. You write into the MRAM. The power goes out and the MRAM is going to have it. So when the power goes up, you can write it to the media. It simplifies the system and eliminates components that tend to be a cost-of-ownership problem or a reliability issue. It frees up space in the drives or a RAID system.

SE: So Everspin is targeting ST-MRAM in what many call the persistent memory market, right?

LoPresti: We believe MRAM has a very solid position to address persistent DRAM. In particular, this is for these high-speed nonvolatile caches that are necessary in many of the high-performance systems out there.

SE: Let's talk about the memory landscape. The industry is not only working on STT-MRAM, but also 3D XPoint, FRAMs, phase-change, ReRAM and even carbon nanotube RAMs. How will this all play out?

LoPresti: Each one of these memories has a different advantage or strength. I don't believe in a universal memory, where one memory solves every problem. Future storage systems and compute products, in general, are going to have various memory technologies that serve different purposes.

SE: What about traditional memory?

LoPresti: Memory technology has not really evolved very much in an exciting way for 20 to 30 years or more. You've had DRAM, SRAM and NAND. The biggest story for all of them has been lower cost-per-bit and higher density.

SE: For some time, Everspin has been the only company shipping standalone MRAM parts. Has that changed?

LoPresti: As far as we know, we are the only ones shipping MRAM products. I know there are smaller companies or private startups that are pursuing MRAM. But I haven't seen anything from them recently, nor am I aware of anything that's shipping. More importantly, I don't think any of them have a production fab partner that would supply an end customer with material if they had a product.

SE: There are many challenges to developing MRAM, right?

LoPresti: If MRAM was easy, everyone would be doing it. It's not simple. Nor is it simple to do **resistive RAM** or some of the other things. They are all challenging.

SE: What is taking the other suppliers so long to develop MRAM?

LoPresti: It's hard to explain why the others are further behind. We have a lot of history, IP and experience. We have access to a production line. That puts us in a good position.

SE: How far can you scale ST-MRAM?

LoPresti: We've produced products at 90nm and 40nm. We're sampling at 28nm. Our partner, GlobalFoundries, plans to deploy it at 22nm. They also believe they have enough confidence in this technology to perhaps be a solution for their 12nm, 14nm and 7nm products. Certainly, when you go to those geometries, you have more challenges that you need to overcome. For MRAM technology in itself, there really isn't anything that says you can't do a 7nm solution.

SE: Will ST-MRAM ever replace DRAM altogether?

LoPresti: I don't know if DRAM will continuously scale and provide the cost-per-bit benefits that it has been able to achieve over the last 10 years. Engineers will find a way to keep making DRAM better. What we see is that the time to improve them is taking longer and costs are not coming down as rapidly. It's a more complicated process. At some point in time, it's certainly possible that another technology like MRAM could supplant it. But there's a lot that could happen between now and then.

SE: What's after the 1-Gbit chip from Everspin?

LoPresti: We haven't made a public announcement about our roadmap beyond 1-gigabit.

SE: What are the emerging apps for ST-MRAM?

LoPresti: We are exploring higher density parts. We are also exploring parts that address specific applications. One area that seems of interest is automotive, such as ADAS systems. We also have artificial intelligence and deep learning.

SE: What about the cost of spin torque MRAM?

LoPresti: With every density increase we provide, there is a cost-per-bit improvement. We don't necessarily say if we will or we won't reach parity with another technology. That's not necessarily our goal. We're designing parts that add value. We don't necessarily think the products should be at parity with a commodity memory. I wouldn't say there is a premium for it. It has its own curve, just like when NAND had its own curve when it was introduced. We also know that higher density and lower cost-per-bit drives the opportunity to address more markets.