Non-Volatile Memory and Its Use in Enterprise Applications

Contributor:

Viking Technology

January 2014

About the SNIA

The Storage Networking Industry Association (SNIA) is a not–for–profit global organization, made up of some 400 member companies spanning virtually the entire storage industry. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, the SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For additional information, visit the SNIA web site at http://www.snia.org or http://www.snia.org/international for a list of International Affiliates

About the SNIA Solid State Storage Initiative

The SNIA Solid State Storage Initiative (SSSI) fosters *the growth and success of the market for solid state storage in both enterprise and client environments*. Members of the SSSI work together to promote the development of technical standards and tools, educate the IT communities about solid state storage, perform market outreach that highlights the virtues of solid state storage, and collaborate with other industry associations on solid state storage technical work. SSSI member companies represent a variety of segments in the IT industry (see http://www.snia.org/forums/sssi/about/members). For more information on SNIA's Solid State Storage activities, visitwww.snia.org/forums/sssi and get involved in the conversation at http://twitter.com/SNIASolidState.

About Viking Technology

Viking Technology is recognized as a leader in NVDIMM technology. Supporting a broad range of memory solutions that bridge DRAM and SSD, Viking delivers solutions to OEMs in the enterprise, high-performance computing, industrial and the telecommunications markets. Viking Technology is a division of Sanmina Corporation (Nasdaq: SANM), a leading Electronics Manufacturing Services (EMS) provider. More information is available at http://www.vikingtechnology.com.

Table of Contents

Executive Summary

This paper reviews different memory technologies, current and future, volatile and non-volatile. The paper also provides an overview of Non-Volatile DIMMs and their integration into industry standard servers using supercapacitor technology. Additionally, the paper will discuss how NVDIMMs enable increased application performance and significantly improved power failure/system crash recovery.

An Overview of Non-Volatile Memory

Over the past 20+ years, numerous memory technologies have been brought to market with varying degrees of commercial success, for example: Static RAM (SRAM), Pseudo Static RAM, NOR Flash, Eprom, EEprom, DRAM and NAND flash.

Generally speaking, these "memory" technologies can be split into two categories, volatile and non-volatile; Volatile memory will not retain data when power is turned off; conversely, non-volatile memory will retain data once power is turned off. The two dominating memory technologies in the industry today are DRAM (volatile) and NAND flash (non-volatile). See Figure 1 below which summarizes emerging, niche, and those memories in mass production:

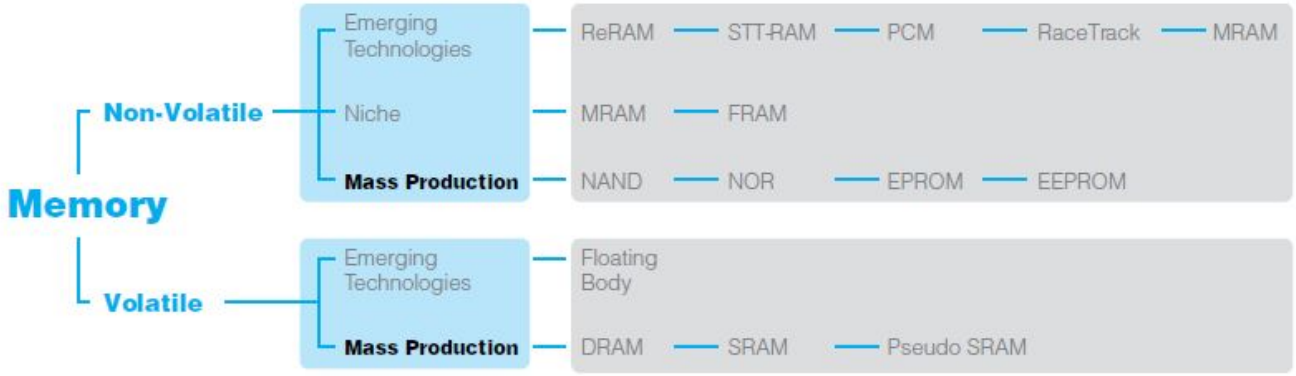

Figure 1. Categories of Memory

The memory technologies that dominate the computing industry today are DRAM and NAND flash, but it should be noted that both of these technologies have their pros and cons:

- DRAM delivers the highest performance (latency /speed), with practically infinite endurance, yet it is volatile and has much lower capacity points than other memories (i.e. NAND flash).
- NAND Flash, on the other hand, scales to high capacity, is non-volatile, and relatively cheap ($/Gbit); but is significantly slower than DRAM. Additionally, endurance and data retention are getting worse as process geometries continue to shrink, meaning that for write intensive enterprise applications. NAND flash, in the long term, may well not be an optimal memory technology.

Table 1: Pros and Cons of DRAM and NAND Flash

Pros and Cons of DRAM and NAND Flash		
TECHNOLOGY	**DRAM**	**NAND**
Speed	✓	
Power	✓	
Capacity (GB)		✓
Endurance	✓	
Data Protection (Non-Volatile)		✓
Cost/GB		✓
Reliability	✓	
Technical Improvements with Geometry Shrinks	✓	

Are New Technologies Emerging?

There is much discussion in the industry as to what new ***universal memory technology*** or technologies will materialize as real contenders to displace either or both NAND flash & DRAM. Some of these newer emerging technologies[1] include: MRAM (Magnetic RAM), FRAM (Ferroelectric RAM), PCM (Phase Change Memory), STT-RAM (Spin-Transfer Torque RAM), ReRAM (Resistive RAM) and Memristor.

Currently FRAM, MRAM, and PCM are in commercial production but still, relative to DRAM & NAND flash, remain limited to niche applications. There is a view that MRAM, STT-RAM, and ReRAM are the most promising emerging technologies, but they are still many years away from competing for industry adoption. Any new technology must be able to deliver most, if not all, of the following attributes in order to drive industry adoption on a mass scale: scalability of the technology; speed of the device; power consumption to be better than existing memories; endurance; densities better than existing technologies, and finally cost. If the emerging technology can only manage one or two of these attributes, then, at best, it is likely to be resigned to niche applications.

One important question is, ***"What will be the next non-volatile memory to replace NAND flash?"*** Currently it is unclear, but due to the limitations of NAND flash (endurance, reliability, speed) when compared to DRAM, it is likely a new non-volatile memory technology will evolve; but perhaps not replacing the current mainstream memories for at least the next 5 to 7 years. However, Non-Volatile DIMMs can deliver an enterprise class solution to enable increased application performance and far improved power failure/system recovery, when compared to current implementations.

[1] Appendix A provides additional information on these technologies

What Is Non-Volatile DIMM?

A Non-Volatile DIMM (NVDIMM) is a module that can be integrated into the main memory of an industry standard compute platform (i.e. server), perform workloads at DRAM speeds, yet be persistent & provide data retention in the event of a power failure or system crash.

A Non-Volatile DIMM, as illustrated by Viking Technology's ArxCis-NV™, is a memory subsystem that combines the speed and endurance of DRAM, together with the non-volatile data retention properties of NAND flash. This marriage of DRAM and NAND technology delivers a high speed and low latency "non-volatile/persistent" memory module. Designed from the ground up to support unlimited read/write activity, it performs at fast DDR3 speeds and can sustain itself from host power failure or a system crash. This solution can be viewed as the first commercially viable "Storage Class Memory" for the enterprise market.

Figure 2. ArxCis-NV from Viking Technology

What makes this module different from standard DRAM modules is that, in the event of a power failure or system crash, the data in the NVDIMM is securely preserved and available almost immediately upon power being restored to the host system (suspend/resume).

A NVDIMM like the ArxCis-NV has been architected to take advantage of the "best attributes" of the DRAM & NAND memory technologies. NVDIMMs like these, together with advanced data movement logic and supercapacitor technology, can deliver exactly what the enterprise market has been yearning for: ***a persistent, Non-Volatile DIMM*** that can be integrated into the main memory of industry standard server hardware.

Why Is NVDIMM Valuable to the Data Center?

If performance is critical to business success and if minimizing down-time is an important issue, then NVDIMMs are valuable to your data center.

Value: Data Center Applications

Some data center applications where Non-Volatile DIMMs would provide significant value are: ***searching/mining, business analytics, digital media creation/transmission, or financial modeling*** – essentially, wherever the bottleneck is storage and I/O, and where downtime costs money

With the recent surge in use of SSDs in the data center and those architectures also utilizing caching software (auto-tiering), most applications have enjoyed significant performance improvements. Therefore the SSD and software bundle has significantly improved the memory/storage gap, which has read applications, not write intensive ones. Therefore, these intelligent caching software solutions when paired with SSDs will utilize other system resources to ensure performance & reliability (CPU, DRAM and HDDs).

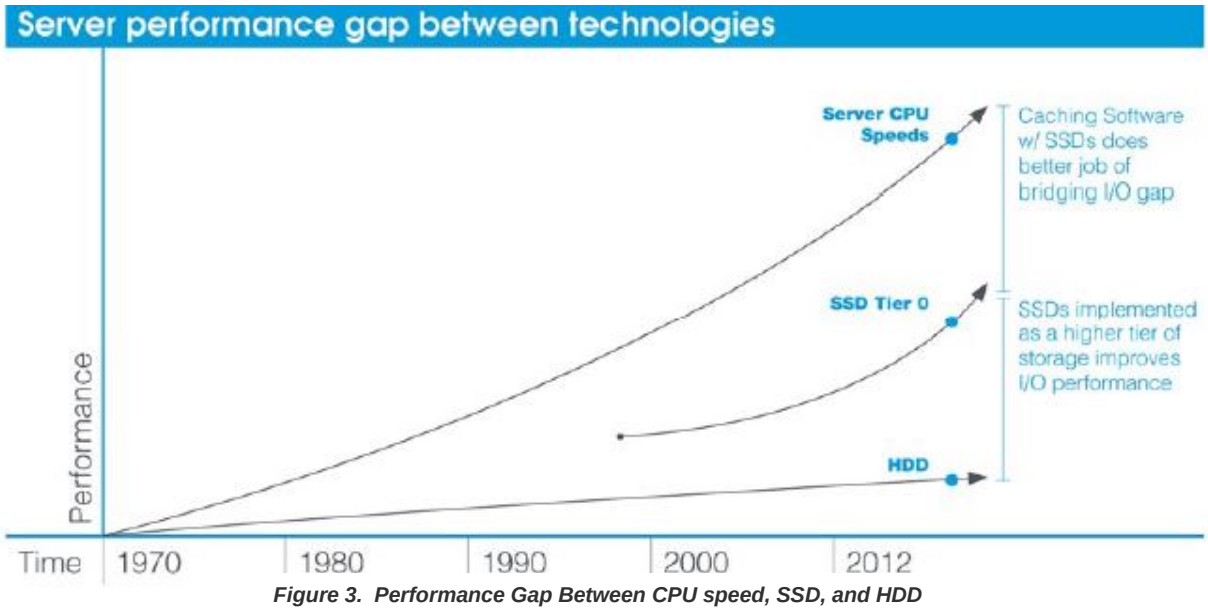

Figure 3. Performance Gap Between CPU speed, SSD, and HDD

Indeed, most SSD caching software will utilize the host system's standard DRAM for intensive write activity to preserve the SSD and also keep the write cache in DRAM; which means this critical data is susceptible to loss in the event of a system crash or power failure. A common work-around to protect the data is to "checkpoint" the write buffer out into slower block based storage (i.e. disk) at regular intervals, but "checkpointing" naturally has a negative impact on the I/O performance.

By integrating NVDIMMs into the host system, data-centric and write intensive applications will enjoy a significant increase in performance. In addition to the benefits of ***increased application performance, recovery*** is another area of significant benefit.

Should a data center experience a power outage, one of two scenarios occur: 1) either this power failure causes a catastrophic loss of the "in-memory state"; or 2) the backup power supplies enable the appliances (servers) to transfer this data held in main memory out to disk; the entire state (which could be 100's of Gigabytes of DRAM) must be saved to a storage back end (i.e. NAS, SAN). Both "***saving" & "recovering/reconstructing***" this amount of data across multiple servers will be extremely slow and place a very heavy load on the storage infrastructure, resulting in severe I/O bottlenecks.

> ***A real world example of this exact scenario happened when a large social networking site experienced a power outage which caused the service to be down for 2.5 hours while in-memory cache servers refreshed their state. This example can be seen as a successful data reconstruction event, but 2.5 hours of downtime would cause negative financial impact to many businesses.***

By utilizing Non-Volatile DIMMs, in the event a data center loses power or experiences a system crash, the NV-DIMMs allow the servers to recover their "in-memory state" almost instantaneously without putting any load on the storage back end (SAN/NAS) – in a sense, ***making the failure appear as a suspend/resume event***. This means that business critical applications such as OLTP (on-line transaction processing) can be up and running again in a matter of minutes rather than hours, and without the need for UPS (uninterruptable power supply) intervention.

Data centers, cloud service providers, and Infrastructure as a Service (IaaS) will have Service Level Agreements (SLAs) in place with their customers to guarantee up-time. Use of NVDIMMs will help reduce any downtime caused by system crash or power failure event, thus providing significant financial value.

Value: Performance

The architecture of NVDIMMs provides a full interconnect on-module that will independently allow transfer of data between the DRAM and the NAND Flash without contention for other I/O or CPU resources. Ultimately, applications can rely on the high speed memory (DRAM) to be "persistent" and not need slow down to "checkpoint" or consume other system resources.

The NVDIMM delivers value that far surpasses a simple DRAM/DIMM and SSD architecture, it is greater than the sum of the two technologies used.

Appendix A. A Glossary of Emerging Technologies

MRAM: Magnetic RAM

MRAM is a non-volatile memory. Unlike DRAM, the data is not stored in a electric charge flows, but by magnetic storage elements. The storage elements are formed by two ferromagnetic plates, each of which can hold a magnetic field, separated by a thin insulating layer. One of the two plates is a permanent magnet set to a particular polarity; the other's field can be changed to match that of an external field to store memory.

Vendors include: Everspin (Freescale spin-off), Crocus Technology

STT-RAM: Spin-Transfer Torque RAM

STT-RAM is an MRAM (non-volatile), but with better scalability over traditional Magnetic RAM. STT is an effect in which the orientation of a magnetic layer in a magnetic tunnel junction or spin valve can be modified using a spin-polarized current. Spin-transfer torque technology has the potential to make possible MRAM devices combining low current requirements and reduced cost; however, the amount of current needed to reorient the magnetization is at present too high for most commercial applications.

Vendors include: Samsung, SK Hynix, Renesas, Toshiba, Everspin, Crocus Technology.

PCM: Phase Change Memory

PCM is a non-volatile random access memory. It utilizes the unique behavior of chalcogenide (a material that has been used to manufacture CDs), whereby the heat produced by the passage of an electric current switches this material between two states. The different states have different electrical resistance which can be used to store data. It is expected PCM will have better scalability than other emerging technologies.

Vendors include: Micron, Samsung.

ReRAM: Resistive RAM

ReRAM is a non-volatile memory that is similar to PCM. The technology concept is that a dielectric, which is normally insulating, can be made to conduct through a filament or conduction path formed after application of a sufficiently high voltage. Arguably, this is a memristor technology and should be considered as potentially a strong candidate to challenge NAND flash.

Vendors include: SK Hynix, HP, NEC, Panasonic, Samsung.

Appendix B. More on Supercapacitors

In order for a Non-Volatile DIMM like ArxCis-NV™ to perform its task, a small energy source is required to ensure 100% data security on system failure. Supercapacitors are an ideal technology for use in this environment, primarily because they provide a superior solution when compared to batteries.

Supercapacitor technology is relatively new to the enterprise arena; any technology, when new and relatively unknown will encounter questions about long term reliability and capabilities. Here follows a few reasons why supercapacitors are suitable for use in this type of enterprise application:

1. **High Efficiency:**: Supercapacitors are highly efficient components. Their efficiency (defined as the total charge removed divided by the total charge added to replenish the charge removed) is greater than 99%, even at very high currents; meaning that little charge is lost when charging and discharging the supercapacitor.

2. **High Current Capability**: Supercapacitors are designed with a very low equivalent series resistance (ESR), allowing them to deliver and absorb very high current. The inherent characteristics of the supercapacitor allow it to be charged and discharged at the same rates, something no battery can tolerate. In battery-based systems, you can only charge as fast as the battery will accept the charge.

3. **Wide Temperature Range:** Since supercapacitors operate without relying on chemical reactions, they can operate over a wide range of temperatures. On the high side, they can operate up to 65°C, and withstand storage up to 85°C, without risk of thermal runaway. On the low side, they can deliver power as cold as -40°C.

4. **Condition Monitoring (SOC & SOH):** Determining battery state of charge (SOC) and state of health (SOH) is a significant factor for robust battery systems, requiring sophisticated data acquisition, complex algorithms, and long-term data integration. In comparison, it is very simple to determine the SOC and SOH of supercapacitors.

5. **Long Cycle Life:** The energy storage mechanism of a supercapacitor is capable of hundreds of thousands of complete cycles with minimal change in performance. They can be cycled infrequently, where they may only be discharged a few times a year, or they may be cycled very frequently.

6. **Long Operational Life:** The energy storage mechanism of a supercapacitor is a very stable process. It is capable of many years of continuous duty with minimal change in performance. In most cases, supercapacitors are installed for the life of the system.

7. **Ease of Maintenance:** Supercapacitors cannot be over charged/discharged, and can be held at any voltage at or below their rating. If kept within their wide operating ranges of voltage and temperature, there is no recommended maintenance.