NANTERO

White Paper

NRAM as Storage Class Memory

NRAM Cache: A Fundamental Improvement in Disk Drive Design

By Bill Gervasi

Abstract: Storage class memory is a term first used by IBM in the late 2000s. It has become synonymous with the majority of next-generation memory technologies given the fact that they all fall somewhere between NAND flash memory and Dynamic Random Access Memory (DRAM) in just about every attribute ranging from cost to performance. Significant investments are underway already to develop storage class memory, due to the enormous customer demand and the fact that both DRAM and NAND are falling off Moore's law. This white paper will discuss the compelling value proposition of storage class memory in rotating magnetic media and solid state disk drives and discuss why NRAM® – based on Nantero's carbon nanotube technology (CNT) - is well suited to provide such value.

About the Author: Mr. Gervasi is Principal Systems Architect at Nantero, Inc. He has been working with memory devices and subsystems since 1Kb DRAM and EPROM were the leading edge of technology. He has been a JEDEC chairman since 1996 and responsible for key introductions including DDR SDRAM, the integrated Registering Clock Driver and RDIMM architecture, the formation of the JEDEC committee on SSDs, and actively involved in the definition of NVDIMM protocols. He chairs the emerging memories committee and authored the JEDEC Automotive SSD standard.

Introduction

The DRAM cache inside a disk drive of either rotating magnetic media or solid state type is a key determinant of the performance of the drive, yet its use also restricts many aspects of drive controller design. Disk drives, both hard disk drives (HDDs) and solid state drives (SSDs), incorporate a DRAM cache for data and commands for many reasons. Keeping up with the data rate of the interface (SATA, PCIe, etc.) requires that data be stored in read and write buffers. Re-ordering commands helps extend the life expectancy of the write-limited Flash memory used for mass storage in SSDs, and for rotating media allows concatenation of sequential operations for performance improvements and power savings. A larger cache means that more data can be buffered and more commands may be reordered, resulting in a fairly direct relationship between increased cache size and increased performance—however cache size is limited by the energy required to maintain its contents and therefore the ideal cache size cannot be achieved today

This cache, typically implemented using DRAM or less typically, SRAM, may also be used for program execution and data for the drive controller device, thereby keeping the cost of the controller down since it requires fewer on-chip resources that consume die area. Data tables, bitmaps, compression intermediate results, etc. can be stored in the DRAM if the interface speed is high enough to prevent performance bottlenecks.

Energy and Capacitance

Because power can fail with very little advance notice, and because data stored in the drive DRAM cache must be preserved to avoid corruption of the media, each drive controller must incorporate an energy source that keeps power stable until a flush operation can commit all the data in the cache to the media.

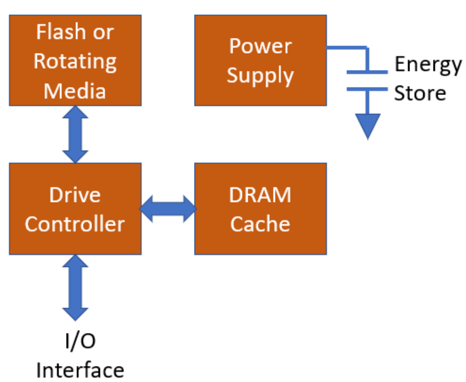

Figure 1: Block Diagram of a Typical Disk Drive

This energy source requirement is one of the most difficult aspects of drive controller design. The amount of energy required to flush a cache to the media is quite high since the DRAM, the controller, and the media all need to be powered continuously until the flush is complete. While the total amount of energy is pretty dependent on the design, it can take a Farad of capacitance to maintain power during a flush of 128MB of DRAM cache to NAND flash, and substantially more to flush the DRAM to rotating media.

This also explains why caches are kept relatively small in drive controllers despite the performance gains that could be realized from a larger cache. If you increase cache size, you must also increase the energy store needed for flushing the cache to the media. Ironically, this penalty is greater for higher capacity drives because the amount of media that must be powered is greater, causing an unfortunate marketing dilemma where lower capacity drives can have higher performance than their more expensive cousins.

Capacitance is a real pain for systems designers. The main options are wet dielectric supercapacitors or dry dielectric metallic types (such as tantalum). The cost and size of a supercapacitor solution can be quite attractive, around $1 and $400mm^2$ of board space, or 7% of available area. However, the long term reliability of a supercapacitor is not very attractive for high end solutions. The wet dielectric tends to evaporate over time, especially in the hot environments most hard drives are used in. This forces greater overprovisioning, and largely restricts supercapacitors to lower end solutions.

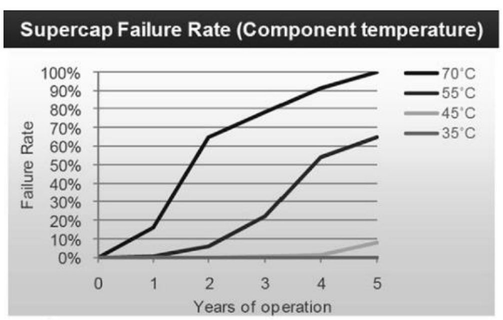

Figure 1: Long Term Reliability of Supercapacitors

The alternative, metallic capacitors, for the equivalent storage capacity as a supercapacitor is a painful tradeoff. The solution cost can jump to $4 and the footprint increase to $1000mm^2$, or 18% of available space, and often require cutouts in the circuit board to fit the radial bodies. This severely impacts not only the layout and routing on the PCBs, it also becomes a mechanical issue to prevent the radial components from vibrating and the solder joints from cracking.

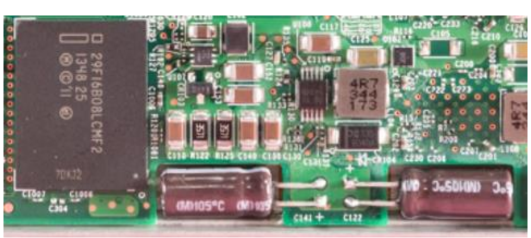
Figure 2: Tantalum Capacitors for Cache Flushing

Benefits of NRAM

Nantero NRAM is a non-volatile memory with DRAM performance and capacity. It connects to the drive controller exactly like a DRAM, providing the necessary throughput to handle data and command buffering as well as controller data access. Based on resilient carbon nanotube technology, NRAM does not suffer from leakage, dielectric fractures, heat sensitivity, or other problems associated with other memory devices. When power fail is detected, the controller need only complete a data burst in process then shut down and wait for power to be restored. This eliminates the need for expensive and space consuming energy sources.

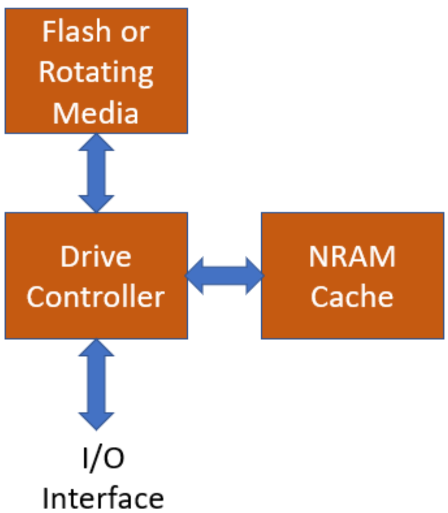

Figure 3: Drive Controller Block Diagram Using NRAM

NRAM offers other features that can help simplify disk drive controller design as well. Today's disk drive controllers using DRAM for a cache must support the needs of the DRAM for going offline regularly to implement the required refresh cycles of a DRAM. If data comes in on the host interface during a refresh cycle, the disk drive controller must buffer all such data until the refresh cycle completes and all buffered data can be written to the DRAM. This additional on-chip buffering increases disk drive controller die size and cost.

NRAM does not require any refresh cycles, therefore 100% of the bus bandwidth is available for data traffic. This means that future disk drive controllers designed to take advantage of NRAM can reduce the required internal data buffering size, further reducing cost, complexity, and increasing reliability.

The DRAM compatible speed of the NRAM is a related benefit. Controllers that want to use the cache for controller code or data may simply access the NRAM at DRAM speeds at any time, without the required times for refresh cycles where the memory is unavailable.

As long as we're tallying the advantages of NRAM over DRAM, perhaps one of the biggest advantages is the temperature insensitivity of the NRAM. Due to the capacitive nature of the storage cell in a DRAM, they tend to be very sensitive to heat, and in fact will fail above 95℃. Disk drives must work extra hard to deal with this in many ways including increasing refresh rates, throttling accesses, or eventually shutting down if the temperature exceeds the DRAM specification.

NRAM's carbon nanotube memory cells do not share this sensitivity to heat, operating in ranges far beyond the other components in the system with the possible exception of the metal screws that hold the SSD together. NRAM eliminates these thermally related problems inherent in DRAM-based designs.

Though it may be obvious from the concepts described above, a little indulgence is to point out that the elimination of the energy source requirements decouples the cache size calculation from system design. Doubling, quadrupling, even growing the cache by 16X or 256X has absolutely no impact on energy requirements for power fail cycles. When power fails, the same procedure is followed: complete any single access in process then shut down and wait for the juice to come back. This opens the way for significant increases in cache size that would lead to dramatic improvements in overall disk drive performance.

Conclusion

In summary, NRAM offers a number of advantages for both HDD and SSD designs, from eliminating energy source requirements to non-volatility, and from thermal insensitivity to improving data availability. The combination of these advantages can lead to a disk drive with a value proposition that is simply not possible today, and one that the users of the disk drives will immediately notice and appreciate. By decoupling the cache size from concerns about an energy store, disk drive designers can provide users with the performance levels expected of tomorrow's computing systems.