OBJECTIVE ANALYSIS
Semiconductor Market Research

Why Wait For Storage Class Memory?

Novel Architecture Provides SCM's Benefits Today

Objective Analysis Brief

Computer architecture is increasing in sophistication at a rapid pace, and some of the fastest changes are occurring in storage. We have moved from systems without solid state storage to systems with flash SSDs, and are now poised to migrate to Storage Class Memory.

In the words of IBM, which coined the term: "Storage-class memory (SCM) combines the benefits of a solid-state memory, such as high performance and robustness, with the archival capabilities and low cost of conventional hard-disk magnetic storage."

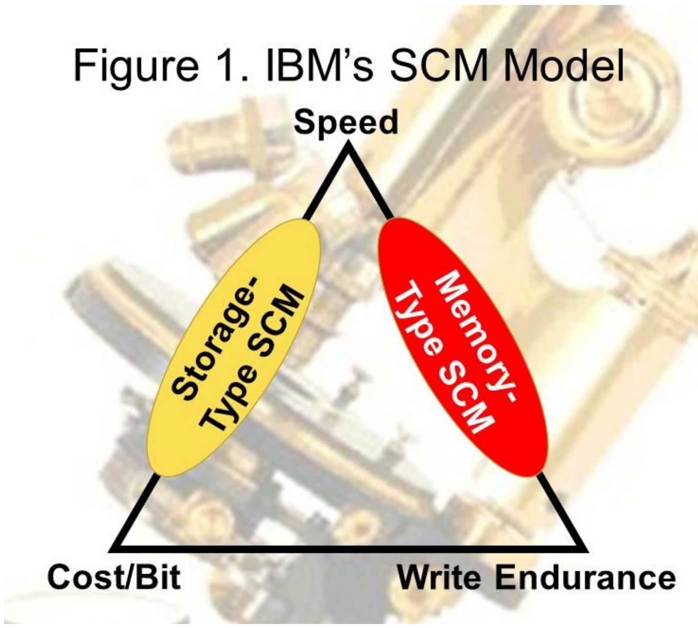

Figure 1. IBM's SCM Model

Figure 1 illustrates IBM's concept of two kinds of SCM, one aimed at storage, and one aimed at memory applications.

The biggest issue with SCM, though, is the industry's focus on using an emerging memory technology to implement the SCM layer. There is currently no sufficiently-economical embodiment of any of these technologies that can meet SCM's goal of providing the low cost of an HDD, or even cost that is lower than DRAM. Emerging memory technologies like MRAM, RRAM, and FRAM have never been put into high volume production, and all of them use new materials that aren't as well understood as silicon. All things considered, these technologies all cost orders of magnitude more than DRAM. Flash-backed NVDIMM-N modules are less costly than emerging memories, and are currently being used to develop SCM-compatible software, but these modules still sell for significantly more than DRAM. While the Intel/Micron 3D XPoint Memory promises to sell at a lower price than DRAM, it is still far from volume production.

What can be done right now to take advantage of the benefits of SCM? This Objective Analysis Brief will outline an SCM module that is being developed by Netlist and other companies, to fill the gap caused by the unavailability of economical emerging memories.

Why Computers Need SCM

Applications are changing the way computers access data, while the Internet has multiplied the amount of data being used in our everyday lives. Data scientists created the term “Big Data” to describe a class of problems that require an alternative approach to computing. Completely new data analytics software architectures have been devised to try and make sense of this data, breeding new database management tools like graph analytics, NoSQL, and myriad others that have helped mine data and extract understanding and insight from vast unstructured data lakes.

Meanwhile, the Internet of Things (IoT) promises to balloon these unstructured data lakes to sizes that will defy any attempt by today’s approaches to analyze this data. Big Data is poised to become so large that data analytics can only be performed in a timely manner by moving storage closer to the processor.

Figure 2. SCM’s Position in the Memory Hierarchy

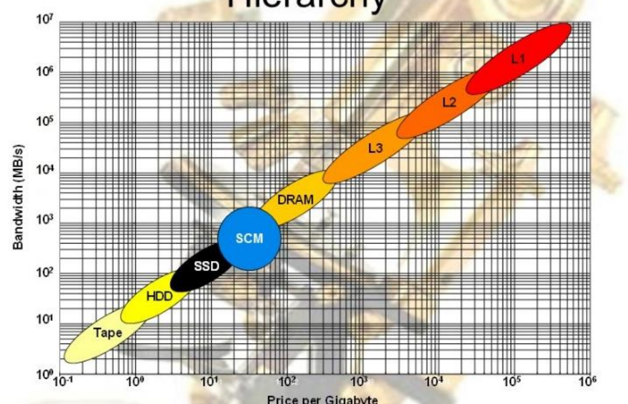

Source: *A Close Look at the Intel/Micron 3D XPoint Memory*, Objective Analysis 2015

Computer scientists hope to mitigate certain speed issues by using In-Memory Databases (IMDBs) like SAP Hana, memSQL, and Microsoft’s Hekaton to analyze the data. IMDBs accelerate data analysis by removing a lot of the fail-safe mechanisms that are a vital part of disk-optimized databases; these mechanisms modify the database’s data while virtually eliminating any potential for data loss.

While an IMDB can use the atomic writes that are supported by certain SSDs as an alternative to the time-tested disk-optimized approach which requires two writes and two reads for every data write, data scientists look forward to the day they can use SCM to store their data safely right where it is accessed by the processor itself.

Imagine the speed improvement made possible by replacing two reads and two writes to network storage by a single write to memory.

SCM Emerges

To this end the concept of SCM was created. Visionaries saw that the DRAM chip technology now used for computer memories would soon reach the point where it could no longer sustain the same annual cost reductions seen over the past four decades. Alternative technologies were being developed by all DRAM makers to supplant DRAM once this point was reached. All of these technologies were “nonvolatile” – their contents were not lost when power was removed. This is the key attribute of disk-based storage. Since the new memory was “persistent” (the storage architect’s word for “non-volatile”) the emerging memory technologies opened up new possibilities for higher-performance computers that did not need to access slower storage media like HDDs or SSDs to assure the permanence of the data.

Storage class memory is burdened with one more very important requirement: It must be cheaper than DRAM and faster than flash to make sense. This is illustrated in Figure 2.

This chart roughly explains how computer architects plan out the memory/storage hierarchy in their machines. Orbs are used to represent different memory or storage types: Tape, HDD, flash-based SSD, DRAM, and three levels of cache internal to the processor (L1-3). The bottom axis measures the price for a byte of each, and the vertical axis approximates their speed. Logarithmic axes are used since a linear axis would result in the L1 orb consuming nearly the entire chart, with all of the other orbs becoming vanishingly small in the lower left corner.

SCM fits between the SSD and DRAM orbs – it must be faster than an SSD and cheaper than DRAM to fit into this hierarchy. Anything else would not make sense: Few designers would choose a memory that is slower than DRAM but costs just as much or more, or a technology that was more costly than an SSD but offered no speed advantage.

Some designers value high-speed persistence enough to pay a premium, and these designers often today address their needs by using a module called the NVDIMM-N (nonvolatile dual inline memory module type “N”) that contains DRAM memory and a similarly-sized flash memory, with a processor and battery or supercapacitor. If power is lost the supercapacitor powers the device long enough to allow the contents of the DRAM to be moved into the flash. When power is restored, the processor moves the data from the flash back into the memory. This device allows users to write to memory at memory speeds without needing to worry about data loss in the event of a power failure.

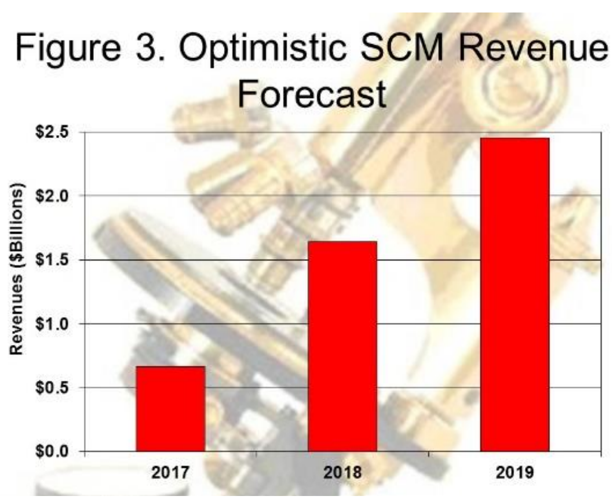

Figure 3. Optimistic SCM Revenue Forecast

The promise of SCM is that it will bridge the gap between DRAM and flash by offering a price per byte lower than that of DRAM with performance higher than that of any SSD. To paraphrase the IBM quote mentioned above, SCM ideally provides memory at storage capacities and storage at memory speeds.

Since it bridges the DRAM-SSD performance gap, SCM will provide a system with higher performance at the same cost, or at a lower cost, for the same performance level. This is the reason that so many members of the computing community are working hard to pave the way for SCM by defining software support standards and even new CPU instructions.

The Market Opportunity

One interesting point about the new Storage Class Memories is that they solve two problems:

1. SCM is cheaper than DRAM and faster than NAND flash.
2. SCM is persistent.

The first of these attributes gives SCM a valid place in the memory/storage hierarchy whether or not it is persistent, simply because systems will realize a better cost/performance ratio with SCM than they can without it. Designers will find they can decrease the amount of DRAM they use in a system, and then apply the savings to add SCM, realizing a net increase in performance for the same cost.

The second attribute is the one that has attracted the most attention. Database management experts now find they can support much faster access to storage to alleviate the risk of data loss in the event of power failure. Today this issue is addressed through slow and complicated storage protocols that, with SCM, can be completely bypassed.

To take advantage of in-memory storage, new software stacks are being developed, nonvolatile memory interface standards were defined by Storage Networking Industry Association (SNIA) in 2013, and the Linux community is developing its PMEM (Persistent Memory) libraries. Steady progress is being made towards taking advantage of storage class memory's persistence, and some software is already able to use SCM and the Linux PMEM libraries to its advantage.

The first problem that SCM solves could, by itself, drive very rapid and strong market acceptance. As soon as SCM ships at a price that is lower than that of DRAM, and as long as it is faster than SSDs, then SCM will find rapid acceptance as a way to reduce system costs by cutting DRAM requirements. SCM may also achieve strong adoption in those applications that require memories much larger than what can be supported by the server's processor.

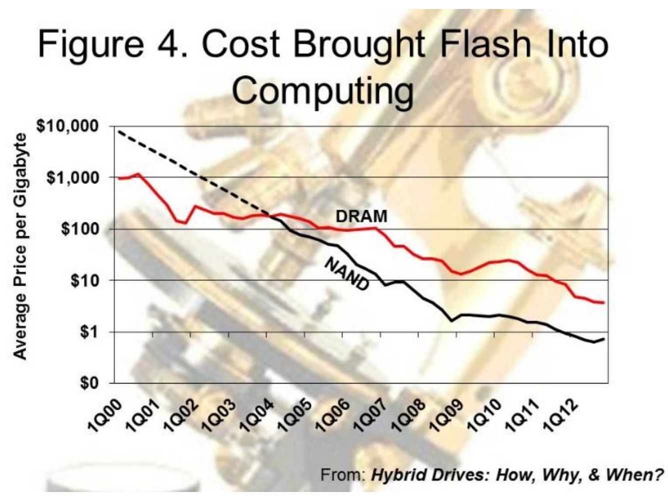

Objective Analysis has modeled the market opportunity for storage class memories used as a layer between SSDs and DRAM memory. This is illustrated in Figure 3.

This is an optimistic outlook for the technology that does not account for buyer sentiment, but simply assumes that anyone who could benefit from this technology adopts it as soon as it becomes available. This approach gives an upper bound for its market adoption.

The assumptions behind this forecast are that:

1. Every application that can use SCM will adopt it rapidly
2. SCM will be adopted in servers first, although it later could be accepted into other DRAM applications like cell phones and PCs
3. The average amount of DRAM per server will flatten, rather than grow as it does today
4. The memory spend per system will remain unchanged, with purchasers buying as much SCM as they did not buy DRAM
5. Meaningful shipments will begin in early 2017.

The market in this forecast grows very rapidly to reach nearly $2.5 billion by 2019. In reality adoption will be slower, as it was with SSDs, since many prospective users will put off adoption until they have witnessed whether the early adopters uncovered any issues.

Netlist's HybriDIMM

One company, Netlist, has found a way to address the need for SCM before any new memory technology becomes cost-competitive with DRAM. Netlist's implementation is built on a standard loadreduced dual in-line memory module (LRDIMM) interface that fits into most computing systems.

The HybriDIMM is similar to an NVDIMM-N in that it packages DRAM and NAND flash into a standard DIMM form factor. Unlike the NVDIMM-N, however, the HybriDIMM's flash bank is significantly larger than the DRAM. Caching and other sophisticated algorithms move data into and out of the DRAM as required by the application program without user intervention.

A rough block diagram of the HybriDIMM appears in Figure 5.

By teaming a smaller DRAM with a large flash memory the HybriDIMM satisfies most of the processor's data accesses at memory speeds while providing a very large addressable space at a price much lower than that of DRAM. Today one gigabyte of DRAM costs roughly twelve times as much as one gigabyte of NAND flash, and this gap has been widening in a relatively steady fashion since 2004 when NAND flash prices first crossed below those of DRAM. See Figure 4. It is very likely that his trend will continue, with the price gap between NAND flash and DRAM widening thanks to NAND's migration to a 3D cell structure.

Figure 5. HybriDIMM Block Diagram

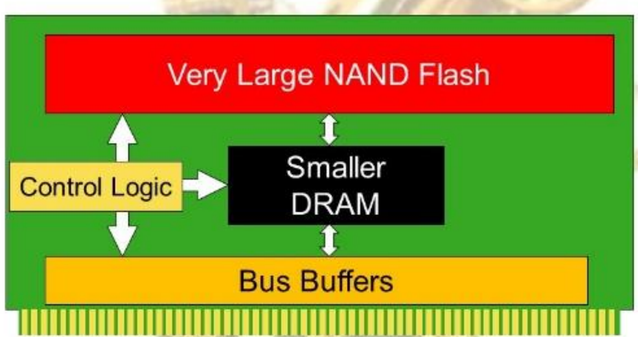

Netlist's intention is to deliver SCM modules today before next generation persistent memories, or emerging memories, become available. This will allow the company to take advantage of NAND's rapid price declines to keep the HybriDIMM's price per gigabyte well below that of any alternative SCM approach that is based on an emerging memory technology.

Netlist has combined established and well-understood technologies, DRAM and NAND flash, with innovative embedded intelligence to provide SCM's speed and cost attributes. Should this approach be readily accepted, HybriDIMM can pave the way for adoption of SCM well ahead of any emerging memory's actual volume production. This could be a game changer for these emerging memory technologies, since the market will be established and ready to accept them well ahead of volume shipments. On the other hand, broad acceptance of the HybriDIMM could serve to slow adoption of emerging memories due to the lack of a compelling reason to convert.

It makes good sense to use high-volume, established components as a substitute for unproven technologies that are not yet in production. This brings an immediate solution to a problem that would otherwise have to wait an uncertain amount of time. It's somewhat similar to the way engineers create prototypes of projects using field programmable gate arrays (FPGAs) as a substitute for the application-specific Integrated circuits (ASICs) their product will use in volume, although the use of FPGAs instead of ASICs comes at a steep cost differential that is not the case in this instance. The largest difference between a NAND + DRAM approach and one based on an emerging memory technology is likely to be one of write speed. However, for the many analytic workloads that are read intensive, this is unlikely to be a critical issue.

What Drives SCM Acceptance?

SCM is also likely to bring change to the way servers are used to manage data, and may even result in a reduction in server count, just as SSDs have done in the past. This is a very difficult factor to estimate and was not included in the forecast, but here's how it works: In many cases slow storage drives system administrators to split their database into "shards," smaller subsections that can be acted upon in parallel. This is a very inefficient process, and requires the use of a number of servers to simultaneously process each of these shards to determine the final result. Not only are costs increased through all this additional hardware, but software licenses may also be required for these additional servers. By using faster storage (in this case storage that runs at memory speeds) IT managers have been able to avoid breaking up their databases into smaller portions that run in parallel across several servers, and can instead run the database on a single server with faster storage. This approach has bolstered the growth of SSDs, and is likely to have a similar impact on SCM.

Another very positive benefit of SCM will be one that also propelled SSDs to success, mainly in OLTP systems, but also in the field of high-frequency trading. In these applications, two of the first to embrace SSDs, the relationship between the use of an SSD and a financial return was very easy to measure. OLTP users increased their transaction volume at a minimal incremental cost. High-frequency traders were able to respond to market changes even faster than they could with HDD-based systems, reaping rewards that were several times as large as the cost of the SSD. This same dynamic is very likely to play out with SCM since it serves to improve performance for a relatively modest price.

We briefly mentioned "caching and other sophisticated algorithms" above. Although Netlist has not disclosed the technology that manages its internal data management, the company is proud of the fact that it can move data into and out of NAND flash in a way that does not interfere with the DDR4 DRAM interface. This is very important to DRAM module users, because it removes any need for exotic support that would require modifications to the system, like bus modifications or BIOS changes.

The data is moved around autonomously within the DIMM by on-board processors. It is extremely important that none of this interferes with the DDR4 interface because the DDR4 bus does not support a delay mechanism – all data transfers must occur at the same speed. The company calls this "PreSight" technology, but otherwise keeps the operation secret. The bottom line is that a HybriDIMM can be plugged into any existing socket that accepts a DDR4 LRDIMM.

Netlist tells us that the application processor opens up new potential performance enhancements. The company is offering its customers deeper management of the HybriDIMM's internal speed than would be available through industry-standard NVDIMM software protocols. This protocol, called AppDirect, will be supported by a system development kit (SDK) to put the faster protocol into the hands of application developers who can use it to squeeze the highest performance out of their designs.

Conclusion

In this paper we have explained how storage class memory came to be, and have detailed the criteria that will lead to its eventual adoption: That it must be faster than SSDs and less costly than DRAM memory.

SCM's persistence (the fact that data remains intact even in the absence of power) was then addressed, with an explanation of how persistent memory can bypass many of the delays in database systems that are designed to assure data integrity whether or not consistent power is available. Persistence is a highly-desirable attribute, but little of today's commercially-available software can support persistent memory. Industry standards have recently been defined to help developers create SCM-aware software, and this will open other significant opportunities for SCM in the future.

Even without software support, the market for SCM could develop rapidly, and could ramp to nearly $2.5 billion by 2019, provided that it meets with broad and unfettered adoption by all prospective users in the server community. Although such rapid acceptance is unlikely, a smaller fraction of that number would still be a significant market.

Finally we discussed an alternative type of SCM that has been developed by Netlist, one that uses a small DRAM and a large NAND flash, both industry-standard high-volume technologies to meet the SCM cost/performance ratio that other companies plan to provide through the use of emerging memory technologies. This alternative, which Netlist has named HybriDIMM, uses commodity DRAM and NAND flash managed by an on-board processor. HybriDIMM supports near-DRAM speeds at below-DRAM prices in a very high density module that is fully compliant with an industry-standard DDR4 LRDIMM memory bus interface.

Netlist plans to use this module as a way to bring the attributes of SCM to users well ahead of the actual availability of those storage class memories that will be based on emerging memory technologies. This approach is likely to be met with good acceptance, since it provides in a timely manner a technology the industry already needs.

The HybriDIMM could even prove to be a hurdle that will get in the way of emerging memory technologies once these new technologies are finally brought to market, since they now must compete not only against NAND flash SSDs and DRAM, but also against inexpensive SCM modules based on DRAM and NAND flash.

Jim Handy*, August 2016*