



OBJECTIVE ANALYSIS

Semiconductor Market Research

NAND CACHE IS BACK! SSD Performance at an HDD Price

OBJECTIVE ANALYSIS WHITE PAPER

The PC market is about to enter a new architecture phase. Although all PCs sold to this time have used DRAM and HDD as adjacent layers on the memory hierarchy, NAND is now finding a natural fit between these two technologies, where it provides performance that is faster than HDD but slower than DRAM at a cost per gigabyte that is cheaper than DRAM but more costly than HDD.

This paper will explore this development, explaining some technical and economic details, giving reasons why it will soon become standard, and showing its impact on the NAND market's growth.

Why SSDs Didn't Happen

It may come as a surprise to some readers that NAND SSDs will never completely replace the HDD in the PC. After all, haven't NAND makers been telling us for a few years that HDDs are doomed, and would soon vanish from the face of the earth? Weren't they constantly showing us forecasts with tens of millions of SSDs shipping by 2009, as is shown by the columns in Figure 1?

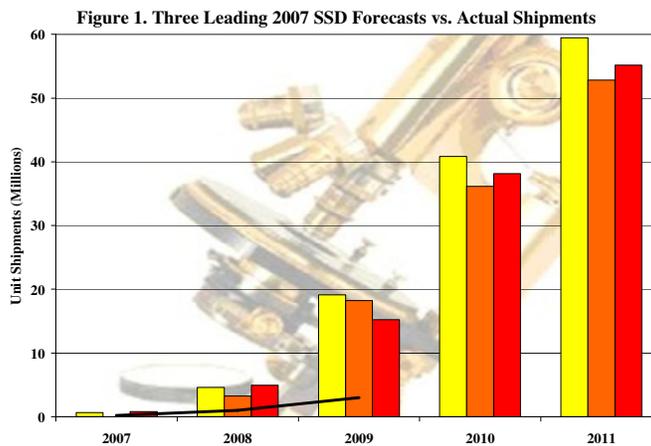
Since 2004 NAND flash companies and other SSD makers have been devoting an enormous amount of energy in an attempt to drive demand for SSDs among PC OEMs and PC buyers. Despite a laudable effort, penetration is still below 1%. Figure 1 shows actual shipments as a black line that overlays the forecasts' columns.

What went wrong?

The unfortunate fact is that consumers are not excited about the price-performance tradeoff that SSDs bring to the PC.

SSDs are simply too expensive or too small to be used as the only storage in a PC, and the current thrust of SSD vendors is to encourage the user to do exactly that. Benefits of faster boot times, longer battery life, and improved ruggedness fade in comparison to the sheer volume of storage available on an HDD at a small fraction of the cost.

The net result of years of SSD vendor marketing efforts has been a penetration of less than 1% of PC sales, with SSD sales stemming mainly from retrofits of existing PCs rather than sales into new PCs. It is generally accepted that



2009 shipments were well below ten million units.

Still Hope for NAND in PCs?

There is an alternative to completely replacing a PC's HDD with an SSD. One such approach that showed some promise was to add a smaller capacity of NAND to serve as a cache between the HDD and DRAM layers of the memory hierarchy. This hybrid approach would provide near-SSD performance at a price that is only a fraction that of an SSD alone. The idea is to leverage the low cost-per-bit of the HDD for capacity, and leverage a smaller amount of NAND for performance.

Caching is a well understood technology already used in PCs where expensive and high-performance SRAM is used as a cache by the processor to avoid the accessing the slower DRAM for data. A great deal of attention is paid to the cache algorithms to determine which data might be "reused" and thus might be a good candidate to store in the cache.

The same approach is used to fit a NAND cache between the DRAM and the HDD – caching algorithms ensure that the most frequently used and most recently used data is stored in the faster NAND cache to avoid accessing the slower HDD for data.

Three key criteria must be considered when choosing flash cache solutions:

- 1) Do the algorithms correctly identify the most frequently used and most recently used data?

- 2) Is there enough NAND capacity to store that data?
- 3) Is the NAND performance significantly greater than that of the HDD?

Ideally the caching algorithm would identify most of the key data and store it in the NAND cache, yielding the SSD-like performance of the NAND cache most of the time. The user would rarely notice those times that the system accessed data from the HDD (which is actually quite fast thanks to today's 7,200RPM consumer drives)

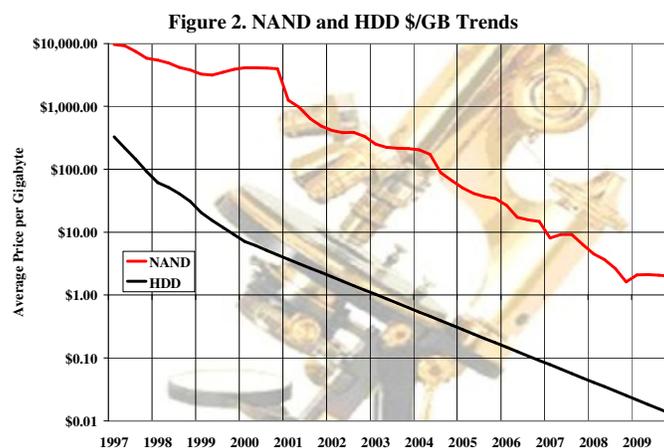
Flash Cache Successes & Failures

Flash caches have been used successfully in a number of larger computing applications, but they have failed to penetrate the PC. A closer look will explain why.

Efforts to place NAND between a PC's DRAM and HDD include the Intel Robson or Turbo Memory, Intel's aborted Braidwood¹ program, the early hybrid HDDs offered by Samsung and Seagate, various attempts to introduce secondary SSDs in addition to an HDD

as was the case with SanDisk's Vaultor Disk, and even Microsoft's ReadyBoost technology which attempted to accelerate the Vista operating system through the simple addition of an external

USB flash drive. The PC market has yet to find success in OEMs' at-



¹ Intel's Braidwood: Death to SSDs? Objective Analysis, August 2009

tempts to insert a NAND layer between the PC's HDD and main memory.

With the exception of ReadyBoost, all these solutions were designed to use a NAND that was extraordinarily small in comparison to the HDD it attempted to accelerate, and was usually smaller than the system's main memory. Typically a successful new layer in the memory hierarchy will be several times the size of the next faster layer.

While the size of the NAND layer may have been one reason that many of these approaches failed to bring a meaningful performance improvement to the system, we understand that inadequate support software was also responsible.

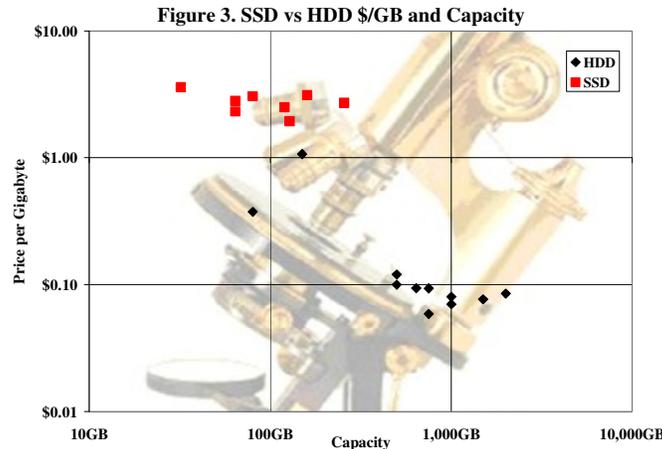
How Expensive is NAND?

The big drawback to replacing HDDs with SSDs in a PC is that an SSD costs so much more than an HDD yet it brings very little in return.

The relative price differences between HDD, NAND, and DRAM pricing provide significant motivation to find the right balance of these three technologies to optimize price and performance. A convenient rule of thumb for PC capacities today is that DRAM costs about \$20.00/GB, MLC NAND \$2.00/GB, and HDD about \$0.20/GB.

The price dynamics of NAND and HDD pricing is illustrated in Figure 2. NAND flash price per gigabyte (the upper line) is about 20 times the price per gigabyte of an HDD. Over time, as NAND prices have declined, HDD prices have declined at the same rate, keeping the 20:1

ratio intact. This trend is most likely to continue into the foreseeable future. This means that users will always get significantly more storage out of a \$50 bargain HDD than they will from a \$200-\$300 SSD.



When purchasing a new PC buyers generally use the PC's price and its HDD capacity as two of 4-5 key metrics. Because of this it is no surprise that very few PC purchasers have (or will ever) opt for an

SSD instead of an HDD.

How does the price of NAND influence the price of an SSD? As a simple example we collected figures on HDD and SSD prices from two on-line sources: Fry's Electronics and NexTag. These appear in Figure 3. The chart presents the data as a scatter plot of capacity (horizontal axis) and price per gigabyte (vertical axis.) The clustering of SSD and HDD data points is very clear in this chart, with SSD prices (red) averaging \$2.75/GB with an average capacity of 113GB, and HDDs (black) priced at an average of \$0.12/GB with an average capacity of 806GB. The SSDs in these two websites were priced an average of 24 times as much as the HDDs while the HDDs had an average of seven times the capacity of the SSDs.

How Much NAND?

Perhaps the trickiest question is: "What is the right amount of NAND to put into the system?" This can be a little hard to answer, since it depends on a number of variables, including the type of software being used,

the size of the main memory, the value of the speed improvements to the user, and even the user's computing habits.

Figure 4 is based upon the same SSD and HDD pricing data that appeared in Figure 3. In this chart we present the cost of the drive as a function of its capacity. The SSDs make up the cluster of red points toward the top center of the chart while the HDDs appear further right, where capacity is high. We drew two lines through these points to illustrate how price and capacity relate to each other.

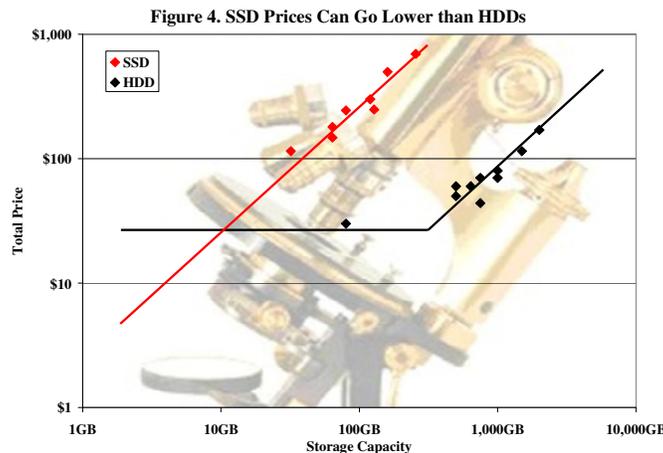
Since SLC NAND flash can still be purchased in chip densities as low as 256Mb (one quarter gigabit, or 32 megabytes), then the line moves straight from the highest capacity SSD in this survey, at 256GB and extends well below the 1GB level at the far left side of the chart. HDDs, on the other hand, have a "floor price" somewhere between \$30-\$50 – it's impossible to make and sell an HDD economically below the floor price. It follows that SSDs can be made cheaper than HDDs to meet a storage budget that is lower than the floor price for an HDD.

This puts some perspective on the problem. It would seem reasonable in many systems to add an amount of flash that costs about the same as the system's HDD. In more cost-sensitive systems the number might be lower than that, and in more performance-sensitive systems the number could be higher. Based upon the data in Figure 4 an SSD of 11GB would have the same price as the floor price of an HDD. The capacity of this crossover point increases over time as

both HDD and SSD prices continue to fall following the trends shown earlier in Figure 2.

In all cases we believe that the price of a NAND cache is likely to range from 20% to 200% of the cost of the HDD. This implies that the size of the cache will range from 4% to 40% of the size of the HDD it accelerates.

We should note that Seagate had just introduced a hybrid HDD as this report was being finalized. The Seagate Momentus XT combines a 4GB NAND read cache and a 32MB DRAM write cache with a 7,200RPM HDD to offer performance that is claimed to approach the speed of an SSD in most cases.



Managing the Cache

It's no easy feat to manage a NAND cache. Manual efforts to decide whether certain files should reside on the SSD or the HDD are beyond the scope of even most PC enthusiasts. Furthermore, such manual methods can become very complex, and are generally ineffective in meeting the goal of improving overall system performance.

One example of such an effort is boot-drives, a new category of small SSD. Boot drives are small, affordable SSDs designed to be paired with an HDD to store a static configuration of the PC's boot image, the operating system, and maybe even certain software applications

While this sounds like a reasonable way to team an SSD with an HDD, in reality it can be tricky to manage.

The rationale for using a boot drive is to achieve higher system performance at a low incremental cost by using only a very small amount of NAND. The smaller the NAND is, the more judgment must be used in managing its contents. To squeeze the NAND down to the most economical size the NAND's contents need to be updated to match the user's work habits which will vary over time. The NAND's contents will also need to be modified whenever there are any updates to the system software.

Hand tuning requires an understanding of which files are in high demand and which ones are seldom used. Even few experienced IT managers understand their systems this well, and the average PC user is unaware that any files exist other than the ones that the user created or downloaded.

Fortunately, this process can be automated. Tiered storage management systems use algorithms that adapt to user behaviors and move frequently-needed data into NAND to optimize efficiency, reliability, and performance, the same way that the cache is managed inside the processor chip. Several companies are either developing or shipping software products that perform this management for enterprise systems. The only currently-sampling storage management software for the PC that we are aware of is NVELO's Dataplex product, which

manages an SSD or other flash cache used in conjunction with an HDD.

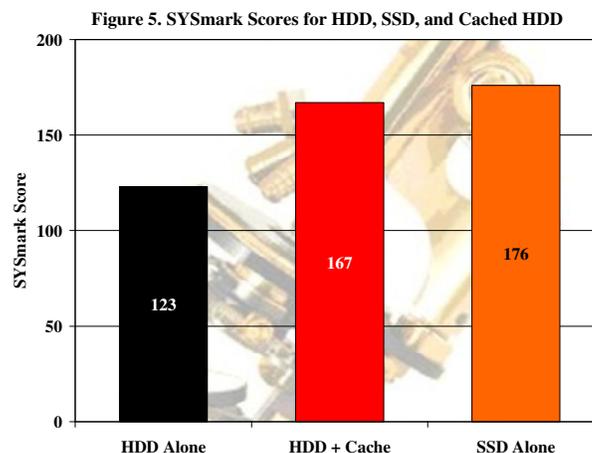
NVELO Dataplex

Although NVELO Software is providing SSD vendors and PC OEMs with early versions of its Dataplex cache software, most benchmarks are covered by strict nondisclosure agreements and are thus unavailable to Objective Analysis for publication. NVELO has given us some limited benchmark data that the company generated by running the Dataplex cache management software on an off-the-shelf PC¹. This information is shown in Figure 5.

NVELO ran the SYSmark 2007 Preview test using the same PC but with three different storage configurations:

- 1) The production 7,200 RPM SATA II HDD,
- 2) A 32GB Intel Extreme SSD,
- 3) The combination of the production HDD and a 16GB cache. The cache was managed by NVELO's Dataplex software.²

While the HDD-only system had SYSmark score of 123, or 69% that of the SSD-based system, the cached system



¹ 2.80 GHz Intel® Core™ 2 Duo Processor T9600, Intel PM45 Express Chipset, 4GB 800-MHz DDR2 DRAM, ATI Mobility Radeon HD 3650 with 256MB VRAM, 320 GB SATA II 7200 rpm HDD.

² The 16GB cache used the Intel SSD but limited the available capacity. The benchmark for this case was run multiple times to allow the cache to learn the system's typical storage patterns and optimize the contents appropriately.

ran at 94% the speed of the SSD-based system, attaining 83% of the performance improvement that the SSD system provided.

Boot-up time was also improved. The Dataplex system booted equally rapidly as the SSD system quite simply because Dataplex automatically caches the boot image.

Although there is a wide range of benchmarks to choose from NVELO selected the SYSmark benchmark suite. SYSmark is generally accepted as a good proxy for real-world system performance since it installs standard software applications, like Microsoft Office, Adobe tools, CAD, video editing suites, etc., and then automatically recreates actual usage scenarios including opening, editing, and saving files.

NVELO's goal was to increase overall system performance. The company claims that most PCs can achieve between 85% - 95% of the performance of an SSD by teaming Dataplex software with a NAND cache as small as 16GB.

A 16GB SLC NAND cache costs approximately \$75 today, so this performance boost would increase the cost over a standard PC by \$75 plus the yet-undisclosed cost of NVELO's Dataplex software. Assuming that the combined cost is less than the \$200+ adder normally associated with upgrading a PC to an SSD, a flash cache looks to have a lot going for it. NVELO has explained that Dataplex will initially be available only through OEMs.

NVELO's Dataplex provides one novel feature that should be comforting to users. SSDs are known to suddenly shut down because of wear. What is really happening is that the operating system has not been designed to accommodate storage volumes whose capacity decays over time: Once the SSD's capacity drops below the size the software expects to see, the operating system de-

terminates that it has failed and stops accessing it. NVELO's Dataplex software has been designed to support graceful decay of the cache, converting the device to read-only mode only when the NAND has reached the end of its life.

Like many SSDs, Dataplex includes a wear reporting mechanism that allows the system to monitor the health of the cache and perform diagnostics so that preventative action can be taken before difficulties arise. We have been told to expect to see OEMs guarantee operation of a 16GB SLC NAND cache for 5 years.

Caches Use SLC NAND

It is reasonable to explain here that PC NAND caches typically use SLC NAND flash because a NAND cache's traffic tends to be higher than that experienced by the NAND in a typical SSD. Also, a NAND cache is intentionally kept completely full over the course of its operating life, while an SSD usually has a significant amount of unused space that its wear leveling algorithm can borrow to improve endurance. In order to reliably accommodate this higher level of traffic and larger wear load, SLC flash is the most sensible alternative.

Although MLC flash would drive the cost of a 16GB cache significantly lower, to about half the cost of the SLC cache, the system would have to compensate for the lower wear rating of the MLC NAND by over provisioning – that is, using more NAND than is absolutely required. If the system were over provisioned by two times, there would be no financial benefit in using MLC flash.

Impact on the NAND Market

The main reason that NAND chip manufacturers have been so eager to grow the SSD market is because the potential market for NAND bits would be so enormous. Hundreds of

millions of PCs ship every year, and an SSD typically uses ten or more times as much NAND as does a flash card, USB drive, smart phone, or music player.

Just how large is the market potential for SSDs compared to that of standard NAND? In 2009 6.7 Exabytes (millions of terabytes) of NAND flash shipped into all applications. The consumption of HDDs in the notebook computer market was 24.2 Exabytes or more than 3½ times as much. Even a modest 10% share of the notebook PC market would give a significant 35% boost to the NAND market.

Since the past few years have shown us that a 10% penetration is not likely to occur, then what of NAND caches? Well, a low-cost NAND cache is significantly more likely to penetrate the PC than an expensive SSD, especially if it can yield 80% or more of the performance of an SSD for only a fraction of the added cost.

In Figure 6 Objective Analysis compares a reasonable penetration of NAND cache in the PC versus likely adoption of SSDs into the PC (based on historic adoption patterns, and assuming that NAND cache did not compete.)

The PC adoption figures for NAND cache were originally forecast in an Objective Analysis report on Intel's aborted Braidwood initiative¹. Braidwood was postponed immediately prior to its early 2010 introduction. We have delayed our original forecast for this technology by one year to account for the fact that there are other non-Intel solutions poised to

fill this market void that should start appearing by early 2011.

By 2013 our models indicate that 6.1 exabytes of SLC NAND will ship into the PC NAND cache market, 90% as many exabytes as the entire NAND market of 2009! This will add nearly \$7 billion to NAND revenues in that year. SSDs, on the other hand, would have only driven \$1.2 billion of NAND revenues representing 2 exabytes of MLC NAND. Note that these SSD figures assume adoption rates that follow historic patterns, and do not account for the fact that NAND caches are quite likely to undermine the market for SSDs in PCs.

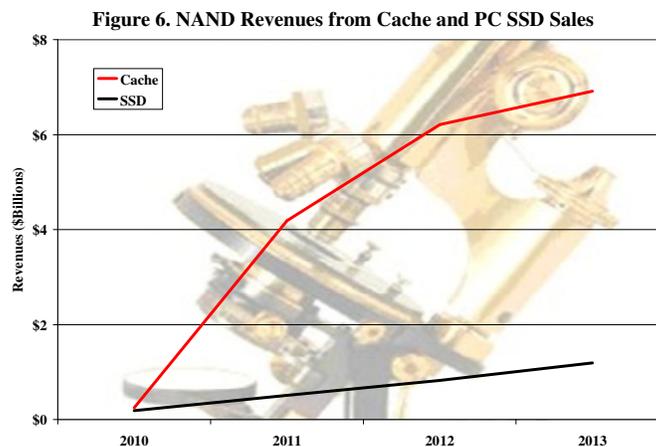
It is quite clear that NAND flash manufacturers would do well to prepare to take advantage of the inevitable adoption of flash cache in the PC, rather than to pin their hopes on an SSD market that shows no signs of springing to life anytime soon.

In fact, Objective Analysis has learned that a number of leading OEMs have entered into relationships with companies who are developing NAND caching technologies. These OEMs are preparing themselves

to fill the void left behind when the Braidwood initiative was halted.

When such products are finally introduced, Objective Analysis plans to detail

them in a report to keep our clients fully apprised of the status of the NAND cache market as it evolves.



In Summary

By now the reader should understand that a market for NAND caches will inevitably develop, and that multiple contenders are already getting ready to participate. This market will prove to be more significant to overall NAND revenues than the SSD market, and the caching technology will provide near-SSD performance to PC users at a cost that is only incremental to the overall cost of the PC.

The NAND cache market will need to follow three rules to achieve success:

- NAND caches must be sized to fit well within the memory hierarchy, remaining smaller than the HDD (to manage costs) yet larger than the main memory's DRAM.
- NAND caches must be designed for the task, delivering both performance and reliability from a small capacity.
- Sophisticated management software will be required to allow the cache to fulfill its role and provide near-SSD performance levels at a cost that is significantly lower than that of an SSD.

Furthermore, the reader should understand why SSDs have failed to displace HDDs, and that the demise of the HDD is unlikely for the foreseeable future, since HDDs will continue to provide relatively fast storage at a very affordable price.

Jim Handy, June 2010

