Solid State Drive Performance

White Paper
David Bartizal & Thomas Northfield
March 24, 2008

Abstract
To those in the computer industry, it's no surprise that each computer component -- from the CPU to the RAM, the video GPU to the storage systems -- needs to be faster to keep up with user demands for their computing platform to do more. Almost every day, there is a new announcement about a faster this or a more efficient that. In addition, as the cost of energy and concern for the environment increases, there is a recent mainstream trend toward lower energy usage in computing.

Mass storage devices are no exception to this demand, but in many ways have not kept up with other computer components in terms of speed and energy efficiency. While capacities continue to increase, in recent years, it appears that hard drive technology has reached a point where higher speed and lower energy usage are at odds and only slight improvements are making it to market. For the user who needs high performance and wants energy efficiency, there is a technology that is finally available, solid state drives.

Solid state drives (SSD) have performance attributes that are considerably better than hard disk drives (HDD). As one would expect, those attributes depend on many things. For example, an SSD's overall performance can depend on the design and quality of the SSD, the application being used and also on other hardware in the system. The primary performance enhancement of an SSD over an HDD is extremely fast random reading from SSD. In several early attempts to provide an SSD, this improvement in random access was gained at the cost of extremely slow random write speeds. This is no longer true and products are now available with really fast random read speeds and with much improved random write speeds.

This white paper identifies the performance attributes of SSDs and provides a comparison to typical HDD performance. In addition the dependencies on application and some particular hardware are explored.

The SSD
An SSD has no moving parts and is essentially an HDD emulator. It is basically comprised of a printed circuit board, a set of NAND flash memory chips, SDRAM cache, a memory controller, an interface controller and an interface connector such as IDE, SATA, SAS or even fiber channel. The drives that are discussed in this paper are all SATA connected devices. In addition to the performance and energy advantages of SSD, the lack of moving parts means that the drives can withstand high vibration and shock forces. In fact, some SSDs can withstand shock force up to 1500G, the equivalent of a drop from 26 feet.

Definitions of Performance Measurements

There are many different ways to measure the performance of a storage device. Key parameters used in this paper are defined here for reference.

Access Time - The time a program or device takes to locate a single piece of information and make it available to the computer for processing. Access time is typically measured in milliseconds (ms).

Sequential Transfer Rate - The amount of data that the device can read or write to adjacent sectors of the storage media in one second. Sequential transfer rate is typically measured in megabytes per second (MB/s).

Random Transfer Rate - The amount of data that the device can read or write to non-adjacent sectors of the storage media in one second. Random transfer rate is typically measured in megabytes per second (MB/s).

Sequential IOPS - The number of input / output operations that the device can complete on adjacent sectors of the storage media in one second. Expressed as sequential read or write inputs / outputs per second (IOPS).

Random IOPS - The number of input / output operations that the device can complete on non-adjacent sectors of the storage media in one second. Expressed as random read or write inputs / outputs per second (IOPS).

All of the transfer rate and IOPS performance measurements also are dependent on the size of the transfer block.

Typical sequential operations include reading a non-fragmented large file (such as a video file) or writing a large file to a contiguous section of the media. When severe fragmentation occurs or when reading or writing includes lots of jumping from one small file to another, random operations are occurring. Loading application or operating system files is often more of a random operation than sequential. In all cases of real world computing, the actual performance of an application is a combination of random and sequential operations.

About the Drives Tested

HDD and SSD drives were selected for testing with various performance attributes. Table 1 summarizes these drives.

Name	Type	Capacity	RPM (if HDD)	Form Factor
WD740	HDD	74GB	10,000	3.5"
WD2500	HDD	250GB	7,400	3.5"
HTS541040	HDD	40GB	5,400	2.5"
SD 5000	SSD	32GB	n/a	2.5"
Imation 7035*	SSD	32GB	n/a	3.5"
Imation 3035**	SSD	32GB	n/a	3.5"
Imation 7025*	SSD	32GB	n/a	2.5"
Imation 3025**	SSD	32GB	n/a	2.5"

Table 1

*Imation SSD PRO 7000 powered by Mtron
**Imation SSD MOBI 3000 powered by Mtron

Access Time, SSD vs. HDD

An HDD contains one or more platters of magnetic media and a read write head that moves on an actuator from inside to outside diameters on the platter surface. To locate and read a particular piece of data, the heads must find and move to the correct location and then wait for the spinning platter to present the physical data. This typically takes 10 ms or more. Some faster HDD models are capable of 7 or 8 ms access time. Due to its solid state components and the ability to address any sector of the NAND memory directly (instead of seeking it), an SSD can access the data in about 0.1 ms or about 100 times faster than an HDD. When this is done tens of thousand or hundreds of thousand of times to complete an operation such as boot up, the user easily recognizes the delay caused by the HDD. Figure 1 shows the access time comparison between Imation SSDs, a commonly available competitive SSD and a variety of HDDs. The much faster SSD times are obvious (shorter bar is faster).

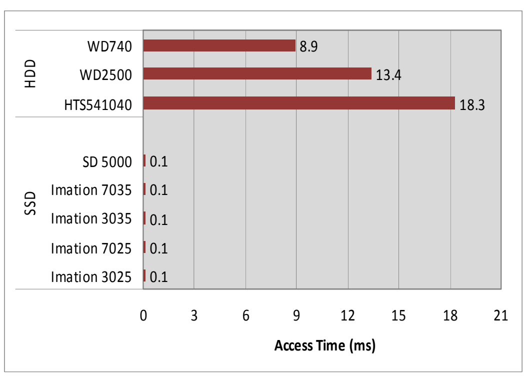

Figure 1

Sequential Transfer Rates, SSD vs. HDD

After the typical computer boots up and software is loaded into RAM, a user may read or create some data file. Eventually, a sequential write is required to save the new or modified file to non-volatile storage such as an HDD or SSD. Any data file such as a document, spreadsheet, presentation, picture or any other data file the user creates will take a finite amount of time to save. Figures 2 and 3 show a comparison of the write and read sequential transfer rates of SSD and HDD devices.

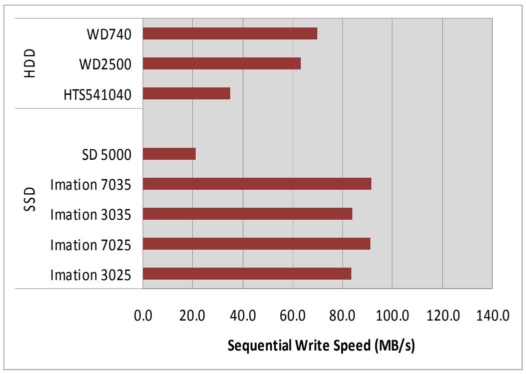

Figure 2

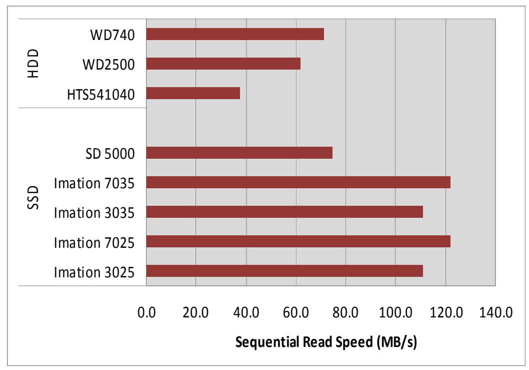

Figure 3

Note that the HDDs have similar read and write speeds while SSDs have read speeds that are faster than the write speed. This is due to the SSD erase process that must precede the write process. In spite of this "write penalty," it is still possible to have an SSD that exceeds the read and write performance of HDD, as seen in the charts above.

File or block size is another dependency. As the block size increases from the very small (512 bytes) to the large (128KB and greater) and on to the very large (8MB), the performance of the drive improves. This is because of the operating system and device overheads that are associated in preparing to read or write a file. Figures 4 and 5 show some examples of how the sequential write and read transfer rate changes as block size changes for SSD and HDD.

Figure 4

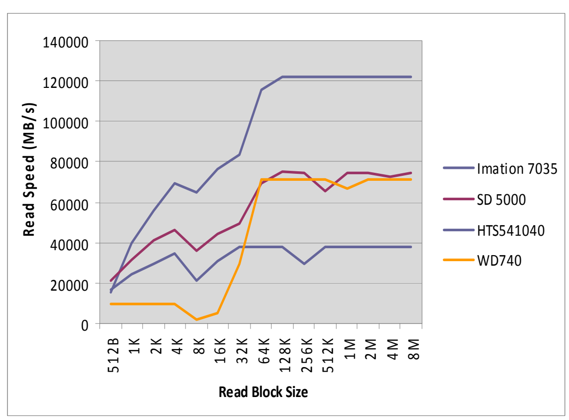

Figure 5

Note that at some point for each drive type, the block size no longer dominates the transfer rate and this is the maximum read or write transfer rate for that drive.

Random IOPS, SSD vs. HDD

In the case of transferring data to or from a highly fragmented drive, or in the case of random access to database files, random transfer speeds are most important. Random speeds are typically lower than sequential speeds because of the overhead associated with moving from one block on the device to another. While it is possible to measure the random transfer rate, it is not as useful as looking directly at the random IOPS. As with sequential transfer rate measurements, random IOPS are dependent on the block size used for the transfer. Figures 6 and 7 show the random write and read IOPS for SSD and HDD.

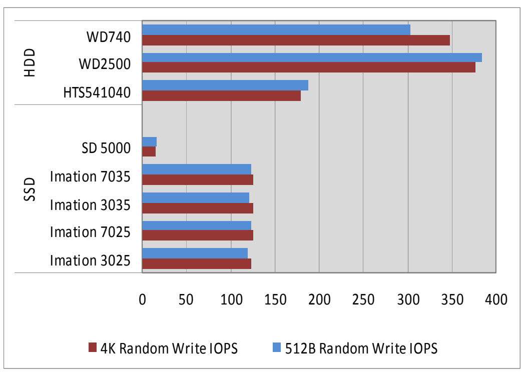

Figure 6

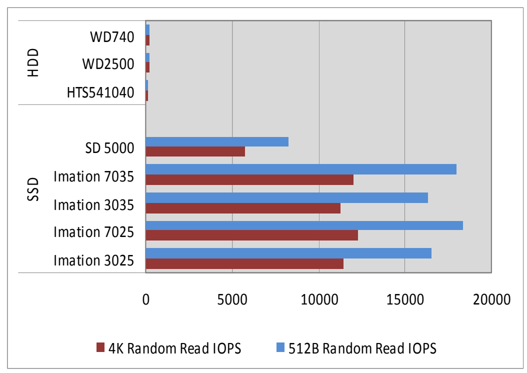

Figure 7

While the write IOPS for SSD is 1/2 to 1/3 that of HDD, the random read IOPS is 100X faster. It is typical that an application does much more random reading than random writing, so this performance enhancement translates into real world speed increases for the user.

Benchmarking / Real World PC Performance, SSD vs. HDD

There are many ways to benchmark the performance of a PC system and to measure the impact that changing one component has on overall system performance. One widely accepted method is to use PCMark05, from Futuremark. PCMark05 provides overall system score benchmarks (called PCMark) as well as capability for a storage only score (called HDDMark). By taking a PC and replacing the HDD with an SSD and running the benchmark, a good picture of real-world improvement can be seen in both scores. The test HDD and SSD drives were mirrored from a master HDD to assure exactly the same system setup. Figures 8 and 9 show these scores.

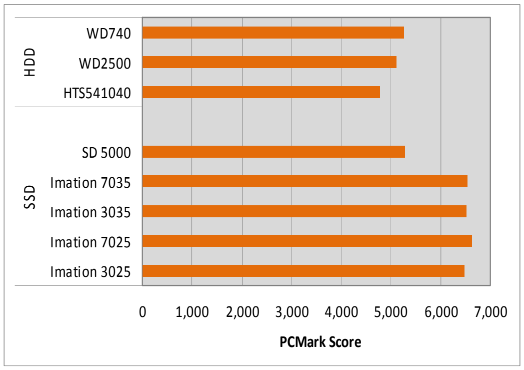

Figure 8

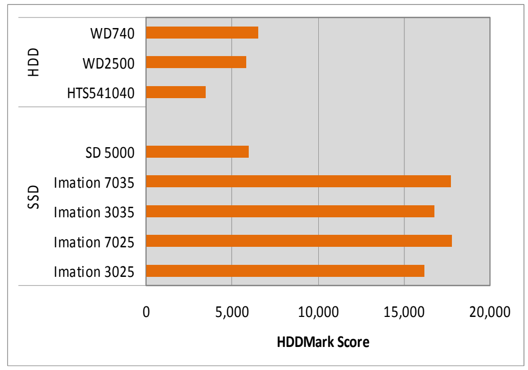

Figure 9

In all cases, the SSD provides for a large improvement in system performance.

PCMark05 also provides a breakdown of the task speeds associated with reading and writing to / from the HDD or SSD. These speeds are shown in Figure 10 with SSD and HDD examples.

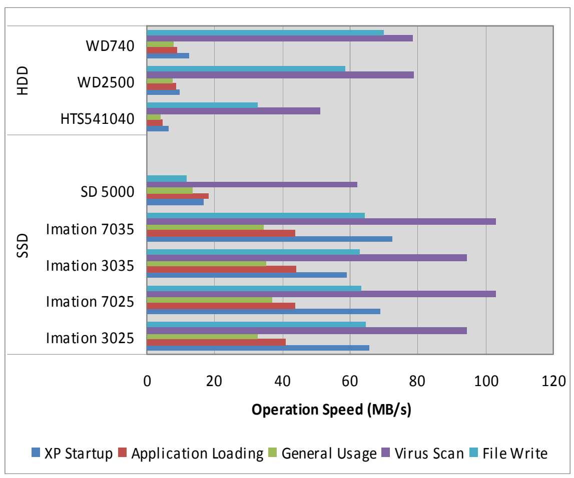

Figure 10

It is again clear that for most cases, the SSD is capable of a performance improvement.

Finally, one may take a PC and compare the actual time it takes to do various tasks. Examples are shown in Figure 11. Note that actual times are very dependent on all of the components of the PC system. The times shown below are on the same PC, (the test HDD and Imation Pro SSD drives mirrored from a master HDD to assure exactly the same system setup).

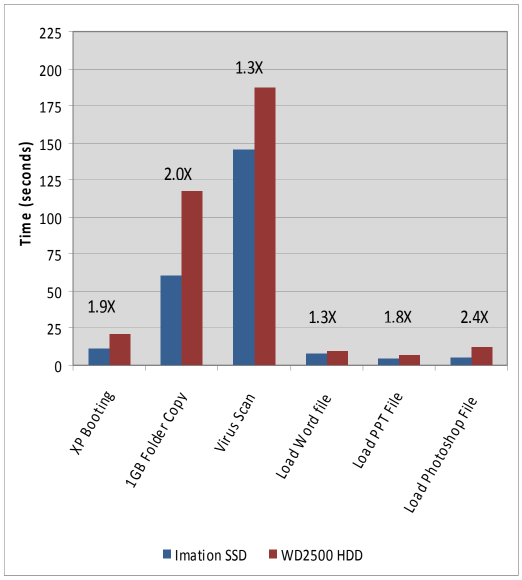

Figure 11

Dependence on Storage Sub-system / Channel

The performance data shown in this report were collected on a Desktop PC that used an NVidia nForce 570-Ultra Chipset. The actual performance that an individual can expect on a given platform is highly dependent on the bandwidth of the chipset that is used for the SSD SATA connection. This is especially true for the south bridge of the chipset.

Notable examples of chipsets with limiting bandwidth are the Intel ICH7R (commonly found on Intel 975X boards) and the Intel ICH9R (commonly found on Intel X38 boards). There may be other examples of limiting chipsets and there may be examples of chipsets with even better bandwidth than the nForce 570. Shown below in Figure 12 are example PC chipsets and the performance they are capable of for sequential read transfer rate.

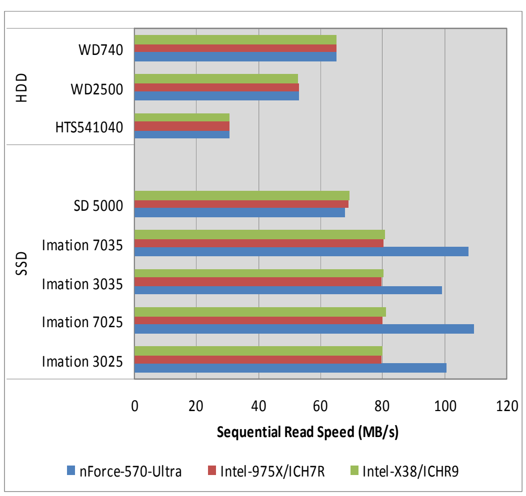

Figure 12

Note that as of the date of this writing, the Intel chipsets are bandwidth limited to around 80 MB/s, whereas the NVidia nForce chipset is not. It is expected that this will be fixed in later releases of the Intel chipsets. As a workaround for owners of Intel chipsets, there are a growing number of reasonably priced add-on SATA controllers that do not suffer this slowdown.

Power Draw

SSD has a much lower power need than an HDD. Lower power also means less heat generated and so has indirect impact on energy usage by reducing the energy needed for cooling.

For a high performance HDD (such as the WD740 10K RPM drive), the power requirement (not including additional savings for reduced heat) for the Imation SSD is:

Idle	Approx. 15% of the WD740
Peak operation	Approx. 30% of the WD740

Figure 13 shows the large reduction in energy needs for an SSD under various operating conditions compared to various HDDs.

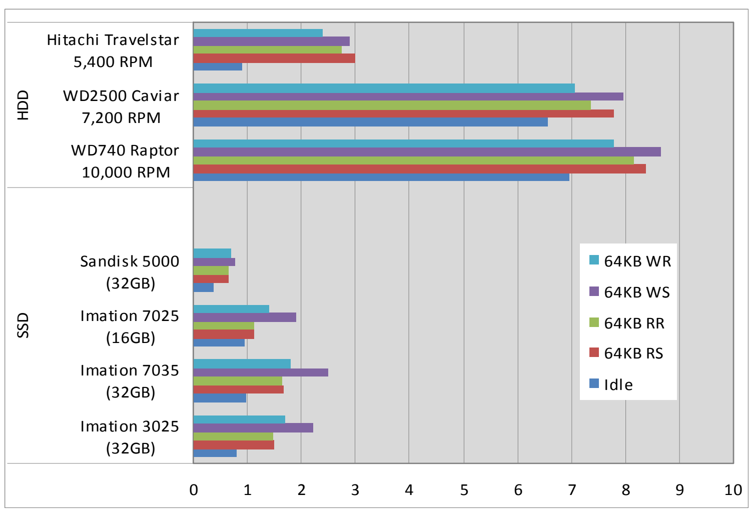

Figure 13

To illustrate the importance of this, along with the benefit of high random read IOPS, we can model the energy savings of the Imation 7035 SSDs vs. enterprise class 15K RPM HDD. Below is an example where an application needs 36,000 IOPS.*

Example Enterprise Class HDD	360 IOPS / Drive (36,000/360 = 100 drives)
Imation 7035 SSD	18,000 IOPS / Drive (36,000/18,000 = 2 drives)

At an approximate 15 watts power draw, an HDD consumes 130 kWh per year. This costs $19.50 per drive per year (assuming $0.15 / kWh). An Imation SSD on the other hands draws only 2.5 watts of power and consumes just 22 kWh per year. This cost is $3.30 per drive per year.

Total HDD Energy Cost	$1,950 per year
Total SSD Energy Cost	$7 per year
===	
SAVINGS	$1,943 per year

* Actual savings are impacted by an application's specific IOPS need. Factors that affect this need include (but are not limited to): Mix of random and sequential access types, the mix of read and write operations, the block size or block size distribution, the type of HDD being considered in the comparison, the actual power draw and IOPS of the HDD and the current cost of energy.

Tools Used to Evaluate SSD Performance

The user may want to determine if they have a component in their system that would limit the speeds that may be realized by SSD. The test conditions on all SSD vs. HDD tests must be well understood. The SSD and HDD must be the slowest component to realize the full potential of each device. SSD and HDD test speeds are dependent on the CPU and motherboard, video card, chipset and any component that may slow down the system. There are many software test suites. Popular examples are shown in Table 2.

Test Software	Test Parameters
HDTach 3.0.1	Sequential sustained and burst read transfer rates (variable zone) and access time
ATTO 2.4.1	Sequential sustained write and read transfer rates vs. block size
IOMeter 2006.07.27	IOPS under different workloads -- includes random/sequential and read/write for various transfer sizes
PCMark	Completes a variety of tests simulating PC usage and provides an overall score relative to other computer systems
SiSandra	Various transfer rate tests including read/write operations/minute based on file sizes
HDBench	Sequential and random read/write speeds

Table 2

Conclusion

Overall performance is limited by the slowest component in the system. In modern computers the HDD is often a bottleneck for data intensive applications. With fast and energy efficient SSD drives now available, this bottleneck can be relieved, overall performance and responsiveness of the PC greatly improved and in cases of high IOPS needs, energy and money saved.

David Bartizal is senior manager of Qualification and Quality Assurance at Imation's Outsourced Products Laboratory. Thomas Northfield is a product engineering specialist at the Outsourced

Products Laboratory. The laboratory is located at Imation's worldwide headquarters in Oakdale, Minn.