Data Storage Technologies

Transition to solid state

Ari Geir Hauksson and Sverrir Ásmundsson
Reykjavík University
Kringlan 1, 103 Reykjavík
{ari04,sverrir04}@ru.is

Abstract

The Solid State Drive will soon become the de facto standard of main data storage for computers. It will replace the Hard Disk Drive in a similar manner the Hard Disk Drive replaced the Magnetic Tape just over 20 years ago. In this paper we will discuss the different technologies behind the two and give arguments for why the Solid State Drive will be the successor of the Hard Disk Drive, both in conventional computing and in other types of devices both portable and stationary. We will also speculate on what difference the penetration of the Solid State Drive will make in the evolvement of computing and the effect it has on peoples every day life.

1. Introduction

Data are continuously piling up in the modern world, at higher rates every day. The content generated and stored in 2006 was 161 exabytes (161 billion gigabytes) and the extra data each year is expected to rise six-times over by 2010 to 988 exabytes [1]. A natural question seems to be, how will these data be stored? Constantly emerging are faster Central Processing Units (CPU) and faster Random Access Memory (RAM), but the trend in Hard Disk Drives is that the storage capacity is the selling point, not so much the performance. This leads to the fact that, performance-wise, HDDs are becoming more and more of a bottleneck in the computer, since the HDDs can't handle sending/receiving data to/from increasingly faster CPU and RAM.

Flash-based Solid-State Disks offer a feasible alternative to the traditional HDDs; they have been used in both enterprise and military applications, due to their speed and ability to perform in extreme conditions. SSDs have been used as cache in high-performance enterprise applications. SSDs have been too expensive in the past for the consumer market but have been becoming more affordable, slowly but steadily. The gap between the price ($/GB) of HDD and SDD is still relatively large but we can still take advantage of SSD.

Hybrid drives are a mix of HDD and SSD in such a manner that the HDD employs SSD as a Non-Volatile Cache (NV Cache) to reduce boot time, decrease electricity usage and improve durability and reliability.

In this paper we will compare the traditional HDD to SSD. In chapter 2 we will loosely cover the history of storage devices and discuss the basic technological factors of the HDD both in terms of it's strengths and weaknesses. In chapter 3 we will compare various factors of the SDD to the HDD and last explain why the SDD will be the successor of the HDD

2. Background

What is data? One could say that all information is data, but not all data is information. Data, and the need to store data has been around from the beginning of humankind, if not longer. Early humans scribbled images on cave walls describing hunting techniques and battles or just telling stories. Later came stone tablets, papyrus and the paper still used today.

When computers became important tools the need to store data in a machine readable and writable form became exigent. For the purpose of this paper we will, from now on, use the term *data* for information stored and read by computers or other types of machinery.

Before the Hard Disk Drive

The earliest known form of data storage medium was the punch-card invented in 1725 by Basile Bouchon. He used a perforated paper loop to store patterns woven into clothes. In 1884 Herman Hollerith adapted this technique to feed instructions to a mechanical tabulating machine he was developing. Punched cards, and later punched tape was used for storing instructions for machines and computers until the mid 1970s.

In 1946 the Radio Corporation of America (RCA) started development work on the Selectron tube by request of John Von Neumann. The Selectron was a volatile medium which means that the data had to be refreshed periodically and was erased when the power was cut from the tube. The Selectron tube suffered production problems and was not released until 1948, but by that time Williams-Kilburn tube had been released. The Williams-Kilburn tube is said to be the medium on which the first ever electronically stored-memory program was ever written[2].

Magnetic tape became commercially available in the 1950s. Magnetic tape is a non-volatile storage medium consisting of a thin plastic strip coated with a material capable of storing positive and negative state of magnetic particles. Magnetic tape was first used to store computer data in 1951 on the UNIVAC I [3].

Since a reel of magnetic tape could store as much data as 10.000 punch cards it became an instant success and remained as the premium data storage medium until the mid 1980s. The latency for random access on magnetic tape is quite long since the tape has to be wound on average one third of the tape's length from one arbitrary data block to another. Most tape systems try to compensate for this, either by keeping a separate index table for looking up physical file locations, or by marking blocks with a tape-mark that can be detected by winding at high speed.

Although the magnetic tape has run it's course in day-to-day computing it is still widely used for backing up and archiving purposes, since it is capable of holding large amounts of data and the backup process supports writing the data in a sequential manner.

A drum memory was invented in 1935 and became widely used in 1950s and 1960s. Many computers used a drum memory as the main working memory, with data and programs being loaded on or off using paper tapes and punch cards.

A drum is a large metallic cylinder coated with magnetic material on the outside. It rotates inside a casing with a number of read/write heads positioned along the long axis of the drum, one for each track.

Drum machines were both expensive and large, but they made fast random access of non-volatile data possible. The read / write heads were fixed for every track of the drum so the only latency was the rotation delay. Since the relatively slow rotation of the drum, there are stories of programmers optimizing the code by locating instructions on the drum so that just as one finished its job, the next would be just arriving at the "read head" and available for immediate execution[4].

The Hard Disk Drive

The Hard-Disk Drive (HDD) is a non-volatile storage medium that stores the data on a rotating platters with magnetic surfaces. HDD has probably been the fastest evolving device used in computers in terms of cost and capacity. The early models were large, heavy, expensive and fragile and had only 4,4 megabytes capacity and weighed over a ton [5].

For many years HDDs were only used in controlled environments such as data centres. Due to their delicacy, size and power consumption they were not suited for use in industrial environments, small offices or peoples homes.

Before 1980 HDDs had typically 8 or 14 inch disks (20 – 35 cm) sometimes removable, and required an equipment rack or large floor space. They often needed high currents and sometimes three-phase power due to the large motors they used. Because of this HDDs were not used with microcomputers until after 1980, when the first 5.25inch HDD was introduced by Seagate Technologies.

The disk systems used by microcomputers in the early 1980s were externally mounted. However the IBM PC/XT, released in 1983, had an internal HDD. That started a trend towards buying bare disks and installing them directly in computers. To this day internal HDDs have been the standard in microcomputers, but external HDDs remained popular much longer in the Apple Macintosh, every Mac until 1998 had a SCSI interface for externally mounting HDDs (and other peripherals). In the late 1990s with the appearance of high speed external interfaces such as USB and FireWire, external HDDs became popular again amongst PC users, but this time as a secondary means of storage often meant for backing up and transferring large volumes of data between computers.

In time the HDD became smaller and faster with more storage capacity. This gradually led to changing the way programmers think. Many programming tasks involve a time-space trade-off where less consideration for data storage space can result in faster execution time. For example, in database applications it is now a common practice to store precomputed views on disk and reducing normalization in order to speed up queries. 20 years ago such use of disk space would be considered at best impractical.

With more and more storage capacity, the HDD has moved the computer from being a single task machinery where each task started with feeding instructions and data to the computer, to being a multi-purpose device where the instructions (programs) and data is readily available whenever you need it to be.

The massively increased storage capacity of the HDD has changed the way people use computers. It's now a common practice for people to store their whole music collections or even feature films on their laptop or desktop computer where only 20 years ago a singe tune would have occupied the lion's share of a computer's disk space.

Today the HDD is no longer bound to computers as we think of them, but is a common component in devices such as cameras, music players and recorders for video and audio. Much effort has been made to make HDDs less fragile and better suited for use in portable devices. Mainly this has resulted in strategies such as buffering and parking the read-write heads in a safe zone when sudden movements are detected. This has given acceptable results, but HDDs are still fragile devices and have the tendency of failure.

The main improvements in HDDs have been in size and storage capacity, but although performance has improved dramatically it hasn't kept up with the performance improvements of computers. As computers become faster, the access time of the HDD has become a bottleneck. Due to physical properties of the HDD the only uptrend possible is by increasing the rotation speed and storing the data in a more compact way increasing the transfer rate.

Lately in sophisticated database management systems, the main emphasis has been to improve the memory manager in an effort to minimize the need to access data stored on disks. This has been possible because of greater speed and lower cost of volatile Random Access Memory (RAM). Some systems go so far as keeping all data in memory and using disk storage only for archiving purposes[6].

HDD Design

A typical design of a HDD consists of a spindle holding one or more disks, onto which the data is recorded. The disks are made from non magnetic material such as glass or aluminium and are coated with a thin magnetic film. The data is recorded and read with read-write heads as the disks are spun at high speed. The position of the read-write heads is controlled by a common arm called the actuator arm.

The data is arranged on a disk in tracks aligned in circles from the disk's centre. Each track has a number of data blocks. A common size for a data block is 512 bytes. A cylinder consists of the corresponding tracks over multiple recording surfaces.

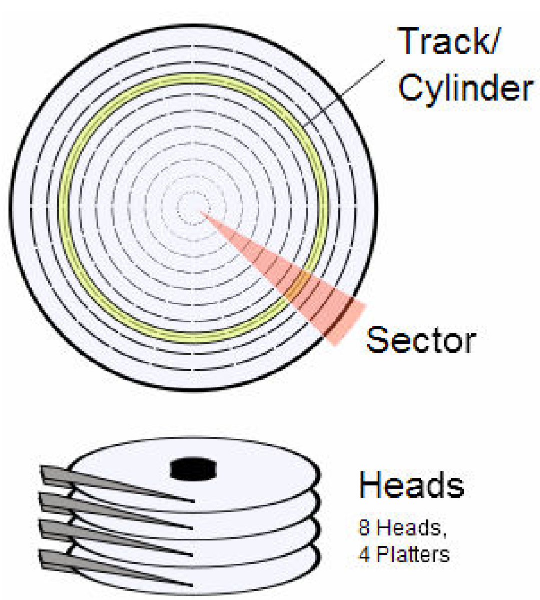

Figure 2.1 HDD Data Arrangement Scheme

In early versions the same head was used to read and write data but in today's disks the read and write elements are separate in a close proximity. The read-write heads float on a cushion or air only nanometres above the disk's surface.

Due to the close proximity of the read-write heads to the disk's surface it's internal environment must be kept immaculate to prevent substance such as fingerprints, dust and smoke particles from damaging the disks. This also makes HDD sensitive to any sudden movement and bumps as the read-write heads can crash into the disk's surface and damage it.

To reduce the risk of damaging the disks the actuator arm is parked at a "landing zone" when not performing reads and writes. The landing zone is a zone on the disk where no data is recorded. When the platters spin down the air pressure beneath the read-write heads decreases resulting in the heads touching the surface of the landing zone which in turn causes wear on the read-write heads thus limiting the disks life span.

The Solid State Disk

The Solid state disk (SDD, also known as Solid State Drive) is a data storage device that uses semiconductors to store data instead of spinning magnetic disks. With no moving parts the SDD eliminates performance bottlenecks such as seek-time and latency along with other mechanical delays and failures associated with the conventional HDD.

To this day there have been two main types of SDD available: The Flash based SDD and the SDRAM based.

SDDs based on volatile SDRAM are known for super fast access times, less than 0,01 milliseconds which is over 250 times faster than the fastest HDD available in 2004 [7]. Due to their volatility they usually contain a battery and a backup disk system. If the power is lost, the battery life is sufficient to write the data to the backup disk for safe keeping. When the power is restored the data is read from the disk into the SDRAM and normal operation is resumed.

The flash based SDD are non-volatile and do not require a backup system in case of power outage. Although they don't offer as fast access as the SDRAM – SDD they are significantly faster than the traditional HDDs.

3. The HDD vs. the SSD

SSDs were originally designed in the 1980s for use as cache in real-time performance critical applications, as well as mass storage in industrial and military systems, where immunity against shock, vibration, and extreme temperatures was required. Until the end of the 1990s, due to high-cost, SSDs were almost only used in avionics and defense industries.

Even though the acronym SSD stands for Solid State *Disk* it does not technically contain any form of disk. The name *Disk* has sticked to them because they are typically used to replace disks in situations where conventional drives are impractical.

Reliability and Mortality Rates

Although the most common F-SSD chips have a limited amount of write cycles per block, it isn't too much of a downside. An F-SSD should last several years, depending on its capacity. They typically withstand greater vibration and temperature ranges than DRAM-based SSDs, and compared to conventional HDDs, which are fragile to those factors, some F-SSD are even considered "ruggedized" by NASA and the U.S. Military. Those drives will withstand intense extremes that would reduce a rack-mount box to rubble.

HDD drives cannot withstand the operating temperature range required in most cases by military applications, +60°C to +95°C. And while SSDs are not very sensitive to temperature changes, temperature deltas as low as 15°C can nearly double disk failure rates [8].

The operational lifespan of a computer HDD is typically just over three years [9], but even with usage patterns of writing/reading gigabytes per day, an F-SSD should last several years, depending on its capacity [10].

Performance

While the disk capacity has been growing at about 60% annually, the price of storage has been dropping, but the speed of data access has grown more slowly, holding total system-performance gains back. The average HDD data-transfer rate, usually measured in megabytes per second (MB/s), had in January 2004 grown at about 40% per year while the more critical metric for most commercial applications, the number of input/output operations per second (IOPS), had grown only 16% per year [11].

As of early 2007 the typical access time for a rotating disk is around 5 to 10 milliseconds (5,000 to 10,000 microseconds) while whereas an SSD's is about 35 to 100 microseconds [10]. This makes an SSD approximately 100 times faster than a rotating disk.

Data Capacities

Data capacity has always been a major selling point for the average user. Many consumers only look at the number of MBs or GB on a disk but fail to acknowledge the importance of the performance issues.

The current state of affairs is that the worlds first 1 Terabyte HDD was announced January 4th 2007 [12], and is to be shipped in the 1st quarter of 2007. The largest consumer SSDs are currently at 64GB. While the consumer market is much restricted by price, the military and industries have a higher budget. The shipping of a 160GB SSD is expected in 2007 [13].

Price Trends

The price of a new technology is almost always high in the very beginning. With time, better production methods emerge, more customers come into play as mass-production goes underway. With mass-production, prices lower gradually.

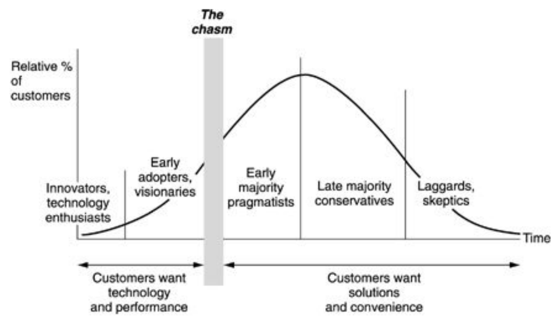

Figure 3.1 A typical penetration cycle of a technology.

A key factor of new technologies is crossing the "chasm" [14]. The chasm is in a way the defining point (in a very simplified look at the market) of a technology, whether it "makes it" or not [Fig. 3.1]. If a technology crosses the gap it is presumed to have got enough maturity for the market to adopt it and the market starts demanding it

Looking at the price trends ($/Mbyte) in HDDs and comparing them with the trends in SSD prices, it is estimated that the evolution of prices in the SSD (F-SSD and DRAM-SSD) market is headed the same way and by 2009 we will see the comparable prices to those of HDDs [Fig. 3.2].

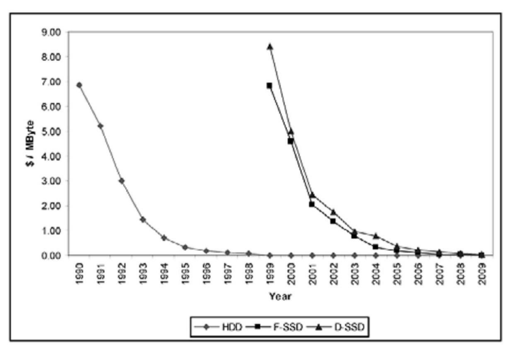

Figure 3.2 Price Evolution Trends Estimate

Heat Generation and Power Consumption

Because of the inherent lack of rotating (or otherwise moving) parts of SSDs their power consumption is much lower than that of conventional HDDs. In turn, this also results in less heat generated, which provides less wear and tear.

F-SSDs do not need as much power as DRAM-SSDs, As a result F-SSDs do not generate as much heat, thus do not need cooling fans whereas DRAM-SSDs do. The fans required for DRAM-SSDs contribute to their power consumption, noise generation and total volume. This fact makes SSDs an ideal storage medium for portable computers and devices where battery life is an important factor.

Improvements

HDDs are constantly being improved, and with microscopic magnetic grains on the disk becoming tinier some challenges in the hard drive indurstry are bound to surface. One of those challenges is overcoming the constraints imposed by the super-paramagnetic effect which occurs when the grains become so tiny that ambient temperature can reverse their magnetic orientations. When that happens, the bit is erased and data is lost.

Longitudinal recording [Fig. 3.3] aligns the data bits horizontally, parallel to the surface of the disk

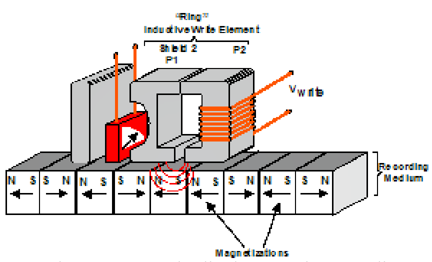

Figure 3.3 Longitudinal Magnetic Recording

Perpendicular Magnetic Recording (PMR) [Fig. 3.4] was first introduced in September 2006. In contrast to longitudinal magnetic recording, PMR aligns bits vertically, perpendicular to the disk. This allows for smaller grains, wherein it is harder to reverse the magnetic orientation, which, in turn, leads to smaller physical bits that are more stable at room temperature.[15]

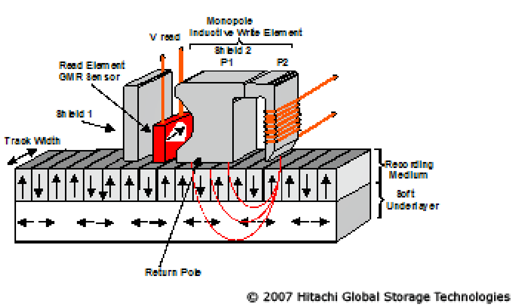

Figure 3.4 Perpendicular Magnetic Recording

The two main types of SSDs, Flash-based SSD (F-SSD) and Dynamic Random Access Memory SSD (DRAM-SSD), have undergone improvements. The DRAM-SSD has a fast sustained read/write and low latency characteristics but a major drawback is that it is a volatile storage device. A volatile storage device is one that loses all its data if its power is cut off. Another drawback to DRAM-SSDs is that they consume huge amounts of power.

F-SSDs on the other hand are non-volatile, that is retain their data on power outage, and can retain the data for up to 10 years without power. They consume much less power and therefore generate much less heat than the DRAM-SSD. The F-SSD's drawback compared to DRAM-SSD has up until now been speed, especially writing speed.

Noise

Traditional HDDs always cause some noise while they are operating. This is due to the fact that they contain actuators and spinning platters. In an F-SSD, there is no moving objects so no noise will be generated. This leads to the fact that laptop computers with F-SSD memory can be operated in places where noise is not allowed, such as libraries. [16]

Users and applications

Because the SSD access time does not depend on a read/write interface head synchronising with a data sector on a rotating disk, the SSD can provide faster random access time.

Airlines have since the beginning of the 1990s started using SSD black boxes for in flight instead of magnetic tapes. Magnetic tape devices of course, have moving parts, but SSDs do not. And with no moving parts, there are fewer maintenance issues and more chance of recovering the information after a crash. While magnetic-tape Cockpit Voice Recorders (CVR) store the last 30 minutes of sound, CVRs that used solid-state storage can record two hours of audio. Solid-state recorders can also track more parameters (such as time, pressure, altitude, airspeed, etc.) than magnetic tape because they allow for a faster data flow. Solid-state FDRs can store up to 25 hours of flight data. Each additional parameter that is recorded by the FDR gives investigators one more clue about the cause of an accident [17].

The gaming industry can also take advantage of SSDs, as seen in the case of Icelandic Company CCP. CCP is the developer of EVE Online which is a Massively Multiplayer Online Role Playing Game (MMORPG) which has over 75.000 subscribers. The Online Transaction Processing (OLTP) of the 10.000+ users was resulting in up to 20 seconds latency, which is unacceptable, especially for a video game. The game developers were able to improve the performance by 4000% by replacing HDDs with SSDs [18].

Hybrids

The cost of SSDs is coming down, but it's still expensive, more expensive for most consumers than to just get a traditional HDD for their laptop. Because of this fact, a number of hard drive manufacturers have formed an alliance called the Hybrid Storage Alliance, to help facilitate the use of hybrid drives among consumers while the prices are still too high for most of them. This will in turn help lower the cost of SSDs without adding too much (about $15 to $20) to the price per drive compared to a regular hard drive.

A hybrid drive contains both a regular rotating disk and a flash memory chip that stores data and applications. This setup allows a laptop to keep the hard disk turned off most of the time, since the hard disk consumes a lot more power than the SSD does, and thus allowing for data to be stored and recovered very quickly. This will decrease the boot time of the computer considerably and Microsoft is already taking advantage of this technology through its Windows ReadyDrive features in Vista. This technology is by no means confined to Microsoft and other operating systems (OS) can use this technology as well.

Of course turning the hard disk off saves power since it curtails platter spin time, which reduces power consumption, which in turn extends battery life.

Yet another perk of keeping the hard disk turned off most of the time is that it makes the laptop more rugged in the sense that if a laptop is dropped and lands while its hard disk is running, the head will most likely bounce off the hard drive and scratch it, resulting in lost data.

For desktop consumers this is not expected to bring much to the table since the advantages of the hybrid drive like fast resume, drop damage reduction and increased battery life do not apply to desktops to the same extent as to laptops.

IDC predicts hybrid hard disk drives will make up 35 per cent of all hard disk drives shipped with mobile devices by 2010 [19]

Conclusions

The Hard Disk Drives have changed the world of computing dramatically. They've evolved from being cumbersome uncouth devices with very limited storage capacity to becoming delicate and refined piece of equipment capable of storing massive amounts of data. It's safe to say that for the last 20 years or so every personal computer has used a Hard Disk Drive for the main data storage.

Although HDDs have in time become faster, more rugged and less subjective to failure, they still suffer from the physical limitations of the rotating disks and other moving parts. For that reason HDDs have not been able to keep up with the evolution in computers and have become a performance bottleneck holding back advances in computing.

Solid State Disks have been known and used in military and aviation applications since the early 1980s. They do not suffer from the physical limitations HDDs do and can tolerate much more physical vibration and changes in temperature and air pressure. They contain no moving parts and are truly randomly accessible.

The reason for why they have not been used in commercial applications to this day is mainly the pricing, they have been far to expensive for the average consumer and busyness. Luckily we are seeing considerable changes there. Recently Hard Disk manufacturers have been marketing hybrid drives that utilize the advantages of both solid state and hard disks for only slightly higher price than traditional HDD.

Looking at history, it's evident that the pricing and capacity trend of the SSD is at a similar spot as the HDD was in the beginning of the 1990s. Given the advantages the SSD have over the HDD, a reasonable assumption would be that the SSD will become highly disruptive to the HDD in a similar manner that the HDD disrupted the use of the magnetic tape in computing.

At this point in time the SSD is crossing the chasm that separates the technology from being used only in specialized applications to being used commercially in devices people use in their every day lives moving the HDD down to back-end processing in legacy applications.

But the HDD still has a role in our future. For the reason that the HDD is still a practical medium for sequential access of large volumes of data, it's foreseen that HDDs will become the medium of choice for backing up and archiving, and will be releasing the magnetic tape of it's duties once again.

We have come to the point in time the power of the ordinary desktop/laptop computer will be unleashed by replacing the HDDs with the faster and more randomly accessible SDD. Both will that result in more possibilities in applications for traditional computing and a plethora of new devices, both portable and set-top devices built for special purpose

References

[1] Martyn Warwick, "Web 2.0 changing the face of the planet – and the data storage industry," vol. 2007,

[2] Wikipedia. Williams tube. 2007(March 22),

[3] Wikipedia. (2007, March 20). *Magnetic tape data storage. 2007(March 22),* pp. 1. Available: http://en.wikipedia.org/wiki/Magnetic_tape_data_storage

[4] Utastro Nather. (1983, May 21). *Real programmers. 2007*Available: http://home.att.net/~rmestel/articles/real_programmers.txt

[5] Wikipedia. (2007, March 18). *Early IBM disk storage. 2007(March 22),* Available: http://en.wikipedia.org/wiki/Early_IBM_disk_storage

[6] Oracle Corporation. (2006, *Berkeley DB reference guide*. Available: http://www.oracle.com/technology/documentation/berkeley-db/db/ref/program/ram.html

[7] Wikipedia. (2007, March 22). *Solid state disk. 2007(March 2007),* Available: http://en.wikipedia.org/wiki/Solid_state_disk

[8] Gerry Cole, "*Estimating Drive Reliability in Desktop Computers Consumer Electronics Systems*," vol. 2007,

[9] Wilson WeiSheng Wang and Joanne De Peralta, " *War of the Disks: Hard Disk Drives vs. Flash Solid State Disks*," vol. 2007,

[10] Nelson Johny, "*SOLID STATE DISK: Time For HDDs To Retire?*," vol. 2007,

[11] David Freund, "*SSDs in the Mainstream*," vol. 2007, 2004.

[12] Melissa J. Perenson, PC World. *Hitachi introduces 1-terabyte hard drive. PC world. 2007(18/03),* Available: http://www.pcworld.com/article/id,128400-page,1/article.html

[13] Wolfgang Gruener - TG Daily, "*Adtron to offer 160 GB 2.5" flash drive*," vol. 2007,

[14] MIT Press, "*The Invisible Computer* ," vol. 2007,

[15] Hitachi Research - Perpendicular Magnetic Recording, "*Perpendicular Magnetic Recording*," vol. 2007,

[16] Samsung, "*SAMSUNG Electronics Launches the World's First PCs with NAND Flash-based Solid State Disk*," vol. 2007,

[17] Kevin Bonsor, "*How Black Boxes Work*," vol. 2007,

[18] Texas Memory Systems, Inc., "RamSan Solutions - The World's Fastest Storage," vol. 2007,

[19] Vawn Himmelsbach, "*Alliance pushes hybrid hard drive technology*," vol. 2007,