Storage

Whitepaper

Next Generation Mobile Hard Disk Drives

The next generation of hard disk drives is here

The evolution to smaller form factors is a natural occurrence in technology and, indeed, in many industries. This evolution has occurred in the hard disk drive (HDD) industry for more than 30 years and continues to do so today. After the usual skepticism, the IT industry ultimately welcomes the change to smaller form factors not merely because of the inherent advantages, but because of the undeniable performance improvements in each successive product.

Traditionally, 2.5” hard disk drives have been viewed as “mobile” or “laptop” HDDs because of their predominant use in laptop PCs. Today, we are seeing 2.5” hard disk drives — and even smaller form factors — available for a wide range of storage designs. We have seen 2.5” hard disk drives accepted in many new applications. This adoption will continue this year and beyond.

This whitepaper focuses on the uses and features of SATA Extended Duty Mobile 2.5” hard disk drives in high-density applications in which power, performance and capacity features are a priority. Examples of these uses are Storage Blades, Near-line storage, Video Servers and compact RAID systems. Enterprise storage systems have also made the transition to 2.5” form factor hard disk drives. These HDDs have higher reliability and higher performance than the traditional 2.5” PATA hard disk drives. They also have the advantages that come with the Serial Attached SCSI protocol. Enterprise systems are not covered in this paper.

When compared to 3.5” hard disk drives, 2.5” hard disk drives offer lower power, lower acoustics, higher capacity per same space, RAID/redundancy capabilities, and overall performance improvements. For example, multi-terabyte video storage systems can be built reliably into an industry-standard 1U rack form factor using 2.5” hard disk drives.

FUJITSU

THE POSSIBILITIES ARE INFINITE

Performance

Smaller form factor HDDs can offer significant performance improvements.

Hard disk drive performance is measured in two ways — transactional and sequential performance. Transactional performance is a measurement of the number of individual commands that can be processed in a second (IOPS — Inputs and Outputs per Second). Sequential performance is the measurement of the amount of data that can be transferred in a second (MB/sec).

The table below is a 2005 comparison between the following drives:

- 2.5” 5400 RPM SATA
- 3.5” 7200 RPM SATA
- 3.5” 10K RPM SATA

	5400	7200	10000
Form Factor	2.5	3.5	3.5
Drive Capacity (GB)	100	400	73
Number of Drives	4	1	1
Total Capacity (GB)	400	400	73
Total Cost ($)	1,400	277	240
Transactional I/O’s	53,789	12,631	34,961
I/O/GB	134.5	31.6	472.4
$/I/O/GB	10.4	8.8	0.5
Sequential Data Rate (MB/s)	105	25	68
Read/Write Power (W)	7.6	12.5	8.4

Higher RPM 3.5” hard disk drives have a higher outer band data rate due to the circumference of the recording track. This provides an advantage in sequential performance when comparing single hard disk drives. However, the 2.5” hard disk drives are small enough that three or four can fit into the same space as a 3.5” hard disk drive and can generate higher data transfer rates if all HDDs are transferring data at the same time. Multiple 2.5” hard disk drives also provide an advantage in I/Os per gigabyte in a transactional environment by effectively reducing the seek time by 1/3 or 1/4 the normal rate. The 7200 3.5" hard disk drive shows an advantage in cost per I/O per GB if the user makes a trade-off between performance and cost.

Power

Power is a key consideration in many systems.

In a system that needs to optimize power, a 2.5” hard disk drive configuration is an ideal choice. Low power is the key to increasing packing density or the number of drives per cubic foot. A 3.5” 7200RPM hard disk drive uses approximately 9.5 watts in seek/read/write operation while a 2.5” 5400RPM hard disk drive uses only 2.1 watts. In addition, 2.5” hard disk drives have power saving features that cause the hard disk drive to automatically drop to a lower power consumption after a certain time with no system access. The table shows the six levels of power management used by 2.5” hard disk drives. It is the combination of low operating power and the ability to enter additional power saving modes that is a compelling argument for 2.5” hard disk drives in power-sensitive applications.

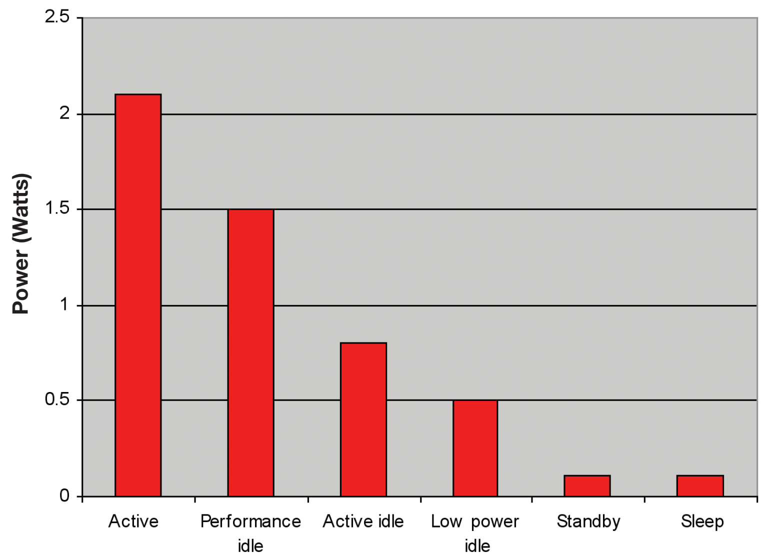

Mean Time Between Failure (MTBF)

2.5” HDDs have lower power and offer improved MTBF and life.

Most 3.5” enterprise hard disk drives today are specified at 1.4M hours mean-time-between-failure (MTBF), while desktop hard disk drives at 500K MTBF, and mobile hard disk drives at 300K MTBF. There are design differences and operating conditions that make these ranges acceptable for the drive industry and specific applications. For example, enterprise hard disk drives are expected to operate non-stop — approximately 700 hours per month — while desktop and mobile hard disk drives have an expected operating time of eight hours per day or 200 hours per month.

All hard disk drives are susceptible to damage by excessive heat. Mobile hard disk drives tend to be operated in a high temperature, low airflow environment and are therefore perceived to have lower reliability. If mobile hard disk drives were operated in an enterprise environment with constant temperature, then the demonstrated MTBF would prove to be an improvement over the mobile specification. The table below represents an example of a Mobile hard disk drive decrease in MTBF as the disk enclosure (DE) temperature increases.

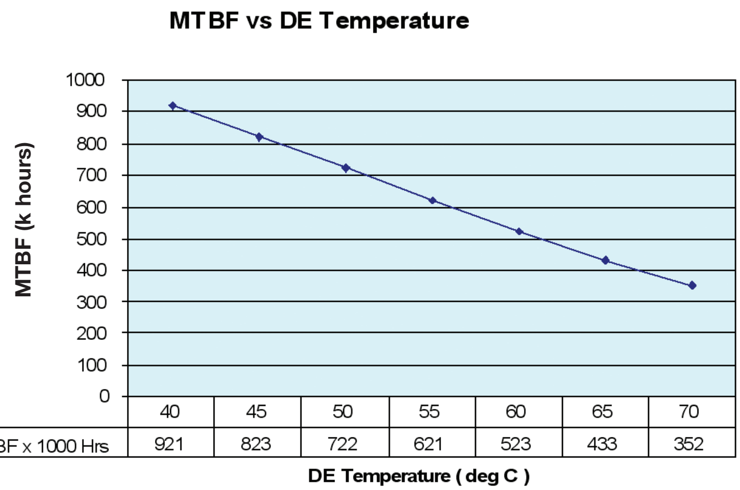

	40	45	50	55	60	65	70
MTBF x 1000 Hrs	921	823	722	621	523	433	352

There is a common misconception that as the number of hard disk drives in a system increases, the failure rate increases proportionally. However, this is not true. MTBF is a statistical result, calculated using a Poisson distribution, or a mathematical statement of the probability that exactly *n* discrete events will take place during an interval of length *t*. In fact, it is this reality that makes RAID a viable and reliable solution. RAID systems might typically employ four to five active hard disk drives with one additional drive providing redundancy. Should any of the hard disk drives fail, the data can be rebuilt from the redundant data drive. If the failure rate increased linearly with the number of hard disk drives, this would not be an economic proposition. From both mathematical calculation and empirical observations, a RAID system is far less likely to fail (two hard disk drive failures) than a single hard disk drive. So the replacement of single 3.5" hard disk drives with multiple 2.5" hard disk drives to achieve performance, capacity, lower power and better acoustics, can also achieve better reliability with appropriate and affordable redundancy.

Life Cycle

Lower temperatures mean longer HDD life.

Testing has shown the largest contributor to the life of a hard disk drive is the operating temperature. Heat from the application environment and sense current exacerbates electron migration in small semiconductors. Specifically, the read head element on a hard disk drive is the smallest geometry and therefore the most sensitive. Testing has shown that head designs that improve the heat dissipation away from the head element have a much improved life. Hard disk drive manufacturers will attack this problem in two ways: first by recommending lower operating temperatures for the device enclosures; and second, by selling only three- or four-headed hard disk drives into this type of market. In multi-headed hard disk drives, the effect of sense current flow for servo on-track is spread over more read elements. These measures will guarantee a minimum life of five years. (See table below).

Most 3.5” desktop hard disk drives are not designed for the rigors of 24x7 operation and generate more heat than the 2.5” hard disk drives. Small form factor hard disk drives are designed for the harsh temperature and shock environments of mobile computing. Through a few important design changes, hard disk drive suppliers have created 2.5” hard disk drives that operate 24x7 at a DE surface temperature of 55°C.

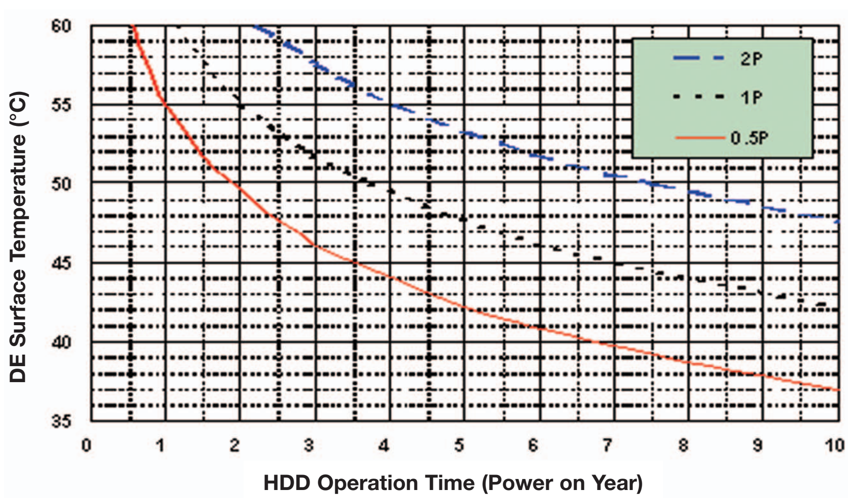

Product life is a serious consideration if an application uses data for much longer than the typical five years. That is not to say the hard disk drives will not perform for a much longer time. ATA-class hard disk drives simply have not been tested in the 24/7 environment for an extended period of years.

Rotational Vibration

Software and hardware improvements can minimize risk of performance loss due to rotational vibration.

Induced Rotational Vibration (IRV) is a phenomenon that occurs when multiple hard disk drives are packaged together and the effect of the seek acceleration is transmitted from the seeking hard disk drive to another hard disk drive that is writing. The transmitted vibration can cause the write head to be pushed off-track and result in a retry and therefore lower performance.

Higher mass hard disk drives — such as 3.5” hard disk drives — tend to induce higher force vibrations on their surroundings. Lower mass hard disk drives — such as 2.5” hard disk drives — are less able to resist the vibration when they are received at the actuator. The worst case scenario would be many 3.5” hard disk drives seeking while a 2.5” hard disk drive is performing a long write.

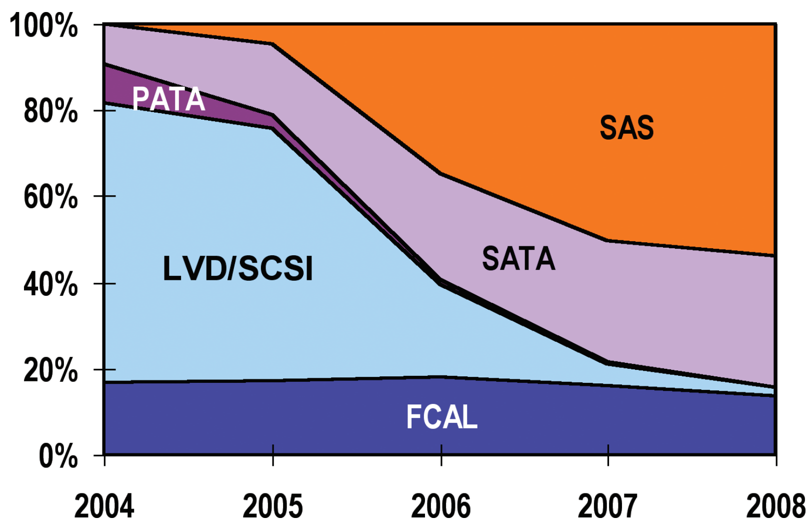

The system we are envisioning has only 2.5” hard disk drives. Experiments have shown that the low mass 2.5” hard disk drives do not generate sufficient vibration during seek to cause anything more than a negligible impact on error rate and overall performance. Testing on a tightly mechanically coupled, four hard disk drive per PCB system has demonstrated less than 0.1% performance degradation.

In a high transactional read environment such as video streams, this negligible effect will be further limited by error correction and RAID configurations.

Induced rotational vibration is another area where smaller form factor hard disk drives will show an advantage over today’s 3.5” hard disk drives.

Interface

Serial interfaces are replacing Parallel interfaces.

Hard disk drive interfaces are moving from parallel to serial — Serial Attached SCSI (SAS) will replace Parallel SCSI, Serial ATA (SATA) will replace Parallel ATA, with Fibre Channel continuing as the serial interface of choice for very high-end storage.

In the low-end market there are two major advantages for SATA in multi-drive environments. First, SATA has shown its superiority in RAID operations because all connections are point-to-point and the design is not complicated by the Parallel ATA master/slave connection. Second, RAID controllers are designed for larger numbers of hard disk drives that can only be connected using the small SATA seven-wire interface on a backplane. It would be impossible to create the same density of packaging with Parallel ATA hard disk drives and have multiple 80-wire busses running through a backplane. SATA also brings the peformance advantage of command queuing and the opportunity to reorder read commands for minimum access time. The Parallel ATA version of this feature has never been accepted by system suppliers. Overall, the choice of interface for 2.5” hard disk drives is not important for functional operation. It is merely a matter of coincidence that small cable interfaces (SATA and SAS) are available at the time of the transition from 3.5” to 2.5”.

SATA also offers a future roadmap. The following chart shows the future of serial interface data rates up to 600MB/s, a speed that cannot be achieved with a parallel connection.

	2005	2006	2007	2008	2009	2010
PATA MB/s	133					
SATA MB/s	300	300	300	300	600	
P-SCSI MB/s	320					
SAS MB/s	300	300	300	600		
FC MB/s	400	400	400	400	400	800

Capacity

Capacity is becoming less important as there are risks of storing too much data on one HDD.

A capacity comparison shows that 2.5” hard disk drives have approximately 50% of the capacity of a 3.5” hard disk drive on a per platter basis. Current technology allows 100GB per platter on 3.5” desktop hard disk drives compared to 60GB per platter on the 2.5” hard disk drives. This puts the smaller form factor HDD at an immediate disadvantage when measuring cost per gigabytes. In addition, most 2.5” hard disk drives are designed to be low power by restricting the number of platters to a maximum of two per HDD while 3.5” desktop hard disk drives are designed for highest capacity without regard to power usage.

However, the capacity advantage overlooks the inherent risk of having too much data reside under one actuator. This can lead to performance issues, reliability issues and recovery issues. The use of 2.5” hard disk drives allows the customers to configure with smaller capacity points and take advantage of the smaller physical size by adding additional spindles. The industry will use new metrics such as total I/Os, I/Os per GB, and cost per I/O per GB as a much better guide when evaluating relative performance in a transactional environment.

Capacity will always remain an important feature when discussing relative hard disk drive characteristics. However, in system configurations such as Storage Blades and Video storage, other factors — such as reliability, power and performance — have to be given equal consideration.

Summary

Small form factor hard disk drives offer the advantage of lower power, higher capacity and higher performance per cubic volume when compared to 3.5” hard disk drives as well as more than adequate reliability and life when operated in a controlled environment with temperatures below 55°C.

2.5” hard disk hard disk drives represent the future of the disk form factor as physical restrictions make it impossible to keep increasing the areal density on 3.5” platters. Solving this problem will inevitably cause a reduction in the disk platter size to less than 3.5”.

About Fujitsu Computer Products of America, Inc.
Fujitsu Computer Products of America, Inc. is a subsidiary of Fujitsu Limited, a leading provider of customer-focused IT and communications solutions for the global marketplace. FCPA provides innovative solutions for the U.S. marketplace. Current product and service offerings include high performance hard-disk drives, magneto-optical drives, scanners, and scanner maintenance.

Fujitsu Computer Products of America, Inc.
www.fcpa.fujitsu.com
1255 East Arques Avenue, Sunnyvale, CA 94085-4701. (800) 626-4686 (408) 746-7000 info@fcpa.fujitsu.com

Printed in the U.S.A. Qty. 3M WP10100106