

USENIX Association

Proceedings of
FAST '03:
2nd USENIX Conference on
File and Storage Technologies

San Francisco, CA, USA
March 31–April 2, 2003



© 2003 by The USENIX Association
Phone: 1 510 528 8649

All Rights Reserved

FAX: 1 510 548 5738

Email: office@usenix.org

For more information about the USENIX Association:

WWW: <http://www.usenix.org>

Rights to individual papers remain with the author or the author's employer.

Permission is granted for noncommercial reproduction of the work for educational or research purposes.

This copyright notice must be included in the reproduced paper. USENIX acknowledges all trademarks herein.

More than an interface — SCSI vs. ATA

Dave Anderson, Jim Dykes, Erik Riedel

Seagate Research
1251 Waterfront Place
Pittsburgh, PA 15222

{dave.b.anderson,james.e.dykes,erik.riedel}@seagate.com

Abstract

This paper sets out to clear up a misconception prominent in the storage community today, that SCSI disc drives and IDE (ATA) disc drives are the same technology internally, and differ only in their external interface and in their suggested retail price. The two classes of drives represent two different product lines aimed at two different markets. In fact, both classes contain a range of products that address a variety of features and usage patterns beyond simply the interface used to talk to the device. The target market and final product specification are taken into account from the earliest design decision through the manufacturing and testing process. This paper attempts to clarify the differences by illuminating some of these design choices and their consequences on final device characteristics. This will hopefully allow the community to build better storage systems with better knowledge of the trade-offs being made and the performance characteristics that result.

1 Introduction

Every manufacturer has different product families aimed at different customer segments. A Smart city coupe from DaimlerChrysler is much different than a Mercedes E-class sedan, although the apparent technology (gasoline engine, four round wheels) may be quite similar.

The disc drives traditionally sold with personal computer systems are quite distinct in appearance, performance and cost from those sold on larger computer systems. We will refer to the former as personal storage (PS) and the latter as enterprise storage (ES).

There are, of course, more than two classes of disc drives. Portable computers and some consumer electronics devices use disc drives that differ in important ways from either of the classes we will discuss here. We will leave as future work comparing the unique features of those drives with their larger cousins.

1.1 ATA versus SCSI

The question addressed in this paper is often phrased in terms of ATA drives versus SCSI drives. This is not accurate, as we will see: *the ATA versus SCSI debate groups the drives by interface, but the interface is perhaps the least significant difference*. Differences in mechanics, materials, electronics, and firmware make for the real distinctions among drive

families and product lines. When choosing a drive for a particular application, system designers must consider these underlying factors, and not assume that the interface distinction alone is sufficient.

The interface difference may appear to categorize the drives correctly, but, in fact, does not. There have been several instances of PS drives equipped with a SCSI interface and ES drives are also used in high-end personal computers. There is no inherent reason why an ES drive could not have an ATA interface.

1.2 Personal storage

The most important quality in PS drives is that a drive have a cost commensurate with the cost of the system in which it is installed. The cost pressure of the personal computer market gave rise to the first low-cost hard discs, and has continued to put pressure on PS drive pricing. As we discuss PS drives, we will come back to this point repeatedly: low cost dominates the design of PS drives.

When the first personal computers appeared, none had a hard drive. The drives of the day were too big and far too expensive. The customer demand for a hard drive based personal computer drove the development of a small-sized, low cost drive.

1.3 Enterprise storage

Since their invention disc drives have been used on large computer systems. At the time, these systems tended to be very big, expensive and were employed to access large quantities of data. Because of the cost, they were used to support many users simultaneously.

This environment gave rise to the essential properties of ES drives. First, they tend to be configured in groups (aggregation), as opposed to PS drives, which are most often the only drive in a system. Second, they are used to randomly access small portions of large data spaces. Third, reliability and performance are critical characteristics. A failure could idle a considerable number of employees and directly impact business operations. In normal operation, the faster the drives can service requests, the more employees can be supported and the more productive those workers can be.

1.4 Key requirements

We will now look at these key requirements and see how they have manifested themselves in drives for each market.

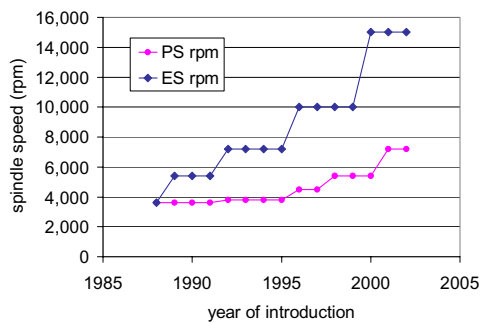


Figure 1: The adoption of higher rotational speeds. Data from Control Data product guides, 1988-1989 and Seagate product guides, 1990-2002.

1.4.1 Cost

There is constant pressure to reduce drive costs, even as drives become more complex to build. Due to the resulting demands on encoding schemes, error correction, and servo processing, it takes considerably more logic to control basic reading and writing with every areal density improvement. It also requires greater precision and lower tolerances for noise and interference of any kind. Each component of a drive must become more complex in order to deliver state-of-the-art capacity, while at the same time being pushed to become less costly to build.

1.4.2 Seek performance

Improving seek performance is the continuous struggle to get the head to move from one location to another faster than in the previous generation product. This involves using more expensive components such as higher performance magnetic circuits, faster microprocessors and lower-mass actuator assemblies. The process of designing an ES drive involves more sophisticated modeling and analysis to optimize the structures for seek movements. The various vibrational modes of the structure can negatively affect seek performance. Fast seeks depend on the ability to rapidly follow the servo patterns on the media in a predictable way. The design must preclude drive seeking being throttled by an obscure resonance of the head/disc assembly [IBM99c].

1.4.3 Rotational latency

Latency is improved by spinning the media faster. PS drives are much slower to adopt the performance improvements first introduced in ES drives. PS performance enhancements are made only when they do not incur any marginal cost. After a given capability has been in ES drives for some years it is practical to move it to PS models; the cost penalty and development cost having been eliminated by the volume of ES market. Figure 1 shows the history of rpm adoption in mainstream products over the last 15 years.

In fact, this history illustrates a general characteristic of the relationship between ES and PS drives. ES drives tend to drive costly innovation - achieving new levels of performance, reliability or function - and PS drives adopt that technology when it becomes cheap enough. This is a model that puts ES drives in a difficult pricing position compared to PS drives, but growth in the ES market depends on these added capabilities.

There is innovation in PS drives as well, but it tends to be in terms of cost savings, such as making a 7,200 rpm motor cheaper, rather than building a 15,000 rpm motor for the first time. ES drive cost comes in the form of higher cost of materials, but also in larger research and development investment.

1.4.4 Aggregation

A notable difference in operating environment between PS and ES drives is the use of ES drives in groups. This is more than simply an interface issue - just being able to electrically interconnect multiple drives. A property of Fibre Channel (FC), SCSI and Serial Attached SCSI (SAS) is that they efficiently attach more drives to a host than the two drive limit of a traditional IDE controller.

That is not, however, all there is to aggregation. If drives are housed together and used at the same time, interactions occur that can dramatically decrease performance if no compensation is included. When one drive is trying to seek or simply stay on track while nearby drives are spinning, there is an energy transfer, known as *rotational vibration*, from one seeking drive to the other drives in the cabinet.

1.4.5 Reliability

Reliability varies significantly with usage patterns and operating environment. Personal computers are designed for active use only several hours per day, while most enterprise systems are active 24 hours a day, every day. This means that design choices made in PS drives for cost reasons will make them less likely to perform well under operational stresses for which they were not designed.

1.5 History of the interfaces

Traditionally, the difference in the two interface was based on how much work was done by the host and by the drive. Until a few years ago, IDE controllers used programmed I/O, where the main system processor was responsible for all interactions with the disc drive, without interrupts or direct memory access (DMA) to offload data transfer. In SCSI, there was always an external control chip on the drive that handled independent operation of the drive.

While a standards group is currently adding a command queuing function similar to that in SCSI to the Serial ATA (SATA) protocol, ATA historically has not added any of the major features of SCSI: multiple CPU support (both failover and simultaneous operation), variable block size support

(that is, the ability to specify and format the drive to a non-512 byte block sector) and dual porting. Note that as this type of ES drive functionality accretes to PS drives, the complexity of implementation also increases.

2 Technology Differences

The differences between PS and ES drives are far-reaching and start from the earliest design choices. The diagram in Figure 2 illustrates the basic components of a modern disc drive. This section will consider each of these items in turn. Note that since the market for disc drives is a very cost-sensitive one, drive designers will not spend an extra penny in material or assembly cost to go beyond the target device specifications.

2.1 Mechanics

The basic component choices in the mechanical portion of the drive affect the overall reliability, seek time, acoustics, and resistance to temperature, shock, vibration, and other environmental variations.

2.1.1 Head/Disc Assembly

The head/disc assembly (HDA) consists of the base casting, heads, actuator, spindle, discs, air handling system, and top cover. The ES drive operates at higher rpm, while also maintaining a higher tolerance for external disturbance. These external disturbances could be the influence of neighboring drives — rotational vibration — or other environmental factors such as temperature. This is complicated by the fact that higher rpm and faster seeking ES drives put more energy into a drive cabinet, creating more disturbance. At the same time the drives are required to be less affected by it. This requires more rigidity in the mechanical structure of the drive, more mass, higher bandwidth servos, and in some cases special support circuitry to offset effects that could otherwise decimate a drive's performance.

Higher rpm drives also require more power to operate, creating more heat that can affect the drive or its neighbors in a cabinet.




Figure 2: Diagram of the major components of a disc drive.

Achieving a million hour MTBF drive is not an easy thing. Every failure mode must be addressed. ES drives will have tighter tolerances and design rules to control externally and internally generated particles and outgassing. These rules include such things as avoiding through holes, greater environmental control and higher quality sealing. The ES drive typically has more environmental protection. An ES drive will have a filter for particles, a desiccant to control humidity, and an active carbon absorbent for eliminating organic substances inside the HDA. The spindle motors have O-ring seals, and the drive cover better gasketing. Each of these little things adds cost but improves reliability. Individually, each one addresses a relatively minor failure mode, but together they help achieve 1,000,000+ hour MTBF.

PS drives are designed for reliability, but they tend to compromise where components can be eliminated to save cost. The O-rings and desiccant, for example, are usually eliminated in PS drives.

An ES drive has more shrouding and air control devices to better manage air flow inside the HDA. This eliminates air turbulence, which would otherwise make it harder to keep the head on track and optimize seek performance. It also directs air to the actuator to help cool it. It also adds cost.

The size and stiffness of the base casting and top cover impact both the acoustic characteristics of the drives, as well as the susceptibility to rotational vibration. Both of these problems become more acute at higher spindle speeds.

2.1.2 Actuator

Larger magnets are key to achieving faster seek times, but they bring additional requirements, along with a higher cost. In order to get the most seek performance and still stay within a tight power budget, the ES actuator coils must have less resistance. This requires thicker coil material with fewer windings. As already mentioned, special HDA design features promote cooling the actuator to prevent overheating.

An interesting example of complex interactions arises with the latch. Inside every drive is a latch to hold the actuator when power is off. The most common method of latching involves a magnetic circuit. However, the latch has a magnetic field associated with it, which can affect seek performance when the actuator is operating near the latch. In a PS drive, there is no compensation for this, as seek performance is not critical. To achieve the optimum seek performance, ES drives will have a bi-stable latch that does not affect performance. This is a more expensive solution, but gives better overall performance.

Both the coil and the bearing cartridge are independently bonded to the arm using a special epoxy in an ES drive. In a PS drive the coil is likely to be attached to the arm with a single molded connector, a less expensive technique. The former makes for a more rigid structure and is necessary to achieving maximum seek performance.

In a PS drive, seek performance is not the high priority it is in an ES drive. Typically a PS drive design must first achieve its cost targets, and then do the best it can with seek performance. The opposite priority holds with ES drives, e.g., the actuator design must prevent its various bending modes and resonances from impacting seek, settle time, and performance in the presence of rotational vibration.

2.1.3 Spindle

For over 15 years drives spun no faster than 3,600 rpm. Since then drives have been sped up first to 5,400 rpm, then 7,200 rpm, 10,000 rpm and most recently 15,000 rpm. Spinning faster is a tremendous engineering challenge. The read/write head must be kept on track, and this is increasingly difficult as rpm goes up. An off-track head during reading can cause a mis-read and a *rotational miss* (requiring a full rotation before the read can be re-tried). An off-track head during writing can cause a mis-write that introduces noise or even overwrites adjacent tracks.

Higher rpm requires more expensive motors. As tracks per inch (TPI) increases the motor becomes a bigger challenge. Disturbances such as *windage* (air movement between the disk and arm) and vibration increase with rpm. At the same time ES drives must be less affected in order to get the best possible random performance. For cost reasons PS drives use a cantilever motor design, where the motor shaft is captured only at the base deck end. An ES motor shaft is captured at both ends, with an attachment to the top cover. With today's TPI, fluid bearing motors are preferred since they minimize runout and acoustical noise (see below for a discussion of runout). For years it was thought impossible to have a fluid dynamic bearing motor captured at both ends. Seagate solved this with a unique conical design that gives ES drives the benefits of both fluid bearings and a motor supported at both ends. This is a more expensive design, but gives better overall performance.

2.2 Electronics

The on-drive electronics are becoming more integrated as improvements in processor technology allow [Matsumoto99]. This means that fewer components are required to provide the same basic functionality.

2.2.1 Control processor

The drive servo system keeps the read/write head on track or moves it from one track to another. The drive determines its position by reading very small fields of information interspersed among the data blocks on every track (servo bursts). Every time the head crosses over a servo burst, the microprocessor suspends what it is doing and takes up the task of identifying where the head is. If it is wandering off track slightly, it must move the head in the appropriate direction and distance to get back in the middle of the track. During seeks, the actuator constantly reads servo bursts as it crosses

tracks. This information is used to determine how close the actuator is getting to the target location and, when it is close, to decelerate the actuator.

2.2.2 Servo processor

As TPI gets higher, more servo processing is needed to keep the head off neighboring tracks. This would not be so hard if the tracks were perfect, repeatable circles. They are not: motor variation, platter waviness (both circumferentially and radially), stacking tolerances and other factors give rise to both repeatable and non-repeatable runout. *Runout* - variation in the radius or circumference of the track - occurs when the head is unable to follow the current track and stay in position above it. Repeatable runout is inherent in the track, and is the same on each rotation, making it easier to compensate for. Non-repeatable runout is due to external influences such as vibration, and varies over time. The servo processor must adjust the head to follow the track wandering underneath it. To get more servo capability, higher capacities require more servo bursts. This requires more processing works against minimizing cost and increasing capacity. A PS drive is a constant balancing act between minimizing cost - including processing power - and tracking the higher TPI's to achieve maximum capacity.

2.2.3 Interface

There is significantly more silicon on ES products. The following comparison comes from a study done in 2000:

- the ES ASIC gate count is more than 2x a PS drive,
- the embedded SRAM space for program code is 2x,
- the permanent flash memory for program code is 2x,
- data SRAM and cache SRAM space is more than 10x.

The complexity of the SCSI/FC interface compared to the IDE/ATA interface shows up here due in part to the more complex system architectures in which ES drives find themselves. ES interfaces support multiple initiators or hosts. The drive must keep track of separate sets of information for each host to which it is attached, e.g., maintaining the processor pointer sets for multiple initiators and tagged commands.

The capability of SCSI/FC to efficiently process commands and tasks in parallel has also resulted in a higher overhead "kernel" structure for the firmware. All of these complexities and an overall richer command set result in the need for a more expensive PCB to carry the electronics.

When the drive processor is busy doing servo work and read/write tasks, it cannot be doing interface work. In order for an ES drive to offer the maximum performance, it is equipped with two processors - one dedicated to servo and the other for interface and read/write handling. Maximizing random access and performance under rotational vibration both depend on that dedicated servo processor.

A PS drive has a single processor, which must handle all three basic processor tasks in a drive. It must run the interface, support the reading and writing of data and do all the servo processing.

2.2.4 Memory

The firmware for the SCSI command set is more than twice as large as that for ATA, requiring more permanent flash for code and increased SRAM at runtime. The more complex command set and larger command queues also require additional memory space. The SCSI command set allows for vendor-specific extensions which require additional code space, allowing greater flexibility in configuration.

2.3 Magnetics

In magnetic componentry there is much similarity between the ES and PS drives since both strive to stretch the same areal density boundary. Differences stem from the performance goals of the ES drives. The higher rpm of ES drives delivers higher data rates.

2.3.1 Heads

Though magneto-resistive head technology has made a profound change in how data is read in a drive, writing is still an inductive process. As such it is sensitive to linear velocity and higher rpm improves not only latency, but data rate, as well. For this reason, ES drives tend to stretch writing capability, and demand constant innovation to keep up with the high rpm and higher areal density. PS drives usually adopt the writer technology proven in previous generations of ES drives.

Reading is just the opposite. Read data rate is generally insensitive to linear velocity, but in some cases it may be adversely affected by higher rotational speed. Signal amplitude does not increase as it does with inductive heads, but noise does. This means that ES drives, with their higher rpm and data rate targets, have a more difficult magnetic environment in which to read data.

The key property to having a system that will read and write reliably is the signal to noise ratio (SNR). It is much harder to reach a given SNR in a high rpm drive. This makes the task of extracting the data from the read signal significantly more difficult. This is sometimes referred to as recording stress, and is usually more pronounced in an ES drive. ES drives must have more expensive read/write electronics to cope with this more difficult magnetic environment and higher data rate.

2.3.2 Materials

The traditional substrate material for media is aluminum, onto which a layer of magnetic material is deposited. The recent use of glass substrates provides a greater uniformity of the magnetic surface and greater stiffness [IBM99], but the magnetic layer is harder to deposit on glass, making it




Figure 3: Diagram of media layers. The base substrate consists of either aluminum or glass, topped with a layer of magnetic material. In anti-ferromagnetically coupled (AFC) media, an additional layer of magnetic material and a layer of ruthenium are added, with the two layers reinforcing each other for better magnetic stability at higher density.

more difficult and expensive to achieve the same read densities [Walker01]. The better shock tolerance of glass must be traded against lower density or data rate. In addition, since glass cannot be textured, the actuator must be removed from the media to land (load/unload ramp) rather than landing on the media (contact start/stop) [IBM99a]. This requires a landing zone at the outer edge of the disc in case there is contact as the heads leave the platters - precisely the area of highest density and data rate.

A recent change in media structure is the use of anti-ferromagnetically coupled media, which contains a second magnetic layer oriented opposite the primary layer to reinforce the magnetic orientation [IBM01]. This is necessary to achieve higher densities, at the cost of increased complexity in both materials and in the manufacturing process [Walker01]. The diagram in Figure 3 illustrates this layering.

2.4 Manufacturing

The build and test times for ES drives are considerably longer than PS drives. Increased test time can make a drive more reliable. During this time, drives also undergo detailed characterization, such as learning precisely how irregular individual tracks are, which allows them to better keep the heads on track during normal operation. More time spent analyzing the media for flaws results in lower probabilities these flaws causing unrecoverable read errors in the field.

3 Performance Differences

We have outlined the design choices possible when designing a disc drive for a particular target market. Most of these choices affect performance in some way, and we will now attempt to quantify the impact of specific choices.

3.1 Capacity

The basic media structures used are the same in both drive types, with the highest areal density used at any given time. The choice of the number of disc platters and the size of the platters changes the overall capacity - for example, 15,000

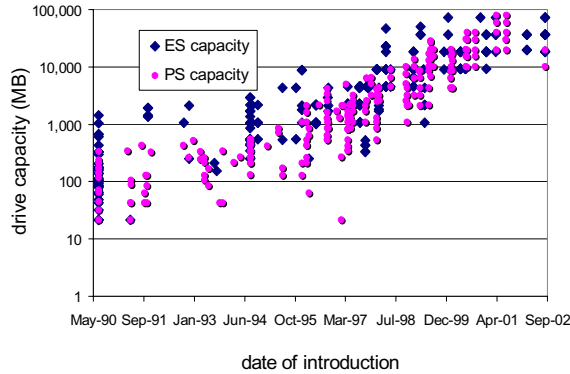


Figure 4: Comparison of capacities. Capacity and introduction date of 10 years of Seagate drives [Seagate02].

rpm drives use 2.5” platters to support the faster spindle speeds, while 7,200 rpm drives use 3.7” platters.

3.1.1 Size of Platters

ES drives spin faster to get better performance. However, power increases almost to the cube of rpm. Smaller diameter platters keep drive power at an acceptable level. This has a cost: an ES drive uses more platters to achieve the same capacity as a PS drive at a given areal density. The smaller platters actually brings two performance advantages - the ability to spin faster and faster seeking. Average seek times are better because the head must traverse a smaller recording band. This, together with the greater investment in actuator capability as discussed earlier, makes for drives that perform random access much faster than their PS counterparts at equivalent areal densities.

The larger diameter platters and the lower rpm give the PS drive a clear advantage in delivering capacity. This is consistent with the primary market requirement of lowest cost. A combination of minimizing the parts cost and delivering the highest capacity yields the lowest dollar per gigabyte. The data in Figure 4 compares drive capacity against date of introduction over the last 10 years.

3.1.2 Number of Platters

Many drives are manufactured with fewer platters than possible, as performance matters more than capacity. The chart in Figure 5 illustrates this trend toward depopulated drives, as well as the more recent use of depopulated heads (only using one platter surface to save the cost of the additional read/write head).

Fewer platters translates into faster seeks because there are less heads, so the actuator has a lower total mass and can move a fraction of a millisecond faster, which can be significant at sub-4 ms average seek times. This also matches the marketplace as users with a requirement for performance will often buy more drives, each of a lower capacity, to spread data across as many actuators as possible.

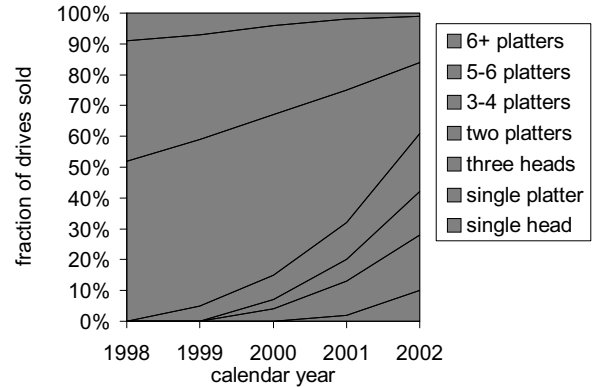


Figure 5: Trend toward depopulated drives. More users are choosing drives with less platters, trading capacity for performance.

3.2 Data rate

The fastest ES drives will always have higher data rate than a contemporary PS drive, due in large part to the higher rpm, as explained earlier. However, the PS drive has an advantage in media size. Typically PS drives use 95 mm (3.7”) platters, compared to 84 mm (3.3”) in 10,000 rpm and 65 mm (2.5”) in 15,000 rpm ES drives respectively. The larger media size helps the PS drives follow closely in data rate.

Another factor favoring the PS drive is that new models tend to come out more frequently than ES drives. Introduction of a new ES drive comes when the new generation is able to double the capacity of the previous generation. Hence successive models over the last several years have been 9, 18, 36, 73 and 146 GB. PS drives, on the other hand, come out as soon as it is possible to deliver an appreciable increase in capacity. Instead of doubling, they have been introduced at 10, 20, 30, 40, 60, and 80 GB per platter. This higher frequency enables PS drives to stay much closer to ES drives in data rate than if they were following the same “jumps” that happen for ES drives. The data in Table 1 compares the data rates of several drives and shows the underlying spindle speed, areal density, and platter size.

		cap	speed	density	dia	int bw (Mb/s)	ext bw	
		GB	rpm	Gb/in ²		calc	spec	
							MB/s	
Atlas 10k 18WLS	ES	18	10000	3.4	3.3”	-	314	24.6
DeskStar 75	PS	30	7200	11.0	3.7”	551	444	35.6
Cheetah 36LP	ES	18	10000	7.3	3.3”	579	427	-
Cheetah X15	ES	18	15000	7.3	2.5”	690	508	39.5
Cheetah X15-36LP	ES	36	15000	17.5	2.5”	969	709	57.7

Table 1: Comparison of drives with increasing data rates. Capacities, speeds, and densities are from published spec sheets. Diameters are typical for those spindle speeds. Internal bandwidths are calculated from the speed, diameter, and TPI as shown in the spec sheets. External bandwidths are as measured by LinuxHardware.org [Augustus01].




Figure 6: Diagram of basic drive parameters. A smaller media (lower r) sacrifices bandwidth and capacity for shorter seeks and lower power. Each additional platter adds to capacity and power consumption.

This table shows the three components to sequential data rate: the rotational speed, the areal density, and the diameter of the platters. Higher speed drives will use smaller platters for lower energy consumption and faster seeks, resulting in lower data rates. The five drives in the table are arranged in order of externally-measured sequential throughput. We see that the 7,200 rpm DeskStar is faster than the 10,000 rpm Atlas due to a much higher areal density, and a larger platter diameter. The DeskStar and the Cheetah 36LP are quite close in data rate because the increased rpm of the Cheetah is only enough to overcome the density disadvantage and the smaller platter diameter. The Cheetah X15 at 15,000 rpm gains data rate, but loses some due to the further reduced platter diameter. Finally, the second generation Cheetah X15 increases the areal density and far outperforms the others even with the smallest diameter platter.

The diagram in Figure 6 illustrates the trade-off among data rate, capacity, seek time and power consumption when choosing a platter size.

3.3 Random performance

Random performance describes the ability of a drive to get from one location to some other unpredicted address to service the next request. There are three components to the performance of this movement - seek performance, controller overhead, and rotational latency.

3.3.1 Seek times

Several of the mechanical items mentioned in the last section directly affect the ability of the drive to seek quickly and to stay on a servo track in response to environmental factors. The data in Figure 7 compares the seek time of drives against their date of introduction.

Seek performance of PS drives always lags that of ES drives, and improves at a slower rate, while ES drives are expected to squeeze out a gain with each new generation of drives. The entire mechanical design of an ES drive is focused on

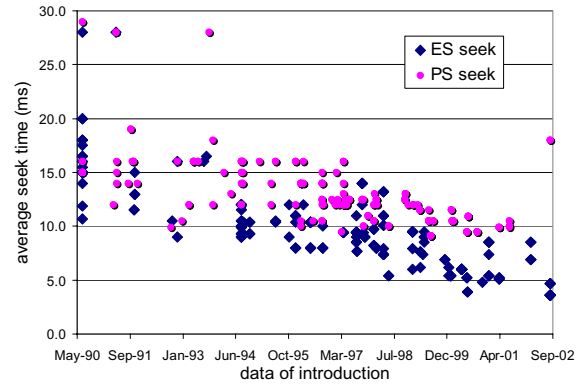


Figure 7: Comparison of seek times [Seagate02].

server	Dell PowerEdge 2550
operating system	Windows 2000 Pro SP3
scsi controller	Adaptec 39160
ata controller	Promise Ultra 100 TX2
ES drive	Seagate Cheetah 73LP - 73 GB
specs	10,000 rpm; 18 Gb/in ² ; 5.1 ms seek
PS drive	Seagate Barracuda IV - 80 GB
specs	7,200 rpm; 31 Gb/in ² ; 9.5 ms seek
benchmark	iometer (jan 2001 code release)

Table 3: Experimental ES and PS drive testbed.

achieving the highest random access performance, as this is critical in the target market.

Table 2 shows the seek performance of a PS against an ES drive under the same workload. This comparison is between a Barracuda IV and a Cheetah 73LP drive in the same system. Table 3 details our experimental setup. The mechanical details of these two drives are different, but quite close - higher density in the Barracuda compensates for the higher spindle speed in the Cheetah. The higher spindle speed will also account for some of the improvement in random performance on the ES drive.

queue depth	read (8 KB)		write (8 KB)	
	PS	ES	PS	ES
1 requests	65 req/s	115 req/s	105 req/s	184 req/s
2 requests	66 req/s	116 req/s	105 req/s	184 req/s
4 requests	71 req/s	146 req/s	105 req/s	187 req/s
8 requests	79 req/s	174 req/s	105 req/s	190 req/s
16 requests	89 req/s	202 req/s	108 req/s	200 req/s
32 requests	101 req/s	235 req/s	108 req/s	213 req/s

Table 2: Comparison of random request rates at increasing queue depth on the same request stream in PS and ES drives. Both drives are run with write caches enabled. If the write cache on the ES drive were disabled, the improvement with larger queue depth would be even larger, as observed in a previous study [White01].

3.3.2 Seek scheduling - queue depths

Seek sorting impacts performance and PS drives generally have shorter queues. In fact, the shorter queue lengths when using the ATA interface also have a direct impact on drive mechanics. The fact that seeks are not aggressively scheduled in PS drives keeps the average seek distance closer to the theoretical average of 1/3 the disc radius. ES drives with more aggressive scheduling can bring this as low as 1/10 of the radius on average. This means that the mechanical *duty cycle* - the total amount of time spent seeking and stressing the mechanical components - of the PS drives could be more than 3x higher for a similar request stream.

The data in Table 2 compares the random performance of a PS drive against an ES drive as the queue depth seen at the drive increases. With a queue of 32 pending requests, the ES drive is able to achieve more than twice the random performance possible with only a single queued request, while the PS drive only improves throughput by 55% for reads and barely at all for writes. This is similar to results in an earlier study [White01] where ES performance increase by 100% and PS performance by only 45%. The improvement comes because seeks are smaller, which leads to both better performance and better reliability. Additional scheduling sophistication could be included in the PS drive as well, but would require some of the additional electronics discussed earlier.

3.3.3 Controller overhead

Controller overhead is optimized by having as much processor performance available as possible to interpret and schedule commands as they arrive. More recently, this has been augmented with custom hardware assist to provide more performance than could be economically realized simply by greater investments in software. Such hardware ensures that data can be moved to and from the interface at rates as close to the internal drive data rate as possible.

3.4 Rotational vibration

When one drive is trying to seek or simply stay on track while nearby drives are spinning, there is an energy transfer from one seeking drive to the other drives in the cabinet. This tends to excite the drives to rotate around their center of mass, throwing the actuator off track. Unless a drive is designed to mitigate this effect, writes will abort or seeks will fail to find the desired track. In most cases this will manifest as a decrease in performance as aborted writes and rotational misses accumulate. At its extreme, this effect can get so bad that any drive, ES or PS, will cease to function. It simply could not stay on track long enough to complete any operation. The key is to understand how much rotational vibration is likely to be present in a server environment and design the drive to withstand it.

Since PS drives are built to be in single drive systems, rotational vibration is not an important factor. Though a CD-ROM drive can create a certain amount of vibration, the

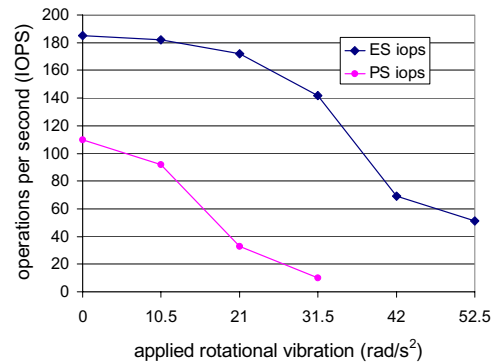


Figure 8: Externally applied rotational vibration can have a major, negative impact on performance. Individual drive cabinets vary widely in the amount of rotational vibration they transfer, and have been measured up to 45 rad/s² [Hall00]. Data for Seagate Cheetah 18LP and Barracuda III.

slight and infrequent effect is not sufficient to produce a noticeable performance problem. PC responsiveness is measured only by what a single user can see, and even a few retries would not create a serious problem in most cases.

ES drives, on the other hand, are explicitly designed to operate in cabinets full of spinning drives. This requires designing a drive to maintain its operation in the presence of considerable rotational vibration. As tracks per inch (TPI) increases, the rotational vibration problem gets worse. It is more difficult to stay on track even in ideal conditions, much less with external vibrations that are difficult to compensate for [Abramovitch96]. Some recent drives have added a rotational vibration sensor that can detect external rotation and compensate in the servo processing.

Earlier we mentioned the performance degradation possible due to rotational vibration, which we will attempt to quantify here. The chart in Figure 8 shows the performance of a single drive on a test stand under varying rotational vibration. Performance of the PS drive is much more affected than the ES drive. The PS drive essentially stops at 30 radians/s² of external vibration, while the ES drive degrades much more smoothly and is able to operate beyond 60 radians/s².

When multiple drives are placed together in the same cabinet, the rotation induced by adjacent discs or other system components affects performance. The design of a cabinet and mountings determines how bad this effect will be in a particular system. Studies of more than 20 drive enclosures and machine designs from a variety of manufacturers show a wide range of vibration characteristics - from the best designs that subject the drives to 5 radians/s² with only minor performance consequences, through cabinets inducing up to 45 radians/s² [Hall00].

3.5 Reliability

One of the trickiest drive characteristics to measure is reliability, which arises from a wide range of factors and consid-




Figure 9: Reliability reduction with increased power on hours, ranging from a few hours per day to 24 x 7 operation [Cole00].

erations in design, manufacturing, and in the operational environment [Kaczeus90, Yang99, Elerath00].

The most significant difference in the reliability specification of PS and ES drives is the expected power-on hours (POH) for each drive type. The MTBF calculation for PS assumes a POH of 8 hours/day for 300 days/year¹ while the ES specification assumes 24 hours per day, 365 days per year. The longer a drive is expected to be running, the lower the MTBF, and the higher the annual failure rate (AFR).

The chart in Figure 9 shows the expected increase in AFS due to higher power-on-hours. Moving a drive from an expected 2,400 POH per year to 8,760 POH per year would increase the failure rate almost two-fold, if there were no compensation elsewhere in the design.

3.5.1 Duty cycle

In addition to the obvious increase with increased power-on hours, the amount of mechanical work the drive has to do is affected by its basic structure and by the workload it is asked to do. A larger number of platters in the drive increases capacity, but also increases the mechanical stresses.

The chart in Figure 10 shows the increase in expected AFR with higher duty cycle. The increase is higher for the drive with the larger number of platters. For a four platter disk, a duty cycle of 40% instead of 100% would reduce the failure rate by almost 50%.

Better seek scheduling leads to shorter seeks on average and therefore a lower effective duty cycle for the same set of user requests. In preliminary measurements on our testbed, we see a mechanical duty cycle of approximately 40% for the ES drive against 75% for the PS drive on the same set of requests.

Adding platters and heads increases the AFR not just due to the additional mechanical stresses, but also due to increased internal heat generation, and the additional head/disc interfaces which might release particles or lead to other negative interactions, such as head crashes.

¹This is the specification used by Seagate, other manufacturers use varying, but similar, assumptions about power-on hours [Vilsbeck02].

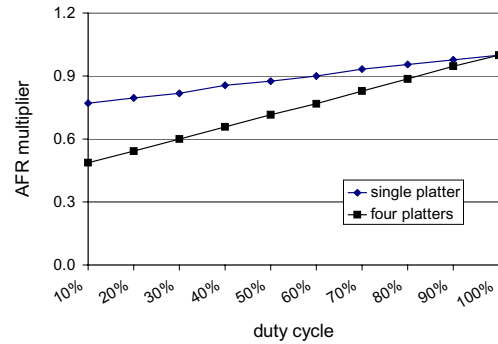


Figure 10: Reliability is decreased with higher duty cycle, and the effect is greater for drives with larger numbers of platters [Cole00].

3.5.2 Temperature

Reliability decreases with increases in ambient temperature. The drive temperature is affected not only by the outside temperature, but also by other components in the system. A high-density server rack with many disc drives grouped close together may experience much higher temperatures than a single drive mounted in a desktop computer.

The chart in Figure 11 shows increased AFR with increased temperature. A fifteen degree temperature rise is expected to increase the failure rate by a factor of two, and an increase of that size is a common assumption in high-density server racks [Patel01].

In order to prevent data corruption and failure at very elevated temperature, some drives contain temperature sensors that provide warnings of temperature outside the specification range [Herbst97].

3.5.3 Overall reliability

Each of these factors makes an individual contribution to drive failure rate, and they may also magnify each other. A capacity-focussed drive with more platters and less sophisticated seek scheduling may have a higher base duty cycle under certain workloads, and may also be more subject to temperature variation.




Figure 11: Reliability decrease due to ambient temperature variation [Cole00].

	iface	cap	price	speed	seek	density	kbpi	ktpi	internal bw			dia [^]	ext bw	dsk	cap	
									3.7"	3.3"	3.0"	spec				raw
UltraStar 36LZX	SCSI	36 GB	\$550	10000 rpm	4.9 ms	7.0 Gb/in ²	352	20.0	645	610	552	452 Mb/s	3.0"	36 MB/s	6	594 Gb
DeskStar 75	ATA	30 GB	\$159	7200 rpm	8.5 ms	11.0 Gb/in ²	391	28.4	551	487	442	444 Mb/s	3.7"	37 MB/s	2	483 Gb

Table 4: Comparison of PS and ES drives from IBM [White01]. The Deskstar drive has a slight advantage in sequential bandwidth, even though the UltraStar has a higher rpm. The authors of the previous study attribute this to overhead in the SCSI interface. In fact, a closer look at the physical discs shows the most likely explanation a smaller platter size in the UltraStar (3.0" instead of the normal 3.7"). This reduces seek time at the expense of lower sequential bandwidth on the outer tracks. Since the UltraStar has a much lower areal density, it must also make up the capacity difference by using additional platters (6 vs. 2). [^]estimated based on the internal transfer rate and raw capacity differences

	iface	cap	price	speed	seek	density	kbpi	ktpi	int bw		dia	ext bw	disks	cap
									calc	spec				raw
UltraStar 36Z15	SCSI	36 GB	\$381*	15000 rpm	4.1 ms	10.7 Gb/in ²	397	27.0	798	647 Mb/s	2.6"	53 MB/s [%]	6	661 Gb
DeskStar 120	ATA	60 GB	\$99	7200 rpm	8.5 ms	29.7 Gb/in ²	547	54.0	771	592 Mb/s	3.7"	48 MB/s [%]	2 [#]	979 Gb

Table 5: Comparison of a newer generation of drives from IBM. In this case, the new UltraStar increases sequential performance over the new DeskStar due to the higher spindle speed, even though the areal density is lower. *from Harddrive.com in August 2002 [%]according to the published specification, not a measured number [#]the 60 GB version of the DeskStar 120 has 2 disks, but only 3 heads, one surface remains unused

Past work comparing the reliability of PS against ES drives reported a failure rate of 25% for 24 IDE drives against 2% for 368 SCSI drives over an 18 month period [Talagala99]. However, these numbers cannot be treated as a controlled study due to the very small sample size for the PS drives.

Another study using data collected during the design phase of two different drives - it is not reported whether they were SCSI or ATA drives - shows a less than 1% annual failure rate for one, and a larger than 4% for the other [Hughes02]. This clearly shows that different design choices can have a significant impact on the final drive failure rates.

4 Related Work

A previous comparison of SCSI vs. IDE [White01] concluded that IDE had slightly better sequential performance, but lagged significantly in random performance. However, the authors of that study did not compare all the mechanical details of the two drives, leading to a conclusion that cannot be generalized to all SCSI and all IDE drives.

The data in Table 4 compares a set of basic characteristics for the two drives considered in their study. The slight advantage of the ATA drive in sequential performance is due to the density advantage of the ATA drive (almost 60% higher) and the larger platter diameter (between 15% and 25% larger), which is not overcome by the rotational speed advantage of the SCSI drive (40% higher). A SCSI drive with a comparable density would perform significantly better, as discussed in Section 3.2.

The data in Table 5 shows the improvement to the next generation of both drives from the same manufacturer. In the newer SCSI drive - comparing the UltraStar 35Z15 to the

DeskStar 75 at the same areal density - the rotational speed advantage, even with smaller diameter platters, still push the SCSI drive to much higher data rates (40% higher).

The advantage of the SCSI drive over the ATA drive in random performance is partly due to the smaller platters, as well as additional differences in the mechanics as explained in earlier sections. Also note that between the two generations of SCSI drives, the seek performance has improved by almost 20%, while the seek performance of the ATA drives has remained constant.

A performance comparison under Windows 2000 [Chung00] shows an IDE drive only 20% slower than a SCSI drive on sequential throughput and 44% slower on random performance. As shown in Table 6, most of this performance difference is again due to the mechanical differences. The higher density and larger platters of the IDE drive almost compensate for the faster spindle speed of the SCSI drive, although both seeks and latencies are significantly lower in the higher rpm drive. If these drives had been introduced with the same density, the higher rpm drive would also have a larger sequential throughput advantage.

	cap	seek	speed	density	dia	int bw (Mb/s)		ext bw
	GB	ms	rpm	Gb/in ²		calc	spec	MB/s
Fireball lct 08	26	9.5	5400	6.1	3.7"	343	257	19
Atlas 10K (SCSI)	18	4.5	10000	3.4	3.3"	444	314	24

Table 6: Comparison ATA vs. SCSI under Windows 2000 [Chung00].

A comparison of SCSI and ATA for end users [Dominguez99] makes many of the high-level points dis-

cussed here. ATA drives are optimized for simplicity and low cost, while SCSI drives must be optimized for performance, reliability and the ability to connect to multiple hosts. Trends in the speed and sophistication of the interfaces have been bringing ATA and SCSI closer together, with ATA gaining complexity as it moves closer to SCSI.

The comparison of ATA and SCSI reliability from the end user perspective, covering many of the factors mentioned above, was discussed extensively in a recent online article [Vilsbeck02].

Trends and recent innovation in disc drive technology, as well as the details of a specific SCSI drive design have recently been published by another disc drive maker [Miura01, Aruga01].

The design of disc drives is a very complex and multi-faceted process, and has been used as an example for students to understand engineering and cost trade-offs [Richkus99].

5 Summary and Discussion

To compare any two individual drive models at a given capacity point, one has to look at the detailed device specifications as these impact every aspect of drive design and determine drive performance. Looking at those factors in turn and comparing the impacts:

Capacity is about the same for both markets, everyone wants the highest affordable density. This is determined largely by the areal density trends. There is some variation in numbers of platters in a drive, but it is possible to build a drive of any chosen capacity for either market.

Data rate is proportional to spindle speed, areal density, and platter size. The data rate for the enterprise market tends to be higher than for personal storage, but higher spindle speeds cost more regardless of the interface used.

Fast seeks cost more and target the enterprise market. This includes larger magnets, better bearings, and stiffer actuators. The challenge is to rapidly find the target track (seek) and then to stay on track (servo) in spite of the harsh electrical and magnetic environment.

Protection from rotational vibration costs extra and targets markets where multiple drives sit next to each other. This includes better motors, top covers, stiffer actuators, and additional mass.

Better scheduling costs extra, requiring more code space, more memory for re-order queues and for algorithms. This is easier to do in the SCSI interface because it has traditionally had queueing and is more mature, but the implementation complexity would exist regardless of the interface used.

Fancier interface electronics cost extra. Because SCSI is richer and more complex, with more customer-modifiable options and host connectivity, it takes more electronics and

more memory space. This is the only difference that truly arises solely from the choice of interface.

Finally, high reliability costs extra. It needs to be considered in every component and material choice along the way, as well as in the overall design. It also has to take into account the duty cycle targets for the expected workload and the expected environment.

6 Conclusions

The differences between enterprise and personal storage disc drives are significant. They derive from the different requirements of the respective markets and offer a range of choices to system designers. Simply separating the products by their external interface - ATA vs. SCSI - misses many of the internal details and design choices that will affect system performance. We have shown that the external interface chosen is one of the smallest contributors to overall performance. The performance and reliability characteristics of a drive are determined by the way the drive is designed - from the smallest mechanical and materials choices in the head-disc assembly, through the seek scheduling algorithms in the interface processing. In order to find the right features and design points for a particular application, the underlying trade-offs must be taken into account across a continuum of specific choices.

7 Acknowledgements

We thank Albertine Flora for help with the performance testing and benchmarking. The details of drive internals reported here are due to discussions with Zip Cotter, Neal Gunderson, and Jim Weispfenning. We thank Kevin Gomez and Mark Lutwyche for explaining some of the terminology. This analysis and comparison builds on previous work led by Heath Miller and Ed Skalko, with contribution from many others. We thank the anonymous reviewers and especially our shepherd, John Wilkes, for many detailed and helpful comments.

8 References

- [Abramovitch96] Abramovitch, D.Y. "Rejecting rotational disturbances on small disk drives using rotational accelerometers" *IFAC World Congress*, July 1996.
- [Aruga01] Aruga, K. "3.5-inch High-Performance Disk Drives for Enterprise Applications: AL-7 Series", *Fujitsu Scientific & Technical Journal* 37 (2), December 2001.
- [Augustus01] Augustus "Seagate Cheetah X15 36LP Review" www.linuxhardware.org/features/01/09/10/065209.shtml, *LinuxHardware.org*, September 2001.

- [Blount01] Blount, W.C. "Fluid Dynamic Bearing Spindle Motors" *IBM Storage Systems Group - San Jose*, February 2001.
- [Chung00] Chung, L., Gray, J., Worthington, B. and Horst, R. "Windows 2000 Disk IO Performance" *Technical Report MSR-TR-2000-55*, Microsoft Research, June 2000.
- [CNET02] CNET Hardware Reviews. May 2002.
- [Cole00] Cole, G. "Estimating Drive Reliability in Desktop Computers and Consumer Electronics Systems" *Technology Paper TP-338.1*, Seagate Technology, November 2000.
- [Dominguez99] Dominguez, R. and Coligan, T. "SCSI vs. ATA: Interface Comparison" *Technology Brief*, Dell Computer, December 1999.
- [Elerath00] Elerath, J.G. "Specifying Reliability in the Disk Drive Industry: No More MTBF's" *IEEE Annual Reliability and Maintainability Symposium*, January 2000.
- [Hall00] Hall, J. "Seagate's Advanced Multidrive System (SAMS) Rotational Vibration Feature" *Technology Paper TP-229D*, Seagate Technology, February 2000.
- [Herbst97] Herbst, G. "IBM's Drive Temperature Indicator Processor (Drive-TIP) Helps Ensure High Drive Reliability" *White Paper*, International Business Machines Corp, October 1997.
- [Hughes02] Hughes, G.F., Murray, J.F., Kreutz-Delgado, K. and Elkan, C. "Improved Disk Drive Failure Warnings" *IEEE Transactions on Reliability*, September 2002.
- [IBM99a] IBM "IBM hard disk drive load/unload technology" *IBM Storage Systems Division*, May 1999.
- [IBM99b] IBM "Higher reliability with IBM glass substrate disks" *IBM Storage Systems Division*, July 1999.
- [IBM99c] IBM "Advanced Servo-Mechanical Design Facilitates Improved Performance and Reliability" *IBM Storage Systems Division*, November 1999.
- [IBM01] IBM Research News "IBM's new magnetic hard-disk-drive media delays superparamagnetic effects" www.research.ibm.com/resources/news/20010518_whitepaper.shtml, February 2001.
- [Kaczeus90] Kaczeus, S. "Disk reliability is a function of design as well as manufacture" *Computer Technology Review* 10 (9), Summer 1990.
- [Mason00] Mason, H. "SCSI, the industry workhorse, is still working hard" *IEEE Computer*, December 2000.
- [Matsumoto99] Matsumoto, C. "Diskcon abuzz over single-chip drives, home markets" *EE Times*, eetimes.com/sys/news/OEG19990917S0013, September 1999.
- [Miura01] Miura, Y. "Information Storage for the Broadband Network Era - Fujitsu's Challenge in Hard Disk Drive Technology" *Fujitsu Scientific & Technical Journal* 37 (2), December 2001.
- [Patel01] Patel, C.D., Bash, C.E., Belady, C., Stahl, L. and Sullivan, D. "Computational Fluid Dynamics Modeling of High Compute Density Data Centers to Assure System Inlet Air Specifications" *ASME International Electronic Packaging Technical Conference and Exhibition (IPACK '01)*, July 2001.
- [Reinsel00] Reinsel, D. "180 Gigabytes in One Drive: Seagate's Barracuda 180" *IDC Bulletin*, November 2000.
- [Richkus99] Richkus, R., Agogino, A.M., Yu, D., and Tang, D. "Virtual Disk Drive Design Game with Links to Math, Physics and Dissection Activities" *29th ASEE/IEEE Frontiers in Education Conference*, San Juan, Puerto Rico, November 1999. Online at bits.me.berkeley.edu/mmcs/disk/disk.html
- [Seagate00] Seagate "Disk Drive Acoustics" *Technology Paper TP-296*, Seagate Technology, April 2000.
- [Seagate02] Seagate Support "Disc Drive Encyclopedia" www.seagate.com/support/disc, Seagate Technology, August 2002.
- [Talagala99] Talagala, N. and Patterson, D. "An Analysis of Error Behavior in a Large Storage System" *Technical Report UCB/CSD-99-1042*, University of California - Berkeley, February 1999.
- [Vilsbeck02] Vilsbeck, C. "Gefahr: IDE-Festplatten im Dauereinsatz" www.tecchannel.de/hardware/964/index.html, tecCHANNEL.de, June 2002.
- [Walker01] Walker, M. "Performance Media: Tweaking Magnetic Capabilities" *Technology Paper TP-577*, Seagate Technologies, August 2001.
- [White01] White, B. Ng, W.T. and Hillyer, B.K. "Performance Comparison of IDE and SCSI Disks" *Technical Report*, Bell Labs - Lucent Technologies, January 2001.
- [Worthington95] Worthington, B.L., Ganger, G.R., Patt, Y.N., Wilkes, J. "On-Line Extraction of SCSI Disk Drive Parameters" *SIGMETRICS*, May 1995.
- [Yang99] Yang, J. and Sun, F. "A Comprehensive Review of Hard-Disk Drive Reliability" *IEEE Annual Reliability and Maintainability Symposium*, January 1999.

9 Appendix

The data in Table 7 shows multiple generations of drives from several manufacturers, including both ATA and SCSI interfaces.

These numbers serve as a reference for the comparisons made in the paper. Detailed data is provided for all the drives discussed in the text, those mentioned in previous studies, and some recently released drives.

	iface	intro	cap	price	speed	seek	density	kbpi	ktpi	dia	int bw		ext bw	disks	cache
											calc	spec			
Quantum Atlas 10K	SCSI	1999	18 GB	-	10000	4.5 ms	3.4 Gb/in ²	256	13.0	3.3"	444	314 Mb/s	25 MB/s ^{@&}	6	2 MB
Maxtor Fireball lct 08	ATA	1999	26 GB	-	5400	9.5 ms	6.1 Gb/in ²	324	19.5	3.7"	343	257 Mb/s	19 MB/s ^{&}	3	512 KB
IBM UltraStar 36LZX	SCSI	1999	36 GB	-	10000	4.9 ms	7.0 Gb/in ²	352	20.0	3.0"	552	452 Mb/s	36 MB/s [^]	6	4 MB
Seagate Cheetah X15	SCSI	2000	18 GB	-	15000	3.9 ms	7.3 Gb/in ²	343	21.4	2.6"	689	508 Mb/s	40 MB/s [@]	5	16 MB*
Quantum Atlas 10K II	SCSI	2000	18 GB	-	10000	4.7 ms	7.7 Gb/in ²	341	14.2	3.3"	591	478 Mb/s	-	3	8 MB
IBM UltraStar 36Z15	SCSI	2001	36 GB	\$365	15000	4.1 ms	10.7 Gb/in ²	397	27.0	2.6"	798	647 Mb/s	53 MB/s [%]	6	4 MB
IBM DeskStar 75GXP	ATA	2000	30 GB	-	7200	8.5 ms	11.0 Gb/in ²	391	28.4	3.7"	551	444 Mb/s	37 MB/s ^{^@}	2	2 MB
IBM UltraStar 73LXZ	SCSI	2001	36 GB	\$239	10000	4.9 ms	13.1 Gb/in ²	480	27.3	3.3"	832	690 Mb/s	57 MB/s [%]	3	4 MB
Seagate Barracuda 180	SCSI	2001	180 GB	\$1369	7200	7.4 ms	15.0 Gb/in ²	490	31.2	3.7"	691	508 Mb/s	-	12	16 MB*
Fujitsu AL-7LX	SCSI	2001	36 GB	\$369	15000	4.0 ms	15.8 Gb/in ²	450	35.0	2.7"	954	734 Mb/s	-	4	8 MB
Seagate Cheetah X15-36LP	SCSI	2001	36 GB	\$395	15000	3.6 ms	17.5 Gb/in ²	482	38.0	2.6"	969	709 Mb/s	58 MB/s [@]	4	8 MB
Seagate Cheetah 73LP	SCSI	2001	73 GB	-	10000	5.1 ms	18.4 Gb/in ²	485	38.0	3.3"	840	671 Mb/s	-	4	4 MB
Fujitsu AL-7LE	SCSI	2001	73 GB	\$529	10000	5.0 ms	19.2 Gb/in ²	485	39.5	3.3"	838	673 Mb/s	-	4	8 MB
Maxtor DiamondMax D540X-4G	ATA	2001	160 GB	-	5400	12.0 ms	25.2 Gb/in ²	442	57.0	3.7"	467	347 Mb/s	38 MB/s ^{<}	3	2 MB
IBM DeskStar 120GXP	ATA	2000	60 GB	\$105	7200	8.5 ms	29.7 Gb/in ²	547	54.0	3.7"	771	592 Mb/s	48 MB/s [%]	2 [#]	2 MB
IBM DeskStar 120GXP	ATA	2000	120 GB	-	7200	8.5 ms	29.7 Gb/in ²	547	54.0	3.7"	771	592 Mb/s	50 MB/s ^{<}	3	2 MB
Seagate Barracuda IV	ATA	2001	80 GB	\$125	7200	9.5 ms	31.3 Gb/in ²	540	58.0	3.7"	761	555 Mb/s	41 MB/s [%]	2	2 MB
Seagate Cheetah 10K.6	SCSI	2002	146 GB	\$1139	10000	5.3 ms	34.0 Gb/in ²	570	64.0	3.3"	-	841 Mb/s	-	4	8 MB
Seagate Cheetah 15K.3	SCSI	2002	73 GB	\$769	15000	4.0 ms	34.0 Gb/in ²	533	64.0	2.5"	1071	891 Mb/s	-	4	8 MB
Western Digital Caviar WD1200	ATA	2002	120 GB	\$179	7200	10.9 ms	-	-	-	3.7"	-	736 Mb/s	50 MB/s ^{<}	2	8 MB*
Seagate Barracuda V	ATA	2002	120 GB	\$185	7200	10.5 ms	42.2 Gb/in ²	542	78.0	3.7"	764	570 Mb/s	44 MB/s [%]	2	8 MB*
Western Digital Caviar WD2000	ATA	2002	200 GB	\$359	7200	10.9 ms	45.0 Gb/in ²	-	-	3.7"	-	525 Mb/s	-	3	8 MB*

Table 7: Comparison of multiple drive generations and manufacturers. All numbers are from manufacturer specifications or product manuals, except where noted. Prices for drives still being sold in August 2002 are from dirtcheapdrives.com. Seek times are for average seek. All values for density and bandwidth are maximums (outer diameter). Internal bandwidth is calculated from the rpm, Kbpi, and disc diameter values and provided for comparison to the published values. [#]the 60 GB version of the DeskStar 120 has 2 disks, but only 3 heads, one side remains unused [@]as measured by Linuxhardware.org [Augustus01] [^]as measured at Bell Labs [White01] [&]as measured under Windows 2000 [Chung00] [%]according to the published specifications, not measured numbers [<]as measured by CNET Hardware [CNET02] ^{*}option, the default cache size is 2 MB