

The Direction of Mass Storage

Stan Pieterick, Cray Research - A Silicon Graphics Company,
890 Industrial Blvd., Chippewa Falls, WI 54729 USA
715-726-5116 spp@cray.com

ABSTRACT: *Mass storage continues to play an important role in in the computer industry. This paper describes the evolution of disk drives in the computer industry and, more specifically, the direction Cray Research has taken.*

1 Introduction

Since its inception, the computer industry has looked for better ways to store information. Initially, it took thousands and thousands of punch cards to store the results. In the early 1950's, a major breakthrough took place with the invention of tape storage. A few years later, disk storage came along and has been the major medium for storage. The first disk drives were very large and primitive, the size of a refrigerator, compared to today's disk drives which can be held in the palm of your hand.

This paper will cover the general direction in which the disk storage industry is going and which paths Cray Research has chosen to follow. This paper will not attempt to discuss in detail, the differences between the individual disk manufacturers.

2 Disk Drive Manufacturers

The disk drive industry is very competitive. The high cost and quick time to market for the individual products is making it harder for the smaller companies to compete. As a result, some of the smaller companies merged with the larger companies. While this may or may not have been good for the companies being swallowed up, it definitely benefited the end-users, who have seen the price of disk storage drop considerably.

The major manufacturers of disk drive today are; Seagate Technology, IBM, Qauntum, and Western Digital. These four companies account for the majority of the disk drives in the market place today, with Seagate and IBM being the dominate forces. The high end is really controlled by Seagate, with IBM being the only other real player.

3 Disk Drives, Past

In the early days, disk drives were designed to fit the computer. That is, the computer company dictated to the disk manufacturer what was needed in a drive and what interface it should have. For example, at Cray Research we entered into a working agreement with Ibis Technology and co-developed the DD-49 disk drive. The DD-49 was a 4-head parallel drive with a transfer rate of 12 MB/s. At the time, this was far greater than the industry standard 3 MB/s, giving Cray a performance advantage over the competition. In those days, most drive interfaces were proprietary to the computer companies, and the DD-49 was no exception. It had a Cray proprietary interface, which limited the market for this drive to Cray's customers, and, in turn, limited the revenue that the disk drive manufacturer could realize.

4 Disk Drives, Present

The disk drive manufacturers were looking at ways to maximize return on investments. The manufacturers decided to proceed with making drives with industry standard interfaces. While this was good for the drive industry (they were able to build drives in volume) this created challenges for computer companies looking for performance advantages over the competitors. The playing field for disk storage became level with industry standard disk drives. The computer companies had to be creative and figure out how to package the disk products to increase performance, decrease cost, or both.

Joining the move towards standardization, Cray Research decided to use the industry leading IPI interface for it's next generation disk drive. However, Cray Research elected to pursue drive manufacturer's who were still willing to collabo-

rate and create a special drive which would outperform the industry norm. Cray picked a manufacturer, Imprimis, with an IPI drive offering, and a willingness to modify the drive's internals to increase performance. The DD-60 drive was the fruit of this relationship. The DD-60 was a 9-head parallel drive, 8 heads data and 1 head parity, with a data transfer rate of 24 MB/s. This drive is still the fastest drive ever to be manufactured. Unfortunately, there are some pitfalls that come along with being the fastest. With 9-heads parallel, the DD-60 disk drive carried a premium price tag and the transfer rates were so fast that few companies made computers that could actually push the drive. In the end, Cray Research was the only company to market the DD-60.

Unfortunately, with the competition increasing between drive manufacturers -- each pushing to be the first to market with new products, trying to keep costs under control, and the burden on resources with doing drives for a niche market -- the drive industry determined that doing 'specials' just didn't make financial sense. Instead, they determined that, in order to be profitable, they had to develop disk products that reached the largest market segment. The largest market segment happens to be the desktop market. This shift by the drive industry required a corresponding shift for the computer companies. In order for the computer company to maintain a performance advantage over their competitors, the company would have to rely on building a faster computer, developing a way to package industry standard disk drives for a superior product, or do both.

Cray Research's next attempt was to package the disk drives creatively. Cray developed a disk controller to handle a Redundant Array of Inexpensive Disks (RAID) which would function 4 drives as data and 1 drive as parity. The controller would give the customer 4x the performance of standard drive offerings. The first controller developed connected to the 60 series disk drives, thereby giving the customer a whopping 80 MB/s transfer rate with the DD-60 drive. This transfer rate was comparable with host to host transfer rates, providing very fast file transfers. Unfortunately, the IPI interface wasn't strong in the desk top application so the interface, along with the drives, has been phased out.

Cray offered SCSI disk drives and controllers as well. Although the SCSI drives were not the performers like the IPI drives, SCSI was, and still is, an industry wide accepted disk drive, running in desktop applications as well as high end computers. Cray presently offers SCSI wide differential disk drives in 5.25 inch and 3.5 inch form factors, with transfer rates from 7 MB/s up to 15 MB/s in certain applications.

5 Disk Drives, Future

Disk drives of the future. Where do we go from here? SCSI disk drives continue to have a strong foot hold. Ultra SCSI has just been introduced into the market place, boosting data rates from 10 MB/s single ended to 20 MB/s, and from 20 MB/s to 40 MB/s for wide differential. This has greatly improved the survival of SCSI. Not too far down the road is Ultra2 SCSI.

Ultra2 SCSI doubles the transfer rates again from 20 MB/s to 40 MB/s and from 40 MB/s to 80 MB/s for single ended and differential, respectively.

On the surface, this appears to be a simple implementation phase for the computer companies and a simple migration path for the customer. However, the issues at hand for the computer companies are cable lengths and a change in the controller. With the faster transfer rates, cable length becomes critical from the controller to the drive (1.5 meters to 3 meters maximum), and, drives added to the string increase the cable length. The Ultra2 SCSI controller makes a shift to Low Voltage Differential (LVD). Low Voltage Differential allows for data switching on the cable to be increased, thereby increasing the data transfer rate. The drawback is LVD drives are not compatible with present wide differential controllers, so customers with wide differential controllers will have to upgrade to newer controllers which support LVD in order to benefit from the transfer rate. Single ended controllers will run LVD disk drives, but at the slower transfer rate, meaning the customer wouldn't gain the performance benefit of the LVD drive.

Cray Research is pursuing all of the options above. We presently offer 9 GB, 5400 rpm, 5.25 inch form factor drives and are presently looking into the next generation of this product; a larger capacity product. We also offer a 4.5 GB and 9 GB, 7200 rpm, 3.5 inch form factor products. The 3.5 inch form factor products offer N+1 power and cooling, and environmental monitoring, as will the next generation 5400 rpm product.

The disk industry is also pursuing a serial channel option for disk drives. A couple of the reasons for a serial channel are to remove the restrictions of the larger connectors (SCSI), and increase distances. The serial options presently available are Serial Storage Architecture (SSA), developed and implemented by IBM, and Fibre Channel - Arbitrated Loop (FC-AL) which has been more readily accepted by the rest of the disk manufacturing industry.

Cray Research has chosen the path of FC-AL over SSA. Cray's implementation of FC-AL has a Fibre Channel Node controller (FCN) which connects to the Cray Research GigaRing on one side, and 5 FC-AL loops on the other side. The FCN connection to the GigaRing has a transfer rate of 375 MB/s half duplexed or 350 MB/s full duplex. Each FC-AL loop has a theoretical transfer rate of 100 MB/s. Actual transfer rates will be lower than 100 MB/s but haven't been determined yet as this product is just starting to roll out.. Each loop in the Cray implementation, is capable of having 80 disk drives but for resiliency purposes, Cray Research recommends that only 40 drives be connected per controller with alternate path implemented. The disk enclosure for the fibre products is called the DSF-1. The DSF-1 can hold up to 10 disk drives, has N+1 power and cooling, and an environmental monitor. The DSF-1 has 2 separate I/O ports connected to a midplane, allowing for alternate path configurations without cables.

Recent announcements indicate that both the SSA and the FC-AL logic designers are looking for a way to incorporate the

advantages of both interfaces into one, tentatively names Fibre Channel - Enhanced Loop (FC-EL), or sometimes referred to Fibre Channel Loop (FCL). FC-EL is in an ANSI committee today and the only requirement that has been placed on it is that it must be backwards compatible with FC-AL products to protect the customer's investment.

6 Cray's Future Directions

Cray Research will continue to be on the leading edge of the disk technology. This will protect Cray's leading edge image for I/O and disk performance. There are many pitfalls along the way. Trying to be on the leading edge puts you at the mercy of the disk drive manufacturers being able to deliver on time. To be on the leading edge, Cray has to pick a certain product or technology that in most cases doesn't exist yet. Cray then designs and manufactures the hardware, writes the software, and then hopes the drives show up to begin debugging. Once the debugging begins, constant dialogue with the drive manufacturer takes place to clear up the problems seen.

Cray will continue to look at larger capacities of the present drive offerings. Cray will also look at the newer, faster interfaces which are in ANSI committees today. There are other ways that the manufacturers are looking to speed up the drives or drive the cost down which Cray will also take a look at. Based on past

industry direction, some of the features expected in the future are; faster rotational speed, smaller form factor drives allowing for denser packaging, higher bit density allowing for larger capacity drives, and faster interfaces.

Cray will also continue to develop products with reliability and resiliency in mind. Cray realizes that our customer's have a lot at stake when running their programs and that some programs require very long run time and therefore cannot tolerate an interrupt to replace a failed part. Disk array RAID solutions will also continue to be in the mix of Cray offerings, and in fact are the main push due to the benefits gained by RAID.

Cray will also be looking to take advantage of volume buying by merging their product demands with Silicon Graphics demand. The larger volume should allow for the companies to negotiate a better price where the same products can be used. This will allow SGI/Cray to price the drives accordingly.

7 Conclusion

Disk technology continues to advance at a fairly impressive rate. The drive manufacturers are always striving to be first to market with faster, higher density drives. Cray Research has the same commitment and will continue to bring emerging technologies to market as quickly as possible to provide industry leading performance to our customers.