Mass Storage Industry Direction

Mike Anderson, Cray Research, Inc., 655A Lone Oak Drive, Eagan, MN 55121 USA

ABSTRACT: *Mass storage technology continues to play a key role in determining the architecture and performance of computing systems throughout the 1990s. This paper briefly describes the overall industry direction in mass storage by examining product trends, as well as the trends in enabling component technologies. In addition, the affect these two key driving forces have on future Cray Research Inc. mass storage product offerings is explored.*

Introduction

Over the past decade, the performance of computing systems has increased by several orders of magnitude. This growth in computational performance has been fueled by evolutionary changes in the fundamental technologies that comprise the memory, central processing unit, and mass storage components of the system. An examination of the industry direction, with regard to disk storage product trends, demonstrates that the evolution of the product offerings will continue during the 1990s. These industry driving forces in turn affect future Cray disk storage product offerings.

This paper does not attempt to review the vast numbers of disk storage products available in the industry today. Instead, it focuses on the building blocks from which these products are designed and built. The disk drive is the primary building block around which manufacturers and integrators design their individual products.

Disk Manufacturers

The disk manufacturing industry has faced many of the same challenges that other facets of industry face today. The high cost of technology development, high manufacturing cost, shortened product cycles and an extremely competitive market has forced several vendors to join forces to survive. In a relatively short time, DEC/Quantum, Micropolis, Conner Peripherals, and Maxtor have all been acquired by another company. Seagate by acquiring Conner has moved to a 31% market share, significantly ahead of Quantum and IBM.

The high volume desktop disk drive market is setting the direction for the industry. Dataquest estimated that the fourth quarter 1995 drive shipments surpassed 24 million units. In 1989, only 20 million drives were shipped during the entire year. Figure 1 indicates the current market share held by the major suppliers of 3.5” disk drives.

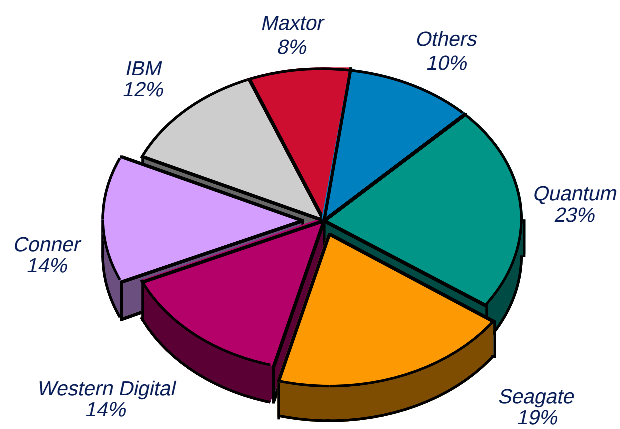

Figure 1: Estimated 1995 Disk Drive Unit Market Percentage
Source: Dataquest

Final 1995 numbers were expected to exceed 86 million units. Annual growth rates for the disk drive business have been 33 and 26 percent for the past two years, respectively. Early estimates for 1996 indicate a demand for at least 100 million disk drives. The lower ‘96 growth rate might be explained by somewhat lower-than-anticipated PC growth rates, the onset of longer disk drive product life cycles, and the absence of a “hype driver” such as Windows 95.

During 1996, allocation of IC’s, heads and disk media is expected to continue. This is partially brought about by the industry’s rapid move to new head and read/write channel technologies.

Product Trends

Disk technology trends fall into four major areas: form factor, capacity, reliability, and device interfaces.

Historically, disk drive manufacturers produced different products for the multiple markets available to them. Today’s manufacturing and engineering costs in this very competitive industry have forced a consolidation of high and low-end products. The disk manufacturers produce less generations of a product and use fewer products to address multiple markets.

In use on Cray systems today are disk products of 14, 8, 5.25 and 3.5 inch form factors. Fourteen-inch disk products dominated during the 1980s, while eight-inch disk products were introduced in 1989 and continue in use today. More recently, disk products have moved from 8 and 5.25 inch to 3.5 inch technology. The bulk of new disk products are the 3.5 inch form factor. The 3.5” drive will be the primary form factor (physical size) drive used on desktop to supercomputer systems over the next few years.

Mobile personal computers such as laptops use 2.5” disk drives due to physical size and weight limitations that are critical to their application.

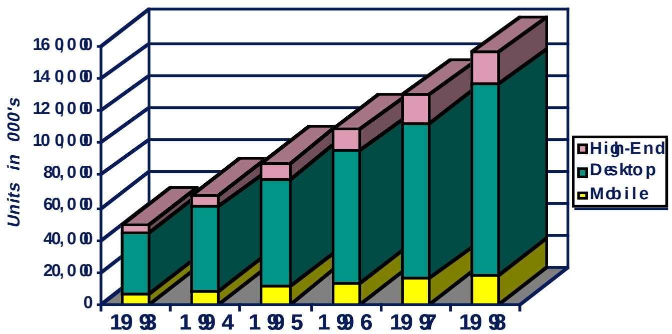

Figure 2: Total Disk Drive Forecast

Another major industry trend of importance is the data capacity of the individual disk products. During the 1980s, disk vendors invested heavily in increasing the capacity of their large form factor disk products. Today the disk industry is clearly focused on reducing the size of the form factor. As a result, this is affecting the storage capacity of each individual disk unit.

As disk drives get smaller, weigh less, and require less power for use in personal computers, including laptops, the smaller drive motors are unable to move the mass and weight of a large number of platters within a head disk assembly (HDA). Therefore, manufacturers are using fewer platters within an HDA, which would reduce the overall capacity of smaller drives if corresponding advances in head and media technology covered later in ths paper were not available.

Reliability of disk drives is clearly an important issue, yet it can be difficult for the consumer to judge a product by. The difficulties arise from the fact that the disk drive Mean Time Between Failure (MTBF) numbers are not measured or predicted in a consistent manner among vendors or products. Inconsistencies arise when one product’s MTBF numbers are based on component failure rate projections while another product’s numbers may be based on actual field installation experience. At least one major vendor uses the MTBF number as the time at which 2/3 of the drives have experienced a failure. Yet many a consumer expects a disk drive to run for the number of hours equal to the MTBF rating before it fails.

Most personal computer disk drives are considered operational and without defect until they are returned to the original manufacturer after experiencing a failure. When personal computer disk drives are disposed of without being returned to the manufacturer, the drives continue to be used in the MTBF measurement for that product. A major percentage of failed personal computer disk drives never get returned to the manufacturer. As prices of new products remain attractive, the computer user typically upgrades to the latest product available when their current disk fails and disposes of the failed drive.

The remaining disk trend that plays an important role within the industry is the device interface. A number of different drive interfaces have been available during the past 10 years. In the past, drive manufacturers sometimes used proprietary interfaces on their own brand of storage products. The SMD interface dominated the medium to the large systems market for many years and is no longer used in new products. The intelligent peripheral interface (IPI-2) is the fastest performance disk drive industry standard interface available, but is too costly for vendors to integrate in a fast changing environment and is approaching end of production.

Large-scale disk drive performance has had a difficult time keeping pace with the advancements in CPU and memory technology. A forward projection of industry trends indicates that mass storage products will be primarily tailored to fit desktop platforms. As a result, it is imperative to understand the enabling technologies from which future disk products for the desktop platforms will be built.

Enabling Component Technologies

Many leading edge technologies are part of today’s competitive disk product offerings. The key disk component technologies include: heads, sliders, discs, channel interface, on-board processor and associated electronics. This discussion is limited to the affect that channel interface and head technology has on Cray’s future disk drive direction.

During the 1980s, mass storage vendors began using thin film heads and media in their products. The thin film technology produced significant gains in reliability, media recording density, and overall disc capacity. Thin film heads use a single read/write head for both reading and writing. This produced a head design that is a compromise between the best design for reading and the best design for writing disc media.

As discs got smaller, it became increasingly difficult to provide high performance and high capacity with thin film heads. Magneto resistive (MR) heads are a revolutionary advance in disc drive technology that enables small form factor drives to reliably use higher recording density to deliver increased capacity and transfer rates without significantly increasing costs.

The MR head is actually two heads, one for writing and another for reading, both mounted on a single head arm assembly. The write head is similar to a thin film head, except that it is optimized for writing only. The read head is a magneto resistive head that enables higher media transfer rates. Reading information from the media is accomplished by constantly passing a sense current through the read element of the head. When the head passes over a magnetic field on the media, the head changes its resistance, which is detected by the change in amperage of the sense current.

With two heads on the same head arm, a repositioning of the head is necessary when switching between reading and writing on the same head to bring the write head into position over the track. This is referred to as microjog positioning. Thin film heads used only one head on the head arm assembly and did not require microjog repositioning.

Furthermore, MR heads have unique properties which make them especially useful for small form-factor drives. The data storage industry is trying to place the maximum amount of storage in the smallest size possible, and MR heads will become one of the important facilitators in this migration.

Increased Recording Densities

A big advantage of MR heads is that they provide increased areal density. This is mainly accomplished through the ability of the MR head to read signals when bits are packed ever more closely together. They are also able to better distinguish bits between adjacent tracks as those tracks are brought closer together. What this amounts to is that drive designers can build drives with higher areal densities than ever before because MR heads can read data in those higher bit-density environments.

MR heads produce a well formed signal crest with no undershoot, simplifying both the design and implementation of the drive. Conventional thin-film heads are also dependent on the rotational speed of the recording media. Specifically, their proper operation relies on the disc rotating within a required linear velocity range. This becomes a significant limitation with smaller form factor designs. As the disc diameter shrinks, the linear velocity is reduced, even though the spindle speed remains constant. Conventional thin-film heads have a harder time coping with the slowdown. MR heads are not speed dependent. The nature of their design allows identical operation irrespective of the rotational speed of the disc. This will become especially significant as higher-capacity drives emerge in 2.5-inch and smaller form-factors. MR heads are ideally suited for use in small form-factor drives, since they simultaneously provide higher areal densities and an insensitivity to linear velocity.

Magneto-resistive heads single-handedly possess many characteristic properties that make them attractive to drive manufacturers. Combining the ability to increase areal density, along with speed insensitivity, they represent the cutting-edge of magnetic recording head technology.

SCSI Fast-40

SCSI Fast-40 is the next major performance advancement to the Small Computer System Interface (SCSI). Formerly called Ultra SCSI, SCSI Fast-40 doubles the Fast Wide SCSI-2 data transfer rate from 20 to 40 Mbytes per second.

The increased data transfer rates are attributable to the faster cycle times for data transfer and the arbitration of SCSI commands. Host systems and devices using SCSI Fast-40 will be able to negotiate optimal parameters for speed, width, offset, etc. These improvements in cycle times are primarily fueled by the higher speed of new semiconductor technologies employed in SCSI chipsets.

Easing the migration from the huge investment in other SCSI versions, SCSI Fast-40 is backward-compatible with previous generations of SCSI and uses the same physical environment. Cables, connectors, and terminators that support SCSI can support SCSI Fast-40. More importantly, SCSI Fast-40 can at least theoretically be integrated without having to modify or change operating systems.

Strategically, SCSI Fast-40 is the logical migration of SCSI to accommodate system-level technology improvements and ever-increasing demand by users for higher performance. While the infrastructure to support the Fibre Channel serial interface is developing, SCSI Fast-40 is a cost-effective solution for servers and workstations in the near term. Over time, SCSI Fast-40 will appear in personal computers as integration costs decline.

Fibre Channel Arbitrated Loop

Fibre channel is an industry-standard interface that has evolved to include electronic (non-optical) implementations and the ability to connect many devices to a host port, including disk drives, in a relatively low-cost manner. This new addition to the specifications is called Fibre Channel Arbitrated Loop (FCAL).

FCAL is a loop architecture as opposed to a bus like standard SCSI or IPI. FCAL has made it possible for Fibre Channel to be used as a direct disk attachment interface, opening new levels of I/O performance for high-throughput, performance-intensive systems such as Cray systems. SCSI-3 has been defined as the disk protocol used with FCAL.

FCAL operates a 100 MB/s and offers connectivity to 126 devices. Cray in its implementation has selected 80 drives, 40 primary and 40 alternte path, as the maximum number of drives supported on a single fibre channel loop. The FCN-1 fibre channel node has 5 FCAL loops providing expandable connectivity to 200 primary drives and 200 alternate path drives on a single controller.

Affect on Cray Product Offerings

New disk storage offerings from Cray Research will use 3.5 and 5.25 inch form factor disk drives and drive packaging that provides considerably higher numbers of Gbytes per square foot. Disk drive unit prices will tend to stabilize while $/Mbyte will continue to be lower on future products. This year’s 9 Gbyte 3.5” drives are expected to be replaced in the future by 18 Gbyte drives of the same physical size. A lower performance, highest capacity drive is expected to remain part of the product offering.

We are planning to offer a 23 Gbyte 5.25” inch drive during the next year.

As disk drives of 9 Gbyte and higher capacities become available, it becomes critical that the number of paths to the disk drives on a system be considered very carefully. Systems requiring 50 Gbytes of user space could be configured with a small number of disk drives but would also suffer I/O performance degradations due to the small number of paths to the data. The new drives will allow more data to be kept on-line, but systems will need at least as many paths to the data as with the smaller capacity drives.

Individual disk drive performance will most likely plateau in the near term. Cray Research will use enhanced controller and channel technology to maximize the I/O performance to CPU performance relationship required of a well-balanced system architecture.

Local disk arrays and network storage disk arrays will be provided with full hardware, software, and service support from Cray Research. RAID technology will be used to meet the reliability and application performance requirements.

Summary

Disk storage technology will continue to play a key role in determining the architecture and performance of future computing platforms.

Cray Research Inc. is committed to providing disk storage products that meet our customer requirements in terms of cost, performance, reliability, and capacity. This requires us to continually work closely with mass storage vendors to ensure their products meet the requirements of the supercomputing environment.