Trends in Disk Technology

Michael D. Dahlin
University of California, Berkeley
dahlin@cs.berkeley.edu

1. Disk Performance

Disk accesses rely on mechanical motions that are slow compared to electronic operations such as DRAM reads. Further, the rate of improvement in disk mechanical performance is lower than the rate of improvement of integrated circuit performance. Disk delays are dominated by three factors: seek time, rotational latency, and transfer time. During the seek time, an actuator moves the disk heads to the disk cylinder being accessed. Rotational latency allows a specific disk block to spin under the disk head, and transfer time allows data to be accessed as it passes under the disk head. Note that for physically sequential transfers, only the first block need pay seek and rotation times.

Table 1 summarizes disk performance trends for high-performance, commodity disks appropriate for workstation servers in 1987, 1990, and 1994. The 1987 and 1990 values reflect performance parameters reported by Gibson [Gibson, 1992] while the information for the more recent drive is from a product data sheet [Seagate, 1994].

During this period, seek and rotational latencies have improved slowly, at about 10% per year, while bandwidth has increased more quickly, at about 20% per year. These improvement rates are similar to those discussed by Hennessy and Patterson [Hennessy and Patterson, 1996]. Bandwidth increases more quickly than rotational speed because of increasing storage densities on the disk — more data spins under the disk head at a given rotational speed.

A consequence of disks' high seek time and rotational latency is that large transfers use disks much more efficiently than small ones. As the table indicates, latency dominates the time to access 8 KB from disk while disk bandwidth dominates the time to transfer 1 MB. However, because disk bandwidths improve more rapidly than other aspects of disk performance, increasingly large transfers are needed over time to maintain a given level of efficiency. For instance, to reduce seek and rotational latency to less than 10% of the total access time required a 562 KB sequential transfer in 1987 and a 988 KB transfer in 1994.

2. Disk Capacity and Cost

Disk capacity per unit cost improves dramatically over time — at 100% per year since 1992. As the data in Figure 1 indicate, between 1983 and 1995 the cost of disk storage has fallen by more than a factor of 400, from about $100/MB in January of 1983 to less than $0.24/MB in July of 1995. For instance, advertisements in *Byte* magazine in January of 1983 offer a 44 MB hard drive for $4,395 while advertisements in the July 1995 edition price a 9,000 MB drive at $2,159. This improvement reflects an increase of 62% per year over the twelve and one-half year period with a faster rate since 1992 as PC manufacturing volumes and competition have ramped up.

Raw Data

The data plotted in Figure 1 (disk cost and capacity over time) and Figure 2 are available at http://www.cs.berkeley.edu/~dahlin/techTrends.

References

Gibson, G. A. (1992). *Redundant Disk Arrays: Reliable, Parallel Secondary Storage*. ACM Distinguished

				Improvement Rates		
	1987	**1990**	**1994**	**1987-94**	**1987-90**	**1990-94**
Model	Fujitsu M2361A	Seagate ST-41600n	Seagate ST-15150n			
Average Seek Time	16.7 ms	11.5 ms	8.0/9.0 ms (rd/wr)	11%/yr	13%/yr	9%/yr
Rotational Speed/Avg. Latency	3600 rpm/8.3 ms	5400 rpm/5.5 ms	7200 rpm/4.2 ms	10%/yr	15%/yr	7%/yr
Bandwidth	2.5 MB/s	3-4.4 MB/s	6-9 MB/s	20%/yr	21%/yr	20%/yr
8 KB Transfer	28.3 ms	18.9 ms	13.1 ms	12%/yr	14%/yr	10%/yr
1 MB Transfer	425 ms	244 ms	123 ms	19%/yr	20%/yr	19%/yr

TABLE 1.**Disk performance parameters.** This table compares the parameters of a 1987, 1990, and 1994 commodity disk drive. Although these disks represent only a single data point for each year, I believe the trends accurately describe high-end commodity hardware available over the specified periods. Average seek time is the time it takes to seek one-third of the way across the disk surface, so it represents a pessimistic estimate of the seek time for a typical request because real requests will benefit from locality [Hennessy and Patterson, 1996]. For the 1994 disk, the seek time is higher for writes than for reads because of more stringent head settling requirements for writes. The rotational latency represents the time for the disk to complete half of a complete revolution at the given rotation speed. The bandwidth is the uncached transfer rate for a large data transfer, ignoring seek and rotational latencies. For the 1990 and 1994 drives, transfer bandwidth is higher near the edges of the disk than near the center because of the higher linear velocity of the media near the edges. To provide an upper bound on disk improvement rates, I use the more aggressive read seek time and outer-edge bandwidths when computing improvements over time. The 8 KB and 1 MB transfer times indicate the time to transfer an 8 KB or 1 MB contiguous block of data including the time for one "average" seek and rotational latency, plus the transfer size divided by the disk's maximum bandwidth.

Dissertations. MIT Press, Cambridge, Massachusettes.
Hennessy, J. and Patterson, D. (1996). *Computer Architecture A Quantitative Approach*. Morgan Kaufmann Publishers, Inc., 2nd edition.
Seagate (1994). *ST-11200N SCSI-2 Fast (Barracuda 4) Specification*. Seagate Technology, Inc.

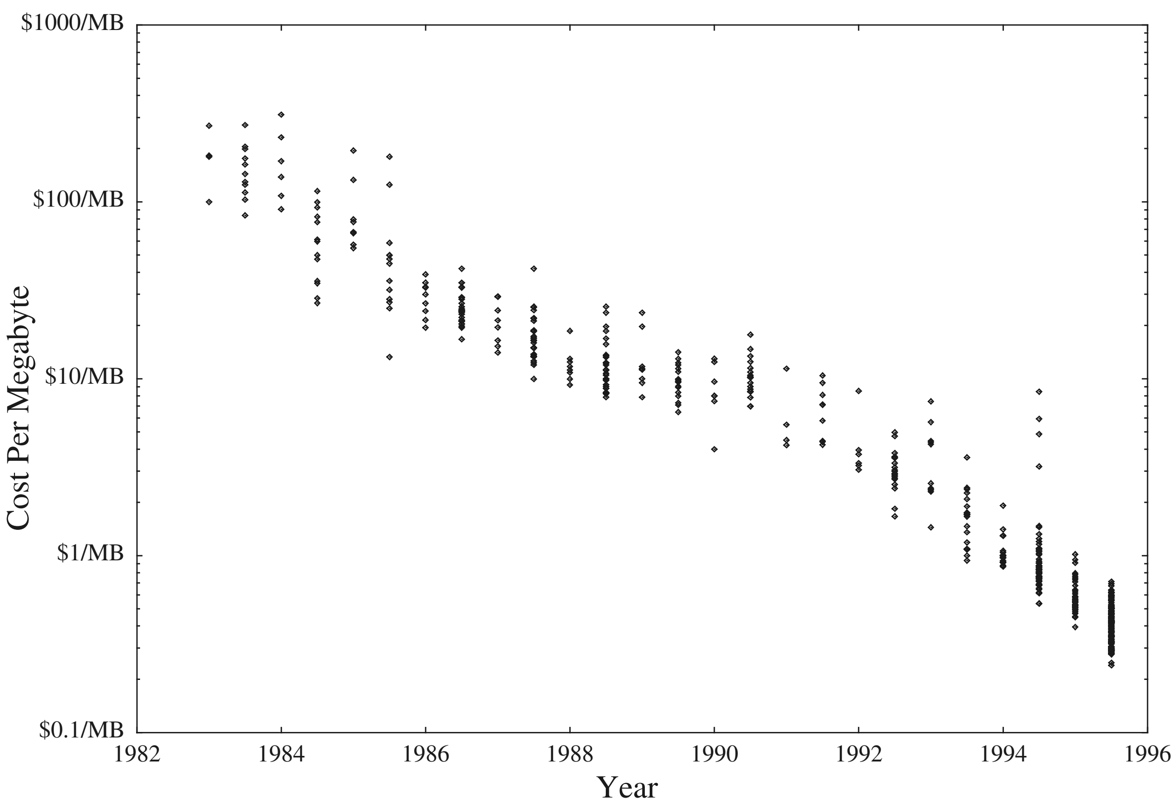

FIGURE 1.**Disk cost expressed in dollars per megabyte of capacity.** Each point indicates the capacity in megabytes divided by price in dollars for a hard disk drive advertised in the January or July edition of *Byte* magazine for the indicated year. Note the log scale for the y-axis.