In More Depth: Disk Arrays

As mentioned in Section 8.2, one method of organizing disk systems is to use arrays of smaller disks that provide more bandwidth through parallel access. In most disk arrays, the data are "striped" across the disks in the array, so that consecutive sectors can be read in parallel. In the first disk arrays, all the spindles were synchronized—sector 0 in every disk rotates under the head at the exact same time—and the arms on all the disks are always over the same track. Let's explore how such a system might work.

8.31 [20] <§§8.3–8.6> Assume that we have the following two magnetic disk configurations: a single disk and an array of four disks. Each disk has 256 sectors per track, each sector holds 1000 bytes, and the disk revolves at 10,000 RPM. Assume that the seek time is 6 ms. The delay of the disk controller is 0.1 ms per transaction, either for a single disk or for the array. Assume that only the disks and the controller limit the performance of the I/O system. Remember that the consecutive sectors on the single disk system will be spread one sector per disk in the array. Compare the performance in I/Os per second of these two disk organizations, assuming that the requests are random reads, half of which are 4 KB and half of which are 16 KB of data from sequential sectors. The sectors may be read in any order; for simplicity, assume that the rotational latency is one-half the revolution time for the single disk read of 16 sectors and the disk array read of 4 sectors. Challenge: Can you work out the actual average rotational latency in these two cases?

8.32 [10] <§§8.3–8.6> Using the same disk systems as in Exercise 8.31, with the same access patterns, determine the performance in megabytes per second for each system.