

Lecture 20: April 6

*Lecturer: Prashant Shenoy**TA: Vimal Mathew & Tim Wood*

20.1 Secondary Storage

To read from a storage device requires physical movements: the disk head must be moved to the right track, and the platter must be rotated so that the proper sector is underneath the disk head. Thus to minimize the time to make a disk access, it is necessary to minimize both the seek time and the rotational latency. Simple solutions like making the disk smaller or spin faster are generally not feasible due to the physical constraints of building a reliable disk device that has sufficient storage space. Instead, the OS and the disk driver must find a way to either schedule disk operations in order to minimize disk head movement (e.g. grouping together reads or writes to the same track), or laying out data on disk in such a way that accesses tend to be made to the same portion of the disk. Another challenge is determining the optimal sector size. Since it is possible that a seek will be needed between every disk access, having a large sector size can lead to more data being read between each seek operation. However, if it is too large, then internal fragmentation can become an issue since the sector is the smallest unit within the disk.

In the future, some of these challenges may be alleviated by the advent of solid state flash drives. SSDs do not require mechanical moving parts, so there is no seek delay or rotational latency.

20.2 Disk Head Scheduling

Disk Head Scheduling algorithms are used by the operating system to try to minimize the time wasted due to unnecessary seeks or rotational latencies. The basic idea is that since at any time there may be multiple disk requests waiting to be serviced, the OS can attempt to reorder those requests in order to reduce the length and number of seeks required.

The **First Come First Served** (FCFS) algorithm simply services requests in the order that they arrive without any reordering. This is the simplest scheduling algorithm, and its performance can be very variable depending on the disk access pattern. If requests happen to arrive in sequential order or if the system is only lightly loaded, then FCFS can perform well. However, if reads are not well ordered, performance can be very bad.

Example Requests: (head starts at 30), 65, 40, 18, 78

Order of Seeks: 65, 40, 18, 78

Distance of Seeks: $|30 - 65| + |65 - 40| + |40 - 18| + |18 - 78| = 35 + 25 + 22 + 60 = 142$

Shortest Seek Time First (SSTF) attempts to optimize the requests by reordering requests to minimize the seek time between each request. This can improve performance by minimizing each seek time, but it does not necessarily lead to an optimal overall cost. The SSTF algorithm can be implemented using a doubly linked list of requests which is sorted by seek time. However, like other greedy algorithms we have seen (e.g. Shortest Job First), it has the potential of starving requests.

Example Requests: (head starts at 30), 65, 40, 18, 78

Order of Seeks: 40, 18, 65, 78

Distance of Seeks: $10 + 22 + 47 + 13 = 92$

The **SCAN** disk head scheduling algorithm tries to alleviate some of the problems of SSTF by simply moving the disk head back and forth across the disk, servicing requests as it passes over them. For example, if there are 100 tracks on the disk it will travel from 0 to 100, then back from 100 to 0. This prevents the disk head from “bouncing back and forth” within a portion of the disk, as is possible with SSTF. A trivial optimization to SCAN is to make the disk head only scan as far as the last request in the current direction of travel. The SCAN algorithm requires only a sorted list of requests in address order.

Example Requests: (head starts at 30), 65, 40, 18, 78

Order of Seeks: (assume head is traveling to lower numbers first) 18, 40, 65, 78

Distance of Seeks: $12 + 22 + 25 + 13 = 72$

A variant of SCAN is called **Circular SCAN** or C-SCAN. This is a modification of SCAN that attempts to make wait times more consistent across requests. Note that for the normal SCAN algorithm, it is likely that once the disk head reaches the end of a disk, most of the remaining requests to be serviced will be at the other end of the disk, causing them to have a high wait time as the disk must travel back to them. C-SCAN treats the disk as a circular buffer, and thus services requests 0 to 100, then immediately loops back to service 0 to 100 again.

Example Requests: (head starts at 30), 65, 40, 18, 78

Order of Seeks: 40, 65, 78, 18

Distance of Seeks: $10 + 25 + 13 + 60 = 108$

20.3 Improving Disk Performance

Previously we discussed how the OS must maintain a free list indicating which disk sectors were free for use in new or growing files. It seems reasonable that it would make sense to allocate disk blocks in contiguous chunks as requested by the OS. Keep in mind however, that disks are always rotating, and thus if the disk may spin past an adjacent sector in a contiguous region before the request for that sector arrives. As a result, **disk interleaving** can be used to allocate blocks that are not physically contiguous, but that are “temporally contiguous” based on the rotational speed of the disk. This means that if the disk will have spun past three blocks during the time it takes for a second disk request to arrive, that blocks should be allocated with that level of separation. As a result, the disk blocks used for several different files may be interleaved in a region of the disk rather than each being given its own contiguous set of blocks.

A second performance optimization is **read ahead**. The goal of read ahead is to reduce the number of seeks needed by trying to predict which blocks a user will request and reading them into a cache ahead of time. If the disk controller guesses correctly and a request is made for one of the blocks stored in the buffer, then it can be instantly serviced without requiring a full disk access. Note that this can be done efficiently since the disk head can read in blocks as the platter rotates to service the next request. However, to be effective, the system needs an effective way of predicting which blocks should be read, particularly since the cache in the disk controller will be orders of magnitude smaller than the full disk. We saw earlier that pre-fetching virtual memory pages was difficult due to the difficulty in predicting which pages will be used next. Fortunately, with a disk it can be easier to make these predictions since many disk requests tend to access sequential blocks within a file.

20.4 Tertiary Storage

In the storage hierarchy there are registers, caches, main memory, secondary storage (e.g. hard disks), and tertiary storage. Tertiary storage devices are typically larger and cheaper than disks, but also slower. They are mainly used for archival data storage—for example tape drives or optical storage like CDs and DVDs.