

Mass-Storage Structure



Practice Exercises

- 11.1 Is disk scheduling, other than FCFS scheduling, useful in a single-user environment? Explain your answer.

Answer:

In a single-user environment, the I/O queue usually is empty. Requests generally arrive from a single process for one block or for a sequence of consecutive blocks. In these cases, FCFS is an economical method of disk scheduling. But LOOK is nearly as easy to program and will give much better performance when multiple processes are performing concurrent I/O, such as when a Web browser retrieves data in the background while the operating system is paging and another application is active in the foreground.

- 11.2 Explain why SSTF scheduling tends to favor middle cylinders over the innermost and outermost cylinders.

Answer:

The center of the disk is the location having the smallest average distance to all other tracks. Thus, the disk head tends to move away from the edges of the disk. Here is another way to think of it. The current location of the head divides the cylinders into two groups. If the head is not in the center of the disk and a new request arrives, the new request is more likely to be in the group that includes the center of the disk; thus, the head is more likely to move in that direction.

- 11.3 Why is rotational latency usually not considered in disk scheduling? How would you modify SSTF, SCAN, and C-SCAN to include latency optimization?

Answer:

Most disks do not export their rotational position information to the host. Even if they did, the time for this information to reach the scheduler would be subject to imprecision, and the time consumed by the scheduler is variable, so the rotational position information would become incorrect. Further, the disk requests are usually given in terms

of logical block numbers, and the mapping between logical blocks and physical locations is very complex.

- 11.4 Why is it important to balance file-system I/O among the disks and controllers on a system in a multitasking environment?

Answer:

A system can perform only at the speed of its slowest bottleneck. Disks or disk controllers are frequently the bottlenecks in modern systems, as their individual performance cannot keep up with that of the CPU and system bus. When I/O is balanced among disks and controllers, neither an individual disk nor a controller is overwhelmed, so that bottleneck is avoided.

- 11.5 What are the tradeoffs involved in rereading code pages from the file system versus using swap space to store them?

Answer:

If code pages are stored in swap space, they can be transferred more quickly to main memory (because swap-space allocation is tuned for faster performance than general file-system allocation). Using swap space can require startup time if the pages are copied there at process invocation rather than just being paged out to swap space on demand. Also, more swap space must be allocated if it is used for both code and data pages.

- 11.6 Is there any way to implement truly stable storage? Explain your answer.

Answer:

Truly stable storage would never lose data. The fundamental technique for stable storage is to maintain multiple copies of the data, so that if one copy is destroyed, some other copy is still available for use. But for any scheme, we can imagine a large enough disaster that all copies are destroyed.

- 11.7 It is sometimes said that tape is a sequential-access medium, whereas a hard disk is a random-access medium. In fact, the suitability of a storage device for random access depends on the transfer size. The term *streaming transfer rate* denotes the rate for a data transfer that is underway, excluding the effect of access latency. In contrast, the *effective transfer rate* is the ratio of total bytes to total seconds, including overhead time such as access latency.

Suppose we have a computer with the following characteristics: the level-2 cache has an access latency of 8 nanoseconds and a streaming transfer rate of 800 megabytes per second, the main memory has an access latency of 60 nanoseconds and a streaming transfer rate of 80 megabytes per second, the hard disk has an access latency of 15 milliseconds and a streaming transfer rate of 5 megabytes per second, and a tape drive has an access latency of 60 seconds and a streaming transfer rate of 2 megabytes per second.

- a. Random access causes the effective transfer rate of a device to decrease, because no data are transferred during the access time.

For the disk described, what is the effective transfer rate if an average access is followed by a streaming transfer of (1) 512 bytes, (2) 8 kilobytes, (3) 1 megabyte, and (4) 16 megabytes?

- The utilization of a device is the ratio of effective transfer rate to streaming transfer rate. Calculate the utilization of the disk drive for each of the four transfer sizes given in part a.
- Suppose that a utilization of 25 percent (or higher) is considered acceptable. Using the performance figures given, compute the smallest transfer size for a disk that gives acceptable utilization.
- Complete the following sentence: A disk is a random-access device for transfers larger than _____ bytes and is a sequential-access device for smaller transfers.
- Compute the minimum transfer sizes that give acceptable utilization for cache, memory, and tape.
- When is a tape a random-access device, and when is it a sequential-access device?

Answer:

- For 512 bytes, the effective transfer rate is calculated as follows.
 $ETR = \text{transfer size} / \text{transfer time}.$
 If X is transfer size, then transfer time is $((X/STR) + \text{latency}).$
 Transfer time is $15\text{ms} + (512\text{B}/5\text{MB per second}) = 15.0097\text{ms}.$
 Effective transfer rate is therefore $512\text{B}/15.0097\text{ms} = 33.12 \text{ KB/sec}.$
 $ETR \text{ for } 8\text{KB} = .47\text{MB/sec}.$
 $ETR \text{ for } 1\text{MB} = 4.65\text{MB/sec}.$
 $ETR \text{ for } 16\text{MB} = 4.98\text{MB/sec}.$
- Utilization of the device for 512B = $33.12 \text{ KB/sec} / 5\text{MB/sec} = .0064 = .64$
 For 8KB = 9.4%.
 For 1MB = 93%.
 For 16MB = 99.6%.
- Calculate $.25 = ETR/STR$, solving for transfer size X .
 $STR = 5\text{MB}$, so $1.25\text{MB/S} = ETR.$
 $1.25\text{MB/S} * ((X/5) + .015) = X.$
 $.25X + .01875 = X.$
 $X = .025\text{MB}.$
- A disk is a random-access device for transfers larger than K bytes (where $K > \text{disk block size}$), and is a sequential-access device for smaller transfers.
- Calculate minimum transfer size for acceptable utilization of cache memory:
 $STR = 800\text{MB}$, $ETR = 200$, $\text{latency} = 8 * 10^{-9}.$
 $200 (X\text{MB}/800 + 8 * 10^{-9}) = X\text{MB}.$
 $.25X\text{MB} + 1600 * 10^{-9} = X\text{MB}.$
 $X = 2.24 \text{ bytes}.$

Calculate for memory:

$$\text{STR} = 80\text{MB}, \text{ETR} = 20, L = 60 * 10^{-9}.$$

$$20 (\text{XMB}/80 + 60 * 10^{-9}) = \text{XMB}.$$

$$.25\text{XMB} + 1200 * 10^{-9} = \text{XMB}.$$

$$X = 1.68 \text{ bytes}.$$

Calculate for tape:

$$\text{STR} = 2\text{MB}, \text{ETR} = .5, L = 60\text{s}.$$

$$.5 (\text{XMB}/2 + 60) = \text{XMB}.$$

$$.25\text{XMB} + 30 = \text{XMB}.$$

$$X = 40\text{MB}.$$

- f. It depends on how it is being used. Assume we are using the tape to restore a backup. In this instance, the tape acts as a sequential-access device where we are sequentially reading the contents of the tape. As another example, assume we are using the tape to access a variety of records stored on the tape. In this instance, access to the tape is arbitrary and hence considered random.

- 11.8 Could a RAID level 1 organization achieve better performance for read requests than a RAID level 0 organization (with nonredundant striping of data)? If so, how?

Answer:

Yes, a RAID level 1 organization could achieve better performance for read requests. When a read operation is performed, a RAID level 1 system can decide which of the two copies of the block should be accessed to satisfy the request. It could base this choice on the current location of the disk head and could therefore optimize performance by choosing the disk head that is closer to the target data.

- 11.9 Give three reasons to use HDDs as secondary storage.

Answer:

HDDs are still the most common secondary storage device.

- They are the largest random-access storage device for the price, providing terabytes of storage at a low cost.
- Many devices, from external chassis to storage arrays, are designed to use HDDs, providing a wide variety of ways to use HDDs.
- They maintain the same read and write performance over their lifetime, unlike NVM storage devices, which lose write performance as they get full and as they age.

- 11.10 Give three reasons to use NVM devices as secondary storage.

Answer:

NVM devices are increasing in size and decreasing in price faster than HDDs.

- High-speed NVM devices (including SSDs and usually not including USB drives) are much faster than HDDs. Secondary storage speed has a large impact on overall system performance.

- b. NVM devices use less power than HDDs, making them very useful for laptops and other portable, battery-powered devices. NVM devices can also be much smaller than HDDs and thus can, for example, be surface-mounted to motherboards in devices like smartphones.
- c. Because NVM devices have no moving parts, they tend to be much more reliable than HDDs.

