Homework #7 Solutions

CMSC 611, Spring 2000

1. Problem 1

We begin by completing the table given in the problem:

Cost	**Rotation rate**	**Seek (avg)**	**Tracks**	**Surfaces**	**Sectors per track**	**Bytes per sector**	**Capacity**	**Rotation latency**	**Sector transfer**
$200	7200	8	5000	4	500	512	5000*4*500*0.5=5GB	1/120*0.5=4.17ms	(1/120)*(1/500)=16.7us
$500	10000	6	10000	6	1000	512	10000*6*1000*0.5=30GB	60/10000*0.5=3ms	6ms*(1/1000)=6us

[a] A single 32 KB request can require two different times, depending on which disk is used. The “overhead” in the system is the same for both disk types:
CPU: 50,000 instructions / 500 million instructions per second = 0.1 ms
Controller: 1 ms
I/O bus: 32 KB / 40 MB/s = 0.8 ms
CPU backplane: 32 KB / 200 MB/s = 0.16 ms
Total overhead = 0.1 + 1 + 0.8 + 0.16 = 2.06 ms

The slower (smaller) disks require 4.17 ms (rotation) + 8 ms (average seek) + 64*0.0167 (64 sectors @ 512 bytes each) = 13.24 ms locally, for a total time of **15.30 ms for a single 32 KB request**.

The faster (larger) disks require 3 ms (rotation) + 6 ms (average seek) + 64*0.006 (64 sectors @ 512 bytes each) = 9.38 ms locally, for a total time of **11.44 ms for a single 32 KB request**.

[b] The maximum capacity of a system using these devices is limited by the number of disks that can be used. The system can have up to 6 controllers, each of which can support 2 buses of 15 drives each. This allows a total of 180 drives, each of which could have up to 30 GB. The maximum capacity is thus 30*180 = **5400 GB = 5.4 TB**.

[c] The maximum data rate for the system can be found by considering the limit placed on the system by each component. The CPU (OS) requires 0.1 ms/request, so can handle 10,000 requests/second. Each controller requires 1 ms per request, and can thus handle 1000 requests/second. An I/O bus can handle 1000/0.8 = 1250 requests/second. Each large disk (they’re faster) can handle 1000/9.38 = 106.6 requests/second. Clearly, the I/O buses are not the bottleneck — each is faster than the controller. The disks on a single bus can do 15*106.6 > 1000 requests per second. Thus far, the controller is the bottleneck. We can have up to 6 controllers in the system, for an aggregate of 6000 requests per second. This is lower than the CPU’s limit of 10,000, so the total number is 6000 requests/second. Because each request is 32 KB, the maximum throughput is 32*6000 = **192,000 KB/sec = 192 MB/sec**.

[d] Clearly, we need to use as many large disks as possible for capacity; however, the smaller disks have higher IOs/second per dollar. If we build a system that has a capacity of 300 GB from large disks only, it’ll have 10 large disks and a request rate of 106.6*10 = 1066 requests per second. We need 434 more requests per second; for each large disk we replace with 6 small disks (to keep the capacity the same), we gain 6*(1000/13.24) - 106.6 = 346.4 requests per second. If we have 9 large disks and 6 small ones, we’ll have the 300 GB we need, but only 1412 requests per second. We can add the remainder with 2 small disks more cheaply than by replacing a large disk with 6 small ones or than simply adding another large disk. The final configuration is thus 9 large disks and 8 small disks, for a total request rate of 9*106.6 + 8*75.5 = 1563 requests per second. Total capacity is 310 GB — more than enough for the minimum requirements. We need two controllers each with one I/O bus to hold all of the disks; a single controller would be limited to 1000 requests/second.

Total cost for the I/O system is 2*$1000 (controllers) + 9*$500 + 8*$200 = **$8100**.

2. Problem 2

[a] A single tape contains 40 GB, and can be read at 9 MB/s. Total time to read a tape is 40000MB/9MB/s = **4,444 seconds**, or nearly 75 minutes!

[b] A single tape could be swapped in 30 seconds, so total time to load and read a tape is 4474 seconds. The system has 8 tape readers, and needs to read 6000 tapes, so each reader must read 750 tapes. This will take 4474 * 750 = **3.335 million seconds**, or nearly 1.5 months. It doesn’t include downtime for repairs....

[c] Seek time is randomly distributed between 0-60 seconds, for an average of 30 seconds. It takes 30 seconds to load the tape as well, for “seek” time of 60 seconds. Reading 200 MB at 9 MB/s takes 200/9 = 22.2 seconds, for a total time of **82.2 seconds** to read a 200 MB file from tape.

[d] If tapes were replaced by fast-seeking disks, the time would drop to **52.2 seconds**, for a savings of 82.2/52.2 = 1.57 times faster.

3. Problem 3

[a] A 6000 tape system with 4 readers will cost $250,000 + $20,000*4 + $50*6000 = $630,000. Its capacity is 30*6000 = 180,000 GB = 180 TB.

[b] A disk system of the same capacity would be built from single-hub units with 16 PCs and 64 disks. Total cost for this would be $1000 (hub) + 16*$400 (PCs) + 64*$200 = $20,200 for 64*30 = 1920 GB = 1.92 TB. We want 180 TB, which will require 180/1.92 = 93.75 of these units. This will cost 93*20,200 + $1000 + 12*$400 + 48*$200 = $1,894,000.

[c] If the disk system needed 180 TB * 8/7 = 205.7 TB of raw capacity, it would cost (205.7/1.92) * $20,200 = $2.164 million.

[d] The inexpensive disks from Problem 1 could each do 75.5 * 32 KB = 2.4 MB/s on relatively small I/Os. The aggregate bandwidth would thus be 2.4 * 205700/30 = 16,456 MB/s = 16.46 GB/s! Of course, this would be limited by the network throughput of the hubs, and perhaps by the performance of the individual PCs. Even so, 16+ GB/s is quite fast, particularly when compared to the 9 MB/s * 4 = 36 MB/s of the tape system.