CMU 18-746	Storage Systems	Assigned: 13 Jan 2009
Spring 2010	**Homework 1**	**Due: 3 Feb 2009**

Solutions are due at the beginning of class on the due date and must be **typed** and neatly organized. Late homeworks will not be accepted. You are permitted to discuss these problems with your classmates; however, your work and answers must be your own. All questions should be directed to the teaching assistant.

Problem 1 : Finding service time.

Find the service time for the following request stream to a very simple disk that has 10000 cylinders, 8 surfaces, 200 sectors per track, and rotates at 10000 RPM. In this case, LBNs are mapped directly to PBNs, requests are serviced in FIFO order, and seek time is a linear function of cylinder distance.

$$\text{seek\_time(ms)} = 0.0006 \times \text{cylinder\_distance} + 2$$

(Remember that a linear seek curve is not realistic, but is a simplification for this problem.) Assume that the disk starts at LBN 0 and that head switches are instantaneous. The layout mapping is shown in Figure 1.

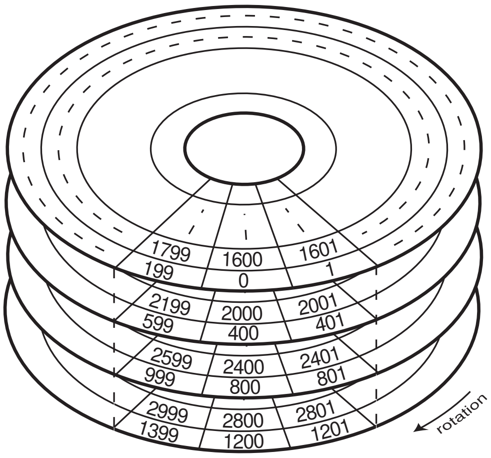

Figure 1: *LBN mappings onto physical sectors on the disk.*

The request stream for which you should compute the service times is shown in Table 1.

Request Stream	
LBN	Size (sectors)
0	4
9600600	8
4801400	4
11200948	16

Table 1: *Request stream.*

This problem is continued on the following page.

Show your work by making a timetable that looks like Table 2. Time (ms) shows the time at which the head is at the given LBN. The LBN is mapped to a cylinder, surface, and offset within the track. The "next event" field shows which part of servicing the request (seek, rotation, or transfer) takes place next.

time (ms)	LBN	cylinder	surface	offset	next event
0	0	0	0	0	transfer 4 sectors
??	4	0	0	4	seek

Table 2: *Sample timetable.*

Problem 2 : Media bandwidth and data buffering.

Modern disks include some amount of on-board memory which is used for keeping firmware state, for speed-matching between the media and the bus, and for buffering data to tolerate host and bus delay. Suppose that a disk streams data at 20MB per second and has 1MB of buffer memory. What is the maximum bus delay that can be tolerated before a media transfer must be interrupted?

Problem 3 : Capacity benefit of zoned recording.

For the disk whose specs are given in the Table 3, find the capacity when zoned recording is used and when it is not. Explain your reasoning.

Seagate Cheetah 4LP		
Cylinders		6582
Surfaces		8
Zone Information		
	cylinder first–last	sectors per track
Zone 1	0–1343	195
Zone 2	1345–2448	187
Zone 3	2450–3541	175
Zone 4	3543–4406	165
Zone 5	4408–5223	154
Zone 6	5225–5956	145
Zone 7	5958–6580	131

Table 3: *Specifications for the Seagate Cheetah 4LP.*

Notice that there is a spare cylinder at the end of each zone. These spares are reserved for remapped bad cylinders and so should *not* count toward the data capacity of the disk.

Problem 4 : Rotational latency vs. seek time.

Disk drive vendors often use average seek time and rotation speed as metrics for disk drive performance. In this problem, we'll compare two real disks and see whether these are the best metrics for measuring disk performance.

Assume a disk access pattern wherein a read request to an LBN in the range $\langle 0, 128 \rangle$ is followed by a read request to the LBN range $\langle 3720000, 3720128 \rangle$. Each request is 2 sectors long and the sectors always fall within the specified ranges. For the disks whose specifications are given in Table 4 and Figure 2, calculate the average service time for each request. Explain the reasoning for your derivation. Which metrics (rotation speed, seek profile, etc.) contributed to the improved performance of one disk over the other?

To simplify your calculations, assume that every sector on the disk is mapped to an LBN. Ignore all other overheads, such as head switch time and bus transfer time. Also, it is okay to eyeball the seek time from the seek curve. Finally, since we are not telling you whether the disks are zero-latency access, what the track and cylinder skew is, or what the distribution of the requests within the ranges is, make a reasonable assumption for calculating rotational latency.

	Seagate Cheetah4LP		IBM Ultrastar18ES	
Year	1996		1998	
Form factor	3.5” half-height		3.5” half-height	
Capacity	4.5 GB		9 GB	
Cylinders	6582		11474	
Surfaces	8		5	
Spindle speed	10033		7200	
Zone Information				
	firstcyl–lastcyl	*sectors per track*	*firstcyl–lastcyl*	*sectors per track*
Zone 1	0–1343	195	0–377	390
Zone 2	1345–2448	187	378–1263	374
Zone 3	2450–3541	176	1264–2247	364
Zone 4	3543–4406	166	2248–3466	351
Zone 5	4408–5223	155	3467–4504	338
Zone 6	5225–5956	145	4505–5526	325
Zone 7	5958–6580	131	5527–7044	312
Zone 8			7045–8761	286
Zone 9			8762–9815	273
Zone 10			9816–10682	260
Zone 11			10683–11473	247

Table 4: *Specifications for the Seagate Cheetah 4LP and IBM Ultrastar 18ES.*

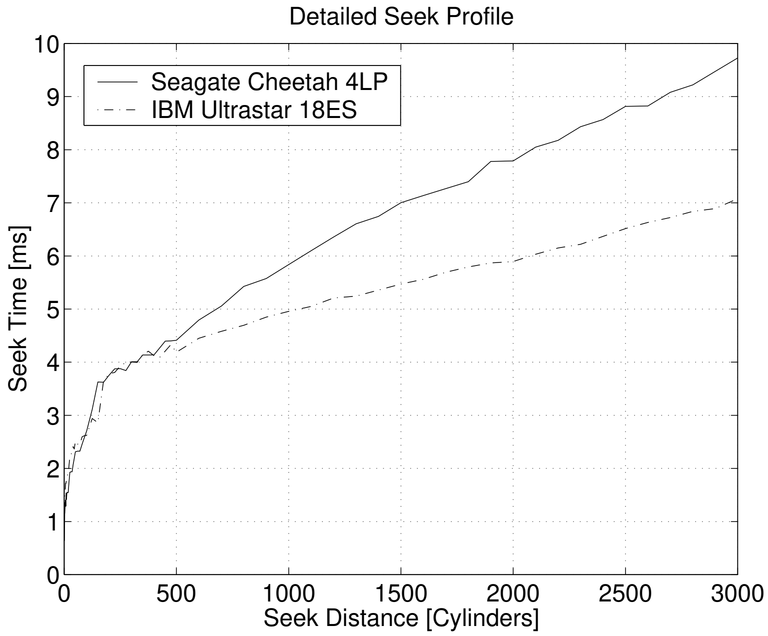

Figure 2: *Seek profile for the Seagate Cheetah 4LP and IBM Ultrastar 18ES.*

Problem 5 : Zero-latency access.

Disk drive access includes three main components: seek, rotational latency and data transfer. Once a seek is completed, a normal disk only transfers in ascending LBN order. As shown in Figure 3, the disk must wait until the first sector arrives at the head to begin the transfer, even if other sectors of the request are already within the bounds of the request.

Figure 3: *In a normal disk, a transfer can only start at the first sector of a request.*

In a zero-latency disk, the transfer can be done in any order, as shown in Figure 4. This allows the transfer to begin as soon as the head is over *any* portion of the request. This can avoid part of the rotational latency that is incurred in a normal disk.

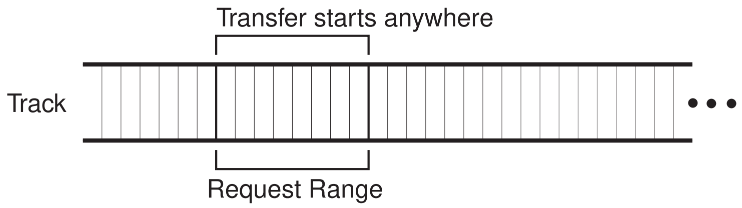

Figure 4: *In a zero-latency disk, a transfer can start anywhere within the request range.*

(a) Imagine that a disk is servicing a request of size S, and that the seek has completed and the head is on track. Give a general formula for the expected number of revolutions for both the non-zero-latency disk and the zero-latency disk once the head is on track. Your answers should be in terms of N and S, where N is the number of sectors per track and S is the size of a request in sectors with $1 \leq S \leq N$. Explain your reasoning.

(b) Plot the expected number of revolutions for both disks as a function of request size S for a disk with $N = 390$.

To help you with solving this problem here are few simplifications and hints:

- Assume that the request of size S is always within one track and never crosses a track boundary.
- Since the head has already arrived on track, remember to ignore the seek portion of the service time – it is just a constant that is added to the service time independent of the request size S or the track size N.
- In expressing time in revolutions, remember that in one revolution N sectors pass under the head.

Problem 6 : A glimpse into the future.

Every year, the disk drive industry releases new disks that are faster, bigger (capacity), and smaller (area, volume, and price). As with many areas of the computer industry, evaluating past trends has some success of predicting future characteristics.

Using data from your course notes (or any other sources), predict the characteristics of the future disks of 2010 and 2015. You should report values for average rotational latency, average seek time, capacity, media transfer speed, and price per gigabyte. You should additionally show how you arrived at these predictions: What reference values did you use? What were the trends for each value? Where did you find this information?