CS 537: Intro to Operating Systems (Summer 2017)

Worksheet 12 - Disk Scheduling

DUE: Aug 2nd 2017 (Wednesday)

This question examines a disk's internal scheduler as a black box, to see if we can learn anything about its behavior. The inputs we give to the disk are a bunch of requests; the outputs we observe are the order the requests are serviced.

For example, we might issue requests to blocks 0, 100, and 200 at the same time. The disk might decide to service 100, then 200, and then 0, depending on its internal scheduling algorithm.

We also know some details of the disk. It has one surface with 100 tracks, each of which have 100 sectors. The outer track (track=0) contains sectors 0...99, the next track (track=1) contains sectors 100...199, and so forth.

In this question, each question gives a workload to a particular disk; you are then asked to determine what you think the scheduler of that disk is doing by servicing requests in that order. In other words, what can you say about the algorithm in use? What about the initial state of the disk, before the requests arrived?

a. Disk Model A

```
Requests           : 0, 1, 2, 3, 4, 5
Completed in order : 0, 1, 2, 3, 4, 5
```

What do you think Disk A's scheduler is doing? What might the initial state of the disk have been?

b. Disk Model B

```
Requests           : 0, 500, 200, 400, 300, 100
Completed in order : 0, 100, 200, 300, 400, 500
```

What do you think Disk B's scheduler is doing? What might the initial state of the disk have been?

c. Disk Model C

```
Requests           : 0, 500, 200, 400, 300, 100
Completed in order : 200, 300, 400, 500, 100, 0
```

What do you think Disk C's scheduler is doing? What might the initial state of the disk have been?

d. Disk Model D

```
Requests           : 0, 50, 110, 600
Completed in order : 0, 110, 50, 600
```

What do you think Disk D's scheduler is doing? What might the initial state of the disk have been?

e. Could these disks (A through D) actually be using the same scheduler? If so, why? If not, why not?