Prepare your answers to the following questions either in a plain text file or in a file that can be opened with Microsoft Word. Submit your file to the Curator system (www.cs.vt.edu/curator) under the heading HW2, by the posted deadline for this assignment. No late submissions will be accepted.

1. Assume a system uses a hard drive with the following physical characteristics:

total capacity	256 GB
# of platters	4
# of tracks per surface	16384
avg # of sectors per track	8192
cluster size	4 KiB
spindle speed	10000 RPM
head start time	0.1 ms
track to track seek time	0.01 ms

The sector size is 512 bytes. In answering the following questions, express all final time values in milliseconds, with precision 4 (e.g., 8.1234 ms).

a) [15 points] What is the average random head seek time for this drive?

b) [15 points] What is the average rotational latency for this drive?

c) [10 points] What is the average total time required to read one randomly-chosen sector from this drive?

d) [10 points] What is the average total time required to read a file of 10 MiB from this drive if the clusters are randomly scattered on the drive?

2. Assume each of the changes described below were made to the design of the drive described in the preceding question, and that no other changes were made. Indicate the effect of the change on the average random head seek time and the average rotational latency time for the drive. The parts are independent. Justify your conclusions if you want credit.

a) [10 points] doubling the rotational speed, and modifying the read/write heads so that they can keep up with the increased transfer rate

b) [10 points] reducing the track spacing by half and doubling the number of tracks per surface (which would alter the total capacity of the device); the track-to-track seek time remains the same

3. Let T be a PR quadtree storing data points that lie in a square with side *S*. According to the survey paper by Samet, if *D* is the minimum distance between any two of the data points, the maximum number of levels in a PR quadtree is:

$$\left\lceil \log\left(S\sqrt{2} / D \right) \right\rceil$$

Suppose that the points in a given data set lie within a square whose side is 2^{20} units, and that the closest pair of data points are a distance of 2^3 units apart.

a) [10 points] What is the maximum height of a PR quadtree that stores the entire data set?

b) [10 points] Explain why the PR quadtree might actually be much shorter than the bound you gave in the previous part.

c) [10 points] Explain, in one clear sentence, precisely why the number $\sqrt{2}$ gets to be involved in the formula for the bound.