ECE 18-316 INTRO TO DATA STORAGE FALL 98 <u>PROBLEM SET #1</u> Due Friday, 9/3/98 In Class or To Jie Zou Before Start of Lab Section (1:30 PM) Late submissions will not get credit

1. <u>Memory Hierarchy</u>

Assume we are building a machine with a memory hierarchy for instructions (do not worry about data accesses)! Assume that the program follows the 90-10 rule, i.e. that 90% of the accessing time is spread evenly over 10% of the code and the other 10% of the time is spread evenly over the other 90% of the code. You have three types of memory for use in your memory hierarchy:

MEMORY TYPE	ACCESS TIME	COST PER WORD
Local, fast	1 clock cycle	\$0.10
Main	5 clock cycles	\$0.01
Disk	5000 clock cycles	\$0.0001

You have exactly 100 programs, each is 1,000,000 words, and all the programs must fit on disk. Assume that only one program runs at a time, and that the whole program must be loaded in main memory. You can spend \$30,000 on the memory hierarchy.

a) What is the optimal way to allocate your budget assuming that each word must be statistically placed in fast memory or main memory?

b) Ignoring the time for the first loading from disk, what is the average number of cycles for a program to make a memory reference in your hierarchy? (This important measure is called the average memory-access time.)

2) <u>Areal Density and Data Rates</u>

Imagine you come home to watch Star Trek and realize that you did not watch the episode you taped last week. You are now presented with a dilemma. You can:

1) Watch the new one then watch the old one in the wrong order

2) Tape the new one, kill an hour, and watch them in sequence later

3) Tape the new one, while going to someone else's house to watch the old one

4) Have two VCR's and watch the old one while taping the new one

... OR ...

5) Use data storage technology to watch the old one while taping the new one... ...onto the same tape.

Philips Research, in the Netherlands, has actually prototyped a system that could accomplish the above task by using a disk drive as a buffer to a tape drive. The disk drive is the input and output stream to the video system, and must have sufficient data rate to handle two sustained compressed digital data streams simultaneously, one incoming and one outgoing, plus the buffer interface stream with the tape device, operating at 3x the video data rate (compressed). The tape drive has sufficient capacity to record two full hours of compressed digital video onto each cartridge, which would not be cost effective with the disk drive alone. It should be noted that the tape drive could handle the sustained data rate of the two data streams without the presence of the disk drive, but, in this configuration, the system would have undesirable delays in searching for data on the tape.

The system, as demonstrated, works by dumping information to the tape and retrieving it at about 3 times the data rate necessary for viewing in real time. In this way it can take the last 15 minutes of digital recording (the new episode) and dump it to tape in 5 minutes, while continuing to record in real time. Following this dump, it repositions the tape to any new location (this must take less than 2.5 minutes). Then, in the next five minute block, it can retrieve the next 15 minutes of the episode you are currently viewing (the previously recorded "old" episode) and then show it to you in real time, while still recording, also in real time, and playing the current 15 minute block. Finally, it has another 2.5 minutes to reposition itself for the next dump.

This sequence can be repeated:

- a) dump 15 minutes of video in 5 minutes to tape
- (while playing and recording with the disk drive)
- b) reposition for 2.5 minutes

(possibly a complete rewind of the tape)

c) retrieve 15 minutes of video in 5 minutes

(while playing and recording with the disk drive) d) reposition for up to 2.5 minutes (possibly another complete rewind of the tape)

In such a configuration, the disk drive needs to sustain a data transfer rate of 5x the real-time rate: 3x for the high speed dumping and retrieving + 1x for viewing and 1x for recording. The tape drive needs to be able to go end to end in search mode in less than 2.5 minutes (worst case scenario), and sustain a transfer rate of 3x the rate of real time compressed digital video.

Your assignment is to verify that the parameters of a proposed system will satisfy these requirements, and select appropriate values for any parameters that are unspecified. These are as follows (by way of review):

- A disk drive that can sustain transfer rates of 5x compressed digital video rate at the ID of the disk (worst case)
- A disk drive that can store at least 1 hour of digital video (one 15 minute block being recorded, one block being played, one block being dumped to tape, and one block being retrieved from tape)
- A tape system that can store 2 hours of compressed digital video on one cartridge
- A tape system that can sustain transfer rates of 3x compressed digital video rate
- A tape drive that can rewind in search mode in less than 2 minutes, end to end.

The attached worksheet has blanks for you to fill in and see what it would take for this system to fulfill the above criteria. Please complete all of the blank areas, while using the above scenario as a guide. Pay careful attention to units.