MC48 Final Exam

Serial #:

This exam is an open-book, open-notes, individual-work exam. Please write your name only on this page. Do not turn the page until instructed, in order that everyone may have the same time. Then, be sure to look at all problems before deciding which one to do first. Some problems are easier than others, so plan your time accordingly. You have two hours to work.

Write the answer to each problem on the page on which that problem appears. You may use the back of that same sheet if you need additional space. You may also request additional paper.

Name: ________________________________

Problem	Page	Possible	Score
1	2	12	
2	3	12	
3	4	12	
4	5	12	
Total		48	

1. [**12 Points**] The MIPS R10000 processor has a primary data cache that can hold 32 kilobytes of data organized in a two-way set associative fashion, with 32 bytes per cache block. Suppose a cache like this was used in a byte-addressed system with 32-bit addresses; what would the total size in bits be of the cache, including the valid and tag bits?

2. [**12 Points**] A particular disk drive has an 11 ms average seek time, and because it rotates at 7200 RPM, has an average rotational latency of 4.17 ms. It has a transfer rate of 50 megabytes/second. The disk is used to read 32 kilobyte blocks, each randomly located relative to its predecessor. To the nearest integer, what is the average number of megabytes per second that the disk can deliver?

3. [**12 Points**] Give a rationale in support of each of the following statements:

   (a) SIMD machines can be built more cheaply per processor than MIMD machines.

   (b) MIMD machines can be built more cheaply per processor than SIMD machines.

   (c) SIMD machines are easier to program than MIMD machines.

   (d) MIMD machines are easier to program than SIMD machines.

4. [**12 Points**] A new system being designed has separate 64KB direct-mapped caches for instructions and data, just like the DecStation 3100 used as an example in section 7.2. However, the designers are considering changing either or both of the caches to use a four-word block size, as opposed to the single-word blocks used in the DecStation 3100. The system's miss penalty will be one cycle per word in the block plus one cycle more, i.e., a total of two cycles for a single-word block and five cycles for a four-word block. Based on the data in figure 7.10 on page 526, will the system be faster if the instruction cache uses one-word blocks or four-word blocks? If the data cache uses one-word blocks or four-word blocks? Assume all other variables, such as the clock rate, are held constant.