1. For the following set of variables {CPI, clock rate, cycle time, MIPS, I, C}, identify all of the subsets that can be used to calculate execution time. Each subset should be minimal; that is, it should not contain any variable that is not needed. An example subset is {CPI, I, cycle time}. Let I = *number of instructions in program* and C = *number of cycles in program*. (15%)

2. One extension of the MIPS instruction set architecture has two new instructions called `movn` (move if not zero) and `movz` (move if zero). For example, the instruction

   ```
   movn $8, $11, $4
   ```

   copies the contents of register 11 into register 8, provided that the value in register 4 is nonzero (otherwise it does nothing). The `movz` instruction is similar but copying takes place only if the register's value is zero. Show how to use the new instructions to put whichever is larger, register 8's value or register 9's value, into register 10. If the values are equal, copy either into register 10. You may use register 1 as an extra register for temporary use. Do not use any conditional branches. (10%)

3. In the embedded market, where cost is crucial, processors sometimes implement floating point only in software. We are interested in two implementations of a computer, one with and one without special floating-point hardware. Consider a program, P, with the following mix of operations:

Floating-point multiply	10%
Floating-point add	15%
Floating-point divide	5%
Integer instructions	70%

   Computer MFP (computer with floating point) has floating-point hardware and can therefore implement the floating-point operations directly. It requires the following number of clock cycles for each instruction class:

Floating-point multiply	6
Floating-point add	4
Floating-point divide	20

Integer instructions 2

Computer MNFP (computer with no floating point) has no floating-point hardware and so must emulate the floating-point operations using integer instructions. The integer instructions all take 2 clock cycles. The number of integer instructions needed to implement each of the floating-point operations is as follows:

Floating-point multiply	30
Floating-point add	20
Floating-point divide	50

Both computers have a clock rate of 1000 MHz.

(a) Find the native MIPS ratings for both computers. (10%)

(b) If the computer MFP needs 300 million instructions for this program, how many integer instructions does the computer MNFP require for the same program? (10%)

(c) Assuming the instruction counts from (b), what is the execution time (in seconds) for the program run on MFP and MNFP? (5%)

4. Two important hurdles make parallel processing challenging.
   (a) The limited parallelism available in programs:
   Suppose you want to achieve a speedup of 80 with 100 processors. What fraction of the original computation can be sequential? (10%)
   (b) The relatively high cost of communications:
   Suppose we have an application running on a 32-processor multiprocessor, which has a 200ns time to handle reference to a remote memory. For this application, assume that all the references except those involving communication hit in the local memory hierarchy. Processors are stalled on a remote request, and the processor clock rate is 2 GHz. If the base CPI (assuming that all references hit in the cache) is 0.5, how much fast is the multiprocessor if there is no communication versus if 0.2% of the instructions involve a remote communication reference? (10%)

5. Average memory access time (AMAT) is commonly used to examine alternative cache designs. Average memory access time is the average time to access memory considering both hits and misses and the frequency of different accesses:

$$\text{AMAT} = \text{Hit Time} + \text{Miss rate} \times \text{Miss penalty}$$

Consider a processor with a 2 ns clock, a miss penalty of 20 clock cycles, a miss rate of 0.05 misses per instruction, and a cache access time (including hit detection) of 1 clock cycle. Assume that the read and write miss penalties are the same and ignore other write stalls.

(a) Find the AMAT for the processor (5%)

(b) Suppose we can improve the miss rate to 0.03 misses per reference by doubling the cache size. This causes the cache access time to increase to 1.2 clock cycles. Using the AMAT as a metric, determine if this is a good trade-off (10%)

6. Consider three processors with different cache configurations
   - Cache 1: Direct-mapped with one-word blocks
   - Cache 2: Direct-mapped with four-word blocks
   - Cache 3: Two-way set associative with four-word blocks

   The following miss rate measurements have been made:
   - Cache 1: Instruction miss rate: 4%; Data miss rate: 6%
   - Cache 2: Instruction miss rate: 2%; Data miss rate: 4%
   - Cache 3: Instruction miss rate: 2%; Data miss rate: 3%s

   For these processors, one-half of the instructions contain a data reference. Assume that the cache miss penalty is 6 + Block size in words. The CPI for this workload was measured on a processor with Cache 1 and was found to be 2.0. Determine which processor spends the most cycles on cache misses. (15%)