# Cpr E 381 Homework 6

1. Problem 3.30. The Big Picture on page 216 mentions that bits have no inherent meaning. Given the bit pattern: 1010 1101 0001 0000 0000 0000 0000 0010, what does it represent, assuming that it is
    a. A two's complement integer?
    b. An unsigned integer?
    c. A single precision floating-point number?
    d. A MIPS instruction?

2. Show the IEEE 754 binary representation in both single and double precision floating point for the following.
    a. 20
    b. 20.5
    c. 0.1
    d. -5/6
    e. $6.022 \times 10^{23}$
    f. $\pi$
    g. -0.0
    h. + infinity

3. Problem 3.35. Add $2.85_{ten} \times 10^3$ to $9.84_{ten} \times 10^4$, assuming that you have only three significant digits, first with guard and round digits and then without them.

4. Problem 3.44. The IEEE 754 floating-point standard specifies 64-bit double precision with a 53-bit significand (including the implied 1) and an 11-bit exponent. IA-32 offers an extended precision option with a 64-bit significand and a 16-bit exponent.
    a. Assuming extended precision is similar to single and double precision, what is the bias in the exponent?
    b. What is the range of numbers that can be represented by the extended precision option?
    c. How much greater is this accuracy compared to double precision?

5. Problem 3.46. While the IA-32 allows 80-bit floating-point numbers internally, only 64-bit floating-point numbers can be loaded or stored. Starting with only 64-bit numbers, how many operations are required before the full range of the 80-bit exponents are used? Give an example.

6. Problem 4.7. Suppose you wish to run a program P with $7.5 \times 10^9$ instructions on a 5 GHz machine with a CPI of 0.8.
    a. What is the expected CPU time?
    b. When you run P, it takes 3 seconds of wall clock time to complete. What is the percentage of the CPU time P received?

7. Problem 4.8. Consider two different implementations, P1 and P2, of the same instruction set. There are five classes of instructions (A, B, C, D, and E) in the instruction set.

   P1 has a clock rate of 4 GHz, P2 has a clock rate of 6 GHz. The average number of cycles for each instruction class for P1 and P2 is as follows:

| Class | CPI on P1 | CPI on P2 |
|---|---|---|
| A | 1 | 2 |
| B | 2 | 2 |
| C | 3 | 2 |
| D | 4 | 4 |
| E | 3 | 4 |

   Assume that peak performance is defined as the fastest rate that a computer can execute any instruction sequence. What are the peak performances of P1 and P2 expressed in instructions per second?

8. Problem 4.9. If the number of instructions executed in a certain program is divided equally among the classes of instructions in Exercise 4.8 (previous problem) except for class A, which occurs twice as often as each of the others, how much faster is P2 than P1?

9. Problem 4.11. Consider program P, which runs on a 1GHz machine M in 10 seconds. An optimization is made to P, replacing all instances of multiplying a value by 4 (multi x, x, 4) with two instructions that set x to x+x twice (add x, x, x; add x, x, x). Call this new optimized program P'. The CPI of a multiply instruction is 4, and the CPI of an add is 1. After recompiling, the program now runs in 9 seconds on machine M. How many multiplies were replaced by the new compiler?

10. Problem 4.12. Your company could speed up a Java program on their new computer by adding hardware support for garbage collection. Garbage collection currently comprises 20% of the cycles of the program. You have two possible changes to the machine. The first one would be to automatically handle garbage collection in hardware. This causes an increase in cycle time by a factor of 1.2. The second would be to provide for new hardware instructions to be added to the ISA that could be used during garbage collection. This would halve the number of instructions needed for garbage collections but increase the cycle time by 1.1. Which of these two options, if either, should you choose?