

# CSCE614 Computer Architecture (Spring 2011)

## Assignment #1

Due: 2/15 11:10AM

1. Consider two different implementations, I1 and I2, of the same instruction set. There are three classes of instructions (A, B, and C) in the instruction set. The clock rates of I1 and I2 are 6GHz and 3 GHz, respectively. The average number of cycles for each instruction class on I1 and I2 is given in the following table:

Class	CPI on I1	CPI on I2	C1 usage	C2 usage	C3 usage
A	2	1	30%	40%	60%
B	3	2	50%	20%	15%
C	5	2	20%	40%	25%

The table contains a summary of average proportion of instruction classes generated by three different compilers. C1 is a compiler produced by the makers of I1, C2 is produced by the makers of I2, and the other compiler is a third-party product. Assume that each compiler uses the same number of instructions for a given program but that the instruction mix is as described in the table.

- (a) Using C1 on both I1 and I2, which machine is faster and by how much?
- (b) Using C2 on both I1 and I2, which machine is faster and by how much?
- (c) If you purchase I1, which compiler would you use?
- (d) If you purchase I2, which compiler would you use?

2. True/False questions. Circle only one of TRUE or FALSE.

- a. (TRUE, FALSE) The fastest computer will be the one with the highest clock rate.
- b. (TRUE, FALSE) Only by looking at the results of benchmarks for tasks similar to your workload can you get an accurate picture of likely performance.
- c. (TRUE, FALSE) The CPI term of the execution time equation can be affected by choice of programming languages.
- d. (TRUE, FALSE) The IC term of the execution time equation can be affected by number of pipeline stages.
- e. (TRUE, FALSE) A RISC ISA typically follows a register-memory architecture model because programs do not access memory very often.

- f. (TRUE, FALSE) Pipelining facilitates high clock frequency designs.
- g. (TRUE, FALSE) Since pipeline is transparent to user programs, software techniques are not helpful to achieve better pipelining performance.
- h. (TRUE, FALSE) Pipeline reduces CPI and execution time of individual instruction.

3. Suppose we enhance a machine to make all floating-point instructions run five times faster. If the execution time of some benchmark before the floating-point enhancement is 20 seconds, what will speedup be if half of the 20 seconds is spent executing floating-point instructions?

4.

Chip	Die size ( $mm^2$ )	Estimated defect rate(per $cm^2$ )	Manufacturing size (nm)	Transistors (millions)
IBM Power5	389	.30	130	276
Sun Niagara	380	.75	90	279
AMD Opteron	199	.75	90	233

Table above shows the relevant chip statistics that influence the cost of several current chips.

- a. What is the yield for the IBM Power5?
- b. Why does the IBM Power5 have a lower defect rate than the Niagara and Opteron?

5. One critical factor in powering a server farm is cooling. If heat is not removed from the computer efficiently, the fans will blow hot air back onto the computer, not cold air. We will look at how different design decisions affect the necessary cooling, and thus the price, of a system. (*Refer to the Figure 1.23 "Power consumption of several computer components" in page 57*)

- a. A cooling door for a rack costs \$4000 and dissipates 14 KW (into the room; additional cost is required to get it out of the room). How many servers with an Intel Pentium 4 processor, 1 GB 240-pin DRAM, and a single 7200 rpm hard drive can you cool with one cooling door?
- b. In a single rack, the MTTF of each processor is 4500 hours, of the hard drive is 9 million hours, and of the power supply is 30K hours. For a rack with 8 processors, what is the MTTF for the rack?

6. Your company's internal studies show that a single-core system is sufficient for the demand on your processing power. You are exploring, however, whether you could save power by using two cores.

- a. Assume your application is 80% parallelizable. By how much could you decrease the frequency and get

the same performance?

b. Assume that the voltage may be decreased linearly with the frequency. Using the equation in Section 1.5, how much dynamic power would the dual-core system require as compared to the single-core system?

7. The main reliability measure is MTTF, and design decisions affect their reliability.

a. We have a single processor with an FIT of 150. What is the MTTF for this system?

b. If it takes two days to get the system running again, what is the availability of the system?

8. Imagine that your company is trying to decide between a single-processor system and a dual-processor system. Figure 1.26 (*Refer to the Figure 1.26 "Performance of several processors on two benchmarks" in page 61*) gives the performance on two sets of benchmarks—a memory benchmark and a processor benchmark. You know that your application will spend 30% of its time on memory-centric computations, and 70% of its time on processor-centric computations.

a. Calculate the weighted performance of the benchmarks for the Pentium 4 and Athlon 64 X2 3800+.

b. How much speedup do you anticipate getting if you move from using a Pentium 4 to an Athlon 64 X2 3800+ on a memory-intensive application suite?

9. Your company has just bought a new dual Pentium processor, and you have been tasked with optimizing your software for this processor. You will run two applications on this dual Pentium, but the resource requirements are not equal. The first application needs 75% of the resources, and the other only 25% of the resources.

a. Given that 60% of the first application is parallelizable, how much speedup would you achieve with that application if run in isolation?

b. Given that 95% of the second application is parallelizable, how much speedup would this application observe if run in isolation?

c. Given that 60% of the first application is parallelizable, how much overall system speedup would you observe if you parallelized it, but not the second application?

e. How much overall system speedup would you achieve if you parallelized both applications?

10. We have a program of  $10^3$  instructions in the format of “lw, add, lw, add, ...” The add instruction depends (and only depends) on the lw instruction right before it. The lw instruction also depends (and only depends) on the add instruction right before it. If the program is executed on the pipelined datapath of Figure 1 below,

a. What would be the actual CPI?

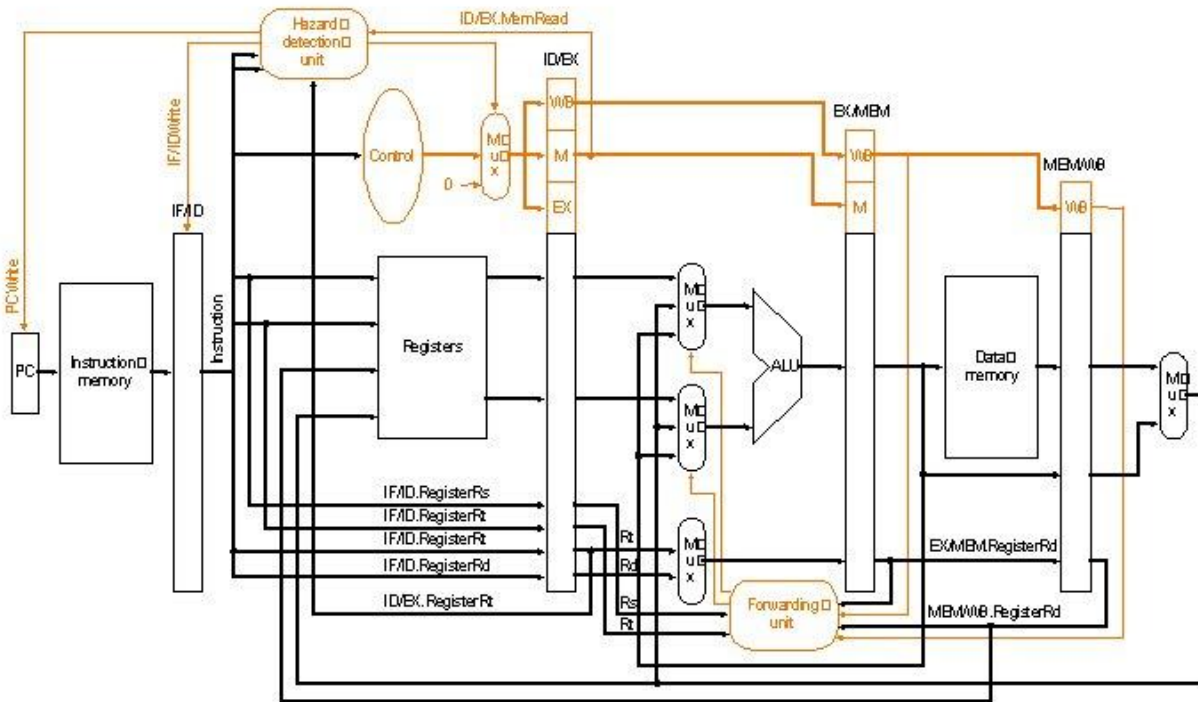
b. Without forwarding, what would be the actual CPI?

11. Consider executing the following code on the pipelined datapath of Figure 1 below.

```

add    $2,    $3,    $1
sub    $4,    $3,    $5
add    $5,    $3,    $7
add    $7,    $6,    $1
add    $8,    $2,    $6
  
```

At the end of the fifth cycle of execution, which registers are being read and which register will be written?



**Figure 1 Pipelined datapath**