

Date: _____

Quiz for Chapter 1 Computer Abstractions and Technology

Not all questions are of equal difficulty. Please review the entire quiz first and then budget your time carefully.

Name: _____

Course: _____

Solutions in Red

1. [15 points] Consider two different implementations, M1 and M2, of the same instruction set. There are three classes of instructions (A, B, and C) in the instruction set. M1 has a clock rate of 80 MHz and M2 has a clock rate of 100 MHz. The average number of cycles for each instruction class and their frequencies (for a typical program) are as follows:

Instruction Class	Machine M1 – Cycles/Instruction Class	Machine M2 – Cycles/Instruction Class	Frequency
A	1	2	60%
B	2	3	30%
C	4	4	10%

(a) Calculate the average CPI for each machine, M1, and M2.

For Machine M1:

$$\begin{aligned} \text{Cycles per Instruction} &= (60/100) * 1 + (30/100) * 2 + (10/100) * 4 \\ &= 1.6 \end{aligned}$$

For Machine M2:

$$\begin{aligned} \text{Cycles per Instruction} &= (60/100) * 2 + (30/100) * 3 + (10/100) * 4 \\ &= 2.5 \end{aligned}$$

(b) Calculate the average MIPS ratings for each machine, M1 and M2.

For Machine M1:

$$\begin{aligned} \text{Average MIPS rating} &= \text{Clock Rate} / (\text{CPI} * 10^6) \\ &= (80 * 10^6) / (1.6 * 10^6) \\ &= 50.0 \end{aligned}$$

For Machine M2:

$$\begin{aligned} \text{Average MIPS rating} &= \text{Clock Rate} / (\text{CPI} * 10^6) \\ &= (100 * 10^6) / (2.5 * 10^6) \\ &= 40.0 \end{aligned}$$

(c) Which machine has a smaller MIPS rating? Which individual instruction class CPI do you need to change, and by how much, to have this machine have the same or better performance as the machine with the higher MIPS rating (you can only change the CPI for one of the instruction classes on the slower machine)?

Machine M2 has a smaller MIPS rating. If we change the CPI of instruction class A for Machine M2 to 1, we can have a better MIPS rating than M1 as follows:

$$\begin{aligned}\text{Clocks per Instruction} &= (60/100)*1 + (30/100)*3 + (10/100)*4 \\ &= 1.9\end{aligned}$$

$$\begin{aligned}\text{Average MIPS rating} &= \text{Clock Rate}/(\text{CPI} * 10^6) \\ &= (100 * 10^6) / (1.9 * 10^6) \\ &= 52.6\end{aligned}$$

2. [10 points] (Amdahl's law question) Suppose you have a machine which executes a program consisting of 50% floating point multiply, 20% floating point divide, and the remaining 30% are from other instructions.

(a) Management wants the machine to run 4 times faster. You can make the divide run at most 3 times faster and the multiply run at most 8 times faster. Can you meet management's goal by making only one improvement, and which one?

Amdahl's Law states:

$$\text{Execution time after improvement} = \frac{(\text{Execution time affected by improvement})/(\text{Amount of Improvement}) + \text{Execution time unaffected}}$$

Assuming initially that the floating point multiply, floating point divide and the other instructions had the same CPI,

$$\text{Execution time after Improvement with Divide} = (20)/3 + (50 + 30) = 86.67$$

$$\text{Execution time after Improvement with Multiply} = (50)/8 + (20 + 30) = 66.67$$

The management's goal can be met by making the improvement with Multiply alone.

(b) Dogbert has now taken over the company removing all the previous managers. If you make both the multiply and divide improvements, what is the speed of the improved machine relative to the original machine?

If we make both the improvements,

$$\text{Execution time after Improvement} = (50)/8 + (20)/3 + (30) = 53.33$$

$$\text{The speedup relative to the original machine} = (100)/(53.33) = 1.88$$

3. [5 points] Suppose that we can improve the floating point instruction performance of machine by a factor of 15 (the same floating point instructions run 15 times faster on this new machine). What percent of the instructions must be floating point to achieve a *Speedup* of at least 4?

We will use Amdahl's Law again for this question.

Let x be percentage of floating point instructions. Since the speedup is 4, if the original program executed in 100 cycles, the new program runs in $100/4 = 25$ cycles.

$$(100)/4 = (x)/15 + (100 - x)$$

Solving for x , we get:

$$x = 80.36$$

The percent of floating point instructions need to be 80.36.

4. [6 points] Just like we defined MIPS rating, we can also define something called the MFLOPS rating which stands for Millions of Floating Point operations per Second. If Machine A has a higher MIPS rating than that of Machine B, then does Machine A necessarily have a higher MFLOPS rating in comparison to Machine B?

A higher MIPS rating for machine A compared to machine B need not imply a higher MFLOPS rating for that machine A. One reason for this can be the following: It is possible that the floating point instructions form a fairly low proportion of the all the instructions in a given program. So if the floating point operations of machine B are far more efficient than the floating point operations of machine A while the other (integer, memory etc) instructions are more efficient on B, then machine B gets a higher MFLOPS rating than A while A has a higher MIPS rating.

5. [6 points] Consider the SPEC benchmark. Name two factors that influence the resulting performance on any particular architecture.

Two factors that influence the resulting performance of a SPEC benchmark on any particular architecture are (a) The compiler flags used to compile the benchmark (b) The input data that is given to the benchmark while measuring performance

6. [5 points] How did the development of the transistor affect computers? What did the transistor replace?

The answer to the first part of this question is subjective: The development of transistor has had a tremendous impact on bring the computer to our homes in the form of PCs and more recently into our hands (Personal Digital Assistants etc). The ability to package and integrate transistors on a chip at a rate that has been increasing exponentially according to Moore's Law has resulted in tremendous performance gains for programs without having to change any of them.

The transistor replaced the vacuum tube in 1951.

7. [25 points] A two-part question:

(Part A)

Assume that a design team is considering enhancing a machine by adding MMX (multimedia extension instruction) hardware to a processor. When a computation is run in MMX mode on the MMX hardware, it is 10 times faster than the normal mode of execution. Call the percentage of time that could be spent using the MMX mode the *percentage of media enhancement*.

(a) What percentage of *media enhancement* is needed to achieve an overall speedup of 2?

We will use Amdahl's Law for this question.

Execution time with Media Enhancement =

$$\frac{\text{Execution time improved by Media enhancement}}{\text{Amount of Improvement}} + \text{Execution time unaffected}$$

Let x be the percent of media enhancement needed for achieving an overall speedup of 2. Then,

$$(100)/2 = (x)/10 + (100-x)$$

Solving for x, we have x = 55.55

(b) What percentage of the run-time is spent in MMX mode if a speedup of 2 is achieved? (Hint: You will need to calculate the new overall time.)

The new overall time is $100/2 = 50$. Now 55.55 of the original program used media enhancement. Let x be the percentage of the new run-time that is spent in MMX mode (for a speedup of 2).

$$x = (55.55 * 50)/100 = 27.78$$

(c) What percentage of the media enhancement is needed to achieve one-half the maximum speedup attainable from using the MMX mode?

The maximum speedup using MMX mode occurs when the whole program can run in media enhancement mode. The maximum speedup in this case is 10. One-half of this is 5. Plugging in 5 instead of 2 in (a):

$$(100)/5 = (x)/10 + (100-x)$$

Solving for x, we get $x = 103.7$

(Part B)

If processor A has a higher clock rate than processor B, and processor A also has a higher MIPS rating than processor B, explain whether processor A will always execute faster than processor B. Suppose that there are two implementations of the same instruction set architecture. Machine A has a clock cycle time of 20ns and an effective CPI of 1.5 for some program, and machine B has a clock cycle time of 15ns and an effective CPI of 1.0 for the same program. Which machine is faster for this program, and by how much?

The CPU Time is given by the equation:

$$\text{CPU Time} = \text{Instruction count} * \text{CPI} * \text{Clock cycle Time}$$

MIPS rating is defined by:

$$\text{MIPS} = (\text{Clock Rate}) / (\text{CPI} * 10^6)$$

For machines A and B:

$$(\text{CPUTime})_A = (\text{Instruction count})_A * (\text{CPI})_A * (\text{Clock cycle Time})_A$$

$$(\text{CPUTime})_B = (\text{Instruction count})_B * (\text{CPI})_B * (\text{Clock cycle Time})_B$$

$$(\text{MIPS})_A = (\text{Clock Rate})_A / ((\text{CPI})_A * 10^6)$$

$$(\text{MIPS})_B = (\text{Clock Rate})_B / ((\text{CPI})_B * 10^6)$$

If clock rate of A is higher than that of B, and MIPS rating of A is higher than that of B,

$$(\text{MIPS})_A > (\text{MIPS})_B \quad \text{and} \quad (\text{Clock Rate})_A > (\text{Clock Rate})_B$$

From the above equations it follows that:

$$(\text{Clock Rate})_A / (\text{Clock Rate})_B > (\text{CPI})_A / (\text{CPI})_B$$

$$(\text{Clock Cycle Time})_B / (\text{Clock Cycle Time})_A > (\text{CPI})_A / (\text{CPI})_B$$

From this it emerges that if the instruction counts are the same, processor A will always execute faster than processor B.

Assuming instruction counts are the same,

$$(\text{CPUTime})_A = (I) * 1.5 * 20\text{ns} = (I) * 30\text{ns}$$

$$(\text{CPUTime})_B = (I) * 1.0 * 15\text{ns} = (I) * 15\text{ns}$$

Machine B is faster by twice as much as Machine A.

8. [6 points] Suppose a program segment consists of a purely sequential part which takes 25 cycles to execute, and an iterated loop which takes 100 cycles per iteration. Assume the loop iterations are independent, and cannot be further parallelized. If the loop is to be executed 100 times, what is the maximum speedup possible using an infinite number of processors (compared to a single processor)?

The sequential part takes 25 cycles. Each iteration of the loop (which takes 100 cycles) can be executed independently and there are totally 100 iterations. Applying Amdahl's law,

Execution time after improvement =

$$\frac{(\text{Execution time affected by improvement})}{(\text{Amount of Improvement})} + \text{Execution time unaffected}$$

$$= (100 * 100) / 100 + 25$$

$$= 100 + 25$$

$$\text{Speedup} = 125/100 = 1.25$$

9. [5 points] Computer A has an overall CPI of 1.3 and can be run at a clock rate of 600MHz. Computer B has a CPI of 2.5 and can be run at a clock rate of 750 Mhz. We have a particular program we wish to run. When compiled for computer A, this program has exactly 100,000 instructions. How many instructions would the program need to have when compiled for Computer B, in order for the two computers to have exactly the same execution time for this program?

$$\begin{aligned} (\text{CPUTime})_A &= (\text{Instruction count})_A * (\text{CPI})_A * (\text{Clock cycle Time})_A \\ &= (100,000) * (1.3) / (600 * 10^6) \text{ ns} \end{aligned}$$

$$\begin{aligned} (\text{CPUTime})_B &= (\text{Instruction count})_B * (\text{CPI})_B * (\text{Clock cycle Time})_B \\ &= (I)_B * (2.5) / (750 * 10^6) \text{ ns} \end{aligned}$$

$$\text{Since } (\text{CPUTime})_A = (\text{CPUTime})_B,$$

we have to solve for $(I)_B$ and get 65000

10. [5 points] Imagine that you are able to perform benchmarking “races” to compare two computers you are thinking about buying. Come up with a list of 5 benchmark programs or usage scenarios you would use to create your own personalized benchmark suite. For each program you select, justify it. For the benchmark suite as a whole, discuss a method for calculated a weighted average of the different program run-times.

The answer to this question is subjective. One possible list of benchmarks could include the following:

Benchmark Program	Reason to choose this benchmark	Program Runtime (s)	Weight
Browser with multimedia plug ins	Opening a lot of media content in a lot of tabs at the same can be used to test the memory performance and parallelism potential	300	5
Financial Application (Stock Value Predictor)	Can be used to measure floating point and vector instruction set performance	200	4
Games that involve a lot of AI	Can be used to test branch prediction performance	100	3
Word processing software	Most common use case. Can be used to test general integer operations and rendering	200	5
Desktop Search Software	Can be used to test I/O performance since a search software builds up a huge index by constant read of file system	300	3

The weighted average of the runtimes is $= (300*5 + 200*4 + 100*3 + 200*5 + 300*3)/20 = 225$

11. [8 points] The design team for a simple, single-issue processor is choosing between a pipelined or non-pipelined implementation. Here are some design parameters for the two possibilities:

Parameter	Pipelined Version	Non-Pipelined Version
Clock Rate	500MHz	350 MHz
CPI for ALU instructions	1	1
CPI for Control	2	1

Name: _____

instructions		
CPI for Memory instructions	2.7	1

(a) For a program with 20% ALU instructions, 10% control instructions and 75% memory instructions, which design will be faster? Give a quantitative CPI average for each case.

$$\text{Average CPI for Pipelined Version} = (0.2 \cdot 1 + 0.1 \cdot 2 + 0.7 \cdot 2.7) = 2.29$$

$$\text{Average CPI for Non-Pipelined Version} = (0.2 \cdot 1 + 0.1 \cdot 1 + 0.7 \cdot 1) = 1.0$$

$$\text{CPU execution time for Pipelined version} = 2.26 / (500 \text{ Mhz}) = 4.5 \text{ ns}$$

$$\text{CPU execution time for Non-Pipelined version} = 1.0 / (350 \text{ Mhz}) = 2.8 \text{ ns}$$

The non-pipelined version is faster.

(b) For a program with 80% ALU instructions, 10% control instructions and 10% memory instructions, which design will be faster? Give a quantitative CPI average for each case.

$$\text{Average CPI for Pipelined Version} = (0.8 \cdot 1 + 0.1 \cdot 2 + 0.1 \cdot 2.7) = 1.27$$

$$\text{Average CPI for Non-Pipelined Version} = (0.8 \cdot 1 + 0.1 \cdot 1 + 0.1 \cdot 1) = 1.0$$

$$\text{CPU execution time for Pipelined version} = 1.27 / (500 \text{ Mhz}) = 2.54 \text{ ns}$$

$$\text{CPU execution time for Non-Pipelined version} = 1.0 / (350 \text{ Mhz}) = 2.8 \text{ ns}$$

The pipelined version is faster.

12. [5 points] A designer wants to improve the overall performance of a given machine with respect to a target benchmark suite and is considering an enhancement X that applies to 50% of the original dynamically-executed instructions, and speeds each of them up by a factor of 3. The designer's manager has some concerns about the complexity and the cost-effectiveness of X and suggests that the designer should consider an alternative enhancement Y . Enhancement Y , if applied only to some (as yet unknown) fraction of the original dynamically-executed instructions, would make them only 75% faster. Determine what percentage of all dynamically-executed instructions should be optimized using enhancement Y in order to achieve the same overall speedup as obtained using enhancement X .

We will use Amdahl's Law for this problem.

Execution time after improvement =

$$(\text{Execution time affected by improvement}) / (\text{Amount of Improvement}) + \text{Execution time unaffected}$$

$$\text{Execution Time using } X = (50) / 3 + (100 - 50) = 66.67$$

$$\text{The speedup is given by} = (100) / 66.67 = 1.5$$

Let the percentage of dynamically executed instructions to which Y is to be applied be x .

$$\text{Execution Time using } Y = (x) / 1.75 + (100 - x)$$

$$\text{SpeedUp} = (100) / (\text{Execution Time using } Y) = 1.5$$

Solving for x , we get $x = 77.78$