

HY425: Computer Systems Architecture

Homework Problem Set 1

Assignment: October 08, 2012

Due: **October 20, 2012 – 23:59:59**

Instructions: Solve all problems in a .pdf file and send them via e-mail to Vassilis Papaefstathiou (papaef@csd.uoc.gr). Use the subject: **HY425 - Homework 1**

Problem 1 (30 points)

For this problem you will need a spreadsheet. The following table illustrates the performance of several processors using two benchmarks. The first benchmark stresses the memory system (how fast can the memory system respond to loads and stores), whereas the second benchmark, Dhrystone, stresses the execution units of the processor.

Chip	# Cores	Frequency (MHz)	Memory Performance	Dhrystone Performance
Athlon 64 X2 4800+	2	2400	3423	20718
Pentium EE 840	2	2200	3228	18893
Pentium D 820	2	3000	3000	15220
Athlon 64 X2 3800+	2	3200	2941	17129
Pentium 4	1	2800	2731	7621
Athlon 64 3000+	1	1800	2953	7628
Pentium 4 570	1	2800	3501	11210
Processor X	1	3000	6000	6050

- Create a table similar to the one above, except express the results as normalized to the slowest processor (in terms of frequency) for each benchmark. Notice that performance is reported as a rate (a-la SPEC).
- Calculate the arithmetic mean of the performance of each processor. Use both the original performance and your normalized performance calculated in part (i). Which processor achieves the highest score in each case?
- Calculate the geometric mean of the normalized performance of the dual-core processors and the geometric mean of the normalized performance of the single-core processors for the memory benchmark.
- Which processor would you choose if you intended to run an application which spends 80% of its time on memory-bound computations and 20% on compute-bound computations?

Problem 2 (20 points)

Your company has just bought a new system with a dual-core processor, and you have been tasked with optimizing your software for this system. You will run two applications on this dual-core machine, but the resource requirements are not equal.

The first application needs 80% of the resources, and the other only 20% of the resources.

- i. Given that 50% of the first application is parallelizable, how much speedup would you achieve with that application if run in isolation?
- ii. Given that 80% of the second application is parallelizable, how much speedup would this application observe if run in isolation?
- iii. Given that 50% of the first application is parallelizable, how much overall system speedup would you observe if you parallelized it, but not the second application?
- iv. How much overall system speedup would you achieve if you parallelized both applications, given the information in parts (i) and (ii)?

Problem 3 (30 points)

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 (a) two or (b) four or (c) eight cores.

- i. Assume your application is 70% parallelizable. By how much could you decrease the frequency and get the same performance in each case?
- ii. Assume that the voltage may be decreased linearly with the frequency. How much dynamic power would each system require as compared to the single-core system? For the capacitive load C of the multi-core systems assume that: $C(N) = C(1) \times \sqrt{N}$.
- iii. Now assume that the voltage may not decrease below 30% of the original voltage. This voltage is referred to as the "voltage floor," and any voltage lower than that will lose the state. What percent of parallelization gives you a voltage at the voltage floor for each case?
- iv. How much dynamic power would each system require as compared to the single-core system when taking into account the voltage floor? For the capacitive load C of the multi-core systems assume that: $C(N) = C(1) \times \sqrt{N}$.

Problem 4 (20 points)

A hypothetical supercomputer has as many as 65536 processors. The processors are organized in 65536 node cards, where each node has one off-chip DRAM memory module and a hard disk. The system administrator of this supercomputer claims that, on average, he loses one processor per two days due to processor failures, a node per five days, due to memory failures, and a disk per thirty days due to disk failure. Assume that the three types of failures are not related.

- i. What is the MTTF of each processor, memory module and disk on the supercomputer?
- ii. What is the MTTF of the system as a whole?
- iii. If each failure requires replacement and reboot of a node on the supercomputer, and this process takes 1 hour, what is the availability of the supercomputer?