**HY425**

**Homework Problem Set 1**

**Assignment: 29/2/2008**

**Due Date: 14/3/2008**

Instructions: Solve all problems in a **.pdf** file and send them via e-mail to Stamatis Kavadias (kavadias@ics.forth.gr), with a copy to the instructor (dsn@ics.forth.gr). Use the following subject in your e-mail: **HY425: Homework 1 Submission**. Please the aforementioned subject only, so that your homework is read and graded.

## Problem 1 (20 points)

The AMD Athlon 64 processor provides a hardware mechanism and a software interface to scale voltage and frequency, in order to reduce dynamic power consumption. The following table lists the voltage/frequency pairs available on the processor. Note that voltage and frequency are set together whenever the software desires to use scaling.

| Processor "gear" | Frequency (MHz) | Voltage |
|---|---|---|
| 0 | 2000 | 1.5 |
| 1 | 1800 | 1.4 |
| 2 | 1600 | 1.3 |
| 3 | 1400 | 1.2 |
| 4 | 1200 | 1.1 |
| 5 | 1000 | 1.0 |
| 6 | 800 | 0.9 |

a. Assume the processor consumes at most 90W of power at gear 0 (maximum frequency), out of which 20W is static (leakage) power and 70W is dynamic power. Give a table listing the expected power consumption of the processor at each gear.
b. Assume the processor executes all instructions in 1 cycle (CPI=1) and the performance of the processor with a nominal workload at gear 0 is R. Give a table showing the performance of the processor relative to R, at the different gears.
c. Assume the user wishes the processor to consume no more than half its maximum power (45W). What gear should the processor use for a workload with CPI=1 and no stalls?
d. A more realistic implementation of the processor executes all instructions in 1 cycle, except from loads and stores to off-processor memory, which require 5 bus cycles each. Assume that the bus runs at the minimum processor frequency (800 MHz). You are running a workload where 10% of the instructions are loads and stores that miss in the on-chip caches and served by off-processor memory. If R is the performance of the processor at gear 0, give a table of the performance of the processor at the different gears.

e. If the user can tolerate up to 20% performance loss from the processor, what gear should the processor use to minimize power consumption while meeting the user's constraint?

**Problem 2 (20 points)**

A program executes 10 million instructions. A user is trying to run this program on various multi-core processors.

a. What is the maximum number of instructions that the program can execute sequentially (i.e. without distributing between cores) to achieve 95% of the maximum speedup (aka 95% parallel efficiency), on a processor with 8 cores, a processor with 64 cores, and a processor with 512 cores?
b. Parallelizing a program involves at a minimum, the overhead of creating parallel tasks and assigning them to different cores for execution. Assume that in the given program, all sequential instructions are used for creating and assigning parallel tasks to the cores of the processor. Assume that the program creates as many tasks as the number of cores. What is the maximum number of cycles that task creation should take in order to sustain 95% of the maximum speedup on a processor with 8 cores, a processor with 64 cores and a processor with 512 cores?
c. Plot a function of the minimum task creation time for sustaining 95% parallel efficiency on an N-core processor, versus the number of cores. You may use a spreadsheet for this answer

## Problem 3 (20 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 | #of cores | Clock frequency (MHz) | Memory performance | Dhrystone performance |
|---|---|---|---|---|
| Athlon 64 X2 4800+ | 2 | 2,400 | 3,423 | 20.718 |
| Pentium EE 840 | 2 | 2,200 | 3,228 | 18,893 |
| Pentium D 820 | 2 | 3,000 | 3,000 | 15,220 |
| Athlon 64 X2 3800+ | 2 | 3,200 | 2,941 | 17,129 |
| Pentium 4 | 1 | 2,800 | 2,731 | 7,621 |
| Athlon 64 3000+ | 1 | 1,800 | 2,953 | 7,628 |
| Pentium 4 570 | 1 | 2,800 | 3,501 | 11,210 |
| Processor X | 1 | 3,000 | 7,000 | 5,000 |

a. Create a table similar to the one above, except express the results as normalized to the fastest processor for each benchmark. Notice that performance is reported as a rate (ala SPEC).
b. 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.
c. Calculate the geometric standard deviation of the dual-core processor performance on the memory benchmark. What does this deviation suggest about how much the choice of the particular processor affects performance?

## Problem 4 (20 points)

We will be using the table from the previous problem. Imagine that your company is trying to decide between a single-core system and a dual-core system. You know that your application spends 30% of its time in memory-intensive computations and 70% of its time on processor-bound computations.

a. Calculate the weighted performance of the benchmarks on the Pentium 4 and the Athlon 64 X2 3800+
b. How much speedup do you expect by moving from a Pentium 4 to an Athlon 64 X2 3800+ on a memory-intensive application suite (that is, an application suite where execution time is dominated by memory stalls)?

c. You are asked to choose between processor X and the Athlon 64 X2 4800+, in workloads with various degrees of memory/computation intensity. Plot a graph of the performance of the processors as a function of the fraction of computation which is CPU-intensive. How would you chose a processor using this graph?

**Problem 5 (20 points)**

A hypothetical supercomputer has as many as 131,072 processors (the IBM BlueGene/L has that many processors indeed, but the exercise uses a hypotherical machine). The processors are organized in 131,072 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 looses one processor per two days due to processor failures, a node per week, due to memory failures, and a disk per month due to disk failure. Assume that the three types of failures are not related.

a. What is the MTTF of each processor on the supercomputer?
b. What is the MTTF of each memory module on the supercomputer?
c. What is the MTTF of the disk?
d. What is the MTTF of the system as a whole?
e. The supercomputer's processor costs 150 Euros. The supercomputer's RAM module costs 300 Euros. The disk on the node costs 250 Euros. The administrator has an annual budget of 30,000 Euros for machine repairs. Is this budget sufficient?
f. If each of the aforementioned failures requires replacement and reboot of a node on the supercomputer, and this process takes 30 minutes, what is the availability of the supercomputer?
g. You have the option to change the scale of the supercomputer, so that its availability is consistently over 99%. How would you change the configuration of the machine?