

ENGN 2910A Homework 01 (75 points) – Due Date: Sept 24 2015 (in office)

Professor: Sherief Reda

School of Engineering, Brown University

---

1. [20 points] A multiprocessor machine has  $P$  processors. On this machine we map a computation in which  $N$  iterate values must be computed and then exchanged between the processors. Values are broadcast on a bus after each iteration. Each iteration proceeds in two phases. In the first phase, each processor computes a subset of the  $N$  iterates. Each processor is assigned the computation of  $N/P$  iterates. In the second communication phase, each processor, in a one-by-one fashion, broadcasts its result to all other processors. Every processor waits for the end of the entire communication phase before starting a new computation phase. Let  $T_c$  be the time to compute one iterate and let  $T_b$  be the time to broadcast a result on the bus. Let  $R = T_c/T_b$  denote the computation-to-communication ratio. Note that when  $P = 1$ , no communication is needed. What is the optimal number of processor? Express your result in terms of  $N$  and  $R$ .

2. [10 points] Consider two possible improvements for a base machine: one improving floating point performance and one improving memory performance. Three programs are simulated: one with no floating-point operations (Program 1), one dominated by floating-point operations (Program 2), and one with balance between memory accesses and floating-point operations (Program 3). The execution time of each program on the three machines is given in the following table.

Machines	Program 1	Program 2	Program 3
Base machine	1 s	10 ms	10 s
Base + FP units	1 s	2 ms	6 s
Base + cache	0.7 s	9 ms	5 s

- Calculate the average execution time for each machine. Which machine is better?
- Compared to the base machine, calculate (i) arithmetic means of speedup, (ii) geometric means of speed, and (iii) harmonic means of speedup. Which machine is better? Please explain any discrepancy to the conclusion compared to question (a).

3. [10 points] In machine M1, which is clocked at 100 MHz, it was observed that 20% of the computation time of the integer benchmarks is spent in a subroutine  $multiply(A, B, C)$ , which multiplies integer  $A$  and  $B$  and returns the result in  $C$ . Furthermore, each invocation of  $multiply$  takes 800 cycles to execute. To speed up the program it is proposed to introduce a new machine, M2, with a special native instruction  $MULT$ , where  $MULT$  takes 40 cycles. Besides the multiplies all other instructions take the same number of cycles. Because of the added complexity, however, the clock rate of M2 is only 80 MHz. How much faster or slower is M2 compared to M1?

4. [10 points] Assume that we are considering enhancing a machine by adding vector hardware to it. When a computation is run in vector mode on the vector hardware, it is 10 times faster than the normal mode of execution. We call the percentage of time that could be spent using vector mode the *percentage of vectorization*.

- a. Draw a graph that plots the speedup as percentage of the computation performed in vector mode. Label the y-axis "net speedup" and label the x-axis "percent vectorization".
- b. What percentage of vectorization is needed to achieve a speedup of 2?
- c. What percentage of the computation runtime is spent in vector mode if a speedup of 2 is achieved?
- d. What percentage of vectorization is needed to achieve one-half the maximum speedup attainable from using vector mode?

5. [25 points] Assume a base-line processor that is single-cycle in which all instructions are executed in a single cycle with frequency  $f$  and  $V$ . Note that without changing a processor's design, you can always run it at lower or higher frequencies while scaling the voltage linearly in proportion to the frequency.<sup>1</sup> For example, if  $f$  is reduced by 10%, then you do not need to operate at the original voltage, you can reduce it by 10%. Please answer the following questions.

- a. Please report the ratio power, delay, energy, and energy-delay product for the following machines with respect to the base-line machine. Consider only the dynamic power.
  - a five-stage pipeline design clocked at  $5f$ .
  - a five-stage pipeline design clocked at  $f$ .
  - a multiprocessor design with 5 single-cycle processors clock at  $f$ .
  - a multiprocessor design with 5 five-stage pipeline processors clock at  $5f$ .
- b. Repeat part (a) for static (leakage) power only.

---

<sup>1</sup>This statement is partially correct as there is a typically a *voltage floor*, and any voltage lower than that will lead to a loss of state.