

ENGN 2910A Homework 07 (50 points) – Due Date: Nov 19th 2015

Professor: Sherief Reda

School of Engineering, Brown University

1. [25 points] Consider the execution of a loop in a single-cycle VLIW processor. Assume that any combination of instruction types can be executed in the same cycle. Note that such assumption only removes resource constraints, but data and control dependencies must be still handled correctly. Assume that register reads occur before register writes in the same cycle.

```
loop: lw    $1, 40($6)
      add   $5, $5, $1
      sw    $1, 20($5)
      addi  $6, $6, 4
      addi  $5, $5, -4
      beq   $5, $0, loop
```

a. [20 points] Unroll this loop once and schedule it for a 2-issue static VLIW processor. Assume that the loop always executes an even number of iterations. Hint: for correct scheduling, you need to identify all potential data and control hazards. Use register renaming whenever possible to eliminate false dependencies.

b. [5 points] What is the speed-up of using your scheduled code instead of the original code? Assume that the loop will be executed a very large number of times.

2. (25 points)

a. Consider a simple 5-stage pipeline that is single-threaded. The pipeline treats every cache miss as a hazard and freezes the pipeline. While executing a benchmark assume that a L1 cache miss occurs every 100 cycles, and that each L1 cache miss takes 10 cycles to satisfy if the block is found in L2 or 50 cycles if L2 misses as well. A L2 cache miss occurs after 200 cycles of computation. Assume that the CPI in the absence of the cache misses is 1. What is the actual CPI, taking into account the cache miss latencies?

b. Consider the same example as in (a.), but assume that hardware is block multi-threaded with support for two HW threads (similar to slide 9 in lecture). Assume that switching overhead is zero, and that there are two threads with identical cache miss behavior as in part (a.). What is the CPI of each of the two programs on the two-way multi-threaded machine? Did the CPI improve? If yes, explain how. If not, explain why one should bother with the two-way multi-threaded machine.

c. Consider the case in (a), but the switching overhead is five cycles. Again compute the CPI of each thread, and explain why it increases, decreases, or stays the same.

d. Consider the case for which the L2 miss latency jumps from 50 to 500 cycles and the switching overhead jumps from 5 to 50 cycles. Compare the CPI in this machine.