COE 308 – Fall 2005
Computer Architecture

Assignment 5: Pipelined Processor Design

Due: Monday, December 19, 2005

1. **(3 pts)** A computer architect needs to design the pipeline of a new microprocessor. He has an example workload program core with 10^6 instructions. Each instruction takes 100 ps to finish.

   a) How long does it take to execute this program core on a non-pipelined processor?
   b) The current state of the art microprocessor has about 20 pipeline stages. Assume it is perfectly pipelined. How much speedup will it achieve compared to the non-pipelined processor?
   c) Real pipelining isn't perfect, since implementing pipelining introduces some overhead per pipeline stage. Will this overhead affect instruction latency, instruction throughput, or both?

2. **(3 pts)** Identify all of the RAW data dependencies in the following code. Which dependencies are data hazards that will be resolved by forwarding? Which dependencies are data hazards that will cause a stall? Using a multiple-clock-cycle graphical representation, show the forwarding paths and stalled cycles if any.

```
add $3, $4, $2
sub $5, $3, $1
lw  $6, 200($3)
add $7, $3, $6
```

3. **(4 pts)** We have a program of 10^6 instructions in the format of "**`lw,add,lw,add,…`**". The **`add`** instruction depends only on the **`lw`** instruction right before it. The **`lw`** instruction also depends only on the **`add`** instruction right before it. If this program is executed on the 5-stage MIPS pipeline:

   a) Without forwarding, what would be the actual CPI?

   b) With forwarding, what would be the actual CPI?

4. **(6 pts)** We have a program core consisting of five conditional branches. The program core will be executed millions of times. Below are the outcomes of each branch for one execution of the program core (T for taken and N for not taken).

   Branch 1: T-T-T
   Branch 2: N-N-N-N
   Branch 3: T-N-T-N-T-N
   Branch 4: T-T-T-N-T
   Branch 5: T-T-N-T-T-N-T

   Assume that the behavior of each branch remains the same for each program core execution. For dynamic branch prediction schemes, assume that each branch has its own prediction buffer and each buffer is initialized to the same state before each execution.

List the predictions and the accuracies for each of the following branch prediction schemes:

a) Always taken
b) Always not taken
c) 1-bit predictor, initialized to predict taken
d) 2-bit predictor, initialized to weakly predict taken

5. **(4 pts)** Moving the branch execution and comparison from the MEM stage to the ID stage can reduce the delay of branches from 3 cycles to 1 cycle. However, it introduces an opportunity for data hazards that cannot be resolved by forwarding. Give a set of code sequences that demonstrate the data hazard cases that must be detected by a branch instruction. Using a multiple-clock-cycle graphical representation, show the forwarding paths and stalled cycles.