

Problem Exercise 1

Due date: December 11, 2012 at 1640 (4:40pm)

This is the first of two problem exercises. It is not obligatory, but it does carry points which may be used to offset a weaker score on the examination, as described in the course syllabus. This exercise is worth in total 50 points. The due date is strict and late submissions will not be accepted, except under extenuating circumstances.

In all cases, answers must be explained. In particular, the formulas which were used must be given. Numerical answers without a clear explanation of the derivation will not receive credit.

Problems similar to these may appear on examinations. It is therefore advisable to know how to solve such problems.

1 Problems

1. A single-cycle implementation of the MIPS architecture has the following parameters:

Component	Delay
Instruction memory read	400 ps.
Data memory read or write	300 ps.
Register read	150 ps.
Register write	175 ps.
ALU operation	200 ps.
Adder (not ALU) operation	125 ps.
All other operations	0 ps.

Using Figure 4.24 of the textbook as the implementation model, compute the critical-path times for each of the following instructions. In all cases, show the names of the components which are included in the path, as well as their delay values.

If Figure 4.24 does not provide all necessary control and communication lines, and multiplexors, sketch briefly that which must be added. It is not necessary to draw a new diagram; just explain those additional items which are necessary to implement the instruction.

- (a) The `addi` instruction.
- (b) The `lw` instruction.
- (c) The `sw` instruction.
- (d) The `beq` instruction.

5DV118, Problem Exercise 1, page 2

- (e) The `j` instruction.
- (f) The `jr` instruction.
- (g) The `jal` instruction.

2. An implementation of the MIPS architecture following Figure 4.33 of the textbook has the following parameters:

Component	Delay
Instruction fetch (IF)	200 ps.
Instruction decode and register file read (ID)	300 ps.
Execution or address calculation (EX)	400 ps.
Data memory access (MEM)	250 ps.
Write back (WB)	150 ps.

- (a) Assuming a non-pipelined implementation, and assuming all delays not explicitly identified above to be negligible, compute the latency for each of the instructions identified in Problem 1 above.
- (b) Repeat part (a), this time assuming a pipelined implementation.
- (c) For a pipelined implementation, if one stage of the pipeline could be split into two stages, each with a latency one-half of that of the original stage, identify the stage, if any, which should be split in order to obtain the least instruction latency.
- (d) Suppose that the instruction mix of a program is as follows: `add`: 30%; `ori`: 10%; `beq`: 25%; `lw`: 20%; `sw`: 15%. For the pipelined implementation, compute the percentage of total cycles which utilize data memory. Assume that there are no stalls or hazards.

5DV118, Problem Exercise 1, page 3

3. Given are the following two instruction sequences.

(i):	<code>lw \$t1, 0(\$t2)</code>	(ii):	<code>sw \$t1, 0(\$t2)</code>
	<code>addi \$t2, \$t1, 50</code>		<code>addi \$t2, \$t7, 50</code>
	<code>add \$t3, \$t1, \$t8</code>		<code>addi \$t1, \$t2, \$zero</code>
	<code>addi \$t1, \$t1, 50</code>		<code>lw \$t1, 8(\$t2)</code>
	<code>sw \$t4, 0(\$t1)</code>		<code>add \$t2, \$t2, \$t6</code>

For each of these instruction sequences, do the following:

- Identify all *read after write* (RAW) data dependencies, regardless of the distance between the instructions.
- For the five-stage pipeline model of the textbook, indicate which of these dependencies results in a hazard in the case that there is no forwarding, and show how `nop` operations may be inserted to avoid these hazards. Find one solution for each program sequence which resolves all hazards, not one solution for each hazard. The solution must not contain any unnecessary `nops`.
- Repeat part (b) in the case that forwarding is employed.

4. A program contains two distinct branch instructions. For the first instruction, the pattern `NYNNN` repeats endlessly, where `Y` indicates that the branch was taken and `N` that it was not. For the second program, the pattern `YNYNN` repeats endlessly. For each of these two patterns, answer the following questions.

- Compute the success rates for the always-taken and never-taken branch predictors.
- Suppose that a one-bit predictor is employed, with the initial state `0` (previous branch not taken). Compute the accuracy of the prediction as the number of passes through the branch pattern becomes large.
- Now suppose that a two-bit predictor is employed, with the initial state `00` (previous two branches not taken). Compute the accuracy of the prediction as the number of passes through the branch pattern becomes large.

2 Submission Rules

Solutions may be developed and submitted by groups of up to three individuals.

A printed copy of the solution must be placed in the appropriate course mailbox on the fourth floor of MIT-huset. The user-id of each group member for the submission must be indicated clearly on a cover page of the printed submission.