# CS/ECE 752 Spring 2024: HOMEWORK # 2 (Due by **11:00 AM** on **Friday, Feb 23**, via upload of PDF to Canvas)

Contact Rajesh Srivatsav (rsuresh6@wisc.edu)

# 1. Problem 1 (6 points)

Assume that in a particular program, 20% of the dynamic instructions are branches, 20% are load instructions, and the remaining 60% of the instructions have a CPI of 1.3. You are considering three microarchitecture designs for handling branches as below:

Design A does not have a branch predictor and has to expend 2 cycles for every branch instruction. The CPI of load instructions is 2. Determine the CPI of design A.

Design B has a branch predictor with an accuracy of 90%. A correctly predicted branch requires only 1 cycle, but a mispredicted branch requires 4 cycles to execute. For this design the CPI of a load degrades to 2.5. Determine the CPI of design B.

Design C has a branch predictor with an accuracy of 70%, and a branch execution time of 1 cycle for correct prediction and 3 cycles for misprediction. The load CPI is 2. Determine the CPI of design C.

Rank the three designs in order from best to worst.

#### 2. Problem 2 (14 points, 4+5+5)

Answer the following questions based on the code sequence below.

LD	F2, 0(R1)	
LD	F4, 0(R2)	
MULD	F2, F0, F6	# F6 is dest reg.
ADDD	F6, F4, F6	# F6 is dest reg.
LD	F4, 100(R3)	
ADDD	F6, F4, F2	# F2 is dest reg.
SD	F2, 20000(R1)	
ADD	R1, #8, R1	# R1 is dest reg.
ADD	R2, #8, R2	# R2 is dest reg.
BLT	R1, R4, loop	
	LD LD ADDD LD ADDD SD ADD ADD BLT	LD F2, 0(R1) LD F4, 0(R2) MULD F2, F0, F6 ADDD F6, F4, F6 LD F4, 100(R3) ADDD F6, F4, F2 SD F2, 20000(R1) ADD R1, #8, R1 ADD R2, #8, R2 BLT R1, R4, loop

The latencies of the different operations are: Integer ADD, SUB = 1 cycle, Memory LD = 7 cycles, Memory SD = 3 cycles, Branches = 3 cycles, ADDD = 3 cycles, and MULTD = 6 cycles.

- (a) What would be the baseline performance (in cycles, per loop iteration) of the code sequence if no new instruction execution could be initiated until the previous instruction execution had completed? Ignore front-end fetch and decode. Assume for now that execution does not stall for lack of the next instruction, but only one instruction/cycle can be issued. Assume the branch is taken.
- (b) Consider a multiple-issue design. Suppose you have two execution pipelines, each capable of beginning execution of one instruction per cycle, and enough fetch/decode bandwidth in the front end so that it will not stall your execution. Assume results can be immediately forwarded from one execution unit to another, or to itself. Further assume that the only reason an execution pipeline would stall is to observe a true data dependence. Now how many cycles does the loop require?
- (c) Continuing the above code, unroll the loop once, merge the two iterations, and schedule the instructions to eliminate as many bubbles as possible. Now how many cycles does the loop require?

### 3. Problem 3 (12 points) Introduction to ChampSim

ChampSim Github Repository - https://github.com/ChampSim/ChampSim

# Step 1: Cloning ChampSim

Running **git clone https://github.com/ChampSim/ChampSim** should create a directory named ChampSim within your current working directory, with all the files and directories. Then you should work from the newly created directory, while following the steps below.

# **Step 2: Installing Dependencies**

Go through the README of the ChampSim Github repository. It primarily involves installation of vcpkg, which is used by ChampSim to manage dependencies.

**NOTE**: If you are using <u>CSL machines (highly recommended)</u>, then the necessary dependencies may not be installed by default and you will have to install them. Installation **does not require** root privileges.

# Step 3: Compile ChampSim

Specified in the README. This involves first setting up the configuration (specified via config file in JSON), followed by running the make command. Note that there exist defaults for all options specified via config file, but feel free to look at parameters that are easily configured via the json config file.

For this assignment, the most important parts are:

- downloading and building ChampSim,
- Setting up a simple configuration.
- how to run ChampSim
- how to parse the ChampSim output and understand the statistics, and
- using/modifying the default configuration file (**champsim.json**).

# Step 4: Running Benchmarks on ChampSim

In this part, we will run a few applications from the SPEC 2017 benchmark suite on ChampSim using *champsim.json* config file.

- Change directories to your ChampSim directory.
- Running benchmark traces on ChampSim is specified in the README, mentioned again below. (run command from ChampSim directory)
  - bin/champsim --warmup\_instructions 2000000
    -simulation\_instructions 500000000
    - ~/path/to/traces/483.xalancbmk-127B.champsimtrace.xz
- All SPEC 2017 traces are present in /u/s/o/sohi/public/html/cs752/Spring2024/traces

You will perform a simple characterization of one of

- (a) 483.xalancbmk-127B.champsimtrace.xz [if your last name begins with A-I]
- (b) 459.GemsFDTD-1169B.champsimtrace.xz [if your last name begins with J-Q]
- (c) 434.zeusmp-10B.champsimtrace.xz [if your last name begins with R-Z]

**NOTE: champsim.json** has lots of parameters relating to the processor core and the memory hierarchy which can be varied. Feel free to explore it.

Submit the following.

Report the IPC, L1-Data Cache and L2 Cache Load miss rates for:

Cache hierarchy for parts (a):

- L1 instruction cache: 16KB 4-way set associative
- L1 data cache: 32KB 4-way set associative
- L2 cache: 512 KB 8-way set associative
- (a) Simulate using the default OOO CPU core with the above cache sizes. (those cache sizes not mentioned here can be left as default size, for example LLC)

Cache hierarchy for parts (b):

- L1 instruction cache: 32KB 4-way set associative
- L1 data cache: 64KB 4-way set associative
- L2 cache: 1MB 8-way set associative

- (b) Simulate using the default OOO CPU core with the above cache sizes. (those cache sizes not mentioned here can be left as default size, for example LLC)
- (c) How much time (roughly) did the simulations take for (a) and (b)?

For all of the above tests, warmup for 1M instructions, and simulate for next 50M instructions. These options are specified in the command line used to run ChampSim on a particular trace.

#### **Optional resources :**

1) Paper - https://arxiv.org/pdf/2210.14324.pdf