Exam in DAT 105 (DIT 051 / DIT 052) Computer Architecture

Time: 26-01-07 at 8:30 – 12:30 Johanneberg

Person in charge of the exam: Angelos Arelakis, Phone: 0763103557

Supporting material/tools: Chalmers approved calculator

Exam Review: More information on this will be available via Canvas

Grading intervals:

- **Fail** : Result < 24
- **Grade 3**: 24 <= Result < 36
- **Grade 4:** 36 <= Result < 48
- **Grade 5:** 48 <= Result

NOTE 1: Only a Chalmers approved calculator is allowed

NOTE 2: Bonus points from Real-stuff studies and Quizzes will be added to the exam results for approved exams used solely for higher grades and will be kept for one year.

NOTE 3: Answers must be given in English

GOOD LUCK!
Angelos Arelakis

[General disclaimer: If you feel that sufficient facts are not provided to solve a problem, either 1) ask the teacher when he visits the exam, or 2) make your own additional assumptions. Additional assumptions will be accepted if they are reasonable and required to solve the problem. Always make sure to motivate your answers.]

ASSIGNMENT 1

You are supposed to characterize the performance of three computers, A, B, and C with data regarding three programs P1, P2 and P3.

The following is known for the two computers and the three programs:

- A and B have **different** instruction-set architectures (ISAs) but A and C have the **same**
- The number of executed instructions on A for P1, P2 and P3 are 1 Billion, 1.5 Billion and 2 Billions, respectively. For B, the numbers are half of those of A. The execution time on A is 1 second, 1.5 second and 2 seconds for P1, P2 and P3, respectively.
- A, B and C are run at a clock frequency of 1 GHz, 2 GHz, 1 GHz, respectively.
- CPI for A, B and C are the same.
- 25% of the CPI is due to cache misses. Cache hits do not contribute to CPI but each cache miss takes 20ns to handle.

1A) What is the arithmetic mean of execution times for P1-P3 of A, B and C? **(4 points)**

1B) What is the geometric mean speedup over the reference machine C of A and B? If the result is not consistent with the result in 1A, explain why. **(4 points)**

1C) Establish the number of Misses Per Kilo Instructions (MPKI) for P1, P2 on A.
(4 points)

ASSIGNMENT 2

In this assignment, we consider how much instruction-level parallelism can be extracted statically through compiler techniques.

We consider in this assignment a VLIW architecture that can issue two memory, two floating-point, one integer, and one branch instruction each cycle according to the pipeline organization below.

- In each instruction word, two memory operations, two floating-point operations, and one integer/branch instruction can be scheduled.
- Forwarding circuitry is provided so that a value can be forwarded from a load, an integer, a floating-point, and a store operation to a subsequent instruction in 1, 0, 2, and 0 cycles, respectively.

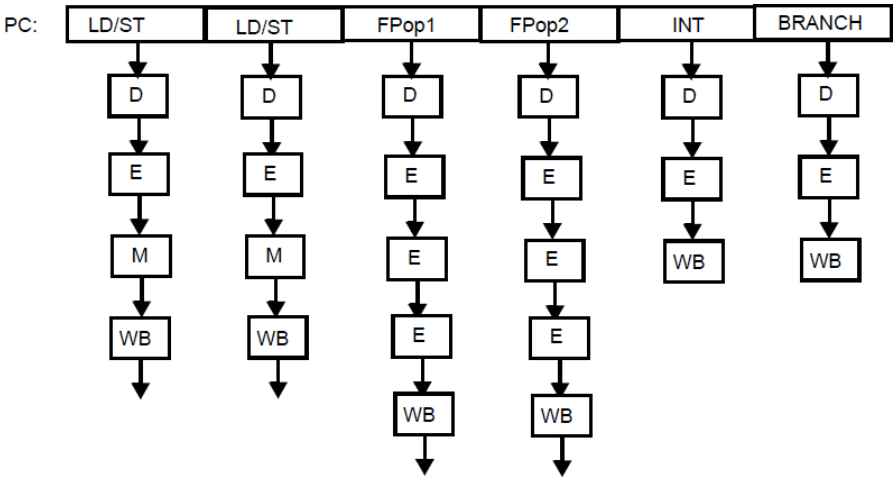

Consider the following simple computation:

```
LOOP:    L.D F0, 0(R1)
         ADD.D F4,F0 F2
         S.D F4,0(R1)
         SUBI R1,R1, #8
         BNE R1,R2, LOOP
```

2A) Unroll the loop three times and show the static schedule for it. **(4 points)**

2B) We want to use software pipelining to execute the loop on the VLIW architecture. Derive the kernel and the prologue code for a *software-pipelined loop.*
Hint: For the kernel, please fill out a schedule table for as many cycles needed to fill the "software pipeline" (i.e., not the VLIW pipeline). (**4 points)**

2C) The trace through basic blocks A and B is found to be more common that through basic blocks A and D. Schedule and optimize the trace to minimize the number of instructions in the common case and also provide fix-up code in case the less likely case of block A and D is executed. **(4 points)**

```
          LW R4,0(R1)
          ADDI R6,R4,#1
/* block A*/
          BEQ R5,R4,LAB
          LW R6,0(R2)
/*block B*/
/*block D* empty/
LAB: SW R6,0(R1)
```

ASSIGNMENT 3

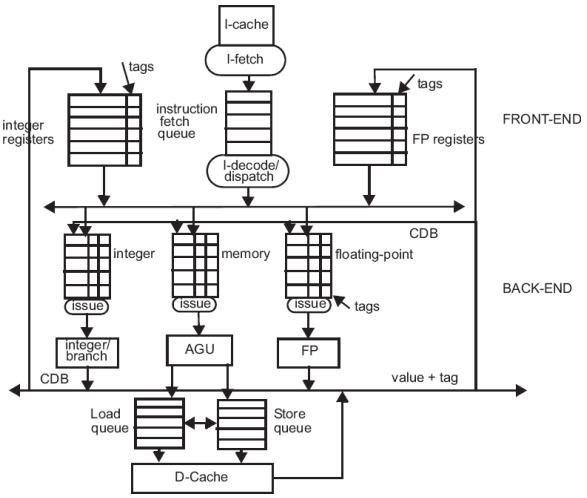

The diagram above shows a pipeline using the Tomasulo algorithm for dynamic instruction scheduling. The pipeline consists of 5 stages: Dispatch, Issue, Execution, Cache access and CDB write. Assume that all integer and branch instructions execute in a single cycle whereas floating-point instructions execute in three cycles, thus the duration of the execution stage varies. A cache access (instruction and data) takes a single cycle whereas stores are decomposed into two instructions (one for address generation and one for cache access).

3A) Describe all the necessary changes that are needed to extend the above pipeline to support speculation. Show the new block diagram. If new pipeline stages are added, explain the use of them. (**5 points)**

3B) <u>For the extended machine of assignment 3A</u>:

i) Determine how long time it takes to execute the first iteration of the code below.

ii) Fill in the appropriate table to explain your answer.

iii) Show also the execution of the first instruction in the second loop iteration. (**5 points)**

```
LOOP: L.S F2, 0(R1)
ADD.S F4, F2, F4
L.S F6, 0(R2)
ADD.S F8, F6, F4
SUBI R1, R1, #8
SUBI R2, R2, #8
SUBI R3, R3, #1
S.S F8, 8(R1)
BNEZ R3, LOOP
```

3C) What is a 2-level branch predictor? Give an example of a 2-level branch predictor to successfully predict the following branches: (**2 points)**

```
if (X==0) then a++;
if (X==1) then b++;
```

ASSIGNMENT 4

4A) Explain the 3C classification of cache misses and propose a methodology by which misses can be classified quantitatively. (**6 points)**

4B) A computer architect wants to establish the relative performance between a system with a blocking and a non-blocking cache using software prefetching. In software prefetching, prefetch instructions are appropriately inserted in the code shown below. Assume that five instructions are executed per iteration where the CPI for each instruction is one assuming that it doesn't cause a cache miss. On the other hand, if the instruction causes a cache miss, CPI is 100. In addition, prefetch instructions result in CPI=2. Caches have a block size of four words and each vector element is a word. In the non-blocking cache, there are 8 miss-status-holding registers.

```
for (i=0; i<1000; i++)
     C+=A[i];
```

First show how the code is annotated with prefetch instructions. How much faster does the program run on the system with a non-blocking cache using software prefetching? A convincing explanation that is easy to follow is needed for full points. (**6 points)**

ASSIGNMENT 5

5A) Consider a multicore system on a chip comprising a number of processors (cores) that are connected to a single-level private cache. The private caches use the write-back write policy. Xi=Ri and Xi=Wi, mean a read and a write request to the same address from processor i, respectively, where Wi=C means that value C is written by processor i. Now consider the following access sequence to a location that initially is assigned the value 0:

R1
R2
W1=1
R1
R2
W1=2
R2

What should be returned by the second read operation from processor 2, what is returned and what is the reason that the correct value is not returned given the cache write policy assumed? **(2 points)**

5B) Assume a write-back cache and the MSI-protocol. Explain the protocol by means of a finite-state machine with proper use of bus-side requests. **(6 points)**

5C) How is a simple five-stage pipeline extended to support fine-grain multithreading? **(4 points)**

****** GOOD LUCK! ******

Solutions to Exam 260107

ASSIGNMENT 1

1A)

We will use the formula T = IC x CPI x Tc
A: T_A is given for the three programs as $T_{A,1}$=1s, $T_{A,2}$=1.5s and $T_{A,3}$=2s. Hence, $\mathbf{T_{A,\,ave}}$ **=1.5 s**
B: IC is half of IC for A. Hence, $IC_{B,1}$=0.5 B, $IC_{B,2}$=0.75 B and $IC_{B,3}$=1 B. CPI is the same with A's. $CPI_{A,1}=T_{A,1}/(IC_{A,1}x\ T_{c,A}\ )$=1. In the same way, $CPI_{A,2}$ = $CPI_{A,3}$ =1. Since the instruction count is half of that of A but the frequency is double and thus the clock cycle half, so we get four times speedup, then we get $T_{B,1}$=0.25s, $T_{B,2}$=0.375s and $T_{B,3}$=0.5s and $\mathbf{T_{B,ave}}$ **=0.375 s**
C: A and C have the same ISA. Hence, IC must be the same. Since CPI and clock cycle time is the same, the execution times must be the same: $\mathbf{T_{C,ave}}$ **=1.5s.**

1B) $T_{geom,A}$ = $SQRT_3(T_{C,1}/T_{A,1}, T_{C,2}/T_{A,2}, T_{C,3}/T_{A,3},)$ = 1. For B it is 4.

1C) Establish the number of Misses Per Kilo Instructions (MPKI) for P1, P2 on A. CPI = CPI_0 + MPI x MP, where CPI_0 is the CPI_0 w/o cache misses, MPI is misses per instruction and MP is miss penalty for a cache miss in cycles. MPKI = 1000*(CPI – CPI_0)/MP. Since CPI is the same for all programs, we get $MPKI_A$=1000(1-0.75)/20 = 250/20 = **12.5**

ASSIGNMENT 2

2A)

Here is the unrolled loop (three times) after renaming and code motion and elimination of replicated subtract instruction and branches.

```
LOOP: L.D   F0, 0(R1)
      L.D   F6, -8(R1)
      L.D   F12, -16(R1)
      ADD.D F4,F0,F2
      ADD.D F8,F6,F2
      ADD.D F14,F12,F2
      S.D   F4,0(R1)
      S.D   F8,-8(R1)
      S.D   F14,-16(R1)
      SUBI  R1,R1, #24
      BNE   R1,R2, LOOP
```

The first iteration uses F0 and F4, whereas the second and third iteration use F6/F8 and

F12/F14, resp. By hoisting loads to the top and moving all stores to the end, we maximize the distance between a load and an add and an add and a store to avoid pipeline bubbles.

2B)

Let's denote the kernel instructions O1, O2 and O3. Given the forwarding mechanism assumed, the floating-point ADD must wait a cycle and the store (S.D) must wait two cycles.

```
LOOP: L.D    F0, 0(R1)     --O1
      ADD.D  F4,F0 F2      --O2
      S.D    F4,0(R1)      --O3
      SUBI   R1,R1, #8
      BNE    R1,R2, LOOP
```

We show below the prologue and kernel code below:

	ITE 1	ITE 2	ITE3	ITE4	ITE5	ITE6	ITE7	ITE8	ITE9	
INST1	O1									
INST2		O1								
INST3	O2		O1							
INST4		O2		O1						
INST5			O2		O1					PROLOGUE
INST6	O3			O2		O1				KERNEL
INST7		O3			O2		O1			
INST8			O3			O2		O1		
INST9				O3			O2			
INST10					O3			O2		
INST11						O3				
INST12							O3			

2C)

```
LW R4,0(R1)
LW R6,0(R2)
/* block A*/
BEQ R5,R4,LAB1
LAB: SW R6,0(R2)
….
LAB1: ADDI R6,R4,#1
J LAB
```

ASSIGNMENT 3

3A)

In order to extend the pipeline, a new stage is introduced called commit. For speculative

execution to be supported, *branch prediction* is needed so that instructions following a branch can be speculatively executed before the branch instruction is validated. A branch predictor is added to the instruction fetch stage and does a branch prediction in a single cycle. Each speculatively executed instruction will be tracked, in program order, in a structure called *reorder buffer*. The reorder buffer may also host all speculatively generated results as they cannot be written back to the register file until an instruction is non-speculative, i.e., committed. This happens when a preceding branch instruction is validated to be correctly predicted. Then the branch instruction will be *committed* (it becomes non-speculative). In subsequent cycles the subsequent speculative instructions are also committed. When an instruction is committed, the value it has generated (if any) will be written back to the register file and the instruction along with its result is removed from the reorder buffer.

3B)

We first enumerate the instructions as follows:

```
I1:    L.S F2, 0(R1)
I2     ADD.S F4, F2, F4
I3     L.S F6, 0(R2)
I4     ADD.S F8, F6, F4
I5     SUBI R1, R1, #8
I6     SUBI R2, R2, #8
I7     SUBI R3, R3, #1
I8     S.S F8, 8(R1)
I9     BNEZ R3, LOOP
```

We fill in the pipeline diagram table for the stages Dispatch, issue, Execution (start/end), Cache, CDB and Retire (Commit). We assume that the previous branch is correctly predicted and has committed so that all subsequent instructions will commit one by one.

#	Instruction	Dispatch	Issue	Execution (Start)	Execution (End)	Cache	CDB	Retire	Comments
I1	L.S F2, 0(R1)	1	2	3	3	4	5	6	
I2	ADD.S F4, F2, F4	2	6	7	9	-	10	11	wait for F2
I3	L.S F6, 0(R2)	3	4	5	5	6	7	12	
I4	ADD.S F8, F6, F4	4	11	12	14	-	15	16	wait for F4 on CDB @10
I5	SUBI R1, R1, #8	5	6	7	7	-	8	17	
I6	SUBI R2, R2, #8	6	7	8	8	-	9	18	
I7	SUBI R3, R3, #1	7	9	10	10	-	11	19	CDB conflict
I8	S.S-A F8, 8(R1)	8	9	10	10	-	-	-	
I8	S.S-B F8, 8(R1)	9	16	17	17	18	-	20	wait for F8 then wait to reach top of ROB
I9	BNEZ R3, LOOP	10	11	12	12	-	13	21	
I1	L.S F2, 0(R1)	11	12	13	13	14	15	22	dispatched right after branch

Most of the time, instructions flow through the pipeline but in some cases, they will be stalled due to either a structural hazard or a RAW hazard. There is a RAW hazard between I2 and I1 as well as I4 and I2/I3. Both will be issued in the next cycle after the dependent operand is written on the CDB by the dependent instruction. Next, instruction I7,1 cannot proceed to the CDB stage in cycle C10 because it is being used by instruction I2 causing a structural hazard. There is also a RAW hazard between I4 and I8 so I8 cannot be issued till cycle 13.

3C)
The conditions in the two if-clauses are not true at the same time; they are mutually exclusive. Hence, if the first branch is taken, the second branch is not taken and vice versa. This can be tracked by a two-level branch predictor. See textbook pages 121 and 122.

ASSIGNMENT 4

4A)
The 3C model refers to the root cause of misses in a cache.

First C: *Compulsory or Cold misses*. These are the misses that happen regardless of the size of the cache and corresponds to the number of unique blocks that are accessed during the execution of the program. The number of C-misses are the number of unique block addresses generated by the execution of a program.

Second C: *Capacity misses*. These are the additional misses that would not have happened in an infinitely sized cache. We can establish the number of capacity cache misses by first counting the number of misses on a fully associative cache with n optimal cache replacement algorithm (referred to as OPT) and subtracting it by the number of compulsory or cold misses.

Third C: *Conflict misses.* These are the additional misses that would not happen in a fully associative cache with the OPT replacement algorithm. We can establish the number of conflict misses by first counting the number of misses on a cache and subtracting it by the sum of the number of capacity and compulsory/cold misses.

4B) Let's first establish how long time the program takes on the architecture with a blocking cache.

We note that four iterations can be executed between cache misses as a block contains four vector elements. The first iteration will experience a cache miss and will take 100 + 4 cycles = 104 cycles. The next three iterations take 3x5 = 15 cycles. Hence, four iterations take 119 cycles and all 1000 iterations take 250 x 119 = 29,750 cycles.

Let's now annotate the code with a prefetch instruction (prefetch(A[i]):

```
for (i=0; i<1000; i++){
    prefetch(A[i+4]);
    C+=A[i];
}
```

The prefetch instruction will prefetch the next block if it is not present in cache. The first miss is unavoidable but subsequent misses can ideally be avoided. It means that the first four

iterations will take 119+2 (considering the prefetching CPI) cycles but the next 996 iterations will only take 5+2 cycles each. Hence the execution time using prefetching will be 121 + 996*7 = 7093 cycles. Hence, the speedup of a non-blocking cache with prefetching is 29,750/7093= 4.19.

ASSIGNMENT 5

5A)
The value 2 should be returned but, unfortunately, original value (0) is returned. The reason is that when processor 1 modifies the block, processor 2's cache is not notified.

5B)
See textbook pages 253-257 (first edition) or 240-243 (second edition).

5C)
A new pipeline stage is added between Fetch and Decode called thread switch or TS for short. This stage will switch to a new thread each cycle. Assuming N threads, we need N program counters where the one used is selected by TS. Moreover, N register files, one for each thread, are also needed.