

CENG 3420 Midterm (2021 Spring)

Name: _____

ID: _____

Solutions

Q0 (0 marks)

1. What is your last digit of your SID (0 is regarded as 10)? This value is defined as **NUM_1** in the whole question paper.
2. What is your last two digits of your SID (00 is regarded as 100)? This value is defined as **NUM_2** in the whole question paper.
3. What is your last three digits of your SID? This value is defined as **NUM_3** in the whole question paper.

Example: if your SID is 12345678, then $NUM_1 = 8, NUM_2 = 78, NUM_3 = 678$.

Q1 (10 marks)

Select and fill the correct answer.

1. RISC-V is a ____ and ____ ISA. (sol: A)
 - A. little-endian RISC
 - B. little-endian CISC
 - C. big-endian RISC
 - D. big-endian CISC
2. Virtual memory systems use ____ mechanism. (sol: A)
 - A. write-back
 - B. write-through
 - C. write-allocate
3. The largest positive integer in IEEE754 single-precision floating-point format is ____ (sol: D)
 - A. $2^{126} - 2^{103}$
 - B. $2^{127} - 2^{104}$
 - C. $2^{127} - 2^{103}$
 - D. $2^{128} - 2^{104}$

4. ____ is a type of single address space multiprocessor in which some memory accesses are much faster than others depending on which processor asks for which word. (sol. B)
- A. UMA
B. NUMA
5. Dividing 10101_2 by 11_2 , the quotient is ____ and the remainder is _____. (sol: C)
- A. 100_2 01_2
B. 111_2 01_2
C. 111_2 0_2
D. 100_2 0_2
6. ____ stores data as electric charge on a capacitor. (sol: B)
- A. SRAM
B. DRAM

Q2 (5 marks)

- Describe the steps that transform a program written in a high-level language such as C into a representation that is directly executed by a computer processor.
- Write C statements that corresponds to the two RISC-V assembly instructions below.

```
srlw t0, t1, 3
andi t0, t2, 0xFEF
```

You should define variables that hold corresponding register values explicitly.

Answer:

- The program is compiled into an assembly language program, which is then assembled into a machine language program.

2.

```
int t0;
int t1;
int t2;
t0 = t1 >> 3;
t0 = t2 & 0xFEF;
```

Q3 (10 marks)

Consider a RISC-V code snippet, and assume that the code is running on a 5-stage RISC-V core.

```

add x15, x12, x11
ld x13, 4(x15)
ld x12, 0(x2)
or x13, x15, x13
sd x13, 0(x15)

```

1. Find all data dependences in this instruction sequence.
2. If the design has not implemented any forwarding or hazard detection, insert NOPs to ensure the correct execution.
3. If there is forwarding, for the first seven cycles during the execution of this code, specify which signals are asserted in each cycle by hazard detection and forwarding units.

Answer:

1. The 2nd instruction depends on the 1st instruction; The 4th instruction depends on 1st and 2nd instruction; The 5th instruction depends on 4th instruction.
2. add nop between(1 and 2 (2nop), 3 and 4(1nop), 4 and 6(2nop))
3. PCWrite = IF/IDWrite = 1; 1st ForwardA = X ForwardB = X; 2nd ForwardA = X ForwardB = X;3rd ForwardA = 00 ForwardB = 00;4th ForwardA = 10 ForwardB = 0;5th ForwardA = 01 ForwardB = 1;6th ForwardA = 01 ForwardB = 10;7th ForwardA = 01 ForwardB = 10;

Q4 (15 marks)

We have a loop as follows.

```

LOOP: ld a0, 0(a3)
      ld a1, 8(a3)
      add a2, a0, a1
      subi a3, a3, 16
      bnez a2, LOOP

```

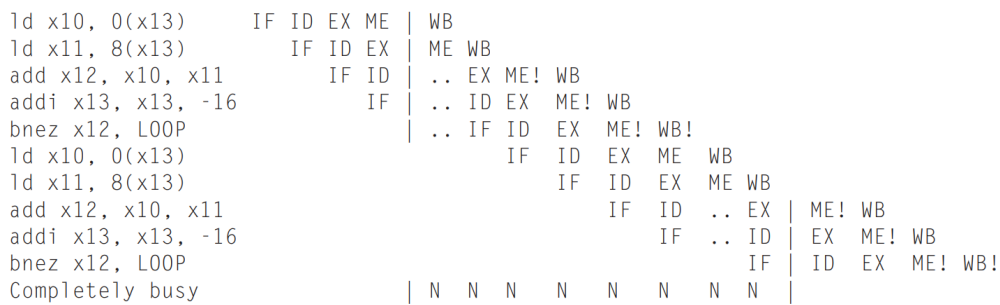
Assume that we have the perfect branch prediction (i.e., no stalls due to control hazards) and there are no delay slots that the pipeline has full forwarding support, and that branches are resolved in the EX (as opposed to the ID) stage.

1. Show a pipeline execution diagram for the first two iterations of this loop.
2. Mark pipeline stages that do not perform useful work. How often while the pipeline is full do we have a cycle in which all five pipeline stages are doing useful work? (Begin with the cycle during which the `subi` is in the IF stage. End with the cycle during which the `bnez` is in the IF stage)

Answer:

1. As shown in the Figure 1

Figure 1: Pipeline execution diagram



2. In a particular clock cycle, a pipeline stage is not doing useful work if it is stalled or if the instruction going through that stage is not doing any useful work there. As the diagram above shows, there are not any cycles during which every pipeline stage is doing useful work.

Q5 (15 marks)

Consider the following loops written in C programming language:

```
for(int i = 0; i < a; i++)
    for(int j = 0; j < b; j++)
        A[j * 4] = i + j;
```

The values of a, b, i and j are in registers t0, t1, t2 and t3, respectively. The register a0 holds the base address of the array A. Translate to the corresponding RISC-V assembly code.

Answer:

```
LOOPI:
addi t2, zero, 0 # init: i = 0
bge t2, t0, ENDI # while i < a
addi t4, a0, 0 # t4 = &A
addi t3, zero, 0 # init: j = 0
LOOPJ:
bge t3, t1, ENDJ # while j < b
add t5, t2, t3 # t5 = i + j
sd t5, 0(t4) # A[j * 4] = i + j
addi t4, t4, 16 # t4 = &A[4 * (j + 1)]
addi j, j, 1 # j++
jal zero, LOOPJ
ENDJ:
addi t2, t2, 1 # i++
jal zero, LOOPI
ENDI
```

Correction: addi t2, zero, 0 should be moved outside LOOPI.

Q6 (15 marks) There is a computer that has a 64MB byte-addressable main memory. Instructions and data are stored in separated caches, each of which has eight 64B cache lines. The data cache uses **direct-mapping**. Now there are two programs in the following form.

Program A:

```

int a[64][64];
int sum_array1() {
    int i, j, sum = 0;
    for(i = 0; i < 64; i++)
        for(j = 0; j < 64; j++)
            sum += a[i][j];
    return sum;
}

```

Program B:

```

int a[64][64];
int sum_array1() {
    int i, j, sum = 0;
    for(j = 0; j < 64; j++)
        for(i = 0; i < 64; i++)
            sum += a[i][j];
    return sum;
}

```

Suppose `int` data is represented in 32-bit 2's complement and `i`, `j`, `sum` are stored in specific registers. Arrays are stored in row-major with the start address NUM_3_{10} (i.e., `NUM_3` in decimal) in the main memory. Answer the following questions.

1. What are the line numbers of the main memory blocks that contain `a[0][31]` and `a[2][2]` respectively? (Cache line number starts from 0)
2. What are the data cache hit rates of program A and B?

Answer:

1. a. $\text{NUM_3} + 31 \times 4 = \text{var}$, $\frac{\text{var}}{64} = \text{result}$.
 b. $\text{NUM_3} + (64 \times 2 + 2) \times 4 = \text{var1}$, $\frac{\text{var1}}{64} = \text{var2}$, $\text{var2} \bmod 8 = \text{result}$
2. The size of array `a` is $64 \times 64 \times 4 = 2^{14}\text{B}$. Since the size of a block is 64B , it takes 2^8 main memory blocks to store it. Under the condition of row-major, 2^8 times of cache miss will appear. Therefore, the hit rate of program A is $\frac{(2^{12}-2^8)}{2^{12}} = 93.75\%$. For program B, the hit rate is 0.

Q7 (5 marks)

Describe two cache replacement strategies.

1. Write-through
2. Write-back

Answer:

1. It always write the data into both the cache block and the next level RAM in the memory hierarchy.
2. It writes to the memory hierarchy when the cache block is **evicted**.

Q8 (15 marks)

The following table shows the different instructions ratios

Table 1: Instruction mix

R-type	I-type (non-Id)	Load	Store	Branch	Jump
$2 \times \text{NUM_1}$	25%	22%	8%	$45\% - 3 \times \text{NUM_1}$	NUM_1

1. What fraction of all instructions use the data memory?
2. What fraction of all instructions use the instruction memory?
3. What fraction of all instructions use the sign extend?
4. What is the sign extend doing during cycles in which its output is not needed?

Answer:

1. Only Load and Store use data memory. 30%
2. 100%
3. Only R-type instructions do not use the Sign extender. $1 - 2 \times \text{NUM_1}$
4. The sign extend produces an output during every cycle. If its output is not needed, it is simply ignored.

Q9 (10 marks)

Given the following specifications of the datapath latencies:

Table 2: Specification of the datapath latencies (unit: *ps*)

Stages	IF	ID	EX	MEM	WB
Latencies (ps)	NUM_2	NUM_3	550	800	600

1. What is the clock cycle time in a pipelined and non-pipelined processor?
2. What is the total latency of an $1w$ instruction in a pipelined and non-pipelined processor?
3. If we can split one stage of the pipelined datapath into two new stages, each with half the latency of the original stage, which stage would you split and what is the new clock cycle time of the processor?

Answer:

1. Non-pipelined (sum of all stages): $\text{NUM_2} + \text{NUM_3} + 550 + 800 + 600$ *ps*; Pipelined (the slowest stage): $\max\{\text{NUM_3}, 800\}$ *ps*.
2. The instruction $1w$ takes all five stages. Non-pipelined: $\text{NUM_2} + \text{NUM_3} + 550 + 800 + 600$ *ps*; Pipelined (takes five stages at the slowest cycle): $\max\{\text{NUM_3}, 800\} \times 5$ *ps*.
3. Split ID / MEM stage (find what is the slowest). If you split ID, and then the new clock cycle will be 800 *ps*. If you split MEM, and then the new clock cycle will be $\max\{\frac{\text{NUM_3}}{2}, 600\}$ *ps*.