

# Worksheet 6

Name: \_\_\_\_\_

UID: \_\_\_\_\_

1. For the following assembly, draw a data flow graph and identify the critical path.

```
.L1:
    vmulsd    (%rdx), %xmm0, %xmm0
    addq     $8, %rdx
    cmpq     %rax, %rdx
    jne      .L1
```

Given the following table, what is the lowest bound latency to execute n iterations of this loop for integer operations? Floating point operations?


Latency Table (CPE)	Integer	Double
Arithmetic (except mul)	1	3
Multiply	2	4
Load Store	1	1

2. Suppose our memory can store 64 bytes of data, addressed between 6'b000000 and 6'b111111, and that we have stored a 4 x 4 array of ints in memory. This could be declared by:

```
int arr[4][4] = {
    {0,1,2,3},
    {4,5,6,7},
    {8,9,10,11},
    {12,13,14,15}
};
```

In addition, suppose we have a **fully associative cache** (only one set) with a single line of **16 bytes**.

Then, our addresses can be partitioned into **2 tag bits** along with **4 block offset bits**. For example, 6'b000000 would be formatted as:



- A. Now, say we accessed the array using the following loop, and our cache is cold (empty) to start. How many cache hits/misses would there be? (Think about how the array is stored in memory, and how data will be cached).

```
let sum = 0;
for (int i = 0; i < 4; i++) {
    for (int j = 0; j < 4; j++) {
        sum += arr[i][j];
    }
}
```

- B. Again, assume our cache is empty to start. How many cache hits/misses would there be if we used this loop instead?

```
let sum = 0;
for (int j = 0; j < 4; j++) {
    for (int i = 0; i < 4; i++) {
        sum += arr[i][j];
    }
}
```

3. Suppose we have a direct-mapped cache with **2 sets, 1 line per set, 4 bytes per line**, and 4 bit addresses that look like:

tag bits: $t = 1$	set index: $s = 1$	offset bits: $b = 2$
-------------------	--------------------	----------------------

Then our address space will look like:

Address (decimal)	Tag bits	Index bits	Offset bits
0	0	0	00
1	0	0	01
2	0	0	10
3	0	0	11
4	0	1	00
5	0	1	01
6	0	1	10
7	0	1	11
8	1	0	00
9	1	0	01
10	1	0	10
11	1	0	11
12	1	1	00
13	1	1	01
14	1	1	10
15	1	1	11

Initially, our cache is cold (empty), and we may represent it like this:

Set	Valid	Tag	Byte 0	Byte 1	Byte 2	Byte 3
0	0					
1	0					

- A. Now, suppose we want to read **1-byte words** from memory, where the word located at memory address  $i$  is indicated by  $M[i]$ . What happens to the cache when the CPU performs the following sequence of reads?

**read words at memory address: 0, 14, 3, 11**

3. Suppose we wish to write a procedure that computes the inner product of two vectors  $u$  and  $v$ . An abstract version of the function has a CPE of 14-18 with x86-64 for different types of integer and floating-point data. By doing the same sort of transformations we did to transform the abstract program *combine1* into the more efficient *combine4*, we get the following code:

```
/* Inner product. Accumulate in temporary */
void inner4(vec_ptr u, vec_ptr v, data_t* dest) {
    long i;
    long length = vec_length(u);
    data_t *udata = get_vec_start(u);
    data_t *vdata = get_vec_start(v);
    data_t sum = (data_t) 0;


    for (i = 0; i < length; i++) {
        sum = sum + udata[i] * vdata[i];
    }
    *dest = sum;
}
```

Our measurements show that this function has CPEs of 1.50 for integer data and 3.00 for floating-point data. For data type double, the x86-64 assembly code for the inner loop is as follows:


```
.L15:                                ;loop:
    vmovsd  0(%rbp, %rcx, 8), %xmm1    ; Get udata[i]
    vmulsd  (%rax, %rcx, 8), %xmm1, %xmm1 ; Multiply by vdata[i]
    vaddsd  %xmm1, %xmm0, %xmm0        ; Add to sum
    addq    $1, %rcx                   ; Increment i
    cmpq    %rbx, %rcx                 ; Compare i:limit
    jne     .L15                       ; If !=, goto loop
```

Latency Table	Integer	Double
Arithmetic (except mul)	1	3
Multiply	2	4
Load Store	1	1

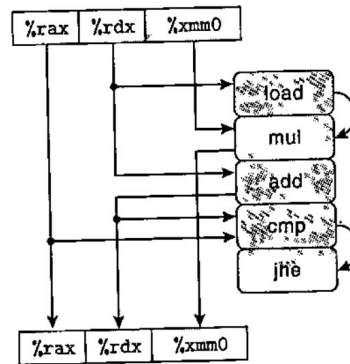
a. Diagram how this instruction sequence would be decoded into operations and show how the data dependencies between them would create a critical path of operations, in the style of the textbook's figures shown below



(a)



(b)



b. For data type double, what lower bound on the CPE is determined by the critical path?

c. Assuming similar instruction sequences for the integer code as well, what lower bound on the CPE is determined by the critical path for integer data?

d. Explain how the two floating-point versions can have CPEs of 3.00, even though the multiplication operation requires 5 clock cycles.