

Computer Architecture Sample Test 1

Question 1. Suppose we have 32-bit memory addresses, a byte-addressable memory, and a 512 KB (2^{19} bytes) cache with 32 (2^5) bytes per block.

- a) How many total lines are in the cache?
- b) If the cache is direct-mapped, how many cache lines could a specific memory block be mapped to?
- c) If the cache is direct-mapped, what would be the format (tag bits, cache line bits, block offset bits) of the address? (Clearly indicate the number of bits in each)
- d) If the cache is 2-way set associative, how many cache lines could a specific memory block be mapped to?
- e) If the cache is 2-way set associative, how many sets would there be?
- f) If the cache is 2-way set associative, what would be the format of the address?
- g) If the cache is fully-associative, how many cache lines could a specific memory block be mapped to?
- h) If the cache is fully-associative, what would be the format of the address?

Question 2. What is the difference between the cache write policies: writeback and writethrough with respect to each of the following:

- a) the relative frequency of memory accesses
- b) the ability of a multiple CPU system to keep caches coherent

Question 3. Assume that a program wants to copy a 500 word block of data from an I/O device to memory. Complete the following table for each I/O technique.

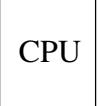
	Programmed I/O	Interrupt-driven I/O	Direct-memory Access
Total number of interrupts for whole block transfer			
Total number of READ I/O commands issued by the CPU to the I/O module			

Question 4. Assume special I/O instructions (i.e., isolated I/O) are used to fill I/O-controller registers. Why can't a user program use these instructions to communicate with the I/O device directly and "by-pass" the operating system's protection checking?

Question 5. When a DMA controllers needs to use the bus to access memory, it performs "cycle stealing" by making the CPU wait if it is also trying to access memory. Even though the CPU is continually executing, why might the DMA device often find the bus to memory free when it needs to access memory?

Question 6. Consider the demand paging system **with 1024-byte pages**.

Running
Process A



Page Table for A

	Frame#	Valid Bit
0		
1		
2		
3		
4		
5		
6		

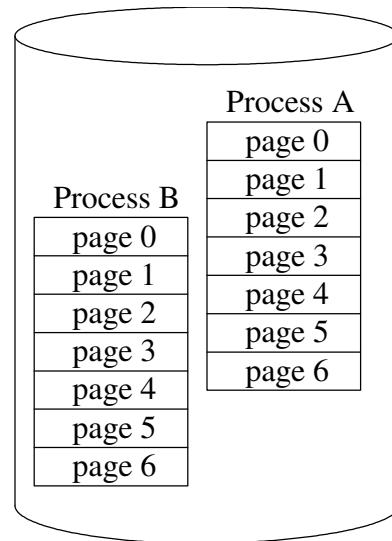
page# offset

--	--

Frame Number	Physical Memory
0	page 2 of B
1	page 3 of A
2	page 4 of A
3	page 5 of B
4	page 0 of A
5	page 2 of A
6	page 4 of B

frame# offset

--	--



Logical Addr.

Physical Addr.

a) Complete the above page table for **Process A**.

b) If process A is currently running and the CPU generates a logical/virtual address of 3100_{10} , then what would be the corresponding physical address?

Question 7. You are to assume the same 5-stage pipeline discussed in class when answering these questions. Assume that the first register in an arithmetic operation is the destination register, e.g., in “ADD R3, R2, R1” register R3 receives the result of adding registers R2 and R1.

- a. What would the timing be **without** bypass-signal paths/forwarding (use “stalls” to solve the data hazard)? (This code might require more or less than 15 cycles)

Instructions	Time →														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
ADD R3, R2, R1	F	D	E	M	W										
STORE R3, 8(R4)															
LOAD R4, 16(R3)															
SUB R3, R4, R1															
MUL R6, R3, R4															
STORE R6, 4(R5)															

(Assume that a register **cannot** be written and the new value read in the same stage.)

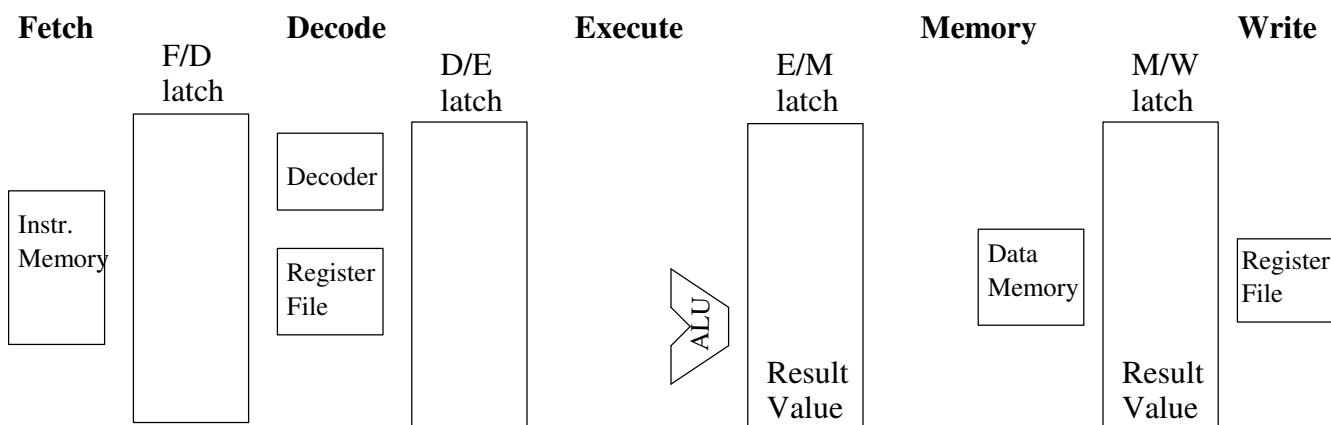
- b. What would the timing be **with** bypass-signal paths?

(This code might require more than 15 cycles)

Instructions	Time →														
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
ADD R3, R2, R1	F	D	E	M	W										
STORE R3, 8(R4)															
LOAD R4, 16(R3)															
SUB R3, R4, R1															
MUL R6, R3, R4															
STORE R6, 4(R5)															

(Assume that a register **cannot** be written and the new value read in the same stage.)

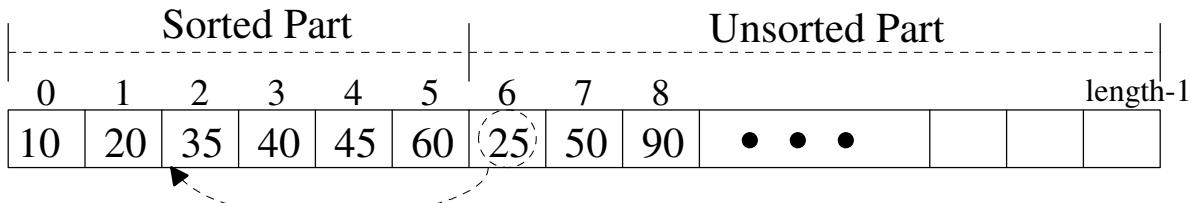
- c. Draw ALL the bypass-signal paths needed for the above example.



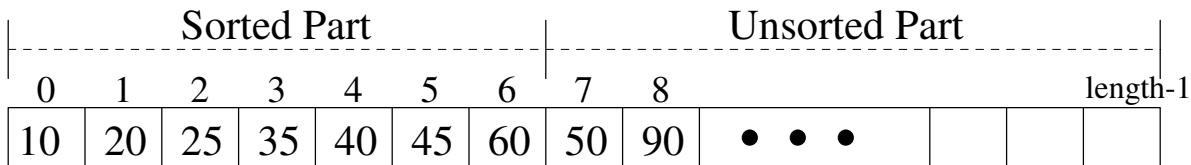
Question 8. Another simple sort is called insertion sort. Recall that in a simple sort:

- the outer loop keeps track of the dividing line between the sorted and unsorted part with the sorted part growing by one in size each iteration of the outer loop. (below the firstUnsortedIndex keeps track of the dividing line)
 - the inner loop's job is to do the work to extend the sorted part's size by one.

After several iterations of insertion sort's outer loop, an array might look like:



In insertion sort the inner-loop takes the "first unsorted item" (25 at index 6 in the above example) and "inserts" it into the sorted part of the array "at the correct spot." After 25 is inserted into the sorted part, the array would look like:



Consider the following insertion sort algorithm that sorts an array numbers:

```
InsertionSort(numbers - address to integer array, length - integer)
    integer firstUnsortedIndex, testIndex, elementToInsert;
    for firstUnsortedIndex = 1 to (length-1) do
        testIndex = firstUnsortedIndex-1;
        elementToInsert = numbers[firstUnsortedIndex];
        while (testIndex >=0) AND (numbers[testIndex] > elementToInsert ) do
            numbers[ testIndex + 1 ] = numbers[ testIndex ];
            testIndex = testIndex - 1;
        end while
        numbers[ testIndex + 1 ] = elementToInsert;
    end for
end InsertionSort
```

- Where in the code would unconditional branches be used and where would conditional branches be used?
- If the compiler could predict by opcode for the conditional branches (i.e., select whether to use machine language statements like: "BRANCH_LE_PREDICT_NOT_TAKEN" or "BRANCH_LE_PREDICT_TAKEN"), then which conditional branches would be "PREDICT_NOT_TAKEN" and which would be "PREDICT_TAKEN"?

c) Assumptions:

- length = 100 and the numbers are initially in **descending** order before the insertion sort algorithm is called
- the five-stage pipeline discussed in class
- the outcome of conditional branches is known at the end of the E stage
- target addresses of all branches is known at the end of the D stage
- ignore any data hazards

Under the above assumptions, answer the following questions:

i) If fixed predict-never-taken is used by the hardware, then what will be the total branch penalty (# cycles wasted) for the algorithm? (Here assume NO branch-prediction buffer)

ii) If a branch target buffer with one history bit per entry is used, then what will be the total branch penalty (# cycles wasted) for the algorithm? (Assume predict-not taken is used if there is no match in the branch-prediction buffer) Explain your answer.

iii) If a branch target buffer with two history bit per entry is used, then what will be the total branch penalty (# cycles wasted) for the algorithm? (Assume predict-not taken is used if there is no match in the branch-prediction buffer) Explain your answer.