1 (10%) Please show, when a system call is invoked, how the control flow is transferred from the user process to the function (inside OS) which implements the system call.

2. (30%) There are three processes, P1, P2 and P3 which born times and executing behavior are shown as follows. The P1 begins with a CPU burst which length is 1000 ms, followed by an I/O burst(using disk) which needs 100 ms to complete. The other numbers can be understood similarly.

	born time(ms)	CPU/disk bursts(ms)
P1:	0	1000(CPU), 100(disk), 100(CPU)
P2:	3	3(CPU), 5(disk), 3(CPU)
P3:	5	3(CPU), 5(disk), 3(CPU)

Assume that the CPU scheduling algorithm is round robin with 4 ms quantum. The ready queue is empty and P1 is running at time 0. Also the first quantum of P1 begins at time 0.

Please show the content of ready queue, the content of disk queue, and the running process at time 6, 9, and 13 ms respectively. Explanation of how you get the answer is required for the score.

Note: When a queue contains more than one items, you have to specify the order among them.

3 (10%) Given the following address mappings (from logical to physical, in decimal), please show how you figure out the maximal possible page size.

123 --> 523, 377 --> 777, 1234 --> 2034, 2013 --> 3213

4 Memory hierarchy (20%)

Assume we have a computer with a 32KB direct-mapped L1 cache, which is indexed by physical memory addresses. The L1 cache has miss penalty of 80ns, 1.8 references per instruction, and the cache miss rate is 2%. The CPI is 2.0 with perfect cache. The clock cycle time is 1ns with this design.

a. (8 pts) Assume we change the cache design to a 64KB 2-way set-associative cache. The new design decreases the cache miss rate to 1.2%. However, it also increases the clock

cycle time to 1.2ns. Which one is better for average memory access time? Which one is better for CPU time?

b. (6 pts) When utilizing virtual memory addresses to index the cache, can we reduce cache hit time? Identify the issues that should be considered to guarantee correct program executions when indexing caches by virtual memory addresses.

c. (6 pts) Compare SRAMs and DRAMs in terms of cost and performance. Which one is better for implementing caches? Which one is better for implementing main memories? You have to state your reasons.

5 Implementation of a datapath (20%)

A conventional single-cycle datapath, and a commonly seen five-stage pipeline are shown in Figure. 1 and Figure. 2.

a. (5 pts) Compare the single-cycle datapath and five-stage pipeline in terms of hardware cost, cycle-time, execution time of single instruction, and throughput.

b. (6 pts) With the MIPS instruction set, and the pipeline as shown in Figure. 2, how many stall cycles would encounter to execute a branch instruction? Why? Provide a method to alter the pipeline to decrease the number of stall cycles for a branch instruction.

c. (3 pts) If an ALU supports 6 different operations, how many bits (at least) does the signal of selecting ALU operation need? Why?

d. (6 pts) Increasing the number of pipeline stage reduces the cycle time and improves throughput. However, most of the processor vendors have turned to utilizing multi-core architecture, rather than more pipeline stages to increase system throughput. Why?

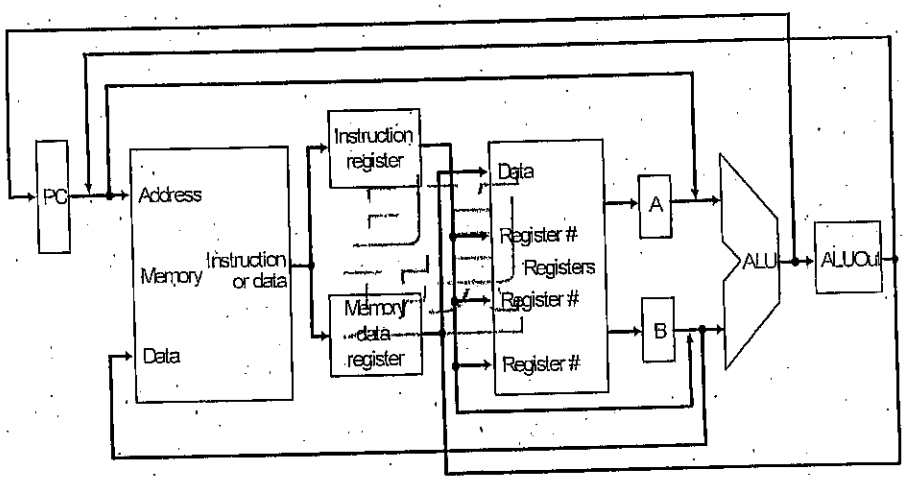

Figure. 1 Single-cycle datapath.

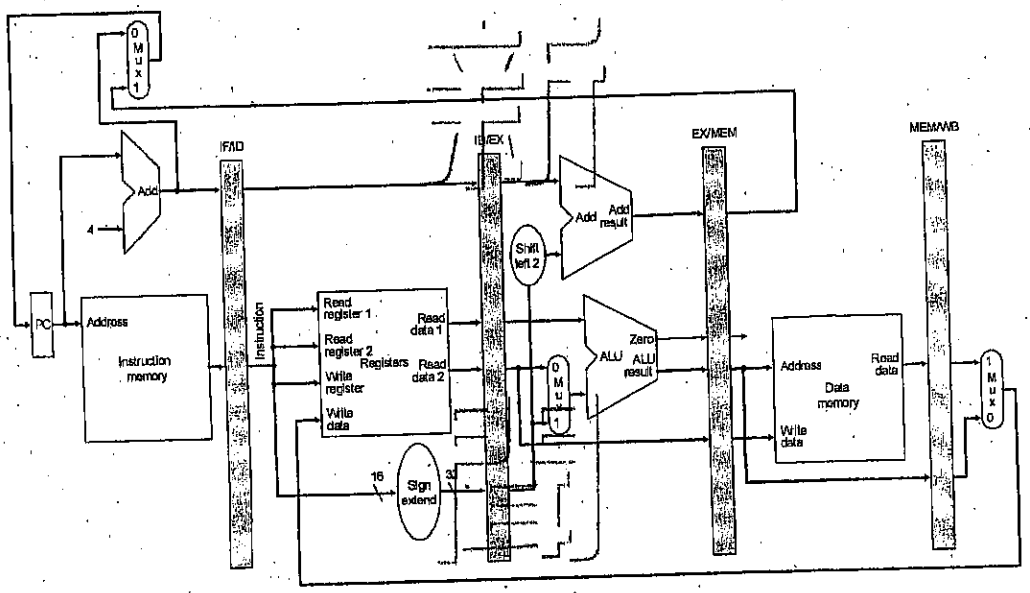

Figure. 2 Five-stage pipeline.

6 Instruction set (10%)

All Instructions in a MIPS instruction set have the same length, and each instruction only performs one operation. For an x86 instruction set, instructions may have variable lengths, and a long instruction may perform more than one operation, e.g., retrieving data from memory and performing arithmetic operation on the data. Compare the two instruction sets, in terms of code size, complexity of the compiler design, complexity of the decoding circuit, and performance.