USN | | | | | | | | | | |

06CS81

Eighth Semester B.E. Degree Examination, June 2012
Advanced Computer Architecture

Time: 3 hrs. Max. Marks:100

Note: ***Answer FIVE full questions, selecting at least TWO questions from each part.***

PART – A

1 a. Define computer architecture. Illustrate the seven dimensions of an ISA. **(08 Marks)**

b. Explain in brief measuring, reporting and summarizing performance of computer system. **(08 Marks)**

c. Assume a disk subsystem with the following components and MTTF:
- 10 disks, each rated at 1000000 - hour MTTF.
- 1 SCSI controller, 500,000 – hour MTTF.
- 1 power supply, 200,000 – hour MTTF.
- 1 fan, 200,000 – hour MTTF.
- 1 SCSI cable, 1,000,000 – hour MTTF.

Using the simplifying assumptions that the lifetimes are exponentially distributed and that failures are independent, compute the MTTF of the system as a whole. **(04 Marks)**

2 a. Explain how pipeline is implemented in MIPS. **(06 Marks)**

b. Explain different techniques in reducing pipeline branch penalties. **(06 Marks)**

c. What are the major hurdles of pipelining? Explain briefly. **(04 Marks)**

d. Consider the unpipelined processor in RISC. Assume that it has a 1 ns clock cycle and that it uses 4 cycles for ALU operations and branches and 5 cycles for memory operations. Assume that the relative frequencies of these operations are 40%, 20% and 40% respectively. Suppose that due to clock skew and setup, pipelining the processor adds 0.2 ns of overhead to the clock. Ignoring any latency impact, how much speedup in the instruction execution rate will we gain from a pipeline? **(04 Marks)**

3 a. What are the basic compiler techniques for exposing ILP? Explain briefly. **(08 Marks)**

b. Explain Tomarulo's algorithm, sketching the basic structure of a MIPS floating point unit. **(08 Marks)**

c. Explain true data dependence, name dependence and control dependence with an example code fragment. **(04 Marks)**

4 a. Explain exploiting ILP using dynamic scheduling, multiple issue and speculation. **(08 Marks)**

b. Explain Pentium 4 pipeline supporting multiple issue with speculation. **(08 Marks)**

c. Suppose we have a VLIW that could issue two memory references, two FP operations and one integer operation or branch in every clock cycle, show an unrolled version of the loop $x(i) = x(i) + s$, for such a processor. Unroll as many times as necessary to eliminate any stalls. Ignore delayed branches.

MIPS	Code
Loop: L. D	F_0, $O(R_1)$;
ADD.D	F_4, F_0, F_2;
S.D	F_4, $O(R_1)$;
DADDUI	R_1, R_1, #-8;
BNE	R_1, R_2, Loop

(04 Marks)

PART – B

5 a. Explain basic schemes for enforcing coherence. **(08 Marks)**

b. Explain performance of symmetric shared memory multiprocessors. **(08 Marks)**

c. Suppose we have an application running on a 32-processor multiprocessor, which has a 200 ns time to handle reference to a remote memory. For this application, assuming that all the references except those involving communication hit in the local memory hierarchy, which is slightly optimistic. Processors are stalled on a remote request, and the processor clock rate is 2 GHz. If the base CPI (assuming that all references hit in the cache) is 0.5, how much faster is the multiprocessor if there is no communication versus if 0.2% of the instructions involve a remote communication reference? **(04 Marks)**

6 a. Explain the six basic cache optimization techniques. **(10 Marks)**

b. Given the data below, what is the impact of second level cache associativity on its mass penalty?

- Hit time L_2 for direct mapped = 10 clock cycles
- Two way set associativity increases hit time by 0.1 clock cycles to 10.1 clock cycles.
- Local miss rate L_2 for direct mapped = 25%
- Local miss rate L_2 for two-way set associative = 20%
- Miss penalty L_2 = 200 clock cycles. **(06 Marks)**

c. What are the techniques for fast address translation? Explain. **(04 Marks)**

7 a. Explain any 3 advanced cache optimization techniques. **(08 Marks)**

b. Explain memory technology and optimizations. **(06 Marks)**

c. Assume that the hit time of a two-way set associative first level data cache is 1.1 times faster than a four-way set associative cache of the same size. The miss falls from 0.049 to 0.044 for an 8 KB data cache. Assume a hit is 1 clock cycle and that the cache is the critical path for the clock. Assume that the miss penalty is 10 clock cycles to the L_2 cache for the two-way set associative cache, and that the L_2 cache does not miss. Which has the faster average memory access time? **(06 Marks)**

8 a. Explain detecting and enhancing loop level parallelism for VLIW. **(06 Marks)**

b. Explain Intel-IA 64 architecture with a neat diagram. **(06 Marks)**

c. Explain hardware support for exposing parallelism for VLIW and EPIC. **(08 Marks)**

* * * * *