Reg. No. : | | | | | | | | | | | |

Question Paper Code : Z 7214

M.E. DEGREE EXAMINATION, NOVEMBER/DECEMBER 2009.

Second Semester

Applied Electronics

AN 1652 — COMPUTER ARCHITECTURE AND PARALLEL PROCESSING

(Common to M.E. – Computer and Communication and M.E. – VLSI Design)

(Regulation 2005)

Time : Three hours Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. List the properties satisfied by information stored in a memory hierarchy.

2. State Bernstein's conditions for parallelism.

3. Which PRAM variant can be the best model for SIMD machine? How?

4. What are the factors affecting cache hit ratios?

5. How is flow control done with buffering method using virtual cut-through routing?

6. What are possible hazards between operations in an instruction pipeline and how is it avoided?

7. Distinguish between various memory organization schemes for vector accesses.

8. State the advantages and shortcomings of context switching policies on a multithreaded architecture.

9. State the features of two language oriented programming models for parallel processing.

10. What are the opportunities for exploiting parallism in Mach kernel.

PART B — (5 × 16 = 80 marks)

11. (a) (i) Explain the applicability and the restrictions involved in using Amdahl's law, Gustafsons's law, and Sun and Ni's law to estimate the speed-up performance of n processor system compared with that of a single processor system ignoring all communication overheads. (8)

(ii) Analyze the data dependencies among the following statements in the given program :

s1 : Load R1, 1024

s2 : Load R2, M (10)

s3 : Add R1, R2

s4 : Store M(1024), R1

s5 : Store M((R2)), 1024

where (Ri) means the content of register Ri and M(10) contains 64 initially.

(1) Draw a dependence graph to show all the dependencies.

(2) Are there any resource dependencies if only one copy of each functional unit is available in the CPU? (8)

Or

(b) (i) Explain the characteristics of Multivector and SIMD computers. (8)

(ii) Explain the various levels of parallelism in program execution on modern computers. (8)

12. (a) (i) The main memory of a computer is organized as 64 blocks, with a block size of eight words. The cache has eight block frames. Draw all lines showing the mappings as clearly as possible for the following.

(1) Show the address mapping and the address bits that identify the tag field, the block number and the word number.

(2) Show the fully associative mapping and the address bits that identify the tag field and the word number.

(3) Show the two way set associative mapping and the address bits that identify the tag field, the set number and the word number.

(4) Show the sector mapping with four blocks per sector and the address bits that identify the sector number, the block number and the word number. (8)

(ii) Consider a cache (M1) memory and memory (M2) hierarchy with the following characteristics :

M1 :16 K words, 50 ns access time

M2 : 1 M words, 400 ns access time

Assume eight-word cache blocks and a set size of 256 words with set- associative mapping. (1) Show the mapping between M2 and M1 (2) Calculate the effective memory-access time with a cache hit ratio of h=0.95. (8)

Or

(b) (i) Explain the difference between superscalar and VLIW architectures in terms of hardware and software requirements. (8)

(ii) A two - level memory system has eight virtual pages on a disk to be mapped into four page frames in the main memory. A certain programme generates the following page trace:

1, 0, 2, 2, 1, 7, 6, 7, 0, 1, 2, 0, 3, 0, 4, 5, 1, 5, 2, 4, 5, 6, 7, 6, 7, 2, 4, 2, 7, 3, 3, 2, 3

(1) Show the successive virtual pages residing in the four page frames with respect to the above trace using LRU replacement policy. Compute the hit ratio in the main memory. Assume the PFs are initially empty.

(2) Repeat for the circular FIFO policy. Compare the hit ratios and comment on the effectiveness of using the circular FIFO policy to approximate LRU policy with respect to this particular page trace. (8)

13. (a) (i) Consider the following reservation table for a four stage pipeline with a clock cycle τ = 20 ns.

	1	2	3	4	5	6
S1	X					X
S2		X		X		
S3			X			
S4				X	X	

One non compute delay stage into the pipeline can be inserted to make a latency of 1 permissible in the shortest greedy cycle. The purpose is to yield a new reservation table leading to an optimal latency equal to the upper bound.

(1) Show the modified reservation table with five rows and seven columns.

(2) Draw the state transition diagram for the optimal cycle.

(3) List all the simple and greedy cycles from the state diagram.

(4) Prove that the new MAL equals the lower bound.

(5) What is the optimal throughput of this pipeline? (10)

(ii) Consider a four stage floating point adder with a 10 ns delay per stage which equals the pipeline clock period. Name the appropriate functions to be performed by the four stages. (6)

Or

(b) (i) Use two input AND and OR gates to construct an $n \times n$ crossbar switch network between n processors and n memory modules. Let the width of each cross point be w bits in each direction. Prepare a schematic design of a typical cross point switch using C_{ij} as the enable signal for the switch in the ith row and jth column. Estimate the total number of AND and OR gates needed as a function of n and w. (10)

(ii) What is a virtual channel? How is it used to avoid deadlocks? (6)

14. (a) (i) Distinguish among High-end mainframes or near-supercomputers among the following vector processing machines in terms of architecture, performance range and cost —effectiveness. (8)

(ii) Consider a vector computer which can operate in one of two execution modes at a time: one is the vector mode with an execution rate of Rv = 10 M flops, and the other is the scalar mode with an execution rate of Rs = 1 Mflops. Let α be the percentage of code that is vectorizable in a typical program mix for this computer.

(1) Derive an expression for the average execution rate Ra for this computer.

(2) Plot Ra as a function of α in the range (0, 1)

(3) Determine the vectorization ratio α needed in order to achieve an average execution rate of Ra = 7.5 Mflops.

(4) Suppose Rs = 1 Mflops and α = 0.7 what value of Rv is *needed* to achieve Ra = 2 Mflops. (8)

Or

(b) (i) Distinguish between static dataflow computers and dynamic dataflow computers. Draw a dataflow graph showing the computations of the roots of a sequence of quadratic equations $A_iX_i^2 + B_iX_iC_i = 0$ for $i = 1,2,...,N$. (8)

(ii) Why are fine-grain processors chosen for future multiprocessors over medium - grain processors used in the past? From scalability point of view why is fine-grain parallelism more appealing than medium-grain or coarse-grain parallelism for building MPP systems. (8)

15. (a) (i) Based on data domain, algorithm used and flow of control in applications distinguish the opportunities for applying domain, control and object decomposition techniques in distributed computing on multicomputers. (8)

(ii) Choose an example program to demonstrate the concepts of macrotasking, microtasking and autotasking on a Cray-like multiprocessor supercomputer. Compare the performances of the three multitasking schemes based on the example program execution. (8)

Or

(b) (i) Explain the multiprocessor UNIX design goals in the areas of compatibility, portability, address space, load balancing, parallel I/O and network services. (8)

(ii) Explain the thread management problem and characterize the various thread scheduling policies with respect to Mach/OS kernel. (8)
