Roll No. | | | | | | | | |

B.E / B.Tech (Full Time) DEGREE END SEMESTER EXAMINATIONS, APRIL/MAY 2012

INFORMATION TECHNOLOGY

VII Semester

CS9043 MULTI-CORE PROGRAMMING

(Regulation 2008)

Time : 3 Hours | Answer ALL Questions | Max. Marks 100

PART-A (10 x 2 = 20 Marks)

1. State Gustafson's law.
2. How is zero-cycle context switch provided in hardware-supported multi-threaded processors ?
3. Identify four important activities in parallelizing an application.
4. Spin-locks are more suitable for multi-core systems than single-core systems. True or false. Justify your answer.
5. How is fine-grained locking useful ? Give an example for any data structure.
6. Why is barrier synchronization needed ?
7. What is the openMP pragma to be used if we want only one process to execute a particular block of code ?
8. Is it possible for an MPI program to deadlock ? Give an example or give reasons.
9. What is meant by a memory consistency model ?
10. What are the parameters to be specified for a Gather function in MPI ?

PART-B (5 X 16 = 80 Marks)

11. (i) Differentiate between the message-passing and shared memory parallel programming paradigms. (6)

    (ii) Discuss with an example the memory and cache related optimizations to be considered when writing programs for a multi-core processor. (10)

12. (a) (i) What is Amdahl's law ? Why is the speed-up estimate given by Amdahl's law considered to be pessimistic in the multi-core scenario ? (6)

    (ii) Explain the advantages and disadvantages of architectures designed for instruction level parallelism, and for thread-level parallelism. (6)

    (iii) Explain the architectural features of any multi-core processor. (4)

OR

(b) (i) Explain Flynn's classification. To which category would the operations in the following ctivity be mapped : your 4-year programme of study at our University. Account for the various activities. (6)

(ii) Consider a big data analysis program of which 90% can be parallelized. Can linear speed-up be achieved by increasing the number of processors ? Justify your answer. (4)

(iii) Differentiate software and hardware multithreading. (6)

13. (a) (i) How can deadlocks be avoided ? Can deadlocks occur when locks are used ? (8)

(ii) What is false sharing ? Assume that words W1, W2 and W3 are in the same cache block, which is in the shared state in the caches of the processors P1, P2 and P3. Assuming the following sequence of events, identify each miss as a true sharing miss, a false sharing miss or a hit : P1 writes W1, P2 writes W2, P1 writes W2, P3 writes W3, P3 writes W1. Explain your answer. (8)

OR

(b) (i) Give a code snippet that can cause a livelock. Indicate how the livelock occurs. Give a mechanism to avoid livelocks. (8)

(ii) Explain a scenario (give an example) where priority inversion occurs. Explain the two solutions to this problem with respect to this example. (8)

14. (a) (i) Write a parallel program using openMP to transpose an nxn matrix. Estimate the speed-up expected with k processors. (8)

(ii) Explain the different scheduling clauses supported in openMP. If 1000 iterations are to be run on 4 threads, calculate the partition sizes for the different cases. (8)

OR

(b) (i) Write a **cache-friendly** parallel code using openMP for matrix multiplication. Provide comments to explain your code. (8)

(ii) What is the effect of the different clauses used for sharing/moving data in OpenMP ? (8)

15. (a) (i) Write parallel code for any sort algorithm using MPI. (10)

(ii) What are the different collective communication functions in MPI ? Explain. (6)

OR

(b) (i) Write parallel code using MPI to perform a word search in a huge dictionary. Explain how you are partitioning the problem, and estimate the speed-up expected. (10)

(ii) Explain how reduction operation can be used. When a reduction operation is performed, what is the initial value to which the reduction variable will be set for (A) a multiply operation and (B) a MIN operation . (6)

****