School of Information Technology
IIT Kharagpur

Course Id: IT60101 Foundations of Computing Systems **End-Semester Examination**

Date: November 28, 2006 **Total Time: 3 Hours (Part A + Part B)**

Answer Part A and Part B in separate answer scripts

PART A

Answer any three questions

1. Answer the following questions in connection with demand paged memory management.

(1+4+4+1=10)

a. Justify or contradict the statement: *it is necessary to have base and limit registers to enforce good memory protection*.

b. A processor generates a 32-bit address, which needs to be mapped to a physical (main) memory of size 16 Gbytes (1G = 2^{30}). The page size is 8096 bytes. Draw a schematic diagram showing the logical to physical address translation. What will be the size of the page table?

c. Consider the page reference stream: 2 3 4 3 2 4 3 2 4 5 6 7 5 6 7 4 5 6 7 2 1. Given a page frame allocation of 3 and assuming the primary memory is initially unloaded, how many page faults will the given reference stream incur under (i) FIFO, (ii) LRU algorithms.

d. What is Belady's anomaly? What is the condition for it to not occur?

2. **(2+2+4+2=10)**

a. Suggest a solution to the mutual exclusion problem using test-and-set instruction.

b. Prove that the shortest job next scheduling algorithm generates the shortest possible schedule among all non-preemptive scheduling algorithms.

c. Suppose that the following processes need to be executed on a single processor.

Process no.	Arrival Time	Length of next CPU burst
P1	0	75
P2	10	40
P3	10	25
P4	80	20
P5	85	45

Assuming that the context-switch overhead is 5 time units, calculate the average waiting time for the processes using round robin scheduling with time quantum of 15 units.

d. Suggest a method of implementing semaphores on a shared-memory multiprocessor system.

3. **(3+3+1+1+2=10)**

a. What is an inode? Where do the inodes reside on a file system? How are the free inodes kept track of in the super block?
b. A Unix-like file system uses 2048 byte blocks, and 32-bit block addresses, and 64-bit file size. Estimate the maximum size of a file in such a system.
c. Why is an interrupt typically acknowledged at the end of an instruction cycle, while a DMA request acknowledged at the end of a machine cycle?
d. How is a system call different from calling a normal user-defined function?
e. Explain with the help of a diagram how DMA block transfer works.

4. **(2+4+4=10)**

a. A synchronous pipeline consists of k stages with stage delays $d_1,d_2,\ldots,d_k$ respectively, which includes the delays of the inter-stage latches. Estimate the total time taken to process n data items in the pipeline. Assume that there are no hazards.
b. Show a typical instruction level pipeline of a RISC processor. Hence illustrate the occurrences of structural and control hazards with the help of one example each.
c. A byte-addressable computer has a small data cache capable of holding eight 32-bit words. Each cache block consists of one 32-bit word. When a given program is executed, the processor reads data from the following sequence of hex addresses:

   200, 204, 208, 20C, 2F4, 2F0, 200, 204, 218, 21C, 24C, 2F4

   The pattern is repeated four times. Assuming that the cache is initially empty, show the contents of the cache at the end of each pass through this loop if we use a direct mapped cache.

PART B

Answer any three questions

1. The knapsack problem can be briefly stated as follows:
   A thief robbing a store finds n items: the i^{th} item is worth v_i dollars and weighs w_i pounds, where v_i and w_i are integers. He wants to take as valuable a load as possible, but can carry at most W pounds. Which items should he take? In the 0-1 version, the thief can either take or leave out any item but cannot take a fraction of it. On the other hand, in the fractional knapsack problem, the thief is allowed to take a fraction of the items.

   (a) Formulate a dynamic programming solution for the 0-1 knapsack problem. **[10]**

   OR

   (b) Write a greedy algorithm for the fractional knapsack problem. You need NOT justify why this would give an optimal solution. **[10]**

2. Consider the matrix chain multiplication A1A2A3A4 with the following dimensions:

Matrix	Dimension
A1	30X35
A2	35X15
A3	15X5
A4	5X10

   Using dynamic programming, derive an optimal parenthesization of the chain. **[10]**

3. Consider the following R-Tree. Assume M=5 and m =2

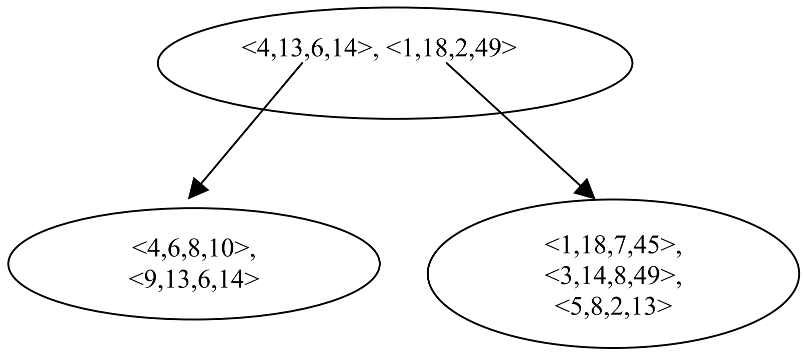

(a) If a new rectangle <3,5,6,8> comes, show the R-Tree after it is inserted. You must justify your choice of the place of insertion.

(b) After insertion of the above rectangle, determine the MinDist and MinMaxDist between the point (2,4) from each of the entries in the root node. **[3+(3+4)=10]**

4.

(a) Consider an undirected graph with the following adjacency list representation. The numbers in parentheses represent edge weights. Draw the graph and using Prim's algorithm show in a step by step manner, generation of its MST.

a → b(4) → h(8)
b → a(4) → h(11) → c(8)
c → b(8) → i(2) → d(7) → f(4)
d → c(7) → f(14) → e(9)
e → d(9) → f(10)
f → c(4) → d(14) → e(10) →g(2)
g → h(1) → i(6) → f(2)
h → a(8) → b(11) → i(7) →g(1)

(b) For the graph shown in the following figure, explain step-by-step operation of the Bellman-Ford algorithm. You need to show the distance of each node from the source as well as the parent of each node at the end of each iteration. What is the final output of the algorithm for this graph? S is the source node. You may name the other nodes in any way you want. What will be the output of Dijkstra's shortest path algorithm for this graph?

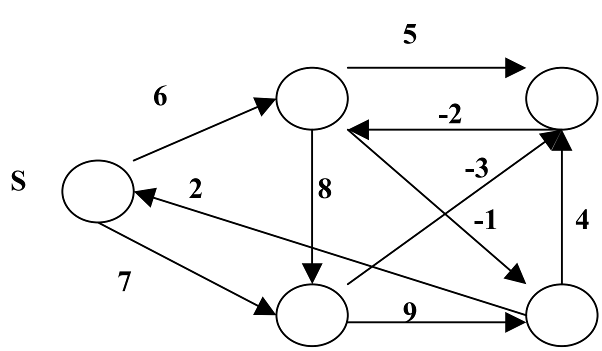

[4+(4+2)=10]