*Date* . . . . . . . . . . . . FN / AN *Time:* 2/3 Hrs. *Full Marks* . . . . . . . . .     *No. of Students* . . . . . . . . .
Autumn / Spring Semester, 20 . . . . . . . .       Deptt. . . . . . . . . . . . . . .       Sub No. . . . . . . . . . . . .
. . . . . . Yr. B. Tech.(Hons.) / B. Arch. / M. Sc.  Sub. Name . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .

**Instructions : Answer Q1. and Any Three (3) from the Remaining Questions (2 - 5)**

1.  (a)                                                                                                            $[\mathbf{8 \times \frac{1}{2} + 3 \times 2}]$

     i. Which of the following would support computation migration, i.e. the initiation of computation on a different node in the distributed system.
        a) Network Operating System, b) Distributed Operating System, c) Remote Procedure Calls, d) a and b, e) b and c, f) All of the above.

     ii. To provide a flexible structure, naming in a distributed file system should NOT be:
        a) Location Independent, b) Location Transparent, c) Bound at compile time, d) Implemented using a naming service.

     iii. It is difficult to coordinate distributed processes because:
        a) Semaphores are computationally expensive. b) Repeated message copying degrades performance. c) The combined state of the distributed system is very large. d) Messages transmission is slow compared to processors.

     iv. In general, problems in a distributed system may be solved in a centralized or fully distributed manner. The advantage of the fully distributed approach is:
        a) Improved performance. b) Improved robustness. c) Easier to implement. d) a & b. e) a & c.

     v. Naming Transparency means:
        I. location transparency: means that the path name gives no hint as to where the file is located.
        II. location independence: means that files can be moved without theiri names changes.
        a) I, b) II, c) I+II.

     vi. What are the benefits of stateful server over stateless server in a distributed file system?
        I. It provides fast service.
        II. It eliminates some network traffic.
        III. It can simply implement consistency and crash recovery mechanisms.
        a. I, b. I+II, c. I + III, d. I+II+III.

     vii. What are the disadvantages of stateful server over stateless server in distributed file system?
        I. slow service,
        II. complex design and implementation,
        III.complex consistency and crash recovery mechanisms.
        a. I, b. I + II, c. II + III, d. I + II + III,

Sig.of the Paper-Setter . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .

viii. What are the advantages of packet switching over circuit switching?
   I. It is faster.
   II. It can handle quality of service.
   III. It gives better utilization of bandwidth.
   IV. It is cheaper.
   a) I & III, b) II & IV, c) I & II & IV, d) III & IV,

ix. Describe Public Key Cryptography, Digital Signature, Pretty Good Privacy.

(b)                                                                              [**5 × 2**]

Give brief answers to the following questions.

   i. What is memory mapping of a file in a Linux environment?

   ii. Why is the management of a swap device different from the management of file system partition?

   iii. An 8-MByte memory is managed by a buddy system. The page frame size is 4-KByte. Possible buddy blocks are of sizes 1, 2, 4, $\cdots$, 1K page frames. Show the number of different blocks immediately after the initialization and after a contiguous allocation of 108-KByte memory.

   iv. In what situation does two entries of the global file table of a Unix system can point to the same inode buffer of the inode table? Is it possible that two entries of the file descriptor table of a process pointing to the same entry of the global file table?

   v. What are the different areas of a disk partition under the Unix file system? Give very brief descriptions of these areas.

2.                                                                              [**6 + 2 + 2**]

Compare the performance of CSCAN and SCAN scheduling assuming a uniform distribution of requests. Consider the average response time (the time between the arrival of a request and the completion of that request's service), the variation in response time and the effective bandwidth. How does performance depend on the relative sizes of seek time and rotational latency. What are RAID levels 0, 1, 4 1nd 5?

3.                                                                              [**10**]

Consider the following situation: One process A is running and is about to make a request for transfer of a large block of data from memory to disk while two other purely compute bound large jobs B and C are in the ready queue. No other user job is present. The DMA/bus/ interrupt controller is connected to both the CPU-Memory bus and a PCI bus to which the disk drive is connected. Clearly write down a possible sequence of events including Kernel calls if any at appropriate points till all jobs are over. Ignore cache and assume all data to be in the memory. Assume the following approximate time requirements data transfer 5 units, rest of A 2 units, B and C 10 units each RR scheduling time slice 2 units.

4. $[4 + 4 + 2]$

The following instruction writes an integer (10) in the virtual memory of a process. The address of the virtual memory location is the content of the register **ebp** + 2000.

**movl $10, 2000(%ebp)**

The data write meets a page fault in the user mode in a Linux on Pentium system. Take swap area, memory region, memory mapping into consideration and answer the following questions.

(a) Explain how the hardware-software combination detects the page fault. Assume two level paging and no segmentation.

(b) Explain clearly the different possible actions that the kernel may have to take to handle the page fault.

(c) What will be the difference in action if there is a page fault to fetch the instruction?

5. $[5 + 5]$

(a) The kernel searches for a disk block with the logical device and block numbers. The block may or may not be present in the buffer cache. Explain the different actions in the following cases.

    i. The block is in the queue and it is free.

    ii. The block is not in the queue but in the free list and marked delayed write.

    iii. The block is in the queue but is busy.

(b) A user process gives the following file system call.

    **fd = open("/home/arvind/data", O_RDONLY, 0) ;**

    i. Explain clearly the way the kernel, in a Unix like environment, parses the path name to get the inode of the file to open.

    ii. Describe the actions taken by the kernel to open a file after it gets the inode.

Sig.of the Paper-Setter . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .