

APJ ABDUL KALAM TECHNOLOGICAL UNIVERSITY

SECOND SEMESTER M.TECH DEGREE EXAMINATION, MAY 2016

Computer Science and Engineering

(Computer Science)

01CS6104 Operating System Design

Max. Marks : 60

Time: 3 Hours

Answer any **Two full questions** from each part

Part A

1.
 - a. How does the Linux process scheduler ensure a fair share of its computational power among the executing processes? What would have been the problem if a binary search tree was used instead of a Red Black tree as the runqueue? (3+2)
 - b. How is process context different from interrupt context? (2)
 - c. Can we use work queues in an Interrupt Context? Justify your answer. (2)
2.
 - a. By default the child process starts execution in Linux. What is the reason behind this? (2)
 - b. How is TASK_INTERRUPTIBLE state different from TASK_UNINTERRUPTIBLE state? Suppose a process is waiting for a specific event to happen. Which state should it be assigned? (2)
 - c. Why are interrupt handlers divided into Top and Bottom halves? Explain with an example. (3)
 - d. What are the differences between Tasklets and SoftIRQs? (2)
3.
 - a. What is the role of system calls in Linux? Explain with an example. (2)
 - b. What is a task list? When are individual members of a task list deallocated? (4)
 - c. When and how are SoftIRQs executed? (3)


Part B

4.
 - a. Explain true concurrency and pseudo concurrency with suitable examples. (6)
 - b. What is the relationship between Tick Rate, HZ and Jiffies? (4.5)

- 5.
- What trouble arises if a kernel code which has acquired a spin lock is interrupted by an Interrupt handler trying to acquire the same spin lock? Is it possible to avoid such a situation? How? (6)
 - What is the difference between `vmalloc()` and `kmalloc()`? Which of these must be used for allocating memory for hardware devices? Why? (4.5)
- 6.
- What is the purpose of using reader-writer locks? (3)
 - How is a mutex different from a semaphore? (3)
 - Explain the following. (2 + 2.5)
 - Per CPU Allocation.
 - Dentry object

Part C

- 7.
- What is the difference between a block device and a character device? Give examples. (2)
 - What is the role of an I/O scheduler? (2.5)
 - How can a client and server communicate with each other? What are the major issues, advantages and disadvantages associated with each of them? (6)
- 8.
- Before the 2.6 kernel, buffer head was used for block I/O whereas in the later kernels, it has been replaced by biostructure as the basic container for block I/O. Why? (4.5)
 - Represent a consistent and inconsistent global state for the following process/event graph. (3)



- Explain the distributed snapshot algorithm. (3)
- 9.
- What drawback of the deadline I/O scheduler is overcome by the Anticipatory I/O scheduler? (4.5)
 - Explain Lamport's timestamp algorithm with an example. (6)