

Dr. Hari Singh

JAYPEE UNIVERSITY OF INFORMATION TECHNOLOGY, WAKNAGHAT

TEST -1 EXAMINATION - February 2020

B.Tech 4th Semester

COURSE CODE: 18B11CI411

MAX. MARKS: 15

COURSE NAME: Operating Systems

COURSE CREDITS: 03

MAX. TIME: 1 Hrs

Note: All questions are compulsory. Carrying of mobile phone during examinations will be treated as case of unfair means.

SECTION A (3x2=6)

Q1. When an interrupt occurs, an operating system

- (A) ignores the interrupt
 - (B) always changes state of interrupted process to 'blocked' and schedules another process
 - (C) always resumes execution of interrupted process after processing the interrupt
 - (D) may change the state of interrupted process to 'blocked' and schedule another process
- Choose the correct choice and provide explanation.

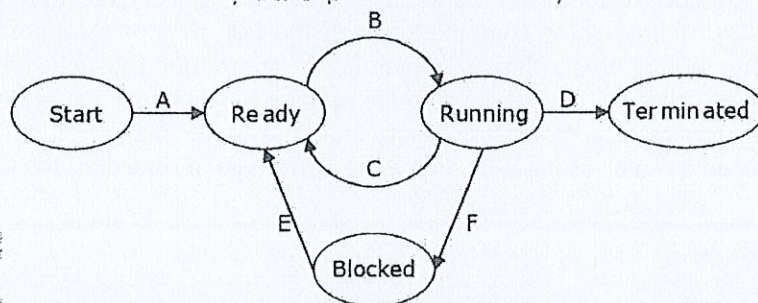
[2] marks

Q2. Using Amdahl's law, calculate the speedup gain of an application that has 60 percent parallel component for (a) two processing cores and (b) four processing cores.

[2] marks

Q3. In the following process state transition diagram for a uniprocessor system, assume that there are always some processes in the ready state. Now consider the following statements:

[2] marks



I. If a process makes a transition D, it would result in another process making transition A immediately.

II. A process P2 in blocked state can make transition E while another process P1 is in running state.

III. The OS uses preemptive scheduling.

IV. The OS uses non-preemptive scheduling.

Which of the above statement(s) are TRUE? One or more choices may be true. Provide explanation.

SECTION B

- Q4. The system generation (SYSGEN) process helps in configuring or generating system for each computer. Describe various system generation approaches adopted by the system administrators to generate systems according to the size and generality, and the ease of modifying it as the hardware configuration changes? [3] marks
- Q5. How many processes are generated by the following program code? Draw a tree showing the generation of parent-child processes. What is the output produced? Briefly explain your answer. [3] marks

```
#include <stdio.h>
#include <unistd.h>

int main()
{
    fork();
    fork() && fork() || fork();
    fork();

    printf("1\n");
    return 0;
}
```

- Q6. Consider three processes, all arriving at time zero, with total execution time of 10, 20 and 30 units, respectively. Each process spends the first 20% of execution time doing I/O, the next 70% of time doing computation, and the last 10% of time doing I/O again. The operating system uses a shortest remaining compute time first scheduling algorithm and schedules a new process either when the running process gets blocked on I/O or when the running process finishes its compute burst. Assume that all I/O operations can be overlapped as much as possible. For what percentage of time does the CPU remain idle? [3] marks