OPERATING SYSTEMS

(Concurrent Processing)

39. Which of the following scheduling algorithms may cause starvation?
    a. First-come-first-served
    b. Round Robin
    c. Priority
    d. Shortest process next
    e. Shortest remaining time first
    (1) a, c and e
    (2) c, d and e
    (3) b, d and e
    (4) b, c and d
    Answer: 2

36. Consider the reference string
    0 1 2 3 0 1 4 0 1 2 3 4
    If FIFO page replacement algorithm is used, then the number of page faults with three page frames and four page frames are ..... and ..... respectively.
    (A) 10, 9 (B) 9, 9
    (C) 10, 10 (D) 9, 10

39. Consider a system with twelve magnetic tape drives and three processes P_1, P_2 and P_3. Process P_1 requires maximum ten tape drives, process P_2 may need as many as four tape drives and P_3 may need upto nine tape drives. Suppose that at time t_1, process P_1 is holding five tape drives, process P_2 is holding two tape drives and process P_3 is holding three tape drives. At time t_1, system is in:
    (A) safe state (B) unsafe state
    (C) deadlocked state (D) starvation state
    Answer: B

50. Which of the following statements is not true for Multi Level Feedback Queue processor scheduling algorithm?
    (A) Queues have different priorities.
    (B) Each queue may have different scheduling algorithm
    (C) Processes are permanently assigned to a queue
    (D) This algorithm can be configured to match a specific system under design
    Answer: C

69. A system contains 10 units of resource of same type. The resource requirement and current allocation of these resources for three processes P, Q and R are as follows :

	P	Q	R
Maximum requirement	8	7	5
Current allocation	4	1	3

    Now, consider the following resource requests :
    (i) P makes a request for 2 resource units.
    (ii) Q makes request for 2 resources units.
    (iii) R makes a request of 2 resource units.

For a safe state, which of the following options must be satisfied?
(A) Only request (i)
(B) Only request (ii)
(C) Only request (iii)
(D) Request (i) and (ii)
Answer: C

70. Consider the following set of processes with the length of CPU burst time in milliseconds (ms) :

Process	A	B	C	D	E
Burst time	6	1	2	1	5
Priority	3	1	3	4	2

Assume that processes are stored in ready queue in following order :
A – B – C – D – E
Using round robin scheduling with time slice of 1 ms, the average turn around time is
(A) 8.4 ms
(B) 12.4 ms
(C) 9.2 ms
(D) 9.4 ms
Answer: A

2. The efficiency (E) and speed up (sp) for Multiprocessor with p processors satisfies :

(A) $E \leq p$ and $s_p \leq p$

(B) $E \leq 1$ and $s_p \leq p$

(C) $E \leq p$ and $s_p \leq 1$

(D) $E \leq 1$ and $s_p \leq 1$

Answer: B

36. There are three processes P1, P2 and P3 sharing a semaphore for synchronising a variable. Initial value of semaphore is one. Assume that negative value of semaphore tells us how many processes are waiting in queue. Processes access the semaphore in following order:
(a) P2 needs to access
(b) P1 needs to access
(c) P3 needs to access
(d) P2 exits critical section
(e) Pi exits critical section
The final value of semaphore will be:
(1) 0
(2) 1
(3) -1
(4) -2
Answer: 1

38. Consider a system having 'm' resources of the same type. These resources are shared by three processes P_1, P_2 and P_3 which have peak demands of 2, 5 and 7 resources respectively. For what value of 'm' deadlock will not occur?
(A) 70 (B) 14
(C) 13 (D) 7

Answer: B

39. Five jobs A, B, C, D and E are waiting in Ready Queue. Their expected runtimes are 9, 6, 3, 5 and x respectively. All jobs entered in Ready queue at time zero. They must run in ..... order to minimize average response time if $3 < x < 5$.

(A) B, A, D, E, C (B) C, E, D, B, A
(C) E, D, C, B, A (D) C, B, A, E, D

Answer: B

40. Consider three CPU intensive processes P1, P2, P3 which require 20, 10 and 30 units of time, arrive at times 1, 3 and 7 respectively. Suppose operating system is implementing Shortest Remaining Time first (pre-emptive scheduling) algorithm, then ..... context switches are required (suppose context switch at the beginning of Ready queue and at the end of Ready queue are not counted).

(A) 3 (B) 2
(C) 4 (D) 5

Answer: A

37. Suppose there are four processes in execution with 12 instances of a Resource R in a system. The maximum need of each process and current allocation are given below:

Process	**Max. Need**	**Current Allocation**
P_1	8	3
P_2	9	4
P_3	5	2
P_4	3	1

With reference to current allocation, is system safe? If so, what is the safe sequence?

(A) No (B) Yes, P_1 P_2 P_3 P_4
(C) Yes, P_4 P_3 P_1 P_2 (D) Yes, P_2 P_1 P_3 P_4

Answer: C

29. A system has four processes and five allocatable resources. The current allocation and maximum needs are as follows:

	Allocated	Maximum	Available
Process A	1 0 2 1 1	1 1 2 1 3	0 0 x 1 1
Process B	2 0 1 1 0	2 2 2 1 0	
Process C	1 1 0 1 0	2 1 3 1 0	
Process D	1 1 1 1 0	1 1 2 2 1	

The smallest value of x for which the above system in safe state is .....

(A) 1 (B) 3
(C) 2 (D) 0

Answer: Marks to all

Explanation:

If Process A's Maximum need is 1 1 2 1 2 instead of 1 1 2 1 3, then answer will be x=1

The needs matrix is as follows:

0 1 0 0 1
0 2 1 0 0
1 0 3 0 0
0 0 1 1 1

If x is 0, available vector will be 0 0 0 1 1, we have a deadlock immediately.

If x is 1, available vector will be 0 0 1 1 1, now, process D can run to completion. When it is finished, the available vector is 1 1 2 2 1.
Now A can run to complete, the available vector then becomes 2 1 4 3 2.
Then C can run and finish, return the available vector as 3 2 4 4 2.
Then B can run to complete. Safe sequence D A C B.

Answer: B

38. Which of the following conditions does not hold good for a solution to a critical section problem ?
(A) No assumptions may be made about speeds or the number of CPUs.
(B) No two processes may be simultaneously inside their critical sections.
(C) Processes running outside its critical section may block other processes.
(D) Processes do not wait forever to enter its critical section.
Answer: C

46. Consider a pre-emptive priority based scheduling algorithm based on dynamically changing priority. Larger priority number implies higher priority. When the process is waiting for CPU in the ready queue (but not yet started execution), its priority changes at a rate a = 2. When it starts running, its priority changes at a rate b = 1. All the processes are assigned priority value 0 when they enter ready queue. Assume that the following processes want to execute :

Process ID	**Arrival Time**	**Service Time**
P1	0	4
P2	1	1
P3	2	2
P4	3	1

The time quantum q = 1. When two processes want to join ready queue simultaneously, the process which has not executed recently is given priority. The finish time of processes P1, P2, P3 and P4 will respectively be
(A) 4, 5, 7 and 8
(B) 8, 2, 7 and 5
(C) 2, 5, 7 and 8
(D) 8, 2, 5 and 7
Answer: B
Explanation:

T	EP	Priority After T			
00	P1	P1[0]	P2[-]	P3[-]	P4[-]
01	P2	P1[1]	P2[0]	P3[-]	P4[-]
02	P1	P1[1]	P2[X]	P3[2]	P4[-] .....P2 completed at 2
03	P3	P1[2]	P2[X]	P3[2]	P4[2]
04	P4	P1[3]	P2[X]	P3[3]	P4[2]
05	P1	P1[3]	P2[X]	P3[4]	P4[X] .....P4 completed at 5
06	P3	P1[4]	P2[X]	P3[4]	P4[X]
07	P1	P1[4]	P2[X]	P3[X]	P4[X] .....P3 completed at 7
08	--	P1[X]	P2[X]	P3[X]	P4[X] .....P1 completed at 8

44. Match the following:

List – I	**List - II**
Process state transition	**Reason for transition**

a. Ready→Running — i. Request made by the process is satisfied or an event for which it was waiting occurs.
b. Blocked→Ready — ii. Process wishes to wait for some action by another process.
c. Running→Blocked — iii. The process is dispatched.
d Running→Ready — iv. The process is preempted.

Codes :

	a	b	c	d
(A)	iii	i	ii	iv
(B)	iv	i	iii	ii
(C)	iv	iii	i	ii
(D)	iii	iii	ii	i

Answer: A

46. A virtual memory based memory management algorithm partially swaps out a process. This is an example of
(A) short term scheduling (B) long term scheduling
(C) medium term scheduling (D) mutual exclusion
Answer: C

44. A set of processors P_1, P_2, ..., P_k can execute in parallel if Bernstein's conditions are satisfied on a pair wise basis; that is
$P_1 \| P_2 \| P_3 \| ..... \| P_k$ if and only if:
(A) $P_i \| P_i$ for all $i \neq j$
(B) $P_i \| P_j$ for all $i = j+1$
(C) $P_i \| P_j$ for all $i \leq j$
(D) $P_i \| P_j$ for all $i \geq j$
Answer: A
Explanation:
Bernstein's Condition:
1. If process Pi writes to a memory cell Mi, then no process Pj can read the cell Mi.
2. If process Pi read from a memory cell Mi, then no process Pj can write to the cell Mi.
3. If process Pi writes to a memory cell Mi, then no process Pj can write to the cell Mi.

14. Given memory partitions of 100 K, 500 K, 200 K, 300 K and 600 K (in order) and processes of 212

30. Which of the following memory allocation scheme suffers from external fragmentation ?
(A) Segmentation (B) Pure demand paging
(C) Swapping (D) Paging
Answer: A

38. A relationship between processes such that each has some part (critical section) which must not be executed while the critical section of another is being executed, is known as
(A) Semaphore (B) Mutual exclusion
(C) Multiprogramming (D) Message passing
Answer: B

39. How many states can a process be in ?
(A) 3 (B) 4
(C) 2 (D) 5

Answer: D

40. A Dead-lock in an Operating System is
(A) Desirable process
(B) Undesirable process
(C) Definite waiting process
(D) All of the above
Answer: C

36. Match the following:
Set-I
(a) Disk scheduling
(b) Batch processing
(c) Time sharing
(d) Interrupt processing
Set-II
1. Round-robin
2. SCAN
3. LIFO
4. FIFO
Codes :
(a) (b) (c) (d)
(A) 3 4 2 1
(B) 4 3 2 1
(C) 2 4 1 3
(D) 1 4 3 2
Answer: C

37. …….. synchronizes critical resources to prevent dead lock.
(A) P-operator
(B) V-operator
(C) Semaphore
(D) Swapping
Answer: C

38. …….. is one of pre-emptive scheduling algorithm.
(A) RR
(B) SSN
(C) SSF
(D) Priority based
Answer: A

39. In order to allow only one process to enter its critical section, binary semaphore are initialized to
(A) 0
(B) 1
(C) 2
(D) 3

36. In the process management Round-robin method is essentially the pre-emptive version of ……
(A) FILO

(B) FIFO
(C) SSF
(D) Longest time first
Answer: B

21. N processes are waiting for I/O. A process spends a fraction p of its time in I/O wait state. The CPU utilization is given by:
(A) $1-p^{-N}$
(B) $1-p^{N}$
(C) p^{N}
(D) p^{-N}
Answer: B
Explanation:
According to the probabilistic model, if a process spends a fraction p of its time in I/O wait state, and if there are N processes in memory, then the CPU utilization is given by $1-p^{N}$

22. If holes are half as large as processes, the fraction of memory wasted in holes is:
(A) 2/3
(B) 1/2
(C) 1/3
(D) 1/5
Answer: D

23. An example of a non pre-emptive scheduling algorithm is:
(A) Round Robin
(B) Priority Scheduling
(C) Shortest job first
(D) 2 level scheduling
Answer: C
Answer: B

23. If a process is under statistical control, then it is
(A) Maintainable (B) Measurable
(C) Predictable (D) Verifiable
Answer: C

29. Pre-emptive scheduling is the strategy of temporarily suspending a gunning process
(A) before the CPU time slice expires
(B) to allow starving processes to run
(C) when it requests I/O
(D) to avoid collision
Answer: A

30. In round robin CPU scheduling as time quantum is increased the average turn around time
(A) increases (B) decreases
(C) remains constant (D) varies irregularly
Answer: D

31. Resources are allocated to the process on non-sharable basis is
(A) mutual exclusion (B) hold and wait
(C) no pre-emption (D) circular wait
Answer: A

41. In which of the following, ready to execute processes must be present in RAM?

(A) multiprocessing
(B) multiprogramming
(C) multitasking
(D) all of the above
Answer: D

36. An example of a non-preemptive CPU scheduling algorithm is:
(A) Shortest job first scheduling.
(B) Round robin scheduling.
(C) Priority scheduling.
(D) Fair share scheduling.
Answer: A

37. There are 'n' processes in memory. A process spends a fraction 'p' of its time waiting for I/O to complete. The CPU utilization is given by:
(A) p^n
(B) $1-p^n$
(C) $(1-p)^n$
(D) 1-np
Answer: B

40. A computer has 6 tape drives with 'n' processes competing for them. Each process may need two drives. For which values of 'n' is the system deadlock free ?
(A) 1
(B) 2
(C) 3
(D) 6
Answer: C

36. How many states can a process be in?
(A) 2 (B) 3
(C) 4 (D) 5
Answer: D

37. In processor management, round robin method essentially uses the preemptive version of .....
(A) FILO
(B) FIFO
(C) SJF
(D) Longest time first
Answer: B

39. ..... synchronize critical resources to prevent deadlock.
(A) P-operator
(B) V-operator
(C) Semaphores
(D) Hard disk
Answer: C

40. The memory allocation scheme subjected to "external" fragmentation is:
(A) Segmentation
(B) Swapping
(C) Demand paging

(D) Multiple contiguous fixed partitions
Answer: A

40. Banker's algorithm is used for ..... purpose.
(A) Deadlock avoidance
(B) Deadlock removal
(C) Deadlock prevention
(D) Deadlock continuations
Answer: A

36. Moving Process from main memory to disk is called:
(A) Caching (B) Termination
(C) Swapping (D) Interruption
Answer: C

38. Bankers algorithm is for.
(A) Dead lock Prevention
(B) Dead lock Avoidance
(C) Dead lock Detection
(D) Dead lock creation
Answer: B

39. Which is the correct definition of a valid process transition in an operating system ?
(A) Wake up : Ready → Running
(B) Dispatch : Ready→ Running
(C) Block : Ready → Running
(D) Timer run out : Ready → Blocked
Answer: B

40. Match the following

(a) Disk scheduling	(1) Round robin
(b) Batch processing	(2) Scan
(c) Time sharing	(3) LIFO
(d) Interrupt processing	(4) FIFO

(A) a-3, b-4, c-2, d-1
(B) a-4, b-3, c-2, d-1
(C) a-2, b-4, c-1, d-3
(D) a-3, b-4, c-1, d-2
Answer: C

40. A scheduling Algorithm assigns priority proportional to the waiting time of a process. Every process starts with priority zero (lowest priority). The scheduler re-evaluates the process priority for every 'T' time units and decides next process to be scheduled. If the process have no I/O operations and all arrive at time zero, then the scheduler implements ..... criteria.
(A) Priority scheduling (B) Round Robin Scheduling
(C) Shortest Job First (D) FCFS
Answer: B

37. Let P_i and P_j be two processes, R be the set of variables read from memory, and W be the set of variables written to memory. For the concurrent execution of two processes P_i and P_j, which of the following conditions is not true?
(A) $R(P_i) \cap W(P_j) = \Phi$ (B) $W(P_i) \cap R(P_j) = \Phi$

(C) $R(P_i) \cap R(P_j) = \Phi$ (D) $W(P_i) \cap W(P_j) = \Phi$

Answer: C

60. Suppose S and Q are two semaphores initialized to 1. P1 and P2 are two processes which are sharing resources.

P1 has statements	**P2 has statements**
wait(S) ;	wait(Q) ;
wait(Q) ;	wait(S) ;
critical section1;	critical section 2;
signal(S) ;	signal(Q) ;
signal(Q) ;	signal(S) ;

Their execution may sometimes lead to an undesirable situation called

(A) Starvation

(B) Race condition

(C) Multithreading

(D) Deadlock

Answer: D

61. An operating system using banker's algorithm for deadlock avoidance has ten dedicated devices (of same type) and has three processes P1, P2 and P3 with maximum resource requirements of 4, 5 and 8 respectively. There are two states of allocation of devices as follows:

State 1	Processes	P1	P2	P3
Devices allocated		2	3	4
State 2	Processes	P1	P2	P3
Devices allocated		0	2	4

Which of the following is correct?

(A) State 1 is unsafe and state 2 is safe.

(B) State 1 is safe and state 2 is unsafe.

(C) Both, state 1 and state 2 are safe.

(D) Both, state 1 and state 2 are unsafe.

Answer: A

62. Let the time taken to switch between user mode and kernel mode of execution be T1 while time taken to switch between two user processes be T2. Which of the following is correct?

(A) T1 < T2

(B) T1 > T2

(C) T1 = T2

(D) Nothing can be said about the relation between T1 and T2.

Answer: A

60. Suppose S and Q are two semaphores initialized to 1. P1 and P2 are two processes which are sharing resources.

P1 has statements	**P2 has statements**
wait(S) ;	wait(Q) ;
wait(Q) ;	wait(S) ;
critical section1;	critical section 2;
signal(S) ;	signal(Q) ;
signal(Q) ;	signal(S) ;

Their execution may sometimes lead to an undesirable situation called
(A) Starvation
(B) Race condition
(C) Multithreading
(D) Deadlock
Answer: D

61. An operating system using banker's algorithm for deadlock avoidance has ten dedicated devices (of same type) and has three processes P1, P2 and P3 with maximum resource requirements of 4, 5 and 8 respectively. There are two states of allocation of devices as follows:

State 1	Processes	P1	P2	P3
	Devices allocated	2	3	4
State 2	Processes	P1	P2	P3
	Devices allocated	0	2	4

Which of the following is correct?
(A) State 1 is unsafe and state 2 is safe.
(B) State 1 is safe and state 2 is unsafe.
(C) Both, state 1 and state 2 are safe.
(D) Both, state 1 and state 2 are unsafe.
Answer: A

65. A thread is usually defined as a light weight process because an Operating System (OS) maintains smaller data structure for a thread than for a process. In relation to this, which of the following statement is correct?
(A) OS maintains only scheduling and accounting information for each thread.
(B) OS maintains only CPU registers for each thread.
(C) OS does not maintain a separate stack for each thread.
(D) OS does not maintain virtual memory state for each thread.
Answer: B

41. Consider a system having m resources of the same type. These resources are shared by 3 processes A, B and C which have peak demands of 3, 4 and 6 respectively. For what value of m deadlock will not occur?
(A) 7
(B) 9
(C) 10
(D) 13
Answer: D

32. A thread is a light weight process. In the above statement, weight refers to
(A) time
(B) number of resources

(C) speed
(D) All the above
Answer: B

43. Consider n processes sharing the CPU in round robin fashion. Assuming that each process switch takes s seconds. What must be the quantum size q such that the overhead resulting from process switching is minimized but, at the same time each process is guaranteed to get its turn at the CPU at least every t seconds?

(A) $q \leq \frac{t - ns}{n - 1}$

(B) $q \geq \frac{t - ns}{n - 1}$

(C) $q \leq \frac{t - ns}{n + 1}$

(D) $q \geq \frac{t - ns}{n + 1}$

Answer: B

13. Which of the following is scheme to deal with deadlock?
(A) Time out
(B) Time in
(C) Both (A) & (B)
(D) None of the above
Answer: A

65. Consider the methods used by processes P1 and P2 for accessing their critical sections. The initial values of shared Boolean variables S1 and S2 are randomly assigned,

P1	P2
while (S1 = = S2);	while (S1 = = S2);
critical section	critical section
S1 = S2;	S1 = S2;

Which one of the following statements describes the properties achieved?
(A) Mutual exclusion but not progress
(B) Progress but not mutual exclusion
(C) Neither mutual exclusion nor progress
(D) Both mutual exclusion and progress
Answer: C

58. Consider the following processes with time slice of 4 milliseconds (I/O requests are ignored):

Process	A	B	C	D
Arrival time	0	1	2	3
CPU cycle	8	4	9	5

The average turnaround time of these processes will be
(A) 19.25 milliseconds
(B) 18.25 milliseconds
(C) 19.5 milliseconds

(D) 18.5 milliseconds

Answer: B

34. Consider a uniprocessor system where new processes arrive at an average of five processes per minute and each process needs an average of 6 seconds of service time. What will be the CPU utilization?

(A) 80 %
(B) 50 %
(C) 60 %
(D) 30 %

Answer: B

52. An operating system has 13 tape drives. There are three processes P1, P2 & P3. Maximum requirement of P1 is 11 tape drives, P2 is 5 tape drives and P3 is 8 tape drives. Currently, P1 is allocated 6 tape drives, P2 is allocated 3 tape drives and P3 is allocated 2 tape drives. Which of the following sequences represent a safe state ?

(A) P2 P1 P3 (B) P2 P3 P1
(C) P1 P2 P3 (D) P1 P3 P2

Answer:

1. One of the disadvantages of user level threads compared to Kernel level threads is

(1) If a user level thread of a process executes a system call, all threads in that process are blocked.
(2) Scheduling is application dependent.
(3) Thread switching doesn't require kernel mode privileges.
(4) The library procedures invoked for thread management in user level threads are local procedures.

Answer: 1

2. Some of the criteria for calculation of priority of a process are:

a. Processor utilization by an individual process.
b. Weight assigned to a user or group of users.
c. Processor utilization by a user or group of processes.

In fair share scheduler, priority is calculated based on:

(1) only (a) and (b)
(2) only (a) and (c)
(3) (a), (b) and (c)
(4) only (b) and (c)

Answer: 3

3. Consider a system which have 'n' number of processes and 'm' number of resource types. The time complexity of the safety algorithm, which checks whether a system is in safe state or not, is of the order of:

(A) O(mn) (B) $O(m^2n^2)$
(C) $O(m^2n)$ (D) $O(mn^2)$

Answer: D

71. Simplest way of deadlock recovery is

(A) Roll back
(B) Preempt resource
(C) Lock one of the processes
(D) Kill one of the processes

Answer: D

Explanation:

Recovery through killing processes
– crudest but simplest way to break a deadlock
– kill one of the processes in the deadlock cycle
– the other processes get its resources
– choose process that can be rerun from the beginning. Kind of an extreme form of "rollback".

3. Match the following :

List – I

a. Critical region

b. Wait/signal

c. Working set

d. Dead lock

List – II

1. Hoares Monitor

2. Mutual exclusion

3. Principal of locality

4. Circular wait

Codes :

	a	b	c	d
(A)	2	1	3	4
(B)	1	2	4	3
(C)	2	3	1	4
(D)	1	3	2	4

Answer: A

40. Two atomic operations permissible on Semaphores are ..... and .....
(1) wait, stop
(2) wait, hold
(3) hold, signal
(4) wait, signal
Answer: 4

36. Consider a system with seven processes A through G and six resources R through W.
Resource ownership is as follows:
process A holds R and wants T
process B holds nothing but wants T
process C holds nothing but wants S
process D holds U and wants S & T
process E holds T and wants V

process F holds W and wants S
process G holds V and wants U
Is the system deadlocked ? If yes, ... processes are deadlocked.
(A) No (B) Yes, A, B, C
(C) Yes, D, E, G (D) Yes, A, B, F
Answer: C

33. Consider a system with five processes P_0 through P_4 and three resource types R_1, R_2 and R_3. Resource type R_1 has 10 instances, R_2 has 5 instances and R_3 has 7 instances. Suppose that at time T_0, the following snapshot of the system has been taken :
Assume that now the process P_1 requests one additional instance of type R_1 and two instances of resource type R_3. The state resulting after this allocation will be
(A) Ready State (B) Safe State
(C) Blocked State (D) Unsafe State
Answer: B

39. Capability Maturity Model is meant for:
(A) Product
(B) Process
(C) Product and Process
(D) None of the above
Answer: B

26. Dijkestra banking algorithm in an operating system, solves the problem of
(A) deadlock avoidance (B) deadlock recovery
(C) mutual exclusion (D) context switching
Answer: A

36. Part of a program where the shared memory is accessed and which should be executed indivisibly, is called:
(A) Semaphores
(B) Directory
(C) Critical Section
(D) Mutual exclusion
Answer: C

37. .....is one of pre-emptive scheduling algorithm.
(A) Shortest-Job-first (B) Round-robin
(C) Priority based (D) Shortest-Job-next
Answer: B

39. Non modifiable procedures are called
(A) Serially useable procedures
(B) Concurrent procedures
(C) Reentrant procedures
(D) Topdown procedures
Answer: C

36. Semaphores are used to:
(A) Synchronise critical resources to prevent deadlock
(B) Synchronise critical resources to prevent contention
(C) Do I/o
(D) Facilitate memory management

Answer: A

65. A thread is usually defined as a light weight process because an Operating System (OS) maintains smaller data structure for a thread than for a process. In relation to this, which of the following statement is correct?

(A) OS maintains only scheduling and accounting information for each thread.

(B) OS maintains only CPU registers for each thread.

(C) OS does not maintain a separate stack for each thread.

(D) OS does not maintain virtual memory state for each thread.

Answer: B

62. Let the time taken to switch between user mode and kernel mode of execution be T1 while time taken to switch between two user processes be T2. Which of the following is correct?

(A) T1 < T2
(B) T1 > T2
(C) T1 = T2
(D) Nothing can be said about the relation between T1 and T2.

Answer: A

48. Pipelining improves performance by:

(A) decreasing instruction latency
(B) eliminating data hazards
(C) exploiting instruction level parallelism
(D) decreasing the cache miss rate

Answer: C