Index- Operating System

Process management

A process has been allocated 3 page frames. Assume that none of the pages of the process are available in the memory initially. The process makes the following sequence of page references (reference string); 1,2,1,3,7,4,5,6,3,1

1) If optimal page replacement policy is used, how many page faults occur for the above reference string?

2 Marks GATE-CSE/IT-2007( )

[A] 7 [B] 8

[C]9 [D]10

Common Data for Q2 and Q3 is given below

Barrier is a synchronization construct where a set of processes synchronizes globally i.e., each process in the set arrives at the barrier and waits for all others to arrive and then all processes leave the barrier. Let the number of processes in the set be three and S be a binary semaphore with the usual P and V functions. Consider the following C implementation of a barrier with line numbers shown on the left.

```
Void barrier (void) {
1    :    P(S)
2    :    process _ arrived ++ ;
3    :    V(S);
4    :    while(process _ arrived ! = 3);
5    :    P(S);
6    :    process_ left ++;
7    :    if (process _ left = =3)
8    :    process _ arrived = 0;
9    :    process _ left = 0;
10   :    }
11   :    V(S);
}
```

The variables process _ arrived and process _ left are shared among all processes and are initialized to zero. In a concurrent program all the three processes call the barrier function when they need to synchronize globally.

2) 24. The above implementation of barrier is incorrect. Which one of the following is true ?

2 Marks GATE-CSE/IT-2006( )

[A] The barrier implementation is wrong due to the use of binary semaphore S

[B] The barrier implementation may lead to a deadlock if two barrier invocations are used in immediate succession

[C]Lines 6 to 10 need not be inside a critical section

[D]The barrier implementation is correct if there are only two processes instead of three

3) Which one of the following rectifies the problem in the implementation ?

2 Marks GATE-CSE/IT-2006( )

[A] Lines 6 to 10 are simply replaced by process _ arrived

[B] At the beginning of the barrier the first process to enter the barrier waits until process_ arrived becomes zero before proceeding to execute P (S)

[C]Context switch is disabled at the beginning of the barrier and re-enabled at the end

[D]The variable process_ left is made private instead of shared

4) What is the maximum number of reduce moves that can be taken by a bottom- up parser for a grammar with no epsilon- and unit-production (i.e., of type $A \rightarrow \in$ and $A \rightarrow a$) to parse a string with n tokens?

1 Marks GATE-CSE/IT-2013( )

[A] n/2 [B] n-1

[C]2n-1 [D]2^n

5) Which one of the following is the tightest upper bound that represents the time complexity of inserting an object into a binary search tree of n nodes?

1 Marks GATE-CSE/IT-2013( )

[A] O(1) [B] O(log n)

[C]O(n) [D]O(n log n)

6) Which one of the following is the tightest upper bound that represents the number of swaps required to sort n numbers using selection sort?

1 Marks GATE-CSE/IT-2013( )

[A] O(log n) [B] O(n)
[C] O(n log n) [D] $O(n^2)$

7) A shared variable x, initialized to zero, is operated on by four concurrent processes W, X, Y, Z as follows. Each of the processes W and X reads x from memory, increments by one, stores it to memory, and then terminates. Each of the processes Y and Z reads x from memory, decrements by two, stores it to memory, and then terminates. Each process before reading x invokes the P operation (i.e., wait) on a counting semaphore S and invokes the V operation (i.e., signal) on the semaphore S after storing x to memory. Semaphore S is initialized to two. What is the maximum possible value of x after all processes complete execution?

2 Marks GATE-CSE/IT-2013( )

[A] -2 [B] -1
[C] 1 [D] 2

8) Which combination of the following features will suffice to characterize an OS as a multi-programmed OS?
(A) More than one program may be loaded into main memory at the same time for execution. (B) If a program waits for certain events such as I/O, another program is immediately scheduled for execution. (C) If the execution of a program terminates, another program is immediately scheduled for execution.

2 Marks GATE-CSE/IT-2002( )

[A] A [B] A and B
[C] A and C [D] A, B and C

9) Consider a set of n tasks with known runtimes r_1, r_2 .... r_n to be run on a uniprocessor machine. Which of the following processor scheduling algorithms will result in the maximum throughput?

1 Marks GATE-CSE/IT-2001( )

[A] Round-Robin [B] Shortest-Job-First
[C] Highest-Response-Ratio-Next [D] First-Come-First-Served

10) Which of the following requires a device driver?

1 Marks GATE-CSE/IT-2001( )

[A] Register [B] Cache
[C] Main memory [D] Disk

11) Which of the following does not interrupt a running process?

2 Marks GATE-CSE/IT-2001( )

[A] A device [B] Timer
[C] Scheduler process [D] Power failure

12) Consider Peterson's algorithm for mutual exclusion between two concurrent processes i and j. The program executed by process is shown below.

repeat
flag[i]=true; turn=j;
while (P) do no-op;
Enter critical section, perform actions, then exit critical section
Flag[i]=false;
Perform other non-critical section actions. Until false;

For the program to guarantee mutual exclusion, the predicate P in the while loop should be

2 Marks GATE-CSE/IT-2001( )

[A] flag[j]=true and turn=i [B] flag[j]=true and turn=j
[C] flag[i]=true and turn=j [D] flag[i]=true and turn=i

13) A system uses FIFO policy for page replacement. It has 4 page frames with no pages loaded to begin with. The system first accesses 100 distinct pages in some order and then accesses the same 100 pages but now in the reverse order. How many page faults will occur?

1 Marks GATE-CSE/IT-2010( )

[A] 196 [B] 192
[C] 197 [D] 195

14) A multi-user, multi-processing operating system cannot be implemented on hardware that does not support

2 Marks GATE-CSE/IT-1999( )

[A] Address translation [B] DMA for disk transfer

[C] At least two modes of CPU execution (privileged and non-privileged)
[D] Demand paging

15) Which of the following actions is/are typically not performed by the operating system when switching context from process A to process B?

2 Marks GATE-CSE/IT-1999( )

[A] Saving current register values and restoring saved register values for process B.
[B] Changing address translation tables.
[C] Swapping out the memory image of process A to the disk.
[D] Invalidating the translation look-aside buffer.

16) Raid configurations of the disks are used to provide

2 Marks GATE-CSE/IT-1999( )

[A] Fault-tolerance
[B] High speed
[C] high data density
[D] None of the above

17) In a resident – OS computer, which of the following systems must reside in the main memory under all situations?

1 Marks GATE-CSE/IT-1998( )

[A] Assembler
[B] Linker
[C] Loader
[D] Compiler

18) I/O redirection

1 Marks GATE-CSE/IT-1997( )

[A] implies changing the name of a file
[B] can be employed to use an existing file as input file for a program
[C] implies connection 2 programs through a pipe
[D] None of the above

19) An operating system contains 3 user processes each requiring 2 units of resource R. the minimum number of units of r such that no deadlocks will ever arise is

2 Marks GATE-CSE/IT-1997( )

[A] 3
[B] 5
[C] 4
[D] 6

20) Each process Pi,i=1……9 is coded as follows
repeat P(mutex)
{critical section}
v(mutex)
forever
The code foe P10 is identical except that it uses v(mutex) in place of P(mutex).
What is the largest number of processes that can be inside the critical section at any moment?

2 Marks GATE-CSE/IT-1997( )

[A] 1
[B] 2
[C] 3
[D] None of the above

21) The process state transition diagram in Fig.1.8 is representative of

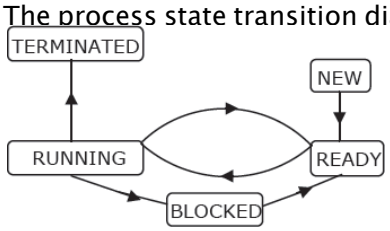

1 Marks GATE-CSE/IT-1996( )

[A] a batch operating system
[B] an operating system with a preemptive scheduler
[C] an operating system with a non-preemptive scheduler
[D] a uni-programmed operating system.

22) For the daisy chain scheme of connecting I/O devices, which of the following statements is true?

1 Marks GATE-CSE/IT-1996( )

[A] It gives non-uniform priority to various devices.
[B] It gives uniform priority to all devices.
[C] It is only useful for connecting slow devices to a processor device.
[D] It requires a separate interrupt pin on the processor for each device.

23) 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?

2 Marks GATE-CSE/IT-1993( )

[A] 7
[B] 9
[C] 10
[D] 13
[E] 15

24) Consider three CPU-intensive processes, whichrequire 10,20 and 30 time units and arrive at times 0,2, and 6,respectively. How many context switchesare needed if the operating system implements a shortest remaining time firstscheduling algorithm ? Do not count the context switches at time zero and atthe end.

1 Marks GATE-CSE/IT-2006( )

[A] 1
[B] 2
[C] 3
[D] 4

25) The atomic feth-and-set x,y instruction unconditionallysets the memory location x to 1 and fetches the old value of x in y without allowing any intervening access to the memory location X. Consider the following implementation of P and Vfunctions on a binary semaphore S.

```
void p (binary_ semaphore *S){
    unsignedy;
unsigned *x = &(S-> value);}
      do {
     fetch- and - set x,y;
    } while(y);
}
void V(binary_semphore *S) {
       {S-> value = 0 ;
}
```

Which one of the following is true ?

2 Marks GATE-CSE/IT-2006( )

[A] The implementation may not work if context switching is disabled in P
[B] Instead of using fetch-and-set, a pair of normal load/store can be used
[C] The implementation of V is wrong
[D] The code does not implement a binary semaphore

26) A process executes the following code for (i=0' I <n ; i++) fork (); The total number of child processes created is

2 Marks GATE-CSE/IT-2008( )

[A] n
[B] $2^n - 1$
[C] 2^n
[D] $2^{n+1} - 1$

Process management

Key Paper

1.	**A**	**2.**	**B**	**3.**	**B**	**4.**	**C**	**5.**	**C**
6.	**B**	**7.**	**D**	**8.**	**D**	**9.**	**B**	**10.**	**B**
11.	**B**	**12.**	**B**	**13.**	**A**	**14.**	**D**	**15.**	**D**
16.	**A**	**17.**	**C**	**18.**	**B**	**19.**	**C**	**20.**	**C**
21.	**B**	**22.**	**A**	**23.**	**A**	**24.**	**D**	**25.**	**A**
26.	**B**								

CPU Scheduling

1)Consider n processes sharing the CPU in a 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?

2 Marks GATE-CSE/IT-1998( )

[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] $q \geq \frac{t-ns}{n+1}$

2) The correct matching for the following pairs is:

(A) Disk scheduling	(1) Round robin
(B) Batch processing	(2) SCAN
(C) Time sharing	(3) LIFO
(D) Interrupt processing	(4) FIFO

1 Marks GATE-CSE/IT-1997( )

[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 – 3 D –2

3) When an interrupt occurs, an operating system

1 Marks GATE-CSE/IT-1997( )

[A] ignores the interrupt

[B] always changes state of interrupted process after processing the interrupt

[C]always resumes execution of interrupted process after processing the interrupt

[D] may change state of interrupted process to 'blocked' and schedule another process

4) Four jobs to be executed on a single processor system arrive at time 0+ in the order A, B, C, D. their burst CPU time requirements are 4, 1, 8, 1 time units respectively. The completion time of A under round robin scheduling with time slice of one time unit is

2 Marks GATE-CSE/IT-1996( )

[A] 10 [B] 4

[C]8 [D]9

5) Which scheduling policy is most suitable for a time shared operating system?

1 Marks GATE-CSE/IT-1995( )

[A]ShortestJobFirst [B] RoundRobin

[C]FirstComeFirstServe [D]Elevator

6) The sequence ..... is an optimal non-preemptive scheduling sequence for the following jobs which leaves the CPU idle for.....unit(s) of time.

Job	Arrival time	Burst time
1	0.0	9
2	0.6	5
3	1.0	1

2 Marks GATE-CSE/IT-1995( )

[A]{3, 2, 1}, 1 [B]{2, 1, 3} , 0

[C]{3, 2, 1} , 0 [D]{1, 2, 3},5

7) Assume that the following jobs are to be executed on a single processor system

Job Id	CPU Burst time
p	4
q	1
r	8
s	1
t	2

The jobs are assumed to have arrived at time 0_+ and in the order p, q, r, s, t. calculate the departure time (completion time) for job p if scheduling is round robin with time slice 1.

2 Marks GATE-CSE/IT-1993( )

[A]4 [B]10

[C]11 [D]12

[E] none of the above

8) Consider the following statements with respect touser- level threads and kernel-supported thread
(i) Context switch is faster with kernel-supported threads
(ii) For user-level threads, a system call can block the entire process
(c) Kernel-supported threads can be scheduled in dependently
(iv) User-level threads are transparent to the kernel
Which of the above statements are true ?

1 Marks GATE-CSE/IT-2004( )

[A](ii), (iii) and (iv) only [B] (ii) and (iii) only

[C](i) and (iii) only [D](i) and (ii) only

9) Consider the following set of processes, with the arrival times and the CPU-burst times given in milliseconds

Process	Arrivaltime	Burst time
P1	0	5
P2	1	3
P3	2	3
P4	4	1

What is the average turnaround time for these processes with the preemptive shortest remaining processing time first (SRPT) algorithm ?

2 Marks GATE-CSE/IT-2004( )

[A]5.50 [B] 5.75

[C]6.00 [D]6.2

10) Consider three processes (process id 0,1,2,respectively) with compute time bursts 2,4,and 8 time units. All processes arrive at time zeo. Consider the longest remaining time first(LRTF) scheduling algorithm. In LRTFties are broken by giving priority to the process with the lowest processid. The average turn around time is

2 Marks GATE-CSE/IT-2006( )

[A]13 units [B]14 units

[C]15 units [D]16 units

11) 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% ofexecution time doing I/O, the next 70%of time doing computation, and the last 10% of time doing I/O again. Theoperating system uses a shortestremaining compute time first schedulingalgorithm and schedules a new processeither when the running process getblocked on I/O or when the running process finishes its compute burst. Assume that all I/O operations can beoverlapped as much as possible. For whatpercentage of time does the CPU remainidle ?

2 Marks GATE-CSE/IT-2005( )

[A] 0% [B] 10.6%

[C]30.0% [D]89.4%

12) An operating system uses Shortest Remaining Time first (SRT) process scheduling algorithm. Consider the arrival times and execution times for the following processes.

Process	Execution time	Arrival time
P1	20	0
P2	25	15
P3	10	30
P4	15	45

What is the total waiting time for process P2 ?

2 Marks GATE-CSE/IT-2007( )

[A]5 [B]15

[C]40 [D]55

13) A virtual memory system uses first In first Out (FIFO) page replacement policyand allocates a fixed number of frames to a process. Consider the following statements :
P: Increasing the number of page frames allocated to a process sometimes increases the page faultrate.
Q: Some program do not exhibit locality of reference.
Which one of the following is TRUE ?

1 Marks GATE-CSE/IT-2007( )

[A] Both P and Q are true, and Q is the reason for P [B] Both P and Q are true, but Q is not the reason for P

[C] P is false, but Q is true [D] Both P and Q are false

14) A single processor system has three resource typesX,Y, and Z, which are shared by three processes. There are 5 units of each resource type allocated to each process, and the columnrequest denotes the number of units of each resource type requestedby a process in order to complete execution. Which of these processes will finishLAST?

	alloc	request
	X X X	X Y Z
P0	1 2 1	1 0 3
P1	2 0 1	0 1 2
P2	2 2 1	1 2 0

2 Marks GATE-CSE/IT-2007( )

[A] P0 [B] P1

[C] P2 [D] None of the above, since the system is in a deadlock

15) Two processes,P1 and P2, need to access a criticalsection of code. Consider the following synchronization construct used by theprocesses :

```
/* P1 */                          /*P2 */
 while (true)                     while (true)
 {                                {
  wants1 =true;                     wants2 = true;
  while(wants2==true);                while (wants1 ==true);
  /*Critical                      /*Critical
   Section */                       Section*/
   Wants 1 = false;                  wants2 = false;
   }                               }
 /* Re mainder section */              /* Re mainder sec tion */
```

Here, wants 1 and wants 2 are shared variables,Which are intitilized to false. Whichone of the following statements is TRUE about the above construct ?

2 Marks GATE-CSE/IT-2007( )

[A] It does not ensure mutual exclusion. [B] It does not ensure bounded waiting

[C] It requires that processes enter the critical section in strict alternation [D] It does not prevent deadlocks, but ensures mutual exclusion.

16) In which one of the following page replacement policies, Belady's anomaly may occur ?

1 Marks GATE-CSE/IT-2008( )

[A] FIFO [B] Optimal

[C] LRU [D] MRU

17) In the following process state transition diagramfor a uniprocessor system, assume that there are always some processes in theready state:
Now consider the following statements:
I.If a process makes a transition D, it would resultin another process making transition A immediately
II. A process P2 in blocked state can make transition E while anotherprocess P1 is in running state
III. The OS uses preemptive scheduling
IV. The OS uses non-preemptive scheduling
Which of the above statements are TRUE ?

2 Marks GATE-CSE/IT-2009( )

[A] I and II [B] I and III

[C] II and III [D] II and IV

18) 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 get 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 ?

2 Marks GATE-CSE/IT-2006( )

[A] 0% [B] 10.6%

[C]30.0% [D]89.4%

19) Least Recently Used (LRU) page replacement policy is a practical approximation to optimal page replacement. For the above reference string, how many more page faults occur with LRU than with the optimal page replacement policy?

2 Marks GATE-CSE/IT-2007( )

[A] 0 [B] 1

[C]2 [D]3

20) A scheduling algorithm assigns priority proportional to the waiting time of a process. Every process starts with priority zero(the lowest priority). The scheduler re-evaluates the process priorities every T time units and decides the next process to schedule. Which one of the following is TRUE if the processes have no I/O operations and all arrive at time zero?

1 Marks GATE-CSE/IT-2013( )

[A] This algorithm is equivalent to the first-come-first-serve algorithm.

[B] This algorithm is equivalent to the round-robin algorithm.

[C]This algorithm is equivalent to the shortest-job-first algorithm.

[D]This algorithm is equivalent to the shortest-remaining-time-first algorithm

21) Consider the following operation along with Enqueue and Dequeue operations on queues, where k is a global parameter

```
MultiDequeue (Q) {
m = k
while (Q is not empty) and (m > 0) {
Dequeue (Q)
m = m − 1
}
}
```

What is the worst case time complexity of a sequence of n queue operations on an initially empty queue?

2 Marks GATE-CSE/IT-2013( )

[A] $\Theta(n)$ [B] $\Theta(n+k)$

[C] $\Theta(nk)$ [D] $\Theta(n^2)$

22) Consider the following statements with respect to user-level threads and kernel-supported threads

(i) Context switch is faster with kernel-supported threads

(ii) For user-level threads, a system call can block the entire process

(c) Kernel-supported threads can be scheduled independently

(iv) User-level threads are transparent to the kernel

Which of the above statements are true ?

1 Marks GATE-CSE/IT-2004( )

[A] (ii), (iii) and (iv) only [B] (ii) and (iii) only

[C] (i) and (iii) only [D] (i) and (ii) only

CPU Scheduling

23) Consider the following set of processes, with the arrival times and the CPU-burst times given in milliseconds

Process	Arrival time	Burst time
P1	0	5
P2	1	3
P3	2	3
P4	4	1

What is the average turnaround time for these processes with the preemptive shortest remaining processing time first (SRPT) algorithm ?

2 Marks GATE-CSE/IT-2004( )

[A] 5.50 [B] 5.75

[C] 6.00 [D] 6.2

24) Consider a system with a two-level paging scheme in which a regular memory access takes 150 nanoseconds, and servicing a page fault takes 8 milliseconds. An average instruction takes 100 nano seconds of CPU time, and two memory accesses. The TLB hit ratio is 99%, and the page fault rate is one in every 10,000 instructions. What is the effective average instruction execution time ?

2 Marks GATE-CSE/IT-2004( )

[A] 645 nanoseconds [B] 1050 nanoseconds

[C] 1215 nanoseconds [D] 1230 nanoseconds

25) Consider three CPU-intensive processes, which require 10,20 and 30 time units and arrive at times 0,2, and 6, respectively. How many context switches are needed if the operating system implements a shortest remaining time first scheduling algorithm ? Do not count the context switches at time zero and at the end.

1 Marks GATE-CSE/IT-2006( )

[A] 1 [B] 2

[C] 3 [D] 4

26) Let W(n) and A(n) denote respectively, the worst case and average case running time of an algorithm executed on an input of size n. Which of the following is ALWAYS TRUE?

1 Marks GATE-CSE/IT-2012( )

[A] $A(n) = \Omega(W(n))$ [B] $A(n) = \Theta(W(n))$

[C] $A(n) = O(W(n))$ [D] $A(n) = o(W(n))$

27) A thread is usually defined as a 'light weight process' because an operating system (OS) maintains smaller data structures for a thread than for a process. In relation to this, which of the followings is TRUE?

1 Marks GATE-CSE/IT-2011( )

[A] On per-thread basis, the OS maintains only CPU register state

[B] The OS does not maintain a separate stack for each thread

[C] On per-thread basis, the OS does not maintain virtual memory state

[D] On per thread basis, the OS maintains only scheduling and accounting information

28) Consider the following table of arrival time and burst time for three processes P0, P1 and P2.

Process	Arrival time	Burst Time
P0	0 ms	9 ms
P1	1 ms	4ms
P2	2 ms	9ms

The pre-emptive shortest job first scheduling algorithm is used. Scheduling is carried out only at arrival or completion of processes. What is the average waiting time for the three processes?

2 Marks GATE-CSE/IT-2011( )

[A] 5.0 ms [B] 4.33 ms
[C] 6.33 ms [D] 7.33 ms

29) Which of the following statements are true?
I. Shortest remaining time first scheduling may cause starvation
II. Preemptive scheduling may cause starvation
III. Round robin is better than FCFS in terms of response time

1 Marks GATE-CSE/IT-2010( )

[A] I only [B] I and III only
[C] II and III only [D] I, II and III

30) Consider the following schedule for transactions T1, T2 and T3:

T1	T2	T3
Read(X)		
	Read(Y)	
		Read(Y)
	Write(Y)	
Write(X)		
		Write(X)
	Read(X)	
	Write(X)	

Which one of the schedules below is the correct serialization of the above?

2 Marks GATE-CSE/IT-2010( )

[A] T1→ T3→T2 [B] T2→ T1→T3
[C] T2→T3→T1 [D] T3→ T1→T2

CPU Scheduling

Key Paper

1.	**D**	**2.**	**C**	**3.**	**D**	**4.**	**D**	**5.**	**B**
6.	**A**	**7.**	**C**	**8.**	**A**	**9.**	**A**	**10.**	**A**
11.	**B**	**12.**	**B**	**13.**	**B**	**14.**	**C**	**15.**	**D**
16.	**A**	**17.**	**C**	**18.**	**B**	**19.**	**C**	**20.**	**B**
21.	**C**	**22.**	**B**	**23.**	**A**	**24.**	**D**	**25.**	**B**
26.	**C**	**27.**	**A**	**28.**	**A**	**29.**	**D**	**30.**	**A**

Inter-process communication

1) Formatting for a floppy disk refers to

2 Marks GATE-CSE/IT-1998( )

[A] arranging the data on the disk in contiguous fashion

[B] writing the directory

[C] erasing the system area

[D] writing identification information on all tracks and sectors

2) A linker is given object modules for a set of programs that were compiled separately. What information need to be included in an object module?

1 Marks GATE-CSE/IT-1995( )

[A] Object code

[B] Relocation bits

[C] Names and locations of all external symbols defined in the object module

[D] Absolute addresses of internal symbols

3) Which of the following system calls results in the sending of SYN packets?

1 Marks GATE-CSE/IT-2008( )

[A] socket

[B] bind

[C] listen

[D] connect

4) Which of the following statements about synchronous and asynchronous I/O is NOT true?

2 Marks GATE-CSE/IT-2008( )

[A] An ISR is invoked on completion of I/O in synchronous I/O but not in asynchronous I/O

[B] In both synchronous and asynchronous I/O an ISR (Interrupt Service Routine) is invoked after completion of the I/O

[C] A process making a synchronous I/O call waits until I/O is complete, but a process making an a synchronous I/O call does not wait for completion of the I/O

[D] In the case of synchronous I/O, the process waiting for the completion of I/O is woken up by the ISR that is invoked after the completion of I/O.

5) Register renaming is done is pipelined processors

1 Marks GATE-CSE/IT-2012( )

[A] as an alternative to register allocation at compile time

[B] for efficient access to function parameters and local variables

[C] to handle certain kinds of hazards

[D] as part of address translation

6) A computer handles several interrupt sources of which the following are relevant for this question.
Interrupt from CPU temperature sensor
Interrupt from Mouse
Interrupt from Keyboard
Interrupt from Hard Disk

1 Marks GATE-CSE/IT-2011( )

[A] Interrupt from Hard Disk

[B] Interrupt from Mouse

[C] Interrupt from Keyboard

[D] Interrupt from CPU temp sensor

Inter-process communication

Key Paper

1.	**D**	**2.**	**D**	**3.**	**D**	**4.**	**A**	**5.**	**C**
6.	**D**								

Deadlocks

1) Let m[0]…m[4] be mutexes (binary semaphores) and P[0] …. P[4] be processes.
Suppose each process P[i] executes the following:
wait (m[i];wait (m[(i+1) mode 4]);
……..
release (m[i]); release (m[(i+1)mod 4]);
This could cause

1 Marks GATE-CSE/IT-2000( )

[A] Thrashing
[B] Deadlock
[C]Starvation, but not deadlock
[D]None of the above

2) Which of the following is NOT a valid deadlock prevention scheme?

2 Marks GATE-CSE/IT-2000( )

[A] Release all resources before requesting a new resource
[B] Number the resources uniquely and never request a lower numbered resource than the last one requested.
[C]Never request a resource after releasing any resource
[D]Request and all required resources be allocated before execution.

3) A computer has six tape drives, with n processes competing for them. Each process may need two drives. What is the maximum value of n for the system to be deadlock free?

1 Marks GATE-CSE/IT-1998( )

[A] 6
[B] 5
[C]4
[D]3

4) A solution to the Dining Philosophers Problem which avoids deadlock is

2 Marks GATE-CSE/IT-1996( )

[A] ensure that all philosophers pick up the left fork before the right fork
[B] ensure that all philosophers pick up the right fork before the left fork
[C]ensure that one particular philosopher picks up the left fork before the right fork, and that all other philosophers pick up the right fork before the left fork
[D]None of the above

5) Consider two processes P1 and P2 accessing theshared variables X and Y protected by two binary semaphores Sx and Syrespectively , both initialized to 1. P and V denote the usual semaphoreoperators, where P decrements the semaphore value, and V increments the

```
P1:                          P2:
whiletrue do{                whiletrue do {
  L1:.....                     L3:.....
  L2:.....                     L4:.....
  X =X + 1;                    Y=Y + 1;
  Y =Y - 1;                    X= Y -1,
  V(SX);                       V(SY);
  V(SY);   }                   V(SX);      }
```

In order toavoid deadlock, the correct operators at L1,L2, L3 and L4 are respectively

2 Marks GATE-CSE/IT-2004( )

[A] P(SY), PaSX); P(SX),P(SY)
[B] P(SX), P(SY); P(SY), P(SX)
[C]P(SX), P(SX); P(SY), P(SY)
[D]P(SX), P(SY); P(SX), P(SY)

6) Suppose n processes,P1,…, Pnshare m identical resource units, whichcan be reserved and released one at atime. The maximum resource requirement of process piis sp where si >0. Which one of the following is a sufficient condition for ensuring thatdeadlock does not occur ?

2 Marks GATE-CSE/IT-2005( )

[A] $\forall i, S_i < m$
[B] $\forall i, S_i < n$
[C] $\sum_{i=1}^{n} S_i < (m+n)$
[D] $\sum_{i=1}^{n} S_i < (m*n)$

7) Which of the following is NOT true of deadlock prevention and deadlock avoidance schemes ?

2 Marks GATE-CSE/IT-2008( )

[A] In deadlock prevention, the request for resources is always granted if the resulting state is safe
[B] In deadlock avoidance, the request for resources is always granted if the resulting state is safe
[C]Deadlock avoidance is less restrictive than deadlock prevention
[D]Deadlock avoidance requires knowledge of resource requirements a priori

Deadlocks

8) Consider the following snapshot of a system running n processes. Process I is holding xi instances of a resource R, for 1 I n. Currently, all instances of R are occupied. Further, for all I, process I has placed a request for an additional yi instances while holding the xi instances it already has. There are exactly two processes p and q such that yp = yq = 0. Which one of the following can serve as a necessary condition to guarantee that the system is not approaching a deadlock ?

2 Marks GATE-CSE/IT-2006( )

[A] $\min(X_p, X_q) < \max_{k \neq p,q} Y_k$

[B] $X_p + X_q \max_{k \neq p,q} Y_k$

[C] $\min(X_p, X_q) < 1$

[D] $\min(X_p, X_q) > 1$

9) Suppose n processes, $P_1, \ldots, P_n$ share m identical resource units, which can be reserved and released one at a time. The maximum resource requirement of process p_i is s_p where $s_i > 0$. Which one of the following is a sufficient condition for ensuring that deadlock does not occur ?

2 Marks GATE-CSE/IT-2005( )

[A] $\forall_i^n s_i < m$

[B] $\forall_i^n s_i < n$

[C] $\sum_{i=1} s_i < (m + n)$

[D] $\sum_{i=1} s_i < (m * n)$

10) Which of the following concurrency control protocols ensure both conflict serializability and freedom from deadlock?

I. 2-phase locking

II. Time-stamp ordering

1 Marks GATE-CSE/IT-2010( )

[A] I only

[B] II only

[C] Both I and II

[D] Neither I nor II

Deadlocks

Key Paper

1.	**B**	**2.**	**A**	**3.**	**A**	**4.**	**C**	**5.**	**D**
6.	**C**	**7.**	**A**	**8.**	**B**	**9.**	**C**	**10.**	**B**

Memory Management

Common Data for Q1 and Q2 is given below

Suppose we want to synchronize two concurrent processes P and Q using binary semaphores S and T. The code for the processes P and Q is shown below.

```
Process P:                    Process Q:
while (1) {                   while (1) {
  W:                            Y:
    print'0';                   print '1';
    print '0'                    print '1';
  X:                            Z:
  }                             }
```

Synchronization statements can be inserted only at points W,X,Y and Z.

1) Which of the following will always lead to an output starting with '001100110011'?

2 Marks GATE-CSE/IT-2003( )

[A] P(S) at W, V(S) at X, P(T) at Y, V(T) at Z, S and T initially 1

[B] P(S) at W,V (T) at X,P(T) at Y,V(S) Z, S initially 1, and T initially 0

[C]P(S) at W,V)(T) at X,P (T) at Y,V(S) at Z, S and T initially 1

[D]P(S) at W, V(T) at X,P(T) at Y,V(T) at Z,S initially 1, and T initially 0

2) Which of the following will ensure that the output string never contains a substring of the form 01n0 or 10n1 where n is odd?

2 Marks GATE-CSE/IT-2003( )

[A] P(S) at W,V(S) at X,P(T) at Y,V(T) at Z,S and T initially 1

[B] P(S) at W,V(T) at X,P(T) at Y,V(S) at Z,S and T initially 1

[C]P(S) at W,V(S) at X,P(S) at Y,V(S) at Z, S initially 1

[D]V(S) at W,V(T) at X,P(S) at Y,P (T) at Z,S and T initially 1

Common Data for Q3 and Q4 is given below

A processor uses 2-level page tables for virtual to physical address translation. Page tables for both levels are stored in the main memory. Virtual and physical addresses are both 32 bits wide. The memory is byte addressable. For virtual to physical address translation, the 10 most significant bits of the virtual address are used as index into the first level page table while the next 10 bits are used as index into the second level page table. The 12 least significant bits of the virtual address are used as offset within the page. Assume that the page table entries in both levels of page tables are 4 bytes wide. Further, the processor has a translation look-aside buffer (TLB), with a hit rate of 96%. The TLB caches recently used virtual page numbers and the corresponding physical page numbers. The processor also has a physically addressed cache with a hit ratio of 90%. Main memory access time is 10ns, cache access time is 1 ns, and TLB access time is also 1ns.

3) Assuming that no page faults occur, the average time taken to access a virtual address is approximately (to the nearest 0.5 ns)

2 Marks GATE-CSE/IT-2003( )

[A] 1.5ns

[B] 2 ns

[C]3 ns

[D]4 ns

4) Suppose a process has only the following pages in its virtual address space: two contiguous code pages starting at virtual address 0 x 00000000, two contiguous data pages starting at virtual address 0 x 00400000, and a stack page starting at virtual address 0 x FFFFF000. The amount of memory required for storing the page tables of this process is

2 Marks GATE-CSE/IT-2003( )

[A] 8 KB

[B] 12 KB

[C]16 KB

[D]20KB

5) The essential content(s) in each entry of a page table is/are

1 Marks GATE-CSE/IT-2009( )

[A] virtual page number

[B] page frame number

[C]both virtual page number and page fame number

[D]access right information

6) A multilevel page table is preferred in comparison to a single level page table for translating virtual address to physical address because

2 Marks GATE-CSE/IT-2009( )

[A] it reduces the memory access time to read or write a memory location

[B] it helps to reduce the size of page table needed to implement the virtual address space of a process

[C] if is required by the translation lookaside buffer

[D]if helps to reduce the number of page faults in page replacement algorithms.

7) A RAM chip has a capacity of 1024 words of 8 bits each ($1K \times 8$). The number of 2×4 decoders with enable line needed to construct a $16K \times 16$ RAM from $1K \times 8$ RAM is

2 Marks GATE-CSE/IT-2013( )

[A] 4

[B] 5

[C]6

[D]7

8) In a system with 32 bit virtual addresses and 1 KB page size, use of one-level page tables for virtual to physical address translation is not practical because of

1 Marks GATE-CSE/IT-2003( )

[A] the large amount of internal fragmentation

[B] the large amount of external fragmentation

[C] the large memory overhead in maintaining page tables

[D] the large computation overhead in the translation process

9) The minimum number of page frames that must be allocated to a running process in a virtual memory environment is determined by

1 Marks GATE-CSE/IT-2004( )

[A] The instruction set architecture

[B] Page size

[C]Physical memory size

[D]number of processes in memory

10) The atomic feth-and-set x,y instruction unconditionally sets the memory location x to 1 and fetches the old value of x in y without allowing any intervening access to the memory location X. Consider the following implementation of P and V functions on a binary semaphore S.

```
void p (binary_ semaphore *S)
{
unsigned y;
unsigned *x = &(S-> value);}        do {
      fetch - and - set x,y;
    } while (y);
}
void V(binary_semphore *S) {
        {S-> value = 0 ;
}
```

Which one of the following is true ?

2 Marks GATE-CSE/IT-2006( )

[A] The implementation may not work if context switching is disabled in P

[B] Instead of using fetch-and-set, a pair of normal load/store can be used

[C] The implementation of V is wrong

[D]The code does not implement a binary semaphore

11) A CPU generates 32-bit virtual addresses. The page size is 4 KB. The processor has a translation look-aside buffer (TLB) which can hold a total of 128 page table entries and is 4-way set associative. The minimum size of the TLB tag is

2 Marks GATE-CSE/IT-2006( )

[A] 11 bits

[B] 13 bits

[C]15 bits

[D]20 bits

12) A computer system supports 32-bit virtual addresses as well as 32-bit physical addresses. Since the virtual address space is of the same size as the physical address space, the operating system designers decide to get rid of the virtual memory entirely. Which one of the following is true ?

2 Marks GATE-CSE/IT-2006( )

[A] Efficient implementation of multi-user support is no longer possible

[B] The processor cache organization can be made more efficient now

[C] Hardware support for memory management is no longer needed

[D]CPU scheduling can be made more efficient now

13) Fetch And Add (X, i) is an atomic Read-Modify-Write instruction that reads the value of memory location X, increments it by the value i, and returns the old value of X. It is used in the pseudocode shown below to implement a busy-wait lock. L is an unsigned integer shared variable initialized to 0. The value of 0 corresponds to lock being available, while any non-zero value corresponds to the lock being not available.

```
AcquireLock(L){
 While (Fetch And Add(L,1))
L = 1;
}
                Release Lock(L){
                L = 0;
                }
 This implementation
```

2 Marks GATE-CSE/IT-2012( )

[A] fails as L can overflow

[B] fails as L can take on a non-zero value when the lock is actually available

[C] works correctly but may starve some processes

[D] works correctly without starvation

14) Let the page fault service time be 10ms in a computer with average memory access time being 20ns. If one page fault is generated for every 10^6 memory accesses, what is the effective access time for the memory?

1 Marks GATE-CSE/IT-2011( )

[A] 21ns

[B] 30ns

[C] 23ns

[D] 35ns

15) Consider a hypothetical processor with an instruction of type LW R1, 20(R2), which during execution reads a 32-bit word from memory and stores it in a 32-bit register R1. The effective address of the memory location is obtained by the addition of constant 20 and the contents of register R2. Which of the following best reflects the addressing mode implemented by this instruction for the operand in memory?

1 Marks GATE-CSE/IT-2011( )

[A] Immediate Addressing

[B] Register Addressing

[C] Register Indirect Scaled Addressing

[D] Base Indexed Addressing

16) An 8KB direct mapped write-back cache is organized as multiple blocks, each of size 32-bytes. The processor generates 32-bit addresses. The cache controller maintains the tag information for each cache block comprising of the following.

1 Valid bit

1 Modified bit

As many bits as the minimum needed to identify the memory block mapped in the cache.

What is the total size of memory needed at the cache controller to store meta- data (tags) for the cache?

2 Marks GATE-CSE/IT-2011( )

[A] 4864bits

[B] 6144bits

[C] 6656bits

[D] 5376bits

17) A main memory unit with a capacity of 4 megabytes is built using 1M×1-bit DRAM chips. Each DRAM chip has 1K rows of cells with 1K cells in each row. The time taken for a single refresh operation is 100 nanoseconds. The time required to perform one refresh operation on all the cells in the memory unit is

1 Marks GATE-CSE/IT-2010( )

[A] 100 nanoseconds

[B] $100*2^{10}$ nanoseconds

[C] 1008^{20} nanoseconds

[D] 3200^{20} nanoseconds

18) Which of the following is not a form of memory?

1 Marks GATE-CSE/IT-2002( )

[A] instruction cache

[B] instruction register

[C] instruction opcode

[D] translation look-a-side buffer

19) The optimal page replacement algorithm will select the page that

1 Marks GATE-CSE/IT-2002( )

[A] Has not been used for the longest time in the past.

[B] Will not be used for the longest time in the future.

[C] Has been used least number of times.

[D] Has been used most number of times.

20) In the index allocation scheme of blocks to a file, the maximum possible size of the file depends on

2 Marks GATE-CSE/IT-2002( )

[A] the size of the blocks, and the size of the address of the blocks.

[B] the number of blocks used for the index, and the size of the blocks.

[C] the size of the blocks, the number of blocks used for the index, and the size of the address of the blocks.

[D] None of the above

21) More than one word are put in one cache block to

1 Marks GATE-CSE/IT-2001( )

[A] exploit the temporal locality of reference in a program

[B] exploit the spatial locality of reference in a program

[C] reduce the miss penalty

[D] none of the above

22) Which of the following statements is false?

1 Marks GATE-CSE/IT-2001( )

[A] Virtual memory implements the translation of a program's address space into physical memory address space

[B] Virtual memory allows each program to exceed the size of the primary memory

[C] Virtual memory increases the degree of multiprogramming

[D] Virtual memory reduces the context switching overhead

23) The process of assigning load addresses to the various parts of the program and adjusting the code and date in the program to reflect the assigned addresses is called

1 Marks GATE-CSE/IT-2001( )

[A] Assembly

[B] parsing

[C] Relocation

[D] Symbol resolution

24) Where does the swap space reside?

1 Marks GATE-CSE/IT-2001( )

[A] RAM

[B] Disk

[C] ROM

[D] On-chip cache

25) Consider a virtual memory system with FIFO page replacement policy. For an arbitrary page access pattern, increasing the number of page frames in main memory will

1 Marks GATE-CSE/IT-2001( )

[A] always decrease the number of page faults

[B] always increase the number of page faults

[C] sometimes increase the number of page faults

[D] never affect the number of page faults

26) A graphics card has on board memory of 1 MB. Which of the following modes can the card not support?

2 Marks GATE-CSE/IT-2000( )

[A] 1600×400 resolution with 256 colours on a 17 inch monitor

[B] 1600×400 resolution with 16 million colours on a 14 inch monitor

[C] 800×400 resolution with 16 million colours on a 17 inch monitor

[D] 800×800 resolution with 256 colours on a 14 inch monitor

27) Suppose the time to service a page fault is on the average 10 milliseconds, while a memory access takes 1 microsecond. Then a 99.99% hit ratio results in average memory access time of

2 Marks GATE-CSE/IT-2000( )

[A] 1.9999 milliseconds

[B] 1 millisecond

[C] 9.999 microseconds

[D] 1.9999 microseconds

28) Which of the following is/are advantage of virtual memory?

2 Marks GATE-CSE/IT-1999( )

[A] Faster access to memory on an average.

[B] Processes can be given protected address spaces.

[C] Linker can assign addresses independent of where the program will be loaded in physical memory.

[D] Programs larger than the physical memory size can be run.

29) A linker reads four modules whose lengths are 200, 800, 600 and 500 words, respectively. If they are loaded in that order, what are the relocation constants?

1 Marks GATE-CSE/IT-1998( )

[A] 0, 200, 500, 600

[B] 0, 200, 1000, 1600

[C] 200, 500, 600, 800

[D] 200, 700, 1300, 2100

30) A language L allows declaration of arrays whose sizes are not known during compilation. It is required to make efficient use of memory. Which one of the following is true?

1 Marks GATE-CSE/IT-1997( )

[A] A compiler using static memory allocation can be written for L

[B] A compiler cannot be written for L; an interpreter must be used.

[C] A compiler using dynamic memory allocation can be written for L

[D] None of the above

31) Thrashing

1 Marks GATE-CSE/IT-1997( )

[A] reduces page I/O

[B] decreases the degree of multiprogramming

[C] implies excessive page I/O

[D] improve the system performance

32) Dirty bit for a page in a page table

1 Marks GATE-CSE/IT-1997( )

[A] helps avoid unnecessary writes on a paging device

[B] helps maintain LRU information

[C] allows only read on a page

[D] None of the above

33) A 1000 Kbyte memory is managed using variable partitions but to compaction. It currently has two partitions of sizes 200 Kbytes and 260 Kbytes respectively. The smallest allocation request in Kbytes that could be denied is for

2 Marks GATE-CSE/IT-1996( )

[A] 151

[B] 181

[C] 231

[D] 541

34) In a paged segmented scheme of memory management, the segment table itself must have a page table because In a paged segmented

1 Marks GATE-CSE/IT-1995( )

[A] the segment table is often too large to fit in one page

[B] each segment is spread over a number of pages

[C] segment tables point to page table and not to the physical locations of the segment

[D] the processor's description base register points to a page table

35) Which of the following page replacement algorithms suffers from Belady's anamoly?

1 Marks GATE-CSE/IT-1995,GATE-CSE/IT-1995( )

[A] Optimal replacement

[B] LRU

[C] FIFO

[D] Both (a) and (c)

36) In a virtual memory system the address space specified by the address lines of the CUP must be_____than the physical memory size and_____than the secondary storage size.

2 Marks GATE-CSE/IT-1995( )

[A] smaller, smaller

[B] smaller, larger

[C] larger, smaller

[D] larger, larger

37) A memory page containing a heavily used variable that was initialized very early and is in constant use is removed when

1 Marks GATE-CSE/IT-1994( )

[A] LRU page replacement algorithm is used

[B] FIFO page replacement algorithm is used

[C] LFU page replacement algorithm is used

[D] None of the above

38) Consider the following heap (figure) in which blank regions are not in use and hatched region are in use.

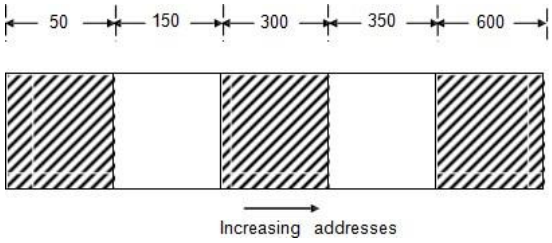

The sequence of requests for blocks of size 300, 25, 125, 50 can be satisfied if we use

1 Marks GATE-CSE/IT-1994( )

[A] either first fit or best fit policy (anyone)

[B] first fit but not best fit policy

[C] best fit but first fit policy

[D] None of the above

39) A Unix-style I-node has 10 direct pointers and one single, one double and one triple indirect pointers. Disk block size is 1 Kbyte, disk block address is 32 bits, and 48-bit integers are used. What is the maximum possible file size ?

2 Marks GATE-CSE/IT-2004( )

[A] 2^{24} bytes
[B] 2^{32} bytes
[C] 2^{34} bytes
[D] 2^{48} bytes

40) A CPU generates 32-bit virtual addresses. The page size is 4 KB. The processor has a translation look-aside buffer(TLB) which can hold a total of 128 page table entries and is 4-way setassociative. The minimum size of the TLB tag is

2 Marks GATE-CSE/IT-2006( )

[A] 11 bits
[B] 13 bits
[C] 15 bits
[D] 20 bits

41) A computer system supports 32-bit virtualaddresses as well as 32-bit physical addresses. Since the virtual address spaceis of the same size as the physical address space, the operating system designersdecide to get rid of the virtual memoryentirely. Which one of the following is true ?
(a) Efficient implementation of multi-user supportis no longer possible
(b) The processor cache organization can be mademore efficient now
(c)Hardware support for memory management is nolonger needed
(d) CPU scheduling can be made more efficient now

2 Marks GATE-CSE/IT-2006( )

[A] Efficient implementation of multi-user support is no longer possible
[B] The processor cache organization can be made more efficient now
[C] Hardware support for memory management is no longer needed
[D] CPU scheduling can be made more efficient now

42) The data block of a very large file in the Unix file system are allocated using

1 Marks GATE-CSE/IT-2008( )

[A] Contiguous allocation
[B] Linked allocation
[C] Indexed allocation
[D] an extension of indexed allocation

43) A processor uses 36 bit physical addresses and 32 bit virtual addresses, with a page frame size of 4 Kbytes. Each page table entry is of size 4 bytes. At hree level page table is used for virtual-to-physical address translation,where the virtual address is used as follows.
- bits 30-31 are used to index into the first level page table,
- bits 21-29 are used to index into the second level page table
- bits 12-20 are used to index into the third level page table
- bits 0-11 are used as offset within the page

The number of bits required for addressing the next level page table (or page frame) in the page table entry of the first, second and third level page tables are respectively.

2 Marks GATE-CSE/IT-2008( )

[A] 20, 20 and 20
[B] 24, 24 and 24
[C] 24, 24 and 20
[D] 25, 25 and 24

Statement for Linked answer Q44 and Q45 is given below

44) A computer uses 46-bit virtual address, 32-bit physical address, and a three-level paged page table organization. The page table base register stores the base address of the first-level table (T1) ,which occupies exactly one page. Each entry of T1 stores the base address of a page of the second-level table (T2). Each entry of T2 stores the base address of a page of the third-level table (T3) Each entry of T3 stores a page table entry (PTE). The PTE is 32 bits in size. The processor used in the computer has a 1 MB 16 way set associative virtually indexed physically tagged cache. The cache block size is 64 bytes.

Q. What is the size of a page in KB in this computer?

2 Marks GATE-CSE/IT-2013,GATE-CSE/IT-2013( )

[A] 2
[B] 4
[C] 8
[D] 16

45) What is the minimum number of page colours needed to guarantee that no two synonyms map to different sets in the processor cache of this computer?

2 Marks GATE-CSE/IT-2013( )

[A] 2
[B] 4
[C] 8
[D] 16

Memory Management

Key Paper

1.	**B**	**2.**	**C**	**3.**	**D**	**4.**	**C**	**5.**	**B**
6.	**B**	**7.**	**B**	**8.**	**C**	**9.**	**A**	**10.**	**A**
11.	**A**	**12.**	**C**	**13.**	**B**	**14.**	**B**	**15.**	**D**
16.	**D**	**17.**	**B**	**18.**	**C**	**19.**	**B**	**20.**	**B**
21.	**B**	**22.**	**D**	**23.**	**D**	**24.**	**B**	**25.**	**C**
26.	**A**	**27.**	**D**	**28.**	**D**	**29.**	**C**	**30.**	**C**
31.	**C**	**32.**	**A**	**33.**	**D**	**34.**	**A**	**35.**	**C**
36.	**C**	**37.**	**B**	**38.**	**B**	**39.**	**C**	**40.**	**C**
41.	**C**	**42.**	**D**	**43.**	**D**	**44.**	**C**	**45.**	**A**

Disk scheduling

1) Which of the following disk scheduling strategies is likely to give the best through put?

1 Marks GATE-CSE/IT-1999( )

[A] Farthest cylinder next
[B] Nearest cylinder next
[C] First come first served
[D] Elevator algorithm

2) The root directory of a disk should be placed

2 Marks GATE-CSE/IT-1993( )

[A] at a fixed address in main memory
[B] at a fixed location on the disk
[C] anywhere on the disk
[D] at a fixed location on the system disk
[E] anywhere on the system disk

3) Consider an operating system capable of loading and executing a single sequential user process at a time. The disk head scheduling algorithm used is First Served (FCFS). If FCFS is replaced by Shortest Seek Time First (SSTF), claimed by the vendor to given 50% better benchmark results, what is the expected improvement in the I/O performance of user programs?

1 Marks GATE-CSE/IT-2004( )

[A] 50%
[B] 40%
[C] 25%
[D] 0%

4) Consider a disk system with 100 cylinders. The requests to access the cylinders occur in following sequence :
4,34,10,7,19,73,2,15,6,20
Assuming that the head is currently at cylinder 50, what is the time taken o satisfy all requests if it takes 1 ms to move from one cylinder to adjacent one and shortest seek time first policy is used ?

2 Marks GATE-CSE/IT-2009( )

[A] 95 ms
[B] 119 ms
[C] 233 ms
[D] 276 ms

5) The enter_CS() and leave_CS() functions to implement critical section of a process are realized using test-and-set instructions as follows:

```
Void enter _ cs (X)
{
  while (test -and-set(X)):
}
Void leave_ CS (X)
{
  X=0 ;
}
```

In the above solution, X is a memory location associated with the CS and is initialized to 0. Now consider the following statements
I. The above solution to CS problem is deadlock-free
II. The solution is starvation free
III. The processes enter CS in FIFO order
IV. More than one process can enter CS at the same time
Which of the above statements are TRUE?

2 Marks GATE-CSE/IT-2009( )

[A] I only
[B] I and II
[C] II and III
[D] IV only

6) Using a larger blocks size in a fixed block size file system leads to

1 Marks GATE-CSE/IT-2003( )

[A] Better disk throughput but poorer disk space utilization
[B] Better disk throughput and better disk space utilization
[C] Poorer disk throughput but better disk space utilization
[D] Poorer disk throughput and poorer disk space utilization.

7) Consider an operating system capable of loading and executing a single sequential user process at a time. The disk head scheduling algorithm used is First Served (FCFS). If FCFS is replaced by Shortest Seek Time First (SSTF), claimed by the vendor to given 50% better benchmark results, what is the expected improvement in the I/O performance of user programs?

1 Marks GATE-CSE/IT-2004( )

[A] 50%
[B] 40%

[C]25% [D] 0%

8) Consider a hard disk with 16 recording surfaces $(0-15)$ having 16384 cylinders $(0-16383)$ and each cylinder contains 64 sectors $(0-63)$. Data storage capacity in each sector is 512 bytes. Data are organized cylinder-wise and the addressing format is . A file of size 42797 KB is stored in the disk and the starting disk location of the file is $<1200, 9, 40>$. What is the cylinder number of the last sector of the file, if it is stored in a contiguous manner?

2 Marks GATE-CSE/IT-2013( )

[A]1281 [B] 1282

[C]1283 [D]1284

9) A Unix-style I-node has 10 direct pointers and one single, one double and one triple indirect pointers. Disk block size is 1 Kbyte, disk block address is 32 bits, and 48-bit integers are used. What is the maximum possible file size ?

2 Marks GATE-CSE/IT-2004( )

[A] 2^{24} bytes [B] 2^{32} bytes

[C]2^{34} bytes [D] 2^{48} bytes

10) A file system with 300 GByte disk uses a file descriptor with 8 direct block addresses, 1 indirect block address and 1 doubly indirect block address. The size of each disk block is 128 Bytes and the size of each disk block address is 8 Bytes. The maximum possible file size in this file system is

2 Marks GATE-CSE/IT-2012( )

[A] 3 KBytes [B] 35 KBytes

[C]280 KBytes [D]dependent on the size of the disk

11) An application loads 100 libraries at startup. Loading each library requires exactly one disk access. The seek time of the disk to a random location is given as 10ms. Rotational speed of disk is 6000rpm. If all 100 libraries are loaded from random locations on the disk, how long does it take to load all libraries? (The time to transfer data from the disk block once the head has been positioned at the start of the block may be neglected)

2 Marks GATE-CSE/IT-2011( )

[A]0.50s [B] 1.50s

[C]1.25s [D]1.00s

Disk scheduling

Key Paper

1.	**C**	**2.**	**A**	**3.**	**D**	**4.**	**B**	**5.**	**A**
6.	**A**	**7.**	**D**	**8.**	**D**	**9.**	**C**	**10.**	**B**
11.	**B**								

Program Evaluation

1) The overlay tree for a program is as shown below:

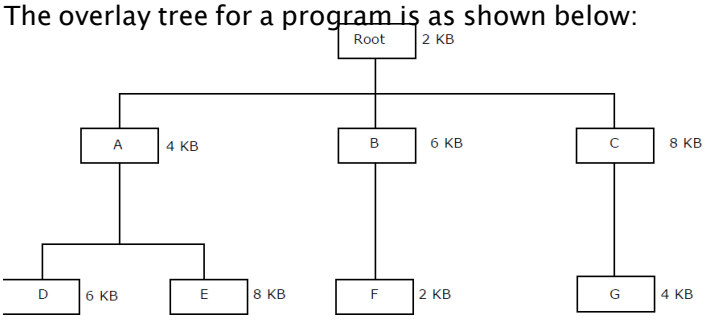

What will be the size of the partition (in physical memory) required to load (and run) this program?

2 Marks GATE-CSE/IT-1998( )

[A] 12 KB
[B] 14 KB
[C] 10 KB
[D] 8 KB

2) A critical section is a program segment

1 Marks GATE-CSE/IT-1996( )

[A] which should run in a certain specified amount of time
[B] which avoids deadlocks
[C] where shared resources are accessed
[D] which must be enclosed by a pair of semaphore operations, P and V

3) The address sequence generated by tracing a particular program executing in a pure demand paging system with 100 records per page with 1 free main memory frame is recorded as follows.What is the number of page faults?
0100 , 0200 , 0430 , 0499 , 0530 , 0560 , 0120 , 0220 , 0240 , 0260 , 0320 , 0370

2 Marks GATE-CSE/IT-1995( )

[A] 13
[B] 8
[C] 7
[D] 10

4) Which one of the following statements is true?

1 Marks GATE-CSE/IT-1994( )

[A] Macro definitions cannot appear within other macro definitions in assembly language programs
[B] Overlaying is used to run a program which is longer than the address space of computer
[C] Virtual memory can be used to accommodate a program which is longer than the address space of a computer
[D] It is not possible to write interrupt service routines in a high level language

5) A part of the system software, which under all circumstances must reside in the main memory, is

2 Marks GATE-CSE/IT-1993( )

[A] text editor
[B] assembler
[C] linker
[D] loader
[E] none of the above

6) Consider the following code fragment :

```
If (fork ( )= = 0)
{
a= a + 5;
print f("%d, %d/n" a, &a);
}
Else
{
a- 5;
print f("%d, %d/n", a, &a);
}
```

Let u, v be the values printed by the parent process, and x, y be the values printed by the child process. Which one of the following is TRUE?

2 Marks GATE-CSE/IT-2005( )

[A] $u = x + 10$ and $v = y$
[B] $u = x + 10$ and $v \neq y$
[C] $u + 10 = x$ and $v = y$
[D] $u + 10 = x$ and $v \neq y$

7) The P and V operations on counting semaphores, where s is a counting semaphore, are defined as follows :

P(s): s=s -1;
Ifs <0then wait
V(s): s= s +1 ;
If s< 0 then wakeup a process waiting on s;

Assume that Pb and Vb, the wait and signal operationson binary semaphores are provided. Twobinary semaphores Xb and Yb are used to implement the semaphoreoperations P(s) and V(s) as follows :

```
P(s):Pb (Xb);
    s =s - 1
    if (s <0) {
       Vb (Xb);
       Pb(Yb);
      }
    elseVb ( Xb) :
P(s) : Pb (Xb) ;
    s= s + 1;
    If(s<=0){
      Vb( Yb);
      Vb(Xb);
```

The initial values of Xb and Ybare respectively

2 Marks GATE-CSE/IT-2008( )

[A] 0 and 0 [B] 0 and 1
[C] 1 and 0 [D] 1 and 1

8) Consider an instruction pipeline with five stages without any branch prediction: Fetch Instruction (FI), Decode Instruction (DI), Fetch Operand (FO), Execute Instruction (EI) and Write Operand (WO). The stage delays for FI, DI, FO, EI and WO are 5 ns, 7 ns, 10 ns, 8 ns and 6 ns, respectively. There are intermediate storage buffers after each stage and the delay of each buffer is 1 ns. A programconsisting of 12 instructions $I_1, I_2, I_3, \ldots\ldots I_{12}$is executed in this pipelined processor. Instruction $I_4$is the only branch instruction and its branch target is I_9 .If the branch is taken during the execution of this program, the time (in ns) needed to complete the program is

2 Marks GATE-CSE/IT-2013( )

[A] 132 [B] 165
[C] 176 [D] 328

9) Consider the following code fragment :

```
If (fork ( )= = 0)
{a= a + 5; print f("%d, %d/n" a, &a); }
Else {a- 5; print f("%d, %d/n", a, &a); }
```

Let u, v be the values printed by the parent process, and x, y be the values printed by the child process. Which one of the following is TRUE?

2 Marks GATE-CSE/IT-2005( )

[A] u = x + 10 and v = y [B] u = x + 10 and v ≠ y
[C] u + 10 = x and v = y [D] u + 10 = x and v ≠ y

Program Evaluation

10) On a non-pipelined sequential processor, a program segment, which is a part of the interrupt service routine, is given to transfer 500 bytes from an I/O device to memory.

```
Initialize the address register
Initialize the count to 500
LOOP: Load a byte from device
Store in memory at address given by address register
Increment the address register
Decrement the count
If count != 0 go to LOOP
```

Assume that each statement in this program is equivalent to a machine instruction which takes one clock cycle to execute if it is a non-load/store instruction. The load-store instructions take two clock cycles to execute.

The designer of the system also has an alternate approach of using the DMA controller to implement the same transfer. The DMA controller requires 20 clock cycles for initialization and other overheads. Each DMA transfer cycle takes two clock cycles to transfer one byte of data from the device to the memory.

What is the approximate speedup when the DMA controller based design is used in place of the interrupt driven program based input-output?

2 Marks GATE-CSE/IT-2011( )

[A] 3,4

[B] 4,4

[C] 5,1

[D] 6,7

Program Evaluation

Key Paper

1.	**B**	**2.**	**C**	**3.**	**C**	**4.**	**B**	**5.**	**A**
6.	**C**	**7.**	**C**	**8.**	**C**	**9.**	**D**	**10.**	**A**