INFORMATION TECHNOLOGY

ONE MARKS QUESTIONS (1-30)

1. A bag contains 10 blue marbles, 20 green marbles and 30 red marbles. A marble is drawn from the bag, its colour recorded and it is put back in the bag. This process is repeated 3 times. The probability that no two of the marbles drawn have the same colour is
   a. 1/36
   b. 1/6
   c. 1/4
   d. 1/3

2. If the trapezoidal method is used to evaluate the integral then the value obtained
   a. Is always > 1/3
   b. is always < 1/3
   c. is always = 1/3
   d. may be greater or lesser than 1/3

3. The determinant of the matrix given below is

$$\begin{bmatrix} 0 & 1 & 0 & 2 \\ -1 & 1 & 1 & 3 \\ 0 & 0 & 0 & 1 \\ 1 & -2 & 0 & 1 \end{bmatrix}$$

   a. –1
   b. 0
   c. 1
   d. 2

4. Let L be a regular language and M be a context-free language, both over the alphabet Σ. Let L^C and M^C denote the complements of L and M respectively. Which of the following statements about the language if $L^C \cup M^C$ is TRUE?
   a. It is necessarily regular but not necessarily context-free.
   b. It is necessarily context-free.
   c. It is necessarily non-regular.
   d. None of the above

5. Which of the following statements is TRUE about the regular expression 01^*0?
   a. It represents a finite set of finite strings.
   b. It represents an infinite set of finite strings.
   c. It represents a finite set of infinite strings
   d. It represents an infinite set of infinite strings.

6. The language $\{0^n 1^n 2^n \mid 1 \le n \le 10^6\}$ is
   a. regular
   b. Context-free but not regular.
   c. Context-free but it s complement is not context-free.
   d. not context-free.

7. Which of the following expressions is equivalent to (A⊕B) ⊕C
   a. $(A+B+C)(\overline{A}+\overline{B}+\overline{C})$
   b. $(A+B+C)(\overline{A}+\overline{B}+C)$
   c. $ABC+\overline{A}(B\oplus C)+\overline{B}(A\oplus C)$
   d. None of these

8. Using Booth's Algorithm for multiplication, the multiplier -57 will be recoded as
   a. 0 –1 0 0 1 0 0 –1
   b. 1 1 0 0 0 1 1 1
   c. 0 – 1 0 0 1 0 0 0
   d. 0 1 0 0 –1 0 0 1

9. A dynamic RAM has a memory cycle time of 64 nsec. It has to be refreshed 100 times per msec and each refresh takes 100 nsec. What percentage of the memory cycle is used for refreshing?
   a. 10
   b. 6.4
   c. 1
   d. 0.64

10. A two-way switch has three terminals a, b, and c. In ON position (logic value 1), a is connected to b, and in OFF position, a is connected to c. Two of these two-way switches S1 and S2 are connected to a bulb as shown below

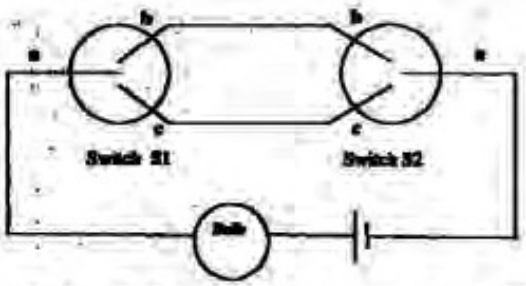

Which of the following expressions, if true, will always result in the lighting of the bulb

a. $S1\overline{S2}$

b. $S1+S2$

c. $\overline{S1 \oplus S2}$

d. $S1 \oplus S2$

11. How many pulses are needed to change the contents of a 8-bit upcounter from 10101100 to 00100111 (rightmost bit is the LSB)?

a. 134

b. 133

c. 124

d. 123

12. The numbers 1,2,…, n are inserted in a binary search tree in some order. In the resulting tree, the right subtree of the root contains p nodes. The first number to be inserted in the tree must be

a. p

b. p+1

c. n-p

d. n-p+1

13. A function f defined on stacks of integers satisfies the following properties. $f(\phi)=0$ and

f(push(S,i)) = max(f(S),0) + i for all stacks S and integers i.

If a stack S contains the integers 2, –3, 2, –1, 2 in order from bottom to top, what is f(S)?

a. 6

b. 4

c. 3

d. 2

14. In a depth-first traversal of a graph G with n vertices, k edges are marked as tree edges. The number of connected components in G is

a. k

b. k+1

c. n-k-1

d. n-k

15. In the following table, the left column contains the names of standard algorithms and the right column contains the time complexities of the graph algorithms. Match each algorithm with its time complexity.

1: Bellman-Ford algorithm	A: O(mlog n)
2: Kruskal's algorithm	B: $O(n^3)$
3: Floyd-Warshall algorithm	C: O(nm)
4: Topological Sorting	D: O(n+m)

a. 1→C, 2→A, 3→B, 4→D

b. 1→B, 2→D, 3→C, 4→A

c. 1→C, 2→D, 3→A, 4→B

d. 1→B, 2→A, 3→C, 4→D

16. A hash table contains 10 buckets and uses linear probing to resolve collisions. The key values are integers and the hash function used is key % 10. If the values 43, 165, 62, 123, 142 are inserted in the table, in what location would the key value 142 be inserted?

a. 2

b. 3

c. 4

d. 6

17. A student wishes to create symbolic links in a computer system running Unix. Three text files named "file1", "file2" and "file3" exist in her current working directory, and the student has read and writes permissions for all three files. Assume that file1 contains information about her hobbies, file2 contains information about her friends and file3 contains information about her courses. The student executes the following sequence of commands from her current working directory :

ln -s file1 file2

ln -s file2 file3

Which of the following types of information would he lose from her file system?

(I) Hobbies (II) Friends (III) Courses

a. (I) and (II) only

b. (II) and (III) only

c. (II) only
d. (I) and (III) only

18. The shell command

find. -name passwd -print

is executed in /etc directory of a computer system running Unix. Which of the following shell commands will give the same information as the above command when executed in the same directory?

a. ls passwd
b. cat passwd
c. grep name passwd
d. grep print passwd

19. A user level process in Unix traps the signal sent on a Ctrl-C input, and has a signal handling routine that saves appropriate files before terminating the process. When a Ctrl-C input is given to this process, what is the mode in which the signal handling routine executes?

a. kernel mode
b. super user mode
c. privileged mode
d. user mode

20. The Function Points (FP) calculated for a software project are often used to obtain an estimate of Lines of Code (LOC) required for that project. Which of the following statements is FALSE in this context

a. The relationship between FP and LOC depends on the programming language used to implement the software
b. LOC requirement for an assembly language implementation will be more for a given FP value, than LOC for implementation in COBOL
c. On an average, one LOC of C++ provides approximately 1.6 times the functionality of a single LOC of FORTRAN
d. FP and LOC are not related to each other

21. Consider the entities 'hotel room', and 'person' with a many to many relationship 'lodging' as shown below

If we wish to store information about the rent payment to be made by person(s) occupying different hotel rooms, then this information should appear as an attribute of

a. Person
b. Hotel Room
c. Lodging
d. None of these

22. A table has fields F1, F2, F3, F4, F5, with the following functional dependencies:

F1 → F3, F2 → F4, (F1, F2) → F5

In terms of Normalization, this table is in

a. INF
b. 2NF
c. 3NF
d. None of these

23. A B-Tree used as an index for a large database table has four levels including the root node. If a new key is inserted in this index, then the maximum number of nodes that could be newly created in the process are

a. 5
b. 4
c. 3
d. 2

24. Amongst the ACID properties of a transaction, the 'Durability' property requires that the changes made to the database by a successful transaction persist

a. Except in case of an Operating System crash
b. Except in case of a Disk crash
c. Except in case of a power failure
d. Always, even if there is a failure of any kind

25. Consider the three commands: PROMPT, HEAD and RCPT.

Which of the following options indicate a correct association of these commands with protocols where these are used?

a. HTTP, SMTP, FTP
b. FTP, HTTP, SMTP
c. HTTP, FTP, SMTP
d. SMTP, HTTP, FTP

26. Traceroute reports a possible route that is taken by packets moving from some host A to some other host B. Which of the

following options represents the technique used traceroute to identify these hosts

a. By progressively querying routers about the next router on the path to B using ICMP packets, starting with the first router

b. By requiring each muter to append the address to the ICMP packet as it is forwarded to B. The list of all routers en-route to B is returned by 13 in an ICMP reply packet

c. By ensuring that an ICMP reply packet is returned to A by each router en-route to B, in the ascending order of their hop distance from A

d. By locally computing the shortest path from A to B

27. Which of the following statements is TRUE about CSMA/CD

a. IEEE 802.11 wireless LAN runs CSMA/CD protocol

b. Ethernet is not based on CSMA/CD protocol

c. CSMA/CO is not suitable for a high propagation delay network like satellite network.

d. There is no contention in a CSMA/CD network

28. Which of the following statements is FALSE regarding a bridge

a. Bridge is a layer 2 device

b. Bridge reduces collision domain

c. Bridge is used to connect two or more LAN segments

d. Bridge reduces broadcast domain.

29. Count to infinity is a problem associated with

a. link state routing protocol.

b. distance vector muting protocol.

c. DNS while resolving host name.

d. TCP for congestion control.

30. A HTML form is to be designed to enable purchase of office stationery. Required items are to be selected (checked). Credit card details are to be entered and then the submit button is to be pressed. Which one of the following options would be appropriate for sending the data to the server. Assume that security is handled in a way that is transparent to the form design.

a. Only GET

b. Only POST

c. Either of GET or POST

d. Neither GET nor POST

TWO MARKS QUESTIONS (31-60)

31. Let f be a function from a set A to a set B, g a function from B to C, and ha function from A to P. such that $h(a) = g(f(a))$ for all $a \in A$. Which of the following statements is always true for all such functions f and g?

a. g is onto $\Rightarrow$ h is onto

b. h is onto $\Rightarrow$ f is onto

c. h is onto $\Rightarrow$ g is onto

d. h is onto $\Rightarrow$ f and g are onto

32. Let A be a set with n elements. Let C be a collection of distinct subsets of A such that for any two subsets S1 and S2 in C, either $S_1 \subset S_2$ or $S_2 \subset S_1$. What is the maximum cardinality of C?

a. n

b. n+1

c. $2^{n-1} + 1$

d. n!

33. An unbiased coin is tossed repeatedly until the outcome of two successive tosses is the same. Assuming that the trials are independent, the expected number of tosses is

a. 3

b. 4

c. 5

d. 6

34. Let $n = p^2q$, where p and q are distinct prime numbers. How many numbers m satisfy $1 \leq m \leq n$ and gcd(m,n) = 1? Note that gcd(m,n) is the greatest common divisor of m and n

a. p(q–1)

b. pq

c. $(p^2-1)(q-1$

d. p(p-1)(q-1)

35. What is the value of

$$\int_0^{2\pi} (x-\pi)^3 (\sin x)\,dx$$

a. -1

b. 0

c. 1

d. π

36. Let P(x) and Q(x) be arbitrary predicates. Which of the following statements is always TRUE?

a. $(\forall x (P(x) \vee Q(x))) \Rightarrow ((\forall x P(x)) \vee (\forall x Q(x)))$

b. $(\forall x (P(x) \Rightarrow Q(x))) \Rightarrow ((\forall x P(x)) \Rightarrow (\forall x Q(x)))$

c. $(\forall x (P(x) \Rightarrow (\forall x Q(x))) \Rightarrow (\forall x (P(x) \Rightarrow Q(x)))$

d. $((\forall x (P(x))) \Leftrightarrow (\forall x (Q(x)))) \Rightarrow (\forall x (P(x) \Leftrightarrow Q(x)))$

37. Consider the non-deterministic Finite automaton (NFA) shown in the figure.

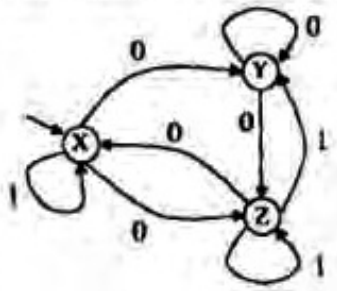

State X is the starting state of the automaton. Let the language accepted by the NFA with Y as the only accepting state be L1. Similarly, let the language accepted by the NFA with Z as the only accepting state be L2. Which of the following statements about L1 and L2 is TRUE?

a. $L1 = L2$

b. $L1 \subset L2$

c. $L2 \subset L1$

d. None of the above

38. Let P be a non-deterministic push-down automaton (NPDA) with exactly one state, q, and exactly one symbol, Z, in its stack alphabet. State q is both the starting as well as the accepting state of the PDA. The stack is initialized with one Z before the start of the operation of the PDA. Let the input alphabet of the PDA be Σ. Let L(P) he the language accepted by the PDA by reading a string and reaching its accepting state. Let N(P) be the language accepted by the PDA by reading a string and emptying its stack.

Which of the following statements is TRUE?

a. L(P) is necessarily Σ^* but N(P) is not necessarily Σ^*.

b. N(P) is necessarily Σ^* but L(P) is not necessarily Σ^*.

c. Both L(P) and N(P) are necessarily Σ^*.

d. Neither L(P) nor N(P) are necessarily Σ^*.

39. Consider the regular grammar:

$S \rightarrow Xa \mid Ya$

$X \rightarrow Za$

$Z \rightarrow SaE$

$Y \rightarrow Wa$

$W \rightarrow Sa$

where S is the starting symbol, the set of terminals is {a) and the set of non-terminals is {S,W,X,Y,Z}.

We wish to construct a deterministic finite automaton (DFA) to recognize the same language. What is the minimum number of states required for the DFA?

a. 2

b. 3

c. 4

d. 5

40. A language L satisfies the Pumping Lemma for regular languages, and also the Pumping Lemma for context-free languages. Which of the following statements about L is TRUE?

a. L is necessarily a regular language.

b. L is necessarily a context-free language, but not necessarily a regular language.

c. L is necessarily a non-regular language.

d. None of the above.

41. Given below is a program which when executed spawns two concurrent processes:

```
semaphore X := 0;
/* Process now forks into concurrent processes P1 & P2
*/
P1 : repeat forever          P2 : repeat forever
     V(X);                        P(X);
     Compute;                     Compute;
     P(X);                        V(X);
```

Consider the following statements about processes P1 and P2.

I : It is possible for process P1 to starve.

II: It is possible for process P2 to starve.

Which of the following holds?

a. Both I and II are true.

b. I is true but II is false.

c. II is true hut I is false.

d. Both I and II are false.

42. Two concurrent processes P1 and P2 use four shared resources R1, R2, R3 and R4, as shown below.

P1:	P2:
Compute;	Compute;
Use R1;	Use R1;
Use R2;	Use R2;
Use R3;	Use R3;
Use R4;	Use R4;

Both processes are started at the same time, and each resource can be accessed by only one process at a time. The following scheduling constraints exist between the access of resources by the processes:

P2 must complete use of RI before P1 gets access to RI.

P1 must complete use of R2 before P2 gets access to R2.

P2 must complete use of R3 before P1 gets access to R3.

P1 must complete use of R4 before P2 gets access to R4.

There are no other scheduling constraints between the processes. If only binary semaphores are used to enforce the above scheduling constraints, what is the minimum number of binary semaphores needed?

a. 1

b. 2

c. 3

d. 4

43. Which of the following input sequences will always generate a 1 at the output z at the end of the third cycle?

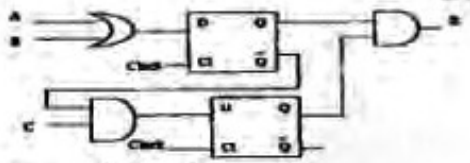

a.

A	B	C
0	0	0
1	0	1
1	1	1

b.

A	B	C
1	0	1
1	1	0
1	1	1

c.

A	B	C
0	1	1
1	0	1
1	1	1

d.

A	B	C
0	0	1
1	1	0
1	1	1

44. We have two designs D1 and D2 for a synchronous pipeline processor. DI has 5 pipeline stages with execution times of 3 nsec, 2 nsec, 4 nsec, 2 nscc and 3 nsec whie the design D2 has 8 pipeline stages each with 2 nsec execution time. How much time can be saved using design D2 over Design D1 for executing 100 instructions ?

a. 214 nscc

b. 202 nsec

c. 86 nsec

d. -200 nsec

45. A hardwired CPU uses 10 control signals silo s1 to s10 in various dine steps T1 to T5 to implement 4 instructions I1 to 14 as shown below.

	T1	T2	T3	T4	T5
I1	S1, S3, S5	S2, S4, S6	S1, S7	S10	S3, S8
I2	S1, S3, S5	S8, S9, S10	S5, S6, S7	S6	S10
I3	S1, S3, S5	S7, S8, S10	S2, S6, S9	S10	S1, S3
I4	S1, S3, S5	S2, S6, S7	S5, S10	S6, S9	S10

Which of the following pairs of expressions represent the circuit for generating control signals S5 and S10 respectively [(Ij+Ik). Tn indicates that the control signal should be generated in time step Tn if the instruction being executed is Ij or Ik] ?

a. S5=T1+I2.T3 & S10 = (I1+I3).T4 + (I2+I4).T5

b. S5=TI+(12+14).T3 & S10 = (I1+13), T4 +(12+14)T5

c. S5 = T1+(12+14).13 & S10 = (I2+I3+ I4).T2 + (I1 + I3).T4 + (I2+I4) T5

d. S5=T1+(I2+I4).T3 & S10 = (I2+I3).T2 + I4.T3 + (I1+I3.T4 + (I2+ I4), T5

46. A line L in a circuit is said to have a stuck-at-0 fault if the line permanently has a logic value 0. Similarly a line L in a circuit is said to have a stuck-at-1 fault if the line permanently has a logic value 1. A circuit is said to have a multiple stuck-at fault if one or more lanes have stuck at faults. The total number of distinct multiple stuck at faults possible in a circuit with N lines is

a. 3^N

b. 3^N-1

c. 2^N-1

d. 2^N

47. $(33.4)_8 \times (23.4)_8$ evaluates to

a. $(1053.6)_8$

b. $(1053.2)_8$

c. $(1024.2)_8$

d. None of these

48. The circuit shown below implement a 2-input NOR gate using two 2-1 MUX (control signal I selects the tipper input). What are the values of signals x, y, & z?

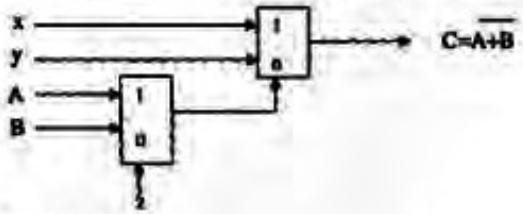

a. 1,0,B

b. 1,0,A

c. 0,1,8

d. 0,1,A

49. An instruction set of a processor has 125 signals which can be divided into 5 groups of mutually exclusive signals as follows:

Group 1 : 20 signals. Group 2: 70 signals, Group 3: 2 signals. Group 4: 10 signals, Group 5: 23 signals.

How many hits of the control words can be saved by using vertical microprogramming over horizontal microprogramming?

a. 0

b. 103

c. 22

d. 55

50. In a binary tree, for every node the difference between the number of nodes in the left and right subtrees is at most 2. If the height of the tree is h > 0. then the minimum number of nodes in the tree is

a. 2^h-1

b. $2^{h-1}+1$

c. 2^h-1

d. 2^h

51. Let T(n) be a function defined by the recurrence

$T(n) = 2T(n/2) + \sqrt{n}$ for $n \geq 2$ and

$T(1) = 1$

Which of the following statements is TRUE?

a. $T(n) = \theta(\log n)$

b. $T(n) = \theta(\sqrt{n})$

c. $T(n) = \theta(n)$

d. $T(n) = \theta(n \log n)$

52. Let G be a weighted undirected graph and e he an edge with maximum weight in G. Suppose there is a minimum weight spanning tree in G containing the edge e. Which of the following statements is always TRUE?

a. There exists a cutset in G having all edges of maximum weight.

b. There exists a cycle in G having all edges of maximum weight.

c. Edge e cannot he contained in a cycle.

d. All edges in G have the same weight.

53. The following C function takes two ASCII strings and determines whether one is an anagram of the other. An anagram of a string s is a string obtained by permuting the letters in s.

```
int anagram( char*a, char*b) {
    int count[128],j;
    for (j = 0;j< 128; j++) count[j]=0;
    j=0;
    while (a[j] && b[j])    {
        A;
        B;
```

```
}
for(j =0;j <128; j++) if (count[j]) return 0;
return 1;
}
```

Choose the correct alternative for statements A and B.

a. A : count[a[j]]++ and B count[b[j]]--

b. A : count[a[j]]++ and B : count[b[j]++

c. A : count(a[j++]]++ and B : count[b[j]]--

d. A :count[a[j]]++ and B count[b[j++]]--

54. The following C function takes a singly-linked List of integers as a parameter and rearranges the elements of the list. The list is represented ac pointer to a structure.

The function is culled with the list containing the integers 1, 2, 3, 4, 5, 6, 7 in the given order. What will he the contents of the list after the function completes execution?

```
struct node { int value; struct node *next; };
void rearrange(struct node *list) {
   struct node *p, *q;
   int temp;
   if (list || !list->next) return;
   p = list; q = list->next;
while(q) {
         temp = p->value;
         p->value = q->value;
         q->value = temp;
         p = q->next;
         q = p ? p->next : 0;
   }
}
```

a. 1,2,3,4,5,6,7

b. 2,1,4,3,6,5,7

c. 1,3,2,5,4,7,6

d. 2,3,4,5,6,7,1

55. A binary search tree contains the numbers 1, 2, 3, 4, 5, 6, 7, 8. When the tree is traversed in pre-order and the values in each node printed out, the sequence of values obtained is 5, 3, 1, 2, 4, 6, 8, 7. If the tree is traversed in post-order, the sequence obtained would be

a. 8, 7, 6, 5, 4, 3, 2, 1

b. 1, 2, 3, 4, 8, 7, 6, 5

c. 1, 2, 1, 4, 3, 6, 7, 8, 5

d. 2, 1, 4, 3, 7, 8, 6, 5

56. Let G be a directed graph whose vertex set is the set of numbers from 1to 100. There is an edge from a vertex i to a vertex j iff either j = i + 1 or j = 3i. The minimum number of edges in a path in G from vertex 1 to vertex 100 is

a. 4

b. 7

c. 23

d. 99

57. What is the output printed by the following program?

```
#include <stdio.h>
Int f(int n, int k) {
   if (n = 0) return 0;
   else if (n % 2) return f(n/2, 2*k) + k;
   else return f(n/2, 2*k) – k;
}
int main() {
 printf("%d", f(20, 1));
 return 0;
}
```

a. 5

b. 8

c. 9

d. 20

58. Let a be an array containing n integers in increasing order. The following algorithm determines whether there are two distinct numbers in the array whose difference is a specified number S >0.

```
i=0; j=1;
while (j<n) {
   if (E) j++;
   else if (a[j] – a[i] == S) break;
   else i++;
}
if (j < n) printf("yes") else printf("no");
```

Choose the correct expression for E.

a. a[j] – a[i] > S

b. a[j] – a[i] < S

c. a[i] – a[j] < S

d. a[i] – a(j) > S

59. Let a and b be two sorted arrays containing n integers each, in non-decreasing order. Let c be a sorted array containing 2n integers obtained by merging the two arrays a and b. Assuming the arrays are

indexed starting from 0, consider the following four statements.

I $a[j] \geq b[i] \Rightarrow c[2i] \geq a[i]$

II $a[i] \geq b[i] \Rightarrow c[21] \geq h[i]$

III $a[i] \geq b[i] \Rightarrow c[2i] \geq a[i]$

IV $a[i] \geq b[i] \Rightarrow c[2i] \geq b[i]$

Which of the following is TRUE

a. only I and II

b. only I and IV

c. only II and III

d. only III and IV

60. We wish to schedule three processes P1, P2 and P3 on a uniprocessor system. The priorities, CPU lime requirements and arrival times of the processes are as shown below

Process	Priority	CPU time required	Arrival time (hh:mm:ss)
P1	10 (highest)	20 sec	00:00:05
P2	9	10 sec	00:00:03
P3	8 (lowest)	15 sec	00:00:00

We have a choice of preemptive or non-preemptive scheduling. In preemptive scheduling, a late-arriving higher priority process can preempt a currently running process with lower priority. In non-preemptive scheduling, a late-arriving higher priority process must wait for the currently executing process to complete before it can he scheduled on the processor.

What are the turnaround times (time from arrival till completion) of P2 using preemptive and non-preemptive scheduling respectively?

a. 30 sec, 30 sec

b. 30 sec, 10 sec

c. 42 sec, 42 sec

d. 30 sec, 42 sec

61. Consider a 2-way set associative cache memory with 4 sets and total S cache blocks (0-7) and a main memory with 128 blocks (0-127). What memory blocks will be present in the cache after the following sequence of memory block references if LRU policy is used for cache block replacement. Assuming that initially the cache did not have any memory block from the current job?

0 5 3 9 7 0 16 55

a. 0 3 5 7 16 55

b. 0 3 5 7 9 16 55

c. 0 5 7 9 16 55

d. 3 5 7 9 16 55

62. Two shared resources R_1 and R_2 are used by processes P_1 and P_2. Each process has certain priority for accessing each resource. Let T_{ij} denote the priority of P_i for accessing R_j. A process P can snatch a resource R_k from process P_j if T_{ik} is greater than T_{jk}

Given the following :

(I) $T_{11} > T_{21}$

(II) $T_{12} > T_{22}$

(III) $T_{11} < T_{21}$

(IV) $T_{12} < T_{22}$

Which of the following conditions ensures that P_1 and P_2 can never deadlock?

a. (I) and (IV)

b. (II) and (III)

c. (I) and (II)

d. None of the above

63. In a computer system, four files of size 11050 bytes, 4990 bytes, 5170 bytes and 12640 bytes need to be stored. For storing these files on disk, we can use either 100 byte disk blocks or 200 byte disk blocks (but can't mix block sizes. For each block used to store a file, 4 bytes of bookkeeping information also needs to he stored on the disk, Thus, the total space used to store a file is the sum of the space taken to store the file and the space taken to store the bookkeeping information for the Mocks allocated for storing the file. A disk block can store either bookkeeping information for a file or data from a file, but not both.

What is the total space required for storing the flies using 100 byte disk blocks and 200 byte disk blocks respectively?

a. 35400 and 35800 bytes

b. 35800 and 35400 bytes

c. 35600 and 35400 bytes

d. 35400 and 35600 bytes

64. The availability of a complex software is 90%. Its Mean Time Between Failure (MTBF) is 200 days. Because of the critical nature of the usage, the organization deploying the software further enhanced it to obtain an availability

of 95%. In the process, the Mean Time To Repair (MTTR) increased by 5 days.

What is the MTBF of the enhanced software

a. 205 days
b. 300 days
c. 500 days
d. 700 days

65. To carry out white box testing of a program, its flow chart representation is obtained as shown in the figure below:

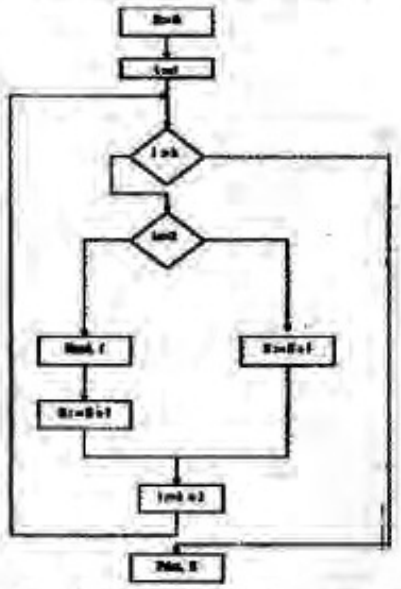

For basis path based testing of this program, its cyclomatic complexity is

a. 5
b. 4
c. 3
d. 2

66. In a data flow diagram, the segment shown below is identified as having transaction flow characteristics, with p2 identified as the transaction center

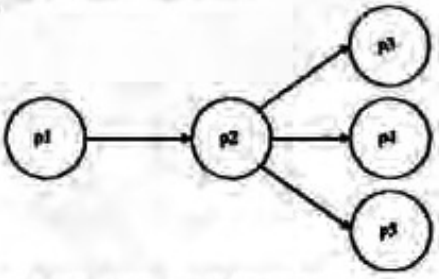

A first level architectural design of this segment will result in a set of process modules with an associated invocation sequence. The most appropriate architecture is

a. p1 invokes p2, p2 invokes either p3, or p4, or p5
b. p2 Invokes p1, and then invokes p3, or p4, or p5
c. A new module Tc is defined to control the transaction flow. This module Ye first invokes p1 and then invokes p2, p2 in turn invokes p3, or p4, or p5
d. A new module Tc is defined to control the transaction flow. This module Tc invokes p2, p2 invokes p1, and then invokes p3, or p4 or p5

67. A company maintains records of sales made by its salespersons and pays them commission based on each individual's total sales made in a year. This data is maintained in a table with following schema :

salesinfo = (salespersonid, totalsales, commission)

In a certain year, due to better business results, the company decides to further reward its salespersons by enhancing the commission paid to them as per the following formula

If commission <= 50000, enhance it by 2%

if 50000 < commission <= 100000, enhance it by 4%

If commission > 100000, enhance it by 6%

The IT stall has written three different SQL script to calculate enhancement for each slab, cacti of these scripts is to run as a separate transaction as follows:

T1 Update salesinfo
Set commission = commission * 1.02
where commission <= 50000;

T2 Update salesinfo
Set commission = commission * 1.04
Where commission > 50000 and commission is <= 100000;

T3 Update salesinfo
Set commission = commission * 1.06
Where commission > 100000;

Which of the following options of running these transactions will update the commission of all salespersons correctly?

a. Execute T1 followed by T2 followed by T3

b. Execute T2 followed by T3; T1 running concurrently throughout

c. Execute T3 followed by T2; T1 running concurrently throughout

d. Execute T3 followed by T2 followed by T1

68. A table 'student' with schema (roll, name, hostel, marks) and another table 'hobby' with schema (roll, hobbyname) contains records as shown below:

Table student

Roll	Name	Hostel	Marks
1798	Manoj Rathod	7	95
2154	Soumic Banerjee	5	68
2369	Gumma Reddy	7	86
2581	Pradeep Pendse	6	92
2643	Suhas Kulakarni	5	78
2711	Nitin Kadam	8	72
2872	Kirati Vora	5	92
2926	Manoj Kunkalikar	5	94
2959	Hemant Karkhanis	7	88
3125	Rajesh Doshi	5	82

Table hobby

Roll	hobbyname
1798	chess
1798	music
2154	music
2369	swimming
2581	cricket
2613	chess
2643	hockey
2711	volleyball
2872	football
2926	cricket
2959	photography
3125	music
3125	chess

The following SQL query is executed on the above tables:

select hostel

from student natural join hobby

where marks >=75 and roll between 2000 and 3000;

Relations S and H with the same schema as those of these two tables respectively contain the same information as tuples. A new relation S' is obtained by the following relational algebra operation:

$S' = \prod_{hostle}((\sigma_{S.roll = H.roll}(\sigma_{marks > 75 \text{ and } roll > 2000 \text{ and } < 3000}(S)) X (H))$

The difference between the number of rows output by the SQL statement and the number of tuples in S' is

a. 6

b. 4

c. 2

d. 0

69. In an inventory management system implemented at a trading corporation, there are several tables designed to hold all the information. Amongst these, the following to tables hold information on which items are supplied by which suppliers, and which warehouse keeps which items along with the stock-level of these items.

Supply = (supplierid, itemcode)

Inventory = (itemcode, warehouse, stocklevel)

For a specific information required by the management, following SQL query has been written

Select distinct STMP.suuplierid

From Supply as STMP

Where not unique (select ITMP.supplierid

From Inventory, Supply as ITMP

Where STMP.supplierid=ITMP.supplierid

And ITMP.itemcode = Inventory.itemcode

And Inventory.warehouse = 'Nagpur';

For the warehouse at Nagpur, this query will find all suppliers who

a. do not supply any item

b. supply exactly one item

c. supply one or more items

d. supply two or more items

70. In a schema with attributes A, B, C, D, and E, following set of functional dependencies are given

$A \rightarrow B$

$A \rightarrow C$

$CD \rightarrow E$

$B \rightarrow D$

$E \rightarrow A$

Which of the following functional dependencies is NOT implied by the above set ?

a. $CD \rightarrow AC$

b. $BD \rightarrow CD$

c. $BC \rightarrow CD$

d. $AC \rightarrow BC$

71. A network with CSMA/CD protocol in the MAC layer is running at 1 Gbps over a 1 km cubic with no repeaters. The signal speed in the cable is 2×10^8 m/sec. The minimum frame size for this network should be

a. 10000 bits

b. 10000 bytes

c. 5000 bits

d. 5000 bytes

72. A channel has a hit rate of 4 kbps and one-way propagation delay of 20ms. The channel uses stop and wait protocol. The transmission time of the acknowledgement frame is negligible. To get a channel efficiency of at least 50%, the minimum frame size should be

a. 80 bytes

b. 80 bits

c. 160 bytes

d. 160 bits

73. On a TCP connection, current congestion window size is Congestion Window=4KB. The window size advertised by the receiver is Advertise Window=6KB. The last byte sent by the sender is LastByteSent=10240 and the last byte acknowledged by the receiver is LastByteAcked=8192. The current window size at the sender is

a. 2048 bytes

b. 4096 bytes

c. 6144 bytes

d. 8192 bytes

74. In a communication network, a packet of length L bits takes link L1 with a probability of pl, or link L2 with a probability of p2. Link L1 and L2 have bit error probability of b_1 and b_2 respectively. The probability that the packet will he received without error via either L1 or L2 is

a. $(1-b_1)^L p_1 + (1-b_2)^L p_2$

b. $[1-(b_1+b_2)^L]\, p_1\, p_2$

c. $(1-b_1)^L (1-b_2)^L\, p_1\, p_2$

d. $1-(b_1^L p_1 + b_2^L p_2)$

75. In a TDM medium access control bus LAN, each station is assigned one time slot per cycle for transmission. Assume that the length of each time slot is the time to transmit 100 bits plus the end-to-end propagation delay. Assume a propagation speed of 2×10^8 m/sec. The length of the LAN is 1 km with a bandwidth of 10Mbps. The maximum number of stations that can be allowed in the LAN so that the throughput of each station can be 2/3 Mbps is

a. 3

b. 5

c. 10

d. 20

76. A company has a class C network address of 204.204.201.0. It wishes to have three subjects, one with 100 hosts and two with 50 hosts each. Which one of the following options represents a feasible set of subnet address / subnet mask pairs?

a. 204.204.204.128 / 255.255.255.192
   204.204.204.0 / 255.255.255.128
   204.204.204.64 / 255.255.255.128

b. 204.204.204.0 / 255.255.255.192
   204.204.204.192 / 255.255.255.128
   204.204.204.64 / 255.255.255.128

c. 204.204.204.128 / 255.255.255.128
   204.204.204.192 / 255.255.255.192
   204.204.204.224 / 255.255.255.192

d. 204.204.204.128 / 255.255.255.128
   204.204.204.64 / 255.255.255.192
   204.204.204.0 / 255.255.255.192

77. Assume that "host1.mydomain.dom" has an IP address of 145.128.16.8. Which of the following options would be most appropriate as a subsequence of steps in performing the reverse lookup of 145.128.16.8? In the following options "NS" is an abbreviation of "nameserver".

a. Query a NS for the root domain and then NS for the "dom" domains

b. Directly query a NS br "dom" and then a NS for "mydomain.dom" domains

c. Query a NS for in-addr.arpa and then a NS for 128.145.in-addr.arpa domains

d. Directly query a NS for 145.in-addr.arpa and then a NS for 128.145.in-addr.arpa domains

78. Consider the following message M=1010001101. The cyclic redundancy check (CRC) for this message using the divisor polynomial $x^5 + x^4 + x^2 + 1$ is:

a. 01110

b. 01011

c. 10101

d. 10110

79. Suppose that two parties A and B wish to setup a common secret key (D-H key)

between themselves using the Diffie-Hellman key exchange technique. They agree on 7 as the modulus and 3 as the primitive root. Party A chooses 2 and party B chooses 5 as their respective secrets. Their D-H key is:

a. 3

b. 4

c. 5

d. 6

80. Given below is an excerpt of an xml specification.

```
<!DOCTYPE library SYSTEM "library.dtd">
<Book>
<title> GATE 2005 </title>
<type value="BROCHURE" />
<accno> 10237623786 </accno>
</Book>
<Book>
<type value="FICTION" />
<accno> 0024154807 </accno>
</Book>
```

Given below are several possible excerpts from "library.dtd". For which excerpt would the above specification be valid?

a.

```
<!ELEMENT Book (title+, type, accno)>
<!ELEMENT title (#PCDATA)>
<!ELEMENT type EMPTY>
<!ATTLIST type value (BROCHURE|FICTION|TECHNICAL)>
<!ELEMENT accno (#PCDATA)>
```

b.

```
<!ELEMENT Book (title?, type, accno)>
<!ELEMENT title (#PCDATA)>
<!ELEMENT type ATTLIST>
<!ATTLIST type value (BROCHURE|FICTION|TECHNICAL)>
<!ATTLIST accno value (#PCDATA)>
```

c.

```
<!ELEMENT Book (title*, type, accno)>
<!ELEMENT title (#PCDATA)>
<!ELEMENT type ATTLIST>
<!ATTLIST type value (BROCHURE|FICTION|TECHNICAL)>
<!ELEMENT accno (#PCDATA)>
```

d.

```
<!ELEMENT Book (title*, type, accno)>
<!ELEMENT title (#PCDATA)>
<!ELEMENT type EMPTY>
<!ATTLIST type value (BROCHURE|FICTION|TECHNICAL)>
<!ELEMENT accno (#PCDATA)>
```

Statement for Linked Answer Question (81.1 and 81.2)

A disk has 8 equidistant tracks. The diameters of the innermost and outermost tracks are 1 cm and 8 cm respectively. The innermost track has a storage capacity of 10MB.

81.1 What is the total amount of data that can be stored on the disk if it is used with a drive that rotates it with (i) Constant Linear Velocity (ii) Constant Angular Velocity?

(a) (i) 80MB (ii) 2040 MB

(b) (i) 2040 MB (ii) 80MB

(c) (i) 80MB (ii) 360 MB

(d) (i) 360 MB (ii) 80MB

81.2 If the disk has 20 sectors per track and is currently at the end of the 5th sector of the inner-most track and the head can move at a speed of 10 meters/sec and ills rotating at constant angular velocity of 6000 RPM, how much time will it lake to read 1 MB contiguous data starting from the sector 4 of the outer-most track?

(a) 13.5 ms

(b) 10 ms

(c) 9.5 ms

(d) 20 ms

Statement for Linked Answer Question (82.1 and 82.2)

A database table T1 has 2000 records and occupies 80 disk blocks. Another table T2 has 400 records and occupies 20 disk blocks. These two tables have to be joined as per a specified join condition that needs to be evaluated for every pair of records from these two tables. The memory buffer space available can hold exactly one block of records for T1 and one block of records for T2 simultaneously at any point in time. No index is available on either table.

82.1 If Nested-loop join algorithm is employed to perform the join, with the most appropriate choice of table to be used in outer loop, the number of block accesses; required for reading the data are

(a) 800000

(b) 40080

(c) 32020

(d) 100

82.2 If, instead of Nested-loop join, Block nested-loop join is used, again with the most appropriate choice of table in the outer loop, the reduction in

number of block accesses required for reading the data will be

(a) 0
(b) 30400
(c) 38400
(d) 798400

Statement for Linked Answer Question (83.1 and 83.2)

Consider the context-free grammar

$E \rightarrow E + E$

$E \rightarrow (E * E)$

$E \rightarrow id$

where E is the starting symbol, the set of terminals is {id, (, +,), *}, and the set of non-terminals is {E}.

83.1 Which of the following terminal strings has more than one parse tree when parsed according to the above grammar?

(a) id + id + id + id
(b) id + (id * (id * id))
(c) (id * (id * id)) + id
(d) ((id * id + id) * id)

83.2 For the terminal string with more than one parse tree obtained as solution to Question 83.1, how many parse trees are possible?

(a) 5
(b) 4
(c) 3
(d) 2

Statement for Linked Answer Question (84.1 and 84.2)

A sink in a directed graph is a vertex i such that there is an edge from every vertex j ≠ i to i and there is no edge from i to any other vertex. A directed graph G with n vertices is represented by its adjacency matrix A, where A[i][j] = 1 if there is an edge directed from vertex i to j and 0 otherwise. The following algorithm determines whether there is a sink in the graph G.

```
i = 0;
do {
    j = i + 1;
    while ((j < n) && E1) j++;
    if (j < n) E2;
} while (j < n)
flag = 1;
for (j = 0; j < n; j++)
if ((j != i) && E3) flag = 0;
if (flag) printf("Sink exists") else printf("Sink does not exist");
```

84.1 Choose the correct expressions for E_1 and E_2

(a) E_1 : A[i][j] and E_2 : i = j;
(b) E_1 : !A[i][j] and E_2 : i = j + 1;
(c) E_1 : !A[i][j] and E_2 : i = j;
(d) E_1 : A[i][j] and E_2 : i = j + 1

84.2 Choose the correct expression for E_3

(a) (A[i][j] && !A[j][i])
(b) (!A[i][j] && A[j][i])
(c) (!A[i][j] || A[j][i])
(d) (A[i][j] || !A[j][i])

Statement for Linked Answer Question (85.1 and 85.2)

Consider a simple graph with unit edge costs. Each node m the graph represents a router. Each node maintains a touting table indicating the next hop muter to be used to relay a packet to its destination and the cost of the path to the destination through that router. Initially, the muting table is empty. The routing table is synchronously updated as follows. In each update ion interval three tasks arc performed.

(I) A node determines whether its neighbours in the graph are accessible. If so, it sets the tentative cost to each accessible neighbour as 1. Otherwise, the cost is set to ∞.

(II) From each accessible neighbour, it gets the costs to relay to other nodes via that neighbour (as the next hop).

(III) Each node updates its routing table based on the information received in the previous two steps by choosing the minimum cost.

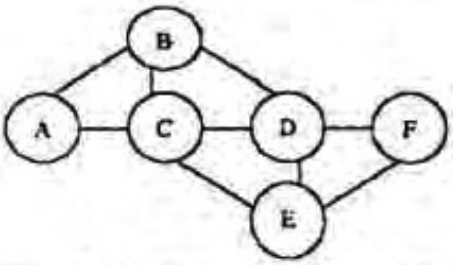

85.1 For the graph given above, possible routing tables for various nodes after they have stabilized, are shown in the following options. Identify the correct table.

(a) Table for node A

A	-	-
B	B	1
C	C	1
D	B	3
E	C	3
F	C	4

(b) Table for node C

A	A	1
B	B	1
C	-	-
D	D	1
E	E	1
F	E	3

(c) Table for node B

A	A	1
B	-	-
C	C	1
D	D	1
E	C	2
F	D	2

(d) Table for node D

A	B	3
B	B	1
C	C	1
D	-	-
E	E	1
F	F	1

85.2 Continuing from the earlier problem, suppose at some time t, when the costs have stabilized, node A goes down. The cost from node F to node A at time (t+100) is :

(a) >100 but finite

(b) ∞

(c) 3

(d) > 3 and ≤ 100