

CS : COMPUTER SCIENCE AND ENGINEERING*Duration : Three Hours**Maximum Marks :150***Read the following instructions carefully**

1. This question paper contains **24** printed pages including pages for rough work. Please check all pages and report discrepancy, if any.
2. Write your registration number, your name and name of the examination centre at the specified locations on the right half of the ORS.
3. Using HB pencil, darken the appropriate bubble under each digit of your registration number and the letters corresponding to your paper code.
4. All the questions in this question paper are of objective type.
5. Questions must be answered on **Objective Response Sheet (ORS)** by darkening the appropriate bubble (marked A, B, C, D) using HB pencil against the question number on the left hand side of the ORS. **Each question has only one correct answer.** In case you wish to change an answer, erase the old answer completely. More than one answer bubbled against a question will be treated as a wrong answer.
6. Questions 1 through 20 are 1-mark questions and questions 21 through 85 are 2-mark questions.
7. Questions 71 through 73 is one set of common data questions, questions 74 and 75 is another pair of common data questions. The question pairs (76, 77), (78, 79), (80, 81), (82, 83) and (84, 85) are questions with linked answers. The answer to the second question of the above pairs will depend on the answer to the first question of the pair. If the first question in the linked pair is wrongly answered or is un-attempted, then the answer to the second question in the pair will not be evaluated.
8. Un-attempted questions will carry zero marks.
9. **NEGATIVE MARKING:** For Q.1 to Q.20, **0.25** mark will be deducted for each wrong answer. For Q.21 to Q.75, **0.5** mark will be deducted for each wrong answer. For the pairs of questions with linked answers, there will be negative marks only for wrong answer to the first question, i.e. for Q.76, Q.78, Q.80, Q.82 and Q.84, **0.5** mark will be deducted for each wrong answer. There is no negative marking for Q.77, Q.79, Q.81, Q.83 and Q.85.
10. Calculator **without data connectivity** is allowed in the examination hall.
11. Charts, graph sheets and tables are NOT allowed in the examination hall.
12. Rough work can be done on the question paper itself. Additional blank pages are given at the end of the question paper for rough work.

Q. 1 – Q. 20 carry one mark each.

Q.1 $\lim_{x \rightarrow \infty} \frac{x - \sin x}{x + \cos x}$ equals

- (A) 1 (B) -1 (C) ∞ (D) $-\infty$

Q.2 If P, Q, R are subsets of the universal set U, then $(P \cap Q \cap R) \cup (P^c \cap Q \cap R) \cup Q^c \cup R^c$ is

- (A) $Q^c \cup R^c$ (B) $P \cup Q^c \cup R^c$ (C) $P^c \cup Q^c \cup R^c$ (D) U

Q.3 The following system of equations

$$x_1 + x_2 + 2x_3 = 1$$

$$x_1 + 2x_2 + 3x_3 = 2$$

$$x_1 + 4x_2 + \alpha x_3 = 4$$

has a unique solution. The only possible value(s) for α is/are

- (A) 0 (B) either 0 or 1 (C) one of 0, 1 or -1 (D) any real number

Q.4 In the IEEE floating point representation the hexadecimal value 0x00000000 corresponds to

- (A) the normalized value 2^{-127}
 (B) the normalized value 2^{-126}
 (C) the normalized value +0
 (D) the special value +0

Q.5 In the Karnaugh map shown below, X denotes a don't care term. What is the minimal form of the function represented by the Karnaugh map?

cd \ ab				
	00	01	11	10
00	1	1		1
01	X			
11	X			
10	1	1		X

- (A) $\overline{b}\overline{d} + \overline{a}\overline{d}$ (B) $\overline{a}\overline{b} + \overline{b}\overline{d} + \overline{a}\overline{b}\overline{d}$
 (C) $\overline{b}\overline{d} + \overline{a}\overline{b}\overline{d}$ (D) $\overline{a}\overline{b} + \overline{b}\overline{d} + \overline{a}\overline{d}$


Q.6 Let r denote number system radix. The only value(s) of r that satisfy the equation $\sqrt{121_r} = 11_r$ is/are

- (A) decimal 10 (B) decimal 11
 (C) decimal 10 and 11 (D) any value > 2

Q.7 The most efficient algorithm for finding the number of connected components in an undirected graph on n vertices and m edges has time complexity

- (A) $\Theta(n)$ (B) $\Theta(m)$ (C) $\Theta(m+n)$ (D) $\Theta(mn)$

Q.8 Given f_1 , f_3 and f in canonical sum of products form (in decimal) for the circuit



$$f_1 = \sum m(4, 5, 6, 7, 8)$$

$$f_3 = \sum m(1, 6, 15)$$

$$f = \sum m(1, 6, 8, 15)$$

then f_2 is

(A) $\sum m(4, 6)$

(B) $\sum m(4, 8)$

(C) $\sum m(6, 8)$

(D) $\sum m(4, 6, 8)$

Q.9 Which of the following is true for the language $\{a^p \mid p \text{ is a prime}\}$?

(A) It is not accepted by a Turing Machine

(B) It is regular but not context-free

(C) It is context-free but not regular

(D) It is neither regular nor context-free, but accepted by a Turing machine

Q.10 Which of the following are decidable?

I. Whether the intersection of two regular languages is infinite

II. Whether a given context-free language is regular

III. Whether two push-down automata accept the same language

IV. Whether a given grammar is context-free

(A) I and II

(B) I and IV

(C) II and III

(D) II and IV

Q.11 Which of the following describes a handle (as applicable to LR-parsing) appropriately?

(A) It is the position in a sentential form where the next shift or reduce operation will occur

(B) It is a non-terminal whose production will be used for reduction in the next step

(C) It is a production that may be used for reduction in a future step along with a position in the sentential form where the next shift or reduce operation will occur

(D) It is the production p that will be used for reduction in the next step along with a position in the sentential form where the right hand side of the production may be found

Q.12 Some code optimizations are carried out on the intermediate code because

(A) they enhance the portability of the compiler to other target processors

(B) program analysis is more accurate on intermediate code than on machine code

(C) the information from dataflow analysis cannot otherwise be used for optimization

(D) the information from the front end cannot otherwise be used for optimization

Q.13 If L and $\bar{L}$ are recursively enumerable then L is

(A) regular

(B) context-free

(C) context-sensitive

(D) recursive

Q.14 What is the maximum size of data that the application layer can pass on to the TCP layer below?

- (A) Any size (B) 2^{16} bytes – size of TCP header
(C) 2^{16} bytes (D) 1500 bytes

Q.15 Which of the following tuple relational calculus expression(s) is/are equivalent to $\forall t \in r(P(t))$?

- I. $\neg \exists t \in r(P(t))$
II. $\exists t \notin r(P(t))$
III. $\neg \exists t \in r(\neg P(t))$
IV. $\exists t \notin r(\neg P(t))$

- (A) I only (B) II only (C) III only (D) III and IV only

Q.16 A clustering index is defined on the fields which are of type

- (A) non-key and ordering (B) non-key and non-ordering
(C) key and ordering (D) key and non-ordering

Q.17 Which of the following system calls results in the sending of SYN packets?


- (A) socket (B) bind (C) listen (D) connect

Q.18 Which combination of the integer variables x , y and z makes the variable a get the value 4 in the following expression?

$$a = (x > y) ? ((x > z) ? x : z) : ((y > z) ? y : z)$$

- (A) $x = 3, y = 4, z = 2$ (B) $x = 6, y = 5, z = 3$
(C) $x = 6, y = 3, z = 5$ (D) $x = 5, y = 4, z = 5$

Q.19 The Breadth First Search algorithm has been implemented using the queue data structure. One possible order of visiting the nodes of the following graph is



- (A) MNOPQR (B) NQMPOR (C) QMNPOR (D) QMNPOR

Q.20 The data blocks of a very large file in the Unix file system are allocated using

- (A) contiguous allocation (B) linked allocation
(C) indexed allocation (D) an extension of indexed allocation

Q. 21 to Q.75 carry two marks each.


- Q.21 The minimum number of equal length subintervals needed to approximate $\int_1^2 xe^x dx$ to an accuracy of at least $\frac{1}{3} \times 10^{-6}$ using the trapezoidal rule is
- (A) $1000e$ (B) 1000 (C) $100e$ (D) 100
- Q.22 The Newton-Raphson iteration $x_{n+1} = \frac{1}{2} \left(x_n + \frac{R}{x_n} \right)$ can be used to compute the
- (A) square of R (B) reciprocal of R (C) square root of R (D) logarithm of R
- Q.23 Which of the following statements is true for every planar graph on n vertices?
- (A) The graph is connected
 (B) The graph is Eulerian
 (C) The graph has a vertex-cover of size at most $3n/4$
 (D) The graph has an independent set of size at least $n/3$
- Q.24 Let $P = \sum_{\substack{1 \leq i \leq 2k \\ i \text{ odd}}} i$ and $Q = \sum_{\substack{1 \leq i \leq 2k \\ i \text{ even}}} i$, where k is a positive integer. Then
- (A) $P = Q - k$ (B) $P = Q + k$ (C) $P = Q$ (D) $P = Q + 2k$
- Q.25 A point on a curve is said to be an extremum if it is a local minimum or a local maximum. The number of distinct extrema for the curve $3x^4 - 16x^3 + 24x^2 + 37$ is
- (A) 0 (B) 1 (C) 2 (D) 3
- Q.26 If P, Q, R are Boolean variables, then $(P + \bar{Q})(P\bar{Q} + P.R)(\bar{P}.\bar{R} + \bar{Q})$ simplifies to
- (A) $P.\bar{Q}$ (B) $P.\bar{R}$ (C) $P.\bar{Q} + R$ (D) $P.\bar{R} + Q$
- Q.27 Aishwarya studies either computer science or mathematics everyday. If she studies computer science on a day, then the probability that she studies mathematics the next day is 0.6. If she studies mathematics on a day, then the probability that she studies computer science the next day is 0.4. Given that Aishwarya studies computer science on Monday, what is the probability that she studies computer science on Wednesday?
- (A) 0.24 (B) 0.36 (C) 0.4 (D) 0.6
- Q.28 How many of the following matrices have an eigenvalue 1?
- $\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$ and $\begin{bmatrix} -1 & 0 \\ 1 & -1 \end{bmatrix}$
- (A) one (B) two (C) three (D) four

- Q.29 Let X be a random variable following normal distribution with mean +1 and variance 4. Let Y be another normal variable with mean -1 and variance unknown. If $P(X \leq -1) = P(Y \geq 2)$, the standard deviation of Y is
- (A) 3 (B) 2 (C) $\sqrt{2}$ (D) 1
- Q.30 Let fsa and pda be two predicates such that $fsa(x)$ means x is a finite state automaton, and $pda(y)$ means that y is a pushdown automaton. Let $equivalent$ be another predicate such that $equivalent(a, b)$ means a and b are equivalent. Which of the following first order logic statements represents the following:
- Each finite state automaton has an equivalent pushdown automaton.
- (A) $(\forall x fsa(x)) \Rightarrow (\exists y pda(y) \wedge equivalent(x, y))$
 (B) $\sim \forall y (\exists x fsa(x) \Rightarrow pda(y) \wedge equivalent(x, y))$
 (C) $\forall x \exists y (fsa(x) \wedge pda(y) \wedge equivalent(x, y))$
 (D) $\forall x \exists y (fsa(y) \wedge pda(x) \wedge equivalent(x, y))$
- Q.31 P and Q are two propositions. Which of the following logical expressions are equivalent?
- I. $P \vee \sim Q$
 II. $\sim (\sim P \wedge Q)$
 III. $(P \wedge Q) \vee (P \wedge \sim Q) \vee (\sim P \wedge \sim Q)$
 IV. $(P \wedge Q) \vee (P \wedge \sim Q) \vee (\sim P \wedge Q)$
- (A) Only I and II (B) Only I, II and III
 (C) Only I, II and IV (D) All of I, II, III and IV
- Q.32 For a magnetic disk with concentric circular tracks, the seek latency is not linearly proportional to the seek distance due to
- (A) non-uniform distribution of requests
 (B) arm starting and stopping inertia
 (C) higher capacity of tracks on the periphery of the platter
 (D) use of unfair arm scheduling policies
- Q.33 Which of the following is/are true of the auto-increment addressing mode?
- I. It is useful in creating self-relocating code
 II. If it is included in an Instruction Set Architecture, then an additional ALU is required for effective address calculation
 III. The amount of increment depends on the size of the data item accessed
- (A) I only (B) II only (C) III only (D) II and III only
- Q.34 Which of the following must be true for the RFE (Return from Exception) instruction on a general purpose processor?
- I. It must be a trap instruction
 II. It must be a privileged instruction
 III. An exception cannot be allowed to occur during execution of an RFE instruction
- (A) I only (B) II only (C) I and II only (D) I, II and III only
- Q.35 For inclusion to hold between two cache levels L1 and L2 in a multi-level cache hierarchy, which of the following are necessary?
- I. L1 must be a write-through cache
 II. L2 must be a write-through cache
 III. The associativity of L2 must be greater than that of L1
 IV. The L2 cache must be at least as large as the L1 cache
- (A) IV only (B) I and IV only (C) I, III and IV only (D) I, II, III and IV

- Q.36 Which of the following are NOT true in a pipelined processor?
- Bypassing can handle all RAW hazards.
 - Register renaming can eliminate all register carried WAR hazards.
 - Control hazard penalties can be eliminated by dynamic branch prediction.
- (A) I and II only (B) I and III only (C) II and III only (D) I, II and III
- Q.37 The use of multiple register windows with overlap causes a reduction in the number of memory accesses for
- function locals and parameters
 - register saves and restores
 - instruction fetches
- (A) I only (B) II only (C) III only (D) I, II and III
- Q.38 In an instruction execution pipeline, the earliest that the data TLB (Translation Lookaside Buffer) can be accessed is
- before effective address calculation has started
 - during effective address calculation
 - after effective address calculation has completed
 - after data cache lookup has completed
- Q.39 Consider the following functions:
- $$f(n) = 2^n$$
- $$g(n) = n!$$
- $$h(n) = n^{\log n}$$
- Which of the following statements about the asymptotic behaviour of $f(n)$, $g(n)$, and $h(n)$ is true?
- $f(n) = O(g(n)); g(n) = O(h(n))$
 - $f(n) = \Omega(g(n)); g(n) = O(h(n))$
 - $g(n) = O(f(n)); h(n) = O(f(n))$
 - $h(n) = O(f(n)); g(n) = \Omega(f(n))$
- Q.40 The minimum number of comparisons required to determine if an integer appears more than $n/2$ times in a sorted array of n integers is
- (A) $\Theta(n)$ (B) $\Theta(\log n)$ (C) $\Theta(\log^* n)$ (D) $\Theta(1)$
- Q.41 A B-tree of order 4 is built from scratch by 10 successive insertions. What is the maximum number of node splitting operations that may take place?
- (A) 3 (B) 4 (C) 5 (D) 6
- Q.42 G is a graph on n vertices and $2n-2$ edges. The edges of G can be partitioned into two edge-disjoint spanning trees. Which of the following is NOT true for G ?
- For every subset of k vertices, the induced subgraph has at most $2k-2$ edges
 - The minimum cut in G has at least two edges
 - There are two edge-disjoint paths between every pair of vertices
 - There are two vertex-disjoint paths between every pair of vertices
- Q.43 Consider the Quicksort algorithm. Suppose there is a procedure for finding a pivot element which splits the list into two sub-lists each of which contains at least one-fifth of the elements. Let $T(n)$ be the number of comparisons required to sort n elements. Then
- $T(n) \leq 2T(n/5) + n$
 - $T(n) \leq T(n/5) + T(4n/5) + n$
 - $T(n) \leq 2T(4n/5) + n$
 - $T(n) \leq 2T(n/2) + n$

- Q.44 The subset-sum problem is defined as follows: Given a set S of n positive integers and a positive integer W , determine whether there is a subset of S whose elements sum to W . An algorithm Q solves this problem in $O(nW)$ time. Which of the following statements is false?
- (A) Q solves the subset-sum problem in polynomial time when the input is encoded in unary
 (B) Q solves the subset-sum problem in polynomial time when the input is encoded in binary
 (C) The subset sum problem belongs to the class NP
 (D) The subset sum problem is NP-hard

Q.45



Dijkstra's single source shortest path algorithm when run from vertex a in the above graph, computes the correct shortest path distance to

- (A) only vertex a
 (B) only vertices a, e, f, g, h
 (C) only vertices a, b, c, d
 (D) all the vertices
- Q.46 You are given the postorder traversal, P , of a binary search tree on the n elements $1, 2, \dots, n$. You have to determine the unique binary search tree that has P as its postorder traversal. What is the time complexity of the most efficient algorithm for doing this?
- (A) $\Theta(\log n)$
 (B) $\Theta(n)$
 (C) $\Theta(n \log n)$
 (D) none of the above, as the tree cannot be uniquely determined.

- Q.47 We have a binary heap on n elements and wish to insert n more elements (not necessarily one after another) into this heap. The total time required for this is

- (A) $\Theta(\log n)$ (B) $\Theta(n)$ (C) $\Theta(n \log n)$ (D) $\Theta(n^2)$

- Q.48 Which of the following statements is false?

- (A) Every NFA can be converted to an equivalent DFA
 (B) Every non-deterministic Turing machine can be converted to an equivalent deterministic Turing machine
 (C) Every regular language is also a context-free language
 (D) Every subset of a recursively enumerable set is recursive

Q.49 Given below are two finite state automata ($\rightarrow$ indicates the start state and F indicates a final state)

Y:

	a	b
$\rightarrow 1$	1	2
2 (F)	2	1

Z:

	a	b
$\rightarrow 1$	2	2
2 (F)	1	1

Which of the following represents the product automaton $Z \times Y$?

(A)

	a	b
$\rightarrow P$	S	R
Q	R	S
R (F)	Q	P
S	Q	P

(B)

	a	b
$\rightarrow P$	S	Q
Q	R	S
R (F)	Q	P
S	P	Q

(C)

	a	b
$\rightarrow P$	Q	S
Q	R	S
R (F)	Q	P
S	Q	P

(D)

	a	b
$\rightarrow P$	S	Q
Q	S	R
R (F)	Q	P
S	Q	P

Q.50 Which of the following statements are true?

- I. Every left-recursive grammar can be converted to a right-recursive grammar and vice-versa
- II. All ϵ -productions can be removed from any context-free grammar by suitable transformations
- III. The language generated by a context-free grammar all of whose productions are of the form $X \rightarrow w$ or $X \rightarrow wY$ (where, w is a string of terminals and Y is a non-terminal), is always regular
- IV. The derivation trees of strings generated by a context-free grammar in Chomsky Normal Form are always binary trees

(A) I, II, III and IV

(C) I, III and IV only

(B) II, III and IV only

(D) I, II and IV only

Q.51 Match the following:

(E) Checking that identifiers are declared before their use	(P) $L = \{a^n b^m c^n d^m \mid n \geq 1, m \geq 1\}$
(F) Number of formal parameters in the declaration of a function agrees with the number of actual parameters in a use of that function	(Q) $X \rightarrow X b X \mid X c X \mid d X f \mid g$
(G) Arithmetic expressions with matched pairs of parentheses	(R) $L = \{wcw \mid w \in (a \mid b)^*\}$
(H) Palindromes	(S) $X \rightarrow b X b \mid c X c \mid \epsilon$


(A) E - P, F - R, G - Q, H - S

(C) E - R, F - P, G - Q, H - S

(B) E - R, F - P, G - S, H - Q

(D) E - P, F - R, G - S, H - Q

Q.52 Match the following NFAs with the regular expressions they correspond to.



1. $\epsilon + 0(01^*1+00)^*01^*$
2. $\epsilon + 0(10^*1+00)^*0$
3. $\epsilon + 0(10^*1+10)^*1$
4. $\epsilon + 0(10^*1+10)^*10^*$

(A) P - 2, Q - 1, R - 3, S - 4

(B) P - 1, Q - 3, R - 2, S - 4

(C) P - 1, Q - 2, R - 3, S - 4

(D) P - 3, Q - 2, R - 1, S - 4

Q.53 Which of the following are regular sets?

- I. $\{a^n b^{2m} \mid n \geq 0, m \geq 0\}$
- II. $\{a^n b^m \mid n = 2m\}$
- III. $\{a^n b^m \mid n \neq m\}$
- IV. $\{xycy \mid x, y \in \{a, b\}^*\}$

(A) I and IV only

(B) I and III only

(C) I only

(D) IV only

Q.54 Which of the following are true?

- I. A programming language which does not permit global variables of any kind and has no nesting of procedures/functions, but permits recursion can be implemented with static storage allocation
- II. Multi-level access link (or display) arrangement is needed to arrange activation records only if the programming language being implemented has nesting of procedures/functions
- III. Recursion in programming languages cannot be implemented with dynamic storage allocation
- IV. Nesting of procedures/functions and recursion require a dynamic heap allocation scheme and cannot be implemented with a stack-based allocation scheme for activation records
- V. Programming languages which permit a function to return a function as its result cannot be implemented with a stack-based storage allocation scheme for activation records

(A) II and V only (B) I, III and IV only (C) I, II and V only (D) II, III and V only

Q.55 An LALR(1) parser for a grammar G can have shift-reduce (S-R) conflicts if and only if

- (A) the SLR(1) parser for G has S-R conflicts
- (B) the LR(1) parser for G has S-R conflicts
- (C) the LR(0) parser for G has S-R conflicts
- (D) the LALR(1) parser for G has reduce-reduce conflicts

Q.56 In the slow start phase of the TCP congestion control algorithm, the size of the congestion window

- (A) does not increase (B) increases linearly
- (C) increases quadratically (D) increases exponentially

Q.57 If a class B network on the Internet has a subnet mask of 255.255.248.0, what is the maximum number of hosts per subnet?

- (A) 1022 (B) 1023 (C) 2046 (D) 2047

Q.58 A computer on a 10Mbps network is regulated by a token bucket. The token bucket is filled at a rate of 2Mbps. It is initially filled to capacity with 16 Megabits. What is the maximum duration for which the computer can transmit at the full 10Mbps?

- (A) 1.6 seconds (B) 2 seconds (C) 5 seconds (D) 8 seconds

Q.59 A client process P needs to make a TCP connection to a server process S. Consider the following situation: the server process S executes a `socket()`, a `bind()` and a `listen()` system call in that order, following which it is preempted. Subsequently, the client process P executes a `socket()` system call followed by `connect()` system call to connect to the server process S. The server process has not executed any `accept()` system call. Which one of the following events could take place?

- (A) `connect()` system call returns successfully
- (B) `connect()` system call blocks
- (C) `connect()` system call returns an error
- (D) `connect()` system call results in a core dump

Q.60 What is printed by the following C program?

```
int f(int x, int *py, int **ppz)
{
    int y, z;
    **ppz += 1; z = *ppz;
    *py += 2; y = *py;
    x += 3;
    return x+y+z;
}
```

```
void main()
{
    int c, *b, **a;
    c = 4; b = &c; a = &b;
    printf("%d", f(c, b, a));
}
```

(A) 18

(B) 19

(C) 21

(D) 22

Q.61 Choose the correct option to fill ?1 and ?2 so that the program below prints an input string in reverse order. Assume that the input string is terminated by a newline character.

```
void reverse(void) {
    int c;
    if(?1) reverse();
    ?2
}

main() {
    printf("Enter Text"); printf("\n");
    reverse(); printf("\n");
}
```

(A) ?1 is (getchar() != '\n')
?2 is getchar(c);

(B) ?1 is (c = getchar()) != '\n'
?2 is getchar(c);

(C) ?1 is (c != '\n')
?2 is putchar(c);

(D) ?1 is ((c = getchar()) != '\n')
?2 is putchar(c);

Q.62 The following C function takes a singly-linked list of integers as a parameter and rearranges the elements of the list. The function is called with the list containing the integers 1,2,3,4,5,6,7 in the given order. What will be 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

- Q.63 The P and V operations on counting semaphores, where s is a counting semaphore, are defined as follows:

P(s): $s = s - 1$;
if $s < 0$ then wait;

V(s): $s = s + 1$;
if $s \leq 0$ then wakeup a process waiting on s ;

Assume that P_b and V_b , the wait and signal operations on binary semaphores are provided. Two binary semaphores x_b and y_b are used to implement the semaphore operations P(s) and V(s) as follows:

```
P(s):  P_b(x_b);
        s = s - 1;
        if (s < 0) {
            V_b(x_b);
            P_b(y_b);
        }
        else V_b(x_b);
```

```
V(s):  P_b(x_b);
        s = s + 1;
        if (s <= 0) V_b(y_b);
        V_b(x_b);
```

The initial values of x_b and y_b are respectively

- (A) 0 and 0 (B) 0 and 1 (C) 1 and 0 (D) 1 and 1
- Q.64 Which of the following statements about synchronous and asynchronous I/O is NOT true?
- (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 asynchronous 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

- Q.65 Which of the following is NOT true of deadlock prevention and deadlock avoidance schemes?

- (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*

- Q.66 A process executes the following code

```
for(i=0; i<n; i++) fork();
```

The total number of child processes created is

- (A) n (B) $2^n - 1$ (C) 2^n (D) $2^{n+1} - 1$

Q.67 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. A three 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, and
- 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

- (A) 20, 20 and 20 (B) 24, 24 and 24 (C) 24, 24 and 20 (D) 25, 25 and 24

Q.68 Let R and S be two relations with the following schema

R(P, Q, R1, R2, R3)

S(P, Q, S1, S2)

where {P,Q} is the key for both schemas. Which of the following queries are equivalent?

- I. $\Pi_P(R \bowtie S)$
- II. $\Pi_P(R) \bowtie \Pi_P(S)$
- III. $\Pi_P(\Pi_{P,Q}(R) \cap \Pi_{P,Q}(S))$
- IV. $\Pi_P(\Pi_{P,Q}(R) - (\Pi_{P,Q}(R) - \Pi_{P,Q}(S)))$

- (A) Only I and II (B) Only I and III
(C) Only I, II and III (D) Only I, III and IV

Q.69 Consider the following relational schemes for a library database:

Book (Title, Author, Catalog_no, Publisher, Year, Price)

Collection (Title, Author, Catalog_no)

with the following functional dependencies:

- I. Title Author $\rightarrow$ Catalog_no
- II. Catalog_no $\rightarrow$ Title Author Publisher Year
- III. Publisher Title Year $\rightarrow$ Price

Assume {Author, Title} is the key for both schemes. Which of the following statements is true?

- (A) Both Book and Collection are in BCNF
(B) Both Book and Collection are in 3NF only
(C) Book is in 2NF and Collection is in 3NF
(D) Both Book and Collection are in 2NF only

Q.70 Consider a file of 16384 records. Each record is 32 bytes long and its key field is of size 6 bytes. The file is ordered on a non-key field, and the file organization is unspanned. The file is stored in a file system with block size 1024 bytes, and the size of a block pointer is 10 bytes. If the secondary index is built on the key field of the file, and a multi-level index scheme is used to store the secondary index, the number of first-level and second-level blocks in the multi-level index are respectively

- (A) 8 and 0 (B) 128 and 6 (C) 256 and 4 (D) 512 and 5

Common Data Questions

Common Data for Questions 71,72 and 73:

Consider a machine with a 2-way set associative data cache of size 64 Kbytes and block size 16 bytes. The cache is managed using 32 bit virtual addresses and the page size is 4 Kbytes. A program to be run on this machine begins as follows:

```
double ARR [1024] [1024];
int i, j ;

/* Initialize array ARR to 0.0 */
for (i = 0; i < 1024; i++)
    for (j = 0; j < 1024; j++)
        ARR [i] [j] = 0.0;
```

The size of double is 8Bytes. Array ARR is located in memory starting at the beginning of virtual page 0xFF000 and stored in row major order. The cache is initially empty and no pre-fetching is done. The only data memory references made by the program are those to array ARR.

- Q.71 The total size of the tags in the cache directory is
 (A) 32 Kbits (B) 34 Kbits (C) 64 Kbits (D) 68 Kbits
- Q.72 Which of the following array elements has the same cache index as ARR [0] [0]?
 (A) ARR [0] [4] (B) ARR [4] [0] (C) ARR [0] [5] (D) ARR [5] [0]
- Q.73 The cache hit ratio for this initialization loop is
 (A) 0% (B) 25% (C) 50% (D) 75%

Common Data for Questions 74 and 75:

Consider the following C functions:

```
int f1 ( int n )
{
    if (n == 0 || n == 1)
        return n;
    else
        return (2*f1(n-1) + 3*f1(n-2));
}
```

```
int f2 ( int n )
{
    int i;
    int X[N], Y[N], Z[N];

    X[0] = Y[0] = Z[0] = 0;
    X[1] = 1; Y[1] = 2; Z[1] = 3;
    for (i = 2; i <= n; i++){
        X[i] = Y[i-1] + Z[i-2];
        Y[i] = 2 * X[i];
        Z[i] = 3 * X[i];
    }
    return X[n];
}
```

Q.74 The running time of $f1(n)$ and $f2(n)$ are

- (A) $\Theta(n)$ and $\Theta(n)$ (B) $\Theta(2^n)$ and $\Theta(n)$
 (C) $\Theta(n)$ and $\Theta(2^n)$ (D) $\Theta(2^n)$ and $\Theta(2^n)$

Q.75 $f1(8)$ and $f2(8)$ return the values

- (A) 1661 and 1640 (B) 59 and 59 (C) 1640 and 1640 (D) 1640 and 1661

Linked Answer Questions: Q.76 to Q.85 carry two marks each.**Statement for Linked Answer Questions 76 and 77:**

Delayed branching can help in the handling of control hazards.

Q.76 For all delayed conditional branch instructions, irrespective of whether the condition evaluates to true or false,

- (A) the instruction following the conditional branch instruction in memory is executed
 (B) the first instruction in the fall through path is executed
 (C) the first instruction in the taken path is executed
 (D) the branch takes longer to execute than any other instruction

Q.77 The following code is to run on a pipelined processor with one branch delay slot:

```

I1: ADD R2 ← R7 + R8
I2: SUB R4 ← R5 - R6
I3: ADD R1 ← R2 + R3
I4: STORE Memory[R4] ← R1
    BRANCH to Label if R1 == 0
  
```

Which of the instructions I1, I2, I3 or I4 can legitimately occupy the delay slot without any other program modification?

- (A) I1 (B) I2 (C) I3 (D) I4

Statement for Linked Answer Questions 78 and 79:

Let x_n denote the number of binary strings of length n that contain no consecutive 0s.

Q.78 Which of the following recurrences does x_n satisfy?

- (A) $x_n = 2x_{n-1}$ (B) $x_n = x_{\lfloor n/2 \rfloor} + 1$ (C) $x_n = x_{\lfloor n/2 \rfloor} + n$ (D) $x_n = x_{n-1} + x_{n-2}$

Q.79 The value of x_5 is

- (A) 5 (B) 7 (C) 8 (D) 16

Statement for Linked Answer Questions 80 and 81:

The subset-sum problem is defined as follows. Given a set of n positive integers, $S = \{a_1, a_2, a_3, \dots, a_n\}$, and a positive integer W , is there a subset of S whose elements sum to W ? A dynamic program for solving this problem uses a 2-dimensional Boolean array, X , with n rows and $W+1$ columns. $X[i, j]$, $1 \leq i \leq n$, $0 \leq j \leq W$, is TRUE if and only if there is a subset of $\{a_1, a_2, \dots, a_i\}$ whose elements sum to j .

Q.80 Which of the following is valid for $2 \leq i \leq n$ and $a_i \leq j \leq W$?


- (A) $X[i, j] = X[i-1, j] \vee X[i, j-a_i]$
 (B) $X[i, j] = X[i-1, j] \vee X[i-1, j-a_i]$
 (C) $X[i, j] = X[i-1, j] \wedge X[i, j-a_i]$
 (D) $X[i, j] = X[i-1, j] \wedge X[i-1, j-a_i]$

Q.81 Which entry of the array X , if TRUE, implies that there is a subset whose elements sum to W ?

- (A) $X[1, W]$ (B) $X[n, 0]$ (C) $X[n, W]$ (D) $X[n-1, n]$

Statement for Linked Answer Questions 82 and 83:

Consider the following ER diagram



Q.82 The minimum number of tables needed to represent M,N,P,R1,R2 is

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

Q.83 Which of the following is a correct attribute set for one of the tables for the correct answer to the above question?

- (A) {M1, M2, M3, P1} (B) {M1, P1, N1, N2}
(C) {M1, P1, N1} (D) {M1, P1}

Statement for Linked Answer Questions 84 and 85:

Consider the following C program that attempts to locate an element x in an array $Y[]$ using binary search. The program is erroneous.

```

1.  f(int Y[10], int x){
2.      int i, j, k;
3.      i = 0; j = 9;
4.      do {
5.          k = (i+j)/2;
6.          if (Y[k] < x) i = k; else j = k;
7.      } while ((Y[k] != x) && (i < j));
8.      if (Y[k] == x) printf("x is in the array");
9.      else printf("x is not in the array");
10. }
  
```

Q.84 On which of the following contents of Y and x does the program fail?

- (A) Y is [1 2 3 4 5 6 7 8 9 10] and $x < 10$
(B) Y is [1 3 5 7 9 11 13 15 17 19] and $x < 1$
(C) Y is [2 2 2 2 2 2 2 2 2 2] and $x > 2$
(D) Y is [2 4 6 8 10 12 14 16 18 20] and $2 < x < 20$ and x is even

Q.85 The correction needed in the program to make it work properly is

- (A) change line 6 to: if (Y[k] < x) i = k+1; else j = k-1;
(B) change line 6 to: if (Y[k] < x) i = k-1; else j = k+1;
(C) change line 6 to: if (Y[k] <= x) i = k; else j = k;
(D) change line 7 to: } while ((Y[k] == x) && (i < j));

END OF THE QUESTION PAPER