



COMPUTER SCIENCE & ENGINEERING

ONE MARKS QUESTIONS (1-20)


1. Consider the polynomial $p(x) = a_0 + a_1x + a_2x^2 + a_3x^3$, where $a_i \neq 0, \forall i$. The minimum number of multiplications needed to evaluate p on an input x is:
- 3
 - 4
 - 6
 - 9
2. Let X, Y, Z be sets of sizes x, y and z respectively. Let $W = X \times Y$ and E be the set of all subsets of W . The number of functions from Z to E is:
- Z^{2^x}
 - $Z \times 2^{xy}$
 - $Z^{2^{xy}}$
 - 2^{xyz}
3. The set $\{1, 2, 3, 5, 7, 8, 9\}$ under multiplication modulo 10 is not a group. Given below are four plausible reasons. Which one of them is false?
- It is not closed
 - 2 does not have an inverse
 - 3 does not have an inverse
 - 8 does not have an inverse
4. A relation R is defined on ordered pairs of integers as follows:
 $(x, y) R (u, v)$ if $x < u$ and $y > v$. Then R is:
- Neither a Partial Order nor an Equivalence Relation
 - A Partial Order but not a Total Order
 - A Total Order
 - An Equivalence Relation
5. Which one of the following reasons does Internet Protocol (IP) use the time-to-live (TTL) field in the IP datagram header?
- Ensure packets reach destination within that time
 - Discard packets that reach later than that time
 - Prevent packets from looping indefinitely
 - Limit the time for which a packet gets queued in intermediate routers.
6. 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
 - 2
 - 3
 - 4
7. Consider the following grammar:
- $$S \rightarrow S * E$$
- $$S \rightarrow E$$
- $$E \rightarrow F + E$$
- $$E \rightarrow F$$
- $$F \rightarrow id$$
- Consider the following LR(0) items corresponding to the grammar above.
- $S \rightarrow S * E$
 - $E \rightarrow F + E$
 - $E \rightarrow F + E$
- Given the items above, which two of them will appear in the same set in the canonical sets-of-items for the grammar?
- 1 and 2
 - 2 and 3
 - 1 and 3
 - None of the above
8. You are given a free running clock with a duty cycle of 50% and a digital waveform f which changes only at the negative edge of the clock. Which one of the following circuits (using clocked D flip-flops) will delay the phase of f by 180°?
-



c.



d.



9. A CPU has 24-bit instructions. A program starts at address 300 (in decimal). Which one of the following is a legal program counter (all values in decimal)?

a. 400
b. 500
c. 600
d. 700

10. In a binary max heap containing n numbers, the smallest element can be found in time

a. $\theta(n)$
b. $\theta(\log n)$
c. $\theta(\log \log n)$
d. $\theta(1)$

11. Consider a weighted complete graph G on the vertex set $\{v_1, v_2, \dots, v_n\}$ such that the weight of the edge (v_i, v_j) is $2^{|i-j|}$. The weight of a minimum spanning tree of G is:

a. $n-1$
b. $2n-2$
c. $\binom{n}{2}$
d. n

12. To implement Dijkstra's shortest path algorithm on unweighted graphs so that it runs in linear time, the data structure to be used is:

a. Queue
b. Stack
c. Heap
d. B-Tree

13. A scheme for storing binary trees in an array X is as follows. Indexing of X starts

at 1 instead of 0, the root is stored at $X[1]$. For a node stored at $X[i]$, the left child, if any, is stored in $X[2i]$ and the right child, if any, in $X[2i+1]$. To be able to store any binary tree on n vertices the minimum size of X should be

a. $\log_2 n$
b. n
c. $2n+1$
d. 2^n-1

14. Which one of the following in place sorting algorithms needs the minimum number of swaps?

a. Quick sort
b. Insertion sort
c. Selection sort
d. Heap sort

15. Consider the following C-program fragment in which i , j and n are integer variables.

for ($i = n; i > 0; i = 2j + i$);

Let $\text{val}(j)$ denote the value stored in the variable j after termination of the for loop. Which one of the following is true?

a. $\text{val}(j) = \theta(\log n)$
b. $\text{val}(j) = \theta(\sqrt{n})$
c. $\text{val}(j) = \theta(n)$
d. $\text{val}(j) = \theta(n \log n)$

16. Let S be an NP-complete problem and Q and R be two other problems not known to be in NP. Q is polynomial time reducible to S and S is polynomial-time reducible to R . Which one of the following statements is true?

a. R is NP-complete
b. R is NP-hard
c. Q is NP-complete
d. Q is NP-hard

17. An element in an array X is called a leader if it is greater than all elements to the right of it in X . The best algorithm to find all leaders in an array

a. Solves it in linear time using a left to right pass of the array
b. Solves it in linear time using a right to left pass of the array
c. Solves it using divide and conquer in time $\theta(n \log n)$
d. Solves it in time $\theta(n^2)$

18. We are given a set $X = \{x_1, \dots, x_n\}$ where $x_i = 2^i$. A sample $S \subseteq X$ is drawn by selecting each x_i independently with probability $p_i = 1/2$. The expected value of the smallest number in sample S is:

- $1/n$
- 2
- \sqrt{n}
- n

19. Let $L_1 = \{0^{n+m}1^m \mid n, m \geq 0\}$, $L_2 = \{0^{n+m}1^{n+m}0^m \mid n, m \geq 0\}$ and $L_3 = \{0^{n+m}1^{n+m}0^{n+m} \mid n, m \geq 0\}$. Which of these languages are NOT context free?

- L_1 only
- L_3 only
- L_1 and L_2
- L_2 and L_3

20. Consider the following log sequence of two transactions on a bank account, with initial balance 12000, that transfer 2000 to a mortgage payment and then apply a 5% interest.

- T1 start
- T1 B old= 1200 new= 10000
- T1 M old=0 new=2000
- T1 commit
- T2start
- T2 B old= 10000 new= 10500
- T2 commit

Suppose the database system crashes just before log record 7 is written. When the system is restarted, which one statement is true of the recovery procedure?

- We must redo log record 6 to set B to 10500
- We must undo log record 6 to set B to 10000 and then redo log records 2 and 3
- We need not redo log records 2 and 3 because transaction T1 has committed
- We can apply redo and undo operations in arbitrary order because they are idempotent.

TWO MARKS QUESTIONS (21-75)

21. For each element in a set of size $2n$, an unbiased coin is tossed. The $2n$ coin tosses are independent. An element is chosen if the corresponding coin toss were head.

The probability that exactly n elements are chosen is:

- $\frac{\binom{2n}{n}}{4^n}$
- $\frac{\binom{2n}{n}}{2^n}$
- $\frac{1}{\binom{2n}{n}}$
- $\frac{1}{2}$

22. Let E, F and G be finite sets

Let $X = (E \cap F) \cup (F \cap G)$ and $Y = (E \cap (E \cap G)) \cup (E \cap F)$. Which one of the following is true?

- $X = Y$
- $X \neq Y$
- $X \cap Y = \emptyset$
- $X - Y = \emptyset$ and $Y - X = \emptyset$

23. F is an $n \times n$ real matrix, b is an $n \times 1$ real vector. Suppose there are two $n \times 1$ vectors, u and v such that $u \neq v$, and $Fu = b$, $Fv = b$. Which one of the following statements is false?

- Determinant of F is zero
- There are an infinite number of solutions to $Fx = b$
- There is an $x = 0$ such that $Fx = 0$
- F must have two identical rows

24. Given a set of elements $N = \{1, 2, \dots, n\}$ and two arbitrary subsets $A \subseteq N$ and $B \subseteq N$ how many of the $n!$ permutations π from N to N satisfy $\min(\pi(A)) = \min(\pi(B))$, where $\min(S)$ is the smallest integer in the set of integers, S and $\pi(S)$ is the set of integers obtained by applying permutation π to each elements of S ?

- $(n - |A \cup B|) |A| |B|$
- $(|A|^2 + |B|^2) n^2$
- $n! \frac{|A \cap B|}{|A \cup B|}$
- $|A \cap B|^2 / \binom{n}{|A \cup B|}$

25. Let $S = \{1, 2, 3, \dots, m\}$, $m > 3$. Let X_1, \dots, X_n be subsets of S each of size 3. Define a

function f from S to the set of natural numbers as, $f(i)$ is the number of sets X_j that contain the element i . That is $f(i) = |\{j | i \in X_j\}|$.

Then $\sum_{i=1}^m f(i)$ is

- $3m$
- $3n$
- $2m - 1$
- $2n + 1$

26. Which one of the first order predicate calculus statements given below correctly expresses the following English statement? Tigers and lions attack if they are hungry or threatened.

- $\forall x[(\text{tiger}(x) \wedge \text{lion}(x)) \rightarrow \{(\text{hungry}(x) \vee \text{threatened}(x)) \rightarrow \text{attacks}(x)\}]$
- $\forall x[(\text{tiger}(x) \vee \text{lion}(x)) \rightarrow \{(\text{hungry}(x) \vee \text{threatened}(x)) \wedge \text{attacks}(x)\}]$
- $\forall x[(\text{tiger}(x) \vee \text{lion}(x)) \rightarrow \{(\text{attacks}(x) \rightarrow (\text{hungry}(x) \vee \text{threatened}(x))\}]$
- $\forall x[(\text{tiger}(x) \vee \text{lion}(x)) \rightarrow \{(\text{hungry}(x) \vee \text{threatened}(x)) \rightarrow \text{attacks}(x)\}]$

27. Consider the following propositional statements:

P1: $((A \wedge B) \rightarrow C) \equiv ((A \rightarrow C) \wedge (B \rightarrow C))$

P2: $((A \vee B) \rightarrow C) \equiv ((A \rightarrow C) \vee (B \rightarrow C))$

Which one of the following is true?

- P1 is a tautology, but not P2
- P2 is a tautology, but not P1
- P1 and P2 are both tautologies
- Both P1 and P2 are not tautologies

28. A logical binary relations \odot , is defined as follows:

A	B	$A \odot B$
True	True	True
True	False	True
False	True	False
False	False	True

Let \sim be the unary negation (NOT) operator, with higher precedence than \odot . Which one of the following is equivalent of $A \wedge B$?

- $(\sim A \odot B)$
- $(\sim(A \odot \sim B))$
- $(\sim(\sim A \odot \sim B))$

d. $(\sim A \odot B)$

29. If s is a string over $(0 + 1)^*$ then let $n_0(s)$ denote the number of 0-s in s and $n_1(s)$ the number of 1-s in s . Which one of the following languages is not regular?

- $L = \{s \in (0 + 1)^* | n_0(s) \text{ is a 3-digit prime}\}$
- $L = \{s \in (0 + 1)^* | \text{for every prefix } p \text{ of } s, |n_0(p) - n_1(p)| \leq 2\}$
- $L = \{s \in (0 + 1)^* | n_0(s) - n_1(s) \leq 4\}$
- $L = \{s \in (0 + 1)^* | n_0(s) \bmod 7 = n_1(s) \bmod 5 = 0\}$

30. For $s \in (0 + 1)^*$ let $d(s)$ denote the decimal value of s (e.g. $d(101) = 5$). Let $L = \{s \in (0 + 1)^* | d(s) \bmod 5 = 2 \text{ and } d(s) \bmod 7 = 4\}$

Which one of the following statements is true?

- L is recursively enumerable, but not recursive
- L is recursive, but not context-free
- L is context-free, but not regular
- L is regular

31. Let SHAM₃ be the problem of finding a Hamiltonian cycle in a graph $G = (V, E)$ with $|V|$ divisible by 3 and DHAM₃ be the problem of determining if a Hamiltonian cycle exists in such graphs. Which one of the following is true?

- Both DHAM₃ and SHAM₃ are NP-hard
- SHAM₃ is NP-hard, but DHAM₃ is not
- DHAM₃ is NP-hard, but SHAM₃ is not
- Neither DHAM₃ nor SHAM₃ is NP-hard

32. Consider the following statements about the context-free grammar $G = \{S \rightarrow SS, S \rightarrow ab, S \rightarrow ba, S \rightarrow \epsilon\}$

- G is ambiguous
- G produces all strings with equal number of a's and b's
- G can be accepted by a deterministic PDA.

Which combination below expresses all the true statements about G ?

- 1 only
- 1 and 3 only
- 2 and 3 only
- 1, 2 and 3

d. $g_s(h_3h_2h_1h_0) = \Sigma (0, 1, 6, 7, 10, 11, 12, 13)$

41. A CPU has a cache with block size 64 bytes. The main memory has k banks, each bank being c bytes wide. Consecutive c -byte chunks are mapped on consecutive banks with wrap-around. All the k banks can be accessed in parallel, but two accesses to the same bank must be serialized. A cache block access may involve multiple iterations of parallel bank accesses depending on the amount of data obtained by accessing all the Ac banks in parallel. Each iteration requires decoding the bank numbers to be accessed in parallel and this takes $k/2$ ns. The latency of one bank access is 80 ns. If $c = 2$ and $k = 24$, the latency of retrieving a cache block starting at address zero from main memory is
- 92 ns
 - 104 ns
 - 172 ns
 - 184 ns
42. A CPU has a five stage pipeline and runs at 1 GHz frequency. Instruction fetch happens in the first stage of the pipeline. A conditional branch instruction computes the target address and evaluates the branch condition in the third stage of the pipeline. The processor stops fetching new instructions following a conditional branch until the branch outcome is known. A program executes 10^6 instructions out of which 20% are conditional branches. If each instruction takes one cycle to complete on average, the total execution time of the program is
- 1.0 seconds
 - 1.2 seconds
 - 1.4 seconds
 - 1.6 seconds
43. Consider a new instruction named branch-on-bit-set (mnemonic bbs). The instruction "bbs reg, pos, label" jumps to label if bit in position pos of register operand reg is one. A register is 32 bits wide and the bits are numbered 0 to 31, bit in position 0 being the least significant. Consider the following emulation of this instruction on a processor that does not have bbs implemented.


temp \leftarrow reg & mask

Branch to label if temp is non-zero

The variable temp is a temporary register. For correct emulation, the variable mask must be generated by

- mask $\leftarrow 0 \times 1 \ll \text{pos}$
- mask $\leftarrow 0 \times \text{ffff} \gg \text{pos}$
- mask $\leftarrow \text{pos}$
- mask $\leftarrow 0 \times f$

44. Station A uses 32 byte packets to transmit messages to Station B using a sliding window protocol. The round trip delay between A and B is 20 milliseconds and the bottleneck bandwidth on the path between A and B is 28 kbps. What is the optimal window size that A should use?
- 20
 - 40
 - 60
 - 80
45. Two computers C1 and C2 are configured as follows. C1 has IP address 203.197.253.1 and netmask 255.255.128.0. C2 has IP address 203.197.75.201 and netmask 255.255.192.0. Which one of the following statements is true?
- C1 and C2 both assume they are on the same network
 - C2 assumes C1 is on same network, but C1 assumes C2 is on a different network
 - C1 assumes C2 is on same network, but C2 assumes C1 is on a different network
 - C1 and C2 both assume they are on different networks
46. Station A needs to send a message consisting of 9 packets to Station B using a sliding window (window size 3) and go-back-n error control strategy. All packets are ready and immediately available for transmission. If every 5th packet that A transmits gets lost (but no acks from B ever get lost), then what is the number of packets that A will transmit for sending the message to B?
- 12
 - 14
 - 16
 - 18
47. Consider the following graph:



Which one of the following cannot be the sequence of edges added, in that order, to a minimum spanning tree using Kruskal's algorithm?

- a. (a-b), (a-f), (b-f), (d-c), (d-e)
 b. (a-b), (d-f), (d-c), (b-f), (d-e)
 c. (d-f), (a-b), (d-c), (b-f), (d-e)
 d. (d-f), (a-b), (b-f), (d-e), (d-c)
48. Let T be a depth first search tree in an undirected graph G . Vertices u and v are leaves of this tree T . The degrees of both u and v in G are at least 2. Which one of the following statements is true?
- a. There must exist a vertex w adjacent to both u and v in G
 b. There must exist a vertex w whose removal disconnects u and v in G
 c. There must exist a cycle in G containing u and v
 d. There must exist a cycle in G containing u and all its neighbours in G
49. An implementation of a queue Q , using two stacks $S1$ and $S2$, is given below:

```
void insert(Q, x){
    push(S1, x);
}

void delete(Q){
    if(!stack-empty(S2)) then
        pop(S2);
    else while(!stack-empty(S1)){
        x=pop(S1);
        push(S2, x);
    }
    x=pop(S2);
}
```

Let n insert and m ($\leq n$) delete operations be performed in an arbitrary order on an empty queue Q . Let x and y be the number of push and pop operations performed respectively in the process. Which one of the following is true for all n and m ?

- a. $n + m \leq x < 2n$ and $2m \leq y \leq n + m$
 b. $n + m \leq x < 2n$ and $2m \leq y \leq 2n$

- c. $2m \leq x < 2n$ and $2m \leq y \leq n + m$
 d. $2m \leq x < 2n$ and $2m \leq y \leq 2n$

50. A set X can be represented by an array x (n) as follows:

$$x[i] = \begin{cases} 1 & \text{if } i \in X \\ 0 & \text{otherwise} \end{cases}$$

Consider the following algorithm in which x , y , and z are boolean arrays of size n :

```
algorithm xyz(a[], y[], z[])
int i;

for(i=0; i<n; ++i)
    z[i] = (x[i] & y[i]) || (~x[i] & y[i]);
```

The set Z computed by the algorithm is

- a. $(X \cup Y)$
 b. $(X \cap Y)$
 c. $(X - Y) \cap (Y - X)$
 d. $(X - Y) \cup (Y - X)$

51. Consider the following recurrence:

$$T(n) = 2T\left(\frac{\sqrt{n}}{2}\right) + 1, T(1) = 1$$

Which one of the following is true?

- a. $T(n) = \theta(\log \log n)$
 b. $T(n) = \theta(\log n)$
 c. $T(n) = \theta(\sqrt{n})$
 d. $T(n) = \theta(n)$

52. The median of n elements can be found in $O(n)$ time. Which one of the following is correct about the complexity of quick sort, in which median is selected as pivot?

- a. $\theta(n)$
 b. $\theta(n \log n)$
 c. $\theta(n^2)$
 d. $\theta(n^3)$

53. Consider the following C-function in which a [n] and b [m] are two sorted integer arrays and c [$n + m$] be another integer array.

```
void xyz(int a[], int b[], int c[]){
    int i, j, k;

    i=j=k=0;
    while ((i<n) && (j<m))
        if(a[i] < b[j]) c[k++] = a[i++];
        else c[k++] = b[j++];
}
```

Which of the following condition(s) hold(s) after the termination of the while loop?

- (i) $j < m$, $k = n + j - 1$, and $a[n-1] < b[j]$ if $i = n$

(ii) $1 < n$, $k = m + i - 1$, and $b[m-1] \leq a[j]$ if $j = m$

- Only i
- Only ii
- Either i or ii but not both
- Neither i nor ii

54. Given two arrays of numbers a_1, \dots, a_n and b_1, \dots, b_n where each number is 0 or 1, the fastest algorithm to find the largest span (i, j) such that $a_i + a_{i+1} + \dots + a_j = b_1 + b_{1+1} + \dots + b_j$, or report that there is no such span.

- Takes $O(3^n)$ and $\Omega(2^n)$ time if hashing is permitted
- Takes $O(n^3)$ and $\Omega(n^{2.5})$ time in the key comparison model
- Takes $\Theta(n)$ time and space
- Takes $O(\sqrt{n})$ time only if the sum of the $2n$ elements is an even number

55. Consider these two functions and two statements S1 and S2 about them.

<pre> work1(int *a, int *b, int *c) { int i = a[1]*2; a[2] = a[i]; return a[1]*3 - 3; } </pre>	<pre> work2(int *a, int *b, int *c) { int i1 = 1*2; int i2 = a[1]; a[2] = i2*i1; return i2 - 3; } </pre>
--	--

S1: The transformation from work1 to work2 is valid, i.e., for any program state and input arguments, work2 will compute the same output and have the same effect on program state as work1.

S2: All the transformations applied to work1 to get work2 will always improve the performance (i.e. reduce CPU time) of work2 compared to work1.

- S1 is false and S2 is false
- S1 is false and S2 is true
- S1 is true and S2 is false
- S1 is true and S2 is true

56. Consider the following code written in a pass-by-reference language like FORTRAN and these statements about the code.

```

subroutine swap(ix, iy)
    it = ix
L1:   ix = iy
L2:   iy = it
end
ia = 3
ib = 8
call swap(ia, ib=5)
print *, ia, ib
end

```

S1: The compiler will generate code to allocate a temporary nameless cell, initialize it to 13, and pass the address of the cell to swap

S2: On execution the code will generate a runtime error on line L1

S3: On execution the code will generate a runtime error on line L2

S4: The program will print 13 and 8

S5: The program will print 13 and 2

Exactly the following set of statement(s) is correct:

- S1 and S2
- S1 and S4
- S3
- S1 and S5

57. Consider this code to swap two integers and these four statements: the code

```

void swap(int *px, int *py) {
    *x = *px;
    *y = *py;
    *px = *py;
    *py = *px;
}

```

S1: will generate a compilation error or may generate a segmentation fault at runtime depending on the arguments passed

S2: correctly implements the swap procedure for all input pointers referring to integers stored in memory locations accessible to the process

S3: implements the swap procedure correctly for some but not all valid input pointers

S4: may add or subtract integers and pointers

- S1
- S2 and S3
- S2 and S4
- S2 and S5

58. Consider the following grammar

$S \rightarrow FR$

$R \rightarrow *S \mid \epsilon$

$F \rightarrow id$

In the predictive parser table, M , of the grammar the entries $M[S, id]$ and $M[R, S]$ respectively

- $\{S \rightarrow FR\}$ and $\{R \rightarrow \epsilon\}$
- $\{S \rightarrow FR\}$ and $\{ \}$
- $\{S \rightarrow FR\}$ and $\{R \rightarrow *S\}$

d. $\{F \rightarrow id\}$ and $\{R \rightarrow e\}$

59. Consider the following translation scheme.

```

S → ER
R → *E{print("*");}R|ε
E → F + E{print("+");}|F
F → (S)|id{print(id.value);}

```

Here E_d is a token that represents an integer and $id.value$ represents the corresponding integer value. For an input $2 * 3 + 4$, this translation scheme prints

- a. $2 * 3 + 4$
 b. $2 * + 3 4$
 c. $2 3 * 4 +$
 d. $2 3 4 + *$
60. Consider the following C code segment.

```

for (i=0; i<N; i++) {
    for (j=0; j<N; j++) {
        if (i%2) {
            x += (4*j + 5*i);
            y += (7 + 4*j);
        }
    }
}

```

Which one of the following is false?

- a. The code contains loop-invariant computation
 b. There is scope of common sub-expression elimination in this code
 c. There is scope of strength reduction in this code
 d. There is scope of dead code elimination in this code
61. The atomic-fetch-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) {
    while (S->value == 0)
        fetch-and-set x, y;
}

void V (binary_semaphore *S) {
    S->value = 0;
}

```

Which one of the following is true?

- 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

62. 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

- a. 11 bits
 b. 13 bits
 c. 15 bits
 d. 20 bits

63. 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?

- 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

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

- a. 13 units
 b. 14 units
 c. 15 units
 d. 16 units

65. Consider three processes, all arriving at time zero, with total execution time of 10, 20, and 30 units, respectively. Each process spends the first 20% of execution time doing I/O, the next 70% of time doing computation, and the last 10% of time doing I/O again. The operating system uses a shortest remaining compute time first scheduling algorithm and schedules a new process either when the

running process gets blocked on I/O or when the running process finishes its compute burst. Assume that all I/O operations can be overlapped as much as possible. For what percentage of time does the CPU remain idle?

- 0%
- 10.6%
- 30.0%
- 89.4%

66. Consider the following snapshot of a system running its processes. Process i is holding x_i instances of a resource R , for $1 \leq i \leq n$. Currently, all instances of R are occupied. Further, for all i , process i has placed a request for an additional Y_i instances while holding the x_i instances it already has. There are exactly two processes p and q such that $y_p = y_q = 0$. Which one of the following can serve as a necessary condition to guarantee that the system is not approaching a deadlock?

- $\min(x_p, x_q) < \max_{k \in \{p, q\}} y_k$
- $x_p + x_q \geq \max_{k \in \{p, q\}} y_k$
- $\max(x_p, x_q) > 1$
- $\min(x_p, x_q) > 1$

67. Consider the relation account (customer, balance) where customer is a primary key and there are no null values. We would like to rank customers according to decreasing balance. The customer with the largest balance gets rank 1. Ties are not broken but ranks are skipped: if exactly two customers have the largest balance they each get rank 1 and rank 2 is not assigned.

```

Query1:
select A.customer, count(B.customer)
from account A, account B
where A.balance <= B.balance
group by A.customer

Query2:
select A.customer, 1+count(B.customer)
from account A, account B
where A.balance < B.balance
group by A.customer
  
```

Consider these statements about Query1 and Query2.

- Query1 will produce the same row set as Query2 for some but not all databases
- Both Query1 and Query2 are correct implementations of the specification
- Query1 is a correct implementation of the specification but Query2 is not

- Neither Query1 nor Query2 is a correct implementation of the specification
- Assigning rank with a pure relational query takes less time than scanning in decreasing balance order and assigning ranks using ODBC

Which two of the above statements are correct?

- 2 and 5
- 1 and 3
- 1 and 4
- 3 and 5

68. Consider the relation enrolled (student, course) in which (student, course) is the primary key, and the relation paid (student, amount) where student is the primary key. Assume no null values and no foreign keys or integrity constraints. Given the following four queries:

Query 1: select student from enrolled where student in (select student from paid)

Query 2: select student from paid where student in (select student from enrolled)

Query 3: select E.student from enrolled E, paid P where E.student = P.student

Query 4: select student from paid where exists


(select * from enrolled where enrolled.student = paidstudent)

Which one of the following statements is correct?

- All queries return identical row sets for any database
- Query2 and Query4 return identical row sets for all databases but there exist databases for which Query1 and Query2 return different row sets
- There exist databases for which Query3 returns strictly fewer rows than Query2
- There exist databases for which Query4 will encounter an integrity violation at runtime

69. Consider the relation enrolled (student, course) in which (student, course) is the primary key, and the relation paid (student, amount) where student is the primary key. Assume no null values and no foreign keys or integrity constraints- Assume that amounts 6000, 7000, 8000, 9000 and 10000 were each paid by 20% of the students. Consider these query plans (Plan

1 on left, Plan 2 on right) to “list all courses taken by students who have paid more than x”.



A disk seek takes 4ms, disk data transfer bandwidth is 300MB/s and checking a tuple to see if amount is greater than x takes 10 μ s.

Which of the following statements is correct?

- Plan 1 and Plan 2 will not output identical row sets for all databases
 - A course may be listed more than once in the output of Plan 1 for some databases
 - For $x = 5000$, Plan 1 executes faster than Plan 2 for all databases
 - For $x = 9000$, Plan 1 executes slower than Plan 2 for all databases
70. The following functional dependencies are given:
 $AB \rightarrow CD$, $AF \rightarrow D$, $DE \rightarrow F$, $C \rightarrow G$, $F \rightarrow E$, $G \rightarrow A$
 Which one of the following options is false?
- $\{CF\}^+ = \{ACDEFG\}$
 - $\{BG\}^+ = \{ABCDG\}$
 - $\{AF\}^+ = \{ACDEF\}$
 - $\{AB\}^+ = \{ACDFG\}$

Common Data Question

Common Data for Questions 71, 72, 73:

The 2^n vertices of a graph G correspond to all subsets of a set of size n , for $n \geq 6$. Two vertices of G are adjacent if and only if the corresponding sets intersect in exactly two elements.

71. The number of vertices of degree zero in G is
- 1
 - n
 - $n - 1$
 - 2^n
72. The maximum degree of a vertex in G is

- $\binom{n/2}{2} 2^{n-2}$
- 2^{n-2}
- $2^{n-1} \times 3$
- 2^{n-1}

73. The number of connected components in G is
- n
 - $n + 2$
 - $2^{n/2}$
 - $\frac{2^n}{n}$

Common Data for Question 74 & 75:

Consider two cache organizations: The first one is 32 KB 2-way set associative with 32-byte block size. The second one is of the same size but direct mapped. The size of an address is 32 bits in both cases. A 2-to-1 multiplexer has a latency of 0.6 ns while a k-bit comparator has a latency of $k/10$ ns. The hit latency of the set associative organization is h_1 while that of the direct mapped one is h_2 .

74. The value of h_1 is
- 2.4 ns
 - 2.3 ns
 - 1.8 ns
 - 1.7 ns
75. The value of h_2 is
- 2.4 ns
 - 2.3 ns
 - 1.8 ns
 - 1.7 ns

Statement for Linked Answer Questions 76 & 77:

A 3-ary max heap is like a binary max heap, but instead of 2 children, nodes have 3 children. A 3-ary heap can be represented by an array as follows: The root is stored in the first location, a [0], nodes in the next level, from left to right, is stored from a [13 to a [3). The nodes from the second level of the tree from left to right are stored from a 43 location onward. An item x can be inserted into a 3-ary heap containing n items by placing x in the location a [n] and pushing it up the tree to satisfy the heap property.

76. Which one of the following is a valid sequence of elements in an array representing 3-ary max heap?
- 1, 3, 5, 6, 8, 9
 - 9, 6, 3, 1, 8, 5
 - 9, 3, 6, 8, 5, 1
 - 9, 5, 6, 8, 3, 1
77. Suppose the elements 7, 2, 10, and 4 are inserted, in that order, into the valid 3-ary max heap found in the above question, Q.76. Which one of the following is the sequence of items in the array representing the resultant heap?
- 10, 7, 9, 8, 3, 1, 5, 2, 6, 4
 - 10, 9, 8, 7, 6, 5, 4, 3, 2, 1
 - 10, 9, 4, 5, 7, 6, 8, 2, 1, 3
 - 10, 8, 6, 9, 7, 2, 3, 4, 1, 5

Statement for Linked Answer Questions 78 & 79:

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 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.

78. The above implementation of barrier is incorrect. Which one of the following is true?
- The barrier implementation is wrong due to the use of binary semaphore S

- The barrier implementation may lead to a deadlock if two barrier invocations are used in immediate succession
 - Lines 6 to 10 need not be inside a critical section
 - The barrier implementation is correct if there are only two processes instead of three
79. Which one of the following rectifies the problem in the implementation?
- Lines 6 to 10 are simply replaced by process_arrived--
 - At the beginning of the barrier the first process to enter the barrier waits until process_arrived becomes zero before proceeding to execute V(S)
 - Context switch is disabled at the beginning of the barrier and re-enabled at the end
 - The variable process_left is made private instead of shared

Statement for Linked Answer Questions 80 & 81:

A CPU has a 32 KB direct mapped cache with 128-byte block size. Suppose A is a two-dimensional array of size 512 x 512 with elements that occupy 8-bytes each. Consider the following two C code segments, P1 and P2.

```
P1: for (i=0; i<512; i++) {
    for (j=0; j<512; j++) {
        x += A[i][j];
    }
}

P2: for (i=0; i<512; i++) {
    for (j=0; j<512; j++) {
        x += A[j][i];
    }
}
```

P1 and P2 are executed independently with the same initial state, namely, the array A is not in the cache and i, j, x are in registers. Let the number of cache misses experienced by P1 be M1 and that for P2 be M2.

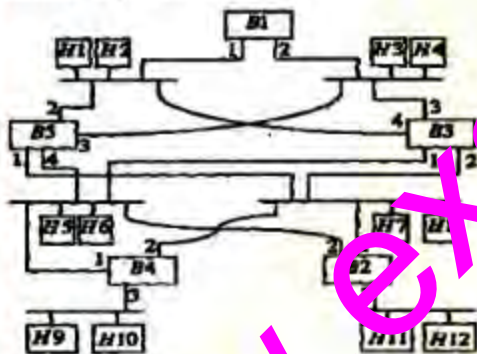
80. The value of M1 is
- 0
 - 2048
 - 16384
 - 262144
81. The value of the ratio M1/M2 is
- 0
 - 1/16

- c. 1/8
d. 16

Statement for Linked Answer Questions 82 & 83:

Consider the diagram shown below where a number of LANs are connected by (transparent) bridges. In order to avoid packets looping through circuits in the graph, the bridges organize themselves in a spanning tree. First, the root bridge is identified as the bridge with the least serial number. Next, the root sends out (one or more) data units to enable the setting up of the spanning tree of shortest paths from the root bridge to each bridge.

Each bridge identifies a port (the root port) through which it will forward frames to the root bridge. Port conflicts are always resolved in favour of the port with the lower index value. When there is a possibility of multiple bridges forwarding to the same LAN (but not through the root port), ties are broken as follows: bridges closest to the root get preference and between such bridges, the one with the lowest serial number is preferred.



82. For the given connection of LANs by bridges, which one of the following choices represents the depth first traversal of the spanning tree of bridges?
- a. B1, B5, B3, B4, B2
b. B1, B3, B5, B2, B4
c. B1, B5, B2, B3, B4
d. B1, B3, B4, B5, B2
83. Consider the correct spanning tree for the previous question. Let host H1 send out a broadcast ping packet. Which of the following options represents the correct forwarding table on B3?
- a.

Hosts	Port
H1, H2, H3, H4	3
H5, H6, H9, H10	1
H7, H8, H11, H12	2

b.

Hosts	Port
H1, H2	4
H3, H4	3
H5, H6	1
H7, H8, H9, H10, H11, H12	2

c.

Hosts	Port
H3, H4	3
H5, H6, H9, H10	1
H1, H2	4
H7, H8, H11, H12	2

d.

Hosts	Port
H1, H2, H3, H4	3
H5, H6, H9, H10	1
H7, H8, H11, H12	4

Statement for Linked Answer Questions 84 & 85:

84. Which one of the following grammars generates the language $L = \{a^i b^j \mid i \neq j\}$?

a.

$S \rightarrow AC \mid CB$
 $C \rightarrow aCb \mid a \mid b$
 $A \rightarrow aA \mid \epsilon$
 $B \rightarrow Bb \mid \epsilon$

b.

$S \rightarrow aS \mid Sb \mid a \mid b$

c.

$S \rightarrow AC \mid CB$
 $C \rightarrow aCb \mid \epsilon$
 $A \rightarrow aA \mid \epsilon$
 $B \rightarrow Bb \mid \epsilon$

d.

$S \rightarrow AC \mid CB$
 $C \rightarrow aCb \mid \epsilon$
 $A \rightarrow aA \mid a$
 $B \rightarrow bB \mid b$

85. In the correct grammar above, what is the length of the derivation (number of steps starting from S) to generate the string $a^i b^m$ with $i \neq m$?
- a. $\max(i, m) + 2$
b. $1 + m + 2$
c. $1 + m + 3$
d. $\max(i, m) + 2$