

SECTION - A

1. This question consists of TWENTY-THREE multiple questions of ONE mark each. For each question (1.1 – 1.23), four possible alternatives (A, B, C and D) are given, out of which ONLY ONE is correct. Indicate the correct answer in the boxes corresponding to the questions only on the FIRST sheet of the answer book.
- 1.1 The minimum number of cards to be dealt from an arbitrarily shuffled deck of 52 cards to guarantee that three cards are from some same suit is
(a) 3 (b) 8 (c) 9 (d) 12
- 1.2 An $n \times n$ array v is defined as follows:
 $v[i, j] = i - j$ for all $i, j, i \leq n, 1 \leq j \leq n$
The sum of the elements of the array v is
(a) 0 (b) $n - 1$ (c) $n^2 - 3n + 2$ (d) $n^2 \frac{(n+1)}{2}$
- 1.3 The determinant of the matrix
$$\begin{bmatrix} 2 & 0 & 0 & 0 \\ 8 & 1 & 7 & 2 \\ 2 & 0 & 2 & 0 \\ 9 & 0 & 6 & 1 \end{bmatrix}$$
 is:
(a) 4 (b) 0 (c) 15 (d) 20
- 1.4 Let S and T be language over $\Sigma = \{a, b\}$ represented by the regular expressions $(a+b^*)^*$ and $(a+b)^*$, respectively. Which of the following is true?
(a) $S \subset T$ (b) $T \subset S$ (c) $S = T$ (d) $S \cap T = \phi$
- 1.5 Let L denotes the language generated by the grammar $S \rightarrow 0S0/00$.
Which of the following is true?
(a) $L = 0^+$ (b) L is regular but not 0^+
(c) L is context free but not regular (d) L is not context free
- 1.6 The number 43 in 2's complement representation is
(a) 01010101 (b) 11010101 (c) 00101011 (d) 10101011
- 1.7 To put the 8085 microprocessor in the wait state
(a) lower the HOLD input (b) lower the READY input
(c) raise the HOLD input (d) raise the READY input

- 1.8 Comparing the time T_1 taken for a single instruction on a pipelined CPU with time T_2 taken on a non-pipelined but identical CPU, we can say that
- (a) $T_1 \leq T_2$ (b) $T_1 \geq T_2$ (c) $T_1 < T_2$
(d) T_1 is T_2 plus the time taken for one instruction fetch cycle

- 1.9 The 8085 microprocessor responds to the present of an interrupt
- (a) as soon as the TRAP pin becomes 'high'
(b) by checking the TRAP pin for 'high' status at the end of each instruction each
(c) by checking the TRAP pin for 'high' status at the end of the execution of each instruction.
(d) by checking the TRAP pin for 'high' status at regular intervals.

- 1.10 The most appropriate matching for the following pairs

| | |
|------------------------------|--------------|
| X: Indirect addressing | 1: Loops |
| Y: Immediate addressing | 2: Pointers |
| Z: Auto decrement addressing | 3: Constants |

is

- (a) X - 3 Y - 2 Z - 1 (b) X - 1 Y - 3 Z - 2
(c) X - 2 Y - 3 Z - 1 (d) X - 3 Y - 1 Z - 2

- 1.11 The following C declarations

```
struct node{
    int i;
    float j;
};
struct node *s[10];
```

define s to be

- (a) An array, each element of which is a pointer to a structure of type node
(b) A structure of 2 fields, each field being a pointer to an array of 10 elements
(c) A structure of 3 fields: an integer, a float, and an array of 10 elements
(d) An array, each element of which is a structure of type node

- 1.12 The most appropriate matching for the following pairs

| | |
|---|---------------------------------|
| X: <code>m=malloc(5); m= NULL;</code> | 1: using dangling pointers |
| Y: <code>free(n); n->value=5;</code> | 2: using uninitialized pointers |
| Z: <code>char *p; *p='a';</code> | 3. lost memory |

is:

- (a) X - 1 Y - 3 Z - 2 (b) X - 2 Y - 1 Z - 3
(c) X - 3 Y - 2 Z - 1 (d) X - 3 Y - 1 Z - 2

1.13 The most appropriate matching for the following pairs

X: depth first search 1: heap
Y: breadth-first search 2: queue
Z: sorting 3: stack

is:

- (a) X - 1 Y - 2 Z - 3 (b) X - 3 Y - 1 Z - 2
(c) X - 3 Y - 2 Z - 1 (d) X - 2 Y - 3 Z - 1

1.14 Consider the following nested representation of binary trees: (X Y Z) indicates Y and Z are the left and right sub stress, respectively, of node X. Note that Y and Z may be NULL, or further nested. Which of the following represents a valid binary tree?

- (a) (1 2 (4 5 6 7)) (b) (1 (2 3 4) 5 6) 7)
(c) (1 (2 3 4)(5 6 7)) (d) (1 (2 3 NULL) (4 5))

1.15 Let s be a sorted array of n integers. Let t(n) denote the time taken for the most efficient algorithm to determined if there are two elements with sum less than 1000 in s. which of the following statements is true?

- (a) t (n) is O(1) (b) $n \leq t(n) \leq n \log_2 n$
(c) $n \log_2 n \leq t(n) < \left(\frac{n}{2}\right)$ (d) $t(n) = \left(\frac{n}{2}\right)$

1.16 Aliasing in the context of programming languages refers to

- (a) multiple variables having the same memory location
(b) multiple variables having the same value
(c) multiple variables having the same identifier
(d) multiple uses of the same variable

1.17 Consider the following C declaration

```
struct {  
    short s [5]  
    union {  
        float y;  
        long z;  
    } u;  
};
```

Assume that objects of the type short, float and long occupy 2 bytes, 4 bytes and 8 bytes, respectively. The memory requirement for variable t, ignoring alignment considerations, is

- (a) 22 bytes (b) 14 bytes (c) 18 bytes (d) 10 bytes

1.18 The number of tokens in the following C statement

```
printf("i=%d, &i=%x",i,&i);
```

is

- (a) 3 (b) 26 (c) 10 (d) 21

1.19. Which of the following derivations does a top-down parser use while parsing an input string? The input is assumed to be scanned in left to right order.

- (a) Leftmost derivation
(b) Leftmost derivation traced out in reverse
(c) Rightmost derivation
(d) Rightmost derivation traced out in reverse

1.20. Which of the following need not necessarily be saved on a context switch between processes?

- (a) General purpose registers (b) Translation look-aside buffer
(c) Program counter (d) All of the above

1.21. Let $m[0] \dots m[4]$ be mutexes (binary semaphores) and $P[0] \dots P[4]$ be processes. Suppose each process $P[i]$ executes the following:

```
wait (m[i];wait (m[(i+1) mode 4]));
```

.....

```
release (m[i]); release (m[(i+1)mod 4]);
```

This could cause

- (a) Thrashing (b) Deadlock
(c) Starvation, but not deadlock (d) None of the above

1.22. B^+ -trees are preferred to binary trees in databases because

- (a) Disk capacities are greater than memory capacities
(b) Disk access is much slower than memory access
(c) Disk data transfer rates are much less than memory data transfer rates
(d) Disks are more reliable than memory

1.23. Given the relations

employee (name, salary, deptno), and

department (deptno, deptname, address)

Which of the following queries cannot be expressed using the basic relational algebra operations ($\sigma, \pi, \times, \bowtie, \cup, \cap, -$)?

- (a) Department address of every employee
- (b) Employees whose name is the same as their department name
- (c) The sum of all employees' salaries
- (d) All employees of a given department

2. This question consists of TWENTY-SIX multiple questions of TWO marks each. For each question (2.1 – 2.26), four possible alternatives (A, B, C and D) are given, out of which ONLY ONE is correct. Indicate the correct answer in the boxes corresponding to the questions only on the SECOND sheet of the answer book.

2.1 X, Y and Z are closed intervals of unit length on the real line. The overlap of X and Y is half a unit. The overlap of Y and Z is also half a unit. Let the overlap of X and Z be k units. Which of the following is true?

- (a) k must be 1
- (b) k must be 0
- (c) k can take any value between 0 and 1
- (d) None of the above

2.2. E_1 and E_2 are events in a probability space satisfying the following constraints:

- $\Pr(E_1) = \Pr(E_2)$
- $\Pr(E_1 \cup E_2) = 1$
- E_1 and E_2 are independent

The value of $\Pr(E_1)$, the probability of the event E_1 , is

- (a) 0
- (b) $\frac{1}{4}$
- (c) $\frac{1}{2}$
- (d) 1

2.3. Let $S = \sum_{i=3}^{100} i \log_2 i$, and

$$\text{and } T = \int_2^{100} x \log_2 x dx$$

Which of the following statements is true?

- (a) $S > T$
- (b) $S = T$
- (c) $S < T$ and $2S > T$
- (d) $2S \leq T$

2.4. A polynomial $p(x)$ satisfies the following:

$$p(1) = p(3) = p(5) = 1$$

$$p(2) = p(4) = -1$$

The minimum degree of such a polynomial is

- (a) 1
- (b) 2
- (c) 3
- (d) 4

- 2.5. A relation R is defined on the set of integers as $x R y$ iff $(x+y)$ is even. Which of the following statements is true?
- (a) R is not an equivalence relation
 (b) R is an equivalence relation having 1 equivalence class
 (c) R is an equivalence relation having 2 equivalence classes
 (d) R is an equivalence relation having 3 equivalence classes
- 2.6. Let $P(S)$ denotes the powerset of set S. Which of the following is always true?
- (a) $P(P(S))=P(S)$
 (b) $P(S) \cap P(P(S)) = \{\emptyset\}$
 (c) $P(S) \cap S = P(S)$
 (d) $S \notin P(S)$
- 2.7. Let a, b, c, d be propositions. Assume that the equivalence $a \leftrightarrow (b \vee \neg b)$ and $b \leftrightarrow c$ hold. Then the truth-value of the formula $(a \wedge b) \rightarrow (a \wedge c) \vee d$ is always
- (a) True
 (b) False
 (c) Same as the truth-value of b
 (d) Same as the truth-value of d
- 2.8. What can be said about a regular language L over $\{a\}$ whose minimal finite state automation has two states?
- (a) L must be $\{a^n | n \text{ is odd}\}$
 (b) L must be $\{a^n | n \text{ is even}\}$
 (c) L must be $\{a^n | n \geq 0\}$
 (d) Either L must be $\{a^n | n \text{ is odd}\}$, or L must be $\{a^n | n \text{ is even}\}$
- 2.9. Consider the following decision problems:
 (P1) Does a given finite state machine accept a given string
 (P2) Does a given context free grammar generate an infinite number of strings
 Which of the following statements is true?
- (a) Both (P1) and (P2) are decidable
 (b) Neither (P1) nor (P2) are decidable
 (c) Only (P1) is decidable
 (d) Only (P2) is decidable
- 2.10. The simultaneous equations on the Boolean variables x, y, z and w,
- $$\begin{aligned}x + y + z &= 1 \\xy &= 0 \\xz + w &= 1 \\xy + \bar{z} \bar{w} &= 0\end{aligned}$$
- have the following solution for x, y, z and w, respectively:
- (a) 0 1 0 0 (b) 1 1 0 1 (c) 1 0 1 1 (d) 1 0 0 0

- 2.15. Suppose you are given an array $s[1..n]$ and a procedure $\text{reverse}(s, i, j)$ which reverses the order of elements in s between positions i and j (both inclusive). What does the following sequence do, where $1 \leq k \leq n$:

$\text{reverse}(s, 1, k);$
 $\text{reverse}(s, k + 1, n);$
 $\text{reverse}(s, 1, n);$

- (a) Rotates s left by k positions
(b) Leaves s unchanged
(c) Reverses all elements of s
(d) None of the above
- 2.16. Let LASTPOST, LASTIN and LASTPRE denote the last vertex visited in a postorder, inorder and preorder traversal. Respectively, of a complete binary tree. Which of the following is always true?

- (a) LASTIN = LASTPOST
(b) LASTIN = LASTPRE
(c) LASTPRE = LASTPOST
(d) None of the above

- 2.17. Consider the following functions

$$f(n) = 3n^{\sqrt{n}}$$

$$g(n) = 2^{\sqrt{n} \log_2 n}$$

$$h(n) = n!$$

Which of the following is true?

- (a) $h(n)$ is $O(f(n))$
(b) $h(n)$ is $O(g(n))$
(c) $g(n)$ is not $O(f(n))$
(d) $f(n)$ is $O(g(n))$
- 2.18. Let G be an undirected connected graph with distinct edge weight. Let e_{\max} be the edge with maximum weight and e_{\min} the edge with minimum weight. Which of the following statements is false?
- (a) Every minimum spanning tree of G must contain e_{\min}
(b) If e_{\max} is in a minimum spanning tree, then its removal must disconnect G
(c) No minimum spanning tree contains e_{\max}
(d) G has a unique minimum spanning tree
- 2.19. Let G be an undirected graph. Consider a depth-first traversal of G , and let T be the resulting depth-first search tree. Let u be a vertex in G and let v be the first new (unvisited) vertex visited after visiting u in the traversal. Which of the following statements is always true?
- (a) $\{u, v\}$ must be an edge in G , and u is a descendant of v in T
(b) $\{u, v\}$ must be an edge in G , and v is a descendant of u in T
(c) If $\{u, v\}$ is not an edge in G then u is a leaf in T
(d) If $\{u, v\}$ is not an edge in G then u and v must have the same parent in T

2.20. The value of j at the end of the execution of the following C program

```
int incr (int i)
{
    static int count = 0;
    count = count + i;
    return (count);
}
main () {
    int i, j;
    for (i = 0; i <= 4; i++)
        j = incr(i);
}
```

is

- (a) 10 (b) 4 (c) 6 (d) 7

2.21. Given the following expression grammar:

$E \rightarrow E * F \mid F + E \mid F$

$F \rightarrow F - \mid id$

Which of the following is true?

- (a) * has higher precedence than + (b) - has higher precedence than *
(c) + and - have same precedence (d) + has higher precedence than *

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

- (a) 1.9999 milliseconds (b) 1 millisecond
(c) 9.999 microseconds (d) 1.9999 microseconds


2.23. Which of the following is NOT a valid deadlock prevention scheme?

- (a) Release all resources before requesting a new resource
(b) Number the resources uniquely and never request a lower numbered resource than the last one requested.
(c) Never request a resource after releasing any resource
(d) Request and all required resources be allocated before execution.

5. A multiset is an unordered collection of elements where elements may repeat any number of times. The size of a multiset is the number of elements in it counting repetitions.
- (a) What is the number of multisets of size 4 that can be constructed from n distinct elements so that at least one element occurs exactly twice?
- (b) How many multisets can be constructed from n distinct elements?
6. Let S be a set of n elements $\{1, 2, \dots, n\}$ and G a graph with 2^n vertices, each vertex corresponding to a distinct subset of S . Two vertices are adjacent iff the symmetric difference of the corresponding sets has exactly 2 elements. Note: The symmetric difference of two sets R_1 and R_2 is defined as $\left(\frac{R_1}{R_2}\right) \cup \left(\frac{R_2}{R_1}\right)$
- (a) Every vertex in G has the same degree. What is the degree of a vertex in G ?
- (b) How many connected components does G have?
7. (a) Construct a minimal finite state machine that accepts the language, over $\{0,1\}$, of all strings that contain neither the substring 00 nor the substring 11.
- (b) Consider the grammar
- $$\begin{aligned} S &\rightarrow aSAb \\ S &\rightarrow \epsilon \\ A &\rightarrow bA \\ A &\rightarrow \epsilon \end{aligned}$$

Where S, A are non-terminal symbols with S being the start symbol; a, b are terminal symbols and ϵ is the empty string. This grammar generates strings of the form $a^i b^j$ for some $i, j \geq 0$, where i and j satisfy some condition. What is the condition on the values of i and j ?

8. A pushdown automaton (pda) is given in the following extended notation of finite state diagrams:



The nodes denote the states while the edges denote the moves of the pda. The edge labels are of the form $d, s / s'$ where d is the input symbol read and s, s' are the stack contents before and after the move. For example the edge labeled $1, s / 1.s$ denotes the move from state q_0 to q_0 in which the input symbol 1 is read and pushed to the stack.


- (a) Introduce two edges with appropriate labels in the above diagram so that the resulting pda accepts the language

$$\{x2x^R \mid x \in \{0,1\}^*, x^R \text{ denotes reverse of } x\}, \text{ by empty stack.}$$

- (b) Describe a non-deterministic pda with three states in the above notation that accept the language $\{0^n 1^m \mid n \leq m \leq 2n\}$ by empty stack
9. Design a logic circuit to convert a single digit BCD number to the number modulo six as follows (Do not detect illegal input):
- (a) Write the truth table for all bits. Label the input bits I_1, I_2, \dots . With I_1 as the least significant bit. Label the output bits R_1, R_2, \dots . With R_1 as the least significant bit. Use 1 to signify truth.
- (b) Draw one circuit for each output bit using, *altogether*, two two-input AND gates, one two-input gate and two NOT gates.
10. Consider the 8085 instruction IN 09H stored as follows:

| Memory Address | Machine Code |
|----------------|--------------|
| 3050 | DA |
| 3051 | 09 |

And the following incomplete timing diagram for the instruction:



- (a) Write the contents of the boxes, A, B, C and D in hexadecimal in your answer sheet. Do not draw any pictures.
- (b) Write the state of both ALE and \overline{RD} pins at time units T1, T2, T3 and T4.
- (c) How do you generate the signal that tells the peripheral to put the data on the bus? Answer by completing the following statement in your answer book:
By combining signals
11. Consider the following 8085 program segment, where registers B and C contain BCD values:
- ```

S1: MVI A, 99H
 MVI D, 00H
 SUB C
 ADD B
 DAA
S2: JC S3

```

```

MOV E, A
MVI A, 99H
SUB E
MOV E, A
JZ S4
MVI D, FFH
JMP S4
S3: INC A
 DAA
 MOV E, A

```

S4: .....

- (a) For the two pairs (B = 44, C = 25) and (B = 33, C = 46) at S1,
- Find the values in register A when control reaches S2.
  - Find the values in registers D and E when control reaches S4.
- (b) What, in general, is the value of D and E as a function of B and C when control reaches S4.
12. An instruction pipeline has five stages where each stage takes 2 nanoseconds and all instructions use all five stages. Branch instructions are not overlapped, i.e., the instruction after the branch is not fetched till the branch instruction is completed. Under ideal conditions.
- Calculate the average instruction execution time assuming that 20% of all instruction executed are branch instructions. Ignore the fact that some branch instructions may be conditional.
  - If a branch instruction is a conditional branch instruction, the branch need not be taken. If the branch is not taken, the following instructions can be overlapped. When 80% of all branch instructions are conditional branch instructions, and 50% of the conditional branch instructions are such that the branch is taken, calculate the average instruction execution time.
13. Suppose a stack implementation supports, in addition to PUSH and POP, an operation REVERSE, which reverses the order of the elements on the stack.
- To implement a queue using the above stack implementation, show how to implement ENQUEUE using a single operation and DEQUEUE using a sequence of 3 operations.
  - The following postfix expression, containing single digit operands and arithmetic operators + and \*, is evaluated using a stack.  
 $5\ 2\ * \ 3\ 4\ + \ 5\ 2\ * \ * \ +$   
 Show the contents of the stack.
    - After evaluating  $5\ 2\ * \ 3\ 4\ +$
    - After evaluating  $5\ 2\ * \ 3\ 4\ + \ 5\ 2$
    - At the end of evaluation.

14. Consider the line  $y = \frac{n}{m}x$ , where  $n$  and  $m$  are positive integers.
- (a) If  $mq - np < 0$ , then is the point  $(p,q)$  above the line, below the line, or on the line?
- (b) Complete the following function, that returns true if the line segment with endpoints  $(p,q)$  and  $(r,s)$  intersects the line  $y = \frac{n}{m}x$ , by writing the line number and the content of each box in your answer book.
- ```

1: function clash (m, n, p, q, r, s: integer): Boolean;
2: begin
3:   clash = false;
4:   if (m*q - n * p)  0 then clash := true;
5:   If (m*s - n * r)  0 then clash := true;
6:   if (m*q - n * p)  0 and (m*s - n * r)  0 then clash := true;
7:   if (m*q - n * p)  0 and (m*s - n * r)  0 then clash := true;
8: end;
```

15. Suppose you are given arrays $p[1.....N]$ and $q[1.....N]$ both uninitialized that is, each location may contain an arbitrary value), and a variable $count$, initialized to 0. Consider the following procedures set and is_set :

```

set (i) {
    count = count + 1;
    q [count] = i;
    p[i] = count;
}

is_set(i) {
    if (p[i] ≤ 0 or p[i] > count)
        return false;
    if (q[p[i]] ≠ i)
        return false;
return true;
}
```

- (a) Suppose we make the following sequence of calls:
 $set(7); set(3); set(9);$
 After these quence of calls, what is the value of $count$, and what do $q[1]$, $q[2]$, $q[3]$, $p[7]$, $p[3]$ and $p[9]$ contain?
- (b) Complete the following statement "The first $count$ elements of _____ contain values i such that $set(_____)$ has been called".
- (c) Show that if $set(i)$ has not been called for some i , then regardless of what $p[i]$ contains, $is_set(i)$ will return false.

16. A recursive program to compute Fibonacci numbers is shown below. Assume you are also given an array $f[0\dots m]$ with all elements initialized to 0.

```
fib(n) {  
    if (n > M) error ();  
    if ( n==0) return 1;  
    if (n ==1) return 1;  
    if (  ) _____(1)  
    return  _____(2)  
    t = fib(n - 1) + fib (n - 2);  
     _____(3)  
    return t;  
}
```

- (a) Fill in the boxes with expressions/statements to make $\text{fib}()$ store and reuse computed Fibonacci values. Write the box number and the corresponding contents in your answer book.
- (b) What is the time complexity of the resulting program when computing $\text{fib}(n)$?
17. An array contains four occurrences of 0, five occurrences of 1, and three occurrences of 2 in any order. The array is to be sorted using swap operations (elements that are swapped need to be adjacent).
- (a) What is the minimum number of swaps needed to sort such an array in the worst case?
- (b) Give an ordering of elements in the above array so that the minimum number of swaps needed to sort the array is maximum.
18. Consider the following program is pseudo-Pascal syntax.

```
program main;  
    var x: integer;  
    procedure Q [z:integer);  
    begin  
        z: z + x;  
        writeln(z)  
    end;  
    procedure P (y:integer);  
        var x: integer;  
    begin  
        x: y + 2;  
        Q(x);  
        writeln(x)  
    end;  
begin
```

```

    x:=5;
    P(x);
    Q(x);
    writeln(x)
end.

```

What is the output of the program, when

- (a) The parameter passing mechanism is call-by-value and the scope rule is static scoping?
- (b) The parameter passing mechanism is call-by-reference and the scope rule is dynamic scoping?

19. Consider the syntax directed translation scheme (SDTS) given in the following. Assume attribute evaluation with bottom-up parsing, i.e., attributes are evaluated immediately after a reduction.

```

E → E1 * T {E.val = E1.val * T.val}
  E → T {E.val = T.val}
T → F - T1 {T.val = F.val - T1.val}
  T → F {T.val = F.val}
  F → 2 {F.val = 2}
  F → 4 {F.val = 4}

```

- (a) Using this SDTS, construct a parse tree for the expression
 $4 - 2 - 4 * 2$
 and also compute its E.val.

- (b) It is required to compute the total number of reductions performed to parse a given input. Using synthesized attributes only, modify the SDTS given, without changing the grammar, to find E.red, the number of reductions performed while reducing an input to E.

20. (a) Fill in the boxes below to get a solution for the readers-writers problem, using a single binary semaphore, mutex (initialized to 1) and busy waiting. Write the box numbers (1,2 and 3), and their contents in your answer book.

```

int R = 0, W = 0;
Reader ( ) {
L1:          wait (mutex);
             If (W == 0) {
               R = R + 1;

```



```

    }
else {
    [ ] _____(2)
    goto L1;
}
.../* do the read */
wait (mutex)
R = R - 1;
signal (mutex);
}
Writer () {
L2:    wait(mutex);
      If ( [ ] ) { _____(3)
        signal (mutex);
        goto L2;
      }
      W=1;
      signal (mutex);
      .../*do the write*/
      wait(mutex)
      W = 0;
      signal (mutex);
}

```

(b) Can the above solution lead to starvation of writers?

21. (a) Suppose you are given an empty B+-tree where each node (leaf and internal) can store up to 5 key values. Suppose values 1,2,..., 10 are inserted, in order, into the tree, Show the tree pictorially
- After 6 insertions, and
 - After all 10 insertions
- Do NOT show intermediate stages.
- (b) Suppose instead of splitting a node when it is full, we try to move a value to the left sibling. If there is no left sibling, or the left sibling is full, we split the node. Show the tree after values, 1, 2,..., 9 have been inserted. Assume, as in (a) that each node can hold up to 5 keys.
- (c) In general, suppose a B+-tree node can hold a maximum of m keys, and you insert a long sequence of keys in increasing order. Then what approximately is the average number of keys in each leaf level node.
- In the normal case, and
 - With the insertion as in (b).

22. Consider a bank database with only one relation transaction (transno, acctno, date, amount)
The amount attribute value is positive for deposits and negative for withdrawals.
- (a) Define an SQL view TP containing the information.
(acctno, T1.date, T2.amount)
for every pair of transactions T1, T2 such that T1 and T2 are transaction on the same account and the date of T2 is \leq the date of T1.
- (b) Using only the above view TP, write a query to find for each account the minimum balance it ever reached (not including the 0 balance when the account is created). Assume there is at most one transaction per day on each account, and each account has had atleast one transaction since it was created. To simply your query, break it up into 2 steps by defining an intermediate view V.