CHENNAI MATHEMATICAL INSTITUTE

M.Sc. / Ph.D. Programme in Computer Science

Entrance Examination, 2026

Part A

1. Six people A, B, C, D, E, F sit around a circular table.

   - A sits opposite D
   - B is not adjacent to A
   - C sits immediately to the right of B
   - E is adjacent to D

   Who sits opposite C?

   (a) B (b) D (c) E (d) F

   Answer. (c)

   Let there be six chairs around the table, numbered $0, 1, \ldots, 6$ clockwise. Suppose, without loss of generality, that A sits in chair 0. Then D has to sit in chair 3. Since B is not adjacent to A, she has to sit in chair 2 or chair 4. C sits to the right of B, which means that either B and C sit in chairs 2 and 1, or they sit in chairs 4 and 3. But 3 is already occupied by D, so B and C are in 2 and 1. E is adjacent to D, so he has to be in chair 4, which is opposite to chair 1, on which sits C. ⊣

2. Dhruva is attending a party. If they play both dumb charades and musical chairs, he will leave the party. If Abhay does not dance or Bharat does not juggle, then also he will leave the party. If Abhay dances and Bharat juggles, then everyone gets free chocolate. Suppose Mahesh, who is also attending the party, did not get free chocolate and musical chairs were played, which of the following can you infer?

   (a) Dumb charades are not played (b) Abhay did not dance
   (c) Bharat did not juggle (d) Dhruva left the party

   Answer. (d)

   Let us use the following letters to denote various sentences of interest.

   - a: Abhay dances.
   - b: Bharat juggles.
   - c: Everyone gets free chocolate.
   - d: Dhruva leaves the party midway.
   - e: Dumb charades are not played.
   - f: Musical chairs are not played.

Also assuming the standard logical notation ¬, ∧, ∨ and → which stand for negation, conjunction, disjunction and implication (respectively), we have the following facts.

$$
\begin{array}{cc}
(1) & (a \wedge b) \to c \\
(2) & (\neg a \vee \neg b) \to d \\
(3) & d \to (e \vee f) \\
(4) & \neg c \\
(5) & \neg f
\end{array}
$$

From (1) and (4), it follows that $\neg a \vee \neg b$. From (2), it now follows that d. and So a is false or b is false. So d (which is the same as option (d)) is true. One can check that each of the other options (a), (b) and (c) can be false along with (d). So (d) is the only option that can be inferred from the given information. $\dashv$

3. Suppose you are writing a 100-question multiple-choice exam with 4 choices per question and exactly one correct answer. You know the answers to 80% of the questions, and you guess the answer to the remaining 20% uniformly at random. What is the expected number of questions you answer correctly?

(a) 82 (b) 85 (c) 90 (d) 95

(a) 82 (b) 85 (c) 90 (d) 95

Answer. (b)

You will anyway answer 80 questions correctly. Out of the remaining 20, you will get $\frac{1}{4} \cdot 20$ correct, which evaluates to 5. $\dashv$

4. In a multiple choice test with 100 questions, each question has 5 possible answers with exactly one of them being correct. A student knows the answer to 40 of these questions. For the rest he will guess the answer uniformly at random.

A question is picked uniformly at random and posed to the student. If he answers it correctly, what is the probability that he actually knew the answer to that question?

(a) $\frac{2}{5}$ (b) $\frac{5}{13}$ (c) $\frac{13}{25}$ (d) $\frac{10}{13}$

Answer. (d)

Suppose a question picked uniformly at random and the students provides an answer. Let K be the event that the student knows the answer. Clearly, $\Pr(K) = \frac{40}{100} = \frac{2}{5}$. Let C be the event that the answer provided is correct.

$$\Pr(C) = \Pr(C \mid K) \cdot \Pr(K) + \Pr(C \mid \neg K) \cdot \Pr(\neg K) = 1 \cdot \frac{2}{5} + \frac{1}{5} \cdot \frac{3}{5} = \frac{2}{5} + \frac{3}{25} = \frac{13}{25}.$$

The question asks us to compute the probability that the student actually knows the answer, provided that the answer was correct. This is

$$\Pr(K \mid C) = \frac{\Pr(K \wedge C)}{\Pr(C)} = \frac{\Pr(K)}{\Pr(C)} = \frac{2/5}{13/25} = \frac{2}{5} \cdot \frac{25}{13} = \frac{10}{13}.$$ $\dashv$

5. 1000 candidates appear in a competitive examination. There are 50 questions, each worth 2 marks. There is no partial credit, and each answer is awarded either 0 marks or 2 marks. Only 10 candidates scored 80 or above. Let n_i $(0 \leqslant i \leqslant 100)$ denote the number of candidates whose score is i. Which of the following scenarios are possible?

(a) $n_0 = 100$.

(b) For all odd numbers $i \in \{0, \dots, 100\}$, $n_i = 0$.

(c) For all $i \in \{0, \dots, 100\}$, $n_i < 25$.

(d) There are at most 30 values of $i \in \{0, \dots, 100\}$ such that $n_i = 0$.

Answer. (a) and (b)

Since each question carries either 2 marks or 0 marks, (b) is always true (and therefore possible). Since (b) holds, (d) is automatically false. Note that we have to distribute 990 candidates among 40 possible marks (all even numbers from 0 to 78). If $n_i < 25$ for each of them, it covers at most 960 candidates. Thus (c) is also false.

(a) is possible (though not necessary). Here is a possible distribution of the marks.

$$n_0 = 100;\ n_2 = 130;\ n_{2i+2} = 20\ (1 \leqslant i \leqslant 38);\ n_{80} = 10;\ n_{2i+80} = 0\ (1 \leqslant i \leqslant 10). \qquad \dashv$$

6. Let $T(n)$ be the number of ways to tile a $3 \times n$ rectangle with 1×2 dominoes; refer to the figure for some examples for the case $n = 4$.

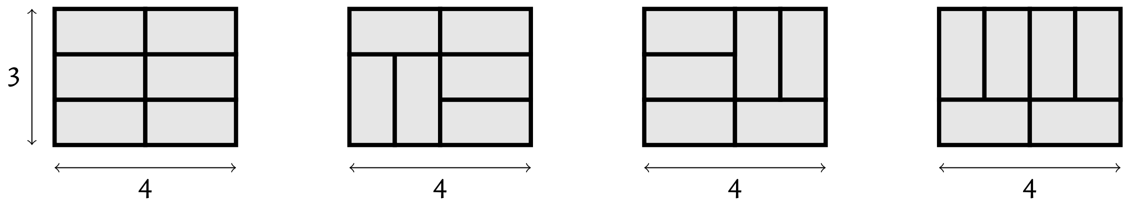

Four distinct tilings of a 3×4 rectangle with 1×2 dominoes.

Which of the following statements is/are true about the growth of $T(n)$?

(a) $T(2) = 3$

(b) $T(n) = 0$ for all odd n.

(c) $T(n)$ is given by a recurrence of the form $T(n) = T(n-1) + T(n-2)$.

(d) $T(4) < 9$.

Answer. (a) and (b)

- Option (a) is true. To see this, observe that a 3×2 rectangle can be tiled in 3 distinct ways using 1×2 dominoes: (i) three horizontal dominoes; (ii) one horizontal domino at the top and two vertical dominoes below; (iii) two vertical dominoes at the top and one horizontal domino below. See the picture below:

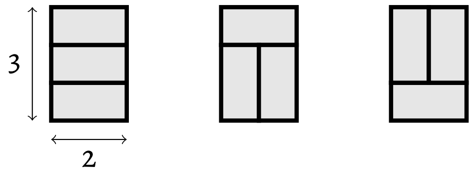

- Option (b) is **true**: Since each domino covers 2 squares, a $3 \times n$ rectangle can only be tiled if n is even. Thus $T(n) = 0$ for all odd n.
- Option (c) is **false**: If the recurrence $T(n) = T(n-1) + T(n-2)$ held, then since $T(1) = 0$, we would have $T(3) = T(2) + T(1) = T(2)$. However, we know $T(3) = 0$ from option (b), and $T(2) = 3$ from option (a).
- Option (d) is **false**. We know that $T(2) = 3$. It follows that there are at least 9 ways to tile a 3×4 rectangle, by considering two disjoint 3×2 blocks. Hence $T(4) \geqslant 9$. ⊣

7. A binary string s of length 2026 starts with bit `0` and ends with bit `1`. Let N_{00}, N_{01}, N_{10}, and N_{11} denote the number of times that the sequences `00`, `01`, `10`, and `11`, respectively, appear as pairs of adjacent bits in s. Which of the following equations are correct?

   (a) $N_{01} = N_{10} + 1$ (b) $N_{10} = N_{01} + 1$
   (c) $N_{01} = 1013 - (N_{00} + N_{11})/2$ (d) $N_{01} = 2026 - N_{00} - N_{11}$

   Answer. (a) and (c)

   The total number of adjacent pairs of bits in s is 2025, and so we have $N_{00} + N_{01} + N_{10} + N_{11} = 2025$. Since the string begins with a `0` and ends with a `1`, every transition from `1` to `0` must be—eventually—followed by a transition from `0` to `1`. It follows that the number of `0`-to-`1` transitions in s is exactly one more than the number of `1`-to-`0` transitions. That is: $N_{01} = N_{10} + 1$, and hence $N_{10} = N_{01} - 1$. Substituting, we get $N_{00} + N_{01} + (N_{01} - 1) + N_{11} = 2025$, from which we get $N_{01} = 1013 - \dfrac{N_{00} + N_{11}}{2}$. ⊣

Questions 8 and 9 pertain to the pseudocode given below, which defines the function $f(x)$ (where x is assumed to be a nonnegative integer). The code uses the function $\mathrm{rem}(n, d)$, which returns the remainder after dividing n by d.

```
function f(x)
    a ← 0
    while x > 0 do
        r ← rem(x, 2)
        a ← (10 × a) + r
        x ← (x − r)/2
    end while
    return a
end function
```

8. What is the value returned by $f(43)$?

   (a) 101000 (b) 101011 (c) 110101 (d) 110100

Answer. (c)

The sequence of values assumed by (a, x) at the start of each iteration of the loop is given below (in the execution of $f(43)$):

$$(0,43) \to (1,21) \to (11,10) \to (110,5) \to (1101,2) \to (11010,1) \to (110101,0).$$

The loop is exited when $x = 0$, and 110101 is returned. ⊣

9. Which of the following best describes the running time of $f(x)$?

(a) Linear in x.
(b) Quadratic in x.
(c) Logarithmic in x.
(d) Exponential in x.

Answer. (c)

At each iteration we replace x by $\lfloor \frac{x}{2} \rfloor$, and we stop when $x = 0$. So there can be at most $\log_2 x$ iterations. ⊣

10. Consider the following two procedures `proc1` and `proc2`, which run in parallel after initializing `x` to `6` and `y` to `3`. Running in parallel means that between any two lines of code in one process, any number of lines of the other process may run. Both `x` and `y` are shared variables which both processes can read/modify.

```
int x = 6, y = 3;

proc1 {                     proc2 {
    x = 2*y;                    y = 2*x;
    x = x - y;                  y = y - x;
}                           }
```

Which of the following values are possible for `(x,y)` after both processes have come to a halt?

(a) (8,5)
(b) (3,3)
(c) (0,6)
(d) (36,-12)

Answer. (b), (c) and (d)

Numbering the lines in `proc1` as 1 and 2, and those in `proc2` as 3 and 4, there are six interleavings possible. The table below lists them, along with the final value of `(x,y)`.

Interleaving	1234	1324	1342	3124	3142	3412
Value	(3,3)	(-6,18)	(0,6)	(12,0)	(36,-12)	(6,6)

Thus we see that (a) is not possible, while the other options are possible. ⊣

Part B

1. You are given a long sheet of (ruled) paper with the string MI written on the first line. You are asked to fill in lines 2, 3, 4, … (in order) with other strings involving the letters M, I and U, subject to the following rules, where x and y stand for any string (possibly empty) formed using the letters M, I and U:

   R1. If the previous line has Mx, you can write Mxx.

   R2. If the previous line has xI, you can write xIU.

   R3. If the previous line has xUUy, you can write xy.

   R4. If the previous line has xIIIy, you can write xUy.

   You can stop at any line, at any point of time, and the string on the last line you have filled in is said to be *derived*. Here are two examples (where we write $x_1 \to x_2 \to \dots x_{n-1} \to x_n$ to mean that the lines $1, 2, \dots, n-1, n$ have entries $x_1, x_2, \dots, x_{n-1}, x_n$ (in order).

   E1. MI. (This shows that you can stop even without adding any line.)

   E2. MI → MII → MIIII → MUI → MUIUI.

   E3. MI → MII → MIIII → MIIIIU → MUIU.

   Questions:

   (a) For examples E2 and E3 above, write down which rule from R1 to R4 is used in each step.

   (b) Show how to derive the strings MIUIU and MIUIUU. (**Hint:** To derive MIUIUU, first derive MIIIIIIII (M followed by 8 occurrences of I) and then proceed.)

   (c) Show that the string MIIMII cannot be derived.

   (d) If x is any string derived using the rules above, and if n is the number of occurrences of I in x, what are the possible values for n mod 3?

   (e) Can the string MU be derived? If so, provide a derivation. If not, explain why not.

   Answer.

   (a) The rules used in E2 are, in order, R1, R1, R4 and R1. The rules used in E3 are R1, R1, R2, R1.

   (b) For MIUIU, we have the derivation MI → MII → MIIII → MIU → MIUIU.
   For MIUIUU, we have the derivation

   $$\text{MI} \to \text{MII} \to \text{MIIII} \to \text{MIIIIIIII} \to \text{MIIIIIIIIU} \to \text{MIUIIIIU} \to \text{MIUIUU}.$$

   (c) In every string derived from MI, the letter M occurs exactly once. So MIIMII cannot be derived.

   (d) n mod 3 is either 1 or 2. It starts with a value of 1, and rules R3 and R4 do not change its value. R2 changes it from 1 to 2 and vice versa.

   (e) MU has no occurrences of I, and hence cannot be derived, by the previous part. ⊣

2. For a set X, X^* denotes the set of all strings formed using X, including the empty string. By xy, we mean the *concatenation* of x and y, i.e., the string x followed by the string y. We say that x is a *prefix* of y if there is some $z \in X^*$ such that $xz = y$. Note that every x is a prefix of itself, and ε is a prefix of every x.

   A function $f : \Sigma^* \to \{0,1\}^*$ is said to be a homomorphism if $f(\varepsilon) = \varepsilon$ and $f(xy) = f(x)f(y)$ for all $x, y \in \Sigma^*$. f is said to be *prefix-free* if $f(a)$ is not a prefix of $f(b)$ for any pair of distinct *letters* $a, b \in \Sigma$. For example, if $f(a) = 01$ and $f(b) = 101$, then $f(aba) = 0110101$. For the following questions, we assume that $|\Sigma| = 8$ and $f : \Sigma^* \to \{0,1\}^*$ is a prefix-free homomorphism.

   (a) Is it possible that $f(a)$ is of length at most 3 for all a. If so, define such an f. If not, provide a proof that this is not possible.

   (b) Is it possible that $f(a)$ is of length 2 for some $a \in \Sigma$ and $f(b)$ is of length 3 for all $b \in \Sigma \setminus \{a\}$. If so, define such an f. If not, provide a proof that this is not possible.

   (c) Show that there are no distinct $x, y \in \Sigma^*$ such that $f(x) = f(y)$.

   Answer.

   (a) We can map the 8 letters to 8 3-bit strings.

   (b) Suppose, w.l.o.g., that $f(a) = 01$ for some $a \in \Sigma$. Then no $b \in \Sigma$ can be mapped to the strings 000 or 001. So at most six elements of Σ can be mapped to binary strings of length 3. Since $|\Sigma| = 8$, we have at least one $b \in \Sigma \setminus \{a\}$ s.t. $f(b)$ is not of length 3.

   (c) This is proved by contradiction. Suppose that there exist distinct x, y such that $f(x) = f(y)$. Let $x = a_1 \dots a_k$ and $y = b_1 \dots b_\ell$ for k, ℓ where each $a_i, b_i \in \Sigma$. One of the following cases holds:

      (i) One of the strings is a proper prefix of the other.

      (ii) There exists $i \leqslant k, \ell$ s.t. $a_j = b_j$ for all $j < i$ and $a_i \neq b_i$.

      For case (i), assume w.l.o.g. that x is a proper prefix of y. Then $f(x)$ has to be a proper prefix of $f(y)$, since otherwise we would have $f(b_{k+1}) = \cdots = f(b_\ell) = \varepsilon$, and no prefix-free map can map a letter to ε, whenever $|\Sigma| > 1$.

      For case (ii), since $a_1 \dots a_{i-1} = b_1 \dots b_{i-1}$, we have $f(a_1 \dots a_{i-1}) = f(b_1 \dots b_{i-1})$. This, along with $f(x) = f(y)$, implies that $f(a_i \dots a_k) = f(b_i \dots b_\ell)$. From this, it follows that either $f(a_i = f(b_i)$ or one of $f(a_i)$ and $f(b_i)$ is a proper prefix of the other. In all cases, we have a violation of the prefix-freedom of f, since $a_i \neq b_i$. ⊣

3. Suppose $G = (V, E)$ is a graph with $|V| = 2k$ for some $k \geqslant 1$. Further, suppose there is no cycle of length 3 in G, i.e., G is triangle-free. Prove that $|E| \leqslant k^2$.

   Answer. Let u be a vertex of maximum degree, with $\deg(u) = d$. There are $2k - d - 1$ vertices in $V \setminus (N(u) \cup \{u\})$. These vertices can contribute at most $d(2k - d - 1)$ edges. By the 3-cycle free property, the vertices in $N(u)$ have no edges between them. So the maximum number of edges possible is $d + d(2k - d - 1) = 2dk - d^2$, which attains its maximum value of d^2 when $d = k$. ⊣

Part C

1. Consider a finite alphabet $\Sigma = \{0, 1\}$. For a string $w \in \Sigma^*$, we write $|w|$ for the length of w. For $u, v \in \Sigma^*$, we write $u \cdot v$ for the string obtained by concatenating u and v.

   Let us define an operation $\pi_{\text{even}}(w)$ inductively as follows:

$$\pi_{\text{even}}(\varepsilon) = \varepsilon$$

$$\forall w \in \Sigma^* \text{ and } \forall a \in \Sigma \qquad \pi_{\text{even}}(wa) = \begin{cases} \pi_{\text{even}}(w) & \text{if } |wa| \text{ is odd} \\ \pi_{\text{even}}(w) \cdot a & \text{if } |wa| \text{ is even} \end{cases}$$

   (a) What is $\pi_{\text{even}}(001011)$?

   (b) For a language $L \subseteq \Sigma^*$, define $\pi_{\text{even}}(L) := \{\pi_{\text{even}}(w) \mid w \in L\}$. What is $\pi_{\text{even}}((01)^*)$?

   (c) Show that if L is regular, $\pi_{\text{even}}(L)$ is also regular.

   Answer.

   (a) 001. It is the projection on to the even positions of the word.

   (b) $\pi_{\text{even}}((01)^*) = 1^*$

   (c) Let $A = (Q, q_0, \Sigma, \delta, F)$ be a DFA accepting L. Construct an NFA (with ε-transitions) $B = (R, I, \Sigma, \Delta, G)$ as follows:

   - $R = Q \times \{0, 1\}$
   - $I = \{(q_0, 0)\}$
   - $\Delta = \{((q, 0), \varepsilon, (q', 1)), ((q, 1), a, (q', 0)) \mid q, q' \in Q, a \in \Sigma, \delta(q, a) = q'\}$
   - $G = F \times \{0, 1\}$

   We can show that any run of A on w from q_0 to $f \in F$ can be mimicked by a run of B on $\pi(w)$, from $(q_0, 0)$ to $(f, i) \in G$, where $i = 0$ if $|w|$ is even, and $i = 1$ otherwise. Hence for every word w accepted by A, the projection $\pi_{\text{even}}(w)$ is accepted by B. Conversely, for every accepting run of B, there exist letters that can replace ε appearing in the run by some letter of Σ, to make it a run of A. If w is the resulting word, then the run of B is on $\pi_{\text{even}}(w)$. ⊣

2. Let Σ be a finite alphabet, and let $L \subseteq \Sigma^*$. Two strings u and v in Σ^* are said to be L-equivalent if for all $z \in \Sigma^*$, ($uz \in L$ iff $vz \in L$).

   (a) Consider $\Sigma = \{0, 1\}$ and let $L = 0^*1^*0^*$. For the following instances of u and v, explain whether u and v are L-equivalent.

      (i) $u = 0, v = 1$

      (ii) $u = 0, v = 00$

   (b) Let L be a regular language. Show that for any infinite set of strings S, there exist at least two strings in S that are L-equivalent.

   Answer.

   (a) (i) $u = 0$ is not L-equivalent to $v = 1$, since $uz \in L$ and $vz \notin L$ for $z = 010$.

(ii) $u = 0$ is L-equivalent to $v = 00$. From the regular expression, we can see that $0z \in L$ iff $00z \in L$ for all z.

(b) Let A be a DFA accepting L. For any two words u, v, if the runs of A on them leads to the same state, then u and v are L-equivalent. In an infinite set of strings S, at least two of them will go to the same state of A, and hence will be L-equivalent. $\dashv$

3. You are given a graph on vertex set $V = \{a_1, \dots, a_k\} \cup \{b_1, \dots, b_k\}$. The edge set is $E = \{(a_i, b_j) \mid i \neq j\}$. An edge coloring of a graph is an assignment of colors to edges such that no two edges sharing a common vertex receive the same color.

(a) Give an efficient algorithm to color the edges of the graph with minimum number of colors.

(b) How many colors does your algorithm use?

Answer. Each vertex in this graph has degree $k - 1$. Since all edges incident with any one vertex must get different colours, any edge colouring of this graph requires at least $k - 1$ different colours.

(a) Go over the edges of the graph one by one. Give the colour $(j - i) \bmod k$ to an edge of the form (a_i, b_j). Since every edge (a_i, b_j) satisfies $i \neq j$, these colours will be from the set $\{1, 2, \dots, (k - 1)\}$. So this colouring uses at most $k - 1$ colours.

Consider two distinct edges $(a_i, b_p), (a_i, b_q)$ incident on a vertex a_i. Since $p \neq q$ and $1 \leqslant p, q \leqslant k$ we get that $(i - p) \mod k \neq (i - q) \mod k$ holds. Thus no two edges incident with a vertex a_i get the same colour under this colouring. By a symmetric argument this holds for edges incident with vertices b_j, as well. So the algorithm gives an edge colouring of the graph using the optimal number of colours.

(b) The number of colours used by the algorithm is $k - 1$. $\dashv$

4. In the *Pots-of-Gold* game, there is a row of n pots $P_1, \dots, P_n$. Each pot contains some number of gold coins. You remove pots and collect the gold coins from the pots that you remove, subject to the rules below. Your aim is to collect as many gold coins as possible.

- You start with no gold coins and remove pots one at a time. You can remove any pot that is *available* for removal.
- In the beginning, all pots are available for removal. Removing pot P_i makes the pots P_{i-1} and P_{i+1} (if they exist) permanently unavailable.

(a) Assuming that we have an array $G[1 \dots n]$ with $G[i]$ storing the number of gold coins in pot P_i, devise an algorithm POTSOFGOLD(G, n) that computes the maximum number of gold coins that you can collect. Your algorithm should run in $O(n)$ time in the worst case. You can assume that each arithmetic operation involving two integers, and each array access (reading/writing one element) takes constant time. **Hint:** Use dynamic programming.

(b) Clearly explain why your procedure of part (a) correctly computes the required output.

(c) Explain why your procedure of part (a) runs in $O(n)$ time in the worst case.

Answer.

(a) **function** PotsOfGold(G, n)
  if $n = 1$ **then**
    return $G[1]$
  end if
  Let $dp[1 \dots n]$ be a new array
  $dp[1] \leftarrow G[1]$
  $dp[2] \leftarrow \max(G[1], G[2])$
  for all $i \in \{3, \dots, n\}$ **do**
    $dp[i] \leftarrow \max(dp[i-1], dp[i-2] + G[i])$
  end for
  return $dp[n]$
end function

(b) Let $dp[i]$ denote the maximum number of gold coins the player can collect from pots $P_1, P_2, \dots, P_i$.
For $i \geqslant 3$,

$$dp[i] = \max(dp[i-1], dp[i-2] + G[i]).$$

This holds because the player either does not take pot P_i (optimal value is $dp[i-1]$) or takes pot P_i (in which case they cannot take pot P_{i-1}, so the optimal value is $dp[i-2] + G[i]$). The base cases $dp[1] = G[1]$ and $dp[2] = \max(G[1], G[2])$ are immediate. By mathematical induction on i, the value returned by the procedure is correct.

(c) The algorithm performs a constant amount of work for the special case $n = 1$ and for initializing $dp[1]$ and $dp[2]$. The **for** loop then executes exactly $n-2$ times (when $n \geqslant 2$), with $O(1)$ work in each iteration. Hence the running time is $O(n)$. ⊣

5. In this question, we only consider real-valued functions from $[0, 1] \to \mathbb{R}$.

(a) Does there exist a function f that is continuous at exactly 0 and discontinuous at every point in $(0, 1]$?

(b) Suppose S is a finite subset of $[0, 1]$. Does there exist a function that is continuous at all points in S, and discontinuous everywhere else?

Justify your answers with proofs.

Answer. Let $S = \{x_1, \dots, x_k\}$ (Set $S = \{0\}$ for part (a)). Define

$$f(x) = \begin{cases} \prod_{i=1}^{k}(x - x_i) & \text{if } x \text{ is rational} \\ 0 & \text{otherwise} \end{cases}$$

f **is continuous at each** $x_i \in S$. Let $p(x) = \prod_{i=1}^{k}(x - x_i)$. This is a continuous function. Hence, for each $\varepsilon > 0$, there exists $\delta > 0$ such that for all x with $|x - x_i| < \delta$, $|p(x) - p(x_i)| < \varepsilon$. But, for all x such that $|x - x_i| < \delta$, $|f(x) - f(x_i)| = |p(x) - p(x_i)| < \varepsilon$ if x is rational; since $f(x_i) = 0$, this quantity is 0 when x is irrational. This proves continuity.

f **is discontinuous at each** $x \notin S$. Suppose x is rational. For any $\delta > 0$, there exists an irrational y such that $|x - y| < \delta$ since irrational numbers are dense in $\mathbb{R}$. Furthermore, when y is irrational, $|f(y)-f(x)| = |f(x)| = |p(x)| > 0$, where the last inequality follows from the definition of p. Thus, for all $0 < \varepsilon < |p(x)|$, there exists no δ such that $|f(x)-f(y)| < \varepsilon$ for all y such that $|x - y| < \delta$.

Suppose x is irrational. For any $0 < \varepsilon < |p(x)|/2$, appealing to continuity of p, choose δ such that $|p(x) - p(y)| < |p(x)|/2$ for all y such that $|x - y| < \delta$. Then, for all rational y such that $|x-y| < \delta$, $|f(x)-f(y)| = |f(y)| = |p(y)| > |p(x)|/2$ by the choice of δ. But this is more than the given ε. Thus there exists ε for which, no matter what $\delta > 0$ we pick, we can find a rational y such that $|x - y| < \delta$ (since rationals are dense in $\mathbb{R}$) but $|f(x) - f(y)| > \varepsilon$.

We have thus proved that f is discontinuous at each $x \notin S$, whether $x \in \mathbb{Q}$ or not. ⊣

6. Let $x_1, \dots, x_n$ be uniform, independent bits; i.e., each x_i is a random variable that takes value 0 or 1 with equal probability, and is independent of the other random variables. Recall: $(x_1, .., x_n)$ are said to be independent if, for all $b_1 \in \{0, 1\}, b_2 \in \{0, 1\}, \dots, b_n \in \{0, 1\}$,

$$\mathrm{P}\left[x_1 = b_1, x_2 = b_2, \dots, x_n = b_n\right] = \mathrm{P}\left[x_1 = b_1\right] \times \mathrm{P}\left[x_2 = b_2\right] \times \dots \times \mathrm{P}\left[x_n = b_n\right].$$

A set of bits $y_1, \dots, y_n$ is said to be *pairwise independent* if, for every distinct $i, j \in \{1, 2, \dots, n\}$, the bits y_i and y_j are independent.

(a) Prove that x_1, x_2, and $x_1 \oplus x_2$ are pairwise independent when $\oplus$ denotes binary addition (exclusive OR); i.e., $x_1 \oplus x_2 = x_1 + x_2 \mod 2$.

(b) Prove that $x_1 \oplus x_2$, $x_1 \oplus x_3$, and $x_3 \oplus x_2$ are pairwise independent.

(c) Can you create 7 pairwise independent bits using x_1, x_2, x_3? Justify your answer.

(d) Generalize this to show that n uniform, independent bits can be used to generate $2^n - 1$ pairwise independent sets. Justify your answer.

Answer.

(a) This is a basic exercise in probability.

(b) This too is a basic exercise in probability.

(c) The answer is 'Yes'. Consider all nonzero linear combinations:

$$x_1,\ x_2,\ x_3,\ x_1 \oplus x_2,\ x_1 \oplus x_3,\ x_2 \oplus x_3,\ x_1 \oplus x_2 \oplus x_3.$$

These are $2^3 - 1 = 7$ distinct bits. Any two distinct such combinations can be shown to be pairwise independent as in part (a) and (b).

(d) We claim that for any distinct nonempty subsets $S, S' \subseteq [n]$, $\oplus_{i \in S} x_i$ and $\oplus_{i \in S'} x_i$ are independent. This is easily seen by noting that at least two of $S \cap S'$, $S \setminus S'$, and $S' \setminus S$ are non-empty, and

$$\oplus_{i \in S \cap S'} x_i, \qquad \oplus_{i \in S \setminus S'} x_i, \qquad \oplus_{i \in S' \setminus S} x_i$$

are independent. Now, setting these parities to y_1, y_2, y_3, we can show the required independence using (a) and (b). This gives us $2^n - 1$ pairwise independent sets. ⊣

7. Let $M_n(\mathbb{R})$ denote the vector space $\mathscr{V}$ of all $n \times n$ matrices over the real numbers $\mathbb{R}$. You may assume that e_{ij} are the basis vectors of this vector space and x_{ij} are the dual basis. Recall that a matrix is said to be singular if its determinant is zero.

   (a) Let E_{ij} be the matrix with 1 in position (i, j) and zeros elsewhere. Describe a linear subspace $\mathscr{L} \subseteq \mathscr{V}$ of dimension $n^2 - n$ containing the elementary matrix E_{ij} such that all matrices in $\mathscr{L}$ are singular. If it helps, you may write down a system of linear equations in the variables x_{ij} whose solution space contains E_{ij} and has the requisite dimension.

   (b) Let $A \in M_n(\mathbb{R})$ be a matrix of rank less than or equal to $n - 1$. Show that there exists a linear subspace $\mathscr{L} \subseteq \mathscr{V}$ such that $A \in \mathscr{L}$, every matrix in $\mathscr{L}$ is singular and $\dim \mathscr{L} \geqslant n^2 - n$.

   Answer.

   (a) Let $k \neq i$. Then a linear subspace containing E_{ij} is $x_{ks} = 0, 1 \leqslant s \leqslant n$. This subspace has dimension $n^2 - n$. Clearly all matrices satisfying these equations have their k-th row zero and are singular.

   (b) W.l.o.g we may assume that the last row is a linear combination of the first $n - 1$ rows. Then there are real numbers $\lambda_1, \lambda_2, \ldots, \lambda_{n-1}$ such that for $1 \leqslant j \leqslant n$, $a_{nj} = \sum_{i=1}^{i=n-1} \lambda_i a_{ij}$. Let us define $\mathscr{L}$ to be the set of matrices satisfying $x_{nj} = \sum_{i=1}^{i=n-1} \lambda_i x_{ij}$ for all j. Clearly $A \in \mathscr{L}$. Furthermore, $\mathscr{L}$ is the solution space to a system of n equations in n^2 variables. Since every matrix M in $\mathscr{L}$ satisfies these equations, the n-th row of M depends linearly on the first $n - 1$ rows, and so M is singular. The solution space has dimension at least $n^2 - n$. ⊣

8. Consider a bivariate function $f : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$.

   (a) What does it mean to say f is continuous at a point $(a, b) \in \mathbb{R} \times \mathbb{R}$? Can you define it formally?

   (b) Consider any straight line L passing through the point (a, b). Show that f restricted to the line is a real function $f_L : \mathbb{R} \to \mathbb{R}$.

   (c) If f_L is continuous at the point (a, b) for each line L passing through (a, b) then is the bivariate function f continuous at (a, b)?

   Justify answers with proofs.

   Answer.

   (a) For every $\varepsilon > 0$ there is a $\delta > 0$ such that if (x, y) is in the δ-disc centered at (a, b) then $|f(x, y) - f(a, b)| < \varepsilon$.

   (b) Points on a line L through (a, b) is of the form $(a + t\alpha, b + t\beta)$ for real α and β (at least one is non-zero). Then $f_L(t) = f(a + t\alpha, b + t\beta)$ defines the desired one-variable real-valued function.

   (c) Not necessarily. The function $f(x, y) = \dfrac{x^2 y}{x^4 + y^2}$ is discontinuous at $(0, 0)$ but along every line L through $(0, 0)$ the function f_L is continuous at $(0, 0)$. ⊣